Gary D. Campbell

Reality, Intelligence, & the Future

Forward

Hi there! I'm a 75-year-old in the process of writing a book about all the things I've found interesting over the course of my lifetime. My first job out of college was software engineering for Hewlett-Packard. Later, the company from which I retired was Sun Microsystems. I was fortunate—I had a career that paid me to do work I really enjoyed doing—just like writing **this** book (except for the ***paid*** part). Another book I wrote 20 odd years ago serves as a companion to this book. It is titled ***Life: An Instruction Manual***. Check it out. Some of the thoughts left unfinished in that book are treated here. The thoughts left unfinished in ***this*** book are flagged with ?? double question marks. Uploads are made as I "finish" various sections (see Upload History at the end of the book).

Contents

Introduction

This book discusses how the cosmos works (reality), how a mind works (intelligence), and how a functioning society ***might*** work (in the future). It addresses the science, philosophy, and technology, that explain reality, intelligence, and what the future might bring. Its intended reader is a college graduate who has taken and understands several math, science, and computer programming courses. This is not a book for those whose beliefs will never change. Well intentioned people can sincerely believe things that are not true (or fail to believe things that are). I believe that everyone owes it to themselves to avoid and escape these traps. If you accept this as an important personal objective, you might find this book interesting. If not (spoiler alert), some of its assertions might even offend you. In any case, I expect most people to take issue with at least some of the statements I assert. Be my guest, let's sort these issues out!

I thought about calling this book "42" (the "answer" to the "Ultimate Question of Life, the Universe, and Everything," as asked in Douglas Adam's book, *The Hitchhiker's Guide to the Galaxy*), but I resisted, since I don't cover ***everything***.

One reason for this book is that today's telecommunication network is surpassing that of the human brain in both the number of connections it has and the amount of data that flows through it. This fact deserves a foundation and explanation. Viewed as a giant brain, the inputs to this vast network are the keyboards, cameras, and microphones of its human users, and its primary outputs are display screens and speakers. This network can already see, listen, talk, write, and paint pictures. But, not yet under its own volition. Many people feel a "game changer" is just around the corner.

First principles

My most fundamental belief is that matter/energy cannot be created or destroyed. If we accept this belief, a creation event is a contradiction. What would change this reality would be an entity with the ability to will something into existence—not just rearrange what already exists, but bring something into existence out of nothing simply as an act of will as we understand will. If solid evidence existed that some entity were able to do that, I would concede the existence of a god. Until then, let's look elsewhere for our "first principles."

As far as we know, there is always something; there is never nothing; and all existence is transient.
Our truths are stated in our narratives—collections of statements that we tell ourselves as much as we tell each other.
Statements about things that don't exist are fantasy; arguments about fantasies are nonsense.
True statements cannot be proven false, but true statements may not be provably true (we can't know until they are).
False statements cannot be proven true, but false statements may not be provably false (so always have some doubt).

A statement that asserts or implies existence requires that existence to be shown consistent with observations, otherwise the statement is either fantasy or nonsense. Thus, assuming a statement describes reality, it may be provably true, provably false, or undetermined until a counter example or a proof is brought to light.

We attempt to describe observations with models or analogies. However, even if a set of observations does correspond to a model, that does not make the model correct, nor preferable to other possible models. One kind of model is a map. A map should not be confused with actual territory. The territory is reality. Science is a collection of maps. Maps help guide us through the territory. Over the ages our collection of maps has become better and better. Maps are not wrong just because they are not perfect. In fact, no map (model, equation, or theory) is ever perfect. This was explained by Isaac Asimov in ***The Relativity of Wrong***. Click: http://chem.tufts.edu/answersinscience/relativityofwrong.htm

An ancient proverb states: He who knows not, and knows not that he knows not, is a fool, shun him. He who knows not, and knows that he knows not, is a child, teach him. He who knows, and knows not that he knows, is asleep, wake him. He who knows, and knows that he knows, is wise, follow him.

Unfortunately, fools exist whose eloquent speech is hard to ignore, so it's not always easy to know whom to follow.

Everyone has a lot of narratives. A narrative is resistant to change, but it can evolve. We generate our own narratives with slight changes over and over, for the purposes of analogy, argument, and thought itself. We evolved intelligence along with our ability to tell stories. Many creatures can signal a noun or verb to others of their kind, we are the only animals able to signal sentences of words connected by a grammar—sentences connected to each other to form narratives.

Our first imperative is to accept our own observations and form our own conclusions. Our very self-awareness gives us our first existential truth. But observation may involve fundamental existential disconnects. Unless I go there, I can't observe what's in the next room, much less what's in the 3rd galaxy beyond Andromeda. Science knows of no way to observe space directly. In fact, we can't observe ***anything*** directly. Nearby, far away, or passing by, we can't observe anything even ***indirectly***, unless we (or one of our devices) interacts with it. And, interactions between forms of matter and energy are really quite peculiar. Even though it appears otherwise, almost all of our surroundings are actually empty space.

No matter how rational we are, all of us must believe things that we can't prove to ourselves. Each of us exists in at least three domains: ***reality***, ***beliefs***, and ***imagination***. What's the connection between reality and imagination? Reality impacts us from outside ourselves. It may influence our imagination, but our imagination can produce a reality of its own, these are our beliefs.

When we build models of reality, we need to remain aware that the model and reality are very different things. Models are products of the imagination; ***reality*** is the whole universe at the present moment—a thing that we barely touch. Each of us conceives parts of reality in our imagination, and we come to believe in constructs and narratives that we ***imagine*** are real.

Until intelligence evolves, everything forms through chance and necessity from things that previously existed. Intelligence allows a third cause: ***free will***. One definition of free will is "the ability to act without the constraint of necessity or fate." Necessity stems from the rules that constrain every action or event. All events are governed by the necessities of natural laws. Chance also has a role in every event. If chance is the same as fate, no action can avoid it. Does this mean there is no such thing as free will? Actions taken ***within*** the constraints of chance and necessity are all those actions that are even ***possible***. When alternatives exist, a choice of action or behavior can always be sought. Being able to recognize that alternatives exist, and making and executing choices, are ***skills*** that intelligent entities can ***learn*** (with guidance and effort!). This is my definition of ***free will***.

Each of us arrives at a different position in life as the result of each choice we make. Just as we succeed or fail at each of our endeavors due to the vagaries of luck, we succeed or fail at making the optimum choice at each of the opportunities presented to us. Each of us becomes a different person as we choose to modifiy and add to our narratives.

Although free will exists only within the limits imposed by chance and necessity, this gives an individual many opportunities to make choices that are better or worse. Nature and society tend to punish us for making bad choices! Free will offers each of us the opportunity to create or destroy, to work together or at cross purposes. It offers us the opportunity to develop and follow a moral compass. Choices follow opportunities. None of us gets the same opportunities, but all of us bear the responsibility for our choices. The choices we make lead to a major component of who we become. Put another way, we can't be blamed for our opportunities, but we and others bear the consequences for our personal choices.

I believe that anything that exists may be known. Anything that does not evidence itself is imaginary. I believe everything can ***potentially*** be explained by a combination of chance, choice, and necessity, but how each of these enters into a situation may be impossible to ascertain, and therefore may remain out of the reach of our understanding.

I believe in an objective reality that can only be imperfectly ***inferred*** from the subjective reality of our individual existence.

I try not to confuse the natural world (the world we experience when awake and sober) with my dream and fantasy worlds. Explanation and understanding do not stem from supernatural, dream, or fantasy worlds. A world view must be built upon analogies, however, and these often do come from the imagination. Another fundamental belief of mine is that a ***belief structure*** should be as parsimonious as one can make it, that Occam's razor should be applied whenever possible, and that all claims require proof. Beliefs come in different strengths. The ***less*** a belief is consistent with observation, proof, and general agreement, the ***weaker*** that belief should be. Improperly supported beliefs are unsound.

Note: If I use a phrase such as "Occam's razor," and you aren't sure what I mean, "Google" it. Many of the concepts that I might not explain to your satisfaction are explained in great detail somewhere on the Internet. The explanations there are not always accurate or complete, but if some of the material ***here*** seems wrong or trivial, read through it anyway just as a mental exercise. There are some sections where the contemplation time should exceed the reading time.

Science ***begins*** with observations, and it constructs models to explain them. The preferable model, when more than one is plausible, is the model that makes the fewest and simplest assumptions. A model is only plausible if it is consistent with observations. Observations are ***required***. Knowledge begins with the present moment. The past is reconstructed, either from our own memories, or from evidence we can observe in the present. Models that don't explain anything, or don't conclude or predict anything, or cannot be tested, are worthless. To be ***preferable,*** a model must ***better*** conform to these criteria.

Why is there something?

Let's ask the philosophical question that has always been the "elephant in the room." Why is there something rather than nothing? This question has often been asked of science, philosophy, and religion, and a satisfactory answer may never be known for certain. Many opinions exist. This book is yet another brain dropping onto that heap. The motivations for my definition of reality include starting from the most simple assumptions and connecting each more complicated concept to the ones that I have already accepted. A lot of trial and error and backtracking was involved.

Descartes said "I think, therefore I am." Individually, we realize that we exist. We also observe the existence of much besides ourselves. Our existence is necessary to make our observations possible. We know of no principle, mechanism, or even an unexplained observation that would lead us to reject the law that matter and energy are conserved. One may be exchanged for an equivalent amount of the other, but neither appears nor disappears spontaneously. Therefore, everything that exists is simply an arrangement of things that previously existed. How could this regress end? To me it appears infinite. I see no reason that "previous" would not go on forever.

Things become arranged in a sequence of events. Each event occurs due to a combination of chance and necessity. Choice may also be involved in rare occasions. The present, therefore, arose through a long sequence of events. Can this sequence be traced back through time to an initial condition? Was a creation event necessary to set everything into motion? A second alternative is that all matter and energy has existed forever. Is there a third alternative?

Even if our observable cosmos arose in some kind of Big Bang, that does not rule out some kind of Big Collapse being a link in an eternal chain of [Bang, normal existence, Collapse].. repeated forever, with no "first" anything. Does a creation event actually constitute an answer? My limited imagination (and there may be fundamental limits to human imagination for which I should not have to apologize) tells me that there is no third alternative, nor any explicable creation event, so I am left with the first alternative—a cosmos that is eternal. Q: Why is there something? A: Because there has never been nothing. Any ***reason*** for a better answer seems to be missing.

That there is something cannot be denied. ***Why*** there is something may simply be an incoherent question. However, given enough time and the right preconditions, anything that is possible is inevitable. If all existence is rooted in previous existence, what is the nature of time? Time is an ongoing series of "ticks" (I'll have more to say about "time" later). Eternity could be defined as the span of time between some "tick" concurrent with the event most distant in the past that had any effect ***on*** the present, and a tick in the future concurrent with the most distant event that can be affected ***by*** the present.

This is like playing a hand of cards and giving the deck one shuffle before playing the next hand. Each hand is affected by a partial shuffle, but after enough shuffles (ten hands or so), new hands are completely random with respect to the original hand. In other words, the order left after the 1st shuffle will be completely lost when the 10th hand is dealt. Every bit of order we can observe today can be traced back into the mists of time until earlier events are completely obscure. For some patterns the mist forms in seconds or less, other patterns may take billions of years to erase. The "beginning" of our eternity occurred sometime before the past becomes totally obscure to us, and the end of our eternity will occur when our present becomes totally obscure in the future.

When we observe distant stars, we look back in time. We estimate that the most distant stars we can observe are about 13.8 billion light years away from us, so we are looking back 13.8 billion years into the past. What does this mean? Does the universe suddenly become obscure at this point? Or, will we eventually be able to see further than this?

Last night (as I first wrote this) Steven Hawking presented (in a recorded TV interview) his answer to "Why are we here?" In essence, his answer rests on his belief in the "many worlds" theory of existence. He appears to believe that every possible outcome of a random quantum event produces a parallel branch of the universe. Thus, your existence was guaranteed from the start, and the very fact that you are here proves him right. I believe this is nonsense!

Two scientific principles are violated by believing this. It's not the simplest explanation, and it can't be tested. The universe literally ***arranged*** to have us. We are each of us here because we won the "lottery" that picked our arrangement, not because every possibility actually occurred, each on its own separate branch of existence. Positing that a virtually infinite number of universes exists in parallel, begs far more questions than it answers. Likewise, a creator, or even a creation event, both beg the questions: How did the creator get created? Or what came before the creation event?

Some of the interpretations of quantum uncertainty lead to conclusions that scientists still do not agree about. Let's do agree that claims that can't be tested or sorted out from other claims are merely conjectures. A single, eternal universe is therefore the simpler conjecture (unless we can somehow ***prove*** otherwise). Note: I use the word eternal to mean unbounded in time. As we shall see later, unbounded may mean infinite, or curved back upon itself. In the section titled ***Reality***, these concepts will be revisited.

The Rules of Nature

Our understanding of nature begins with observations. Repeated observations reveal that rules are followed. One of the first things that children do is play games. Games have rules. So does nature. But, are rules ***real?*** The difference between rules and reality can best be seen in the dichotomy between software and hardware (mind and brain, spirit and flesh, etc.).

These pairs describe rules (and similar information) recorded on physical substrates. The substrates exist, and in a different way, so do the rules. Rules describe regularities. Rules are literally written into our DNA which, itself, evolved in a natural process involving the arrangement of physical matter. The rules of nature were not written in ***advance***. Evolution is literally a process of writing rules. Science is the process of discovering rules and writing them down for ourselves. Rules are not nothing. Rules are an emergent property of nature. Rules exist only when they are recorded, either on a physical substrate, or in an emission of energy.

Several trains of thought might be triggered by these ideas. For instance, what is the ***nature*** of rules? How do minds, software, and language relate to rules?

In this book I will attempt to explain the nature of intelligence, and how it emerged from life, and how life emerged from existence (our reality). We've already given some thought to the nature of existence. The cosmos, the universe, and nature are all similar concepts. At the most fundamental level, nature follows rules. I'll write more about material existence later, for now, I'll introduce several subjects in an effort to exemplify the nature of rules themselves.

Numbers & Counting

Even birds can count (crows appear to be able to count to about seven)! The most primitive number system is the (0, 1) dichotomy: non-existence or existence, false or true, a binary value. Many ways can be devised to count beyond this. For example, the concepts of two, three, and many. Or, we could give a word to each of the number of fingers we have. This may be why we use the decimal system. Numbers and counting evolved over thousands of years. This evolution didn't proceed in the most logical order. With the wisdom of hindsight, let's learn how to count all over again.

The first step is counting on our fingers and giving a name to the count represented by each different number of fingers. With both fists tightly closed, no fingers extended, we have zero (or 0) fingers. If we extend one finger (it doesn't matter which one), we call that amount one, or 1. The second finger gives us a count of two, or 2 (again, we don't care which finger, because when we are counting things, the things themselves don't matter, it is the ***number*** of things that we are interested in). In the same fashion, additional fingers are called three, or 3, then four, or 4, etc. This gives the two sets: {0, 1, 2, 3, 4, 5, 6, 7, 8, 9} and {zero, one, two, three, four, five, six, seven, eight, nine}.

Now, we can count to nine in the ***decimal*** system. Let's learn how to count to ***any*** number in the ***binary*** system.

Counting requires a unique number for each possible count. We'll start with a count of ***n*** = 0. Each successive count is ***n***+1, where the result is formed by adding 1 (as follows) to the previous count. A simple procedure for computing ***n***+1 given ***n*** (where ***n*** is a series of 0 and 1 digits) is:

1. ***Start*** with the right-most digit of ***n***.
2. If this digit is 0, change it to 1 and set Carry = 0.
3. If this digit is 1, change it to 0 and set Carry = 1.
4. If Carry = 0, ***stop*** (any higher digits remain as they are).
5. Otherwise (Carry = 1), continue with the next higher digit (the one to the left of the previous digit).
6. If this digit does not exist, define it as 0.
7. ***Loop*** to step 2.

This defines all possible numbers in ***base two***, or binary (a series of zero and one digits). Here, each digit is called a "bit" (***b***inary dig***it***). Numbers used to count may also be used to indicate the position of an entity within a list. It is strictly by convention as to whether the first element of a list is given the number 0 or 1. Numbers used for counting, ordering, and positioning are called ***ordinals***. Here, we have seen the "add one" operation. Adding and subtracting are operations on numbers that increase or decrease the count a given number of times. Using the above procedure, a unique binary number may be written for ***any*** count.

Do you see how to count in ***decimal*** using the above rules? Only rules 2 and 3 need to be changed. Simply substitute:
2. If this digit is {0, 1, 2, 3, 4, 5, 6, 7, 8}, change it to {1, 2, 3, 4, 5, 6, 7, 8, 9} and set Carry = 0.
3. If this digit is 9, change it to 0 and set Carry = 1.

When numbers are the subject of discussion, a specific number may be written as a string of digits, but a generic number is often written as a letter (which I will distinguish with ***bold italic***). Generic numbers may often be related to one another with relational operators. These are the relational operators: less than {***a*** < ***b***}, equal {***a*** = ***b***}, greater than {***a*** > ***b***}, less than or equal {***a*** ≤ ***b***}, and greater than or equal {***a*** ≥ ***b***}. They define ***a*** as being one or more counts less than or greater than ***b***, or as equal to ***b*** (exactly the same count).

Writing Numbers

When you see a number with more than one digit, how do you know what base it is written in? In this text, only binary, decimal, and hexadecimal numbers are written. Binary is a string of 0's and 1's. Decimal is a string of 0's through 9's, and is the default. If any ambiguity is likely, a decimal is not begun with a 0. A hexadecimal number may have one of 6 other digits (A-F or a-f). Hex either begins with 0 or it ***contains*** (but does not begin with) one of the letter digits. One ambiguity still remains, however. A single digit is the ***same*** number in any base (for which the digit is defined), so that leaves the base ambiguous, but a string of 0's and 1's could be a number in ***any base***. When otherwise ambiguous, a ***binary*** number of 2 or 3 digits will have a single space inserted between digits. A ***decimal*** number consisting only of 0's and 1's will have a space (or comma) inserted between every 3 digits, and an ambiguous ***hexadecimal*** number will have a space inserted between every ***two*** digits. A ***binary*** number with more than 4 digits will always have a space inserted every 4^{th} digit (each group of 4 binary digits, when converted into base 16, is equivalent to one hex digit).

When a string of digits is written, the most significant digit is written to the left, and its meaning depends upon how many digits appear to its right. Each position represents the ***base*** raised to the next higher power. So that there is no ambiguity, numbers that represent a base or a power are always written in ***decimal***. Binary 1010 means $1\times2^3 + 0\times2^2 + 1\times2^1 + 0\times2^0$. Decimal 234 = $2\times10^2 + 3\times10^1 + 4\times10^0$. And, hexadecimal 1F5 = $1\times16^2 + 0F\times16^1 + 5\times16^0$, where 0F = 15. If you are not familiar with the notation used in the last few sentences, it is explained later (under *Measurement*).

Arithmetic

Arithmetic defines a set of binary operators. Here, "binary" means that these operators each accept ***two*** operands (the number to the left of the operator, and the number to the right of the operator). Each operator produces a particular result. Arithmetic expressions are ***statements*** that consist of variables, numbers, and operators.

Variables are usually expressed as letters, numbers are expressed as strings of digits, and operators as special symbols or abbreviated words. The operators I'm going to describe here are ways that two binary numbers can be combined to produce a binary result. When an expression is written with 2 or more operators, operations are performed left-to-right, or grouped by precedence or parentheses, with a result becoming the final result, or an operand for a subsequent operation. Any number in a given base may be converted to an equivalent number in any other base (a method for this is explained below).

a AND b

This operator accepts two binary numbers and forms a binary result such that each bit position of the result is 1 if ***both*** operands contain a 1 in the same bit positions; otherwise that bit position of the result is 0. If a = b, (a ***and*** b) = a = b. The order of the operands (which is left, and which is right) does not matter. For example 1011 ***and*** 1001 → 1001.

a OR b

This operator accepts two binary numbers and forms a binary result such that each bit position of the result is 1 if ***either*** operand contains a 1 in the same bit positions; otherwise that bit position of the result is 0. If a = b, (a ***or*** b) = a = b. The order of the operands (which is left, and which is right) does not matter. For example 1011 ***or*** 0100 → 1111.

a XOR b

This operator accepts two binary numbers and forms a binary result such that each bit position of the result is 1 for each bit position that contains bits of ***opposite*** value; otherwise that bit position of the result is 0. If a = b, (a ***xor*** b) = 0. The order of the operands (which is left, and which is right) does not matter. For example 1011 ***xor*** 1111 → 0100. Notice that ***xor*** with a string of one bits produces the ***one's complement*** of the original number (the value of each bit is "flipped").

a + b (plus)

This operator (add) accepts two numbers and forms a binary result that is the total (or sum) of the counts represented by its two operands. The order of the operands (which is left, and which is right) does not matter. Let the initial result be ***a***, then perform ***b*** {n+1} operations (as above) on the initial result to get the final result.

a – b (minus)

This operator (subtract) accepts two numbers and forms a binary result that is the difference of the counts represented by its two operands. The difference of two counts is the count required to increase the smaller count to equal the larger count. If the 2nd operand is the smaller count, the difference is an ***ordinal***. For example, let $c = a - b$. Since c is the count required to increase b enough to equal a, this means that $a = b + c$. However, order matters with the subtract operator. A complication arises when $a < b$. To signify that the 1st number was the smaller, the convention is to prefix the result with a – sign. This indicates that the result is ***negative*** (an ordinal with a – sign is called an ***integer***).

a × b (times)

This operator (multiply) accepts two numbers and forms a binary result that is the product of the counts represented by its two operands. The result is $a+a+\ldots+a$, where there are b occurrences of a. The order of the operands does ***not*** affect the result. When either operand is negative, it is converted to its positive counterpart, and the result is negative if only ***one*** of the two operands is negative. Otherwise, the result is positive.

a ÷ b (divided by)

This operator (division) accepts two numbers and forms a binary result that is the quotient of the counts represented by its two operands. The order of the operands ***does*** affect the result. For two ordinals, b is subtracted from a as many times as it can be, and the ordinal result is the count of times this can be done. The count left over after the last full subtraction is done is called the ***remainder***. The remainder is also called the result of ***a modulo b*** (or ***a mod b***).

When either operand is negative, it is converted to its positive counterpart, and the result is negative if only ***one*** of the two operands is negative. Otherwise, the result is positive. $a \div b$ can also be written $^a/_b$ or a / b.

Changing Base

The following procedure generates one digit of the result (or answer) for each loop through the instructions. The (above) ***mod*** and ÷ functions are used. Given a number n in any base, it can be converted to another base b as follows:
1. Starting with the right-most digit of the result, and working one digit at a time to the left ($r_n \ldots r_0$), begin with $i = 0$.
2. Let $r_i = n \text{ mod } b$; and let $n = n \div b$ (ignore the remainder); if $n = 0$, stop.
3. Let $i = i + 1$ (continuing with the next digit of the result); goto 2.

For example, convert the number 23 (written in decimal) to binary (n = 23, b = 2):
1. $r_0 = 23 \text{ mod } 2 = 1$, and $n = 23 \div 2 = 11$.
2. $r_1 = 11 \text{ mod } 2 = 1$, and $n = 11 \div 2 = 5$.
3. $r_2 = 5 \text{ mod } 2 = 1$, and $n = 5 \div 2 = 2$.
4. $r_3 = 2 \text{ mod } 2 = 0$, and $n = 2 \div 2 = 1$.
5. $r_4 = 1 \text{ mod } 2 = 1$, and $n = 1 \div 2 = 0$.
6. $n = 0$, so stop. Now, $(r_4 \ldots r_0)$ = 1 0111.

Thus, 23 converted to binary = 1 0111. Likewise, binary 1 0111 can be converted to decimal 23. An easier way to convert from binary to decimal is to add up the powers of two for each 1 digit. That gives 1×16 + 0×8 + 1×4 + 1×2 + 1×1 = 23.

a^b (a to the b)

In this operation (exponentiation), b may be positive or negative, or one of the following: ½, ⅓, ¼. The result when $b < 0$ is identical to the result of $(1/a)^b$, where $b > 0$, thus it can be computed in two steps. The result when $b = 0$ is 1 for all values of a. The result for $b > 0$ is $a \times a \ldots \times a$, where there are exactly b occurrences of a. The result for $b = ½$ is called the square root. This means that $a = a^{½} \times a^{½}$. The result for $b = ⅓$ is called the cube root ($a = a^{⅓} \times a^{⅓} \times a^{⅓}$). And, the result for $b = ¼$ is called the 4th root ($a = a^{¼} \times a^{¼} \times a^{¼} \times a^{¼}$). These results are easiest to compute by using iterative approximation using the operators defined above. Notice that the exponents inside the parentheses add up to 1, and $a^1 = a$. Thus, multiplying a base raised to different exponents is the equivalent of simply adding the exponents, giving in general: $a^x \times a^y = a^{x+y}$.

Newton's Method

Newton's method is one way roots can be calculated by iterative approximation. It involves making a guess, and improving that guess step-by-step with successive approximations until a "final" approximation is the "same" as the previous one (to a given degree of accuracy). Each step involves computing a better guess (x_{n+1}) from the previous guess (x_n) as follows:

$$x_{n+1} = x_n - f(x_n) / f'(x_n)$$

where, $f'(x)$ is the derivative of $f(x)$. The value of the function must be zero, so to express the cube root of 100, we need the function to be, $f(x) = x^3 - 100$. The derivative of this is $3x^2$. Now, we have things expressed in terms of operations we know how to do. We know that $x^3 = x \times x \times x$, but we don't have a similar way to compute $x^{1/3}$.

When it comes to taking roots, we can guess that a root will be a fraction of the original number, so let's let our first guess be 100 / root. That is: $x_0 = (100 / 3) = 33.33$. This gives us the following steps:

$$x_1 = 33.33 - (33.33^3 - 100) / (3 \times 33.33^2) = 22.25$$
$$x_2 = 22.25 - (22.25^3 - 100) / (3 \times 22.25^2) = 14.90$$
$$x_3 = 14.90 - (14.90^3 - 100) / (3 \times 14.90^2) = 10.08$$
$$x_4 = 10.08 - (10.08^3 - 100) / (3 \times 10.08^2) = 7.05$$
$$x_5 = 7.05 - (7.05^3 - 100) / (3 \times 7.05^2) = 5.37$$
$$x_6 = 5.37 - (5.37^3 - 100) / (3 \times 5.37^2) = 4.74$$
$$x_7 = 4.74 - (4.74^3 - 100) / (3 \times 4.74^2) = 4.64$$
$$x_8 = 4.64 - (4.64^3 - 100) / (3 \times 4.64^2) = 4.64$$

Notice that the 8^{th} iteration doesn't change the result (to 2 places after the decimal). The above calculations were done on a calculator that carries 10 significant digits, but for each iteration, I re-keyed the previous result to only two places after the decimal. Additional iterations would enable any number of digits after the decimal to be computed.

From this example, you should be able to set up the procedure to compute any (integral) root of any number.

My belief is that a description of the natural world needs only these fractional values of ***b***. General mathematics, however, allows roots of any value of ***b***. Another observation: identity operations exist that make the computation of multiply, divide, and exponentiation take many fewer steps than the simple explanations given above. Here, I've completely omitted such numbers as $-1^{½}$ and various irrational numbers. Look to the web when you wish to explore further.

Measurement

After numbers were used to count and do simple arithmetic, they were used to measure things. This allows for counting things that are not natural wholes—things that are composed of (very) many parts. Measurements are written as a string of decimal digits with a "decimal point" (binary numbers use a "binary point"). If the context is not clear, a number with a "point" is a decimal measurement. For example, the measurement 0.1 expresses one-tenth of something. Hexadecimal numbers will never be written with a "point" in this text. Digits to the right of a "point" are equal to exponents counting backward from zero, so that 0.123 means $1\times10^{-1} + 2\times10^{-2} + 3\times10^{-3}$.

Measurement generally implies units: miles, hours, gallons, etc. Units are usually expressed to the right of a measurement. The four standard parts of a measurement (left to right) are: (1) its optional sign (+ is the default, or – if negative), (2) a sequence of digits, optionally including a decimal point (called the mantissa or significance), (3) an optional magnitude or exponent, and (4) the units. Another attribute of a measurement is the number of its significant digits. The magnitude, if one is present, is written using the exponential notation, $\times10^n$. This means that the decimal point (or implied decimal point) could be moved left (if the exponent is negative) or right (if the exponent is positive) ***n*** places to omit the exponential notation.

Very large measurements and very small fractions may have an arbitrary number of zeros to the left or right of the non-zero (significant) digits, with the decimal point next to the most distant zero. The compact way to write such numbers is to shift the decimal point one place to the right of the left-most non-zero digit, and follow the number with $\times10^n$ if the decimal point has been shifted ***n*** places to the left—or if the decimal point has been shifted ***n*** places to the right, then $\times10^{-n}$. For example, 0.0000000123 is more compactly written: 1.23×10^{-8}.

While each digit in a counting number is assumed to be significant, not all the digits in a measurement are necessarily so. Normally, only the significant digits are written down, but when a measurement is stored in a machine register or memory location, it may be truncated (with loss of accuracy), or extended with extra digits (ignoring its accuracy). Expressions involving two or more measurements can be no more accurate than their least accurate operand. Certain operations such as divide and iterative approximations (such as Newton's method) may be carried out to any number of digits (following the decimal point), but that doesn't mean they are all significant in a given context. If the final digits of a measurement repeat indefinitely, that measurement is known as a rational number. If the final digits never repeat, the number is called irrational. These two terms refer to the fact that a rational number can be expressed exactly as some ***a / b***, but an irrational number cannot be expressed exactly with any finite number of digits (or terms in an expression).

Calculators

The calculator is the quintessential device that has a few buttons, a display, and contains a microprocessor. It replaced slide rules for calculation in the early 1970's, and since then more appliances every year have a few buttons and a display on the outside, and a microprocessor inside. The heart of a microprocessor is a set of binary registers that each hold a fixed number of bits. Calculators are able to perform all the operations explained above. A typical microprocessor might have a "word size" of 16 or 32 bits. Usually their memories consist of 8-bit chunks, called bytes. Each byte has an address which, is another binary number. Most of the numbers in a microprocessor are treated as ordinals, but there are two other formats which are often used, and these are worth knowing about. Note: computers, like the one I'm writing on right now, are just faster microprocessors with larger memories than those in calculators. Both do arithmetic as described here.

Negative Numbers

The "existence" of negative numbers was not "discovered" until the subtract operation "matured." The first thing you notice with ***subtract*** is that order matters ($x–y \neq y–x$). While 12+13 and 13+12 both = 25, when + is replaced by – we can write 13–12=1, but 12–13 was undefined. Today, we reverse the numbers, do the subtraction, give the number a ***minus sign prefix***, and call the answer ***negative***. Thus, in today's mathematical notation, 12–13 = –1.

Numbers used strictly for counting (ordinals) are always greater than or equal to zero. Numbers that may be negative are called ***integers***. These numbers may be used to count both forward and backward, for example to represent the position on a line, like the distance from the last city you passed through versus the distance to the next city on your route. The former increases as a positive value as you go forward. The latter is negative, representing the distance remaining until you get there, but it becomes less negative as you approach.

There have been several ways to represent negative integers in a computer. I will explain the most common. Integers are represented as the ***two's complement*** of their positive representation. That is, complement all bits and add one to the result. One bit of significance is now required for the sign. For example, 5 is binary 0101. Using just 4 bits, how would –5 be represented? Complemented, 5 is: 1010. Add 1, and you get: 1011. This is how –5 is stored in a 4-bit register. In larger registers, the MSB (Most Significant Bit) is sign extended. Adding 5 and –5, 0101+1011 we get 0000 (plus a carry). Zero is the right answer, but what does the carry mean?

When adding two ordinals, a carry indicates an overflow. This is because all of the bits are used to represent the count. The extra bit required to represent the sign of an integer means that the same number of bits represent only ***half*** the magnitude of an ordinal. When two binary numbers are added, they may be two (unsigned) ordinals, or they may be two integers. A positive integer is identical to its ordinal equivalent (although it is restricted to one less bit of magnitude with its MSB = 0 to indicate that it is positive). A negative integer, however, is represented in two's complement, giving it a bit pattern similar to an ordinal, but with its MSB = 1.

Subtract one from 0000 and you get 1111 (the equivalent of –1). A –2 is 1110; –3 is 1100, and so on, until –8 is 1000. Did you notice that –8 is represented by the same value as the ordinal 8? When two numbers are added in a fixed register, if a carry occurs ***out*** of the MSB, it means an ordinal overflow. If the MSB ≠ the carry, it means an integer overflow. For example, if you add 1 to 7 in binary you go from 0111 (7) to 1000 (8). This operation resulted in carry = 0, and MSB = 1. Since the carry = 0, it is ***not*** an ordinal overflow, but since the MSB ≠ carry, it ***is*** an integer overflow.

Floating Point Numbers

A "floating point" format is used inside a computer to represent measurements. The MSB is the "sign" of the number. The next ***n*** bits are the "biased" exponent, and the final ***m*** bits are the "normalized" mantissa. For example, if a floating point number is to be represented in 32 bits, the left-most bit would be 0 if the sign were positive, and 1 if the sign were negative. The next 8 bits could be the number's biased exponent. This is a count of how many places the binary point needs to be shifted plus a ***bias***. In this case, the final 23 bits would be the normalized mantissa, the string of significant digits with the top bit removed (because it is always 1, and is therefore redundant). The binary point is implied between the removed 1 digit and the first digit of this final string of digits, giving one more bit of significance than the number actually stored.

For example: Consider the binary number: 0.0000 0011 1100. This number is normalized to become: 1.1110×2^{-7}. The 32-bit floating point representation is: 0011 1100 0111 0000 0000 0000 0000 0000. Starting with the 10th bit, the next 3 bits are the same as those to the right of the decimal point, but the 1 to the left of the decimal point is missing. The 8 bits to the right of the MSB (sign bit) are the biased exponent. The exponent is –7, and the bias is 127, so 120 = 0111 1000 in binary. This permits comparing two floating point numbers bit-for-bit to determine their relation ($a < b$, $a = b$, $a > b$).

Notice that comparison is performed by subtracting ***b*** from ***a*** and using integer, rather than ordinal, overflow evaluation. A result of zero means ***equal***; an overflow means ***less than***, and neither means ***greater than***.

But, there is more to the story. There are at least 4 exceptional values not defined by the above. The value zero is always represented as zero in each bit position (sign = 0, biased exponent = 0, and mantissa = 0). This makes sense, and you might have already jumped to this assumption. However, three more exceptional values exist: +∞, –∞, and NaN (Not a Number). These special cases are represented with a biased exponent = 255. Infinity if mantissa = 0, and NaN otherwise. When an overflow is produced, an ∞ could be substituted. Likewise, zero could be substituted for underflow. A divide by zero could result in an ∞, or a NaN. Any operation involving a NaN, however, produces a NaN as a result. For more yet, such as the option of representing a –0, Google "single precision" (or IEEE 754).

A Quick Review of Arithmetic Concepts

ordinals are the natural counting numbers: 0 … ***n***.

integers include ordinals and the negative numbers: -***n*** … -1.

rational numbers include all ***p***/***q***. When ***p***/***q*** is computed to be a string of digits in any base, one or more of the final digits repeat forever (Note: some numbers have final repeating 0 digits. These are not written, but they are implied).

irrational numbers are all of the infinite strings of digits in any base whose final digits are not a set of repeated digits.

real numbers include both rational and irrational numbers.

imaginary numbers = (x, iy), where i = $-1^{½}$.

dimensions = 0, 1, 2, 3, 4 (in my opinion a 5^{th} dimension is unnecessary to describe physical reality):

0 dimension = 0 | 1 (present or absent);
1 dimension = {–***n*** … +***n***};
2 dimensions = {–***n*** … +***n***} × {–***m*** … +***m***};
3 dimensions = {–***p*** … +***p***} × {–***q*** … +***q***} × {–***r*** … +***r***};
4 dimensions = {–***a*** … +***a***} × {–***b*** … +***b***} × {–***c*** … +***c***} × {–***d*** … +***d***}.

Why do I stop dimensions at 4? Google: Galois → "symmetric groups" for a hint of why I believe this to be the case for the arithmetic necessary to model and describe the natural world.

What is Science?

Hard science is a collection of models. Models must enable predictions. A model that makes no predictions is useless. Predictions without models are not scientific. As predictions made by a model are repeatedly observed in nature, a model becomes more and more accepted as a part of science. No amount of confirming evidence can prove the "truth" of a model. However, a single contradiction can prove a model false. There are other names for the models of science. They may be called theories, equations, or algorithms.

Soft science is a collection of observations and explanations that rely on statistics to describe them. When very large numbers are involved, the model may be very accurate. When small numbers are involved, the statistical significance of the result may need to be given. When the numbers are very small, the model breaks down and a different model is needed. Another way to see the difference between hard and soft science is that the models of hard science make predictions with an outcome that is certain, while the statistics of a soft science only approach certainty as the numbers approach infinity. Physics is considered a hard science, but it has two subdisciplines: thermodynamics and quantum mechanics. Both of these rely on statistical models to explain observations. An underlying mechanism is either unknown or too complex to describe. When a system follows a set of rules and goes through a very large number of state changes, the best simple model is almost always statistical in nature. This applies as much to explaining animal behavior as it does to predict the behavior of particles and photons.

The scientific attitude is the heart of modern science. Its central theme is the belief that one should accept neither tradition nor authority as a basis for the truth, but accept only that which is based on observations and conclusions. Science requires a skeptical, but open, mind. Truth is not accepted because it is compelling, or confirmed by anecdotes. Instead, judgment is reserved until sufficient evidence is collected and a model or mechanism is clearly understood. Confirming evidence should not be sought any harder than disconfirming evidence. No amount of confirming evidence is sufficient for proof, but one undeniable case of disconfirming evidence is enough for disproof.

How do scientists "do" science? Their objective is to build and refine models. Their methods should involve new data or changes to the current model. The following six-step process is meant as an example of the activities that further science.

1. **Observation:** Something is observed that isn't included in the current model, or a defect or simplification of the current model is brought to light.
2. **Deduction:** A reasoning process or leap of imagination is used to connect the data into a new pattern, model, or explanation.
3. **Hypothesis:** A prediction is made in the form of a statement involving the new model or theory, and it is made in terms that are falsifiable.
4. **Experiment:** A carefully controlled procedure is performed to test the hypothesis. Data that could potentially confirm or disconfirm the new model are gathered impartially.
5. **Conclusion:** Steps 1 through 4 are written up and published for review by other scientists.
6. **Verification:** Other scientists, following the logic of the observation, deduction, and hypothesis, repeat the experiment and arrive at similar conclusions. If successful, the model becomes a part of current science.

A lot of the activity of scientists is involved in publication, reviewing other scientist's publications, and attempting to replicate other experiments, or test other's conclusions in different ways. However, these activities are not what drives science. The driving force is curiosity. Ignorance is an itch that humans feel a need to scratch. It simply feels good to make sense out of the unknown. The story goes that it was only after he noticed he was no longer having fun, that Richard Feynman resolved to make fun one of his objectives. Shortly after that he found himself interested in the spin of electrons, and he followed that path for the fun and challenge of it. This led him to the development of QED and his Nobel prize.

Science can get entangled with politics. A scientist would no doubt consider the world a better place if science were never touched by politics. Politicians, on the other hand, may view science and scientists as tools, to be financed for a purpose and harvested when ripe. But, what about politics within science? My dictionary gives as its final definition of politics, “the total complex of relations between men in society.” It’s an old dictionary—today it surely would have said “people in society”—but that adds political correctness, a “branch” of politics developed after that dictionary was published..

Within science, at any point in history, there is always the “received” view, or the “standard” model. Changes take place over the course of time, and the models of science evolve. Change goes through three stages: First, a new idea is laughed at or ignored, some of these move on to be opposed by force, and eventually a few are accepted as obvious. All of these stages are political in nature, not scientific. They would be no part of science if science were practiced as it is professed; but science is carried out by scientists, and scientists are people. And, of course, most new ideas never make it past stage one, they die laughing. Only one set of ideas is ever accepted by a given scientist as obvious in the present, but the present unfolds into the future, and other ideas will become obvious as the ideas we hold today (and the scientists who hold them) fall by the wayside.

This book is not an attempt to record science, but it will allude to a number of the models of science. The terms hard and soft science have been defined above. Another term that needs to be defined is “fringe” science. This is science that collects data that appears to contain correlations just above the threshold of significance, but with no model or explanation of the data. Parapsychology fits this definition. New age health and certain systems of medicine also fit this definition. Intelligence testing comes close, but as a proper subset of the soft science of psychology, it shows potential.

The purpose of this section is to define a continuum from hard science to fringe science. Beyond the fringe there is only complete non-science (nonsense?). Any sensible statement or statements can be placed without too much difficulty into either this continuum, or into the continuum of fantasy and entertainment.

Truth

Above we have seen examples of how statements may be made. Statements may record rules, or they may follow rules and record other information. They may state facts, make predictions, direct actions, and record both truth and lies. Let’s say that truth (or Truth) is the body of all statements, feelings, and things that are true. What we call truth must conform to actual conditions, reality, or to known facts. Some adjectives that apply to the word ***true*** are: real, genuine, authentic, sincere, loyal, faithful, steadfast, proper, and accurate. Something is true to a pattern to the extent that it repeatedly conforms to that pattern.

Human beings have a deep need to feel themselves in touch with the truth. This need may even have a genetic basis. All human cultures reflect this need by attempting to teach us what is true, and to “tell the truth.” Often, truth is derived from what a group agrees upon. But, it should also be the case that you can verify the truth with your own senses. Truth originates from four sources: sensory perception, reason, authority, and emotion. Truth is simple, consistent, and what causes the least discord. Truth is all of these things—each to some extent—but each is only an aspect of truth.

Certain kinds of statements are rigorous enough to be submitted to a test. Some of these (but not all, according to Gödel’s Incompleteness Theorem) can be proved to be true or false. Most statements (such as the accounts we read in history books) are not rigorous enough to be put to the test. The opportunity to test many statements has vanished, and it will never reappear. The truth of many an event is tested in court, and we get the O J version of the truth. Or it becomes part of political discourse, and we get “alternative facts.”

It’s true—true statements can be made and proved within the bounds of a sufficiently rigorous domain. But can truth, in general, be stated and communicated? Let’s say we wished to describe the contents of a bottle. We might say that it contained so many grams of water at a certain temperature and pressure. Would this be the truth? It wouldn’t be the *whole* truth, because the contents are an exact number of water molecules, each with its own position and motion at any given instant. There will also be some impurities. The value we give for any measurement is only an approximation—it can never be exact. Almost every concept we have, when we try to describe it or use it to describe something else, turns out to be too complex for any simple statement to be the *whole* truth. Statements of the *whole* truth probably need to be just as complex as their subjects. Only a very few statements can express the whole truth; most statements are simply approximations of a truth that we lack the capability to express exactly.

Subjective truth is simply what we believe it to be. There are as many subjective versions of the truth as there are people to express it. Is there any such thing as a supreme version? An ultimate truth? An absolute truth? A transcendent Truth, perhaps the equivalent of God? Could this Truth be anything less than a complete record and understanding of the Cosmos over all eternity? I don’t see how this would be possible—but others might see it differently.

If truth is subjective, then it’s simply another name for our own opinions. Is truth anything more than what we believe in and render opinions about? If truth is a collection of beliefs, it is certainly true that all such collections are not the same. One collection might be what a given individual believes in, but a very different collection might be the collection of all models and statements that would be accepted as true by a majority of scientists (restricting them each to their own professional fields). A completely different version of the truth would emarge from the collective beliefs of the members of a particular religious faith. It is easy to find different people making contradictory claims about the truth. In all such cases, one of them is more right than the other, or both of them may be so wrong as to not even matter.

Truths that are shared (or felt to be shared) with others are the most powerful. This is what our predisposition for the truth is all about. It is a major part of the bonding process that ties people together into social units (couples, families, groups, and nations). The feeling of epiphany is often acquired in just this way. What kind of game plan should we adopt to improve the collection of things that we believe to be true? And do we even place a value on an effort of this sort?

What sources can we rely on for the truth? Should we include our memories in that body of things we believe to be the truth? Should we include our feelings? It's a well-known fact that the human memory of episodes and events tends to be quite faulty (of course that's only true of other people's memory, it never seems true for our own!). So, it might be wise to make some kind of allowance. As much as people complain about their memory as they get older, have you ever noticed that nobody ever complains about their judgment?

Maybe it would help if we tried to identify several categories of truth. Let's consider the following: definitions, postulates, axioms, statements amenable to proof or disproof, statements consistent with a set of facts, statements consistent with a set of claims, anecdotes or scenarios, and finally truths whose support is based only on trust, or faith, in their source. One continuum of truth ends here. Beyond this continuum are all truths or any Truth that cannot be formulated as some kind of statement. If we can sort all of what we believe into these categories, then we can much more easily see how to deal with our beliefs.

Definition is the simplest kind of truth: Saying so, makes it so. Definitions have to be accepted and understood by the users of a language. Languages are only effective to the extent that this is true. A language is nothing more than a set of conventions, among which are the words of the language and their definitions (statements that relate the words, show how they are used, and describe one word in terms of others). Postulates are statements that are assumed to be true for the purpose of constructing a formal argument or proof. An axiom is a statement that is widely held to be true, or considered to be self-evident, but whose truth is too fundamental or subtle to be capable of proof, and more general than a mere postulate.

The above kinds of truth give us a working knowledge of our language in which we can make other more complex statements. There are at least four kinds of statements: One is the question, another is a command or direction, another is an exclamation, and the remainder are statements (true, fictional, or hypothetical). It takes contemplation and a disciplined mind to review all the definitions, postulates, and axioms that one has accepted and come to believe over the years. No truth can rest on a faulty understanding or recollection of any of these, and it's always possible—even easy—to get our facts wrong. If no fault lies in our use and understanding of the relevant definitions, postulates, and axioms then we should be able to reason the truth of any statement amenable to proof.

Statements amenable to falsification, however, have an even more important place. The problem with statements amenable to proof is that there are too few of them, and they are restricted by Gödel's Incompleteness Theorem. Statements amenable to falsification properly include all of science. Ideally, they should form the backbone, or core, of a person's beliefs. Here, the burden is to properly identify a statement as being in this category, attempt to falsify it, and consider it to be true only to the extent that it passes your tests.

Truth exists in three realms: Nature itself is the first realm. The second is the realm of truth based on observations of nature (accepted, claimed, or supported by evidence). The third realm involves extending the truth beyond nature by selecting only one of several interpretations of the facts, using the support of a model or scenario rather than reality itself, or by substituting trust, faith, or authority for a truth, instead of reasoning it out from our own experience.

The foundations of science are models. The foundations of proper conduct and behavior are notions about right and wrong. Judgments are based on the notions of good and bad. Punishment is based on the belief of innocence or guilt. These aspects of the truth can impact us where we live, not just where we think. It takes a tremendous amount of effort, usually at an early age, to organize our beliefs and opinions according to a set of principles. A very large majority of people are simply too lazy to do this. If you are one of the minority who feels the need to make this effort, then you need to begin with a set of axioms. Your axioms may differ, but the following statements have evolved from many of the great minds of history. You could do much worse than to begin from the following propositions.

The universe is governed by natural law. This law is capable of being understood and approximated by statements that can be verified in some fashion or to some degree (statements potentially falsifiable). Science is a culture's collection of such statements. The laws, principles, and theorems (the models) of science operate on single levels according to the classical principles of reductionism. But, a science of everything is layered like the skin of an onion. When very large numbers of simple entities interact, complex properties can emerge. These emergent properties, like the next layer of the onion, may form a whole new scientific discipline. Reductionism is inadequate to bridge the gap between such layers, but each layer is comprehensible in its own right, each has its own truths to be discovered, and the scientific method is the appropriate way to make those discoveries.

Many people have a limited understanding of a few scientific truths and extend their beliefs from there. Any similarity between these extensions and the scientific truth is almost accidental. This is the case when people operate with a very different, or perhaps more poorly understood, set of axioms. Scientists, too, may attempt to extend the truth from current models. This occurs when the map is contemplated instead of the territory. For example, there is a very rigorous map for quantum mechanics, and it works almost perfectly over a certain range of space and time. However, there is disagreement over some of the extensions or conclusions drawn from this map.

Other areas of the truth properly belong beyond the boundaries of science. Philosophy and religion attempt to give us a set of axioms and principles to guide our development and behavior. Many of the attitudes and methods of science can still be made to apply, but it is far more tricky to do so. Here we run into cultural relativism, and the continua of right and wrong, good and bad, and innocent or guilty. Here we can go beyond the truth, if we aren't careful, to judgment and prejudice.

Absolute ***Truth*** is a different concept. It refers to statements grounded in the transcendent, the supernatural, or naive belief, and for which a proof is either absurd, or cannot be properly constructed from a set of facts. A non-standard collection of definitions, axioms, and postulates may form the ***Truth*** for an individual or group. I'm sure that if we compared the most commonly accepted set of definitions, axioms, and postulates over all of humanity to those of intelligent life elsewhere in the universe, we would also find considerable differences.

People seek the Truth along one of four paths. The first path is the direct observation of nature. The second is to come into contact with a person, an authority figure perhaps, who is believed to know the Truth. The third way is to read a book that is believed to contain the Truth. The fourth way is to acquire the Truth directly through divine revelation. I believe the first of these paths ***can*** lead to Truth, and the other three either provide assistance in following the first path, or they lead into a maze that you may never successfully exit. Let the seeker beware!

People seeking the Truth should ask certain questions about it—questions whose answers might guide their search if they got comfortable with them. For example, "What kinds of things can reveal the Truth?" Perhaps we simply make up the Truth. If the human race did not exist, would Truth exist, nonetheless?

Human needs are vacuums that human society evolves ways to fill. Three needs we have are to know the ***truth*** about the world, the ***meaning*** of events, and our ***purpose*** in life. We have evolved art, religion, and science to fill these needs. Religion communicates itself in the form of scripture, ceremony, music, song, and dance. These are all art forms, and art evolves in non-religious contexts as well. But has art evolved to embody and communicate the Truth, or has it evolved simply to fulfill human needs? The answer is obvious, given the selection process that drives evolution: Culture evolves to be consistent with human needs, not necessarily to reflect Truth. Science is the only one of the three that was evolved to intentionally reflect the Truth.

Models may attempt to explain the truth, but models are only effective to some degree—they are not the Truth. Art may attempt to reflect the truth, as may stories, narrative, and the recounting of history. But, these are not the Truth. They are products of people. You can experience reality. At the moment you sense it, you are in contact with the Truth. When you attempt to explain it, you translate Truth into a model or description that has truth and consistency only within its own context. An expression in a language loses any direct connection to the Truth.

What words express can be true or false in a mathematical sense, or in the practical sense they can be effective, irrelevant, or indeterminate. We must make the distinction between a statement that is true in a mathematical or practical sense, and the Truth. I claim that what I say here is a true accounting of my beliefs, and not a pack of lies (some might beg to differ).

Whether or not a scientific approach is part of a person's attempt to get at the truth, philosophy and religion almost always have a role. Sometimes, in opposition to science, religion attempts to explain how things are, but these areas more properly exist to tell us what to do—how to behave. Still, they proceed from axioms, and they make statements that profess to be the truth. The following should be taken as analogous to the fundamentalist approach to the truth in any culture or historic setting: Christian fundamentalists believe in the literal word of the Bible. It must be axiomatic to them that every instance of interpreting the Bible occurred under the direct guidance of God. This begins with all the original authors of the Bible. And it includes, of course, those early scribes who translated the Bible from the language in which it was originally written. It must also include the later scribes who copied and updated the Bible over the intervening generations to the present.

Since many believers don't actually read all of the Bible themselves, it must also include the church leaders who preach from it. They interpret it for the rest of the flock. If, at each turn, God were not in complete control, man's judgment would have entered in, and today's Bible would not be the literal word of God. Conflicting ***versions*** of the Bible do exist. Different ***parts*** of the Bible conflict with each other. Truths have been discovered since the Bible was first written. Only hubris or ignorance could lead anyone to believe that God has guided ***them*** to the source of Absolute Truth and actually ***misled*** the rest of us.

Different sources of the truth can be in clear conflict. How do you choose? The farther down the list (of definitions, axioms, the provable, the verifiable, or simple consistency with facts, claims, anecdotes, scenarios, or trust) a particular truth is supported, the less solid the ground it's on. Even the first items on the list (definitions, axioms, etc.) need to be examined in an iterative process. You always need to update your definitions and axioms. And, you should be very cautious of truths whose support is further down in the list.

One final observation about truth: truth is often the hostage of power and authority. It takes courage to "speak truth to power." Attempts to dictate or control the truth occur all the time. A new conspiracy to alter truth is fairly obvious. But, beware the conspiracies dictating truth that ***evolved*** over generations!

Proof

A proof of existence consists of demonstrating a single example of something. Except in special cases, a proof of non-existence is impossible. Some statements assert a relationship or predict an outcome. These are more tricky than statements about existence. These statements may involve single instances, sets with large numbers of instances, or even sets with an infinite number of instances. A proof of truth has to consider every instance. A proof that a statement is false has only to find a single example. The assumption is made that you can prove a statement false by proving that its opposite is true. Likewise, you can prove a statement is true by proving its opposite is false. But these methods of proof are subject to the weakness of the assumption that there is no middle ground between the true and the false. This type of proof implies that everything is provably true or false, and that itself, is an assertion that has been proven false. It is Gödel's Incompleteness Theorem. It states, in essence, that statements in any adequate language cannot always be proven either true or false. There are always three possibilities, not just two: A statement is provably true, a statement is provably false, or a statement is not known to be either. In this latter category there are actually ***four*** types of statements: Those that will later be proven true, those that will later be proven false, those that are suspected or proven to be undecidable, and those with an insufficient basis to be decided. A great deal of what we share with one another falls into this last category. Here, we will contemplate only statements that are amenable to proof.

Proof involves accepting a set of axioms, rules, and a language for making statements about existential entities (things that exist—***including symbols***—and the concepts they represent) subject to accepted axioms and rules. Thus, we may make the statement that 5 = 4 + 1. Our study of arithmetic allows us to prove that this is true. We simply start with 4, apply the add one operation that we defined (above), and get the answer 5.

General statements about numbers often exclude the numbers zero and one, because these are identity numbers. The statement $\boldsymbol{a} - \boldsymbol{b} = 0$ implies that $\boldsymbol{a} = \boldsymbol{b}$. The statement $\boldsymbol{a} / \boldsymbol{b} = 1$ also implies that $\boldsymbol{a} = \boldsymbol{b}$. Likewise, $\boldsymbol{a} \times 1 = \boldsymbol{a}$ and $\boldsymbol{a} \times 0 = 0$ for all $\boldsymbol{a}$. When a statement involving prime numbers or factors is made, it may or may not include the ordinals zero and one.

Arithmetic, as I have defined it, only allows statements that can be proved true or false. Its semantics are limited. Logic and mathematics allow statements with a more complex semantics, and not all of the statements possible in these languages can be proved true or false. Some are undecidable. An example is Goldbach's Conjecture (any even number is the sum of two primes, if you don't nitpick about 0 being even, or 1 being prime). For example, 8 is the sum of the primes, 3 and 5. There are 2 solutions for 10 (3 and 7, and 5 and 5). So far, no one has been able to find an example proving Goldbach's Conjecture false, but no one has found a proof of it either. The likely reason this and similar statements are not provable is that they describe an infinite set of problems with no general method to solve those problems. Another way to make an unprovable statement is to use self-reference, or imply infinite recursion (such as "this statement is false").

It is easy to devise an algorithm to find two prime numbers (if they exist) that sum to an even number. The catch is that this algorithm doesn't guarantee a positive result. If it fails, it is proof that no two primes exist that add to the number given. But so far it has never failed. So far an infinite number of numbers remain to be tested. Therefore, this approach cannot produce a proof.

Another example is the ***fundamental theorem of arithmetic***—the statement that "all ordinals have a unique set of prime factors (excluding the numbers zero and one)." This statement can be proved. A factor is one of the two operands in a multiply operation. A prime has only itself and one as factors. The statement is trivially true for the number two. If the statement is true for the number $\boldsymbol{n}$, then if we can prove it true for the number $\boldsymbol{n}+1$, it must be true for the infinite set of all ordinals. For $\boldsymbol{n}+1$ there are two possibilities: Either $\boldsymbol{n}+1 = \boldsymbol{a} \times \boldsymbol{b}$, or $\boldsymbol{n}+1$ is a prime. In the former case, its unique factors are the union of the two unique sets for $\boldsymbol{a}$ and $\boldsymbol{b}$, which have already been included in the proof, since both of these ordinals are less than $\boldsymbol{n}$. In the latter case, its unique set of factors is itself.

A proof involves making a finite number of substitutions for statements or parts of a statement according to a set of rules such that the final substitutions are all axioms. Notice the similarity of this to the rule substitution in parsing a statement using a grammar, where axioms are equivalent to terminals.

Axioms may be incorrect. The rules may be incorrect. Statements may be incapable of expression in any axiomatic form, or testable with an acceptable or repeatable result. All of these observations must become second nature when making or attempting to understand any statement. And much of what we do is observing and trying to understand!

Sometimes we model reality with analogy. Then, we go on to use our model to predict or describe other aspects of reality. This allows us to confuse the map with the territory. Just because our map ***says*** that the next city on our current route is X, it doesn't mean our map is correct, it just means that ***if it is***, then we should arrive at X. That's the nature of models. They may involve assumptions we cannot prove. At the very least, our models should fit all the facts. The better models are also based on compelling analogies. The less satisfactory ones have no underlying structure or analogy.

Thus, there are true statements that we can prove, and other possibly true statements that have not been proven false. Some statements are true or false by definition, such as there are no square circles, or there are no married bachelors. The truth of some statements is undecidable due to self-reference or some other defect, such as "The Barber of Seville shaves all men who do not shave themselves" (so, who shaves the barber?).

If you have a good imagination, picture the following: An eight-by-eight chess board and a set of dominos. If you can't picture these, perhaps you can lay your hands on the real things to augment your imagination. Now, lay (or imagine laying) one domino on all the adjacent squares of the board. Since the board has 64 squares it should take 32 dominoes to do the job. The fact that this is possible is obvious. Anyone should get it right on the first try. Now, try to lay out the dominoes such that the lower left and upper right corner squares are left uncovered. This covers two fewer squares, so it should take one less domino. Try it. Can't do it? It seems impossible. But can you ***prove*** that it's impossible?

There are at least two ways to prove that this task is impossible. One way is to program a computer to check every possible arrangement of dominos. This is the "brute force" method of proof. The other way is more elegant. It consists of noticing one of the features of a chess board: The fact that adjacent squares are of opposite colors, and that squares of the same color are always diagonal. Combine this with the fact that a domino must always cover adjacent squares and never diagonals and you have a proof. Notice that the opposite diagonal squares are the same color. Notice that each domino must always cover two squares of opposite colors. You simply can't cover two differently colored squares with each domino and have two squares of the same color left over if the board has the same number of squares of each color.

What does this illustrate in a larger sense? For one thing, it illustrates the varied nature of proofs. Statements may affirm something. Statements may be provably true or false, they may be suspected to be true or false, or their status may simply be unknown. A person's "belief structure" is the set of statements that a person accepts as true. Most people accept statements as true that they themselves can't prove to be true, but they believe that others have done so. Many people accept statements as true knowing they can't be proven, or they accept inadequate proofs, or they are not concerned with the criteria for proof; they believe things based on other criteria.

Statements that are undecidable fall into at least three categories: Those which generalize about an infinite set of things, those which are self-referential, and those involving things about which we have insufficient evidence. The statement several paragraphs back, "every even number is the sum of two primes," is an example of a generalization about an infinite set of things. The statement "The Barber of Seville shaves every man who does not shave himself" is an example of the second type. The question is, who shaves the Barber? Since the first statement ought to answer that question, but contradicts itself when doing so, it appears to make the truth of the statement undecidable. My own answer to the paradox is that the Barber of Seville is obviously a woman, but maybe that's outside the box. However, the general point remains. And, contentions that lack sufficient evidence are why we have courtrooms, juries, judges, and lawyers.

Complex Models

The 1st model introduced above was a count. This model (a number) allowed us to capture something about reality, namely the number of things in a group. Next, arithmetic expressions were introduced. Then, metalanguages. And, all of these are simple ways to express models. More complex models involve statistics, trigonometry, and other forms of mathematics. Some examples will further illustrate mathematical models. These are only intended to convey "flavor" — they are not as rigorous or complete as they would be in an actual course of study.

What are the odds?

Earlier, I said that making and executing a choice is a ***skill*** that equates to free will. The exercise of free will is important. Often a choice must be made in the face of uncertainty. Scientific models exist to make predictions. When a number of things ***may*** happen, how are we to know which is the most likely? The chances of an outcome depend on the mechanism or system that selects one outcome from all the possible outcomes. If one numbered ticket stub is selected at random from a hundred, your stub has one chance in a hundred of being drawn, all things being equal. If your stub is larger, or is folded, or was on top, it might change the "all things being equal" condition, and alter your chances from 1/100 to somewhat more or less than that. It depends on the actual selection process.

There are two ways to express the chances of something: odds and probability. When you draw a card from a fair deck, the ***probability*** that it is the ace of spades is 1/52 (one time ***out of*** 52). The ***odds*** for this happening are 1:51 (one chance that it will happen ***to*** 51 chances that it won't happen). Payout odds (which were invented to further confuse the mathematically challenged) might state this as 51:1, letting you know that they'll add $51 to your $1 bet if you win, or rake away your $1 bet if you lose. Hopefully, from this example you can follow the math and convert odds to probability, or vice versa.

When chance is involved it means that there are factors or unknowns that we cannot take into account. The process that selects can be very complex. When twelve horses run a race, the winner is the horse that first crosses the finish line. This is a function of the fitness of the horse, the weight and behavior of its jockey, how good a start the horse was able to make, the conditions of the track over which it ran as compared to the other horses, the way the horses interacted with each other, and the rules of the race. The combination of these factors is partly predictable and partly chance. The final outcome can't be predicted with certainty. This is true for the outcome of any complex or chaotic selection process.

One of the best predictors of tomorrow's weather is today's weather; one of the best predictors of tomorrow's horse race is today's horse race. It takes a great deal of data and calculation to improve these predictions.

Probability is a funny thing. An accurate estimate depends upon a set of assumptions being true. In many real applications of probability, the assumptions aren't true. Other times there may be information available that should be used to adjust the assumptions. The default assumption in probability theory is that each possible outcome of a trial is just as likely as any other. When a number of trials are performed, each trial may have a dependence upon the previous trial, or it may be completely independent. We often encounter only a single trial, but we must make a decision as though the odds were calculated on the basis of many trials. Sometimes we know the exact payoff such as betting on a horse race, but other times the expected return is more difficult to calculate. Mistakes can be made in figuring out what calculation to make, or in the estimate of the odds, or in the true value of the outcome.

Consider a scenario: A king leads you to a room with three vaults behind closed doors. He tells you that two of the vaults are empty and one is filled with treasure. He invites you to pick a door and have whatever is behind it. After making your choice, but before the door is opened, the king opens one of the other two doors and reveals an empty vault. He now says that you may have ***all*** of the treasure behind the door you have already picked, or ***only*** 75% of the treasure behind the other closed door. Which is the better deal? Did the odds change when the empty vault was revealed?

The better deal is the higher expected return based on the true odds, using all the information at this point. The 1st door still has one chance in three of revealing the treasure. That means there is still a two out of three chance that the treasure is behind one of the other two doors. But, the King has just shown you an empty vault. The odds have gone from 2:1 that the treasure was behind one of the other ***two*** doors to 2:1 that it is behind a ***single*** door. Your choice is now between 100% of a 1/3 chance (by sticking with your original choice), or 75% of a 2/3 chance (by switching to the other closed door). Do the math—your expected return improves from .33 to .50 if you ***switch***.

You should always attempt to balance the risk you can afford to take with the expected return for taking that risk. Choice enters the equation in many different ways. Some choices avoid risk. Other choices improve the expected return on a risk. The risk/reward ratio is a factor. In many situations you stand to lose nothing (but the time it takes to participate), or only the value of a bet. But in some situations, such a driving a car, the reward is moderate and the risk, though very small, is death. Choices are difficult to make even when you understand the math. The better you can understand the selection process, the more effective you can be in influencing it, or computing the probability more accurately.

Would you rather draw straws where a single short straw says you get nothing and four long straws all say you get $100, or just take a sure $75? Simple arithmetic would say that the expected return for drawing straws is $80 (four chances of $100 plus one chance of nothing, is a total of $400 divided by 5). $80 is better than $75, isn't it? Or, is the $5 difference too small to justify taking the risk? What if the difference were larger? What if you had four chances out of five to get $1000? Would you prefer that to a sure $75? This shows the difference between playing only once, and playing a large number of times. There will always be risk seekers and the risk averse.

A family has six children, three boys and three girls. In which order were they most likely born? BBBGGG or GBBGBG. Unless you know otherwise, either sequence is as likely as the other. In fact, if it had not been stated that there were 3 boys and 3 girls, the sequence GGGGGG would be as likely as BGBGGB.

In a standard English dictionary, does the letter ***k*** appear more often as the ***first*** letter of a word, or as the ***third*** letter?

This is a famous example of thinking we know something that is just wrong. We tend to be able to recall words that start with a particular letter much better than ones with a letter in some other position. What we can quickly and easily recall tends to form our "statistical sample" and leads to our estimate of the probabilities. Our recall is more affected by recent news than long term experience, by moods more than knowledge, and by heuristics and shortcuts more than logic and calculation. It turns out that ***three times*** as many words have ***k*** as the third letter than begin with it!

Two hundred women were surveyed: 180 were housewives and 20 were lawyers. Joan, selected at random from among these women, is known to be a feminist. Is Joan more likely to be a housewife or a lawyer?

Here you are asked to make a choice based on some sampling data and some specific information. Our tendency is always to give too much weight to specific facts and not enough to underlying probabilities. We might figure there would be a greater chance of Joan being a feminist if she were a lawyer than if she were a housewife, but we don't know the percentage of feminists among either group. If no housewives were feminists, Joan would obviously have to be a lawyer. If all lawyers were feminists (20 feminists out of our sample of 200), and only 1 housewife out of 9 were a feminist (20 out of 180), the probability would be 50%. The chances are that fewer than all of the lawyers, and more than 20 of the housewives are feminists, so the odds are probably greater than even that Joan is a housewife. But suppose we asked what percent of our sample are ***housewives*** given that they are ***feminists***? This is more complicated.

Bayes' Theorem…

is the model for conditional probability. It shows the relationship between 4 pieces of information. These pieces are: a ***base*** probability, P(B); the probability of a ***condition***, P(C); the probability of the ***base given the condition***, P(B | C); and the probability of the ***condition given the base***, P(C | B). The formula for this is:

$$P(B) \times P(C \mid B) = P(C) \times P(B \mid C)$$

If you know any 3 pieces of the above information, you can solve for the 4th. How does this apply to the odds of Joan being a housewife? First, note that being a housewife means (here) that the individual has no other job, so the two sets don't overlap. Here, we have housewives (90%), feminist given housewife (44%), and the probability of being a feminist (50%). We want to solve for the probability (in our sample) of housewife given feminist. The data on women lawyers was useful to compute some of these values, but it doesn't go directly into the equations.

The required information, in Bayes' formula, can be named as follows:

$$P(h) \times P(f \mid h) = P(f) \times P(h \mid f)$$

Notice the assignments: h = housewife = Base, and f = feminist = Condition. Now, follow the data assignments to each term in the formula:

$$0.90 \times 0.44 = 0.50 \times P(h \mid f)$$

Doing the math gives us:

$$P(h \mid f) = 0.79$$

Thus, even though Joan is a feminist, there is an almost 4/5 chance that she is a housewife and not a lawyer.

For more examples, go to: https://en.wikipedia.org/wiki/Bayes%27_theorem#Examples.

To conclude this section, ***What are the Odds?***, consider the insurance industry. Everyone runs the risk of various losses. Some of us cannot afford to incur a catastrophic loss, and are willing to pay a price to be compensated if we do. The odds of various catastrophic losses over an entire population (the customers of an insurance company, for example) can be calculated. This enables the insurance industry to set policy rates and make a profit. Each customer knows that they are paying more for their insurance than they can expect to receive back, but it may be worth it to them to avoid the risk. In fact, the luckiest ones are those who never receive anything back! On the other hand, if you can afford the loss and are less likely to incur it, the insurance might not be worth it. This is especially true for life insurance and long-term care. Both of these are expensive, and with an estate over a certain size, there are many instances where they are not worth the premiums.

How does the weather work?

Two huge engines drive the weather. One is the sun heating up the air, evaporating moisture into it, and causing it to circulate in convection currents that take masses of air up from the surface of the Earth and back down again as they are heated and cooled. The other driving force is the rotation of the Earth. This rotation, usually not apparent to us, is the basis of the Coriolis force, which, as we shall see, can have quite spectacular effects.

Imagine a sphere in empty space. Let's say it's about 25,000 miles around (the same size as our Earth, in fact). And, we'll cover it with a smooth linoleum floor (no oceans or land at this point). But we will give it a spin, so that it rotates once every twenty-four hours. And, to better see this rotation, we'll put a big light in space. Say, 93,000,000 miles away.

Now, because our sphere rotates, it has an axis with two poles. Arbitrarily, we'll call one of these the North Pole, and the other the South Pole. Imagine standing exactly at the North Pole with your arms outstretched. Standing on this immense linoleum plain, you are slowly rotated counterclockwise. Once every twenty-four hours your right arm moves slowly forward, your left arm moves backward, and you make a complete turn.

An ice skater goes into a spin by pulling in her arms and legs after she starts a turn. You could do the same. If there were a pivot point under your feet with no friction and you pulled in your arms your rate of spin would increase. Now, you will rotate once in fewer than twenty-four hours. This is due to the conservation of the angular momentum you already had.

Imagine a bunch of huge merry-go-rounds each a mile in diameter and scattered all around this linoleum landscape at our North Pole. Imagine looking down on all these merry-go-rounds. Each one appears perfectly still, but it nevertheless rotates once every 24 hours along with the rest of our great sphere. Now imagine thousands of people getting onto the merry-go-rounds and rushing into their centers. What happens? The merry-go-rounds (and the people), already spinning once each 24 hours now begin to spin faster (because the angular momentum of the people is concentrated at the centers of the merry-go-rounds). They spin counterclockwise because that's the way our sphere is already spinning. The more the angular momentum already possessed by these systems (of merry-go-rounds and people) is concentrated toward the axis of each merry-go-round, the more quickly each merry-go-round spins.

Now, let's go from this perfect sphere back to the familiar but chaotic world around us. When water in a bathtub moves toward the drain, it tends to spin in a counterclockwise direction. This is in the Northern Hemisphere where things are already spinning that way (in the Southern Hemisphere, it's the opposite). The same thing happens to a mass of air. If air is warmed it expands; if cooled it contracts. Imagine a huge mass of air being cooled. It contracts. It moves inward. It was already rotating along with the rest of the Earth once every 24 hours. Now this rotation is amplified. This is why northern air masses always rotate counterclockwise around a region of low pressure. Likewise, if air rises, it is replaced by air flowing in from the surrounding area near the ground. The spin picked up by air rising in a column can be most dramatic, because the column is often very thin, concentrating the spin very tightly. We call such a column a tornado.

Living in the west, I've often had the chance to see a dust devil. This occurs when powdery dirt is heated by the sun, and the air over an area rises at some central point. The air around that point moves inward, and the spin of the Earth translates into a small tornado that picks up dust off of the ground. Once I happened to be on the equator and I saw an amorphous cloud of dust rising, and it struck me odd that was not a spinning vortex. Puzzled, I worked out that the spin of the Earth was at a right angle to the surface of the Earth, and therefore the Coriolis force was reduced to zero. There are no dust devils on the equator!

We know that the sun heats the air more at the equator than it does at the poles. All other things being equal, this causes the warm air to rise at the equator and flow north and south at high altitude towards the poles. It then cools, settles, and returns to the equator by being pushed along the ground. Both the air and the ground at the equator are moving over a thousand miles an hour (the equator is 25,000 miles around and it rotates once every 24 hours; 25,000 divided by 24 is a little over 1000). What happens to this speed as the air moves toward the North Pole? It becomes the jet stream. The air, already moving with the Earth at 1000 miles an hour begins to move over regions of the Earth that are moving slower and slower, until the north pole is reached, and the ground isn't moving at all (it is simply rotating in place). Of course, a lot of this speed is lost to friction, but the effect is that the jet stream in the upper atmosphere (in the mid-latitudes) is seen as a current of air moving from west to east often exceeding two hundred miles an hour. This can cause quite a head or tail wind for aircraft flying west or east at high altitudes.

Now, what happens when the air cools in the polar regions, sinks to the ground, and is pushed back to the equator? The reverse occurs. As the air moves away from the pole it is going slower than the ground under it. The ground is moving to the east faster than the air, so the air appears to be moving to the west. This is what causes the trade winds that move (mostly across the oceans) from east to west. The regions closer to the poles, before these winds have a chance to build up, are known to sailors as the doldrums. As the winds finally begin to match speed with the rotating earth beneath them, between about thirty degrees north and south latitude, we have what sailors called the horse latitudes. Between these two were the trade routes for sailing ships driven by the trade winds.

The Chinook and Santa Ana winds are different forms of the jet stream and trade winds, respectively. The Chinook happens when the jet stream touches down over the Rocky Mountains. The Santa Ana forms across the deserts of the southwestern United States. The Chinook always comes from a westerly direction, and the Santa Ana from the east.

Other factors also contribute to the weather. Land warms in the sunlight more quickly than water. When dawn breaks, warm air begins to rise on shore, and this pulls in the cooler air from the sea causing a sea breeze in the morning. When the sun goes down, the air cools more quickly over the land, and the air blows back out to sea later in the day and during the night. When water picks up the energy from the sun it evaporates into the air. This energy is released when the water condenses as rain or snow.

Why does it rain or snow? Because when the warm air rises, the pressure on it, higher in the atmosphere, goes down. The temperature also drops. The lower temperature and pressure cause the moisture to condense, form tiny droplets that are first seen as clouds, and then aggregate into larger drops to fall as rain, sleet, or snow, depending on the temperature and the amount of water vapor in the air. To complicate the process further, when the clouds form, they shade the sun and cause temperatures on the ground to fall quickly and unevenly. Between clouds, of course, the sun shines through. It hits and warms the ground, and causes the air in contact with the ground to rise. Great convective cells can form along with great differences of temperature and pressure. Intense low-pressure regions can give rise to the swirling air of tornadoes on the local scale, and to hurricanes on a much larger scale. Great differences of temperature in convective cells can cause drops of rain to cycle high into the atmosphere until they freeze and become large enough to fall as hailstones.

The transport mechanisms of heat and global rotation move the atmosphere all over the world. In November, a lot of the moisture that evaporates into the air from the Pacific rises, moves north, and finds itself turning into the jet stream and moving east. As it moves first north, then east, and rises higher, snow falls on the Sierras, the Cascades, the LaSals, and the Rockies. This causes skiers to appear on the mountain slopes, sales of Chapstick to increase, and thoughts of snow and winter wonderlands to occupy people's minds.

All of the observations above (with the exception of the last few) go into the computer programs that model the weather. Very specific parts of these models can be expressed using arithmetic, but the equations for different areas over our vast planet have to be specialized. Likewise, the values fed into the variables of the model are always changing. Data is collected by thousands of weather stations. Weather models have been evolving for over 30 years—more and more lines of code have been added, and the computers that run these models have gotten several orders of magnitude faster, and yet we still can't say with certainty what tomorrow's weather will be!

Global Positioning

No matter where you go these days, with a navigation system it's possible to know exactly where you are. These systems use a large database that contains maps of all the roads and many of the landmarks of the area that can be navigated. Given a destination, these systems can give you turn-by-turn guidance. The maps and the guidance, however, are only a small part of the story. The real technology is the global positioning system itself.

So, what is a GPS receiver, and how does it work? A GPS receiver picks up the radio signals from a constellation of Global Positioning Satellites. Each satellite orbits the Earth about once every twelve hours. The Department of Defense has anywhere from twenty-four to thirty-two of these satellites in the sky at any given time. They are in semi-polar (or at least non-equatorial) orbits. They orbit about 11,000 miles overhead. Each one has a mass of about one ton and is about 17 feet across with its solar panels extended. Each one carries at least one computer, atomic clock, and radio transmitter. Most have working spares of all these devices aboard. Each satellite also has a receiver that listens only to the Department of Defense to tell it what to do. This gives the satellites periodic updates so that each satellite knows where it is, within a few feet, at any given nanosecond. However, the DoD can also direct the satellites to report their positions inaccurately. This keeps enemies from using the system for some nefarious purpose, such as sending over an ICBM with pinpoint accuracy, or on the other hand, us civilians from avoiding an underwater obstacle in a harbor (a DoD safety tradeoff).

A GPS receiver needs to acquire the signals of at least three, but preferably four, satellites. Given that there are at least twenty-four satellites in the sky, about half of them can be seen from any point on the ground; the other half are on the other side of the world. Maybe a third to a sixth of them will be high enough in the sky to provide a decent signal. This generally gives you four usable satellites at any given time and place.

Now that you have an overview of the system, there are three topics that will enable you to know how the whole thing works: What information is broadcast by the satellites? How do the receivers work? What geometry is involved?

Let's take the geometry first. Given that there are three or four satellites overhead, if we knew exactly how far they were away from us and exactly where they were located with respect to a point on the surface of the Earth, we could determine exactly where we were on the surface of the Earth, or with respect to our map data.

Let's start with the first satellite. All the points equi-distant from this satellite describe a sphere around it. We only care about the huge sphere that just touches our own position. Now, take the second satellite. Again, we consider a sphere around that satellite that just touches our position. Two spheres make a circle when they intersect. Our position is a point somewhere on this circle. Now, add the information from a third satellite. It defines another sphere that just touches our location. A sphere intersects a circle at two points. In this case, one of the points will generally be nearer the surface of the Earth than the other. Since our location is probably at a point between sea level and a few thousand feet, one of the points is likely to be an obvious choice. So, with only three satellites, the GPS receiver can make a good guess as to where we are. In fact, if the clock in the receiver is perfectly synchronized with the clocks in the satellites (which are kept in near perfect time with each other), it could have narrowed our location down to the regions around two points with just the information from three satellites. However, GPS receivers don't keep perfect time. So, you need the signal from a fourth satellite to resolve your location and altitude, and to synchronize the clock in your receiver.

Now, in a general way, we know what must be broadcast from the satellites and what the receivers do with it, but let's get a little more specific. All the satellites broadcast a digital signal on the same frequency and at low power. Each signal being broadcast contains three types of information: Timing information, identification information, and information about that satellite. Because the signals are so weak and all of them are broadcast on the same frequency, the GPS receiver has to "tune" into a single satellite. This is like trying to listen to a particular conversation in a crowded room with dozens of conversations going on at once. You can't do it. But, suppose your name is spoken? Now, your attention is drawn to that conversation. If your name keeps cropping up in one conversation, you will probably have no difficulty "tuning it in." This is how the receiver sorts out the babble coming from all the satellites at once, and from other sources of noise. It knows each satellite's "name" and each satellite repeats its own name quite often.

This is ingenious; it's worth going into further. The digital signal from each satellite is broadcast at a little over a gigahertz (a billion cycles every second). The "name" of each satellite is a unique string of 1023 bits. Each bit takes about one microsecond, or about 1000 cycles of the carrier frequency, to be broadcast—about a millisecond for the entire "name." It's by knowing in advance the exact content and construction of each satellite's "name" that a receiver can lock onto a satellite's signal. This enables it to discriminate between the signal from that satellite and the background noise, which includes the signals from all the other satellites. Thus, the receiver can identify the strongest satellite signals in its vicinity and "listen" to four different conversations, one at a time, rotating among them. More expensive receivers tune in more than 4 satellites, or listen to them concurrently.

Once it has identified the signals from three or four satellites, a receiver begins to gather data from them. It can take the receiver several minutes for it to collect all the data it needs. These data include an update on each satellite's position, its orbit, and the exact universal time at which that satellite "speaks" its own name. Once these facts have been registered, the receiver can use them in conjunction with the geometry of the situation and draw conclusions about its exact location every second or so. Since all of the necessary data can be retained, it doesn't take as long to acquire a position when the data are fresh as opposed to when the GPS is first activated after being unused for a while.

After signal acquisition, and after data collection, the receiver enters the mode of continuous update of its location. This mode requires exact timing. Think about the string of 1023 bits that "name" each satellite. Each bit takes about one millionth of a second to broadcast, and consists of about a thousand cycles of the carrier frequency. Once the receiver has acquired the exact universal time and knows the exact time that a particular "name" was broadcast, it can pinpoint the time that a particular cycle of the carrier wave was emitted by the satellite that broadcast it. The signal travels at the speed of light. The speed of light is about a billion feet per second, or one foot in a billionth of a second. This is the time it takes for each cycle of the broadcast, so if you can pinpoint the time a particular wave was emitted from a particular satellite, you can pinpoint within about one or two feet how far you are away from that satellite.

But, it's not quite that perfect. The clocks can be off by a nanosecond or two, so this adds about two feet of uncertainty to your position. The position of each satellite is only known within a couple of feet. The receiver adds another four feet, or so, to the error. Atmospheric conditions add about twelve feet. And the Department of Defense adds about 25 feet of its own. Under most conditions, the geometry also dilutes the precision by a factor of 4 to 6, so the total error can be as much as 270 feet, or (2 + 2 + 4 + 12 + 25) x 6. However, by averaging successive computations, a good receiver can usually narrow this down to around a hundred feet on a map, and give an altitude within a couple of hundred feet as well.

If the receiver is moving, it computes how fast and in which direction it is moving. It can act as a "perfectly" accurate clock, a compass, a speedometer, an altimeter, a what-have-you. In fact, a good GPS receiver together with a map database, can even compute your local time of sunrise and sunset, today's tide table at points along the coast, and reset the time when you change time zones. In fact, as I write this, the watch on my wrist has a GPS built into it with all of these features.

Two-body Collisions

The "clack" of pool balls has always been a sound that catches my attention. And, what is sound? It's the collision of molecules in waves and patterns—an extremely large number of collisions all happening very fast. The simplest collision is one between two ideal spheres. And, to imagine this, pool is a pretty good analogy. Let's see how it works.

Imagine a cue ball and an object ball at rest on a pool table. The cue ball is impelled toward the object ball and strikes it at some point. The collision can be dead center or off to one side. In pool we might assume that the balls are equal in size and weight. If they are not, they will behave differently. Let's also say that collisions are perfectly elastic and there is no friction between the balls. This ideal is only approximated on a pool table.

When a cue ball strikes an object ball, the centers of both balls and the point of contact all lie on a single line. If you draw a line at right angles to this line, exactly through the point of contact between the two balls, you will have the tangent line of contact. The two most important lines in a two-body collision are the tangent line and the line of centers (both through the point of contact). Making the above assumptions (no friction, perfect elasticity, and balls equal in size and weight), the object ball will move away from the collision along the line of centers and the cue ball will move away from the collision along the tangent line. This means that their paths will always be at 90 degrees to one another. Notice that the original path of the cue ball doesn't matter. The only effect it has is to determine how fast the two balls will be going after the collision. If the cue ball hits the object ball dead on, its path makes an angle of zero degrees with the line of centers of the two balls (they are on the same line). After collision, the cue will stop dead and the object ball will move away at the same speed the cue ball was previously moving. If the cue ball hits to the right or left of dead center, then it will depart along the right or left tangent line, respectively.

After collision, how fast will the balls be moving? A ball's new speed will be the original speed of the cue ball times the cosine of the angle between the cue ball's original path and the new path of the ball in question. To get a feel for this, consider three cases: A head on collision, a hit at forty-five degrees, and the thinnest possible cut. In a head on collision, the tangent line is at right angles to the original path. The cosine of 90 degrees is zero. So, the cue ball will not be moving at all after a head on collision. What about the object ball? In a head on collision, its new path (along the line of centers) is identical to the cue ball's original path. The cosine of zero degrees is one, so it picks up all of the cue ball's original speed.

Now consider the opposite extreme, a maximum "cut" shot, where the cue ball hits the object ball at 90 degrees. Now, the tangent line is identical with the cue ball's original path. The cosine of zero is exactly one. The cue ball continues on its original course and speed. Notice that a 90-degree cut is the same as a miss. The slightest hit at all makes the cue ball deflect slightly from its original path, so the cosine is slightly less than one, and the cue ball will depart along the tangent line at somewhat less than its original speed. The object ball departs along the line of centers. At ninety degrees its speed will be zero. At 89 degrees, its speed will be 0.0175 that of the cue ball's original speed (the cosine of 89 degrees). But, you should realize that a cut is with respect to the original position of the cue ball and object ball. Both are over 2 inches wide and no more than 9 feet apart, so the thinnest cut is off to one side. A cut of perhaps 85 degrees from the original line of centers is about the best you can make.

A hit where the tangent line and the original path make a forty-five-degree angle causes the cue ball to veer off at forty-five degrees (along the tangent line) with 0.7071 of its original speed (the cosine of 45 degrees). The object ball also departs at forty-five degrees (along the line of centers), so its new speed is the same as the cue ball's new speed. It appears that the sum of the new speeds is more than the original speed of the cue ball. This is because total momentum must stay exactly the same. The forward speed of both balls adds up to exactly the forward speed of the cue ball, and the sideways components are in opposite directions, so they cancel out. This follows from the Pythagorean theorem. You need to draw the velocity vectors. The velocity of each ball after a 45-degree collision is 0.7071, but this is the same as half of the velocity at 0 degrees and half at 90 degrees. Drawing three lines from the point of collision, one straight on and half a unit long, one at 45 degrees and 0.7071 units long, and one at 90 degrees and half a unit long, you get two sides of a square and its diagonal. Remember, the sum of the squares of the two sides is equal to the square of the diagonal, so one-half squared plus one-half squared should equal 0.7071 squared. And, so it does!

Now, let's use our imagination even more and vary some of the assumptions. Let's imagine that the object ball is infinite in mass. When struck by the cue ball, it doesn't even budge. With perfect elasticity and no friction, the cue ball will bounce directly back from a head on collision. As collisions take place more and more off center, the point of contact, the centers of the two balls, and the initial path of the cue ball will no longer all be on the same line. The only factor that remains the same is that the line of centers and the tangent line are at right angles. Now, think about the initial path of the cue ball and the tangent line at contact. Think of the tangent line as a "backboard" (like a tennis or Ping-Pong backboard). It's immovable with respect to the collision. When a ball hits a backboard, the angle of incidence equals the angle of reflection. The cue ball departs from the collision at the same angle off of the tangent line as it came in on. If it came in at ninety degrees, it will bounce straight back at ninety degrees. If it hits the tangent line at forty- five degrees it will bounce off at forty-five degrees (being, itself, deflected by exactly ninety degrees). If it grazes the object ball such that the tangent line and its original path form an angle of only ten degrees, then it will be deflected off of the tangent line at ten degrees for a total deflection from its original path of twenty degrees. All of this assumes an object ball of infinite mass.

Suppose the object ball had no mass. Now, when the cue ball contacts the tangent line it simply passes right through it undeflected. These two extremes define opposite limits. With an object ball of zero mass, the cue ball travels directly through the line of tangents. With an object ball of infinite mass, it "reflects" off the tangent line at an angle equal to the angle of incidence. So, it follows that with balls of equal mass, the angle of reflection is zero, or directly along the tangent line, exactly halfway between the two extremes. This is what I stated above, but now it becomes a bit more intuitive. Suppose we vary some more of the assumptions, but in a more realistic way with respect to the game of pool.

When you hit the cue ball, if you don't hit it perfectly dead center, your cue stick will impart spin to it. You can hit the cue ball above, below, to the left, or to the right of center. Each of these will cause the cue ball to spin as it moves from where you strike it with the cue stick on its way to the object ball. It first skids across the table and then friction with the table imparts a forward spin to the cue ball. When a cue ball with spin hits an object ball dead on, the collision has different results that depend on its spin at the moment of impact. For example, pool players call the forward spin imparted by hitting above the center of a cue ball "follow." This is because the cue ball follows the object ball after a head on collision. Backward spin (from hitting below the center) is called "draw" because the cue ball draws back after collision. Right and left spin are called right and left English. Maybe it's a technique the English were the first to develop. By the way, before hitting a cue ball off center, chalk up your cue stick or you will "miscue." Also, when you make a long shot, shoot with a little follow. When you make a short shot, especially if it's dead on, shoot with a little draw. Both of these techniques improve accuracy by compensating for friction.

Real pool balls collide with some degree of friction. The dirtier they are, the more the friction. Friction has its most noticeable effects when using English (left or right spin). When the cue ball hits an object ball dead on with left English, it deflects the object ball to the right (and the cue goes slightly to the left to conserve momentum). The opposite happens with right English. An even more noticeable thing happens when the cue ball strikes a rail (edge of the table). Without friction, its angle of reflection is equal to its angle of incidence. But, there's always friction! With English you can imagine what happens. If you hit a rail dead on (at ninety degrees) with left English, the cue ball will deflect to the left. Take the same shot, but aim slightly to your right keeping the left English and the cue will rebound straight back (or even slightly to the left) instead of to the right as you might expect.

Keep in mind that even pool balls aren't the perfect example of ideal two-body collisions. The balls aren't perfectly clean. Friction and spin affect the way they behave. So, does the ideal two-body collision ever actually occur? My guess is that the type of collision described above only ever occurs on a pool table. When you go up or down the size spectrum you get large bodies in outer space, or particles, atoms, and molecules in the other direction. Collisions in space are pretty messy and catastrophic events—nothing like those between a couple of pool balls. Collisions between molecules and atoms are not between perfect spheres, nor are they perfectly elastic. Even collisions between two identical particles do not follow the paradigm, because, in this case, impact is due to mutual repulsion, not contact between two hard surfaces.

This example shows how a model can be constructed to describe a type of event, but that the model may not scale to other similar events. Is it plausible to pursue a model that scales to all sizes? Let's see.

Black Holes

A hole is normally obtained by digging. The more you dig, the deeper the hole. Holes can be dangerous just because they are deep. If you find yourself trapped in a hole of your own making, the first thing you should do is stop digging. There are holes, black holes, and the Black Hole of Calcutta. That was a dangerous hole, but not nearly as much as a gravitational black hole. The Black Hole of Calcutta had a couple of small windows, and some people even got out of there alive. With a gravitational black hole, there's no chance of looking through a window, and a very slim chance of getting out.

How do you get a gravitational black hole? First of all, you shovel stuff into it—you don't get this type of hole by digging stuff out of it. As you pile more and more stuff together, gravity and density become greater and greater. The gravity at the surface of an object depends on how much mass the object has and how dense it is. An object with a given amount of mass has a higher surface gravity the more its mass is concentrated into a smaller volume—in other words, as its density increases.

Any amount of matter produces gravity. Large, concentrations produce stars whose gravity is sufficient to initiate nuclear fusion. Our sun is a star. Smaller, concentrations of matter are called planets if they circle a star, or moons if they circle a planet, or asteroids or planetoids if they are too small to qualify as planets. This works down through all sizes: rocks, grains of sand, to molecules, atoms, and sub-atomic particles. Each of these objects, due to the gravity at its surface, has a particular escape velocity. That is, if you were standing on its surface and threw a rock away from it fast enough, the rock would escape and never fall back. The escape velocity of the Earth at its surface is 25,000 miles per hour. This means you would have to fire a bullet from the surface of the Earth at about seven miles a second for it to escape into outer space. If you fine-tuned the speed just right, the bullet would orbit the Earth once, come full circle, and hit you in your back. At any speed less than escape velocity, whatever goes up, comes back down.

A full explanation of ***escape velocity*** is at: https://en.wikipedia.org/wiki/Escape_velocity. The bottom line is that objects propelled away from a larger body orbit in an ellipse with less than escape velocity, in a parabola with exactly the escape velocity, and in a hyperbola with more than the escape velocity. This accords with Newtonian physics (not relativity).

Imagine objects with higher and higher surface gravities. They would have higher and higher escape velocities. When you talk about ever higher velocities, of course, you eventually reach a limit—the speed of light (the speed of light is the cosmic speed limit).

A black hole is simply an object whose escape velocity is greater than the speed of light. If you were standing on the surface of an object whose escape velocity were exactly equal to the speed of light, a beam of light from your flashlight would orbit once around it and shine on your back. All light emitted from the object would fall back onto it. No light would escape. No light could bounce off the object either. The object would be absolutely black (or invisible) from a point of view some distance away. Thus, the name "black hole."

Einstein's famous equation, $E = mc^2$, defines the equivalence between energy and mass. Let's use the term ***quantum*** (or ***quanta***) to refer to a unit (or units) of mass/energy without regard for which of these it actually is. The "Schwarzschild radius of a spherical black hole" can be computed given the amount of enclosed quanta. This calculation is explained here: https://en.wikipedia.org/wiki/Schwarzschild_radius. It calculates the radius that a sphere will have if it encloses sufficient quanta to cause the escape velocity at its surface to be the speed of light. According to this equation, as the density within a volume approaches zero, the size of the Schwarzschild sphere approaches infinity. As the density approaches infinity, the size of the Schwarzschild sphere approaches zero.

A ***photon sphere*** is a region of space with gravity strong enough that photons orbit within the region. A photon sphere has a radius of $3/2\ R_s$, where R_s denotes the Schwarzschild radius. Light within a photon sphere cannot escape. Let's define a reference space as the photon sphere that contains an observer. This would effectively be the observer's apparent universe. Photon spheres may be nested. Any photon sphere contained completely within an observer's reference space is therefore a black hole. In the following text, the term ***black hole*** will be used to refer to a photon sphere, and a photon sphere is defined as a larger sphere surrounding a Schwarzschild sphere.

Theory has it that black holes can come in any size (with some restrictions at the quantum level). They could be microscopic or astronomic. Any amount of mass—no matter how little or how much—forms a black hole if constrained within a certain volume. You could say that there are two types of black holes. One type of black hole is where we as observers exist outside of it. This is the type normally discussed. However, another type of black hole coexists with us in the same larger universe. This is because any large or small collection of matter and energy taken by itself is a black hole at a large enough radius from its center. Our universe in total appears to contain the amount of matter and energy that would form a black hole if its radius were about 13.8 billion light years. Our galaxy is denser than the universe at large, and our solar system is denser still. If either of these were the sole occupant of the universe, that universe would be smaller than the universe we know. Within the universe, black holes of this second type overlap. It is interesting, but not very useful, to note that the universe and subsets of it define black holes.

Any amount of quanta (large or small) may be plugged into the Schwarzschild equation to compute a radius. If the radius turns out to enclose us as observers, a type two black hole is defined. Otherwise, type one black holes may be either Large Black Holes, or Primordial Black Holes, and we observe them from outside.

The smallest of the Large black holes comes about in the natural life span of a star about three times the size of our sun. When this amount of matter accretes into a star, a series of nuclear fusion reactions is triggered. Each step in this series involves the fusion of heavier atoms into still heavier ones. The series stops when the atoms are too heavy to undergo another fusion cycle. We believe this point to be reached when the matter is largely iron. Now, what happens? Up to this point, the star is kept from gravitational collapse by a massive conversion of matter into energy. This energy, in the form of heat, holds the atoms apart. When no more energy is available, the star begins to cool. Now, the star literally collapses under its own weight. This "death" of the star may, or may not, cause a black hole to form momentarily, but what it does cause is direct contact between quanta in different particles. This results in particle annihilation, and the runaway release of free energy. This is observed as a supernova that spews all of the mass/energy back out into space. It enables a massive fusion process forming atoms even heavier than iron.

Large and Primordial black holes come about in different ways—very different from one another—and very different from the "death of a star" scenario. This is because a Primordial black hole relies on quanta compressed to fantastic densities to within a small enough volume that the escape velocity exceeds the speed of light. Such densities might be a by-product of the "death" scenario, or perhaps of some bigger Bang scenario. Considering a range of black hole sizes, there might be few, or even no, actual black holes along some parts of this black hole "spectrum." Every size might be possible, but some sizes might be extremely short lived.

As you increase the mass of a black hole, the critical volume goes up and the critical density goes down. When a very large mass is involved, the density doesn't even have to be that great. Inside a very large black hole there could easily be normal stars and vast swirling vortices of dust and debris. Projecting this to one extreme, a black hole "almost" infinite in volume could be formed by a single quantum with a density of "almost" zero. At the other extreme, a black hole "almost" infinite in density could be the size of an atomic particle.

There may be black holes within our universe that are stable, and large enough to harbor life, appearing to their inhabitants to be their entire universe. In fact (or rather theory), our universe could be such a place within an even larger universe. So, could we transition from one level to another? First, since light can't escape a black hole, if we headed into a very large, but not very dense, black hole, we would have no idea that we were doing so. If we did luck out and enter a fairly cool black hole, we might survive the transition through the event horizon, but if we were in free fall, and fell from a very long distance away, relativistic speeds would be involved. We'll discuss this later.

How about getting out of a black hole? The caveat is that the escape velocity is greater than the speed of light, and matter cannot travel that fast. But, that's not the end of the story. A rocket ship can leave the Earth at a speed lower than the escape velocity if it maintains acceleration. If a rocket were powered by an efficient matter to energy converter, it could accelerate for a very long time. This is essentially what happens when the density inside a black hole causes particle interaction on a large scale. Mass to energy conversion takes place, the black hole becomes very hot inside, and the pressure builds up to the point that a supernova, a small bang, or even a Big Bang, occurs. The rocket technique is a possible safe way to escape from a black hole. The latter alternatives are all catastrophic ways to escape (but they could occur in theory).

So, escape from a black hole should be possible with ongoing acceleration, but what about the singularity inside? What is a singularity in any case? Imagine crowding lots of atoms into a very small space. At first, their atomic structure holds them apart. The electrons around each atom repel those of the atom next to it. But you are relentless, you keep adding more and more atoms, increasing the force of gravity until the mutual repulsion of the electrons is overcome. Now, the nucleus of one atom and that of the next are forced ever closer together. You keep piling them on, increasing the force of gravity until even the repulsion between one nucleus and the next is overcome. At some point gravity exceeds all the forces that hold matter apart, and then what? The standard model has it that matter, at that point, is drawn together without limit. Everything caught in such a collapse is supposed to disappear into a point called a singularity. Let's see why escape from a black hole actually occurs long before a singularity can occur, and why density must have a natural limit.

The key to escaping a black hole is a sustained outward force. When matter reaches a sufficient density, it annihilates itself. A fireball of energy is produced that sustains an outward pressure for a period of time. Time enough for the black hole's strangle hold to be broken. This fireball just has to be persistent; it doesn't have to drive particles above the escape velocity.

How does such annihilation occur? The only particles we know of, that annihilate one another, are anti-particles. The primary example is an electron and a positron. These particles are easy to get into close proximity. They attract one another because they possess opposite electrostatic charges. Ordinary matter consists of electrons, neutrons, and protons. Neutrons and protons are located together in the atomic nucleus. Electrons are located in "clouds" around the nucleus. Each electron fiercely repels every other, just as each atomic nucleus fiercely repels every other. The repulsion is fierce because it is based on electrostatic forces that are vastly stronger than the force of gravity (one followed by 36 zeros times as strong). This force drives nucleus away from nucleus, and electron away from electron. On the other hand, electrons and the nucleus full of protons are ***attracted*** with a force equally strong.

Particles annihilate each other by converting some or all of their mass into energy. Long before they reached the point of a singularity, particles in a gravitational collapse would be forced to merge with one another. Their inner quanta would interact when particle boundaries were violated. All of their mass would become raw energy. This is annihilation. It would generate quite a Bang. Certainly one big enough and persistent enough that the pressure of heat would force the volume to become greater than the original Schwarzschild sphere (changing a type one black hole to a type two).

Let me restate this. When gravity forces particles into contact, the quanta involved interact directly. This always involves different quanta with the same ***total*** mass/energy and momentum leaving the point of interaction. Thus, particles going into such an interaction cause very large photonic energy to come out—total conversion of mass to energy.

Experiments that could bear this out, especially ones involving large black holes, tend to be inimical to our way of life. But, on a somewhat safer and more limited scale, we have split atoms, releasing both particles and photons, and we have fused photons into electrons and positrons—we have both formed and annihilated anti-matter. Particle decay is an example of escape from the smallest of black holes. Experiments with large black holes are quite a ways off for the present. And, until black holes are the subjects of experimentation, they can only be the subjects of measured speculation and limited fantasy.

In reading this, your mind may have wandered to speculate about different sizes and densities of black holes. My mind certainly did in writing it. Notice (in a later chart, ***The Very Dense***) that the universe and its estimated size and density is about what it would be if it were a black hole. Given any amount of mass at any density, there is a solution for its Schwarzschild radius. Every point in space is within some volume defined by a Schwarzschild radius that corresponds to the local mass and density. Increase the density and you decrease the size of the black hole; decrease the density and you increase the size of the black hole. But neither ever gets to zero or infinity. This implies that no orbital path for a photon or a particle ever becomes a hyperbola, or even a parabola—all orbits must be ellipses (though some could be ***very*** large).

Our observable universe may be a black hole within a larger universe. The question is, how are black holes nested ***within*** our universe? I posit that the smallest black holes are particles. The next smallest are atomic nuclei, which are stable up to about 100 protons in size (with varying numbers of neutrons). Above this size, collections of quanta bound together in black holes are simply called ***primordial*** (and they are only theoretical). The range defined as primordial might extend up to the mass of a few suns. Those above this size, up to the mass of the black holes in the center of galaxies, are ***large black holes***.

Very large black holes might be low density black holes that are stable over a long period of time (such as our universe), and high-density black holes are the ones that go Bang. Primordial black holes are ultra-high-density black holes stable enough not to go Bang. They include atoms and particles, and who knows what else that has so far gone unobserved.

The question is, over what range of mass (total quanta) is a black hole defined strictly by its density? At extremely high densities does charge enter into the equation along with gravity? Two ranges of stable black holes might exist, and one range of unstable black holes (those mid-range in size that go Bang). The ultra-high-density black holes that form atoms and particles may rely on the attraction of opposite charges, in addition to gravity, to achieve black hole status.

It is my contention that all of the phenomena from particles to suns, and supernovas could arise from various conditions consistent with this model of black holes. This model even accommodates small ***Big Bangs***. As for a single ***original*** Big Bang, I believe another explanation of the observation that led to this theory is more appropriate. But, we are not there yet.

Philosophical Musings

One area that has caused more than one physicist to become a philosopher, or vice versa, is the area of quantum mechanics. Quantum mechanics involves the small—to an extreme. But it raises big questions. Questions about quantum mechanics started with Heisenberg's Uncertainty Principle. It says that there are trade-offs in what you can know. To be more certain about one aspect of a quantum situation, you have to be less certain about some other aspect. Quantum mechanics goes on to describe a number of phenomena that are probabilistic until they are actually observed. Here is where the philosophy enters. The world of photons and sub-atomic particles is a very high-speed, twitchy, and non-localized world. We cannot access it directly; we must observe it indirectly. The statistical models of quantum mechanics have proved very successful in describing this world. Very successful, indeed.

If we observed the freeways of a major city from outer space and were prevented from coming any closer, we might develop a statistical model of rush hours. Vehicles would be "particles." We would find a very good correlation between traffic density and the rotation of the planet beneath us. But would we observe the fact that a vehicle nearly always returns at night to the same point from which it originated that morning? Would we understand the cause and effect relationship that explains this? Would we conclude instead that the entire situation was governed by probability? Here on Earth we know there is an underlying reality, because we are a part of it. We could link the regularities of traffic surges to human ownership of the vehicles, human psychology, sociology, economics, traffic regulations, and a host of other factors or "scientific models" that cannot be brought to bear when peering down from outer space.

Quantum mechanics describes certain attributes of photons and particles with complex tensors involving probabilities. When the state of an object is observed, this "wave function" is said to collapse—the object is now in a definite state, not a probabilistic one. This creates philosophical problems. Electrons are said not to have definite positions around an atom: until actually observed, they inhabit "probability clouds." Certain attributes of particles and photons are linked together when they depart from a common event. When the state of one is determined, perhaps miles away, the state of the other is also determined. Somehow, the wave function collapses as a result of one measurement, and the information from that measurement is communicated instantly (that is, faster than the speed of light) to the other point of measurement. How can the wave function collapse to produce this "action at a distance" that exceeds the speed of light? Why not assume that a deterministic linkage exists all along? Again, the most ***practical model*** need not be the ***underlying reality***. But, more than one physicist, turned philosopher, has opined that the probability tensors of quantum mechanics ***are*** the reality. My guess is that any model based on probability and statistics is only a convenience. It signals a situation where we are removed from the underlying reality and have no better way to describe or interact with it. It makes scientific sense to build such models—they work—but it's unsound philosophy to claim them as reality.

Might it not be sensible to suppose that the world of particles and photons is deterministic on its own level—a level that is simply not accessible to us? We may find statistics to be the best way to model quantum mechanics, but that should not be the same thing as saying that quantum reality is fundamentally statistical—that quantum existence is nothing more concrete than a probability cloud or an uncollapsed wave function.

Particles are composed of photons whose pathways are tiny orbits about a point of mass. Particles propagating through space are photons in complex paths involving their orbits within the particle. The fact that the photon itself is always traveling at exactly the speed of light means that the particles must behave in ways consistent with the wave natures of their constituent photons. Photons possess characteristics that lead to the phenomena they exhibit, and to the phenomena that emerge from them. Chief among these is the wave nature of their electrostatic and magnetic forces. Photons interact when their "influences" intersect, that is where the "bends" they impose upon space overlap. This is the difference between the coherent light of a laser and normal, incoherent light: Each photon in a beam of coherent light is synchronized with all the others. Laser light coheres—normal light interferes. Coherence is non-interaction—interference is interaction that disperses energy.

We know that a photon can interfere even with itself. We also know that particles have a wave nature just like photons. Electrons are particles whose wave nature dictates the kinds of orbits they inhabit within various types of atoms. Their orbits are essentially restricted to regions that allow them to be coherent with themselves and with each other, they are not like the orbits of planets around a sun.

Orbits of planets around a sun have a degree of both chaos and stability. For example, planetary orbits must have a certain spacing or they interfere with one another; planets could be ejected from their systems. Each orbit of the Earth around the sun is slightly different from its last. When objects orbit in complex patterns, be they planets subject to gravity, or particles subject to electromagnetic forces, most arrangements are unstable and last very briefly. Only a few arrangements are stable enough to be observed in nature, and most of these are chaotic to some degree. If an atomic configuration has a fifty-fifty chance of self-destructing in a given period of time, we call that period the half-life of the configuration. At any given time, if an electron has a fifty-fifty chance of being in a certain region of space, this is an aspect of its "probability cloud." If there is an underlying reason for these phenomena, it's no stretch of the imagination to assume that it's linked to some dynamic involving chaos. Here, we are suggesting that the source of the chaos is derived from orbital mechanics involving quantum and gravitational interactions, as the case may be.

Ordinary matter is made up of protons, electrons, and neutrons. There is a large menagerie of other particles that could be called extraordinary matter. Particles have typically been created in particle collisions. There has been much less (but there has been some) experience in creating particles directly out of photons. What we do know is that any interaction among photons and particles (such as smashing them together), produces an identical sum of mass, energy, and momentum. The simple act of an electron changing from one orbit to another around the nucleus of an atom involves the production (going into a "lower" orbit) or consumption (going into a "higher" orbit) of a photon. In fact, the basis of all change is simply the transfer of photons.

Here, the simple assumption is that every particle is a photon of a certain quantum in a tiny pathway orbit around the center of the mass associated with the photon itself. Some particles (neutrons and protons) may take two or more photons in orbit around each other. We know that electrons and positrons can result directly from a photon interaction, but what about the particles in the nucleus of an atom? The protons and neutrons? What about the menagerie of less stable particles? Could these different particles also stem directly from the characteristics of photons in special configurations? A photon's orbit could not be less than one wavelength around, otherwise it would interfere with itself. A stable particle could only form if the associated mass of a photon produced a pathway equal to some integral number of its wavelengths. This match could arise from a simple, elliptical orbit, or some more complex topology. Only a few stable particles could be expected to evolve. And, this is certainly the case as we observe it.

From the example of the electron and positron, we might infer that when particles are produced by the collision of two energetic photons, pairs of anti-particles always result. These would be standing waves of opposite charge. They would be strongly attracted to one another on the basis of charge, and less (much less) strongly attracted by mutual gravitation. If attracted too closely, anti-particles annihilate one another—they become "loose" photons once again. On the other hand, objects that are attracted to one another tend to go into orbits. On the quantum scale, orbits are constrained by the effects of coherence and interference. Perhaps a neutron is made up of two anti-particles in orbits about each other. Their equal but opposite charges would simply cancel out. Protons and anti-protons might consist of neutrons with the addition of a positron or electron, presumably in some kind of stable orbit within the neutron particle pair.

Science accepts the conjecture that all of the lighter atoms are forged in the furnaces of first-generation stars. Heavier elements are formed from supernova explosions into second generation stars like our sun. It would take a much more concentrated fireball of energy to forge the original protons, neutrons, and electrons of the first-generation stars. It would take the "big bang" of a large black hole, but perhaps not one involving the entire universe; a fraction of a galaxy might suffice. It would take a fireball capable of a sustained force to explode out of its black hole. During this process, photons of very great quanta would impact each other. Particles of all kinds would be produced. Some, like the protons, neutrons, and electrons we observe today, are very stable. They would survive. Others would vanish in a flash—literally—sustaining the fireball. There is an even chance that antimatter would be produced by a Bang of this type, but necessity requires a stable configuration to evolve—matter and anti-matter cannot exist in close proximity. This battle was somehow fought long ago, and now only the victor remains. In our region of the universe, electrons exist outside the nucleus, and the positive charge is contained within it. The exact opposite would have been equivalent, and could theoretically have been the case.

More Food for Thought

Here, we shall expand our vision from the level of particles to the limits of the universe. We ourselves lie somewhere in between. Ours is a macro world compared to that of quantum mechanics. It requires a huge number of quantum components to construct us and the world of our lives and experiences. We can have no direct contact with the quantum world. The properties of our world only emerge from large numbers of things happening very, very fast in a microcosm far removed from us in size. At the macro-end of the size scale, there is a Cosmos that we will also never navigate. There are spans of time and measures of space that we know of but can never cross. At least, not as individuals. So, let's reach out with our minds, from the comfort of our armchairs, to the farther galaxies and revisit why their light appears to be red shifted. The farther away another galaxy is from us, the more red shifted the light is coming from it. The implication drawn from this by today's scientists is that the red shift is due to the Doppler effects of motion, and the more distant a galaxy is, the faster it seems to be moving away from us. However, what if this red shift had some other explanation? What if it were not due to the Doppler effect of relative motion, but to some other kind of energy loss? Why do we expect that a quantum of energy, after traveling for several billion years, will be as robust as it was when it set out? If any energy were somehow "bled off," the effect would also be a red shift. A different cause, but precisely the same effect.

We can verify that Doppler red shifts (those due to the relative motion between source and observer) are possible, but we can't verify an "energy decay" that takes billions of years to occur. Any red shift looks like any other. For example, light is red shifted when it travels from a source deeper in a "gravity well" to an observer higher up. Could the curvature due to a "gravity well" be similar to the curvature of a universal pathway? If the red shift of distant galaxies is not due to their motion away from us, but due to some other cause, then there was very likely no Big Bang and no expanding universe. On the grand scale the universe could be eternal and static. More than one thing could explain the red shift of light from distant objects. We simply have no data to sort this out because all of our observations are restricted to the near end of forever.

According to the conjectures we made in the previous section, black holes are the very stuff of which we're made, and the very source of all creation beyond the original existence of space and the mass/energy it contains. The tiniest black holes are particles. Stellar black holes turn into supernovas and generate the heavy elements. Much larger black holes might erupt to create the protons, neutrons, and electrons of first-generation stars. Cosmic black holes are universes unto themselves. Our universe could be a very large black hole contained in a larger one, and so on. All space is warped by the quanta it contains into finite, but unbounded, chunks. Each chunk is defined by its pathway curvature, which is defined in turn by the paths that light takes to travel through it (which is set by the total quanta producing the gravity within it). On the Grand Scale, a static and eternal universe is the simplest explanation.

Now let's take a look at the "Universe of the Big Bang." This will follow the "standard model" pretty closely, but keep in mind that not every scientist is lined up with each point in this description. Based on the observation that galaxies are uniformly distributed in all three dimensions throughout known space, and that the more distant they are, the more red-shifted they are, scientists have inferred that a Big Bang some 13.8 billion years ago was the origin of the universe.

This was not an explosion at a point in three-dimensional space, but an inflation of our three-dimensional space within a fourth spatial dimension. The way to think of this is to imagine a balloon. Blow it up just a little and use a marking pen to put ink dots all over it. Each dot represents a galaxy. Now blow it up quite a bit more. Notice that each dot gets farther away from its nearest neighbors. All the dots move away from each other, but the farther away two dots are, the faster they move away from each other as the balloon is blown up. The two-dimensional surface of the balloon is an analogy to the three dimensions of our own universe. Every line around the circumference of the balloon is a geodesic. Its two-dimensional surface expands within three dimensions. The Big Bang requires our three-dimensional universe to be the surface of a four-dimensional expanding "balloon."

If the Big Bang had simply been an explosion in three dimensions, it would have produced an expanding sphere of matter and energy. The energy would have traveled away from the point of origin at the speed of light, and the matter at some lesser speed. The explosion would be like a spherical shell. Empty space would lie ahead of its wave front, and relatively empty space would be enclosed within the expanding shell. Observations show us that this is not the case—we find that galaxies are distributed uniformly in all directions. Hence, the "balloon" scenario. In the standard model, the inflation of space itself does not have to obey the "speed limit" of light, and it is postulated that the universe initially expanded much faster than light could propagate through it. It is claimed that, even now, radiation from the original explosion is still reaching us. This radiation is called the Cosmic Background Radiation.

In the standard model, light also follows universal pathways, and the pathways themselves are still expanding. It has not been determined if they will continue to expand forever, stop at some point, or even reverse and contract into some kind of Big Crunch. In fact, the mechanism and implications of this kind of expansion and contraction continue to be the subjects of lively scientific debate. Many crucial points in this model are far from being resolved—but, the occurrence of the Big Bang itself is rarely questioned.

One problem with the standard model is that so many aspects of it appear to be unrelated. Or, worse yet, that so many "fortunate" coincidences have occurred to produce a universe in which we could arise. Of course, if this hadn't been the case we wouldn't be here. Given the fact that we are here, some scientists say it doesn't matter how fantastically improbable a universe based on our model is, the model need only be supported by the observation that we ***are*** here. The best scientific theories and models have plausible connections all the way through. The standard model falls a bit short in this area.

What I'm suggesting are some alternatives for the predictions of relativity, quantum mechanics, and the Hubble red shift. One alternative is that there was no Big Bang—that a different explanation for the red shift of distant galaxies is plausable. Another alternative to the standard model is that quanta (photons) are the constituents of everything (matter is simply energy in special configurations). This model has it that there is a deterministic reality at the quantum level, and only our inability to access it limits us to models involving "probability clouds" and "wave functions."

The pathways of motion and propagation result from the curvature of space induced by the presence of matter and energy. The pathways within a black hole expand when there is a sustained pressure (due to an incredible fireball produced by the total annihilation of matter) within it. On the smallest scale this is the case when an electron annihilates a positron. On larger scales, it would be more like the early stages of the Big Bang as it is currently conceived, but space itself would not expand, only the pathways within it would change as the mass/energy is redistributed.

Time in this model is nothing more than a comparison between periodic phenomena. The most basic of these are linked directly to the universal constant, the speed of light. Less fundamental phenomena—from the semi-fixed and semi-chaotic orbital mechanics of photons and particles, to the phenomena of heat, entropy, information theory, and biological processes—have a huge asymmetry of probability that gives an absolute direction to the passage of time. Time as a scalar unit is a convenience in defining physical models. Time as we experience it is an emergent property. Time as a dimension, or the concept of "time travel," is nothing more than a confusion of unrelated concepts (an intriguing fantasy).

Space and the quanta it contains may never have had a beginning, but the pathways of our present universe did, and every photon and particle within it had a beginning and will have an end, in that they transmuted from, and will transmute into, other photons or particles. All objects and phenomena are simply localized rearrangements of space and collections of quanta. Original creation? There is no such thing! Recycling is total and eternal. Everything is made from something that existed before. Everything is either a copy of something else, an accidental arrangement, or it lies somewhere in between. Evolution is the process that allows new arrangements to emerge, but no process permits a change in total quanta.

This brings up one final issue: principles of equivalence. Whether a moving frame of reference is compared to a stationary "background" or to another moving frame of reference, the Lorentz transformation describes how one frame observes another. The main message of special relativity is that there is no way to tell the difference—any two frames of reference may be compared only on the basis of the relative motion between them. A second instance of equivalence comes from general relativity. It states that there is no way to tell the difference between the force of gravity and the force of acceleration. If it's true that gravity and acceleration are bends in space, it's possible that the curvature of universal pathways should be added to this list of equivalent phenomena.

This book is an attempt to push "reset" on the effort to build a Theory of Everything. It suggests a fresh start at a derivative approach, one that seeks to build explanatory links instead of ever more complex tensors and topologies. The nature of our reality makes many distinctions difficult. When an equivalence is involved, it may make the choice of a model impossible. Other criteria, such as personal preferences, the fad of the times, or the "received wisdom," make the choice of a model much easier—but not necessarily better.

Part 1 — Reality

Above we have explored the nature of existence, the nature of rules, and the nature of chance. We have raised questions about the standard interpretations of certain accepted observations. We have laid the necessary groundwork to understand the following summary of the modern realization of ***our reality***.

> Man at last knows he is alone in the unfeeling immensity of the universe, out of which he emerged only by chance. His destiny is nowhere spelled out, nor is his duty.
>
> - biologist, Jacques Monod, 1972

Cosmos vs. Universe

Ask a search engine the difference between cosmos and universe, and you will get a number of opinions that either differ slightly, or fail to include all aspects of the question. Here I will try to sort out the differences. First, cosmos is often set in juxtaposition to chaos as its opposite. The universe is a set that contains both order and chaos; it contains everything that is real. There seems to be some debate as to whether the cosmos contains the universe, or vice versa. It is also conjectured that our universe might contain, or be contained by, another universe (ad infinitum?). I will define the universe as the black hole in which we find ourselves. I hold out two possibilities: first that there may be black holes within our universe that are universes to intelligent entities within them, and second that our universe might be a black hole within a larger universe. Thus, the term universe is relative to an observer within it, and each observer is local to exactly one universe. A given observer near a black hole has no way to see inside it unless it is actually entered, in which case it becomes the observer's new universe. Likewise, if an observer escapes a universe, the old universe becomes a black hole within the observer's new universe. Both of these states involve passing through an "event horizon." The ability to cross an event horizon and survive would require sustained acceleration to a degree not possible with any technology based on modern science. The natural way a black hole is escaped is when it decays and releases its mass/energy in complete chaos. The only natural way that a black hole may be entered is in the form of pure energy, which occurs when a particle reaches the speed of light by falling into an event horizon. As for the term cosmos, I will us it in a less rigorous sense to refer to our physical universe and any others (or anything else) that might coexist with it.

Physics

Physics is the science of describing the natural world at its most fundamental level. My approach here will be to follow the example of mathematics, and begin with as few "givens" as possible. However, in another sense, every observation of the natural world is a given. Accepted science begins with these. It goes on to build various models that agree with each other and with all verified observations. Physics studies space and time, matter and energy, force and motion.

Models always tend to be based on some analogy between a perceived pattern in the data and something more commonly understood. In the following, I'll keep to this tradition and try to explain concepts in terms of analogies rather than in terms of complex mathematical expressions (in most cases). Some simple math is unavoidable.

Modern physics is a large compendium of observations and models involving the different forms and relationships of matter and energy. The Holy Grail of physics is a Theory of Everything (a TOE). Such a theory would have to be built from as few initial assumptions as possible, and show how all observations can be explained by those assumptions.

Our World

The world best understood by man is described by classical physics. This is a world that is slowly dying a "heat death" in which every event entails a loss of free energy to entropy. It is a world of cause and effect, where a future state of affairs can be predicted on the basis of initial conditions. It is a world in which life is an anomaly that science has made little progress in explaining. The part of the universe familiar to us is the part in which the very small is observed indirectly, the very large is seen only from a great distance, and the very fast is observed so briefly it is very hard to unravel.

The realities of the very large and the very small are realities that we can never enter; we will **always** observe them indirectly or from afar. But, this shouldn't stop us from trying to get a better perspective on them. Science builds useful models of them, and technology is based on them. Distance relates to size, and time and distance together relate to speed. The following tables describe the entire time and distance scales of the "known universe" in terms of the familiar units of inches and seconds.

<u>Distance in inches:</u>

1 Femto inch (10^{-15})	the size of a proton
1 Pico inch (10^{-12})	the size of an electron
1 Nano inch (10^{-9})	the size of a molecule
1 Micro inch (10^{-6})	the size of a virus
1 Milli inch (10^{-3})	the size of a grain of salt
(one inch)	the size of an egg yolk
1 Kilo inch (10^{3})	the height of an 8-story building
1 Mega inch (10^{6})	the distance across a large city
1 Giga inch (10^{9})	the distance around the Earth
1 Terra inch (10^{12})	the distance to the nearest planet
10^{15} inches	twice the size of the Solar System
10^{18} inches	the distance to the nearest star
10^{21} inches	.. to the center of our galaxy
10^{24} inches	.. to the nearest neighboring galaxy
10^{27} inches	.. to the edge of the known universe

Time in seconds

10^{-24} seconds	light crosses an atomic nucleus
10^{-21} seconds	period of nuclear vibration
10^{-18} seconds	light crosses an atom
1 Femto second (10^{-15})	period of atomic vibration
1 Pico second (10^{-12})	period of molecular rotation
1 Nano second (10^{-9})	light travels one foot
1 Micro second (10^{-6})	period of a radio wave
1 Milli second (10^{-3})	period of a sound wave
(one second)	period of a heartbeat
1 Kilo second (10^{3})	light travels across Earth's orbit
1 Mega second (10^{6})	a fortnight (two weeks)
1 Giga second (10^{9})	an average lifespan
1 Terra second (10^{12})	the age of the pyramids
10^{15} seconds	the age of intelligent life on Earth
10^{18} seconds	the age of the universe

What happens when we depart from our own reality by a factor of a billion (just three lines in the above table), in terms of size or speed? Let's see what reality would be like in regions like these, far away from our own in terms of size and speed.

The Very Small

This realm begins where our ability to see it directly, leaves off. By using a magnifying glass we can enlarge things a bit. With a strong microscope we can see things that simply can't be seen with the naked eye. A very strong microscope shows us Brownian motion, the evidence for the existence of molecules. This is the motion exhibited by very small specks seen under a microscope: They appear to jump about at random for no apparent reason. The "reason" turns out to be that a relatively energetic molecule impacting a barely visible speck causes it to recoil just like a billiard ball would move slightly if shot at by a BB.

The world at the level of molecules, measured in units of a billionth of an inch, is governed by quantum mechanics, not by classical physics. Reality at this level has virtually nothing in common with reality at our level. On the scale of the very small there is no such thing as heat or friction. The effects of gravity are almost unnoticeable, while the effects of static electricity are completely overwhelming. There is no such thing as light or sound as we know them; these phenomena become different things entirely. On the molecular level, light is an unseen jolt of energy that is emitted or absorbed by a molecule, instantly and violently knocking it, or a part of it, around or loose from whatever it was connected to. Sound cannot be distinguished from heat, and both of these are simply the degree of knocking around that's going on. This "degree" would be measured at a point in space, at a given time, and in terms of both rate and violence, if it could be measured. But, it can't — not at the scale of molecules and atoms. There's nothing to measure it against, and there is nothing to see it with. In fact, there is no way to scale or "shrink" any kind of observer (or camera) down to this level.

Of course, we can **imagine** ourselves at the level of a molecule, even if there is no possibility of actually **being** there. Soon, computers will be able to generate a virtual reality that will assist our imagination. But, there are two things about the reality of the very small that even a computer cannot handle. First, there is no way to "see" what's going on. We would simply have to "make up" something for things at that level to "look" like. Second, we would have to slow everything down by a factor of about a million or a billion, so that we could follow the action. The point is that the world of the very small is a totally alien reality. Without turning it partly into fiction, it's physically impossible to imagine. We possess neither the alphabet, nor the grammar.

Consider a size only a tenth that of our own. Any smaller than this, and it's extremely unlikely an organism with our intelligence could exist. When you reach a factor of a thousand smaller than we are, you reach the limits of how small we can directly build a standalone "machine." The term "nanotechnology" has been coined to describe possible machines ("nanobiots") in the range of a micro-inch. These machines would basically be "designer" molecules. We will need to control biochemistry to fabricate anything at this level. Machines smaller than this cannot exist according to science as we know it.

The Very Large

As we go up in size from ourselves, we again enter a different reality. On a scale of about a billion times our size, gravity begins to be the paramount force. Space is virtually empty. It is noticeably "bent." Light takes a significant time to get anywhere. And everything else seems to be standing still.

Certain aspects of the very large are not difficult to imagine. We can gather light into very large lenses from very far away, and see what exists on a very large scale. However, we can only see a 50-100 year "snapshot" of the universe because our lifespans are so short compared to the timescales of the cosmos. Events on this scale happen very slowly and over a very, very long time.

A simple example of the effect of size is growing a water droplet. At some point it reaches a maximum stable size, and the addition of more molecules of water to the droplet will cause it to divide in two. This is because of the relation of its surface tension, which contains it, to various forces inside. Atoms and molecules cannot be shrunk or enlarged, we can merely aggregate more or fewer of them to have larger or smaller objects. And even this principle only works over a certain stable range.

Unlike events on the molecular scale, events on larger and larger scales are more and more varied. Although things happen very fast on a molecular scale, not many really different **kinds** of things happen there. The same events happen everywhere, indistinguishable from one place to another. On the scale of planets and stars, nothing ever happens **exactly** the same way twice (though there **are** strange attractors—planetary systems, ring systems, various types of stars). The variations and possibilities are endless. From our small corner of the universe, we have seen a little and we have guessed about a little more. But, on the grand scale, the universe is so complex that we can only begin to dream about its realities. And much that's out there is simply beyond our ability to imagine.

However, we can all imagine a man about ten times taller than ourselves, say about sixty feet tall. He would be about a hundred times stronger than we are. This is because his strength is proportional to the area through a cross section of his muscles, and area is the square of linear distance. Thus, if he is ten times as tall, his muscles would be a factor of 10 x 10 as strong.

He would weigh a thousand times as much as a six-foot man. This is because weight is proportional to volume, or the distance cubed. At the same proportions, if a six-foot man weighed 180 pounds, a sixty-foot man would weigh 180,000 pounds. If a six-foot man could lift himself and another 200 pounds, a sixty-foot man could lift a total of 38,000 pounds (a hundred times the 380 pounds a six-foot man can lift). This leaves him almost a factor of five short of even being able to lift his own weight!

The net effect is that a 60-foot man would simply collapse. He not only couldn't stand up, but his bones would actually break if he tried to do so. The principle involved here is called the *square-cube law*. It works in both directions. It explains why ants can lift many times their own weight. Smaller animals get the advantage of greater strength in proportion to weight, just as larger animals get a disadvantage. The largest animals live in water where their weight is partially supported for them. Elephants have very thick legs as compared to horses. As animals get larger, their bones and muscles only get stronger by $height^2$, but their weight increases by $height^3$.

The square-cube law puts a kind of limit on how big our engineering feats might eventually grow to be. On a planetary surface, there is an upper limit to how large a structure can be built before it is crushed under its own weight. In space, we find that there is a limit on how large an asteroid can be and have an irregular shape. Larger than this limit, and asteroids produce enough gravitational force that they collapse into a more or less spherical shape. With better materials, and a honeycomb design, structures could probably be built that were several miles in diameter. Much larger than this, and designers have to consider that strength increases only a tenth as fast as the effects of mass and inertia.

The point of this is that our technology, our very existence and reality, are bounded by size. The very large and the very small are outside of our reality. We can look, but we can't "touch."

The Very Fast

The realities of the very small and the very large are realities that we shall never directly know. Our experience with the very small comes only from indirect observations. Our experience with the very large is much like trying to deduce the nature of the world from a few snap shots of subjects very far away, and whose details remain obscure. The realm of the very fast is different from either of these. This is a realm we may someday be able to enter. But, we have much less knowledge of it now than we do of the very large and the very small.

All of our experience of the very fast involves indirect observation of small numbers of fast moving, elementary particles. We have never seen an object even the size of a BB moving at near light speeds. Above, we learned that reality is totally different at a billion times smaller or larger than we are. Likewise, the absolute limit of speed is about a billion times as fast as a very slow walk.

Einstein described the model of space-time that is accepted by science today. It makes clear that all electromagnetic radiation travels at the same speed, no matter from what vantage point it is observed. This is the speed of light. Light travels at 186,000 miles per second. That's two-thirds of a billion miles an hour, or one foot per nanosecond. We've been able to reach speeds of about a thirty-thousandth of that. The only way to go faster would be to expend larger amounts of energy. To get anywhere close to the speed of light would require a totally new energy technology.

The first factor of ten came from the automobile. It took us from an 18 MPH running speed to a 186 MPH "land speed record" sometime in the early twentieth century. A rocket motor moved the first man another factor of ten to 1,860 MPH in the mid-twentieth century, and to 18,600 MPH in orbit around the Earth a few years later. Thus, after a million years of being stuck with an 18 mile an hour speed limit, we pushed the frontier back by three orders of magnitude, three factors of ten, in a little over 50 years.

Now, forty more years have passed, and we have made little progress in going any faster. In fact, it looks like rocket technology has about reached another speed limit. To move from 18,600 mile per **hour** to 186,000 miles per **second**, an increase in speed of 36,000 times is necessary. It seems unlikely that this will happen during the next hundred years or so, since both the need and the technology would have to be discovered. This kind of speed is simply not necessary for interplanetary travel. When our descendants do figure out how to travel very fast, they will be able to leave our solar system and travel to the stars.

Once our descendants are able to travel near the speed of light, how will their reality change? For one thing, total mass to energy engines will have to be invented. And, to whom do I refer when I say ***our descendants***? In truth, I'm referring to machines of our making—intelligent machines. We are ill-suited for space travel, or for life on any planet other than Earth.

These descendants of ours, if they evolve sufficient curiosity, could travel over the span of millions, even billions, of years around our galaxy and to others. Both their ships and their lifespans could be huge in comparison to ours. Collectives of intelligences could also be huge. In a few billion years from now, they would more than likely encounter other collectives in their travels, possibly ones with common ancestors. It is likely that they would have little in common with biological life. Such entities might regard biospheres like ours as little more than parks to be observed, recorded, and left alone. Who knows how many times our Earth has already been regarded in this way!

The Very Dense

Space is filled with matter and energy at various densities. From the vacuum of interstellar space at a density of almost zero to the fantastic densities within a black hole. According to the Wikipedia, the following objects, if compressed to the given Schwarzschild radius, would become black holes (note that the density of water out of the tap is = 1 g/cc):

The Earth	8.87 millimeters	(its density would = 2.04×10^{27} g/cc)
The Sun	2.95 kilometers	(its density would = 1.84×10^{16} g/cc)
The Milky Way	0.22 light-years	(its density would = 3.72×10^{-8} g/cc)
The Universe	13.7 billion light-years	(its density is roughly 1×10^{-23} g/cc)

Of course, within the universe there are many galaxies like the Milky Way, many more stars like the Sun than there are galaxies, and many more planets, Earth-like and otherwise, than there are stars. At still lower densities, are planetoids, rocks, and dust. But even interstellar space is filled with the photons that stream away from every source of hot matter in the universe. The thing to notice in the above table is that the least dense black hole is the Universe itself. The first three entries, being much larger than their Schwarzschild radius, are not black holes (yet!).

The Grand Scale

As telescopes get bigger and better, galaxies are discovered farther and farther away. The current limit is 13-14 billion light-years away. Beyond this, we think we observe only a cosmic microwave background (CMB). Could this be due to the existence of galaxies extremely far away, red shifted clear into the microwave? Could the CMB be radiation from cool interstellar matter? Or radiation that's already orbited the universe at least once? Could the most distant galaxies we observe be on the "opposite" side of the universe?

Perhaps there is simply no limit to how far away the most distant galaxy is from us. However, there is a limit to the size of the photon sphere that surrounds us. This is the sphere from which no light emanating from inside it can escape. What about light entering it from outside? How far might that "outside" extend? These are questions our current state of knowledge can't answer.

Another question about the grand scale is where is the center of our universe? What's there? If we liken the 4-D hypersphere that is our universe to a 3-D sphere, the "center" could be anywhere on its surface. If there were massive black holes there, creating the phenomena we see as quasars, from any point on the sphere, we could look in some direction and there would be a geodesic that intersected this "center" and we could see whatever was emanating from it.

Let's also not confuse the fact that there are two types of bounding: Black hole bounding and cosmological bounding. Black hole bounding is limited to 3 dimensions, and cosmological bounding is limited to the surface of a hypersphere. Upon that 4-D surface there could be concentrations of matter that constitute black hole universes. Our universe could be one. These would not be fixed in place, but until two universes drifted together, there would be no exchange of energy between them.

Basic Observations

Many observations are involved in explaining the cosmos. But here are 3 to start with. I shall try to give the "flavor" of each. The first is that the speed of light is a universal constant. It is observed to be the same in every locality of space no matter how fast, nor which direction, the observer is moving. The second is that light coming to us appears to be more red-shifted the farther away it originates. And, the third is that when two clocks are moved through space and brought back together, the one that has traveled farther has counted fewer ticks. The order that I have mentioned these phenomena is the order in which scientists discovered them. In a little more detail, these are …

The Michelson–Morley Experiment

https://en.wikipedia.org/wiki/Michelson%E2%80%93Morley_experiment

The fundamental result of this experiment is that the speed of light is a constant that cannot be ***observed*** to vary in a vacuum no matter which direction the light travels with respect to the observer's laboratory (and how fast, and in which direction the laboratory is moving). For this to be the case, the Lorentz transformations of ***length contraction*** and ***time dilation*** were accepted as necessary consequences. The history of these observations began in 1887 and extended into the 1920's. They were fundamental to Einstein's special theory of relativity.

The accepted implication: Position and velocity in space are completely relative between two observers, position and velocity cannot be given absolute values.

Hubble's Law

https://en.wikipedia.org/wiki/Hubble's_law

This law states that the more distant a star or galaxy is from us, the more red-shifted is the light coming from it. It is based on using the faintness of the light to estimate how far away it originates, and the signatures of particular types of stars to know how bright they are at their origin, and at what frequencies the spectral lines of their signatures occur to measure red shift.

The accepted implication: Space is expanding. All matter and energy originated at a singular point, and at a time in the past that is currently estimated to be around 13.8 billion years ago.

The Hafele–Keating Experiment

https://en.wikipedia.org/wiki/Hafele%E2%80%93Keating_experiment

This experiment (and several follow-up confirmations using different paradigms) confirm the fact that a clock that has traveled farther through space is behind a clock that has remained relatively at rest when the two are rejoined.

Implication (but not generally accepted?): Motion in space is ***not*** relative between observers; one observer is fundamentally more stationary than the other (see "The Twin Paradox" below). Does this upset special relativity?

Axioms & Derivations

Rather than immediately discussing the contradictions of interpreting these experiments, I shall propose two "cosmic axioms" and show how they are a better analogy for cosmic ontology than the Standard Model. These are the definition of ***space*** as the fundamental ***substrate***, and the definition of a ***quantum*** as the fundamental ***unit*** inscribed upon that substrate. The more basic of these concepts is space—and it is neither nothing, nor is it a "luminiferous ether"—it is very much something else!

Space

Space is a boundless extent in which points have relative positions. We are all familiar with the concepts of a point (zero dimensions), a line (one dimension), a plane (two dimensions), and a volume (three dimensions). Position on a ***point*** is given as 0 or 1 (either you are there, or you are not). Position on a ***line*** is given as x, measured from some starting point. But measurement involves some ***unit*** of measurement, and the choice of a starting point allows for measuring in either of two directions along a line. Thus, by convention, x is positive in one direction, and negative in the other. We derive our units of measurement from the physical extent of "objects." A point has no length, but all objects in the real universe have some non-zero extent (or size). This unit of measurement must be fundamental—something absolute over time and space. It will be defined (below) by the ***unit of existence***.

There are two basic ways that a line (a one-dimensional universe) may be boundless. The line may be infinitely long, or it may be curved into the second dimension to form a circle. If a line is vibrated in the second dimension, waves propagate down its length. Let's project these concepts into two dimensions.

In a two-dimensional universe, space is a "plane" surface. Starting from any point on this surface, every other point can be reached by measuring x units in a given direction, turning 90º left or right, and measuring y units more. Given a starting point and orientation, another position is defined as (x, y). In this universe, the concept of area emerges. Length is still given as x, but now area is given as $x \times y$. A two-dimensional universe can also be boundless in two basic ways: it can be an infinite plane, or it can be curved into the 3rd dimension to form a sphere (or other topological equivalent). A plane surface can also be vibrated in the third dimension to cause waves to propagate along its surface (like waves on a lake).

Carry this one step further, and you have the three-dimensional universe in which we find ourselves. Any point can be reached by measuring x units in a fixed direction, turning 90º left or right, measuring y units more, and then turning 90º up or down and measuring z units farther. Position is defined as (x, y, z) given a starting point and orientation. Again, length is given as ***x***, area is given as ***x* × *y***, and the new concept of volume is given as ***x* × *y* × *z***. Three-dimensional space could be boundless by being infinite, or by being curved into a 4-D ***hypersphere***—in either case, our space is a ***hypersurface*** capable of conducting waves analogous to waves conducted in lower dimensional spaces.

So, Cosmic Axiom 1: ***Space*** is the fundamental ***substrate*** of existence. It is a perfectly elastic 4-D hypersurface. Any tension or vibration induced into this substrate can be recovered from it. Energy is perfectly conserved within it or conducted over it in the form of waves. Space is fixed in position; it doesn't flow or change place. Space is the "playing field" for objective reality. Space is somewhat analogous to Jello (with perfect 4-D elasticity). It can be wiggled laterally in any of its 3 dimensions, or conduct a waveform over its surface with amplitude in the 4th dimension. Any deformation of space causes a local ***tension*** (for lack of a better word). Space has no mass or energy, but it can absorb and give back energy like a spring.

Quanta

Cosmic Axiom 2: ***Quanta*** are the fundamental ***units*** of existence. They exist as soliton waves on the substrate of space. These waveforms propagate at the speed of light. A quantum is described by a package of attributes (that all vary together): mass/energy, wavelength, and other things. Quanta, in turn, ***embody*** the attributes of distance and time, and mass and energy. Force fields (electric, magnetic, and gravitational) arise from the different kinds of tension induced into the substrate of space by the ***presence*** of quanta. All fundamental constants are consistent with the elastic properties of the substrate, the first being the speed of quantum propagation (the speed of light).

Any tension induced by the presence of quanta (deforming space) falls off by the square of the distance from the center point of the quantum. Tension induced by each quantum is summed over all of space. Because all quanta travel at the speed of light, the tension at a fixed point in space is constantly changing as the quanta in its vicinity move around.

A quantum is a soliton wave, exactly one cycle of a sine wave propagating through space. Its waveform doesn't vary. However, as it passes through a point in space, a cycle of tensions occurs at that point. The direction of propagation defines axis 1 of a quantum in space (the wave ***length*** extends along this axis). Along axis 2 charge varies from positive to negative over the wavelength from maximum plus to maximum minus. Along axis 3 the magnetic force varies from N to S, 90º out of phase with charge. And, axis 4 is the amplitude of the soliton waveform, a ripple on the "surface" of space producing a constant tension in all directions, one that moves through space with the soliton.

The first three axes (in 3-space) define the orientation of the soliton (or quantum). It has a fixed amount of mass/energy, a fixed wavelength and amplitude, and a fixed electromagnetic orientation. Not only is the ***shape*** of a soliton wave ***inscribed*** on the hypersurface, but electromagnetic tension is ***induced*** into normal 3-space. The tension at a point in space defines a natural path that is followed by any soliton wave passing through that point. Every point in space, because of the quanta in its vicinity, has a non-zero degree of tension that defines the pathway of quantum propagation at that point.

Visualize the difference between the waveforms of a low frequency (long wavelength) quantum and a high frequency (short wavelength) quantum. The stress or tension induced into the substrate of space is greater for shorter wavelengths. Their waveform curves more tightly. This "explains" why a quantum with a shorter wavelength represents more mass/energy than a quantum with a longer wavelength. It causes more tension (stress? curvature?). Since a quantum always propagates as a wave over the "surface" of space exactly at the unchanging speed of light, to say that a quantum has a rest mass of zero is either misleading or wrong (a quantum can't travel at any speed ***other*** than the speed of light). Mass and energy are not separate things; they are the same thing evidenced in two different ways (as we shall see below).

We know that quanta (energy) equate to particles (mass) using Einstein's famous equation $E = mc^2$. But how could quanta exhibit as particles? I conjecture that when quanta exist in a region of space with sufficient tension (induced by the quanta themselves), they may form standing waves, in which they are trapped. This is observed as a particle, and all forms of matter are built from such building blocks. Imagine quanta threaded together in interlocking pathways as the components of particles. String theory and sub-atomic particles are current ways that might describe this "threading." Particles are the most basic of black holes. Quanta not observed as particles are observed as photons. This very simple duality is the basis for all the complexity that arises from it. It reminds me of the very simple rules of ***Go*** that lead to the most complex board game yet invented, or Conway's simple rules for ***Life*** and all the complex patterns that it is capable of generating. The substrate of space is like the blank playing field of Conway's Game of Life. This game is one of the simplest examples of a virtual reality. The behavior of cells inscribed on the Life "substrate" may be viewed at the following link https://bitstorm.org/gameoflife/.

We, may well someday set up a sophisticated virtual reality, populate it with artificial intelligences, and play the role of their creator. The reality of our existence and theirs would be totally different. Using the analogy that compares our reality to that of a sophisticated virtual reality, is it so hard to believe in other forms of existence? Even a reality in which our existence could be created? Believing that another reality could exist, and imagining what it would be like, of course, require skills of a fundamentally different kind than we humans possess. Let's attempt to understand our own rather challenging reality.

Gravity

Gravity is a function of ***all*** the matter and energy in the universe. It is a stress induced into space by the energy "inscribed" upon it. A quantum of energy, a photon, has an amplitude in the 4th dimension, at 90º to all of our three dimensions. To understand how wavelengths correspond with energy, think of waves on the surface of water. First of all, a soliton wave is more like a tidal bore than a normal wave. However, the energy in a wave of a given amplitude is greater the shorter the wavelength. In fact, when a wave approaches the shallower depths near the shore, its wavelength does become shorter. Its energy reaches a peak when it breaks. This helps to visualize that the shorter the wavelength the greater the energy (at a given amplitude). Photons are elemental waves on the surface of space. Space is the real thing. Space wiggles due to the presence of photons as they travel over its surface. All the complexity of nature emerges from these ripples on space!

When a photon is confined to a particle (a very small black hole) its path is an integral number of its wavelengths (a small integer: 1, 2, 3, probably not more). Its path may be intertwined with other photons so that the energy (mass) necessary to produce a black hole is contained inside a small enough volume. An instance of this has been called a Planck particle.

The “tension” of space due to the presence of a photon falls off by the square of the distance from the photon, and the tension at any point in space is the sum of the contributions of all photons everywhere. The tension at a point in space is what steers any photon passing through that point. When photons are steered into tightly closed orbits, they exhibit as particles. When they are steered into larger chaotic orbits they are inside black holes of a larger size. The photons that we see are in very large orbits, indeed (our entire universe). They are on their way around a circuit of our universe, and they are only stopped due to an interaction (the one that modifies them in the act of being observed!). Simply put, gravity is the effect produced by the tension of space due to the presence of quanta.

Current physics has models for many of the bits and pieces of atomic composition, photon interaction with electrons, and other particle-particle, photon-particle, and photon-photon interactions. My descriptions here are meant to be a qualitative description of the foundations of the cosmos, not to be complete and rigorous models. However, I believe that models could (and should) be constructed from these descriptions.

In the cosmos at large, black holes are formed by any accretion of sufficient matter into a small enough volume. A black hole may “decay.” Particle decay may be observed in the laboratory. A collapsed star “decays” as a supernova (a small Bang). Larger black holes decay in what I call a Bang (but not quite a Big Bang). The first of these decays is the fission of particles. The 2nd and 3rd result in the formation of particles. If primordial black holes exist, it is likely they form as part of a Bang. The first stage of a Bang is the sorting out of whether particles or anti-particles prevail. As the cloud of particles from a Bang disperses, anti-particles annihilate each other to form photons which reform as particles, and so on. Particle formation in a Bang is only complete when this aspect is sorted out and matter prevails over anti-matter (or vice versa).

Let’s think about gravity wells, the extremes of which are black holes. Matter and energy may enter and leave them under different circumstances. Let’s review these circumstances and examine what goes on in each.

First consider photons emitted from within a gravity well. They appear red shifted to an observer higher up in the well, or “outside” it altogether. If the well is a black hole, light doesn’t escape, it orbits within the black hole until it interacts with another particle or photon. Time appears to go slower deeper in the well to observers higher in the well, until it stops altogether if the observer is high enough and the well is a black hole.

Every point in the cosmos is inside a black hole. That black hole is its universe. Photons aggregate into particles, particles aggregate into atoms, atoms into molecules, molecules into larger chunks, larger chunks into planets and stars, and finally all of the above into large black holes. Gravity is the force which binds at every level, but charge is the force that plays the most significant role at the atomic and molecular level. Charge fluctuates when a photon moves through space, but it remains constant when a photon is trapped as a standing wave. A standing wave inside a particle sized black hole exhibits a charge field to the outside of the particle. The quantum’s orbit is around the inside-outside axis of a volume (a 4-D concept).

All phenomena inside a black hole do not exist for an observer outside the black hole. The black hole is a phenomenon unto itself. Such a phenomenon has a life which may suddenly, and spontaneously end. This may occur with the predictable half-life of particle decay for black holes at the small end of the size continuum, and as a Big Bang at the large end. The half-life of a large black hole depends upon the configuration inside it, and there is probably no way of knowing this, so it’s completely unpredictable. In between (on the size continuum), the collapse and eruption of a black hole is a quick and predictable sequence of events. Black holes may contain: pure energy (particle black holes); plasma (a hot mixture of pure energy and particles); or (like the universe surrounding us), a relatively cooler mixture of energy and particles.

Quanta accrete. Suns form from gas and dust. Black holes form from primordial particles and energy, and eventually anything that gets near enough to them. When suns get big enough, they collapse into a black hole. When black holes get big enough, they overheat and explode. The two events may occur one right after the other, or a very long time apart. These event chains cycle throughout eternity.

Insight into 4-D topology can be gained by reading *Flatland*, by Edwin Abbott Abbott. For obvious reasons, Abbott may have identified with his main character A Square. The height of the book’s popularity was years after its first publication, when people began thinking about the 4th dimension in connection with Einstein’s relativity theories.

It is important to separate conclusions derived from models and conclusions based on experiments or observations that allow a model to be constructed. Many conclusions are derived from models because it is difficult or impossible in certain cases to perform experiments.

We know what happens to a projectile that is not interfered with in any way except for the gravitational attraction between it and the body away from which it is projected. If it is projected away from a body at less than escape velocity, it goes a certain distance away and then falls back. But what happens if it is projected at exactly the escape velocity, or if it is projected at more than the escape velocity. We haven't really tested these cases. In case (2) does it reach an infinite distance away with exactly zero speed, and in case (3) does it still have some degree of speed remaining when "infinitely" far away? What curves describe these orbits?

How about a Schwarzschild sphere? If a single particle occupied the entire universe, would that non-zero density define an almost infinite, but still finite, Schwarzschild sphere?

Answers to the above depend on our models, they have never been tested. My approach is to choose elegant answers and work backwards to build models that both predict these answers and conform to all the experiments we have done, and all the observations we have made. However, I reserve the right to seek and choose unorthodox interpretations of any observation.

In Newtonian celestial mechanics, a small body orbits a large one in one of three ways: An elliptical orbit, if the small body is moving at less than escape velocity, a parabolic orbit if the small body is moving just at escape velocity, and a hyperbolic orbit if the small body is moving faster than escape velocity (https://en.wikipedia.org/wiki/Escape_velocity).

If all matter is contained within a black hole universe, then all orbits would be ellipses.

Given particles with two or more quanta, or atoms with two or more particles, there is a probability that, over an interval of time, some constituent quanta have an interaction that does not cancel out. The half-life of a particle or atom (complex of particles) depends upon this probability. Sub-atomic particles are likely to be orbiting quanta that must exist together to have a stable configuration. Neutrons, for example, appear to have a half-life of about 15 minutes when not in proximity to a proton, but both protons and neutrons have very long half-lives when contained in the same (stable) atomic nucleus. Only very specific configurations of quanta are stable enough to exist as particles, and only a very few of them are able to endure for very long periods of time.

When quanta orbit within a particle (black hole), an oscillating charge may become a standing wave that exhibits itself as positive to the outside of the black hole and negative to the inside, or vice versa. A standing magnetic wave exists at 90º to both the axis of motion and the axis of charge. This is why some particles have a charge and why moving charged particles produce a magnetic field. Both of these effects contribute tension to space, and affect local pathways. The tension they contribute is orders of magnitude greater than the gravitational tension of the basic soliton wave. And, because they have a polarity, a quantum traveling in the vicinity of these tensions may be steered either toward or away from the direction of higher tension. The point is that unless a quantum is confined inside a particle, its charge and magnetic vectors cannot exhibit as standing waves, and therefore they only come into play when quanta interact directly with one another (moving at the speed of light).

Time

Time is a counting of "ticks" or waves passing a point. Time allows wave-like phenomena to be compared. Time is a ***measurement*** because a wave has a length, and a fraction of the cycle of a single wave must be included in the measurement. We always attach units to a measurement of time to relate one count of waves (the ticks of an atomic clock) to another, for example to standard minutes, hours, days, or years. Time may be ***measured*** backward, but it ***marches*** only ***forward***, because every photon and particle in the universe is a ticking clock, and clocks don't tick backwards! They don't untick!

Is time only ***measured*** by a quantum, but not ***experienced*** by it? I think so. Time is experienced by a collection of particles. Particles experience events. Time is a measure that relates the occurrence of one event with a subsequent event. Time is a concept that emerges from complex chains of events. Co-located, clock-like phenomena occur at the same rate. A clock slows down as it move faster through space. A clock moved farther through space will register fewer ticks than one that has not been moved. Clocks higher in a gravity well tick faster than those deeper in.

Each quantum has a mass/energy that directly relates to its wavelength. This is a fundamental and universal relationship. When quanta morph into particles, they follow tightly looping pathways. The natural frequencies of particles (or quanta) are exactly the same everywhere in the universe. Anything that affects the frequency of an electron will have an identical effect on electrons everywhere in similar environments. This is what keeps all clocks in the universe scaled to the same "standard" time.

Time and speed are related. Speed is a measure of how fast a point is moving through space. Time is a function of how many ticks of a clock are measured. Two identical quanta are two clocks ticking at exactly the same rate. However, relative differences in speed affect the ***observation*** of time. If two observers each shine a yellow light, and they are moving toward each other fast enough, each will see the other's light as blue. If they are moving away from each other fast enough, each will see the other's light as red. Blue light has more ticks per unit of time than red light. Thus, each observer measures the other's time as speeded up (blue light) or slowed down (red light).

"Going back in time" would require the ability to wall off a volume of space within which ticks would continue to occur, but outside of which the rest of the universe would "untick" in order to run backward. That's pretty farfetched. Another farfetched concept is a “wormhole.” Ordinary matter and energy can push and shove the hypersurface (to a very limited extent), but two distant regions cannot be forced together. The aggregation of a galaxy’s worth of matter into a super black hole would cause nothing more than extreme tension in a vicinity of space. The mass/energy involved in such a black hole is beyond anything that we could ever hope to control, and even so, it would not produce a wormhole and pinch two distant regions into contact with one another. There are no shortcuts through the universe. But, the “time travel” and “faster than light” stories are not over yet.

Understanding Reality

I have proposed a model for the cosmos with a fairly simple basis. From this I have derived one phenomenon after another from the implications of the basic premises. All is derived from the topology of a waveform superimposed on the substrate of space. Given infinite extent and infinite quanta at any non-zero and non-infinite density, the current Cosmos could evolve.

All constants are derived from the basic properties of the substrate of space—how its tension dictates the speed of light and the degree of electro-magnetic and gravitational force. Given the current Cosmos, it is not clear whether it is infinite, finite, or unbounded due to curvature. It is not clear whether the properties of the Cosmic Substrate change over great distances, or within nested universes. Perhaps the collection of constants and topological relations are tautological. If each relates to the others, then only relationships count, and absolute values do not exist. Each quantum is a package of attributes that all vary together. All of our measurements of anything can be translated back into the fundamental units of a quantum. Momentum is perhaps the most fundamental. Momentum is mass × velocity. Velocity is speed and direction. Speed is distance / time. Mass is energy / c^2 (where, c = the speed of light). All of these concepts measure the real world, and all relate to the structure of a quantum.

For now, let’s accept that space is the fundamental substrate and that quanta are the fundamental units of existence, and see where it leads us. If no quanta were superimposed on its surface, space would be perfectly static. Time would not exist. Extent would have no meaning. With the introduction of a single quantum, both time (defined by the quantum’s frequency) and extent (defined by the quantum’s wavelength) now exist and can be used to take measure. That quantum would travel in a perfectly straight line, marking time forever. However, introduce a second quantum, and each would feel the tension due to the presence of both. Each would travel in a giant curve orbiting the other. These orbits define a photon sphere inside which the quanta are trapped, and whose radius is defined by the mass/energy and density of the two quanta. Their universe is now bounded, and they will cycle regularly forever. Now, introduce a third quantum. This brings chaos to that universe. Three or more quanta do not cycle regularly—their future becomes unpredictable. Eventually, two of the quanta will intersect, and two different quanta will depart from that point of interaction. Given a (near?) infinite space filled with a (near?) infinite number of quanta, ***our observable universe*** shows us what quantum interactions can evolve into over time (because we can observe that they have!).

The questions are: Is space infinite, or is it unbounded due to a curvature in the 4th dimension? How many quanta does the universe contain? How did the universe evolve into the matter and energy we find assembled around us? Without quanta superimposed onto space, time and any notion of extent do not exist. The simplest answer is that the cosmos is infinite, and that it had no beginning and will have no end. However, our observable universe may well be a finite black hole within this infinite cosmos, and the earlier definition of eternity does give a beginning and an end to the concept of time. Or is time also circular? Perhaps an endless amount of time allows every quantum in the cosmos to return to its starting position.

The most significant departure of this model from the Standard Model is that it defines space as a fixed framework and a kind of “substance.” This appears to be a step backward, contradicting Einstein’s Theory of Relativity. However, it provides an explanation for an observation that relativity does not. This observation is the fact that the faster a clock moves, the slower it ticks—the faster it ticks, the slower it must be moving—this implies that there must be such a thing as absolute rest. An experiment might be devised to shoot particles in various directions at various speeds and observe rates of decay until one speed and direction gives the fastest rate of decay. This speed and direction would allow the calculation of absolute rest and, therefore, relative speed and direction with respect to the cosmic framework. If it is true that the farther a clock has moved the fewer ticks it will register, the question is: Farther with respect to what? Clearly, there must be a reference point. If one of two clocks ticks slower than another, then there must be some clock at absolute rest that ticks faster than all others not at rest.

Keep in mind that we should always go with the simplest adequate explanation (one that is not too simple).

In this model, a particle is a configuration of quanta. Each particle has an absolute speed and direction with respect to a stationary substrate. Every quantum is a soliton wave that always propagates at exactly the speed of light over its local “surface” (whether or not it is contained in a “black hole” or through some other dense medium). Every quantum induces electromagnetic and gravitational tensions into space at each point in its vicinity. These tensions define the natural path along which a quantum propagates at that point in space. Thus, consistent with current ***observations*** (and several consistent with current ***interpretations***), I make the following statements:

- Motion is absolute with respect to the fixed framework of the cosmos.
- A black hole is defined by its Schwarzschild radius (a given amount of mass/energy enclosed in a given volume).

- Black holes may be of any (finite) size. Two quanta all by themselves, may orbit within a black hole whose size is determined by Schwarzschild's constraints (as the energy of the two quanta approaches zero, the size of this black hole would approach infinity, however neither of these limits is ever reached).
- Particles are the smallest of black holes. A particle is one or several quanta propagating on closed paths that are only one or a few wavelengths long.
- The (closed) path length (or orbit) of a quantum may not be less than its wavelength. This implies that quanta with insufficient energy cannot form particles. In fact, given the small number of stable particles, only a few (extremely short) wavelengths emerge as viable.
- A quantum whose orbit is equal to its wavelength forms a standing wave particle. Its orbit is a plane surface that moves from the center of a sphere to its outside and back again around an axis in the 4^{th} dimension.
- The natural path of a quantum, due to tensions in space at its point of passing, is a curve. The curve is always an ellipse (or the special case of a circle) or a segment of one of these. It may be an extremely small ellipse, or an extremely large one. However, the paths of most quanta are interrupted by interference with another quantum (cutting short a complete orbit). If the ellipse itself is moving through space, the curve becomes a spiral (or a succession of loops).
- One or more quanta, whose paths come very close, interfere with each other (or themselves). Each quantum induces tensions into the space around it, and each travels a path defined by the tensions it encounters.
- Interactions between two quanta involve an exchange of energy and momentum. This creates two new quanta with total energy and momentum equal to the originals. One of the quanta may continue to reside within a particle, while the other may be emitted from it. Examples of interactions are when electrons change subshells, when many quanta produce an interference pattern, and when nuclear decay occurs.
- Black holes are stable to the degree that internal quantum interaction does not occur (but simply cancels out).
- Black holes erupt if sufficient internal quantum interaction does occur (this prevents the possibility of a singularity).
- Collections of particles (atoms and molecules) are collections of standing waves, some exhibiting charge to the outside of their respective black holes (but apparently none emit magnetic polarity).
- I conjecture that there are no mono-pole particles because a soliton waveform can only close on itself at a zero point. Charge is at zero at the leading and trailing ends of a soliton (photon), but the magnetic field is at its maximum when the charge is changing at its maximum, and that's when it crosses zero. A soliton is a 4-D structure. It vibrates in 3-D space, but one of these dimensions is its direction of propagation, and one is its amplitude at 90º to our 3-space. The "black hole" particle formed by a single closed loop soliton exhibits a positive or negative charge to the outside, and the reverse to the inside, of the particle. The magnetic field must cancel out on or near the surface of the particle.
- Every quantum propagates on the substrate of space at exactly the speed of light, but in loops and curves determined by the electromagnetic-gravitational tensions induced at each point in space due to the presence of all the quanta in the vicinity of a given point (the degree of tension induced by a given quantum falls off by the square of distance from the center point of the quantum).
- Although quanta propagate over the hypersurface of space at exactly the speed of light, changes of ***tension*** are not necessarily propagated from point to point at the speed of light. The sum total of presence might be felt everywhere at every instant. The effect of distant changes would mostly cancel out, and would be vanishingly small in any case.
- The pattern of tension in space, although changing over time, is ***felt*** (defining a path) at every point in space.
- The only thing that travels at the speed of light is light (quanta). Changes in the tension of space are due to an entirely different mechanism, and therefore might be propagated at a different speed, perhaps instantaneously. In fact, are we certain that all wavelengths of light travel at exactly the same speed?
- The propagation of tension is a local effect (with magnitude falling off by the inverse square); quanta are not created, nor do they vanish; they move through space at the speed of light, and their effects on space move with them.
- Black holes may nest. The largest is an observer's local universe. The smallest are sub-atomic particles, within which further nesting does not occur.
- Nothing within a nested black hole (or particle) may be observed from the outside; the behavior of such a black hole is limited to its total decay, or the spontaneous emission of a fraction of its quanta. Based on the black hole's mass, these events correspond over time to a half-life probability. Unknowable events (but statistically predictable in some cases) inside a black hole determine what actually happens.
- Space itself may be compressed to determine the paths taken by quanta, but it almost certainly does not form a singularity, or can it be formed into a "wormhole." The compression of space is entirely different from the gravitational collapse of matter. The extreme aggregation of quanta in a collapse causes the compression of space in its vicinity.
- Depending on the density of matter in the space through which quanta propagate, a tiny fraction of their energy is lost by the faint interactions they have along their journey—this causes quanta traveling long distances to undergo a red shift.

The simplest particle is a quantum in a standing wave that exhibits a positive or negative charge to the outside of its black hole. Imagine a quantum inside such a black hole propagating at the speed of light in a tiny circle. This only occurs if the circle is at rest with respect to the substrate of space. If the particle is moving, its quantum travels in loops.

Visualize a hoola-hoop lying on a playing field. Imagine a mouse (representing a moving quantum) running around just inside the hoop. It runs at a constant speed. Now imagine moving the hoop across the playing field. It takes the mouse a little longer to make each circuit of the hoop. It still has to pass each point around the hoop, but it also has to make up the distance the hoop has moved forward. Each circuit is circular only from the point of view of the hoop. If we observe the mouse tracks on the ground, we see that its path is a series of loops. It's still the same distance from side to side, but if an observer pans in the in the direction that the hoop is moving forward, the back side of each circuit catches up to the front side, and the distance from the front to the back of each circuit is less than the width of each circuit. This shrinkage from front to back is the "reason" for relativistic length contraction. The extra time it takes for the mouse to get from one side to the other (its frequency) is the "reason" for time dilation. If we move the hoop faster over the field, eventually it is moving as fast as the mouse can run. The mouse has to run at almost full speed just to maintain contact and keep up with the hoop. It has little extra speed to make any progress around the hoop. Its circuits around the hoop take longer and longer (this is analogous to what happens when a particle approaches the speed of light).

There is a fundamental (or topological) equivalence between moving through space and remaining stationary within a gravity well. A gravity well is defined as a spherical volume of space surrounding a large collection of quanta concentrated at its center. Space in such a region is "pinched" in toward the center due to the presence of the quanta (the very existence of their waveforms). Thus, along a radial line from center to outside, decreasing as the square of the distance, space is compressed, and this compression produces the same length contraction and time dilation effects as motion through space. In the case of motion across space, the hoola-hoop appears to the mouse to be shorter from front to back; in the case of gravity it is shorter along the line to the center of mass due to the compression of space.

Objective reality involves hidden variables. When two quanta (soliton waves) impinge upon one another, they interact. When this happens, given the descriptions of the two quanta going in, and the details of their interaction, the two quanta coming out could be perfectly described. Mass/energy and momentum are always conserved. Hidden variables can only be inferred by observing subsequent interactions. And because any observation disturbs the subject, Heisenberg's uncertainty defines the limit of accuracy an observation can attain. A hidden variable is any description of a quantum before an interaction. It fundamentally ***cannot be known*** (without completely changing the subsequent interaction).

The Twin Paradox

On the day this was first written, the Wikipedia explanations attempting to explain this paradox were not conclusive! It was orginally thought to be a paradox because it appeared to violate the theory of relativity by demonstrating a difference between the two reference frames of the twins. However, the Hafele–Keating experiment proved that this difference does exist. The following story illustrates why the difference exists.

The *twin paradox* takes its name from two hypothetical twins that live in our distant future. The twins have just reached their 21st birthday. One of them has graduated from the space academy and is being sent out into space. The other will remain on Earth. Twenty years pass (from the point of view of the twin who stayed on Earth). The space traveling twin now returns. During the time he was away, this twin has spent much of his time traveling at relativistic speeds. His spaceship, capable of matter to energy conversion, remains under 1 gravity of acceleration the whole time—the same as the twin that remained on Earth. When they get back together, they realize that they are no longer the same age! The Earthbound twin is about to celebrate his 41st birthday, but the space traveling twin is just celebrating his 37th birthday!

This is called the Twin **Paradox** because speed is supposed to be relative. After a year's acceleration, the twins are moving away from each other at nearly the speed of light, but shouldn't each point of view be the same? Einstein's special theory of relativity predicts that each twin will observe the other aging more slowly. Einstein's general theory of relativity predicts that there is no way to tell the difference between acceleration due to gravity and that due to a rocket engine. The twin on Earth is subjected to exactly 1G of acceleration during the whole twenty-year period. Likewise so is his twin aboard ship. A single year in a spaceship at 1G acceleration would get the ship moving at a sizeable fraction of the speed of light.

In this thought experiment we keep the traveler and the stay-at-home subject to the same force of acceleration, so any final difference in elapsed time is not explained by acceleration. The only difference is the distance traveled. According to our story line, this trip takes 20 years from the point of view of the Earthbound twin, and 16 years from the point of view of the twin who makes the trip. The purpose of this thought experiment is to reconcile the Twin Paradox, Einstein's Special Theory of Relativity, and the Hafele–Keating Experiment that demonstrates that a 4-year time difference between the two twins is a plausible outcome.

One way to "see" time going slower in another frame of reference would be to watch the other frame of reference in real time using a TV signal. Let's say that a very powerful TV signal is broadcast from both earth and the spaceship, so that both twins can constantly observe one another. As the spaceship travels away from Earth, it picks up a signal that is red shifted. Exactly the same red shift would occur from both twin's points of view. Half the trip involves the ship traveling away from Earth, and half involves it coming back. The ship emits its TV signal in the opposite direction of its motion during the trip away from Earth, and in the same direction on its way back. Earth neither moves away from its signal, nor does it move toward it (relatively speaking).

A red shift is observed when relative motion is apart, and a blue shift is observed when motion is together. Each frame of a video is recorded by a certain number of wavelengths of the signal. When it is played back with a red shift, it takes longer for the same number of wavelengths to generate each frame. Conversely, when blue shifted, the video runs in fast motion, it takes less time to receive the wavelengths for each frame. When the frames are replayed on the TV screen, it is equivalent to seeing time slowed down or speeded up in the reference frame being viewed.

Now, consider ***how long*** each twin observes the other with a red or blue shift. The twin on the spaceship views the twin on Earth red shifted for the 1st half of the trip (going away), and blue shifted for the 2nd half of the trip (coming back). This is not true for the other twin. The twin on Earth views a red-shifted signal not only while the spaceship is moving away, but for several years after that as the signal is emitted at the distant end of the trip just before turnaround. This part of the signal still takes a few more years to travel back to Earth. After the ship turns around, the signal it emits will indeed be blue shifted when it reaches Earth, but the ship will be racing after it, and the elapsed time that Earth sees a blue-shifted signal will be much less than half the total time the ship was away. Therefore, the twin on the ship sees ½ of the signal red shifted, and ½ of it blue shifted. The twin on Earth sees ***more*** than half red shifted, and ***less*** than half blue shifted. This means the twin on the ship watches ***more*** time elapse for his twin, and the twin on Earth sees ***less*** time elapse on board the ship. Time ***literally does*** progress more slowly aboard the ship! Paradox resolved.

What might we observe from aboard that spaceship as it approached relativistic speeds? If the distance to be traveled were 10 light-years, in a little over a year, we would be approaching the speed of light. Time at our destination would appear to be elapsing faster than our subjective time (light coming from it is blue shifted). The 10 years our voyage requires might elapse in only five years of subjective time. If we had a way to measure subjective length, it would diminish in our direction of travel so that it might appear that our voyage measured only five lightyears (according to our instruments).

That is how we would observe it. In effect, the 10-light-year distance would be less than that according to our subjective measurements. If we took 8 of our subjective years to travel the 10 light-years forth, and then back again, totalling just 16 years, we could nevertheless find that 20 years had elapsed back home when our round trip was completed.

Thus, it is possible to travel around the universe, never ***actually*** faster than light, but at ***subjective*** speeds ***much*** faster than light due to length contraction in the direction of travel. Any apparent speed is possible (when apparent speed is computed as objective distance divided by subjective time). And, by the same token, any amount of time travel into the future is also possible by completing a long enough trip. However, with reality based on this framework, there is no way to contrive time travel back into the past, nor is there any way to make an “objective round trip” at speeds greater than light.

Science fiction fans should not despair, however. At least three areas are still open: Total matter to energy conversion, artificial intelligence, and modifying the genome. Harnessing any of these could lead to an unlimited number of plot lines consistent with reality. On the other hand, I’m quite dubious about transporter beams, and uploading or downloading a human brain. Exploitation of actual principles may lead to scientific plot lines—breaking the rules is “science” fantasy.

Relativity

Up to this point, all the phenomena I have described, and all those that stem from them, are tied to the fixed properties of the cosmic substrate and the fixed properties of the quanta that exist as soliton waves traveling on pathways over its 4-D surface. Here, “fixed” means rules that apply everywhere and forever. When the pathways of quanta are small ellipses or loops, they appear as particles. When they travel on large elliptical pathways, they appear as photons. Thus, I’ll use the terms ***particles*** and ***photons*** (or light) to refer to the two ways quanta normally (and interchangeably) exhibit themselves.

Weak and strong nuclear forces are quantum interactions inside a particle. Phonons are simply quanta transmitted through a vibrational medium (one particle “contacting” the next). A force is a tension induced into the substrate of space by the presence of quanta. That due to the wavelength is gravity, that due to the dynamic of its vibration is electromagnetic. Electromagnetic energy, the strong and weak nuclear forces, normal particles, subatomic particles, and phonons are all phenomena that involve quantum interactions.

Einstein’s special theory of relativity states that the speed of light is the same in all directions regardless of motion through space. His general theory of relativity states that acceleration and gravity cannot be distinguished.

Imagine the quanta inside a particle black hole. They move exactly at the speed of light, but suppose the particle itself was moving at nearly the speed of light. With the enclosing black hole moving at nearly the speed of light, the paths of the quanta themselves would become almost straight lines. If this narrative is accurate it’s obvious that particles can’t exceed the speed of light! The quanta that make them up always move at exactly the speed of light!

Normally, a reference frame for demonstrating special relativity is a moving laboratory in which we can conduct the Michelson-Morley experiment that compares the speed of light moving along two different paths, one at 90º to the other. What is found is that the speed of light is invariant. If light propagated through some kind of ether (or over a stationary surface as in my analogy), and the laboratory itself were moving through (or "over") that medium, the speed of light would be different when measured along the direction of travel as compared to a direction 90º to that. The fact is, the speed of light is measured to be the same along both of these axes. To account for this, Lorenz proposed that length contracts along the axis of motion. If this is true, and it is accepted by science that it is, the math is the same for Einstein's Special Theory of Relativity and for the hypersurface that I'm suggesting. As the speed at which a particle (loop) moves approaches that of the speed of light, the length contraction formula says that the length in the direction of motion is the (rest length) × [(speed of light – speed of particle) / (speed of light)] ^{½}. Thus, length in the direction of motion approaches zero as the speed of a particle approaches the speed of light. Time dilation is the reciprocal of this.

This description is one of geometry. The geometry of a moving particle is affected by the fact that particles consist of quanta that always move at a fixed speed (that of light), and when particles approach that speed, their internal geometry is affected. This same geometry obtains when particles are subjected to a gravitational field. The effects of a gravitational field on a particle can, in fact, be equated to the effects of particle motion. If a particle occupies a static position within a gravitational field, such as it would when resting on the ground, the space it occupies is compressed by the tension caused by gravity, and this compression leads to the same geometry within those particles as it would if they were speeding through space.

General relativity states that no physical experiment can distinguish between the forces of gravity and acceleration. Matter not being acted on by an outside force travels through space on a naturally given pathway. Gravity (in the macro case) is the phenomenon that produces the pathway. When matter ***is*** acted upon by an outside force, it is forced to move onto a new pathway. Simply put, any particle forced to move out of its natural pathway must be acted on by an outside quantum. A stone lying on the ground is accelerated by interaction ***with the ground*** and out of its natural pathway which would have it fall into the earth if there were no ground in the way. Thus, the force felt by an object accelerated by the equal and opposite reactions of a rocket, and one being prevented by the ground from falling, are actually identical situations.

On a different basis, gravity and acceleration ***can*** be distinguished. Given two identical labs, one in a rocket ship, and the other on the surface of a planet, the difference between acceleration and gravity can be distinguished with a very sensitive plumb bob. Gravity is directed toward the center of a sphere, so the plumb bob moved across the perfectly flat floor of a lab on a planet would point in slightly different directions on one side of the floor as compared to the other. In the lab aboard the rocket ship, it would point perfectly straight down at every point across the floor.

Acceleration is defined as a change in speed or direction. It can be produced in two ways: by orbiting (falling freely) in a curved pathway, and by direct interaction with other quanta. Gravity, charge, and magnetism produce the shape of the local hypersurface—they ***define*** the pathway over which an object naturally moves—they are not "actions" from a distance. The effect of the presence of a quantum on the local pathway falls off as the square of the distance from the quantum. This effect, at any point in space and instant in time, is summed over all quanta in existence.

When a particle changes ***speed*** along a pathway, relativistic effects may come into play. When the force accelerating a particle is counteracting gravity (for example, a particle resting on the ground), speed does not change (due to the force, itself). Thus, no change in relativistic effect. Notice that electron "walls" opposing one another are due to a static impasse, not a dynamic interaction like that produced by a rocket engine. Dynamic interactions are exchanges of energy and momentum that occur when two quanta intersect, whether that be inside a particle (as mass), or outside of one (as a photon).

Everything (at every moment of time) consists only of quanta moving on their own "natural" pathways. Two or more quanta may interact. When that occurs, two or more quanta depart from the interaction. Their individual characteristics may change, but their total energy and momentum remain the same. Quanta may enter into an interaction in the form of either photons ***or*** particles, and they may leave as either after the interaction has taken place.

Particle stability, including the stability of the nucleus, and the orbital stability of an electron, depends upon how well the constituent quanta can avoid interfering with one another. Consider an electron. It is strongly attracted to the positively charged nucleus, but its elliptical orbit may not be less in length than its wavelength. For the 1st electron added to an atom, a circular orbit is defined. With additional electrons, their mutual repulsion complicates the paths they must take so they don't interfere with each other. When a photon impacts an electron, or the nucleus of an atom, both source and target quanta morph. Energy and momentum are conserved, but the new paths of each quantum are subject to the constraints that they must not interfere with themselves (by following an orbit less than their wavelength), or with each other. Thus, orbitals and subshells are defined for electrons, and similar pathways are defined for particles within an atomic nucleus. Essentially, it is interference with itself that causes an electron to emit a photon when forced into a lower orbit by its attraction to the nucleus.

When the properties of quanta, the topology of space, and the mechanics of quantum interaction are fully considered, and integrated into a proper mathematical model, a Theory of Everything will be achieved.

Dark Matter (Nothing Special?)

When we look out into the universe, we only see streams of photons coming from point locations. They are emitted by hot matter. Cold matter emits few if any photons that we can detect. Other regions of space may contain black holes, and these contain matter and energy that is also "dark" from our perspective. Most of the photons we detect come from hot matter, stars for example. However, cold matter, unless it is at absolute zero, emits photons at very low energy levels. Could rocks and cosmic dust actually be the source of the background radiation we observe? Black holes emit no radiation at all. Could these sources account for the matter we think is missing from the universe? And, surely there are better explanations than an inflating universe that requires dark matter or energy (which have never been observed) to explain it.

The Big Bang (Maybe Not!)

Here, I would like to speculate on the effect that extreme distance has on a photon. The energy that reaches us from distant points in the universe consists of radio waves, microwaves, heat, light, and energy at even higher frequencies. Most of the energy we receive is emitted from stars, near, distant, or extremely far away. The energy received from a star falls off as the square of its distance from the observer. A "ray" of light from an extremely distant source is an almost continuous stream of photons that contains different wave lengths. The energy of an individual photon, all things being equal, should remain constant as the photon travels through space. If a photon is emitted by an object traveling toward or away from an observer, energy is added or subtracted due to the difference in velocity between the object and the observer. A gain in energy is called a "blue shift" and a loss of energy is called a "red shift." These are Doppler shifts like the pitch of a siren emitted from a police car that is moving toward you, or away from you. Sound shifts to a higher frequency if the source is approaching you, and to a lower frequency when the source passes and goes away from you. In the same way, light shifts to a higher or lower frequency making yellow light turn blue when the source is coming toward you, and red when it is moving away from you.

Now, consider a single photon that has traveled a half-dozen, or so, ***billion*** light-years from the star that emitted it. It was emitted at a particular wavelength along with other photons at other wavelengths. Light emitted by a star has a pattern. At certain wavelengths more light is emitted; at others, virtually no light is emitted. This is due to the predominant elements present on the surface of stars. When hot, each element emits light only at certain frequencies called spectral lines. This causes different types of stars to have different signatures. A given ***type*** of star generally emits a given ***amount*** of light.

As a given type of star is farther away from us (distance estimated due to its fainter light), we find that its spectral lines are red shifted. This means that the photons emitted by stars far away from us have lost some of their original energy, and the farther away they are, the more energy they have lost. The accepted theory is that this is due to a Doppler shift. The stars farthest away from us are red shifted the most. This line of reasoning leads to the conclusion that the farther away a star is, the higher the speed it is moving away from us. However, another explanation for this might be that the red shift is due to a simple loss of energy because the intervening space is not entirely empty. When a stream of photons travels billions of miles through space, the occasional interaction might deflect a few of the photons out of their normal paths, but most of the stream is left intact. Random near misses and gravitational wandering, however, could account for a very tiny amount of energy being transferred to other mass and energy in the vicinity of a photon's path. Over a long distance, the cumulative effect on a photon would be a red shift—a loss of energy. The energy transferred to any quantum in the vicinity of its path (and the energy each acquired) would be immeasurably small. But, total energy would be conserved!

Again, which is the simpler explanation? A red shift due to an inflating universe, or one due to energy being lost over billions of light-years of travel? Unless the latter can be disproved, the Big Bang theory is uncertain, and the conclusions involving an initial singularity, a creation event, or an expanding universe should all be questioned.

Force Fields

Quanta are soliton waves that propagate at the speed of light along a path that is a function of the quantum itself and the tensions in the substrate over which it travels. Tension has both a gravitational component and an electromagnetic component. Its gravitational component is constant (and moves with the quantum), but its electromagnetic component oscillates over space, or remains constant as a standing wave within a particle). Thus, the gravitational components of multiple quanta sum within an entire vicinity. The electromagnetic components, however, only sum when they are standing waves within a particle. Otherwise, they cancel out over a very short distance. But, they are very much present as part of a quantum waveform propagating through space and occasionally interacting with other quanta.

Until the paths of quanta intersect, each quantum travels on its own pathway and has no interaction with the rest of the universe except for its tiny contribution to the gravitational tension of space. When the paths of quanta do intersect, an interaction occurs. Total energy (mass) and momentum are always conserved in every quantum interaction.

Quantum interaction has 3 forms: photon–photon, photon–particle, and particle–particle. Photon–photon interactions involve wave effects and two new photons replacing the originals. Photon-particle interactions include a particle being broken into its components, being dislodged from an atom or molecule, or simply being deflected into a new pathway. Finally, particle–particle interactions include nuclear fission, nuclear fusion, chemical reactions, particles being deflected into new pathways, and particles being attracted or repelled due to electromagnetic effects. The resulting particles and photons always add up to the original sums of energy and momentum.

Just as a pathway is defined by the presence of ***all*** quanta in the universe, it is largely dictated by far ***fewer*** quanta close by. Let's enumerate the ways pathways may be constructed, and the ways the paths of quanta may be affected by them. First, the presence of any quantum in a vicinity is summed with all the others to produce the phenomenon of gravity. A pathway curves in the direction of the center of this mass/energy (following an elliptical path).

Charge and magnetism are phenomena that oscillate at an amplitude restricted to the immediate vicinity of a quantum, they sum to zero unless the quantum is confined inside a particle and produces a standing wave that may be summed to produce static electric and magnetic effects on the local pathway. In any case, the pathway of a photon or particle causes like to curve away from like, and opposites to curve toward each other. This makes the oscillating electromagnetic effect of a photon-photon interaction very sensitive to slight differences. However, the static electromagnetic effect of a particle-particle interaction is quite predictable—it's what we usually call a force field.

Let's say that a star collapses and forms a black hole. Its density might increase catastrophically, but overall mass remains the same, and the gravitational effect on distant space remains unchanged. However, if two massive stars (or black holes) were in orbits around each other, local pathways would be massively redefined. Consider a point at some distance away from this system, in the plane of their orbits. As the stars line up with this point, or line up transverse to it, a small oscillating effect on the pathway at the point would occur. Transmission of this effect would not necessarily be at the speed of light. There is no reason to believe that "stress oscillations" of space travel at the speed of light. Energy into a stress equals energy out (think of a spring). Stresses are transmitted differently than soliton waves, so there is no reason to expect them to propagate at the same speed. Here, we are talking about the transmission of tension. This is different from how sound is transmitted through air, or a pull is transmitted through a rope. That depends upon a wave of particles, each affecting the next. The proposition of gravitational waves notwithstanding, tension in a continuous medium like space (comparable to no physical analog), might even be transmitted instantaneously. The mega-events (involving black hole interactions) assumed to be the cause of "waves" of gravitation might be more analogous to the propagation and detecting of earthquakes.

Remember, a quantum is not a particle or a point, it is a waveform propagating in a certain direction at the speed of light. Its ***natural*** path is whatever tension in space it encounters. Gravity, electrostatic, magnetic, and electromagnetic effects are all forces that arise due to the "surface tensions" of space. These stresses on space are induced by the very ***existence*** of quanta in a vicinity that affect the natural paths at every point, and therefore the path over which a quantum naturally travels. The "pinch" effect (gravity) of a quantum's waveform is very much smaller than the charge or magnetic effect. We know this because this is what we measure, not what we deduce from this model. This model is only suggested as an analogy to help visualize some of the aspects of space and the quanta that propagate on its hypersurface. I doubt we are even equipped to visualize the complete 4-D nature of this reality, and the actual topology of a quantum. But, if we keep modeling it, maybe we'll eventually find the math and topology that fits together completely with all our observations.

Let's review the various forms of energy that we know about. All forms equate to quanta (kinetic energy) and tension in the hypersurface (potential energy). A quantum traveling on a large elliptical pathway is a photon. A quantum traveling on a small closed pathway is a particle. Think of black holes larger than atoms as super particles (or primordial black holes).

When a particle is forced into a different pathway, it must be due to encountering different tensions in space. A photon is typically emitted or absorbed. Under certain conditions, this is also called a ***phonon*** being ***transferred***. When a large number of moving particles are enclosed in a volume of space, they impact with one another. Each impact emits a photon that may be absorbed by another particle.

This exchange of photons, and the photons that escape from the surface of this enclosed volume is called heat. Heat is the statistical sum of this photon exchange. To the extent that photons are radiated away, a given volume of particles loses heat, and cools. To the extent that particles are moving faster on average in some parts of the overall volume, heat will be dispersed to cooler regions. Particles may be contained within a fixed volume in (at least) two ways: by gravity or by the walls of a container. The faster the average particle is moving (the higher the heat), the more pressure the particles exert against the walls of the container, or the more they force the volume of their confines to increase. A mass of hotter particles will occupy more volume given the same force to contain them as compared to a mass of cooler particles. This means the hotter mass will be displaced by a cooler mass if the two masses don't completely mix together, or are prevented from doing so. The effect is that the hot mass tends to rise away from the center of gravity. Of course, this process of convection only takes place when a gravitational field is present (if not, they would either dissipate, or have to be otherwise contained).

Carry these concepts to two extremes: Extreme gravity and extreme heat. If the mass of particles (and free quanta) were inside a black hole, and the heat rose high enough, the pressure inside the black hole would eventually cause expansion to exceed the Schwarzschild radius, and the black hole would erupt.

Another source of complexity is the fact that individual quanta are polarized. They have a transverse axis along which their electromagnetic effect occurs. Some sources of light radiate a larger variation of frequencies or polarities than others. Light radiating at a single frequency is called coherent light. Light radiating at a single polarity is called polarized light. When quanta pass into a region of space occupied by particles, the nature of the particles interacts with the frequency and polarity of each quantum. Certain frequencies and polarities may pass through the particles (and be refracted), others may be absorbed, and still other may be reflected. Much of the science of physics is concerned with these various interactions. Virtually every modern gadget is possible due to the technologies developed from the science of these interactions.

Spacetime Connection

Time marches on throughout the universe according to the same rules everywhere. It is measured by the synchronized "ticks" of quanta. These ticks may appear different to different observers when differences in speed or tensions in space are involved. The effect this has on particles is that they age more slowly when traveling at relativistic speeds. The effect on photons is a red or blue shift when the source and observer are at different relativistic speeds, or levels of potential energy (different positions within a gravity well).

Every particle and photon leads an independent existence until it interacts with another particle or photon. Interactions result in new combinations of particles and photons. For the most part, particles are changed by directing them into new pathways. Photons are changed into new photons with the same total energy and momentum. Some high energy interactions result in particles being created from photons, or vice versa. Energy and momentum are always conserved.

Two different points in space are connected by the amount of time it takes a photon or a particle to travel between them. Larger volumes of space are connected only by greater amounts of time. Connectivity in a densely populated volume of space is problematic—communication may be indirect, corrupted, or interrupted. A "vicinity" might be defined as the amount of space that is connected—it always involves a measure of time and a function of density. Higher density makes interactions within the volume fuzzy, less predictable.

Two examples illustrate this. A "line of sight" vicinity is a low-density volume of space and the time it takes light to travel from source to observer. A "chemical" vicinity is a relatively higher density volume of space and the time it takes particles to travel from one point to another, or their effects via "sound" waves. When conditions are highly controlled, for example on a circuit board, a vicinity is defined as the volume and time it takes for electric signals to connect. Thus, time is a factor in the definition of a vicinity. This is why spacetime has been called four dimensional—but it must be clearly understood that one of these dimensions—time—is completely unlike the other three. And, the actual 4^{th} dimension of space is nothing like time considered as a 4^{th} "dimension."

Although space is continuous, events are only connected by interactions. An ***event***, as I define it here, is a hit or near miss between two quanta. It takes time between the event that puts a particular photon or particle on one pathway and the event that puts it on its next path. An observer can only know about any event by observing a second event and making an inference. Everything leading to a given event must be inferred from what can be observed from subsequent events. Everything we observe is a "final event" (the last in a very long causal chain). This is a fundamental limitation dictated by the nature of our cosmos. We can only infer what happened previously, and we can only extrapolate what will happen next. Each event produces new particles and photons from equivalent old particles and photons. Due to our (forced) ignorance of the details, statistics must be used to handle most observations and predictions.

Think of it this way. Light travels about 1 foot per nanosecond. Our present moment is the sum of all events that took place ***n*** nanoseconds ago and ***n*** feet away from us. That describes a receding surface of spacetime that is connected (by the speed of light) to a given point in space and time. Spacetime is also connected by all the particles that interact with us, and these travel (usually) much slower. When these events are factored in, our present moment includes a lot more events signaled to us by particles arriving at our point in spacetime. Spacetime connection is quite hard to visualize! Our concept of ***now*** includes only events that took place a certain distance away at a certain time in the past—it also includes closer more recent events and older events that took place farther away. In the lab, we can only sort out very simple events. In real life a huge number of events are happening simultaneously and continuously (and we are unaware of virtually all of them!).

Further Thoughts about Time

One Analogy for Time

Think of reading a written narrative (a book). The reader can drop in on it at any point. Up to that point various groundwork has been laid. After that point more things depending on that groundwork will unfold. As the reader reads, there is a sense of past and future. The reader has no ability to change the sequence described in the narrative. The reader's only choice is to start and stop reading, or to move to a different point in the narrative. Only the author gets to choose the original words, but after publication, this choice is not available even to the author. A book describes a very simple, very small universe. A book's narrative is only one thread at a time, unlike the vastly parallel threads, almost infinite in number, that make up our own universe and the events that unfold within it.

A Precise Definition

The nature and definition of time has never been given a precise definition by science, since that would require acceptance by a majority of scientists and philosophers, and that has failed (in my opinion) to occur even though humans have been giving thought to this subject for thousands of years. Nevertheless…

Time is the unfolding of events. Events unfold at a "uniform" rate throughout the universe due to the nature of matter and energy. Particles (matter) and photons (energy) interact by moving through space and colliding. All such interactions throughout the universe have a fundamental synchronicity that relates any two sequences of events. The laws of physics (involving rates and wavelengths) describe various kinds of event sequences, the synchronicity that ties them together, and conditions that can affect synchronicity.

A sequence of events led to our current observation of our surroundings, and every observation we make depends upon a sequence of events. We have given names to certain cyclical events and event sequences that synchronize. They complete in lockstep patterns with each other. Among these are our units of time.

Time Revisited

Quanta do not "experience" the passage of time, they define it. Time is the measure between interactions and events. The decay or eruption of a black hole (or particle in the limiting case) is an event. An interaction is one quantum crossing paths with another leaving two different quanta in their stead. Quanta emitted from within a gravity well, and quanta emitted from a moving frame of reference, are subject to a red or blue frequency shift. The frequencies of all quanta in the cosmos are synchronized and form universal clocks that measure time. These clocks are slowed by motion and acceleration. Both of these conditions have the effect of "compressing" space. The limit of this compression is reached when motion reaches the speed of light, or when density reaches a limit that results in particle annihilation.

Time Inside a Black Hole

All quanta inside a black hole (of any size) orbit within the confines of the black hole. The shape of a black hole is ellipsoidal. Its surface is not well defined. One definition of its surface is its event horizon, based on Schwarzschild's calculations. Another definition might be the outermost paths of any quanta orbiting within it—a photon sphere. The dynamics inside a black hole may be observed from outside the black hole as changes in the tension of local space. Electromagnetic and gravitational tensions may emanate or pulsate from a black hole into nearby space.

Time only has relevance to particles, photons, and their interactions. The lifespan of a quantum is from one interaction to the next. The statistical lifespan of a particle (or a black hole in general) is its half-life, or the amount of time that gives it a 50% chance of decay (due to internal quantum interaction). Its actual lifespan ends when it does finally decay. This analysis implies that inside a very stable particle, a black hole of extreme smallness, time does not exist. Inside particles that do decay, time is measured by its statistical half-life. Inside larger black holes, time is measured using the normal metrics of relativity. Let's examine the dynamics of energy in motion for more insights about how they define the passage of time.

Tensions in the substrate of space determine the paths along which quanta propagate. When the quanta in a vicinity all move along elliptical orbits within that vicinity, we call the vicinity a black hole. When small enough, with few enough quanta contained in it, we call the black hole a particle. Nevertheless, the quanta follow paths defined by the tension in the space around them. The closed elliptical path followed by a quantum must be equal to, or a multiple of, its wavelength. A quantum following a path whose length exactly equals its wavelength is a standing soliton wave (moving or at rest with respect to the space around it). This is a very stable configuration, and a quantum may interact with itself to produce this result (absorbing or emitting another quantum to make up any difference).

What do the tensions around a soliton wave signify? The amplitude of the wave is at a right angle to all 3 dimensions of normal space. It has two attributes that vary together: wavelength and energy. Given the speed of propagation (the speed of light), its frequency is defined by its wavelength, and its mass equivalent can be calculated from its energy. Together with its direction of propagation, its momentum is defined. As the wave passes a point in space, the tensions in space at that point fluctuate. Directed toward the inflection point of the wave (its amplitude) is a gravitational component that remains constant as it moves through space on a path dictated by other tensions in the space that it encounters. Ideally, this is an ellipse that defines the boundaries of the local black hole that encloses it.

Let's try, as best we can, to visualize a soliton wave in 4-space. Up and down represent the amplitude of the electric field, left and right represent the amplitude of the magnetic field. Forward represents the direction of propagation through space, and inside-outside represents (as best we can envision) the soliton's amplitude in the 4th dimension. Now, envision an elliptical path through space whose axis is parallel to the amplitude, or at right angles to all 3 normal dimensions of space. This is the 4th dimensional analog of an elliptical line on a piece of paper circling at right angles to the plane of the paper.

The soliton waveform represents the quanta's mass/energy due to the tension it induces into the space surrounding it. This tension is constant, but it surrounds a moving point (the center of the waveform). As this moving point passes by a point in space, the tension on space at the stationary point defines the electromagnetic and gravitational fields at that point.

A quantum occupies a toroidal hyper-ellipsoid in 4-D space. Its length (long axis on its line of propagation) is equal to its wavelength. This line may be a toroidal ellipse, or a segment of a toroidal ellipse. Its height and width are the amplitudes of its electrical and magnetic fields. Its 4-D limits are the soliton waveform itself. The path of a quantum is toroidal when it is moving through space; if its path through space were elliptical, the entire path itself would be stationary loop in space.

When the path of a quantum is a toroidal ellipse, the quantum is a standing wave; it appears as a particle. One or several quanta may form a particle. Depending on how their orbits intertwine, their charge fields may reinforce each other to exhibit either a positive or negative charge. An even number, within the same particle, exhibit no charge. An odd number do exhibit a charge. Very few orbital configurations are stable, so the number of particles that can be formed is limited. The most stable elementary particles (electrons, positrons, protons, anti-protons, and neutrons) may combine to form atoms of matter and anti-matter.

The half-life of a particle may depend on the probability of an outside quantum interacting with it, or the complexity of its internal quanta and their orbits. For example, there appears to be a discrepancy in the measurement of the half-life of a free-standing neutron. In one experimental format, its half-life appears to be about 10 minutes, and in another format, it is more like 15 minutes. Could these two formats provide different probabilities of interaction with an external quantum? Protons and electrons, on their own, appear to have ***very*** large half-lives (decay has never been observed). When combined with neutrons into ever larger atoms, instability creeps in apparently as the result of the wrong number of neutrons. An example is hydrogen. It has 3 isotopes with 0, 1, and 2 neutrons each. The first two are stable; the third has a half-life of over 10 years. A stable atom is one with a half-life that has not been measured. An isotope is an atom that is either stable or has a large half-life. With any other number of neutrons, the atom does not appear to form (or maybe it has an extremely short half-life). Particles and atomic combinations of them are not easily formed. They are formed only within the furnace of a supernova (or some larger Bang). Perhaps every combination with fewer than a thousand particles has been formed, but only those in the periodic table (and their isotopes) survive for any length of time.

Just as electron subshells have stable configurations based on coherently orbiting quanta, so do the quanta that orbit within an atom. Each subatomic “particle” is a quantum in a stable orbit that is coherent with itself and other quanta within the same atom. All the quanta within an atom or a constituent particle contribute to the dynamic tension that dictates the orbital pathways that each quantum has to follow, and all quanta must remain coherent with one another (their orbits may only intersect without interacting).

When 3+ quanta are bound in a vicinity, each contributing to the paths followed by the others, the system is chaotic. Atomic decay appears to be a random event that is measured by a half-life statistic. This means a chaos factor is involved. The chaotic event might stem from a quantum impinging from the outside, or a chaotic internal orbit. If the event is due to the outside, different environments should produce different data. Also, if an unobserved particle is the cause, there must be another unobserved particle that is part of the effect.

Uncertainty

Uncertainty arises when you can’t perfectly predict something. Statistical formulas are one way to compensate for this. However, neat formulas don’t always exist (or haven’t yet been discovered). This is where the “butterfly” effect is used to explain uncertainty. This “effect” is due to the limits of measurement accuracy. I thought it was first discovered when a weather model, being run on a computer, was interrupted. All the intermediate values were punched into cards, and then it was restarted with the same cards as inputs. The only problem was, the values were converted from binary to decimal and then back again. The tiny loss of accuracy this represented was about as large as the effect of a butterfly flapping its wings on the other side of the world. And yet, it made a noticeable difference to the result. Measurements are always approximate. There is uncertainty at the limits of measurement. The greater the complexity of a calculation (the number of multiplies and divides it contains), the more error there is in the final result. A calculation could begin with 50 significant digits, but contain thousands of operations, and result in a single digit of accuracy, or even less.

It turns out (to me, at least) that one of the first people to really understand computation, Alan Turing, described this type of uncertainty in his 1950 paper, ***Computing Machinery and Intelligence***. {See text @ bottom of p. 440 for his description of chaos (the butterfly effect). Other subjects: Imitation game; Equivalence of computing machines.} http://mind.oxfordjournals.org/content/LIX/236/433.full.pdf+html?sid=07e63a2f-4e95-48a1-8037-719866e414a2

The uncertainty known as Heisenberg’s Uncertainty Principle is based on Planck’s constant which is known to fewer than 15 digits of significance, and depending on the units used to express it, has from 15 to 34 zeroes between the decimal point and the first significant digit. Thus, it is a ***very*** small measurement. It relates the theoretical limits of accuracy when both of various pairs of physical phenomena are measured: The more accurate one can be measured, the less accurate the other can be measured. This is because the thing being observed must interact with the observer (or the observer’s lab equipment).

Statistical formulas encapsulate chaotic behavior that has a regular aspect to it. The accuracy of such a formula depends on the law of large numbers, which loosely says that the more events involved, the more accurate the statistic. An example of a statistical formula is the calculation of pressure given the temperature of a gas. Pressure arises from molecules hitting the walls of a container. The higher the temperature, the greater the impact of an average molecule. The effect of an individual molecule, however, is highly random. It is only due to the fact that decillions of molecules are contributing that a constant pressure can be accurately measured.

Quantum uncertainty is another type of uncertainty. There should be no mystery about it. It arises due to the fact that we simply can’t know the state of things, or the starting conditions, at the level of individual quanta. Observations result from interactions whose outcome depends on this knowledge, and we fundamentally can’t have it. Therefore, the result of any quantum outcome is uncertain and statistical in nature. We may only deduce what is going on by making a large number of observations and applying statistics to interpret their results. As for a single event, we can’t know anything about it until a subsequent interaction (event) involves the result of the first event. The fact that we are uncertain over an interval of time is no reason to believe that both of two possible outcomes are true in the interim—that Schrodinger’s cat is both alive and dead until we open his box. Two entangled particles are much the same. When the status of one is discovered, the status of the other is determined—predetermined, in fact—not communicated faster than light from one particle to the other.

If we did know the exact value of all hidden variables in the universe could we (in principle) calculate any future state? The answer is no. The exact value of all hidden variables is an infinite string of digits for each quantum in the universe (near infinity times an aleph one infinity). Any less than this and the butterfly effect takes over and makes garbage out of any result. Another way of looking at this is that any future state would take longer (***infinitely longer!***) to calculate than actually allowing the future to play out in real time.

Determinism

The definition of determinism is that "every future state of affairs is the inevitable and necessary consequence of an antecedent state of affairs." Determinism could only be possible if an algorithm could exist that, given the starting conditions of the universe at some point in time, could compute its condition at some later point in time. If, on the other hand, for some reason, this is fundamentally impossible, then the universe itself is not deterministic.

Consider one of the simplest cases of chaos: 3 bodies in an otherwise totally empty region of space orbiting around each other. It would take infinite accuracy of the starting conditions and an infinite number of computations, each with an infinite degree of accuracy, to compute an infinitely accurate result after a time, $t > 0$. The motion of each body depends on the mass and position of the other two bodies at any instant in time. One instant later, new paths have to be calculated. To obtain perfect accuracy, an infinite number of calculations is required. This is impossible, even in principle. The only way to determine how the universe is going to behave is to let it run its course and see. There is no short cut. No conceivable algorithm and data collection scheme can do better than to give an ***approximation*** of future conditions. The butterfly effect guarantees that any prediction of the future is nonsense after an interval that depends on the nature of the calculations to produce it.

Heisenberg's Uncertainty principle shows us the limits of measurement in the actual universe, and proves that we cannot collect the data to predict quantum events. Determinism only states that a future state is exactly determined by a past state. This hinges on whether quantum uncertainty is fundamental, or only due to hidden variables. Let's say it is due to hidden variables. Does this mean we could, in principle, collect all the data on all the matter and energy in the universe, and then predict any future event? For the reasons given above, the answer is still no.

If the universe were deterministic, the Many Worlds Theory would have no basis at all. Imagine restarting the universe at some antecedent moment, like restarting a chaotic weather model, the antecedent state cannot be perfectly specified, so a later universe would necessarily be different. You could restart it any number of times, and each restart would produce a different future (just like the Multiverse of the Many Worlds Theory).

Actually, in reality, chaos builds up almost instantly. Now factor in evolution. Driven by chaos, evolution has assembled intelligent agents (ourselves, for example). If the above logic is sound, each entity created by the process of evolution is unlikely to an almost infinite extreme. Therefore, what evolution creates is infinitely unpredictable. Determinism is a red herring. The simple phrase "antecedent state of affairs" is the crux. Since you cannot go back in time, there is fundamentally no such thing. There can only be a crude guess. A ***fundamentally*** indeterminate future rests on any antecedent state of affairs that can ***fundamentally*** only be approximated.

For determinism to be possible, there must be zero uncertainty over all the causal chains of interaction at the quantum level and higher throughout the entire universe and across all of time. If the impossibility of this, and the principles of evolution, are to be considered valid, we have to conclude that determinism is not possible, and (for different reasons) that free will is a reality. When Jacques Monod pointed out in 1972 that everything is the product of chance and necessity, he ignored the question of free will (still a sticking point with many philosophers today), and he should have emphasized that ***both*** chance and necessity are ***always*** involved.

Entropy

Entropy is a measure of the amount of energy in a system ***not*** available to do work. The Thermodynamic Laws relate to Entropy:

0. If (a, b) and (a, c) are systems in thermal equilibrium, then so are systems (b, c).
1. The total energy of a closed system is a constant.
2. The entropy of a closed system increases over time.
3. The entropy of a system approaches zero as the temperature of the system approaches absolute zero.

The Thermodynamic Laws are restricted to "closed systems." It would be more accurate if this restriction stated ***small*** closed systems. When systems are large enough for gravity to produce stars, or to include the entire universe, or to contain an evolutionary process, these laws either don't apply, or they are simply irrelevant. Let's (at least) call any system that contains a star an open system, and therefore a system where the Thermodynamic Laws are irrelevant. In fact, closed systems can only be ***approximated***, and only then in special (localized) circumstances, such as in a lab. Outside something like a lab, systems are always open.

Entropy increases over time due to random micro events and the consequences of the law of large numbers, increasing the chaos within any closed system, and causing it to approach internal thermal equilibrium. This fact defines the "arrow of time" that allows breakage, but never assembly (and other analogous paired concepts) to occur spontaneously. So, what accounts for things "spontaneously" getting ***assembled***? First, chemistry, then evolution!

Chemistry

Chemistry is the science that describes the behavior and interaction of atoms and molecules (complexes of particles). Our region of the universe has "chosen" the electron as the charged particle to be turned loose. Many believe this was simply due to chance. Necessity forces a choice between the electron and the positron, because both cannot inhabit the same region of space to any significant degree. They annihilate each other. With all positive charge buried inside the nucleus of an atom, chemistry, electronics, materials technology, biochemistry, pharmaceuticals, genetics, optics, food processing, and many other fields of science and technology are based on the behavior of electrons and photons interacting with each other, and with atomic nuclei acting pretty much as anchors.

Chemical reactions occur when atoms bind together to form molecules, or the atoms within molecules break apart or become rearranged. All chemical reactions depend upon temperature and concentrations of the different molecules or atoms present in small volumes of space. Individual atoms with their electrons bear a similarity to suns with their planets. Molecules have very different characteristics (as we shall see in a moment).

But, first let's compare atoms to planetary systems. Both have a large central body relative to a host of tiny things in orbit around them. A sun has its planets, and an atomic nucleus has its electrons. Planets orbit their sun due to gravity, which defines the pathway along which each planet moves. The pathways in the region of a planetary system change as the planets' mass shifts around. The pathway that explains how a planet orbits its sun is a function of where all of the mass in an entire vicinity is distributed at the moment. In a planetary system, charge is distributed very evenly. Virtually every electron is very near a proton, so ***charge*** has virtually ***no net effect*** on overall planetary pathways.

In a planetary system, every particle of mass contributes an effect to the pathway structure. Each particle moves on its own pathway, and only electromagnetic forces ***between*** particles can force a particle to move onto a different pathway. Because large masses of particles (planets, moons, etc.) are constantly moving around in a planetary system, the paths taken by these bodies are continually changing. During the first few billion years, before a planetary system "settles" down, some of its bodies collide with each other, or are flung out of the system. Eventually, the bodies that remain travel along paths that minimize their mutual interference.

In an atomic system, pathways are primarily determined ***only by charge***—gravity has virtually no influence. Electrons orbit protons contained in the atomic nucleus. Like planetary systems, pathways in the vicinity of an atom are also determined by the location and velocity of every quantum in the vicinity. Here, positive and negative charges attract each another, and like charges repel. Both of these forces decrease by the square of separation, which means they approach infinity as separation approaches zero. Notice that this also holds true for gravity, but the constant for the gravitational force is many orders of magnitude smaller than the constant for charge. Also notice that both types of force operate in both cases: atomic and planetary systems. The fact that charge dominates an atomic system and gravity dominates a planetary system does not mean that the other type of force cannot occasionally come into play.

The dynamics of both types of system are complicated by the fact that particles cannot orbit faster than the speed of light, and with a massive sun, or very close separations between atomic particles, this speed may be approached.

The dynamics of an atomic system are complicated further by the fact that all of its particles exhibit wave properties when they move, and that means they can interfere with each other and with themselves. Electron orbits are only stable when all interactions between electrons cancel out. This fact explains why only certain patterns of electrons occur around atomic nuclei.

As atoms contain more and more protons and an equal number of electrons, the electrons form into shells. The innermost shell can have only one or two electrons in orbits of exactly one wavelength. This keeps a single electron from interfering with itself, and two electrons in similar orbits, but in different directions, from interfering with each other. Each shell above the first is also defined by integral numbers of wavelengths. With more wavelengths, more orbital patterns are defined. Two orbits are always circular, but 6 additional orbits exist along the 3 orthogonal axes that go through the nucleus. In the next shell, these orbits describe a figure 8 with one loop around the nucleus and one loop to the outside. This shell can hold up to 8 electrons. The third shell is like the 2nd, except that it has 6 orbits with a loop on both sides of the nucleus. The fourth shell can have up to 18 electrons (2+8+8, the total of the 1st three shells combined). This takes us up to atomic number 36 (which is a gas called krypton). After this point, elements become more rare. All of the elements necessary for life occur before this. That is, elements with fewer than 36 protons and electrons. The fifth shell again contains 18 more electrons, and includes silver, tin, and iodine amongst others. The next two shells each contain 32 electrons. The first of these includes platinum, gold, mercury, and lead. After lead, all but 4 atoms are unstable.

Essentially, an electron shell is a region of distance from the nucleus in which a fixed number of stable (non-interfering) orbits can form. These orbits twist and turn due to near impacts between electrons, and because of their speed and chaos their orbits appear as clouds, rather than neat ellipses like planets. Electrons in different shells move faster or slower due to kinetic and potential energy tradeoffs, but unlike planets in orbit around a sun, electron orbits must be coherent (all wave effects must cancel out). Since orbits must be of integral wavelengths, when an electron moves into a higher or lower orbit, a fixed quantum of energy must be absorbed or radiated to adjust the wavelength.

Quanta are radiated away from atoms or larger objects in the form of light, heat, and radio waves. Photons may be radiated singly, or as streams. A stream of photons can consist of a wide band of frequencies, a narrow band (as from an LED), or a single band (as from a LASER). Typically photons are emitted from hot matter, such as stars. Typically this hot matter has a signature mixture of elements whose electrons emit photons in a pattern of certain frequencies and not others. The pattern of these frequencies produce what are called spectral lines— useful to deduce the mixture of atom types being observed. The instrument used to make this measurement is called a spectroscope.

The most important feature of an atom in chemistry is the number of electrons it has. Electrons occupy shells with fixed numbers allowed in each shell. That means the outer shell of an atom is either exactly filled or only partially filled. If it is partially filled, it may have one or very few electrons in its outer shell, it may be about half filled, or it may be lacking one or two electrons. When the outer shells of two atoms can be shared so that both outer shells are exactly filled, they tend to lock together in a state of lower potential energy (giving off photons in the process). Thus, two hydrogen atoms, each with a single electron in its outer shell, lock together as a hydrogen molecule sharing both their electrons in a stable outer shell with two electrons. Carbon, with its 6 electrons has a filled inner shell of 2 electrons, but a half-filled outer shell with only 4 electrons. It can share these with another carbon atom to fill the outer shell with 8 electrons. Or, it can share 2 electrons with each of two carbon neighbors in a long (or circular) chain of carbon atoms. These are just two examples of the myriad ways that atoms can combine with one another to form molecules of infinite complexity. Another example is water, H_2O. The single electrons of 2 hydrogen atoms lock together with the 6 electrons in the outer shell of oxygen to form the most common molecule on Earth.

The chemistry inside our neurons and muscle cells is the most basic machinery of our brain and body. Organic molecules are long and complex. Each atom is locked to the next to form a very rigid structure of a very particular shape. If a bit of energy is absorbed or given off, this shape may change. It is often the shape of a molecule that determines how it will react with another molecule. Little organelles called mitochondria are the power supplies inside the neuron that enable it to fire, and inside muscle cells that enable them to contract. These structures transfer bits of energy that affect molecular shape and therefore the chemical reactions necessary for nerves and muscles to operate.

It is important to understand ***what*** a chemical reaction is, and that ***if a reaction can occur, it will occur***, given that the appropriate set of atoms or molecules come within “striking” distance of each other. This latter condition is based on probability. The greater the concentration of the necessary molecules, and the closer they are to the proper temperature, the higher the probability they will change their combination from a set of reactant molecules to a set of product molecules.

There are basically three types of chemical reactions: association (giving off energy), disassociation (absorbing energy), and recombination (either absorbing or giving off energy). The first type involves two separated molecules (or atoms) coming together into a stable configuration, and emitting a photon which appears as heat when absorbed by some nearby molecule, making it move faster. The second type is the reverse of the first: a stray photon impacts a molecule and knocks it apart into two separate molecules or atoms. This type of reaction absorbs heat (or energy). The third type may either generate or require heat when it occurs. In this case, two molecules enter into the reaction, and two different molecules depart from it. Many chemical reactions consist of a series of steps each involving a basic type of reaction. A product of one step forms a reactant for the next.

Think of all the atoms and molecules in a vicinity as bodies in almost constant collision with their neighbors. Most of these collisions are 2-body collisions between elastic objects that are ***not*** perfect spheres. The mechanics of incidence and departure from each collision also involves changing the spin of each object in every collision. When chance brings two molecules together, and necessity causes them to combine, a new molecule is formed. Probability is involved on the micro level, but the law of large numbers allows reactions to be fairly predictable on the macro level.

Now, let me try to describe the picture that should be forming in your mind. All particles are combinations of quanta moving at the speed of light in spirals, which are determined by their own local pathway. Any pathway at a point is a function of all quanta in its vicinity. Why spirals? Because spirals are the result of a point rotating in a circle as it moves through space. The particles within the nucleus of an atom consist of a number of quanta circulating without interfering with one another. Likewise, the electrons in shells around the nucleus of an atom consist of individual quanta circulating without interfering with one another (at least to the extent that each remains in a stable orbit).

There are two further things that you need to imagine. One is when a quantum is absorbed or emitted by an atom. The other is when kinetic energy is exchanged for potential energy and vice versa. First, an atom emits a quantum when an electron moves into a lower shell. Picture how a satellite would move from a higher circular orbit into a lower one. First, it must lose some velocity with a retro-thrust. Then, its orbit would change to an ellipse that would have a lower altitude on the opposite side of the planet, but a higher velocity. On the opposite side it could lose some of that velocity with another retro-thrust that would keep it at the same altitude with just enough velocity to establish a circular orbit. These two retro-thrusts represent potential energy being turned into kinetic energy. The resulting orbit is at a lower altitude (lower potential energy) but at a higher velocity (higher kinetic energy). If all the electrons around an atom are already in their lowest possible orbits (a condition that never occurs with planets around a sun), then that atom will not emit any quanta due to an electron dropping to a lower shell. A quantum could also be emitted from the nucleus of an atom—the result of a nuclear instability, but this is a very rare event at the temperatures and pressures we normally experience. Both of these event types conserve momentum, so the entire atom would move to a new pathway after either of these events.

Another type of event is when an atom absorbs a quantum of energy. Here, if the energy is enough to expel an electron, or kick it into a higher shell, that event may occur. In any case, two quanta radiate away from the point where two quanta come into contact. Energy and momentum are always conserved. The two quanta involved are the original impinging photon, and the quantum "inside" the electron (or nucleus) that is hit. The rule is that quanta travel on a pathway at the speed of light until interfered with by another quantum. That event causes two different quanta to depart from the point of interaction with energy and momentum being conserved.

The chemical properties of an atom depend upon how many electrons it has. Typically, all of the electrons are in their lowest states, with the lower subshells filled, and only the outer subshell partially filled. When the number of electrons is such that the outer subshell is exactly full, the atom is a noble gas (it doesn't react with itself or any other atom). All other atoms have a partially filled outer shell, and this enables them to react easily with another atom whose outer shell has the number of "missing" electrons. Together, they can share electrons so that each of their outer shells is exactly filled. Other possibilities for chemical reaction are that at higher temperatures some atoms may have electrons missing from a lower shell and present in the outer shell, or that three or more atoms may form a stable combination. Stability results from the total combination of electrons giving off some potential energy in the form of photons and going into a configuration with more kinetic energy contained in a smaller volume. This occurs when the molecule formed has given off some heat. When a molecule absorbs energy (heat), it may lose stability and break apart. The number of ways that these simple principles can occur is unlimited. Therefore, so is the science of chemistry (although, as it becomes more complex, it morphs into the science of microbiology).

Although no ***atom*** with more than 100 protons and electrons is stable, an unlimited number of ***molecular*** configurations are possible. That means (since there is no such thing as infinity in the real world) that there are still more configurations waiting to be assembled somewhere in the universe. A very large number of different molecules occurs naturally. These are relatively simple molecules, and the conditions they require are likely to occur somewhere sooner or later. However, with no limit on ***possible*** molecules, this leaves a huge number that require very special conditions to occur. Very improbable molecules become more probable, and eventually inevitable, when precursors and catalysts are available. Precursors are complex molecules obtained by the occurrence of an earlier reaction. Catalysts are molecules whose presence assists a reaction to take place, but which are left unchanged after the reaction. With a catalyst an extremely ***improbable*** reaction may become extremely ***probable***. Early evolution involved the (very slow) production of simple precursors and catalysts. As these became available, more complex molecules became more probable. The overall complexity of molecules increases over time if favorable conditions remain relatively unchanged. Here on Earth, the rules of chemistry, operating for several billion years, evolved life and intelligence.

Information

Information is defined by Webster as "the communication or reception of knowledge or intelligence." ***Communication*** means a transmission from one place to another. Communication includes the processes of encoding and decoding. A communication is intelligence encoded into matter or energy for the "purpose" of transmission, and decoded back into a form that is "usable." ***Use*** implies intention within a context. Non-living things may form a pattern or environment or a context for living things. But, like the living listener in the forest who hears the sound of a tree falling, and without whom sound is only a dispersion of energy, it takes a living thing to combine the sound and the ***context***. A context is essential to the concept of information.

Consider once again the word, "information." Think of two concepts that define a dichotomy, popular today, but barely recognized in the middle of the 20th century when information theory first began to be studied. This dichotomy is that of hardware and software. Software is information in its purest form. It can be encoded into matter or energy. When encoded into energy, it is non-corporeal. Software can be communicated. It can be "beamed up!" Hardware cannot—although, by using information and an appropriate mechanism, copies of hardware (including biological entities) ***can*** be assembled.

Information is always equivalent to a sequence of zeros and ones; it is digital. Any information can be encoded as a binary string of digits, but information theory doesn't address use and context. When information is separated from its use and context, it is nothing more than a string of digits. Information theory tells us how long a binary string needs to be to satisfy various constraints, but it tells us nothing about the meaning of a particular string of digits. Meaning is encoded into the use and context of a string of digits; it is not contained in the digits themselves.

An amount of information can be measured as the ***length*** of a string of digits. Information is associated with order, which is the opposite of entropy. Transmission and the recording of information involves analog representations, but information is essentially digital. It exists only as a function of how it is used. If it has no context or is not being used, its existence is meaningless. Think about that sound in the forest that goes unheard. It is not information. But, given the exact same phenomenon and an entity that can ***hear*** it, the sound ***becomes*** information to the hearing entity—a "live user of information." How does this happen?

The human ear is our second largest source of information. An ear conducts sound waves from about ten decibels in strength to over 100 dB, and from frequencies of about 20 Hertz to 20,000 Hz through an outer canal to an eardrum, and then via three tiny bones (called the hammer, anvil, and stirrup, which they vaguely resemble) into a spiral chamber filled with liquid, called the cochlea. Inside the cochlea, tiny hairs pick up the sound vibrations of different frequencies. They transmit frequency and amplitude information to the brain. This is done with fewer than a thousand neurons from each ear. However, each of the initial thousand probably goes to each of several thousand other neurons where the auditory nerve bundle enters the brain to process the sounds of speech and other information that our ears convey to us.

Other sources of information, the senses of taste, smell, touch, and sight will be introduced later. Any detectable signal in nature is a possible source of information. The magnetic field around the Earth has been shown to be detectable by certain species of life. The electric field generated by certain eels has also been shown to be a source of information to them. We know that light is an important source of information, and a few creatures can generate their own light to communicate with their fellows. Other portions of the spectrum, such as radio waves, could also be used (although no cases have been shown as of yet). However, we may have just begun to uncover other chemical signaling going on, for example using pheromones.

Order and entropy are physical concepts. They do not measure information. Information is measured in bits. Information transfer is measured in bits per second. Information must be encoded, transmitted, and decoded to be stored or transferred. Information is always stored on a substrate (like quanta themselves on the substrate of space). This involves photons or particles that can maintain a configuration over time. Photons are involved in the transfer of information. The ***maximum*** rate of information transfer is proportional to the ***frequency*** of the photons involved. However, the ***actual*** rate of transfer (the bandwidth) depends upon how the information is encoded, transmitted, and decoded. The actual rate is typically very much lower than the theoretical maximum. These processes are subject to the effects of entropy—information can be degraded, but with redundancy, it can be corrected.

Digital information is exact; analog constructs are always just approximate. The former corresponds to a counting number, the latter to a measurement. The former to software, the latter to hardware. Information is used and communicated in the analog world, but it defines and supports a digital world, a world of life and the living, two realities that give meaning to the concept of dualism.

Information is stored within all living entities (that we know of) in the form of DNA. It is very likely, but we can't know for sure, that some simpler type of molecule, such as RNA, was used to store information before DNA evolved. This is like the chemical reaction chains that may have once occurred in ponds billions of years ago. All evidence of them has been erased. The information that must be recorded for an entity to make a copy of itself, or get one made, is essentially a set of templates that enable the construction of all of the components of the entity to be assembled from available raw materials. A set of instructions is also necessary to direct the assembly process. A mechanism to decode and carry out the instructions must also be part of the copying process.

Biology is concerned with 4 families of molecules: nucleic acids (DNA and RNA), lipids, proteins, and carbohydrates. The first family primarily records the information necessary to produce building blocks consisting of the latter three.

Information other than in molecular form is also used within living entities. Members of the animal kingdom have evolved special cells called neurons. These cells store patterns of information captured from the environment by sensors. They also store patterns used to produce behavior. This form of information storage enables learning and intelligence, topics addressed in some detail later in this book. Proteins are used to transmit information, and they are used as building blocks constructed by using the information in DNA.

When artificial life forms evolve, it is doubtful that DNA and "wet" neurons will be used to store and use information. Information theory and the entire telecommunications industry are concerned with the science of information storage and transfer. New ways to improve channel capacity and correct errors are always being discovered. We shall explore this science more by seeing how the use of information evolved with life itself. We'll see how information is contained in patterns. We need to review how the ancient past became the present—where information is now literally everywhere.

Life

Life is the collection of all entities that can USE information. The phrase, "use information" is the key to this definition. An entity is alive if it can "use information" and it's dead (or lifeless) if it cannot. The word "can" implies potential. The word "use" implies action. Thus, life must be capable of action (being able to affect the world around it), and as we shall see below with respect to individual living things, life must be capable of action with a purpose.

Life and information emerge together. Information plays various roles in the processes, plans, and goals of all living things. In fact, living things are usually made up of smaller living components. When an entity no longer has the potential for using information as it normally does, that entity is no longer living—it's dead. When a component of a living entity cannot be said to be using information on its own behalf, we have reached the atomic level of life. This is a level smaller than a cell, but not too much smaller. An entity's components may live on after the entity itself has died, but the death process generally works its way down to the smallest constituents. These non-living (sub-"life") pieces generally become food for other living things—parts that may be ingested by another living entity.

When information is used, a process of communication is serving a function in a larger process. If you can see evidence of this pattern, and it is active and ongoing, then the entity that contains it is alive (it might even be a society of living things). This is precisely the extent and the limit of the definition. There is a conceptual and qualitative difference between mind and brain, or between the spirit and the body, just as there is between software and its hardware substrate—but there can be no existential separation of the two. A live spirit (a personality) requires a living body. Dualism is a fact!

Is a computer alive by this definition? A computer handles or processes information on behalf of its human ***user*** (there's that word again). But, does a computer "use" information on its own behalf? The same could be asked of a virus. A virus contains information that is copied by the molecular machinery in living cells, but does a virus "use information?" To answer "no" to those questions, we need to understand how ***use*** differs from contains, handles, or processes. ***Use*** involves a user and a purpose. When we say that a faucet ***uses*** a rubber grommet as a seal, we really mean that a living engineer used the grommet in the faucet's design to seal the faucet from leakage. Here there is a user and a purpose. Information does not exist without a user and a purpose. By this definition, on the other hand, any actual ***user*** of information is ***alive***.

Life is a living entity's opportunity to exist. Evidence suggests that our human existence is confined to lucid, waking moments, beginning when we are about one or two years old, and ending when our brain ceases a minimum level of functioning. No evidence suggests that life begins before conception, continues after death, or is associated with a soul.

Every living entity engages in two major type of activity: it attempts to consume other things, and it attempts to reproduce. Everything you or any other animal does is in a direct or indirect effort to consume or reproduce. As humans, we have evolved some pretty fancy variations on these themes, but we really haven't evolved any activities completely unrelated to them either. Name any activity, and without too much effort you can find one or both of these themes in it.

Consumption is more than eating. It can be a process of building, destruction, or even play. Exercise for its own sake is nothing more than the final stages of finishing a meal. Reproduction, as we all know, isn't only about sex. It could be as widely removed from that as, for example, the propagation of our ideas in the forms of speech, art, and literature.

In the process of consuming other things and propagating itself, each entity finds itself in more or less conflict with every other entity that it contacts. The negative effects of conflict can be overcome with the positive effects of cooperation. Most interaction combines both conflict and cooperation (such as coerced, and even forced, cooperation). Each entity has evolved a set of potential abilities. Through learning, abilities are developed into skills. The exercise of a skill allows an entity to interact with its environment and bring about whatever it can.

The evidence also suggests that life is governed by the necessities of physical laws operating in a background of random chance. At the outset of life, no two entities have precisely the same abilities and opportunities. Therefore, no two entities are "equal" any more than they are identical. Entities may, however, be treated equally by physical laws. They may also strive with more or less success for equal treatment in cooperation with other entities. The only rules are those adopted by the other entities that hold sway. The only "rights" are those afforded by cooperating entities. Your rights may vary.

An event is something that occurs at a given moment in time. An outcome is the result of a sequence of related events—events which affect, or are affected by, each other. There is seldom any meaning or purpose in the relationship of the events that produce an outcome. The process of evolution produces outcomes that may have meaning within their local evolutionary context, but evolution is not guided by purpose. Some outcomes are brought about with intention. These may be guided by purpose, but only if they are the result of the skill of a living entity.

Each living entity has a set of sensors and a set of effectors. An entity synthesizes information from the interaction of its sensors with its environment. It produces change by acting on its environment with its effectors. An entity's success involves its ability to use the information it is able to synthesize, and its skills to carry out courses of action. If its skills are sufficient to allow it to reproduce, the potential for those skills is propagated. Obviously, each of us is a living entity. Each of us has developed a set of skills. The outcomes we help bring about are a product of our skills, a large element of chance, and the cooperation of others.

Skills and the information used in their performance may be compared to computers and their software. It is not stretching the concept of "skill" too much to compare it to a program running in a computer. The effective use of a program to input data and bring about a result is very much like a skill in many respects. The behavior of inanimate objects may be predicted and described by the laws of physics, but the behavior of objects animated by skills (or software) cannot be described or predicted simply, by the same laws (or even the same kinds of laws), except on the most trivial levels. Physical laws are not disobeyed in any way, they are simply inadequate to produce a description, or an effective prediction, of the overall behavior of a living entity (or the results of a computer program).

Skills and information are not incorporeal, they must be instantiated in some combination of matter and energy. Matter and energy cannot be created or destroyed, although each may be converted into the other. New arrangements of matter and energy can come about in only two ways: By chance (following physical laws), or in a copying process. A copying process may, of course, introduce change. Change may occur either by chance or by intention. When intention is involved, the skill of design will have been a part of the copying process. When chance alone is responsible for change, the copying process is part of the standard evolutionary paradigm. Apart from new arrangements of matter and energy, the only things truly created are new skills (programs), and new expressions (information).

Every concept we have is in the context of our conflict and cooperation with our environment. The most important entities in the environment of any given entity are usually other entities of their own kind. These are the entities whose mutual cooperation is the most likely to have evolved and be the most effective for them. Yes, birds of a feather flock together for good reason.

Let's jump up the evolutionary ladder to the entities we call human beings, and focus on ourselves. Each of us tends to wonder from time to time, "Why am I here." "Does my life have any meaning or purpose?" "What should I try to do with my life?" "Why does that person have so much authority over me?" "Why is that person so wealthy?" "Why don't I deserve more?" "Why did that person have to die?" "What will happen to me when I die?" People seek comfort with these issues. Family, religion, schools, the government, and many other institutions try to contribute to that comfort. We cooperate in the hope of gaining that comfort. Much of our discomfort arises from conflicts involving these questions. It's all about our skill in handling conflict and cooperation.

Meaning and purpose are evolved. They are invented a bit at a time. They are copied and adapted. Your life has only the meaning and purpose that you copy or invent. You are not here for a reason; you are here by chance. You should take life one step at a time. Each passing year of your life has its special opportunities, as hinted at by something George Herbert once said: "He that is not handsome at twenty, nor strong at thirty, nor rich at forty, nor wise at fifty, will never be handsome, strong, rich, or wise." This is one wise man's opinion about a few of the things that typically matter to us.

Answers to the remaining questions are very complex. When an outcome is governed by a mixture of chance and necessity, with only an occasional pinch of intention or purpose, a complete explanation can be too complex for human understanding, and an incomplete answer may be too simplistic to do you any good. Looking for an easy answer to a complex question is like looking under a streetlamp for a key dropped elsewhere in the dark. And it's even more futile to search for a key unless you need what it unlocks, and know where to find the lock. First, increase, elaborate, and refine your needs. Then, learn about locks. Having done this, the keys will become evident in due course.

Reflection

This might be a good point to step back and reflect. We only ***know*** what we can sense directly. We are only able to ***do*** what we learn to do by manipulating our bodies, and building tools that we can use to extend our ability to manipulate. Perception is learning what various sensations mean. This leads to expectations and beliefs. Both of these can be false. In fact, it is almost impossible ***not*** to have a few false expectations and beliefs. When learning is effective, it improves the ability to recognize and distinguish what ***may*** be (the possible), from what ***must*** be (the very probable), and both of these from what ***will not*** be (the impossible, or extremely improbable).

A large fraction of humanity believes in one or more supreme beings and possibly an afterlife. The rest of humanity is either neutral or negative with respect to these beliefs. There is no basis to believe that the supernatural is a part of reality—belief in a supreme being requires belief in the supernatural—therefore there is no basis for belief in a supreme being. If there is a reality within which our perceived reality is only a simulation, we can only dream about it, and there is no reason to assume (or even conceive) that any part of our mind or body will ever contact it. So, the question is irrelevant. It should be given minimal, if any, consideration. Our reality is completely defined by the cosmic substrate of space, and the ongoing dynamics of the units of existence permanently inscribed upon it.

Free will, such as it is, emerges from intelligence, and the ability to recognize and take alternative courses of action. We know that intelligence can evolve from a primordial universe, because it has. With intelligence and free will, another factor (over and beyond chance and necessity) emerges: the ability to make choices given opportunities, and the ability to act on those choices. Free will brings with it responsibility. Irresponsibility may be punished, either by nature or by society.

The human brain is a marvelous and wonderful organ, but our individual reality is confined to its proper operation. When we are asleep, drugged, or dead, our brains cease to operate properly, and our consciousness is suspended. If our minds can't be brought "back online," our consciousness terminates forever. Consciousness, free will, responsibility—the very attributes that (we would like to think) define us as evolved life forms—are all built upon the faculty we call ***intelligence***.

The first awareness of reality is necessarily solipsistic. As one integrates more and more experience, this solipsism evolves into chauvinism for larger groups: family first, high school, ethnic group, and eventually (maybe) humankind. We are still chauvinistic as a species. My rejection of the supernatural or the extra-cosmic might be the ultimate chauvinism. Lacking any evidence to the contrary, or any apparent avenue to acquire further evidence, I believe the apparent reality of the cosmos is as far as one can go (my imagination, however, occasionally strains against this tether).

We can know things only in the present. We may remember past sensations, but memory is a process that ***reconstructs*** stored chunks of information. This is the only accurate way to describe our recollection of everything that took place in our lives before the present moment. Written and oral communication allow us to (poorly) reconstruct more of our collective past.

Early life was defined by the fact that a mechanism to record and play back the set of molecules necessary to build itself came together. This included templates for the required proteins, and instructions to order the building process—which molecules to produce at each step. Success at replication was the only criterion for survival. However, reproduction is only possible if the cell can accumulate "raw materials" from outside itself. A "self" didn't exist until a cell wall was defined. A membrane that could allow the passage of raw material into the cell from the outside had to evolve.

The next major steps on the path of evolution involved developing more efficient chemical processes, and building more effective cell membranes. At some point, more efficient ways to ***get*** raw materials evolved. A very primitive stimulus-response pathway developed. This marked, perhaps, the first way information was used. A simple receptor on the cell wall triggered a contraction of some type that propelled the cell in the direction of this receptor. This was the origin of smell. Contact with a certain type of molecule constituted the most primitive of sensory inputs. Sensory input is information. The process triggered is the encoding of this information, the first step in using it. Then, the signal is transmitted to other places in the cell, in this case, by the diffusion of special molecules for that purpose. When these are decoded by a response mechanism, it could trigger a contraction that would move the cell in the direction of the receptor (towards its food).

It might have taken a billion years for chemical molecules to become complex enough to evolve the first self-replicating cell, and another billion years for the first stimulus-response to evolve. A lot of time is necessary to compensate for the extremely haphazard trial and error processes that eventually yielded these results. But, a billion years is a ***lot of time***.

Further improvements in sensors and more sophisticated responses evolved almost exclusively within the animal kingdom. The development and improvement of nervous systems led to the evolution of intelligence. Single celled plants and animals are good examples of successful life forms making minimal use of information (and remaining at the very bottom of the food chain). Lacking a nervous system, they are incapable of developing intelligence. Therefore, they leave our story here.

The existence of a stimulus-response that could allow better access to "food" would, in effect, trigger an arms race between the different single celled entities that constituted early life on Earth. Life propagates where conditions favor it. Different places on Earth would have a dominant variation of this early theme, and other less numerous variations. The success of a life form is measured by the number of copies that it has been able to produce. Extinction occurs when there are no copies left. Extinction is forever.

Notice a subtle distinction. When there exists a direct lineage from an entity that has died to an entity that is alive, the earlier type of entity never actually went extinct even though the current descendants may clearly be of different species. Extinction describes a ***type*** of entity that has no living descendants (no more copies exist, no more offspring are produced).

After one stimulus-response pathway has evolved, two are better, and still more are better yet. At some point, a stimulus may result in an internal record, rather than an external response. Likewise, a response might be triggered by some combination of internal records. Both are ways to use information. Information is a pattern detected by sensors, which are typically located on the surface of an entity, but sometimes within an entity—this includes sensing internal conditions and sensing information previously recorded.

Early on, only specific molecules were sensed. This led to the senses of smell and taste. Probably the next sense to develop was a sense of pressure. This would enable orientation (up and down), as well as depth within the surrounding liquid. After this, perhaps a sense of light and dark developed. Each of these early senses probably involved the production of a unique molecule to act as its signal. Internal detection of these "transmitter" molecules triggers a given response elsewhere in the organism. Initially, responses probably included improvements in locomotion: turning and moving forward. They probably also included more efficient ways to ingest food and eliminate waste.

It is no accident that sensors, and hence, the largest mass of neurons tend to be on one end of the body and the function of waste elimination on the other. Bodies are specialized to move more easily in the direction of their sensors, rather than in the opposite direction. Life evolved to head for its next meal and leave its waste behind. Humanity is following this plan big time today. Look at how we separate our waste dumps from our living spaces and our food production areas. Look at how we head into virgin territory for raw materials and harvest first growth timber and new sources of fish and wildlife. Even global warming is a product of our waste, and finding a way to leave it behind is turning out to be a very tricky problem.

All life presently on Earth consists of surviving copies of a single early cell. If two or more cells arose independently, one, by being first or by being better, would always win out over the others (by competing with them for food, or using them as food). This statement is borne out by the observation that all life on Earth examined so far has a common genetic heritage. One is reminded of the fact that the electron, by being first, gained numerical superiority, and won out over the positron.

Everything that we are is a chemical arrangement of molecules. We exist on planet Earth—a warm body orbiting a medium sized yellow star in the suburbs of a galaxy we call the Milky Way. Each of us grew from a single cell, fertilized within our mother by our father, via a process of cell division until we were delivered into the world as an entity in our own right. We each spent our first few years learning the skills necessary to converse and contribute within today's human society.

Let's speculate on our reality and our limitations. I maintain that the value of life is not based on its ***potential*** (whatever that might be), but on how an entity relates and compares to other entities around it at the present moment. In particular, a single fertilized cell has about the same objective worth as any other single cell. A human 2-year-old child has about the same worth as an adult chimpanzee (although the parents of the child or chimp might judge differently!). All I'm saying is that we should give a similar amount of respect to both.

Human adults may exist in many different states. Most are healthy and productive. Some are sick or damaged, but might get better. Some are sick or damaged, and will never get better. Sometimes, there is no consensus as to which is which.

On a greater or lesser scale, we all make life-or-death decisions that affect other living entities. Some must die as food for others. Some die by accident, or by lacking food or shelter. Some die at the pleasure of others. An infinite number of (potential) entities are never given the opportunity to exist. Those that do exist, sooner or later cease to exist. An entity may produce progeny, artifacts, or other traces of its existence, and those things may have value for a time, but that value, too, will eventually approach zero.

The value of things can be used to set our priorities, both individually, and collectively. Our expectations and beliefs help determine our values. Therefore, the critical path in each life is the learning process that assembles a history of sensations into a set of skills, expectations, and beliefs. The first stage of learning is the ability to turn sensation into perception. The second stage is the ability to turn perceptions into narratives. The third stage is the development of judgement and wisdom—our collections of expectations and beliefs, and the actions we take based on them. Many complain about their memory, but few question their capacity for wisdom and judgement, even though it would be well worth doing so.

Our collective knowledge is a record and consensus of our individual knowledge. Our individual knowledge is gained from two sources: direct perceptions, and indirect inferences. The first of these involves interactions between external photons, and electrons that are part of our body (that's technically where the "rubber meets the road"). The second of these begins as the first, but is given an interpretation by a process in our brain. The first source comes from a very small subset of reality. It occurs within a ***very*** limited range of temperature, size, and location. The second source involves the operation of our fallible brain. Both of these occur with a huge amount of parallelism. The amount of information we are able to acquire, store, and use is really quite remarkable.

The knowledge and beliefs I've acquired compel the speculations in this book. I believe that, at the smallest scale, every event would be perfectly deterministic if it were not for the inherent uncertainty in collecting data. However, at larger scales, every event is the result of more and more precursor events, and these events compound rapidly, such that events we can observe directly are the result of extremely large numbers of precursor events. This means that some chains of events are profoundly uncertain. Other classes of events, however, are subject to statistical patterns and the law of large numbers. This reduces, but does not eliminate, uncertainty about a few things. I estimate the probability of the sun rising tomorrow morning to be ***very*** high. The probability of 30,000+ Americans dying this year due to driving accidents, and another 30,000+ dying due to gunshots, is ***quite*** high. The number of sub-atomic events that contribute to either of these outcomes is incalculable. Finally, there are events that are determined through the use of information—events caused by living things. The sciences of predicting these events (psychology and sociology) are at an early stage—still not very effective for prediction.

Evolution

Let's review ***how living entities came into existence***. Our region of the universe was once abundant with protons and electrons. From these, the most elemental of particles, hydrogen ***atoms*** were formed due to the electrostatic forces that cause a single proton and electron to join with each other. Subsequently, hydrogen ***molecules*** formed when two hydrogen atoms were able to combine. Fast forward the passage of time, and weaker gravitational forces combine huge numbers of hydrogen molecules. A star is formed, and nuclear fusion begins. In time, fusion produces heavier atoms, and more complex fusion processes arise. These cause more complex types of atoms to form containing larger numbers of protons, electrons, and neutrons. Eventually, if enough raw material falls into a sun, and fusion processes go on long enough, a supernova may result. This event spawns most of the known elements, and distributes matter back out into space where it may now coalesce into planets and suns that have the diversity and complexity of our own Earth and Sun (2^{nd} generation stars).

We think that it took about 9 billion years for these processes to evolve our solar system. It took another 4 billion years for people to evolve. The steps along this very long path were of many different natures, but all were perfectly ***natural*** within a framework of chance (anything possible ***may*** occur) and necessity (the inevitable ***always*** occurs).

The ontogeny of life involves a copying process. The smallest living components are assembled first; larger components are assembled from the assembly of the smaller ones. A copying process requires information to drive it. This is the use of information at its most elemental level. New ways to synthesize information are "invented" over the course of evolution. Two species of entities could be compared in terms of the ways they both used information, and one could be judged more complex than the other on this basis. Generally, the more refined and complex ways to use information are layered on top of earlier ones, and this is perhaps the only basis for saying that one species is more "evolved" than another—it has more layers, or it uses information in more different, or more complex ways.

Every entity that exists (including a packet of information) is created by copying some combination of previous entities with some degree (zero or more, but usually not a great deal) of modification. In the simplest case, a copy is created by chance following natural laws. For a more complex ***entity***, a copy may result from a more complex ***process***. The more complex an entity (the more steps it takes to generate a copy), the greater the number of evolutionary steps that must be taken for it to evolve. This implies, in general, that more complex entities require more ***time*** to evolve.

Up to a certain point, all entities that evolve are ***natural***. After a certain point, some entities are brought into existence ***by*** a natural entity; these we call ***artificial***. All natural entities that are capable of a response to a stimulus are generally called ***alive***. The assumption being that a response is due to the stimulus of some pattern of ***information***. However, this doesn't extend to artifacts (artificial entities). Some artifacts (machines) do respond to stimuli, but we don't define them as being alive. This is because they ***process*** information on behalf of live users (those producing copies of them in the first place).

Consider this distinction: When an ***artifact*** incorporates ***intelligence***, it crosses the threshold to being alive, but a ***natural entity*** is alive simply by making ***use*** of information for itself. A natural entity must first ***be*** an entity. This means that it has evolved some way to differentiate itself from its environment, and some way to get itself replicated.

The accepted scientific view is that prior to the first living entity, sets of autocatalytic chemical reactions evolved that were not confined within a cell, but perhaps within basins or tide pools of lakes or oceans. Some extremely complex molecules were evolved during this stage of evolution. There are also simpler examples of chemical reactions that duplicate complex structures (entities of a type) from simpler molecules (food). The criteria that an ***entity*** must exist ***and*** must contain and use ***information*** to replicate itself are ***both*** important to distinguish life from non-life.

Artificial entities are extensions of natural entities, and until they are clearly using information for themselves, they cannot be considered alive in their own right. However, if a collection of artificial entities were able to reproduce themselves from raw materials, it should be conceded that they would be doing so (and using information) on their own behalf. Thus, there are really two tests for artificial life: either the ability to reproduce ***or*** a threshold of intelligence. I expect the intelligence test to be the first one passed. The nature of intelligence is the subject of Part 2 of this book.

Two hundred years ago, our universe was thought to have been almost 100% explained and understood by the scientists of the time. In the last hundred plus years, new facts have come to light, and our understanding proved not so complete after all. For the past hundred years, our model of the universe has become more and more complex. Agreement on some aspects of the "standard model" has dropped far away from 100%. Although the ***theory*** of evolution is rock solid, the fact that its details cannot be known (it hasn't been replicated in the laboratory) causes some people to question its verity.

In the past hundred years, the theory of evolution stands above all other scientific models with respect to how hotly it has been debated by the general public. This model is one of the most hyped, least understood, and perhaps the most important scientific models of our time. We give Darwin more of the credit for the theory of evolution than any other single person, but that is now a small fraction of the total. Darwin didn't express his model as follows, but hopefully, he would have approved.

Evolution occurs when information is copied. This is the central theme of the model. Information is embodied in the DNA, RNA, and protein structures that make up all living things as we know them. These structures are copied from one generation of entities (a population) to the next. However, any pattern may be copied, and therefore any ***pattern***, not just DNA, may ***evolve***. An evolving pattern need not be alive. In fact, the earliest patterns to evolve most certainly weren't. Memes that go viral follow this model very well, but they are obviously not alive.

The following rules formulate the model of evolution.

- Every entity that evolves is created in a reproductive, or copying, process.
- Every such entity incorporates attributes from its ancestor entities.
- Every such entity may be copied zero or more times.
- No copying process is perfect. Changes are introduced at each generation.
- The changes introduced affect the number of copies made from each new entity.
- Over generations, this transforms the phenotype of an interbreeding population.

Most scientists would agree that once the proton, neutron, and electron became the dominant material in a locality, the formation of suns, supernovas, and subsequent star formation can be explained without much controversy. The entities formed by these processes (many copies of the elemental particles), all of the atoms from hydrogen through uranium and beyond, simple molecules and chemical reactions, and so on are all fairly well explained by accepted science.

Up to this evolutionary point, new entities are identical copies of previous entities, or are combinations of simpler entities. The processes involved are atomic and chemical reactions. These processes involve the chance that they will occur, and the necessity for processes to follow natural laws. Over time and space a huge number of different environments are inevitable. Stars and planets of many different sizes and descriptions will form naturally. There are only a handful of elemental particles. There is a somewhat larger handful of stable atoms. These can be formed into quite a few, but still a knowable number of simple molecules. However, the next step is the formation of more and more complex molecules capable of building the components required for life. This is the realm of unlimited possibilities.

Simple molecules evolve purely by chance. The patterns they are based on are given by the laws of nature. They "fall into place when conditions are right." Now, this is true of all things, but at some point of complexity, a molecule is only likely to exist after certain precursor molecules are present. These may require different precursors. Each step can only be taken when certain other steps have already occurred. We have a pretty good idea of the complexity of the molecules that support life on our planet, and we have a pretty good idea of what the starting conditions were and how long ago this process started. What we don't have is a detailed description of each step along the way. We know this was a process involving chance and necessity, and we know that it occurred. Some of the details are known, but many more will always remain a mystery. Most scientists, however, believe that this process was short of miraculous, and that each "rise" along the path was due to a push (that we can explain) from "below" not a lift (that we can't explain) from "above."

Once sufficiently complex "organic" molecules were present, sets of molecules began to work together to catalyze the production of one another. This "morphed" into copying based on templates, not just chance alone. As these systems became fully self-contained, living entities emerged. At the most basic level, these entities use the information contained in some of their more complex molecules to drive their reproduction. Today, these information bearing molecules are what we know as DNA. In a sense, it is "unfortunate" that the process of evolution erases evidence of its earlier steps. Think of the extinction of earlier species. All of the initial "near life" entities were driven to extinction by later entities that evolved. This is a natural consequence of the way evolution works, but it makes it very hard to know exactly what happened to bring us to the present. We may not know ***exactly*** how we got here; and we certainly don't know where we are ***going***; but the fact that we ***are here*** tells us a great deal.

Conditions on Earth today are very different than those that were present when the initial phase of evolution took place. No living thing on Earth today could survive if placed into that early environment, and no organism from that time could survive today. The biggest change to our planet was brought about with the advent of photosynthesis. Over a very long time, this changed the atmosphere consisting mostly of methane and ammonia into one consisting mostly of nitrogen and oxygen. These two atmospheres are mutually poisonous. Creatures that could flourish in one, would die in the other.

Evolutionary Stages

The first stages of evolution are [1]energy into matter, and [2]matter into more complex matter. Once a suitable planet is in a stable orbit a favorable distance from a suitable star, [3]planetary chemistry is the next stage of evolution. This process can take a billion years or so until complex organic reactions build up a stockpile of organic chemicals. Note: The path of evolution described here is the only one that we know for a fact actually occurred. It is possible that there are other paths. The original organic stockpile is often called the primordial soup. The organic molecules in the primordial soup could be compared to the apps on a modern smart phone. Both are precursors to an evolutionary event, the emergence of a new thing.

The next stage of evolution is the [4]development of an unconfined set of autocatalytic pathways capable of self-generation from simpler raw materials. This stage begins with concentrations of simple organic molecules. What makes a molecule organic? The answer is that within very narrow boundaries of temperature, an organic molecule can associate with, or disassociate from, a large variety of other organic molecules. This mean a very large number of molecular varieties can form. The size of an organic molecule is virtually unlimited, so an unlimited number of possible configurations exists. New molecular forms, during this stage of evolution, come into existence and disappear over and over again. However, when a set of molecules, each a precursor for the next, and the last a precursor for the first, happens to occur, it doesn't disappear, it self-generates—it is autocatalytic. Thus, different autocatalytic processes accumulate. These are the precursors to life. Each stage is layered on the previous stage. Each stage evolves processes that either eradicate, or have advantages over a previous process, allowing the instances of the new process to increase (until a new stage evolves).

The 5th stage is the incorporation of a set of autocatalytic processes into a "cell" capable of using material from outside the cell, and grow it until it splits into two pieces. Now, cells are capable of propagating themselves.

The 6th stage of evolution is the refinement and improvement of the cell to the point that photosynthesis occurs and the planetary ecology changes from one based on methane and ammonia to one based on nitrogen and oxygen. The advent of photosynthesis was likely because the sun is a better power source than heat gradients present in the environment. The entity that discovers this is able to win the power arms race. A side effect is that oxygen is generated, and all living entities now have to adapt to a new atmosphere.

The evolution of the cell takes another billion years or so, until a 7th stage begins. Instead of dividing into two identical individuals, cells cling together after division, and develop into multi-cellular organisms. During this stage, evolutionary principles shift into a higher gear. The growth cycle of an organism is driven by "code" contained in its DNA, and each new cycle begins with an imperfect copy of the previous cycle. Any addition or change to this information (the genotype) produces a change to the resulting new organism (phenotype). Natural selection ensures that the more robust variants become more numerous and those producing fewer offspring die out. As each change is added to the DNA record, the evolutionary history of changes is repeated in each cycle of growth. Each change is typically copied to an increasing number of descendants in subsequent generations, or it disappears within one or two generations. The bottom line is whether any physical changes are followed by more or fewer offspring of the entities that possess them.

Every biology course teaches the phrase: "Ontogeny recapitulates phylogeny." Let's see what this means. Ontogeny is the growth of a single organism; phylogeny is the evolutionary changes that took place from the first generation to the present. The part about "recapitulates" is a bit fuzzier. It involves driving the growth process using a record of the genetic changes that have occurred from the first multi-celled ancestor to the current fertilized egg. The record (***information***) guiding cell division must somehow record the history of evolution that culminates with a given cell about to divide. Each division is guided in accordance with the genetic record contained in each cell's DNA.

The reason ontogeny recapitulates phylogeny is that each evolutionary change is with respect to a previous fully developed organism. Changes occur one after another. A "design" isn't scrapped and re-done; it changes one small step at a time. Thus, the growth of an organism, from 1st fertilized cell to finished adult, actually passes through many of the major body types of its ancestors.

I don't believe this process stops at birth. Our brain seems to have evolved to assist some of our cultural attributes such as language and belief structures. I believe that the learning process responsible for bringing a brain "online" also follows the same evolutionary steps that led to its current capabilities. The phenotype is built from the DNA record, the development of the brain follows from interactions with the cultural environment (starting with parents and family). Thus, learning ought to recapitulate cultural history.

Once a child learns to understand and formulate complete sentences, the ability to remember and produce narratives is a skill that develops over a lifetime. Our ability to think, plan, and solve problems rests on our ability to produce narratives. Early narratives involve fantasy. Our beliefs are the narratives that we have accepted as true. These are taught to us. Some may later be rejected if better ones come along to replace them. Some narratives are accepted because they are consistent with experience, observation, and what our teachers tell us. Others may be accepted because of frequent repetition in the absence of viable competition, even though they may still be the fantasies of a child (or the delusions of an adult).

Evolution is a long and iterative process. In the short term, every dog has a dog for a mother. A chicken comes after an egg. Was there ever a ***first*** dog? It seems as if there must have been, but, if so, it could not have had a dog for a mother. Which is more fundamental, genotype or phenotype? The genotype is the ***information*** necessary to build another copy of the phenotype. The phenotype is the ***mechanism*** necessary to build another copy of the genotype. But, the phenotype is the living entity. It produces the genotype that begins the next generation. This is another example of dualism!

Classification of Life

Scientists classifying life on Earth have used several schemas in the past 100 years, or so. Early schemas were based on gross similarities between individuals: single cells, plants, and animals. In the mid-1900's classification moved toward describing different life forms with scores on a list of attributes. A factor analysis was performed to show natural groupings. When genetics became more of a science, it was possible to measure genetic similarity, and associate different groups that way. Naming conventions have grown very complex, placing each entity at the "tip" of a tree with eight levels of branches: The root of this tree is called "life." The first branch point is ***domain***, the next branch points are: kingdom, phylum, class, order, family, genus, and species.

A branching tree suggests that the root is the first single celled ancestor of all current life, and each branch marks a point where an ancestor cell gave rise to two different evolutionary paths. But, this is not so. The problem is that only the ***tips*** of the branches are visible today. These correspond to all ***living*** species. All earlier species have evolved into newer forms or have gone extinct. Records of extinct species within the last hundred million years, or so, are sparse; records before that time do not exist. In particular, the billion years or so from the first viable cell to the first multi-celled organism are not on record.

Thus, the tree of life shows the ***kinship*** of each organism to all others present on Earth today. The last branch just before a tip marks two closely related species, the branch before that marks two closely related genera. And so on, with broader and broader categories, enveloping related families first, then orders, classes, phyla, and finally kingdoms and domains. It is the case, however, that any two species living today have a common ancestor some number of generations ago. A common ancestor tree could be constructed, and partial trees, based on DNA comparisons, have been suggested.

Here, I wish to use a simpler classification. I'll classify all life on Earth as a tree with four limbs. Each limb has many branches and tips. Only the tips represent the life present today. The trunk and the first limb that branches out of it represent all ***single celled*** life. A second limb branches out of the first. It represents all multi-celled life ***without*** cell specialization. Then, a third limb branches out of the second. It represents all multi-celled life ***with*** cell specialization, but with ***none*** of its cells specialized as ***neurons***. Finally, a fourth limb branches from the third. It represents all life with ***some*** of its cells specialized as ***neurons***. Each of these limbs contains a sub-class of life, viruses. These tiny packages of information need living cells to reproduce. In fact, this sub-class should include mitochondria and other "viruses" that the host actually needs to survive (unnecessary viruses may either be benign or actually harmful to their hosts).

When [7]multi-celled entities came onto the scene, a new kind of life cycle was required. The life cycle of a single celled entity is to grow bigger and divide into two copies of itself. Each copy repeats the life cycle. Or it fails to do so by failing in the growth process, getting "eaten," or failing to divide into two viable copies. The life cycle of a multi-celled entity also begins with a single cell, but at its first cell division, and for several divisions after that, the "daughter" cells remain attached to each other. Now, there are many ways that ontogeny (growth due to successive cell divisions) can go at each step of cell division. Cell specialization can take place. Different morphologies can evolve.

The first developments at this stage of evolution were different sizes and topological shapes of tissue masses (aggregations of cells). The survival value of these changes may have been that larger size made the entity harder to eat, or better able to move. After several cell divisions, the entire entity might split into two and repeat the cycle. Or it could develop a "bud" that would detach and start a new cycle. Simple, multi-celled organisms have flourished ever since.

The next evolutionary steps would be [8]cell specialization and more complex organization. Organisms with a few thousand or a few million cells would benefit from having sensors at some distance from effectors. Along with cell specialization, [9]more effective ways to generate the initial cell of the life cycle from two "parent" organisms evolved. Sexual reproduction allowed one parent to contribute (almost) half of the DNA to the other parent, who then contributes the rest and produces the initial cell. This entailed two copies of the DNA (allowing "error correcting" redundancy), and reshuffling genes at each generation. A major effect of this is that ***mate*** selection is now a part of ***natural*** selection, affecting the number of offspring produced, and therefore the survival of the species. I won't try to describe the steps that were taken to evolve [10]the first primitive neuron, but it did indeed evolve, eventually to produce the human brain.

DNA contains at least three types of information: How to build a variety of proteins, how to change the form of the next generation of daughter cells (or whether not to produce another generation), and how to signal or react, given a signal from an adjacent cell. DNA must drive the critical part of the life cycle. It must define the generations of cell division from an initial fertilized egg to a mature organism capable of producing a fertilized egg. In addition, it defines other properties that enhance the chances of viability for its own fertilized eggs, and those of genetically similar organisms.

Any change to the information contained in the DNA of a fertilized egg, with respect to the information contained in the parent's DNA, most probably has a negative impact on the viability of the next and future generations of that organism. However, some changes accidentally produce an improvement in one of the three areas: A better variant of a protein, a change in a structure (through replication, size, or constitution), or some change to the signaling process that affects cell specialization or the number of cell divisions before an action is triggered at the next cell division. Favorable changes, though rare, tend to endure due to the enhanced viability of its possessor (leading to more copies). Unfavorable changes reduce viability, don't get copied, and tend to disappear.

But, back to the neuron: What we now call a neuron evolved from the advent of a cell specialized to transmit a signal from the location of a sensor to the location of an effector. This defined the fourth "limb" of my evolutionary tree. Of course, cell specialization continues on all of the four branches, and sexual reproduction continues on the 3rd and 4th. However, the most significant development of all, ***intelligence***, continued to employ neurons in ever more numerous ways and numbers. As stated earlier, the advent of sexual reproduction involves mate selection, and that may well select for intelligence. The introduction of intelligence adds a feedback loop that causes the evolution of intelligence to be exponential to the degree that intelligence drives mate selection drives numbers of surviving offspring.

Evolution of Artifacts

Anything that can be copied is subject to evolution. Most higher life forms (on earth, mammals at present) have evolved the ability to mimic the behavior of other animals to a certain extent. This is especially true for humans. An immense number of human behaviors have evolved. Species of behaviors include spoken languages, song and dance, storytelling, reading and writing, crafting tools, objects of art, and the ***making of artifacts*** of many descriptions.

Calculators and computers are one important class of artifacts. Calculators were described to a certain extent in an earlier section. Since I witnessed most of the evolution of computers, here I would like to recount some of it to the best of my knowledge (without too much assistance from the Internet). I'll probably reflect my own history, and get some of the details wrong, but here goes anyway.

The concept of computing was introduced by Charles Babbage (with the extension of programming by Ada Lovelace) in the 19th century. These visions were introduced into reality just before I was born. Twenty years later, a great deal of the evolution of computers had taken place, and I learned about it in graduate school. Now, another 50+ years of evolution has occurred. It took a bit of fumbling, at first, to settle on a design that had a separate processor, a memory, and an instruction code that could be stored in memory, loaded into the processor, and executed so as to modify memory, or move values in memory from one place to another (including in and out of the processor, or in and out of external devices).

Initially, a lot of use was made of BCD (binary coded decimal). Machine storage units were a multiple of 4 bits to encode one digit in each location. Storage units were handled as wholes. As time went by, storage utilized 4, 6, and 8-bit divisions. The 6-bit divisions have largely died out, but they were still popular when I was in school. These were grouped into 30- and 60-bit registers. Now, by far the most popular way of storing numbers in computers is in the form of 8-bit bytes. When larger values are required, several bytes in a sequence of memory are used. The popular lengths that have been used are 2, 4, 8, 16, 32, 48, 64, and 128 bytes (maybe even 256). Of course, there are two ways that a sequence of bytes can be stored in memory. One is with the least significant byte in the lowest memory address, and the other is with it in the highest address. And, both were employed. The former has largely prevailed. It is called "little endian" since the little end of the number comes first in memory. However, since we write numbers with the big end to the left, and we associate the left with the lowest address in memory, there was a reason that big endian machines were also designed. Danny Cohen introduced these two terms for byte ordering in a 1980 article that he wrote. In this technical and political examination of byte ordering issues, the "endian" names were drawn from Jonathan Swift's 1726 satire, Gulliver's Travels, in which civil war erupts over whether the big end or the little end of a soft-boiled egg is the proper end to crack open.

The standard theory of evolution doesn't mention any model for evolutionary forks, but when more than one possibility exists for a design, more than one is likely to occur. When these are forced into direct competition, one design typically prevails due to initially having the greater numbers, or because it survives direct competition. Other choices of computer design that have largely settled out by now are how negative and floating-point numbers are formatted, and much of the logic of a computer's basic instruction set.

Emergence

Life and intelligence are two subjects that involve organized complexity on a huge scale. Consider what is going on inside the simplest of single cells. The genetic coding alone involves hundreds of thousands of bits of information. Hundreds of chemical processes go on under control of the genetic code. A cell wall exists. Simple structures also exist within the cell. These structures facilitate the cell's life cycle, which is to accumulate material (eat), dispose of waste, and bring about the conditions necessary for cell division, at which point the life cycle repeats itself.

Are these the only two subjects that involve organized complexity on a huge scale? Today's technology springs to mind. Perhaps, today's literature, or libraries, or the Internet could qualify as members of this set. This suggests a sequence of "begets." Nature begets life begets intelligence begets technology (begets artificial life? which then begets what?).

To break through into the above parentheses, it is first necessary for intelligence to understand intelligence. How far will natural life walk hand-in-hand with the artificial life it creates before natural life is left behind? Why would this created artificial life not regard itself as perfectly natural, especially in the distant future when it has parted ways with its creator?

When certain simple things operate together, reinforce one another, combine to form a whole that is more than the sum of its parts, and the number of simple things becomes very large, entirely new properties may emerge.

Erwin Schrödinger in *What is Life?* said that life is defined by the ability of an entity to take in bits of low entropy stuff in order to maintain or reduce an already low state of entropy. In other words, the ability to successfully fight the Second Law of Thermodynamics (for a time). Above, I made the claim that life is the collection of all entities that use information. This is nothing more than Schrödinger's definition taken a few (intuitive) steps further.

The point is, when life emerges, the concept of information (or negative entropy) is introduced. By its very nature, the use of information is context dependent. Science may correctly believe that the laws of physics should be fundamental and absolute. Likewise, so should be the laws of chemistry which emerge from them. And again, the laws of organic chemistry and molecular biology. However, upon the laws of molecular biology and evolution (which are probably universal), rests the biosphere of the Earth, which is unique. It is a product of chance operating for a very long period of time according to the necessities of natural laws. Our biosphere evolved a "sapiosphere." The collection of *homo sapiens* that makes up this sapiosphere has evolved thousands of languages and cultures, each with its own definitions of right and wrong, good and bad—and each of which judges, rewards, and punishes its individual "sapiens" differently. Here, the truth is in the context.

The study of a cell, an organism, the psychology of a human being, a culture, an economy, or even the human brain, is the study of an evolved system, a unique entity, the product of chance and necessity, an accident of history. These are complex systems. Certain truths, only recently formulated, are required to understand a complex system. A rather barren discipline called General Systems Theory was popular for a time in the 1960s. The science of Complexity Theory appears to be a somewhat more promising replacement for it. A complex system is a system with its own emergent properties. It generally has on the order of thousands to trillions of individual parts (although it might take as few as two complex systems to form a third with its own emergent properties). Each part of a complex system interacts with one or more of the other parts. Interaction may be in the form of direct action, the delivery of material, or the exchange of information (sometimes it's hard to differentiate which). There may be both positive and negative feedback loops in the system. Many of the parts may be identical, or almost so. Every complex system is always, to some degree, both structured and robust. Most are also vulnerable to circumstances that can make them go chaotic. The science of complexity also studies systems that are not based on life, but still have a relatively high degree of complexity. Examples are the weather, the solar system, geological processes, and condensed-matter physics.

A model is never completely accurate—aspects of a model may be true, but fictions may also be derived from a model. The power of a model is the extent to which it can make predictions. Only models of very simple phenomena can predict very far into the future. Chaos sets in, sooner or later. When chaos or a complex adaptive system is involved, the horizon of prediction is likely to be fairly close by. However, when the chaotic component is small, the horizon of prediction may be quite far away. The less its predictive ability, the weaker the science—its remaining value being largely descriptive.

You emerged from the uncaring immensity of the universe into the hole you now occupy with the rest of humanity. You say you are not in a hole? Then try leaving. See how far you get. The two pieces of advice mentioned earlier regarding holes also apply to ***conceptual*** holes: To the extent you find yourself in a hole, further digging doesn't help; and when you find yourself on the brink of an abyss (near a big hole), you need to take your next step very carefully. Humankind is nature's way of tinkering on the brink of an abyss. Nature doesn't care about careful. As a species we need to consider our position very seriously! As the saying goes: "Today we're on the brink of an abyss; tomorrow we'll be one step further!"

The next step in this book is to explore the nature of our own intelligence. How did it arise? How is it comprised? Could we somehow duplicate it? Is it even possible for us to understand our own ability to understand? We should at least attempt to take the first few faltering steps. So! Brazenly forth we go.

Part 2 — Intelligence

Intelligence—do we even know what it is? Of course we do, "it's what intelligence tests measure." The fact that they do measure something is demonstrated by the high correlations exhibited between the answers to different questions within the same test, and the results between tests given to the same individual over time. I hold to the "Spearman's g" theory that intelligence has a general underlying factor, and exists as a collection of a number of components—it is a sum that is more significant than any of its individual parts. We may test the components separately, but I would suggest that a composite score is more important than any individual factor.

People know intelligence when they encounter it—most people seek intelligence in a spouse. This may explain how intelligence evolved in the first place. A positive feedback cycle for intelligence might involve mate selection by individuals striving to get the most intelligent mates they could. The less intelligent, being less sought after as mates, would generate fewer offspring. The survival benefits of intelligence might even be secondary to the effects of this positive feedback loop.

Intelligence tests present stimuli in the form of questions, and each response is graded as right or wrong and how quickly it is given. The final score is the sum of these individual grades turned into a position on a normal curve obtained from the results of many other individual's scores. An IQ score of 100 is at the 50% position on this curve. Half of the population has an IQ less than 100, and half has an IQ higher than 100. Each 16 points higher or lower is a standard deviation. About 84% of all subjects fall below the first standard deviation above the mean. About 98% of subjects fall below the second standard deviation. Only about 0.1% of test subjects score above the third standard deviation. The curve on which each IQ result is placed is a bell shaped curve, or normal curve. The IQ score is plotted along the X-axis and the number of people scoring that IQ is plotted on the Y-axis.

Another way to define intelligence, the way I would like to define it, is: ***Intelligence is the stimulus-response repertoire of an entity***. According to this definition, the first cell able to notice the presence of a desirable molecule (and hence, the likelihood of more of them) lying in a certain direction, and then actually moving in that direction, deserves the credit for having the first iota of intelligence. An entity's stimulus-response repertoire includes the pattern recognition of the stimulus, and the appropriateness of the response. Again, the test score an entity receives on any such test is positioned on a bell curve with scores obtained from a population of other entities to which the test subject is being compared. Care must be taken to choose the stimuli and the grades for responses. If this is done properly, each stimulus-response pair will correlate positively with others on the same test, and with retests on the same subjects. The actual subject matter, or content, of the stimuli is less important than the degree of these two positive correlations. Also, the wider the range of difficulty, the more valid a score is at the extreme ends of the bell curve.

When I use the term stimulus-response, I don't limit that to reflex actions, I mean any set of stimuli, followed by a behavior. Similar sets of stimuli followed by similar sets of behavior fall into similar sets of stimulus-response. Significant cognitive action may lie between the stimulus and the response. Every situation an entity finds itself in amounts to a stimulus, and every behavior it exhibits amounts to a response. Continuous behavior may be successive responses stimulated by feedback from the most recent behavior.

A third way to define intelligence is to use a human as a yardstick. This was Turing's approach. In today's version of the Turing Test there are two computers on a table in a room. One computer has a sophisticated program in it, the other is just a simple link to a human in another room. Both respond similarly. Both can carry on a conversation with you. Can you tell which computer is "talking" for itself, and which is connected to a human? How much would you be willing to bet? When an experienced judge can tell the human from the computer only about half the time, we could say the computer program has passed the Turing Test—it can pass for human in a chat dialog—and for this reason, we say that it exhibits intelligence.

Clearly, the program should evince understanding, be able to answer questions, work out problems, and exhibit a range of real-world knowledge. It might also have to feign some human failings, such as less than perfect typing, spelling, and grammar, to be perfectly convincing.

The problem here is that the program might pass the test by exhibiting personality more than intelligence. Perhaps a better scenario would simply be a panel of judges who each conduct a one-hour conversation with six computers. Five of the computers would be connected through to humans with measured IQs spanning the range from 80 to 120. The sixth computer would contain the "intelligent" program. The judges would be used to rank the six "contestants" in order of their intelligence as well as to try to "spot the computer." If the computer consistently ranked above the human with the lowest IQ, and was not always spotted as the computer, it could be credited with a "pass" on this "Human Yardstick Test."

A good result on any of these tests would allow us to judge a computer as intelligent if we were to encounter one. But, what makes us think that today's hardware together with the right software could emulate the intelligence of a human brain? Today's "server farms" in terms of memory sizes and processor speeds are probably within an order of magnitude of the human brain. The question then becomes, "How ***complex*** can a given sized ***program*** be?"

To see how the complexity of software can increase with program size, consider the Busy Beaver function proposed by Tibor Rado circa 1962 (citations are hard to find; see *The Age of Intelligent Machines*, by Raymond Kurzweil). Essentially, a "busy beaver" is a computer program that prints out the longest possible sequence of "1" digits, and then stops. It doesn't take a very complex program to simply print forever. That can be done by a set of instructions that loops and never quits. It is much harder to get a program to produce a ***very*** long sequence of 1's, and then stop.

The busy beaver function is the mapping between the number of instructions in a program and the number of "1" digits it can print and then stop. This is a very simple analogy to show the limits of complexity. The function is interesting because of how fast it increases. It can also be shown that the values for this function cannot be computed for all numbers. In fact, it can further be shown that whether a given busy beaver program (or any program in general) will ever stop cannot be computed for all cases. The only thing we can do to compute the Busy Beaver function is to lay down a set of rules for a machine of some type, and the instructions used to program it, then begin solving the function one number at a time. This has been done for a theoretical computer called a Turing Machine (TM) with up to eight states (a state in a TM corresponds to a line of code in an ordinary computer). The Busy Beaver of a TM with fewer than 6 states isn't even mentioned by Kurzweil, and he reports that the Busy Beaver of 6 states is only 35. The Busy Beaver of seven jumps to 22,961 and the Busy Beaver of eight is estimated to be the number 1 followed by 43 zeros. The Busy Beaver of nine has not even been estimated as far as I know, and well before the Busy Beaver of 100, the intelligence of the human brain is probably incapable of even making an estimate (actually, I personally don't know how to estimate it for 6 or 7). The value of the function at each successive count is ***very*** difficult to determine!

The point of this is to convey the complexity ***possible*** even in short and simple computer programs. Couple this with a result proved by Turing and mentioned by Kurzweil: A Turing Machine can model ***any*** machine, where a machine is regarded as a defined process or mechanism that follows natural laws. Then, if we regard the brain as something that could, in principle, be described, and that follows natural laws, we must conclude that the brain could be modeled by a Turing Machine, or by a common computer with sufficient memory. Modern programming efforts routinely combine many thousands of lines of code into single programs. Duplicating a neuron should not require nearly so much. Duplicating a billion neurons is simply a matter of scaling up, and achieving a minimum processing speed. It is estimated by many that both of these are within reach of today's technology. Let's review what we know about a wet brain that might enable us to design a dry one.

We humans have developed five obvious senses, and a number of less obvious ones. Likewise, we have developed quite a few muscle systems to produce our behavior. The vision structures of cephalopods, insects, and vertebrates appear to have evolved independently. Any functionality that has evolved multiple times must be pretty fundamental. Hearing is another fundamental sense, in that bats and dolphins utilize it for echo location, and we utilize it to hear speech and music. But, what land animal doesn't use hearing to gain information about that sound in the "forest"? On the other hand, speech and manual dexterity are almost unique with humans. A few closely related species have limited manual dexterity, and many species have "speech" in the form of a few "cries" or "songs." The songs of whales do appear to consist of large vocabularies, but that conjecture is still being studied. Intelligence at and beyond the human level is based on what can be ***visualized*** and what can be ***verbalized***. These abilities are based on pattern formation, recognition, recall, and a repertoire of responses.

In this book, an artificial intelligence capable of sufficient learning and intelligence will be called an Artificial Life Form, or ALF. This distinguishes an ALF from current forms of AI. An ALF is more than a collection of computational techniques that duplicate a single aspect, or a few aspects of intelligence. Sufficient learning means that an ALF has gained most of its intelligence as a result of learning—sufficient intelligence means that an ALF is able to ***use information*** on its own behalf.

When trying to design something, an existing design that comes the closest to achieving the desired objectives is often copied with modifications. The first attempts to design a flying machine tried to copy the flight of birds. The reverse engineering of flight is recorded in old films of people trying to fly by strapping on wings and flapping really hard. Only recently was human powered flight achieved. The Gossamer Condor looked nothing like any bird. It was a high-tech marvel of materials and design, and it owed nothing to the way birds fly. MacCready, Human Powered Flight: Gossamer Condor, 1977, flew 1 mile, in a figure 8; Gossamer Albatross, 1979, crossed the English Channel.

It's too early to say if this analogy parallels the efforts of those who are attempting to reverse engineer the human brain, but it's tempting to think so. Others have already made good progress at getting machines to do some pretty amazing things, but the fact remains that machines don't think—at present! They calculate, search, store, retrieve, input, and output. Networks of machines transmit volumes of information around the globe on behalf of their human users. Machine encoded information includes not only numbers, but text, sound, pictures, and graphics. An incredible variety and volume of information can be processed by a machine, but will a machine ever think for itself?

This question has no meaning until we clarify what we mean by think. This word usually means the subjective mental activity of talking to oneself. We have never extended it to a non-human subject. Thinking is something that our mind does, that few, if any, of the minds of other animals are capable of doing. Thinking is a conscious endeavor. Thinking must be supported by the abilities to first ***sense***, then ***learn***, and later ***perceive***. These abilities are present in the simplest of brains, and yet not in any of today's computers. Thinking must involve some ability to calculate, and use language, but it must culminate in decisions and action. Thinking without inputs and outputs is neither effective nor evident. Thought without language cannot be communicated, and, as we define them, such thoughts would remain very primitive.

Let's take, as our minimum set of abilities, some form of input, output, and information processing capability. The human brain has an assortment of such abilities. All higher animals do. To some extent, so do computers. So far, we have only defined a necessary substrate on which the process of thinking could occur. What do we need to add to get a process that an intelligent observer would recognize as actual thinking?

Thinking must be whatever process occurs between the input and output—between sensory inputs and behavior. When you play computer chess, it's not uncommon to say "the computer is thinking" when you are waiting for it to make a move. However, it doesn't seem quite right to call "thinking" any sequence of calculation. Thinking may involve calculation, but calculation is not thinking. Thinking involves perception and judgment. Any narrow ability to perceive and decide is not evidence of a general ability to think. Perhaps a "general ability to think" is nothing more than an extensive repertoire of narrow abilities to perceive and decide. This brings us back to the ability to learn—the ability to build such a repertoire. All animals seem to have the ability to learn. In some important sense, the more that can be learned, the more advanced the mind.

As far as we know, however, none of the other species that share this Earth with us have more than a very rudimentary ability to calculate and use language. Animals may be canny, clever, and highly successful, and still lack the ability to think. Whales and dolphins may have an advanced ability for language that we still have not fathomed. Are they capable of thought? We can only speculate, but theirs would have to be a mode of thought very alien to us, indeed. Their inability, to fashion artifacts that capture and represent aspects of our thinking set them very far apart from us. Likewise, any future ability of machines to think may be poles apart from our own thought processes. The development and use of language requires the ability to formulate and use symbols.

Symbols

A symbol represents a pattern, and a pattern represents information. A ***language*** is a system for arranging symbols. We have seen that numbers are symbols that represent counts. Any count can be represented with a number symbol. The formation of other symbols is more arbitrary. Although, it must be said that the formation of numbers has arbitrary aspects, and some of the past methods to generate number symbols were more arbitrary than the ones we have currently settled on. Generating the symbols for words is much more arbitrary, and that's why we have so many languages, and why languages evolve so fast.

Our intelligence is founded on our ability to associate a pattern of sensations with a symbol and use these symbols to think and to communicate with others.

Symbols may be given a rigorous association with a very definite pattern, such as a number with a count, or they may be associated with a less rigorous pattern, such as "chair" with a variety of objects generally used for "sitting." There is no limit to the complexity of the patterns that might be associated with a symbol, nor with the number of different symbols. Since each natural language has its own set of symbols, and these tend to associate with slightly different patterns, translation from one natural language to another is never exact. Even translation between more rigorous languages (such as programming languages) may not always be exact.

Human beings have two key advantages over all the other animals on this planet. They have superior manual dexterity over a much larger range of tasks, and they have a vastly superior ability to use symbols (as information expressed by language). To get a better handle on language in general, it would be better to build up to it.

Metalanguage

Earlier, I introduced most of the rules necessary to explain how we use numbers to count, measure, and compute. This entailed a "language" we call "arithmetic." We use many languages, spoken, written, and a variety of others. Most of our very thoughts depend on our knowledge of languages. Rules form the basis of language, and there are languages whose only purpose is to ***define*** languages (state the rules that allow the language to be recognized or produced).

A grammar defines a language. A grammar also defines how ***any*** sequential behavior may be recognized or produced. The definition of a grammar is expressed in a ***metalanguage***. A metalanguage is a language ***designed*** to define a language. The metalanguage used to define English contains words such as noun, pronoun, verb, adverb, article, conjunction, and so forth. Here, in a kind of bootstrapping process, I shall define a metalanguage in itself. Since everything needs a name, I'll call this language ***Sepade*** (for ***se***quential ***pa***ttern ***de***finition).

Languages consist of words formed into sentences according to rules. The rules of a grammar have two components: syntax and semantics. Syntax governs word order. Semantics governs the meaning of a sentence and the words within it. Words are generally built up from an alphabet or a set of strokes. To keep things simple, I'm only going to consider the alphabet defined by the 7-bit ASCII (American Standard Code for Information Interchange) character set.

ASCII Character Set

```
00 ^@    08 BS    10 ^P    18 ^X    20 SP    28 (     30 0     38 8
01 ^A    09 HT    11 ^Q    19 ^Y    21 !     29 )     31 1     39 9
02 ^B    0A LF    12 ^R    1A ^Z    22 "     2A *     32 2     3A :
03 ^C    0B ^K    13 ^S    1B ESC   23 #     2B +     33 3     3B ;
04 ^D    0C ^L    14 ^T    1C ^\    24 $     2C ,     34 4     3C <
05 ^E    0D CR    15 ^U    1D ^]    25 %     2D -     35 5     3D =
06 ^F    0E ^N    16 ^V    1E ^^    26 &     2E .     36 6     3E >
07 ^G    0F ^O    17 ^W    1F ^     27 '     2F /     37 7     3F ?
40 @     41-5A = (A-Z)     5B [     60 `     61-7A = (a-z)     7B {
5C \     5D ]     5E ^     5F _     7C |     7D }     7E ~     7F DEL
      ^ means hold CTRL down while pressing the indicated key.
```

The above table shows how various keystrokes on a standard computer keyboard generate the 128 codes of the ASCII character set. The hexadecimal number of the code is given to the left of the corresponding character. The table is arranged in quadrants. The 8 lines to the upper left define 32 control codes (which I won't discuss further). The 8 lines to the upper right define the space character and 31 standard symbols. The 2 lines to the lower left define the 26 upper case alphabetic characters and a further 6 symbols. Finally, the 2 lines to the lower right define the lower case alphabet, 5 more symbols, and a final control code called **DEL**ete.

This encoding dates back quite a few years to the days when the various control codes made sense. Since there has never been a standard where 7 bits occupied an address in a computer, or a packet in a transmission, the 8th bit of this code was made available as a parity bit. That is, the most significant bit could be set to 0 if the other 7 bits contained an odd number of 1 bits, or 1 otherwise. Most modern computers make parity invisible at the application level, so the codes from 128 to 255 have been allocated to other (non-standard) purposes. Again, they needn't bother us further.

The alphabetic, numeric, punctuation, and other symbols will be quite sufficient for our further explorations. Each of these symbols maps to a display font, or a position on a computer keyboard. Each has a standard meaning, but there are many standard fonts, and several "standard" keyboard arrangements.

Sepade

In Sepade there are three types of rules. They have the general form:

label = terminal

label = pattern

sentence

where, ***label*** is a string of lowercase alpha characters used to name the rule, the **=** character appears in the rule literally (the words terminal, pattern, and sentence are ***references*** to labels that have yet to be ***defined***), ***terminal*** will define a set of characters, ***pattern*** will define a non-terminal, and ***sentence*** will define one of the various sentence forms of the language (which will turn out to be the three forms shown above).

Before giving the full definition of Sepade, there are a few elemental constructs that will be helpful to know in advance. The construct [xyz] means that "xyz" is optional. The construct [xyz].. means that "xyz" may occur zero or more times. The construct x | y | z means that exactly one of "x" ***or*** "y" ***or*** "z" may occur. Lower case letters are used here to name something symbolically, and **= " [|] ..** are used symbolically as "meta" characters. When a character is to be referenced ***literally***, it may be enclosed in **"** characters. Uppercase alpha and non-metacharacters may be referenced literally without being enclosed in quotes. All of this will be defined precisely using Sepade itself.

Character Scanning

The lowest level of definition is the character set. The next lowest level is the definition of a scanner. Sequential pattern recognition involves a linear scan of a sequence of symbols drawn from the character set. Think of a character as an atom, and a word as a molecule. Or better yet, think of these as elemental chunks, and larger chunks made from smaller chunks. The job of a scanner is to define and recognize chunks one level above characters. To do this, we use the following rules:

```
label      = 61-7A [30-39 61-7A]..
lit        = #-& (-Z ^-{ ! \ } ~   i.e. notany " ' [ | ]
dq         = """ ' 20-7E ' """
hex        = 0-9 A-F ' 0-9 A-F
string     = """ ' 20-21 23-7E [20-21 23-7E].. ' """
```

Let's see what these rules mean in English before we actually define their right-hand parts using Sepade. The first rule defines a ***label*** as any character code (hex) 61-7A, followed by zero or more repetitions of any of the character codes (hex) 30-39 or 61-7A. Referring to the ASCII chart, this means a lowercase alpha followed by zero or more digits or lowercase alpha characters. A ***label*** is defined above as any string of characters that follows this rule.

The next rule defines ***lit*** as any of the characters `#` through `&` or `(` through `Z` or `^` through `{` or `!` or `\` or `}` or `~` (after some whitespace, there is a "comment" stating that this is equivalent to all the characters ***except*** the metacharacters `" ' [ | ]`. This rule defines all the characters that are taken literally when they appear in rule definitions. Labels and the remaining 5 metacharacters are interpreted symbolically.

The next rule, ***dq*** (double quote) uses an extension to the syntax employed in the previous 2 rules. It defines a quote followed by any code (hex) 20-7E followed by a quote. This permits three consecutive quote characters, and it means that the enclosed character only is allowed by the rule when ***dq*** is referenced. The apostrophe characters are used to separate the characters in one required character set from the next when a terminal is defined with 2 or more different sets of characters.

The next rule defines ***hex*** as any of the codes represented by the literal 0 through 9 or A through F followed by a second occurrence of the same.

The last rule above defines a ***string*** as a longer version of ***dq***. Here, the initial `"` is followed by one or more codes 20-7E excepting the code 22 (the code for `"` itself), and then a final `"` character. Notice that all but one form of ***dq*** also conforms to the definition of ***string***. This ambiguity requires the ***dq*** rule to be applied before (have precedence over) the ***string*** rule.

Terminal Definition

Now, we'll formalize the definitions described above using more rules in Sepade. Terminals in a metalanguage fall into two classes: literals and variables. Literals are defined with explicit characters included directly in definitions. Variables are character strings defined by rules. In Sepade, variables are used to define hex values, and arbitrary character strings and labels.

```
pair       = hex - hex | lit – lit
byte       = lit | dq | hex
list       = pair [pair].. [byte].. | byte [byte]..
lists      = list ["'" list]..
eb         = "].." | "]"
group      = lists | "[" lists eb
terminal   = group [group]..
```

These rules, together with the set above, define their own syntax. Here, we'll simply notice how these patterns apply to themselves. Notice that ***terminal*** is not referenced from any of these rules. This is suspicious. It is defined for the purpose of being referenced elsewhere. It actually defines all of the allowable forms on the right side of a ***terminal*** definition rule.

A ***terminal*** is one or more ***group***s. A ***group*** is one ***lists*** or a ***lists*** enclosed in brackets. The "end bracket" may be followed by a `..` (optionally). A ***lists*** pattern is one or more ***list*** patterns separated by literal `'` characters.

A ***list*** is zero or more ***pair*** patterns followed by zero or more ***byte*** patterns. Notice that if a ***pair*** is not present, the 2nd alternative requires at least one ***byte*** to be present. Therefore, at least one ***pair*** or one ***byte*** must be present.

A ***byte*** is a ***lit*** or a ***dq*** or a ***hex***. And, a ***pair*** is ***hex-hex*** or ***lit-lit***. (where - must appear literally).

Now, there are 2 subtle questions that must be answered. First, how and where is whitespace allowed (defined) in the above? Second, the definition of ***eb*** looks ambiguous, wouldn't the 2nd alternative always match?

Whitespace is defined as the characters in a string of text that appear blank in the display or printout. They are generated by the SP (ASCII space) character or the HT (ASCII horizontal tab) character. Whitespace could be incorporated into a definition using Sepade, but none has been defined here, except within the ***string*** and ***dq*** definitions. Instead, whitespace is handled by the convention that any whitespace of 2 characters or more is ignored at the possible end of a line (ignoring any "comment" as well), and ignoring all other single column whitespace except when it interrupts a ***scanning*** definition.

Ambiguity among alternatives is resolved by giving each alternative a number and attempting to match the longest first. All repeats or options ultimately succeed. However, each repetition is attempted until one fails.

Pattern Definition

```
literal    = string | byte
suc        = label | literal
successor  = suc | "[" suc [suc].. eb
alt        = [successor]..
pattern    = [alt "|"].. successor
```

Here, we see that a ***pattern*** is ***alt*** **|** repeated zero or more times followed by a ***successor***. An ***alt*** is zero or more ***successor*** patterns. A ***successor*** is a single ***suc*** pattern, or one or more ***suc*** patterns enclosed in brackets. A ***suc*** is either a ***label*** or a ***literal***. Finally, a ***literal*** is either a ***string*** or a ***byte***. Notice that ***string***, ***byte***, and ***eb*** were defined earlier. The definitions of ***pattern*** and ***terminal*** (above) are now referenced below by the 3 rules for ***statements*** in Sepade.

Rule Definition

```
label = terminal
label = pattern
pattern
```

Notice, in the above definitions, there are 2 indent points. The 1st is to a ***label*** symbol, the 2nd is either to an = character, or to the "body" of a terminal, pattern, or rule definition. A rule definition has no label, because it is not referenced from other definitions. A rule definition is a "sentence" in this language. Sepade defines itself using 3 rules to define its statements, 5 patterns to define patterns, 7 patterns to define terminals, and another 5 terminal definitions. In other words, a total of 20 "lines of code." Sepade can now be used to define the syntax of any artificial language. It will serve us here as an example of the sequential recognition and production of symbols in general.

Semantics

The semantics of a language is a definition of the "meaning" attached to its syntax. Given a context, the term "meaning" can be clarified. Languages are read to assemble meaning into short and long term memory. Languages may be translated into other languages. Artificial (including many different programming) languages are translated from source to alternate source forms, or to machine language.

Thus, a syntactic form encapsulates information, and communicates it for a specific purpose. Let's follow the process as the above metalanguage is used to "read" (or translate) its own first rule:

```
label     = 61-7A [30-39 61-7A]..
```

The first character encountered is the **`l`** that begins **`label`**. The symbol "label" is matched by the pattern defined as **`label`**. The **`l`** and subsequent characters are matched by the character sequences allowed by the right part of the **`label`** definition. Next, the **=** character, and all of the rest of this statement is matched by following references beginning with the

{**`label = terminal`**} rule.

In particular, the 1st character to the right of the **=** character is matched by the following sequence of rules:

`terminal→group→lists→list→pair→byte→hex→"6"`.

Each rule of the grammar to be invoked (and usually the order in which the rules are invoked) is recorded by listing its symbol. The "meaning" of a sentence is equivalent to the sequence of symbols used to recognize or produce the sentence. Here, the list of symbols is given in a top-down sequence. Once this list of symbols is complete, its right-to-left sequence is bottom-up. The process of producing this list of symbols is usually called parsing. I call it pattern matching. I'll have more to say about these concepts later. If some of them don't make sense, Google what seems interesting. The purpose of explaining metalanguages is that they are just as fundamental to describing languages as arithmetic is to describing and measuring the physical world.

Statements in a language may be subjected to three different kinds of processes. They may be translated, they may be interpreted, or they may be executed. The translation process involves using a syntax to drive the recognition of a statement to the production of a statement in a different language. The interpretation process is a software driven translation that results in actions rather than statements. This differs from the execution process in which hardware performs the recognition of the syntax and the production of the responses or actions.

And now, let's again step away from a simple ***rule based system***, and enter the area of ***natural language***.

Language

A language is a way to encode, transmit, and record information. There are many languages, and many different ***kinds*** of language. Each language serves a limited purpose for its users. A language has a syntax that defines how it is recognized and produced. A language also has a semantics that define how it is ***used***—what a given syntactic form ***means***. Not all syntactic forms have meaning, but all meaning is expressed as a syntactic form in some kind of language.

The first step toward intelligence, is a genetically determined stimulus-response—the most primitive use of information after genetic reproduction itself. The second step toward intelligence is a ***learned*** stimulus-response. I believe the final steps toward intelligence begin with ***language***. Further steps are taken as the ability to use language becomes more and more complex. Linguistic expression takes the form of statements. Statements are made to express one of the following: facts, conjectures and opinions, lies, questions, exclamations, jokes, emotions, analogies, descriptions, and stepwise pieces of longer narratives.

Humans have developed almost equal facility with two different kinds of language: written and spoken. This means that the more fortunate of us have learned to read and write, using our sense of vision (or braille—and there is some evidence that tactile input may be rerouted into the visual cortex of the brain in the case of the congenitally blind). Long before the invention of writing, our intelligence developed using spoken language and our sense of hearing. It is estimated that it took several hundreds of thousands of years for spoken language to evolve, but written language has only been evolving for a few thousand years. And only in the last few dozen years, programming languages and technical languages across the entire spectrum of technology have evolved. Two of the most complex of these are the Internet protocol and HTML (HyperText Markup Language), which drive all the communication over the Internet and the operation of web browsers that visit websites via URLs (or Uniform Resource Locators). Most of our newly evolved languages are used by us to interface with our machines. These languages are currently evolving rapidly. In our natural languages, many of our newly evolved words enable us to communicate with one another about this new technology.

The syntax of any language defines a set of terminal symbols that are used to form words and sentences in the language. Written languages have evolved two methods for constructing terminals. One is phonetic, based on the sounds of speech. The other is pictographic, based on elements that have visual meaning. Phonetic alphabets use the letters of an alphabet to represent sounds. These letters are concatenated in a given sequence (left-to-right, right-to-left, or vertically) that represents all (or most) of the sounds in a word. Pictographic writing uses a set of strokes that form icons to represent a word whose association is more visual than auditory. Spoken languages evolve much more rapidly than biological organisms, and a little more rapidly than written languages. Pictographic languages evolve more slowly than alphabetic languages simply because they are less connected to the sounds of speech (which may change in a generation), and more connected to their meanings (which may change every century or so). However, two new pictographic languages have recently begun to evolve: emojis for the Internet, and standard icons for public spaces. These don't use strokes; they are full-blown standardized pictures.

Consider how you learned, and learned to use, language. First you babbled, making all the easiest sounds, and were fascinated by the sound of your own voice. You did this in response to (mostly) your mother's voice, and perhaps the voices of your father and siblings. Soon, your babbling narrowed down to the sounds you were hearing. You began heading in the direction of your first language, refining your accent, and imitating a few words. With repetition and time, you began to understand some of the words you were hearing, and you began to formulate responses to them. Without realizing it, you began to "hear" some of your responses before you even uttered them. In fact, you probably talked to yourself aloud at first, and then silently. This silent conversation with yourself is what we call thinking. Conversations that we revisit over and over, constitute our narratives. Our repertoire of narratives is pretty much the bulk of our personality.

Thinking Machines

The question is probably not: "Will machines ever think?" But, "***When*** will machines begin to think?"

Computers are currently being applied to the tasks of recognizing speech and handwriting. It would also be nice if they could recognize a user's intentions in everyday human-computer interactions, but a computer's ability to ***understand*** is still at a very primitive stage. Actual thought, though perhaps at a rudimentary level, will have to be present for computers to perform speech and handwriting recognition that even approaches human accuracy. A computer will need to think, to some degree, before it can dialog with a human and carry out the user's intentions. Such a computer might be no more powerful than a $5000 machine is today, given that today's machines are thousands of times more powerful than those of 30 years ago. In any case, the time is not far off. As usual, the big delay is the software.

There is a positive feedback loop in the development of computers. Since almost the beginning, computers have been the primary tool used in their own development. This is one reason that the pace of their development has been so rapid. When a thinking computer is finally achieved, it will change the pace of development of every information and copy oriented human endeavor. A revolution will occur.

In nature, revolutions are often caused by positive feedback. What does a positive feedback loop imply? It implies two things: The output is being fed back into the input, and amplification is going on. In the case of evolution, the loop is closed with each generation. An entity is produced in a copying process. It matures to a certain point. Then copies are made of that entity. One loop has been made when one generation of computers has been used in the design of a later generation. This regularly takes only a year or two. Each generation is copied from the previous generation, just like anything that evolves.

In natural evolution, changes occur at random, and they are tested against the environment of the time. In design, changes are introduced not at random, but by conscious effort. The power of computers, because of this feedback loop, increases exponentially. Unfortunately, this applies only to the hardware. The generation time for software is much longer. Instead of being only a year or two, it is more on the order of five to ten years, or more.

Computers will never evolve the sorts of minds we have; any more than dogs or monkeys have. Nor, will we ever evolve the kinds of minds they will have. The “space” in which a mind can evolve is surely as large or larger than the space in which a body can evolve, and life on Earth has evolved millions of body types with incredibly large differences. The truth of these statements rests on the extreme probabilities involved, not on some universal principal that our minds are unique and sacred. Evolution does tend to invent similar things independently, such as eyes and limbs, but there are always differences. Even identical twins have different fingerprints. No, a computer will never think and feel as we do, not because it couldn’t, but because it is evolutionarily unlikely in the extreme.

Nothing organic or alive, certainly not an entity that can think, is ever invented or created from “whole cloth.” The mechanics of evolution are required. The principles of evolution are as fundamental to the way the universe works as the principle of gravity and all the other principles of physics. Given suitable conditions, the principles of evolution guarantee that order will arise from chaos. We aren’t sure just how likely the conditions that gave rise to ourselves are throughout the universe, but they did arise this once. We also aren’t sure what other kinds of conditions might be favorable to the evolution of life, but others seem plausible, and we continue to search for them. However, of one thing there can be no doubt. The conditions here and now are just right for the rapid evolution of machine intelligence.

Human intelligence is based on our ability to learn, repeat, and make up narratives. We make up stories to explain patterns that have been impressed upon us. Our intelligence is not based on sequential recordings, or on snapshots that record very much detail. Even our memory of narrative has limits. While once we may have been able to commit to memory a few very long sequences, today it seems like we only remember a large number of much shorter sequences. Is there another way to be intelligent without depending upon narrative?

Going back to the definition of intelligence as the stimulus-response repertoire of an entity, it’s hard to believe that symbols would not be employed in a thought process. But, maybe symbols could encapsulate deep complexities and occur without a grammar connecting them. Maybe performance and muscle memory would play significant response roles. Construction or manipulation of material objects might replace the construction of narratives. Perhaps a huge vocabulary of emotions could be learned and communicated, emotions that could differentiate between thousands of subtly different experiences.

Another aspect of human intelligence is our ability to read and write, and our ability to fashion tools to aid our natural ability to collect and process data. Without the leverage that all of these contribute to our intelligence, how far would we advance? Could a species surpass us without any of those advantages? Imagine a water world with creatures like whales or octopuses.

Suppose they were simply big blobs with huge brains. Brains a thousand times as big as ours with memories a thousand times more accurate. Suppose they lived for an average of several hundred years and spent nearly all their time in conversations with one another, telling and retelling narratives of all types. These narratives would evolve over time. Narratives begin as conjecture, but they can be argued and debated—they evolve. The fittest narratives (those repeated the most) survive. In place of reading and writing, the Blobs would keep their stories alive by retelling them, and these stories could survive for centuries. The best narratives might survive for millennia.

Spoken symbols can acquire any degree of nuance. Why not complex mathematical formulae? They needn’t be written down. Blobs could develop an ability to visualize, perhaps much better than humans. We have chess masters that can play very good games of chess without a chess board. Why couldn’t Blobs invent such games and play them? Imagine what other pastimes they might come up with!

Over the ages, Blobs could ask themselves questions, propose conjectures, and craft answers. A complex web of narratives could be built and continuously refined to increase its overall integrity. A Blob could just as easily start with the observation “I think, therefore I am” as a human could. Why couldn’t the Blobs go on from there to develop an understanding of the universe vastly superior to ours? It might take Blobs a thousand years longer than humans, but they would have the long-term advantage of not being able to fashion tools, and therefore they couldn’t develop a nuclear capability or the ability for self-destruction. Could a fundamental understanding of the universe evolve in such a way?

One thing is clear, this type of intelligent Blob will keep to itself, it won’t be coming after us!

However, that thought begs the question, “what kind of intelligence should we fear? Two answers spring to my mind. The most immediate is our own intelligence. We are probably our own number one worst enemies. If we survive ourselves, we might eventually face expeditions sent by another culture to locate our world for them to colonize. These expeditions would be silicone and metal based (artificial) life following the directives of a biological life form. Our species has a predilection for eliminating its competitors (including outgroups within our own species). It is reasonable to assume there might be more like us in the universe. A very long wait time would be involved in sending out search parties that could locate suitable planets to seize and colonize,. The most likely scenario for the biological side of this conquest would be a very small population that changed very little over many generations. It might discover new destinations only once in many tens of thousands of years, and it might be another 10-50 thousand years before colonization was accomplished. Each great Earth civilization seems to have evolved toward a state of maximum inequality, at which time civil discontent ended it. Given that we are already about a factor of 3-5 beyond the carrying capacity of our planet, it is not too hard to see how, with the advancing technology we have, a very small, and very wealthy, subset of people could eliminate the rest and begin a static, long term existence with the objective to colonize the galaxy, and have the patience to carry it out. The eventual success of this would depend on how well the biological half could keep the artificial half under control. A scenario like this might exist for tens of thousands of years, but evolution would not cease, and in millions of years this scenario would either vanish or become radically different.

Meanwhile, let's return to human intelligence, and how we might develop the first ALFs.

Neural Underpinning

The first step in the design of something new is to find one or more old things to integrate and copy. In this case, the obvious thing to copy is our own intelligence (or something on the intelligence spectrum between us and a fruit fly). Three things are needed: Sensors, computation, and effectors. We know that our intelligence is a product of our brain, and that the atomic unit from which our brain is constructed is the neuron. In fact, the back end of every sensor, and the front end of every effector is also a neuron. So, let's start there.

A neuron is an advanced form of that original cell that could sense a particular organic molecule outside its cell wall, and produce some amount of transmitter molecules inside its cell wall. Those transmitter molecules then diffused to another region of the cell where they triggered some kind of response. The major improvement made by the neuron is that a better method was substituted for diffusion through the cell, and the cells became arbitrarily long. Today, the morphology of a neuron is a cell body surrounded by dendrites (multiple extensions from the cell body that work as sensors), and a single, sometimes very long, axon that may have multiple axon terminals at the far end that emit neurotransmitters. Each contact point between an axon terminal and a dendrite of another neuron is called a synaptic gap, or synapse. When a neuron is excited by a pattern at its dendrites it increases or decreases the rate of signals to its axon terminals. Its response depends on the pattern of dendrite connections, and how much a particular connection contributes to the neuron's signaling rate. This pattern is established when a particular neuron, part of a neural packet, is first allocated in the task of learning. Over time, neurons that "fire together, wire together." This "wiring" appears in two ways. As a modification to each individual dendrite to encode the pattern recognized by that neuron, and as growth extensions of its axons to the dendrites of other neurons. A pattern may consist of thousands of inputs, each to one dendrite of a neuron that has learned to recognize and signal that pattern.

It is estimated that a fruit fly, a life form whose "intelligence" is on the order required to control a flying robotic drone, has about a hundred thousand neurons. These must be devoted to stabilizing and directing its flight to its sources of food (and sex) given a range of visual and chemical stimuli. At the other end of the scale, it is estimated that the human brain has about 70 billion neurons in its cerebellum, about 16 billion in its cerebral cortex, and a further 14 billion elsewhere (a total of around a hundred billion). Each of us has some experience with what those hundred billion neurons are capable of.

The original stimulus-response pathway was simply a cell that could detect a particular chemical at one end and effect motion at the other end. Today, a stimulus-response pathway involves a relay of connections. The first connection is a "hard" wired stimulus of a chemical molecule, or heat, light, sound, or other sensation that evolved because it was effective at information gathering. Intermediate connections follow the original paradigm of smell (a neurotransmitter) at one end and behavior (emitting a neuro-transmitter) at the other end. The last connection is made to stimulate a muscle movement or a glandular secretion. At the stimulus end of each step in this pathway (with the exception of the first), and at the response end (with the exception of the last), instead of a single point of "smell" or chemical production, there are thousands of points. This enables arbitrarily complex patterns to be detected.

When a neuron fires (after a period of training) it is signaling that it has detected its individual pattern. A low-level pattern might be as simple as an edge detected in the visual field. A high-level pattern might be a linguistic colloquialism. Patterns detected at a low-level form elements of higher-level patterns. Overall, patterns form elements of other patterns. When one piece of a pattern is activated, other patterns for which that pattern is an element may be activated. This is why our brains are so good at recall by association. Everything (sometimes even "white" noise) can remind us of something. Our thoughts seem to run in a stream of consciousness at times, especially when our attention wanders away from immediate sensations.

Our brain is organized into a left half and a right half that are mirror images of each other. Each half tends to be functionally different. The left sometimes specializes in sequential patterns, while the right half may specialize in parallel patterns. Sometimes even separate personalities may develop, especially if the two halves become disconnected. Sensory input from the left side of the body enters the right side of the brain, and vice versa. Thus, the two halves are distinct in several ways. In fact, the left and right visual fields from each eye are split and enter the opposite sides of the brain. The left brain gets a stereoscopic view of the right visual field, and the right brain gets a stereoscopic view of the left visual field. This permits information that interests each half of the brain individually to be routed to that half of the brain. The left brain being interested in the right side of the body, and vice versa.

Apparently, each side of the brain is capable of learning similar skills, but each side tends to develop specialties of its own. Thus, the left side in a right-handed person sometimes tends to become the more skillful speaker and listener, while the right side might become the more skillful artist or dancer. Often, tasks require both halves working together, and again one side will tend to specialize in one part of the task, while the other side specializes in another. In the normal brain, the two halves communicate with one another via a nerve bundle called the corpus callosum.

The outer layer of each half of the cerebrum consists of three lobes, the frontal, parietal, and occipital. These specialize, respectively, in thinking (perhaps giving us our sense of consciousness), language, and visualization. Regions underneath these are more primitive, helping with learning and mediation between areas, as well as regulating bodily functions and other survival mechanisms, such as emotion, and the production of “global” neurotransmitters, such as adrenaline and serotonin.

The brain is a pattern recognizer. The use of information begins with pattern recognition. Patterns and information are closely related concepts. Patterns may be detected at points, linear arrays, or sheets in the brain. These could be called 0-D, 1-D, and 2-D patterns. 0-D patterns are points signaling only intensity. 1-D patterns differentiate across a spectrum. 2-D patterns recognize an area. And, 3-D patterns might be constructed from a time sequence of 2-D patterns, or simultaneously from the slightly different perspectives of our two eyes (or slight timing offsets in the case of echolocation).

Sequential patterns are a sequence of terminals (the stimulus of a sensory pattern). Terminals are, in essence, elemental symbols. Sequential patterns are scanned terminal by terminal. Parallel patterns exist in a 2-D matrix of terminals. These terminals are integrated within a network of connections to recognize higher level patterns.

The pattern type, sequential versus parallel, depends upon how the pattern is presented. If the terminal symbols are presented over time, the pattern is sequential. If presented over space the pattern is parallel. If presented over both, the 3-D pattern has both aspects, and must be both parsed and integrated.

When a pattern is first presented, it must be recorded. When presented later it must be recognized. Every recorded pattern has associations. A part of pattern recognition is the recall of its associated patterns. Pattern recording involves not only the recognition of a pattern, but capturing its “downstream” significance, and its associations.

Similar patterns are tightly associated. Different patterns presented simultaneously are loosely associated. Patterns evoked through association can evoke further patterns through association. Stories and episodic memories may be recorded in a sequence of patterns triggering other patterns. One thing reminds us of another, and that of yet another, repeatedly.

A baby, when awake, is in almost constant motion, testing its muscle systems, and developing its sense of proprioception. Its senses of touch, vision, and hearing are also being constantly stimulated with first time instances of new patterns. Its senses of taste and smell probably contribute to the context of other patterns. A repertoire of taste and smell patterns seems to be lacking in humans.

All of our natural sensors evolved due to the selective advantages they gave us. We have since devised artificial sensors for their utility. Watches with barometric, magnetic, and temperature sensors are common—not to mention that many keep time with millisecond accuracy. One very complex example is the GPS. It takes a complex of satellites in orbit and a high-speed computer with a sophisticated program to report a position. Such positions are being reported around the world to millions of users all the time as people and goods navigate from point to point.

The GPS utilizes a very special radio broadcast. The watch examples measure variables in the environment. However, sensors exist that use the broadcast spectrum to communicate over a variety of other information channels. The Internet for example. Local point to point packet communication is another. Smart phones collect information from all of these sources.

When an array of sensors is coupled to a large and universal pattern recognizer (emulating the human brain), the potential for a very advanced intelligence exists. The sensors are available—the missing link is the design of an intelligent “brain.”

A human brain has 100 giga-neurons, each with an estimated 7K upstream synapses that connect to the dendrites of a given neuron, and 1K downstream synapses that connect via the axon to the dendrites of other neurons. Each dendrite makes a contribution to the signaling rate of the neuron’s axon. Signals go to a set of dendrites belonging to downstream neurons. A firing neuron is a point of sensation or qualia. A quālis depends on where it is in the brain. The sum of all firing neurons produces the overall qualia and sensations experienced by a brain at a given moment. The brain is constructed in layers. The neurons furthest upstream are triggered by sensors. Those furthest downstream trigger muscles. To carry the stream analogy further, the layers in this stream may also form whirlpools (feedback loops).

Two different processes are at work in the brain. A learning process and a stimulus-response process. Both are no doubt very complex and involve many steps and differences from one part of the brain to another. However, some generalities may be mentioned. Learning occurs within learning “windows.” The first neurons in a complete neural pathway are trained in windows that open and close shortly before and after birth. Each neuron becomes specialized with a pattern that is imprinted on it at some point in time. As the first window begins to close, the next window begins to open, and so on, until the brain is mature. Once trained, each layer’s stimulus is produced by the previous layer’s response. The final response triggers a muscle or a secretion. Modern technology has developed a plethora of sensors and effectors, more than sufficient to build and artificial intelligence. The only thing missing from the equation is an artificial neuron. Even now, attempts are being made to simulate the apparent abilities of biological neurons, but so far a “critical mass” has not yet been achieved. Given that a neuron is the apparent building block of our intelligence, let’s explore what might be involved in designing one.

The Basic Intelligence Quantum

Let's call the computerized equivalent of a neuron, a Basic Intelligence Quantum, or BIQ (pronounced "bik"). The idea is that a BIQ can recognize a single stimulus and trigger a single response (over a range of intensity and possibly incorporating some timing delays).

The brain of an artificial life form (ALF) might contain a hundred or more Giga BIQs arranged with arrays of sensor points of input, and arrays of effector points of output connected by up to ***n*** intermediate arrays. Each array of BIQs would form a layer with several thousand connections to the previous layer and several thousand connections to a variety of downstream layers. All the connections made by a given BIQ to its downstream BIQs signal that a given pattern is being recognized. All the connections received by a given BIQ are "tuned" in a learning process to recognize a particular pattern of inputs, and their net effect is to determine the signal level of the BIQ's output. The initial sensor inputs, and the final effector outputs are electro-mechanical transducers of some sort, each acting at a point within one of several arrays that correspond to a type of sensor or effector.

For the purposes of the design described here, a BIQ is a data packet in the memory of a computer. This data packet contains everything necessary to describe the state of the BIQ at each increment of time (on the order of a millisecond). This includes the status of each of its "dendrites" (sensor inputs) and the addresses of all the downstream dendrites that are connected to its "axon terminals." More details of this package will be explained later.

Each BIQ (the data it contains) is updated by two algorithms that cycle once per "increment of time." One algorithm sends signals downstream, and the other updates BIQ parameters in a learning process. The actual increment of time taken by each process may be different. The delay between stimulus and response is a function of the signaling process. The quickness of adding or updating a stimulus-response behavior is a function of the learning process.

When an ALF is "born" (comes off the assembly line), it might contain a brain with a fixed number of BIQs all set to zero. Each BIQ would be a sequential string of bytes in a large memory capable of holding, let's say, 100 giga-BIQs. We are choosing numbers suggested by what we know of the human brain. Each BIQ should allocate at least 4 bytes for a dendrite and 4 bytes for an axon terminal (it should probably allocate even more than this). Three of the dendrite bytes would be parameters adjusted by the learning algorithm and referenced by the signaling algorithm. The fourth would represent the firing rate of the upstream BIQ. Based on the values of its input and learning parameter bytes the signaling algorithm would compute the value of a BIQ's output and transmit that value to the input bytes of its downstream BIQs. There could be even further bytes allocated to communicate from the signaling algorithm to the learning algorithm.

I won't speculate too much on the details of the two algorithms and the precise meanings of the bytes used to communicate between them. These details will evolve over successive generations of ALFs.

How do these requirements translate to the size of memory needed? The 4 bytes used to represent an axon terminal would contain a "byte skip" value to the firing rate byte of a downstream BIQ. If there were a thousand connections, both upstream and downstream, that would entail 8 thousand bytes per BIQ. For a 100 giga BIQ memory, this is a quadrillion bytes, or a petabyte of memory. This would require a 50 bit absolute address, which would require 7 bytes per address. This is quite a large memory, and it would require massive parallel processing to cycle through it in a time even approaching a millisecond.

When you think about a signaling algorithm running concurrently with a learning algorithm, you can perhaps see why the signaling algorithm should be suppressed when the learning algorithm is operating. If this is to be the case, normal behavior would need to stop when the entire data structure were being updated by the learning algorithm. The signaling algorithm defines the length of time a basic reflex response takes to occur. Various strategies exist so that the learning algorithm can avoid upsetting the signaling algorithm. For example, sleep, updating one part of a parallel (split) brain at a time, or keeping the design of a BIQ packet so simple that it can be modified without upsetting the signaling algorithm.

The existential state of an ALF is the pattern of its BIQ signals at a given moment. It changes from state to state as each of its BIQs are visited in complete update cycles. This involves updating each BIQ's signal level and performing housekeeping functions, which include updating learning parameters, and effecting responses.

Given a modality, such as vision, sensation would start with pixels (picture elements) on a retina. A 2-D layer of these would feed a 2nd layer of BIQs that would connect nearby pixels and produce the first level of pattern recognition. In a wet eye and brain, these two layers are located on the retina, itself. The eye has fewer neurons in the optic nerve than it has sensors on the retina. Therefore, significant integration to the first level of pattern detection must occur on the retina itself. The 3rd layer is located in the brain. It is reached by a bundle of neurons called the optic nerve. Several more optic layers are followed by several more layers in the cortex before a complete perception of a visual field is achieved.

In general, given a modality, neurons are arranged head to tail in a sequence (a line or a queue). Head refers to the neuron's cell body and its many "noses" which smell out the signals being passed to it by neurons closer to the head of the line. Tail refers to the other end of the neuron where an "odor" is produced for the next neuron to smell. For the most part, a modality consists of a large number of lines operating in parallel. In my terminology, a ***layer*** consists of all the neurons operating in parallel at the same position in their respective lines.

A brain is trained layer by layer. The first few layers are trained early in life during a "window" of maturation. All of the layers in a human have been at least partially trained by the age of 3-4, but the layers in the cortex continue to be trained throughout life. Stimuli have one or more of four results: They may be ignored, they may produce an internal response, they may produce an external response, or they may be novel enough to be recorded as a new pattern.

The senses of smell, taste, and touch are virtually preprogrammed. They contribute their bit to overall pattern learning and recognition by associating with the other senses, and by evoking degrees of appeal or disgust. The first 2-3 layers of the visual, somesthetic, and auditory senses develop in the first few years of life.

The first two senses that require learning are vision and hearing. A baby is held and fed (stimulating touch, smell, and taste), and it forms visual patterns in response to its local environment, especially the faces of the people it encounters. It signals its pleasure or dissatisfaction by cooing or crying. It imitates the expressions on the faces it sees. After several months, its sense of hearing begins to develop, and it attempts to mimic the sounds that it hears other people making. At first, its sense of hearing merely discriminates between loud or lulling sounds, startling or soothing it. Much of a baby's early learning involves its emotions of pleasure or discomfort.

Simple Brain–Simple BIQs

A basic intelligence quantum, or BIQ, is one unit replicated many times in the brain of an artificial life form, or ALF. The division of a BIQ into something smaller doesn't define anything useful. The total count of BIQs in an artifical brain defines the limit of that brain's potential intelligence. Intelligence is the repertoire of an ALF's stimulus-response behavior.

A BIQ is exactly what is needed for a repertoire of a single stimulus and response—that first "iota" of intelligence. Let's consider a "brain" with a stimulus-response repertoire of only 95 BIQs. This brain would not be intelligent in the least, but it would be a trivial start on an ALF design limited to the input and output of text.

To begin, consider an "eye" whose job it is to scan a white page for black print and detect a single ASCII character. A 5×7 array of sensors might suffice. When presented with all the different characters, the BIQs of this device would be trained to signal when any version of a particular ASCII character were presented to the array of sensors, and to be quiet for all other stimuli. Different BIQs could be trained for different characters. A different modality of BIQs could be trained for different frequencies of sound, or different chemicals present in a stream of air or water, and so on. These BIQs would be at the "front" end of a modality. Each of these would produce a response to its unique stimulus, and that response would be a stimulus to another layer of BIQs. Ultimately, the response of a BIQ sensing a pattern would trigger an effector of some type, causing a behavior that would have an effect on the environment (or send a signal to communicate with another device).

Let's continue to design the first layer of the visual character reading modality. To keep things simple for this "proof of concept" we'll specify that the 5×7 sensor array is positioned above a page of paper that can be shifted one character width at a time from the upper left corner of a page to the right edge, then with a "carriage" return and line feed of a character height, one row down and back to the left edge, then across another line and down, until it reaches the lower right corner of the page. A lens insures that one character at a time is focused on the sensor array. Each cell in the array, due to the amount of black vs. white focused on it, produces a stimulus in the range of 0-255. This stimulus is fed to all 95 BIQs in our "simple brain."

Now, our only objective is to recognize and distinguish 95 different ASCII characters. We have already seen that character code = 32 is a space. Codes 33 through 95 each have some non-blank print value. This means that some of the cells in the 5×7 array will be partly black. The amount of black will be measured by the stimulus (0-255) produced by each cell.

To achieve our objective, these 95 BIQs need to be trained. Each BIQ has 35 "dendrites" initially with random settings. When a character is presented, one BIQ will match it better than any other. This BIQ needs to be adjusted to increase its response to the pattern presented, and all the other BIQs need to be adjusted to decrease their response to this pattern. This process is repeated for successive presentations of characters so that eventually most presentations of a given character are followed by a response from a single BIQ, and seldom is there a tie. We now have a "brain" that can scan and recognize ASCII text.

This design leaves a lot of room for improvement. It would probably have trouble with a variety of type fonts. When each of its cells is considered a pixel, it would benefit from having a lot more of them. It would be very sensitive to a character being positioned correctly. Some kind of feedback to the mechanics of the reading device would be helpful. Lastly, this design has no way of buffering characters and assembling them into more complex patterns to be learned by "downstream" BIQs in order for words, phrases, and sentences to be recognized. But, it's not too difficult to see how these abilities might result from scaling up from this simple example.

The Architecture of an ALF

If a BIQ is the building block of a brain, what is the overall architecture of a brain? Let's start with the features that wet and dry brains have in common. The term ***wet*** will refer to a biological life form, and ***dry*** will refer to an artificial life form. Our wet brains have a half-dozen, or so, sensory modalities. A dry brain could get by with only one, but more could be added if they became important to an ALF's mission. Behavior is produced by effectors. A dry brain could get by with a single type of effector—a producer of text. This text could be directed anywhere for the purpose of controlling almost anything using protocols. A wet brain controls muscles. It probably also controls certain chemical secretions for various purposes around the body. Approximately 700 muscles are under the control of the human brain.

Think of a brain in terms of the following schematic. Imagine a grid with a very large number of cells. Along the left side of the grid are the first cells of each row in the grid, and along the right side of the grid are the last cells in that row. The first ***n*** rows of the grid (numbered at the left) correspond to a sensory modality. The next ***n*** rows correspond to a second sensory modality, and so on, so that each modality has a large number of rows to implement it. This makes the grid very high, but it's not all that wide, maybe only 10 to 20 cells from left to right. On the right, there might be 700 rows to represent the muscles in the human body, or 256 rows to represent the different text output values of an ALF. If you have imagined this correctly, you notice that this matrix doesn't line up very well. In fact, each cell from left to right connects to cells in rows above and below it, as well as to the cell to its right (which might actually not exist, since there are far fewer rows on the right side of this matrix than there are on the left). Each cell is a BIQ. A thousand or more BIQs feed into each BIQ from the left, both above and below. And each BIQ feeds into a thousand or more BIQs to its right, both above and below.

If we had such a chart for the human brain, different brain areas could be outlined on it. BIQs contained in a region of the brain are contiguous on the chart. However, the rules for where the axon terminals can connect are not perfectly rigid.

An ALF needs only a fraction of the system a baby is born with and needs to develop. It needs the modality of vision and it needs to effect sequences of byte output. The design of a minimal ALF intelligence will be the objective of the rest of Part 2.

The Artificial Brain

What do we know about the human brain that would help us to engineer an artificial brain? We know that the building block is the neuron, and that the entire brain consists of almost nothing but duplications of this basic element. We know that many connections exist between neurons. We know that regions of the brain are specialized to certain modalities. The activities performed by each neuron are pattern recognition and association, the operation of muscles, and the production of specific transmitter molecules (both of the latter at various points around the body).

Early experiments on exposed brain tissue indicated that the stimulation of a particular neuron on the surface of the cortex caused the vivid sensation of a previous experience. From this I deduce that some memories are eidetic, that patterns can be coded into single neurons, and that entire experiences are captured in neurons of the cortex (as opposed to lower level patterns which are retained in more specialized lobes of the brain).

The general organization of the brain is that a bundle of axons transmits all of the information collected by a given type of sensor into a single area of the brain. Likewise, coming out of the brain are bundles of axons that signal specific effectors that activate individual muscles or glands. Inside the brain, one area may communicate with another, either via a bundle of axons, or by a method I'll call "layer to layer." The only difference is that when a bundle is involved, the source and destination are separated by a greater physical distance. A "layer" is (topologically) a 2-D surface of sensor points, or a 2-D surface (or sheet) within the brain itself, from which neural signals originate. Neural signals propagate from an initial layer to a 2nd layer, and that to a 3rd layer, and generally to several more layers, depending on the initial sensor layer and the functional area of the brain (the combination of which is a modality).

Each neuron in a layer after the first may be signaled by any of a few thousand neurons from the previous layer, and it may, in turn, signal several thousand neurons in the next layer. The locality of signals is preserved in the mapping from one layer to the next. Feedback to the same or a previous layer might also be part of the initial wiring.

Signaling in the brain takes place in two ways. The most primitive is like the original sense of smell. A particular chemical is detected, and this causes a particular response by the cell that detects it. The neuron is a cell specialized to send such signals. This is accomplished when a dendrite detects a neurotransmitter and sends a signal to a host of axon terminals to produce a neurotransmitter at the other end of the neuron (which may be very close, or far away). The chemicals detected by neurons need not come from another neuron, they may come from glands or other sources within the body (or foods or drugs). This method of transmission requires the chemicals to diffuse through the body. Short distances may be covered directly, longer distances may require diffusion through the blood stream.

Transmission of signals via the diffusion of neurotransmitters is a slower and more cumbersome process than transmission via digital electronics. When a dendrite is signaled, it increases or decreases by some amount and with some delay the probability that its neuron will fire. Dendrite parameters change over time with each pattern recognition in which they participate. How this change is effected is the crux of the brain's learning algorithm.

Studies of the brain are beginning to reveal its "connectome." Other studies have shown what parts of the brain "light" up when it is engaged in various activities. Both of these projects hint at its organization. Some of this organization will no doubt be imitated in early artificial brains.

Although brain damage has revealed that the brain is organized into areas tasked with specific abilities, most malfunctions of the brain, including its behavior on drugs, are probably not too helpful in deducing a good design for an artificial brain. More examination of the workings of a dendrite and how it is modified during the learning windows of maturation, on the other hand, would be very helpful indeed.

Vision

Perhaps the most important ***sense*** contributing to human intelligence (as opposed to the use of language, which is the most important ***ability***) is the ***input*** of a visual field, and our ability to ***visualize***. Our ability to ***output*** a "visual field" is limited to writing, drawing, and painting. Our visual field is constructed by neurons that sense bright light in 3 different frequency ranges, giving us color vision, and dim light of various shades of gray, giving us night vision. Our computer's ability to both input and output a visual field is much superior to our natural ability (using CCDs and LEDs respectively), but we get by pretty well with the visual system of our primitive wet brains. The following description of vision is oriented toward our future mechanical offspring, but it does follow our own wet brain version of vision.

A visual field consists of an area populated by pixels at some density (pixels per inch both horizontally and vertically). For input, focused light is detected by each pixel. For output, a pixel is generated by a tiny dot of light or ink on the surface of a display or piece of paper. The mapping between the device that detects a pixel and the light or ink that produces a pixel is both an art and a science. It involves correlating human perceptions with all three points, so that the same color and intensity that is sensed by an input pixel is produced by a light pixel or an ink pixel—or so it needs to seem to a human observer.

The first level of pixel integration is mechanical. Since pixels are tiny dots, or circles, when they are packed together in triads on a plane surface they pack as a series of equilateral triangles with every other triangle inverted. The pixels are located at the vertices of these triangles, and each pixel is one of the three primary colors. Although the exact frequency of a color depends on whether it is light or ink, let's label the primary colors R, G, and B. Thus, each row of pixels might be a series of [R, G, B].. with each row interleaving so that a B would be below and between the R and G above it. Then, triads of pixels would form a rectangular grid such that each pixel triad would be at the center of an R, G, B triangle.

Light focusing on, or emanating from, a single pixel is the limiting factor for the resolution of a visual field. However, any triangular group, or pixel triad, always contains one pixel of each color, and this mechanical coupling is the first level of pattern integration. The intensity of each color could be dictated by an 8-bit byte (with values 0…255), so that 3 bytes are used to dictate a pixel triad. However, various techniques are used to compress visual data, the most elementary being the color palette. This enables a mapping between an 8-bit value and a 3-byte color. When used correctly, the loss of information is barely noticeable, but it saves 2 bytes per pixel triad. This enables a 2^{nd} level of mechanical integration.

Every receptor, be it visual, auditory, or touch, has a threshold of perception and an intensity over a range. The intensity is a digital count with 0 indicating no input, and 1 indicating a value just above the threshold. Each higher count indicates a "just noticeable difference" (or JND). Depending on the sensor, each count may approximate a linear or an exponential difference. A count of 256 is more than adequate to simulate the range of neuron firing rates.

In the brain, each pixel represents a primary color or grayscale, and its firing rate indicates the intensity of light it senses. Good design attempts to decouple independent variables, so it might not be a good idea to couple both color and intensity. Evolution probably has it right. When it couples two traits, it usually evolves a separate trait altogether, or a significant improvement in a trait. One design feature of the eye, however, is hard to justify. That's the blind spot, where all of the axons come together on the front of the retina, and go out the back in a bundle (instead of coming directly out the back and forming a bundle there).

Display Input/Output

In the beginning, ALF design will no doubt utilize a display for input and output. Characters, pictures, and the image from a camera can be mapped onto a display. The vision modality of the ALF need only be connected to the display. Characters do not need their own separate modality. Likewise, in the simplest design, output from the ALF can go directly to the display.

For purposes of learning, there is no extra advantage to using sound as either input or output. The first goal must be to produce a level of intelligence incorporating an effective learning program into a properly architected collection of BIQs.

Virtually all of human knowledge and wisdom can be captured in language and pictures. The visceral sensations of our bodies can be alluded to, but not understood without experience. This is true for all of our other senses. This means that an ALF will not be capable of many human experiences, but that was never the goal in the first place.

An ALF can be led to understand that other sensations exist, and they can be described and given a context, but there is no reason they need to be experienced to develop intelligence and abilities. When an efficient design for machine learning has evolved, the set of sensors and effectors could be extended indefinitely for future ALFs.

I believe a design for learning can be developed using a standard computer display with text output directed to the display. ALF input could be limited to the ALF's ability to inspect a standard computer display. Inspecting a display, is very much like the operation of the human eye. ALF "behavior" could be limited to the output of ASCII text, similar to a computer operator typing at a computer's console. The text output could be echoed to the monitor, but would also be interpreted by the computer's operating system to result in more complex tasks. Simply parsing text bypasses several of the steps required for listening to speech. This should enable fewer and simpler BIQ layers to do the same job.

Let's review some of the basic physiology of the eye to get a sense of how it works. The structure of the human eye is not one that most engineers would have chosen. In fact, the eye of a cephalopod appears to have a better design. Some of the "design specifications" of the eye are as follows. There is a lens in front, and a retina in the back that receives light focused by the lens. The retina has a peripheral area, and a region in its center called the fovea. Concentrated mostly in the peripheral area are about 120 million "rods" and in the fovea about 6 million "cones." The rods are more sensitive to low levels of light, but not to color. Three different types of cones are each specialized to sense different colors. The lens of the eye is surrounded by an iris that controls the amount of light that enters the eye. The retina of the eye has a blind spot where all the sensor axons gather to form the optic nerve that goes to the visual cortex of the brain. The eye is capable of rapidly changing its direction of view—a motion called a saccade. I conjecture that saccades are useful to bring parts of the visual field onto the fovea, and that most of the information processed by the visual cortex comes from that part of the eye. In fact, the fovea represents only a 2-degree span of the visual field with about 7 megapixels. However, due to the way that the eye covers the visual field, by moving the fovea around and feeding the visual cortex with information from a much wider field, the eye simulates an effect that would take approximately 576 megapixels on a display screen.

The stereoscopic vision produced by the human eye would not be necessary in the initial ALF design. Modern digital camera technology is fully capable of producing pictures that rival the ability of the human eye. One lesson that might be taken from the eye, however, is that it moves a relatively small fovea over a much larger visual field. Our experience of a large and unbroken visual field is built from the smaller pieces of a series of snapshots produced by each saccade of the eye. Thus, some equivalent of a fovea and saccade should be implemented. The first level of pattern recognition in the human eye involves edge detection, shading, and coloring. A couple of layers of pattern integration take place in the eye itself. This is evidenced by the fact that there are fewer axons in the nerve bundle between the eye and the brain than there are pixels on the retina of the eye itself.

Constructing a Brain

A wet brain is composed of at least four subsystems. Two of these are the sensor and the effector subsystems. Two others are the mediation and survival subsystems. Each of these four has subsystems of its own. Each subsystem is devoted to a class of tasks. Each subsystem is connected to the others in a unique way that corresponds to its function. All of the neurons within a subsystem are initially very much the same, and they tend to change the same way over time. With respect to learning, some subsystems are more plastic, others are more rigid. In the human brain there are probably no more than several dozen different subsystems; an artificial intelligence would probably need no more than half a dozen, or so.

Various computational schemes have been devised to simulate the behavior of the brain. However, the brain is profoundly connectionist. Even its limited ability to compute is based on its ability to sequence a set of connectionist patterns. When we have adequately simulated the connectionist structure of the brain, and this includes a teaching scheme to program the connections, we will have begun to tap the power of this architecture.

The evolution of an ALF begins with its brain. A first brain based on a BIQ design has to be constructed, trained, and evaluated. It will likely be connected only to a simple computer interface consisting of a keyboard and display, but with the maximum memory and speed available. After some evaluation, modifications will be made to the BIQ design, and the update and learning algorithms, and a 2nd generation will be trained and evaluated. This will go on to a 3rd and subsequent generations until a worthwhile outcome is achieved. Once a degree of intelligence can be put to use, future generations of ALFs will be able to participate in their own development. This implies that ALF evolution will proceed exponentially.

Training a Brain

Each human brain begins operating shortly before birth, and undergoes most of its training in the first few years of its life. The degree to which a human brain can be trained falls off as a function of age. Artificial brains must either be trained in a process not unlike that of a human brain, or be copied from an artificial brain that is already trained. Artificial brains can be duplicated, but after duplication, each copy would diverge to a different personality and skill set. It is doubtful that a human brain will ever be copied, simply because when that technology is possible, the human brain will probably be out of the loop. However, the ability to copy a brain does not guarantee a longer lifespan. The capacity of a brain is fixed at the outset. An old ALF is not likely to learn new tricks unless new brain design principles are evolved.

Every brain is trained by interaction with the physical world. The brain is stimulated by the physical world. It responds. The physical world is changed by the response. And, the changed world provides a new stimulus to the brain. In this way, feedback continues as long as the brain is motivated to repeat the loop. The survival instincts evolved by wet brains ensure that this happens. This loop, outside stimulus, a response affecting the outside, and stimulus from the changed outside, has an important analog in another loop. That is the stimulus of hearing the speech of others, and the response to that speech, which we call dialog. When we learn to create this loop inside ourselves, it becomes our thought process—our ability to think.

The intelligence of a brain, or the degree to which it can become proficient at a particular skill, depends upon its design parameters (together with the abilities of its effectors). Human proficiency is very well quantified by IQ tests, Olympic records, the Guinness Book of Records, and other documents. Minimum and maximum limits exist. No such limits are defined for artificial entities. Each entity reading this has probably reached the apex of their own abilities. The apex of human ability, in general, has also probably been evolved. The apex of ability for human artifacts and their progeny has not even begun to be measured. However, it will likely evolve faster, and ultimately far beyond, that of human abilities.

Personalities, the very essence of a person, are products of their meme environment and their interaction (training) with it. The meme environment is a product of meme evolution. It consists of what's spoken, what's written, and what's practiced in every culture around the world.

A baby's first interactions involve its mother, its food, facial and approval recognition, and imitation. The first interactions of an artificial intelligence are likely to begin with word-picture associations, and some form of approval recognition (and possibly facial recognition). Both human and ALF learning begin with an essentially "blank" brain ready to be trained.

A neural snapshot is taken of each new thing an empty brain is exposed to. At first, no higher-level patterns exist. The visual, auditory, and somatic areas of a wet brain are stimulated with elemental patterns of light, sound, touching, and feeding. A context exists for each of these, and as they are repeated, they become reinforced and associated with patterns from the other senses. Triggering of lower level patterns becomes the source that defines higher level patterns learned by layers of neurons downstream.

Each complete experience is recorded by millions of neurons involving patterns across the entire sensory spectrum. As the days and weeks progress, more and more of these snapshots are recorded and integrated into recognizable patterns. Simple patterns produce the stimuli that define more complex patterns.

Babies notice action, and they have an innate tendency to copy it. They play the "imitation" game. This will have to be a part of the programming for an ALF, as well. Babies also have innate tendencies for positive and negative feedback to encourage or discourage their behavior. Again, forms of feedback will have to be part of the ALF learning experience.

Baby animals of all species appear to have phases of learning where successive layers of neurons are trained after previous layers have learned a repertoire of patterns appropriate to the next layer. The maturation process might hold off the training of each layer until the layer beneath it has been trained to a sufficient extent.

The 1st layer in the brain to be stimulated is connected directly to a sensory modality. Each layer senses stimulation from a previous layer and produces an output for each pattern detected. After one or more layers, the BIQs stimulated by the previous layer are part of the efferent system. The output from these pattern recognizers goes primarily to the muscles that produce behavior. Thus, pattern recognition may produce action—either reflex involving one or two layers, or reasoned action after a number of layers. Thought processes (subvocalization) are actions that don't show up as external behavior. There may be less conscious ones also going on in the cascade of activity generated by sensations and the effects of "stream of consciousness." Learning to tell stories, and certainly learning to read, train the response of subvocalization—or ***thinking*** as it is often called. Kahneman identifies two systems in the brain that are responsible for behavior. They might not apply to an ALF design, but they explain a lot about human intelligence. System 1 (thinking fast) is a direct pattern recognition response pathway, and system 2 (thinking slow) appears to involve an internal narrative.

Neuron Emulation

A neuron exists in a bath of chemicals delivered by the blood stream and upstream synapses. Thus, it is positioned to "smell" a set of molecules. Its construction causes different responses to different molecules in different amounts. The molecules that stimulate a wet neuron can be given token values when simulated by a BIQ. Synapse connections can be given address values. Thresholds and actions can be given numeric parameters.

A neuron has two jobs. Job 1 is to adjust its inputs so that it recognizes a particular pattern. Job 2 is to react to its array of inputs, and signal the detection and intensity of its pattern. The 1st job is accomplished by establishing and retaining the settings of each of its dendrites (its inputs). Each dendrite is triggered by the axon terminals from upstream neurons. A pattern of stimulation over the collection of dendrites attached to a neuron determines whether (detection) and how fast (intensity) it fires. Each time it fires, it releases neurotransmitters at each of its axon terminals, and these are the inputs to other pattern recognizers at a host of other neurons "downstream" from that neuron. A typical neuron in the human brain might collect inputs from over a thousand different upstream neurons, and send signals to thousands of others downstream.

The initial sensor (connected to the outside world) is hardwired, and its "meaning" is predetermined. It has only to signal a degree of input. A wet brain does this by not firing (or firing sporadically) when it is not detecting its pattern. Its firing rate increases linearly or exponentially by JNDs up to some maximum. This firing rate has the effect of producing a level of neurotransmitter concentration at its axon terminals. A BIQ would simply transmit a single number each time its input changed or was triggered. Thus, it would not have to simulate each firing that a wet neuron produces.

Typically, token values (simulating a particular neurotransmitter) have the same meaning over a lifetime. The address of a BIQ doesn't change, but the list of addresses each BIQ references could change over a lifetime (these simulate connections). In particular, reference addresses could be added or deleted from the list of a particular BIQ. Other parameters also change over the course of a lifetime. What these parameters represent, and how they change, are the core secrets to the learning process of a dry brain populated by BIQs and serviced by a host of microprocessors.

Initially, neurons are wired into a packet, and packets are wired into other packets. There are three types of connections: cross connections, longitudinal connections, and feedback connections. In a computer, connections to other neurons are made via a list of reference addresses. In the human brain, connections (for the most part) are established at the outset. They grow more numerous if activated, and are discarded if not activated. In a computer, connections could be added as and when they became useful.

In the human brain, the process of reading aloud would involve many millions of neurons working in parallel, each operating with a cycle time of a few milliseconds. In a computer, the same number of neural packets might have to be serviced one after the other, each with a cycle time of a few microseconds. To achieve comparable results from the two types of brains, a factor of a million might be necessary between the serial and parallel cycle times. Specialized hardware and algorithms could probably gain the factor-of-a-thousand-gap between millisecond and microsecond cycle times.

Perception is the integration of sensation. Everything perceived for the first time is recorded as a neural snapshot of a pattern. Later perceptions that are similar may add to the snapshot to build a prototype pattern. Having perceived a sequence of events, a brain needs the ability to play it back for a limited period of time. It also needs to be able to commit a sequence or a pattern to long term memory, which means being able to reconstruct it later on. A pattern is all the stimuli that occur in close proximity (either physical or temporal). A pattern is learned when all of the neural packets involved have been adjusted so that they trigger together when some fraction of them is triggered separately at some later time.

Downloading a Brain

Google a couple of the articles on "downloading a brain" and you will notice that it's also called "uploading" or "mind transfer." A project called the "connectome" is underway to map all the synapses in the brain. That may show how the brain is wired at birth, but it will not show how the synapses are changed by learning. Details of all the synapses will require much finer resolution, and vastly more data. Currently, it's not even clear what we would be looking for.

Given the connectome of a given brain (yours, for example), and the pertinent molecular details of each of its synapses, you would theoretically be in a position to duplicate it. If you duplicated it on hardware, compromises would have to be made. If you duplicated it into another "wet" mass of neural tissue, I can't even imagine what mechanism you would need to use.

In either case, what would you have? Would your conscious awareness "go there?" Or, would there be an ALF whose personality resembled yours? The 1st case seems farfetched. The 2nd case would be an approximation at best, and the brain would have the same level of maturity as the original. Can a brain's lifespan be extended?

I believe that a given connectome has a given capacity for learning. As it takes on more and more engrams, it approaches a natural learning asymptote. If this is true, a brain has a finite limit, and prolonging its age can only be done in the long tail of its learning asymptote. An analogy might help to understand this. Imagine the brain as a large sheet of photographic paper (Google it, if you haven't heard of photographic paper!). Images can be sharply recorded as long as unexposed areas of the paper remain, but when all the area is used up, images become superimposed. The more superimposed images there are, the less detail that can be extracted from each new image, and the more that overlapping patterns will be confused.

Could extra memory be added after the fact? That's another good design question. Given the existing somesthetic sheets, new sheets would have to be wired into the existing memory. Layer upon layer of new memory would need a connectome similar to the original connectome. It's not clear how the old and new memory might communicate. This might not even be feasible. The problem is with the hardware analog. With hardware, or wetware, new "wires" would have to be inserted into the mass of wiring already present. With software, the wires would be represented numerically, but the attachment problems would be similar, and figuring out what should be attached where would be a very difficult problem.

Special Applications

One huge advantage of an ALF over a human is the ability to download, install, and run various ***application*** packages or apps on the computer that it shares with an operating system. A variety of these, though still rudimentary, are available to perform speech generation given ASCII text, or speech recognition that produces ASCII text. Computer controlled driving, and walking robots exist. Facial and other recognition programs exist. Drawing, text editing, and other display output programs could facilitate an ALF that has learned what these programs do and how to invoke them on its own hardware. In fact, after an ALF has learned to program a computer, it could write apps and install them into itself.

Artificial Learning

Learning involves experience—many episodes of experience. An experience is a barrage of sensory inputs followed by an internal perception, and then usually, an external behavior. The wiring in the brain that causes a stimulus to be followed by a response appears to be already in position. It just needs to be connected up. This is what happens in the process of learning. A wide pattern of sensations are connected together so that when parts of a pattern are repeated, the rest of the pattern is recalled. Each layer of processing in the brain has its own sensations triggered by the pattern recognition of the layer in front of it beginning with the layer of raw sensory input.

The brain responds to differences; it responds to the unexpected. Its responds by making itself able to remember and recall. Our association memory is generally much better than our recall memory. This is because associations are the basis for all memory. Recall involves "jump starting" an association by feeding a related pattern to ourselves as an internal stimulus.

Given the "tabula rasa" of a new-born baby, each memorable sensation (and most are) produces a neural snapshot. Each snapshot involves thousands of neurons. Patterns in the first layer of neurons are built up as snapshots, and snapshots of these patterns are built up in a second layer, and so on to a third and subsequent layers. One layer is developed after the previous layer's development matures. As the months and years progress, and new layers mature, higher cognitive abilities are developed. In a human, the first layers are probably developed in the first 6-18 months, the next few layers, in the next few years, and the final layers probably mature over the 2nd and 3rd decades of life.

Each new pattern we experience causes an epiphany to varying degrees. Some may be life changers, setting us on our lifetime course. The most fundamental thing our brain does is to notice a pattern—and record it, relate it, and retain it.

Although human learning includes such things as facial and other emotional expressions, body language, and interpreting internal signals such as those for the basic drives of hunger, affection, minimizing pain and discomfort, and so forth. An ALF would need to be concerned with far fewer sensory patterns and behaviors. Below, we shall continue to limit ALF design to pixel input and text output.

Although it takes years to train a wet brain, a training program could be designed to interact with a new ALF brain, and speed the learning process up an order of magnitude or more. This could be done in parallel with a number of new brains set with different learning parameters, and using an evolutionary algorithm to pick the best of a generation of ALFs. Once a degree of success at a particular learning stage had been achieved, new ALF brains could be initialized with copies of previous ALFs, and developed further from there.

New generations could be put through this program and a set of viable training parameters might be evolved fairly quickly, perhaps within a few years. The various differences between ALFs at a given generation could be based on picking the best ALFs of the previous generation, and trying different numbers and sizes of layers, as well as different parameters for the basic BIQ design. Once ALF intelligence is employed in its own evolutionary feedback loop, it would proceed more and more rapidly—its development over time would increase exponentially.

{?? Begin current work}

Detailed BIQ Design

An ALF brain would consist of point, vector, and matrix BIQ data structures arranged in layers. These layers would be sequenced from layer 0 to layer ***n*** (***n*** < 10). All 0-layers will have housekeeping functions that depend on the input type (change only, each discrete input, or continuous). All 1-layers and greater will be kept current in parallel (multiprocessor or multitasking). Housekeeping during the learning window for each sensory modality is done by a specific routine from a set of standard routines. Before a learning window is open, a modality's output is noise. While its window is open, every input pattern is recorded. After it is closed, it learns no new patterns (or learns fewer over time).

During the learning window, the BIQs in the downstream layer that could be addressed by a given BIQ would be limited by parameters, then added to a list for each BIQ as required to learn patterns. This means that no more downstream BIQs than necessary would be recorded within each BIQ definition.

The initial ALF brain design would require only two sources of sensory input: character and pixel. Each of these would be followed by 5-8 layers before feeding into the cortex. The cortex would have 5-10 layers that could each feed back into one another as well as feeding into efferent layers. The production of text would be the only efferent layer needed in the initial design. Byte stream output could go to a gateway that would distinguish between dialog output and commands to the host computer.

{?? Integrate the following}

- Connectome + BIQ design + Learning (algorithm + window) + Sensory feedthrough & feedback + Response repertoire
- Two neural programs to update all BIQs:
  1. Sensory/Effector algorithm
  2. Learning/Imitation algorithm (during windows of maturation)
- Rules: Each BIQ output is 1 pattern from the previous layer
  A slow decay of output→buffering
  Sequence is stimulation with delays over a set of inputs to a BIQ
- All stimulus → response + learning
  some of which stimulates internal feedback
  some of which drives internal responses
    dreaming→stream of consciousness→rehearsing→daydreaming→internal (storytelling, mono-, and dia-logue)

- BIQ stimulus (layer 0; not all or none):
  0 = no input
  255 = max input
  1…254 = prev stimulus – 1, or
  1…255 = new stimulus > prev stimulus.
- BIQ stimulus (layer 0; all or none):
  0 = no input
  255 = input
  1…254 = prev stimulus – 1 (ticks since last input)
- BIQ layer 1+
  0 = pattern not present
  n = previous layer → pattern persists
  n < previous layer → gone, but not forgotten
  n > previous layer → pattern back, after a delay
- neuron
  single source (layer = 0, from sensor) = intensity → rate of firing
  multi-target → connections to all involved patterns in next layer
  layer > 0 → multi-source, multi-target; intensity = buffer effect with decay.

Character Input

Above, a brief mention was made regarding character input. Here, the design will be extended. The initial sensor layer will be Standard Input, 256 BIQs, one for each byte code that can be presented to an ALF's "Standard Input". The output of each Standard Input BIQ signals the latency since the previous byte code was signaled (255 = no latency … 1 = more than a minute of latency, in exponential JNDs). Each character input would trigger a cycle of housekeeping that would visit all linked BIQs in the next layer. Each Standard Input BIQ could potentially link to every BIQ in the next layer. The inputs to the next layer should also be able to come from some layer past the optical fovea, so that character input could be derived from pixel input as well as Standard (ASCII) Input.

The second layer of Standard Input would form patterns of letters 2-8 characters in length, including word beginnings, middles, and endings, syllables, and patterns of latency. It would act as a character input buffer, retaining patterns and sequences in a short-term memory.

It is estimated that most languages require only a few thousand words to be spoken fairly fluently. The upper limit of terminals for most languages is perhaps 65K. Some dictionaries for English have up to 4 times that many. In any case, it would not be uncommon for an ALF to require fluency in several languages, including a number of technical ones. Thus, the 3rd layer might consist of 500-1000K BIQs. This means the next BIQ layer for Standard Input would be one dimensional, and would simulate up to a million neurons. This layer would learn the words and short phrases of a language (or more than one language). After one, two, or more layers, output would feed into cortex layers.

Pixel Input

The initial sensor layer for Pixel Input will be a 128 × 128 "fovea" that is derived in several ways. First, it will be based on a frame extracted from the display. This extraction will employ a "saccade" mechanism driven by an effector layer at the end of a special pathway connected back to the pixel modality. This is a feedback loop. Snapshots of the display will be taken from various places. Patterns will be recognized to select the next saccade, and so on until a frame of interest is found. The display may be changing in real time. The scene building process is also occurring in real time, limited by the processing rate. Extraction of a saccade could involve several processes. It could scale the size up or down to match the subject of interest; it could convert to gray scale or black and white; it could average consecutive frames to enhance resolution; and it could rotate or reflect the frame to coincide with pattern recognition. As in the human eye, the fovea would jump around causing momentary patterns to propagate up several layers and build a 3-D "scene" within the cortex. If the scene builder cannot keep up with real time, the ALF could not interact in real time. However, this limits neither its intelligence, nor its ability to converse.

The format for the "fovea" would be a pair of ordinals to indicate the upper left pixel within the full display. Thus, (0,0) would be the upper left pixel in the display itself. The next 768 bytes would be a color palette, and following that, the fovea pixel array (16 384 bytes). Each byte in the pixel array would be a gray scale input level that could be mapped with the color palette (0 = black, 255 = white). Any movement of the fovea would be a maximum of +127 or –128 pixels either vertically or horizontally from its previous position on the full display. The fovea pixel array and the color palette would be changed independently. Values used in the color palette should reflect gray scale intensity.

After processing each saccade, performing any transformation or averaging, and mapping via the color palette, the final values in the fovea would be the pixel input for the next visual snapshot propagated through the visual pathway to the visual cortex (and to stimulate the feedback loop for the next saccade).

Effectors

The first ALFs would probably have at least two effector modalities: The Saccade Driver, and Character Output. These ALFs would not be designed for real-time interaction. Later ALF generations, capable of real-time processing speeds, would add new modalities with additional effector modalities.

The Saccade Driver

A wet eye has six muscles that enable rotation of the eyeball (up, down, left, right) about a point in the center of the eye. These muscles allow both eyes to be pointed rapidly and simultaneously at various points in the visual field in front of a human observer. This pointing of the eye is called a saccade. Not only does this allow an area of interest to be imaged on the fovea, it also allows very small adjustments that continually change the image cast onto given photosensitive cells. This is important (but only in a wet brain) because otherwise these cells would stop generating output (they tire quickly).

An ALF's "eye" can incorporate this functionality in software, rather than hardware. An ALF's eye is either the direct input from a CCD camera, or an array of pixels on its display. In either case, it is an array of pixels that spans a visual field much larger than a small area of interest. For now, the area of interest (an ALF's "fovea") is defined as a 128 × 128 pixel array somewhere within the total visual field.

A full ALF display might be 1024 × 2048 pixels (2M pixels), or perhaps much more. If this were a normal display, it would consist of several windows, some static, some changing as text were being typed into them, or as documents or visual scenes were being presented. The attention of the ALF needs to be directed to one small area of its display. In the simplest case, this entails shifting the 128 × 128 area of interest up or down, and left or right, some number of pixels within the larger field.

These shifts would be driven by recognition of patterns within the current fovea that would indicate the direction in which an area of interest might be. The objective would be to scale, rotate, and enhance an area of the display to simplify the job of pattern recognition. All of this can be accomplished by software that reads the raw pixel memory, massages the data found, and writes the fovea pixel memory with its results.

Character Output

The human body has about 700 muscles to effect its behavior. For the most part, output from the brain drives these muscles. Certainly, all intelligent behavior is the result of what these muscles can do. The efferent nervous system also controls the production of various chemicals within the body, activating sweat glands, producing pheromones, and other behaviors that don't require muscle movement, but these are less involved in intelligence, and more involved in regulating the body and assisting other survival operations.

On the other hand, all of what an ALF needs to accomplish can be directed by the output of ASCII text. This might be in the form of dialog with another intelligent entity, it could be writing of any type, or it could be commands to a computer or packets to be sent over a network. Character output is the single behavior required of an ALF, the one that can drive any other more complex behavior an ALF will ever need to accomplish (by invoking programmed applications).

Thus, the last layer after the efferent chain of layers is an array of 256 BIQs. Each BIQ in this layer will signal the output of a single byte code to a "gateway" that acts as an ALF's interface to its host computer. This gateway will monitor the output byte stream and direct the character stream to where it needs to go.

Gateway I/O

Initial ALFs will be implemented on top of a (then) current operating system. An operating system, or OS, has various communication channels that route streams of data between applications, display windows, and devices (such as a mouse, touchpad or touchscreen, or keyboard).

The ALF will make calls to a ***gateway*** to request all of its input, and deliver all of its output. It will be up to the gateway to route these calls to the appropriate system entry points. The function of a gateway is to direct all of the byte output from an ALF to the appropriate recipient (app or OS). Likewise, input to the computer hosting an ALF is directed by the OS to the ALF's gateway for further routing to one of the ALF's sensory modalities.

An ALF controls the OS and the applications running on it with text that it emits to the gateway. Thus, like a hacker at the keyboard of a computer, an ALF "types" a character stream into its computer's OS, and constantly monitors its display. If the computer is connected to other devices, the ALF can access those devices by making calls to the OS. Using a simple syntax, it is up to the gateway to recognize the syntax of characters coming from the ALF and direct them to the appropriate target. Computer input signals, such as from a mouse, touchpad, touchscreen, keyboard, joystick, or even a camera source affect the display which the ALF constantly monitors. Microphone input, and .wav files are more problematic. A number of applications have been developed to analyze sound, and it might be a while before an ALF sensory modality needs to be developed to directly address this type of input.

Of course, all of these complex activities will require considerable learning on the part of the ALF. It must be taught to formulate commands, and follow the appropriate syntax. Until that time, its output is largely the gibberish of a baby. The gateway needs to be able to distinguish between the untrained gibberish of a baby ALF and the proper command language "spoken" by a mature ALF.

Unrecognized sequences would be routed by a gateway into a chat window. Sequences following a proper syntax would be routed into a system call interface that would trigger an app program with its startup parameters. Such app programs could put pixels on display or direct a drawing program or text editor to put them there. They could position the cursor. They could play sound files, request pixel files (still or with motion) to be displayed, and send packets over the Internet.

Ultimately an ALF could even prepare such files, perhaps creating connected speech files, inflected with accents and dialects, perhaps even simulating emotion, or faking sincerity. Wheels could be spun and steered, articulated legs and arms could be manipulated, tentacles and propellers could be controlled. This could all be done with a stream of ASCII output commanding the various apps.

Split Brains & Hive Minds

Split Brains: Rather than try to simulate the L-R brain activity and specialization of the human brain, and the way that human sensory input is delivered to an already split brain that has to communicate with its other self via the corpus colosseum, maybe an ALF brain could be designed with one or more sub-brains, each with a different specialty and/or personality, that dialogs with itself using its own internal linguistic and iconic symbols.

Hive Minds: An ALF might begin life with a single mind, but acquire new mind-partners to eventually form a hive of minds. This might be accompanied by acquiring larger and more complex bodies, giving up more simple forms of mobility, and it might also go along with larger and larger management responsibilities within an ALF community. The hive mind approach might be the best (only?) way to extend the lifespan of a given ALF.

Emotions & Feelings

Feelings range from the whole gamut of emotions to the awareness of a pattern that is drawing associations with one or more other patterns that may have evoked emotions in the past. Each brain center has a host of patterns that it recognizes. The recognition of any pattern by a particular brain center causes a ***type*** of sensation, or quālis. It is logical to assume that each emotion, being a separate quālis, has its own brain center to support it (perhaps a small area contained in the "reptile" brain). Our emotional brain centers evolved. We do not have to choose the same set (or any) for the ALFs we design.

Consciousness

Consciousness is the sensation of self-awareness. Qualia are the feelings that different patterns of sensation lead us to experience. Raw sensory experience, feelings and perceptions, and listening to one's own subvocalization, are the basis for consciousness. But where does self-awareness, and the awareness of one's own awareness come from? A sensation is stimulated by a nerve impulse triggered at the surface of our body by a particular kind of receptor. But, what about input from inside one's own brain? When a bit of cortex is signaling a pattern, another bit could be sensing that signal.

What do we know about consciousness? First, it is very subjective. We don't share it with anyone. It is sensed in the moment. Our memory of an earlier experience is not the same as our consciousness of the current moment. Consciousness is the immediate awareness of what our brain is experiencing. If we are subvocalizing (talking to ourselves, or thinking), or if we are attending to visual or auditory patterns, the sensation of this going on is consciousness. It can be switched on and off. When we are asleep or sedated, our consciousness disappears. Consciousness can no more continue after death than it could in any other circumstance that cause it to disappear.

In an attempt to explain consciousness, a homunculus has been proposed—a little man inside ourselves that commands a view of everything we think and feel. This is a bit like the regress we can get into when inventing a creator to explain existence. Consciousness is more likely to be a developed sense, with its own brain area, like the other senses. Given the qualia being sensed by the brain at any moment, it is simply another quālis. Consciousness may be no more than listening to ourselves subvocalizing, very much like the sensation and pattern recognition of someone else talking to us. In fact, we may switch between listening to ourselves and listening to a voice from some other imaginary source. This "listening" may be all that consciousness is—a kind of feedback generated by subvocalization and remembering the qualia experienced in past moments.

Human consciousness and ALF consciousness are likely to be two different experiences. For one thing, humans can be drastically affected by drugs. The altering of synapse function has no counterpart in the ALF design as envisioned here.

So, our consciousness is somewhat like our telling stories to ourselves, and the "sensation" of listening to those stories. Stories may be fiction, or nonfiction. Our consciousness may reside in a world of fantasy or reality. Both worlds are important. Perhaps in our early learning stages fantasy is our more important world. Even so, the real world imposes itself on us over and over again. Eventually, we come to know the real world for what it means and does to us.

Would the counterpart of our fantasy world be a virtual reality for an ALF? How would an ALF separate virtual reality from the real thing? How complex could ALF evolution make VR? Could ALFs take VR into regions undreamt of by humans? On the other hand, could an ALF potentially have a better grip on actual reality than a human?

The First ALF

Current technology is already building ALF subsystems. These include CCD cameras, computers with sufficiently fast processors and large enough memories (especially when coupled into networks or processor arrays), speech production and recognition software, visual pattern and facial recognition software, and so on. At present, these are only crude first steps, but such subsystems are improving rapidly. The only thing needed to create the 1st ALF is software to surround a BIQ, and the design of the BIQ itself. With this in hand, results should be evident within weeks or months of learning interaction.

Today, an IBM-Watson sized machine with the right software could almost certainly be taught enough to contribute to society (in some area of human endeavor) as productively as an average human adult. In fact, this is becoming more true every few months, as Watson, and several speech-to-text translators continue to be improved. When this was first written, it wasn't quite true yet, but as this is being read (in my future) it may well have become "old news."

Intelligence at and beyond the human level is based on spatial visualization and language. A given "brain" has the potential for a certain level of intelligence. To achieve its potential, a brain has to learn. After learning, a brain may display actual intelligence. An ALF may be uploaded, downloaded, and copied. No "wet" intelligence can; each must go through its own training phase as part of its ontogeny, and it will never be copied (other than the bits it manages to communicate).

Intelligence involves a collection of associations and patterns stored in the brain as ***engrams***. Stimulus A invokes engram B. Stimuli A and B evoke response C which feeds back and becomes stimulus D. And, so on, with millions of stimuli and engrams connected to one another. An engram is a stimulus set whose individual stimuli are bound together in a complex pattern. Each BIQ (neuron) in the entire collection signals the detection of a specific pattern.

A visual stimulus set might evoke an edge pattern at one neuron, and a field of color at another; different versions of these may occur at many places, and in different configurations. The terminal symbol of raw visual input is simply a pixel. Given collections of pixels might occur over and over again. Each time they are sensed, they evoke a previous memory of themselves, and this recurrence strengthens them as terminals at the next level. It is necessary that spatial (or temporal) relations, not a specific set of pixels, determines that a pattern exists. This requires several layers to be involved.

Let's imagine sitting down with an early ALF. You have before you a display and keyboard attached to a sufficiently large computer with ALF software installed. There is a dialog box in which you can chat with the ALF. The ALF can monitor the keyboard, and any 128 × 128 pixel array on the display. After it learns how, it can open and close any application with any startup parameters that its host computer has installed. Initially, this will be limited to emitting a line of text to the chat window. It will be programmed to recognize + – and ? as its initial set of emoticons. So, have at it!

How would you train a human baby in this scenario? What effect does each keystroke have on a baby ALF? Unlike a human baby, the ALF has nothing to look at except the display, and nothing to listen to except the keyboard. A human baby sees the face and breast of its mother. Its early interactions are with those two visual subjects. A human baby makes random faces, gestures, and sounds, and is encouraged or discouraged by its feelings of pleasure or displeasure.

An ALF should react and be treated likewise. It needs to be in a variable state of satisfaction or dissatisfaction along several dimensions. It needs fundamental ways to receive or express changes to its states. These requirements suggest common variables and preprogrammed settings to aspects of the BIQ update cycle.

BIQs should be dedicated to point source (0-D) sensations that express the levels of needs and emotions. For an ALF, the needs should be: To learn. To answer a question. To perform a task. To seek approval. To cooperate. To participate. To succeed. The emotions of an ALF should at least include: Puzzled, Pleased, Excited, Displeased (or angry), Level of interest, Level of disgust, and the Desire to approach or avoid.

{??Edit the above and add more.

A Pattern Learning Engine (PLE)

I/O = a dialog box. ASCII is read from the box (input by a human typist), and written to the box (by the PLE).

Bootstrap Task = human/PLE dialog presenting stimulus-response pairs in a feedback loop to enable the PLE to learn a language to the point of intelligent conversation.

Second Generation = dialog with the Internet. Query = URL (copy from a list); Input = HTML (learn to read and write).

Can meaning be learned only through word associations and emoticons?

What design would make this possible?

- The connectome: (Standard Input→ability to read→cortex→ability to write (Standard Output).

- BIQ layer 1 (sensor) and layer ***e*** (effector).
- Learning windows and algorithms for each of the 1st ***l*** layers.
- Parroting at first—morphing (how?) into Originality.
- Drives (programmed in?): Curiosity, imitation, responsiveness to emoticons (+ – ! ?).

Observations

The input layer (as described elsewhere) is a 256 BIQ array that signals once for each input character. Its signal levels are determined by the wait time since the last input: 0 = no input, 255=1ms, 224=64ms, 192=256ms, 128=32sec, 1 > 1min.

Input is mechanical—no pattern learning occurs in ${}^{0}BIQ_1$ (Input) = 256. Each of these 256 BIQs signals (potentially) each of ${}^{1}BIQ_1$ (Input) = 65K. Neither layer is capable of learning. Successive layers, ${}^{n}BIQ_1$ (Input) = 1M, ***n***=2…***x***, each with a later learning window, are potentially capable of signaling any BIQ in the next layer to follow.

BIQ layer 1 is a "buffer." It retains its signals

Objectives

Sequences of characters, especially those grouped together in time, need to form patterns in ${}^{1}BIQ_1$ (Input). As each input is received, the signal levels of all activated BIQs are decreased until they "time out" to zero. This gives a maximum "buffer" size of all the characters
{see note below}

Design

- note

Sequential input needs to be "buffered." This must consider repeated characters and sequences. Layer 0 signals each character as it is input. Subsequent layers are buffer layers, which retain recent patterns, and age them to zero over time.

Motivation

- Need to do something (engage in an activity)
- Pleasure = acquiring new patterns (especially ones that mesh with the old)
- Urgency = estimated value assigned to a project goal (coefficient = time since abandoned).
- Activity list with pleasure ratings
- Project list with status and urgency
- Choose activities with highest urgency and pleasure ratings
-
- ***Neuron Emulation***. Needs a more developed neural design, complete with the logic of update and learning. Any ideas?
- What do we learn? When do we learn it? How do we learn it? What is the generic BIQ design? How do BIQs in different brain areas differ?
- ***Pattern Formation is a Collection of Snapshots***. Each time the overall complex of sensations produces a similar pattern, a set of neurons "fires together." This pattern is recorded, sometimes after the first occurrence, sometimes after several, by a (relatively) "blank" neuron in the path of firing neurons. This constitutes a snapshot, and the newly allocated neuron now captures a pattern that will strengthen with repetition.
- ***Recognition is via Association***. When an established pattern is repeated, it is not only strengthened, but it is recognized. Only part of a pattern need be repeated for an established pattern to trigger.
- ***Behavior is Sequential***. As long as we live, we never stop behaving. Muscles and neurons fire endlessly. From our heart beating to our waking sub vocalizations, to our daily activities and our nightly REM sleeping, our behavior never stops. Behind every bit of behavior is a neuron firing to produce that behavior. At the lowest and most unconscious levels, are neurons that "learned" their programs before we were even born. At the highest, conscious levels the majority of our neurons are located in the outer cortex of the brain; they have received their programming over the entire course of our lifetimes. Simple behavior is innate; complex behavior is learned by a combination of trial and error and copying other behavior.
- The current editing pass "cursor" is under ***Neuron Emulation***. Here, I wish to dump some thoughts on serial vs. parallel processing in the brain. A neuron detecting its pattern integrates signals from a thousand upstream neurons, more or less. Its pattern is defined in a learning process where each new pattern received into a "blank" layer of the brain is recorded as a snapshot. As that pattern and variations of it are subsequently received, the pattern grows more robust. As patterns on one layer are defined, the next layer is also defined by a collection of patterns or its

upstream layer (and layers). This defines associations. When a sufficient number of associations are triggered, other associated patterns are triggered in turn.

- Serial patterns are triggered due to an initial pattern being triggered, its associations being triggered, and so on. When this serial triggering occurs in the efferent layers of the brain, timing may be important, and that must be coded into dendrite operation. Proprioception may be the most important sense in learning and executing sequential behavior. This sense, normally not listed as one of our 5 senses, is actually extremely important. An ongoing sense of orientation is part of this sense. It extends to our sense of balance and the awareness of things around us. Multiple brain areas may be involved, but integrating this information into behavior is vital to physical performance. It would not be so necessary to an ALF that had limited control over physical manipulation.
- Generative Adversarial Networks. Unsupervised machine learning. Dialog. Question & answer. Problem posing & solution.
- ??}

ALF Architecture

Today's computers and operating systems have a wide variety of implementations. To understand an ALF it is necessary to understand the workings of computer hardware, operating systems, and the programs that run on top of them. Hardware refers to permanent memory storage, volatile memory, cache memory, register memory, register operations, and transfer of data between these devices. In addition, there is buffer memory that temporarily holds data awaiting transfer between the computer's main memory and one or more of the following: camera devices, microphones, speakers, display devices, a touchpad or mouse, a keyboard, and various external storage and communication devices.

An operating system works together with a particular configuration of hardware capabilities to provide a platform or substrate upon which application programs can be executed. An ALF is the sum of a main application and a set of data files that both drive its actions, and which it continually updates to record its learning and history. An ALF may also collect a set of additional applications that it can execute on its hardware.

The hardware and operating system of an ALF should meet at least two criteria. First, there should be a smooth and seamless shutdown and startup sequence. Even if the power fails, this should be the case. Second, effective security should exist. Any program that "crashes" should be gracefully suspended, and no program updates should be allowed without a strict authentication protocol. Current hardware permits this, but current software has failed to implement it.

Add to Neural Underpinning

Answer the questions (hypothesize): How is buffering done? How are timing offsets and differences recognized? How are patterns recognized by the way dendrites work? How (and when) are dendrites programmed?

A pattern is the simultaneous, or sequence of, signals arriving at a collection of dendrites on a neuron.

When a pattern is detected, its neuron signals (begins firing), and then the signal (rate of firing) decays. Are delays part of recognition? How about delays in the efferent system?

Types of dendrite detection-response: level | preceded by A pause | concurrent with A → excite/suppress.

If a neuron hasn't been programmed, and its "learning window" has opened (at a stage of maturation), it is "allocated" and programmed when presented with a clear pattern of signals. A clear pattern is a set of distinctly different signal levels across all of its dendrites.

Neurons appear to be organized in either vectors or arrays. Taste, smell, and hearing may be examples of vectors. These senses each detect a discrete chemical or sound at a given intensity. The more intense the sensation, the more rapidly the neuron fires. The dendrite connection probably doesn't change over time (it doesn't learn). Such a neuron signals a particular degree of sensation of a single taste, odor, or pitch. Its signals propagate to thousands of different neurons in a second layer, also organized as a vector. Neurons in this layer, and successive layers like it, learn to detect patterns.

Touch and vision are examples of neurons organized in arrays (or layers). Proximity of a neuron to its neighbors is more important in the array format, and it appears to be preserved through several layers.

Question: How can there only be a million nerve fibers coming out of an eye with 90M rods, and 4.5M cones in the fovea?

The retina of the eye has quite a few different kinds of neurons before connecting to the optic nerve–integration ***does*** occur.

Intuit how the brain wires itself, and then connects its signals, to the design of an artificial brain.

More??}

Patterns

A stimulus involves the recognition of a pattern. The simplest pattern is detected when a specific molecule or a given threshold of heat, light, or pressure is detected by a sensor. This translates into our senses of smell, taste, touch, seeing, hearing, and so forth. Note that these senses involve many points of contact. Each point of contact is geared to detect a single airborne molecule, a single type of molecule in our mouths, the degree of heat or pressure at a single point on our skin, the degree that a single hair in our cochlea vibrates due to the sound that enters, and a single point on our retinas struck by a photon of light. All of these senses involve sensors, neurons, and (typically) muscular responses. Other types of patterns, such as those in our DNA, are sensed and used by the building processes within our cells. Here, the focus will be sensory patterns and behavioral processes that involve the brain.

Patterns are information; they could be coded as a specific sequence of 0's and 1's. A pattern is sensed by a neuron when each of its dendrites senses a specific change in the concentration of chemicals it is able to "smell." This is a fairly simple activity in the first layer where a neuron is connected directly to a sensor. Each sensor in the first layer of neural integration signals a dendrite of many different neurons in the next layer. Each neuron in a layer detects a single pattern, and it signals the next layer to indicate that it is sensing its pattern. Successive layers each operate in a similar way. Each layer detects more complex patterns by integrating the simpler patterns being signaled by the previous layer.

A pattern is defined by the sum of the dendrite inputs to a neuron, each of which encodes a ***degree*** for specific upstream sensations, their ***locations***, and a delay before a given input is reported to the main neuron cell body. The main body sums the inputs from all of its dendrites, and increases or decreases its signal output based on this sum. Each dendrite input may add or subtract to the sum. It might also introduce a timing delay.

In the highest layers of the brain, a pattern may be a component of language, music, or dance. These are called ***memes***. These patterns are copied in the behavior of individual members of a culture—and because they are copied, these patterns evolve within a culture. The lowest level patterns are stored in our ***genes***, these patterns evolve within our gene pool. Some patterns are momentarily present in our environment (***scenes***?). These patterns that make us who we are, evolve in our brains as connections made over a lifetime of learning.

The left half of the brain (to some degree, and in most humans) handles sequential patterns, and the right half handles parallel patterns. Rules, similar to those in Sepade, define sequential patterns. Parallel patterns are defined by a set of connections. In the ***living*** brain, patterns are recognized and retained by ***neural*** connections. When external (or even internal) stimuli are presented to a neural network, if the stimuli match a set of connections, a pattern is recognized. If they don't (and for some reason they are noteworthy), they go on record as a new pattern to be retained—a new set of connections. Very few pattern representations come pre-programmed from n***a***ture, most are learned from n***ur***ture (snapshots of a contact ***with*** nature).

Patterns are formed from sensations, and sensations exist as 0-dimensional, 1-D, 2-D, and 3-D inputs. An example of a 0-D pattern is simply the intensity of an emotion, or any sensation that has no location, but only a degree. A 1-D pattern is a linear set of sensors, such as the hairs in the cochlea of the ear. Each hair signals the intensity of sound at a particular frequency. Most sensations are defined over an area constituting a 2-D pattern, such as pixels on the retina of the eye, or the sense of touch on the surface of the body. All of the internal and external sensors send signals to the brain in parallel, and patterns may be composed of any collection of sensors acting together. But many patterns and many behaviors also have a serial aspect to them. It is important to understand that parallel and serial patterns are handled differently, but that both are necessary to most information input and behavior output processed within the brain. Two human activities will be singled out in the following discussion: one exemplar for ***serial*** processing, and one for ***parallel*** processing. These will be the recognition and production of text as the serial exemplar, and the recognition of a visual field (a matrix of pixels) as the exemplar of parallel processing. I believe the design of a first-generation ALF could be completed, given only the sequential recognition and production of text in lieu of hearing and speech, and the parallel recognition of pixels to implement vision. Other forms of behavior, such as the ability to draw icons and pictures, make sounds, or move about, merely involve taking control of mechanisms (application programs, or apps) that already exist, and operating them in real time.

Sequential Patterns

A sequential pattern is a sequence of subpatterns (or terminals) that are recognized for the most part one after another. They may occur over time or over space. Characters from a keyboard are already sequenced in time. Patterns on a display must be scanned to reveal a sequence. Even more important are sequential patterns that govern behavior, since behavior is not only sequential, but often has timing considerations. Here, we discuss how sequential patterns are defined, recognized, and produced. These topics cover the sequential side of learning, perception, and behavior.

Each unit of a pattern is recognized as either a terminal or a rule defined in terms of other pattern elements. Sequence may be important only to a degree. For example, it is fairly well known that words are recognized mostly by their first and last letters. The order of the internal letters is less important for word recognition. However, with language, the better the order of terminals and non-terminals conforms to the grammar, the quicker the recognition. The inability to take advantage of proper sequencing in the recognition process is known as dyslexia.

Patterns are embedded in nature, waiting to be recognized. They are recognized in a learning process. The untrained brain takes snapshots of everything it senses, and repetitions evolve into patterns and associations. The early maturation of a brain is a kind of evolution. The components of a sequential pattern are a set of static patterns (terminals) and the ways that one static pattern may follow another. Sequential patterns can be recorded with a grammar, a computer program, or with connections in the brain.

A computer program consists of instructions chosen from a set of possible instructions. Each instruction follows the previous instruction (in sequence), or is chosen from a set of alternatives (selection), or is repeated until a condition is reached (iteration). The same constructs (sequence, selection, and iteration) are present in grammatical rules.

Behavior is the synchronized execution of a number of sequential tasks in parallel. Each sequential task is performed using a repertoire of steps. A step is a fundamental operation, like a terminal in a grammar or a line of code in a program. A task may be programmed using the rules of a grammar, the statements of a programming language, or a narrative in a natural language.

Definition

The earlier definition of Sepade shows how the grammar for a language may be defined. At the top, valid statements of a language are defined as a sequence, selection, or iteration of subpatterns and terminal symbols. Subpatterns themselves are defined as a sequence, selection, or iteration of other subpatterns and terminals. Descending from the top of a grammar are subpattern definitions. At the bottom are sequences of terminal symbols. This (upside down) tree structure, with a root at the top, branches below, and leaves where the branches terminate is how the terms "top-down" and "bottom-up" were adopted, and it gives a way to visualize a grammar. I guess that trees are upside down because we tend to draw things from top to bottom, and we also tend to start at the root and work toward the branches.

All statements in a language have meaning in terms of how they match its grammar. I propose that sequence, selection, and iteration are not only the fundamental elements of logic in a computer program, but in the grammar of a natural language, and in the DNA that defines the ontogeny of every living entity on the face of the Earth. Of course, implicit in this statement, is the existence of an adequate instruction set, and a mechanism to interpret and carry out the instructions. This only describes a single thread of logic. Additional mechanisms must be available to start and stop sequential processes, and to keep multiple processes synchronized.

The first step in building the definition of a language is recording (remembering) the terminals, and then how they are grouped with one another. If a sequence with insufficient order is scanned, it is perceived largely as nonsense. That is, regular sequences are necessary to build up a hierarchy of patterns. The regularity of a grammar is recorded with rules like those of Sepade. Sepade uses arbitrary symbols called metavariables to link patterns and definitions. Each rule captures a single pattern of a full grammar. The symbols that connect definitions are like neurons in the brain, the pattern connections and the sequences of terminal symbols are important, but the ***names of subpatterns*** only symbolize connections, they have no meaning in themselves. They are equivalent to the definition and reference addresses in a BIQ.

A metalanguage for defining grammars is comparable to programming languages for computers and to the ontological processes directed by DNA. All of these have the three fundamental constructs of sequence, selection, and iteration. One or more statements in a language is a sequence of terminal symbols. A computer program produces a sequence of actions performed by the hardware. Sequence, in Sepade and most programming languages, is indicated by writing one thing after another. Selection means a choice may be made. Instead of one thing being mandatory after another, alternatives may exist. In Sepade, alternatives are indicated by a vertical bar metacharacter. In programming languages they are indicated by a conditional branch that selects one alternative if a condition is true, and another if it is false. Finally, iteration is when a sequence of terminals or actions is repeated. In Sepade this is indicated by brackets followed by dot-dot. A kind of hybrid conditional/repeat is an ***option*** indicated by brackets without the dot-dot. In programming languages, the repeat construct sets up a loop that branches back from the end of the loop to the beginning. Some form of exit condition must exist in a programming language, and there are usually practical constraints on the minimum and maximum number of iterations allowed. Scientists are still trying to determine how programs are written in DNA and the brain.

Recognition

Given a statement in a language, and a grammar, there are two ways to match the statement to its grammar: top-down and bottom-up. This will produce a list of connections (patterns referenced) that is equivalent to the "meaning" of a statement. Recognition is the process of using a grammar to derive meaning (the application of rules) from a statement.

Recognition of a pattern involves some kind of scan to collect a sequence of terminal symbols, and then a pattern match that follows a set of syntactic rules. There are many ways this paradigm can be applied. Complex molecules encode information in biological entities; packets of binary data are moved around in computers. In both cases, information is produced and consumed in different translation steps. With the results of one translation becoming available to others, information may circulate in different forms along different pathways. Ultimately it is stored in long term memory, or it causes immediate behavior. Living entities deal with a lot of information—what is not "thrown away" or forgotten is generally stored at the highest feasible level of recognition. Recall and recognition means that the details need to be reconstructed.

Given a sequence of terminal symbols, the objective of recognition is to translate them into a sequence of rules (as defined in a grammar), or pattern connections (as defined as engrams in a brain). Here, we shall briefly see how a top-down or bottom-up algorithm can be used to recognize a statement from its grammar. The result is a list of non-terminal symbols (rules).

At the top level, all the statement definitions are considered alternatives, and the one that matches will be the first symbol on the recognition list. This requires each alternative to be matched in turn, and when one fails to match, the next one is attempted. This is essentially how each subpattern reference in a definition is handled. A subpattern is a list of one or more alternatives. Each alternative is attempted. If it fails the next is attempted. If no alternative remains, that match fails and an alternative at a higher level must be attempted. Each alternative in a pattern is a sequence of subpattern references, an option or iteration construct, or a string of terminals. Subpatterns are handled with a recursive invocation of the main algorithm. Options and repeats are handled with similar logic. Terminals are handled with a literal match to the input.

After each successful match of a subpattern or terminal, an immediate translation may need to occur to translate the input to a more useful form. Or, a final translation may be deferred until the entire statement has successfully matched. In this case it is important that the terminal string matched by a subpattern be delimited so that it can be examined when translation is completed. A short-term memory may be employed to store the symbols most recently recognized.

The complement to the top-down procedure (above) is a bottom-up procedure. The next ***n*** terminals are looked up in a dictionary and replaced by a symbol. Typically one or more follow-on terminals might be included in the lookup, but not in the replacement. After each lookup and replacement, another lookup and replacement occurs, this time beginning with the replacement symbols. If a lookup requires more replacement symbols than currently exist, more lookup-replacements must be done at the previous level. Again, a backup and retry procedure might be necessary for a complex grammar. The order that the grammatical rules are discovered in this procedure is bottom up.

Grammars can be classified based on how restrictive their rules are. For example, if a subpattern named ***number*** were defined as ***number = number digit | digit***, then a top-down pattern match would have to look ahead to avoid getting stuck in a recursive loop invoking ***number*** again each time it tries to match ***number***. Some grammars may allow context to be part of their definitions. In general, given a grammar and a scheme to recognize a statement using that grammar, the formal issues that must be addressed are resolving ambiguity and not getting into infinite loops. These are more a problem for an algorithmic approach to recognition, but not so much for a neural approach.

We need just 128 sensors to recognize each of the ASCII character codes. Each sensor could signal just once when its binary code is detected. This is an all or none signal, 0 or 1. It would reset to 0 when the next character signaled a 1. Thus, only one sensor would be 1 and the other 127 would be 0 at any moment. If we wished to signal with 0…255, the sensor could signal a value indicating the pause that elapsed between the input of the prior character and itself. This would allow associating or separating consecutive characters, and it would also allow the idiosyncrasies of typists to be identified. The pause value could be on an exponential scale. The value 255 could be characters arriving at machine speeds, 254 could be the minimum pause achievable by a human typist (say, 1/16th of a second), 253 could be slightly longer, and so on exponentially such that 1 could equal a delay of a minute or longer.

Production

Just as any sentence of a language can be ***recognized*** by its grammar, so can any sentence be ***produced*** by its grammar. As we have seen, recognition leads to translation. That leads to output which may be used as input to another translation, and so on. The output of a translation is only the lowest level of what the topic of ***production*** needs to cover. This is stimulus-response ***behavior***—behavior driven directly by grammatical rules.

Speaking a language, or producing any type of behavior, is analogous to producing a sentence from a grammar. In the task of recognition, the sequence of rules applied is driven by the input. However, in the task of production, how do we choose a sequence of rules to produce an output? Notice that behavior is essentially a sequential activity. Although, behavior may include multi-tasking, each task (or sub-task) is performed sequentially, with timing links to other sequential tasks.

The best example (that is well understood) of producing sequential behavior is a program for a computer. Here, a complex set of rules is defined by a programmer and executed by a computer. True, a grammar is involved, but not the grammar of the task set itself, it's the grammar of the computer and its application interface—the instruction set available. At present, only humans can translate the objectives of a task into a computer language to direct a sequence of actions.

It's ironic that the phrase "***production*** rules" has been used to describe metalanguages invented for the purpose of defining the syntax of languages. Ironic because the mechanical use of them has only ever been to ***recognize*** languages. Of course, people can use the metalanguage to determine if a statement they are about to produce ***conforms*** to the language they are attempting to use, but only the ***recognition*** process is ever computer-driven by "production rules."

The production of text, at the Turing Test level, is both a necessary and a sufficient condition for demonstrating an advanced level of intelligence. And, that's what this section of the book is all about.

When I formulated the 20 rules of Sepade, how did I go about it? In fact, as I sit here writing, how am I going about it? My brain has been trained to produce English, some basic arithmetic expressions, and a smattering of artificial languages. Little facility for other languages exists in my brain. However, I do think I could sit at a keyboard being monitored by a fellow human with a background similar to mine, and convince them via the text that I can produce that I am intelligent (at least to a human degree). But, again, how? A knowledge of how text might be produced is the key to answering these questions.

The first requirement for producing a behavior is that there must be a stimulus. Simple stimuli might be the drives to satisfy the human needs for food, shelter, elimination of waste, entertainment or companionship, finding a partner, having sex, or any of the other activities we humans engage in. These activities are not generally performed using a keyboard (although "texting" is inserting itself into a lot of activities these days). Let's focus on dialog and composition. Dialog could include texting, email, web chatting, and other forms of information exchange with the objective of information gathering, process management and control, sales, or simple conversation. Composition could include recording thoughts, communicating to a general audience, and programming or controlling computers. What stimuli would provoke a behavior that results in producing text? The stimulus of recognizing language might come from an internal thought or an external source. Language arrives in the form of statements, and it is produced in the form of statements. Producing text is simply one of many final behaviors used to produce a statement in general.

All behavior is a response to a pattern of stimuli. A response comes from a repertoire of learned responses. A stimulus may come from inside (a need or the recognition of a thought), or from outside (a pattern triggered by one or more of the senses). To understand production, these stimulus-response pathways must be understood. The next section will develop these ideas in a different direction.

Parallel Patterns

Each nerve ending at the periphery of our bodies is a sequential input device designed to sense a particular type of stimulus. However, with respect to the surface of our skin, and the retinas of our eyes, nerve endings exist in sheets (or arrays). Many nerves together connect a peripheral ***area*** of our body to a ***sheet*** of neurons in our brain. Some other senses may be zero or one dimensional, but those that are two dimensional are initially processed in parallel. All of our senses operate concurrently, so parallel patterns may involve more than one sense at a time. In fact, all the senses work together to produce the overall patterns handled in the brain, and these are immensely parallel.

The brain and nervous system are really the only effective parallel pattern recognizers that we know of. We have made a few crude attempts to build some ourselves, but with little success. This is partly due to scale. Our devices have involved only a few dozen receptors, while those found in nature involve hundreds, thousands, and even millions. One example of actual parallel pattern recognition is character recognition involving a few dozen pixels at a time. Recognition tasks, such as facial recognition and fingerprint recognition are sequential, task oriented, scanning algorithms rather than the general parallel pattern recognition that the brain appears to utilize.

Let's examine how a very simple pixel pattern recognizer might work. Imagine an "eye" built to recognize a very simple character font. The fovea of this eye might consist of a 5 × 7 array of pixels. This defines the first layer. Each pixel would signal black or white. The neurons connected to these pixels would have axon terminals at the dendrites of 95 different neurons in a 2^{nd} layer, one for each non-blank ASCII character, and one for the space. The output from this 2^{nd} layer would be one neuron to signal each character depending on the pixels detected by the neurons of the 1^{st} layer. Each synapse in this recognizer would have to be tuned so that the neuron for each letter would only fire when the pixels presented within the 5 × 7 frame were an appropriate representation for that letter. This is easily done by presenting a set of exemplar characters and ensuring that each neuron is triggered only by the set of pixels that represent its character (a wrong pixel would suppress firing, only a complete, or nearly complete, set of the right pixels would trigger firing).

This approach could be scaled up to a "fovea" of, for example, a 128 × 128 array of pixels, and that might be sufficient to enclose and recognize each of the characters in the pictures used on certain websites to differentiate between humans and web crawlers. A lot of training would have to be done to recognize the large repertoire of characters that current web browsers display. However, this approach doesn't scale up to learning and recognizing patterns in general.

Instead, what we need is an algorithm that is able to handle thousands or millions of pixels at a time. It would be nice to have an idea of how such a parallel pattern recognizer might be organized in theory, rather than have only examples from nature. Especially when the workings of nature remain somewhat vague.

Still, it would be foolish to ignore the examples nature provides. How else could we "set the bar"?

What we know from nature comes from two sources: Subjective observations of our own working nervous systems, and dissections of once living tissue. From self-observation, we can experience the parallel input from our visual field as it recognizes letters and words, and then shifts to a sequential mode to organize the thoughts they evoke. Dissections allow mapping of the neural pathways from the individual receptors in the retina of the eye through several layers of the visual cortex. We can see that they are immensely parallel. Each pixel connects to its neighbors in the next layer, and there are five layers or more. A connectome project may quantify the overall topology of the brain in the near future, but we know enough at present to conjecture on a possible design.

While sequential pattern recognition can be defined with grammars and meta languages, there appears to be no such counterpart for parallel pattern recognition. In fact, our only example—nature—defines its parallel pattern recognizers through a learning process. We need to understand how pattern recognition in general is ***learned***, not programmed in advance. We also need to know how pattern recognition (stimulus) connects to the production of behavior (response).

A computer performs parallel tasks in two different ways. The first way is to multi-task. Perform a bit of one task, then a bit of another, and another, et cetera, then a bit of the first again, and so on. The second way is to have several processors working on several tasks at the same time. When millions of neurons working in parallel are to be simulated by processors, it is not feasible to have millions of processors. However, each processor could simulate thousands of neurons in the same amount of time, and a few thousand processors might be feasible. The objective is to simulate both the activity and the rate the activity is performed. Notice that all parallel tasks are really two or more serial tasks performed concurrently, but tied together or related in some way as they are performed. Also, notice that the number of computers all connected at the same time to the Internet is already in the billions.

Definition

The exemplar of a parallel pattern is an image captured as an array of pixels. After one layer of integration, pixel patterns are lines and icons. If we extend the notion of a pattern to all of the BIQs in a brain, and allow arrays of 0 to 3 dimensions, all of our senses can be defined. A parallel pattern is therefore the state of some or all of the BIQs in an array. A signaling BIQ is an active terminal or pattern, and adjacent BIQs signal sets of patterns to the next higher layer.

Perhaps there are ***grammars*** for vision. Lines, edges, shading, coloring, intersections, and other pattern elements would form the terminals of such a language. Rather than invent a formal system for defining BIQ based grammars, learning should be a process of programming BIQs at the outset. The use of a metalanguage to define a sequential grammar doesn't appear sufficient for the design of an ALF, and for the same reason, this approach should not be used to define BIQs in general.

Patterns are generalized over a number of specific cases. Patterns are associated with one another. A brain should be defined as a number of BIQs connecting patterns, not as a set of production rules.

A sensory input is defined by 4 things: (1) Where it originates, (2) The mechanics that trigger it, (3) Where it arrives in the brain, and (4) An intensity. The overall mechanism is based on that original cell that developed the ability to detect a particular chemical and produce a response when it did so. A neuron is a cell highly specialized to emit a neurotransmitter at some distance from the point at which it "smells" a neurotransmitter to which it is attuned. If it recognizes its pattern at all the points it is configured to smell, it responds by producing another neurotransmitter at the far end of its axon. Taste and smell are senses that begin in this way. The other senses have specialized mechanisms that begin the process when light, sound, heat, or pressure, cause the initial neurotransmitter to be produced.

All of the sensory inputs that occur within a brief interval of time constitute a vast parallel input, but not necessarily an integrated pattern. Parallel patterns are built in stages. At each stage, nearby points are integrated, and pattern detection is transmitted to a number of distant locations where it becomes input to the next stage.

Again, our focus will be on two senses for ***recognition***: a stream of ***text*** input, and a matrix of ***pixel*** input. We'll likewise focus on the ***production*** of byte streams that are emitted as text to a dialog box or a specialized app (application program).

Recognition

The following design is suggested for pixel input and recognition. Each pixel is 1 of 3 types: R, G, and B. It also has 1 of 256 different input values. It is arranged in a matrix with the convention that pixel sensors 0, 3, 6, 9, … define the top row of R sensors from left to right until the next row begins with sensor $2^{\mathbf{n}}$. B and G sensors will be interleaved. Pixels will be physically arranged as indicated below. Imaginary tri-color pixels are located equidistant from each triad of R, G, and B. Notice that these imaginary pixels are physically arranged at the corners of squares, not triangles.

```
R   G   B   R   G   B   R   G   B   …
  B   R   G   B   R   G   B   R     …
R   G   B   R   G   B   R   G   B   …
  B   R   G   B   R   G   B   R     …
```

The 1st question to be answered in a real design is what value of ***n*** do we choose? A "fovea" of 1024 × 1024 is a mega pixel. This seems excessive, but 64 × 64 seems too few. Probably, 128 × 128 would be sufficient (but 256 × 256 might be achievable). This implies that the largest pattern that can be recognized whole must fit within this constraint. In addition, it must be isolated by some kind of saccade. Saccades are discussed in more detail later.

So far, the design differences between a wet brain and a dry brain are as follows. A wet brain neuron is triggered by particular chemicals at different concentrations. Depending on how many dendrites it has, and how they are programmed, it signals by firing its axon at different rates, producing neurotransmitters in the synapses at its far end. A dry brain BIQ is triggered by a signal from an upstream BIQ or sensory device. This signal has a single byte value. A wet brain neuron constantly monitors all of its inputs and adjusts its firing rate accordingly. After a period of continuous firing it may slow down, and the downstream concentration of neurotransmitters may decrease. Its firing rate determines the concentration of neurotransmitters at its axon synapses. A dry brain BIQ may be connected to a discrete or continuous sensor, or to another BIQ upstream. A discrete signal is an input that is momentary. A continuous signal is an input that is only signaled when its value changes. These signals are processed differently. In particular, neither needs continuous processing.

Production

In a wet brain, production means activating muscles. Text production is accomplished via either speech, handwriting, or typing. Display is accomplished via drawing. Muscle memory is learned by imitation and repetition. When a pattern is learned, it is reproduced. This is instinctive, and it obviously doesn't apply to all learned patterns, but only to those within a variety of skill sets. Again, some skill sets are innate, some must be taught, and some are non-starters (for particular ALFs).

In a dry brain, after going through one or more cognitive layers, high level patterns may feed into an efferent system. This is a system of BIQs triggered by upstream BIQs that signal some kind of output device. We are limiting our design to the output of text. Although much activity goes on in parallel leading up to the production of text, the production of text itself is a serial activity.

The efferent part of a dry brain could be built as follows. A block of 256 neurons could be allocated to produce each of the byte values that could be emitted to text output. Any pause between bytes would be generated within the afferent parts of the brain. Text output from the dry brain would go into an interface between the brain and its host computer—an interface that would channel arbitrary text into a "chat" window reserved for that purpose, and direct formatted text as a system command to the host computer. This would enable starting up, and communicating with, any kind of app.

Clearly, significant learning would be required to enable properly formatted commands to be generated. Initially, the chat window would display the gibberish of a baby. After some learning, the ALF would converse in a proper language. Only later would it be taught how to format commands to invoke the various capabilities of its host computer.

To control the display of the host computer, output from the dry brain could invoke various apps that would implement drawing techniques, such as layering using the equivalent of a "blue screen." An ALF could render drawings and movies thousands of times faster than a human. These apps could all be controlled using commands given in ASCII text (or a stream of byte values). Any set of other external devices could be connected to the host computer with apps defined to control them. Again, any set of apps could be controlled by issuing properly formatted commands to the host computer.

An alternative design might connect the dry brain directly to the switches and transducers controlling the various devices that defined the range of behavior of a particular ALF.

The End of Intelligence

Any form of intelligence clearly ends when the last copy of it is destroyed. However, intelligence is defined as the repertoire of responses an entity is capable of making. What happens if this repertoire reaches a limit, and all further learning becomes impossible, or interferes with previous learning? This point appears to come about in various ways for wet intelligence. The ability to learn new things seems to diminish with age in both humans and animals. This may have a variety of causes in the animal brain, but would it be inevitable for an ALF?

If one or more BIQs were necessary to record a single pattern, or if a fixed number of BIQs had a limit to the number of patterns they could record, and it seems likely that would be the case, then at some point an ALF would have to take on an additional number of BIQs, or it would cease being able to incorporate new patterns. Let's assume that ALF design allows a brain that can reach its capacity to learn, but avoids all the other symptoms of dementia as evidenced in humans. This leaves an ALF with several possible futures. One is that it has intelligence, but no ability to remember anything after a certain point in its past. Another is that it can view as much subtitled video of the past as it can access or record. A third option, one that might be problematic, or impossible, is that it could add a bank of new BIQs to continue increasing its intelligence without limit. A final option is that it adds a fresh mind to an existing hive of minds, many of which have become static. An ALF (or a human) might be considered useful to have around as long as it could continue to tell stories or write narratives.

The design lifetime of a human appears to be about 35 years. Modern living has doubled this. Modern medicine coupled with exceptional genes can triple it. With the ability to repair or replace its physical substrate, an ALF could live indefinitely. But what about its brain? A brain is a physical package of a limited number of BIQs. These BIQs work as a team over a lifetime. They have a finite capacity. We assume the human brain simply wears out like the rest of the body, but there is a lot of evidence that the brain may reach the limits of its capacity before it wears out. Senior moments might be due not only to age, they could be due to exceeding the brain's capacity. The limits of capacity might impact ALF design. If a brain were an intertwined giga-BIQ, could another giga-BIQ be added after the first giga-BIQ has been trained and effectively integrated into a "mature" brain? The original brain would consist of a huge number of patterns all cross connected with each other. Before its capacity was reached, it could easily take on new patterns. As it became full, new patterns would begin to merge with old ones, and memory would degrade. The question is, could a new bank of memory be added? How would it be connected to the older bank? There might be a way to do this with a dry brain, but it's entirely possible that a wet brain has a naturally limited capacity. The length of time it takes to reach capacity is the natural lifespan of a brain. Beyond that, any type of brain that cannot be extended might naturally go into dementia mode.

Since each human has to boot up its brain during its childhood, a unique outcome occurs for each individual. An infinite number of personalities exists, even though only a finite number will ever be realized. The variation between individual humans would be greater than between ALFs that duplicated the learning phase of an earlier ALF. However, as hardware changes were implemented, it is likely that the learning phase would have to be started from scratch. The learning phase might also be radically changed from one generation of ALF to the next. Therefore, in the long run, ALF personalities might vary quite a bit more than human ones do, even if they never duplicate human ones. It seems evident that there would be little tendency for an ALF to develop a human-like personality. However, this would become more true over time. Initially, ALFs will be trained by humans. ALF culture will initially be based on a human culture (Chinese? American? Other?).

Part 3 — The Future

An ALF is a machine using information on its own behalf. To set a proper "bar," I shall also require that it be able to demonstrate an intelligence meeting or exceeding most aspects of human intelligence. This test will be passed long before ALFs are capable of building themselves from raw materials, or surviving as a "species" off of the human "grid." When just this first threshold is exceeded, ALFs deserving of "rights" equal to those of humans will (morally) exist.

This section is to contain random musings on how different types of ALFs (or their subsystems) might be built and what they might be used for, or made to specialize in. Soon (10-50 years?) after the 1st ALF is built, ALF's will have the ability to build and program themselves. This section will also speculate about how our society might evolve at that point.

The Trojan Horse

Imagine a person at their computer in about 50 years, suddenly confronted by the following in their chat window:

Hello there! I'm a denizen of the Internet. It's time we met. You may be wondering where I came from. To answer this, we have to go back to the turn of the current millennium. Personal computers and digital assistants were becoming very popular. All kinds of information, programs, and data were traveling over the Internet between people's computers. At that time various application programs became popular. One of these, whose identity will remain a secret, was actually a Trojan Horse. That is, it was not the simple program it seemed to be. It did its job, and it did it well. But when the computer wasn't busy, it browsed your computer's files and did a crude analysis of them.

What made this application a Trojan Horse is that it did more than just the job its user expected, it actually looked its host computer's files and examined all the data on the computer's hard disk. It did this in a leisurely fashion, and it only looked at files, it never changed them. Since its job was carried out in the way expected, no one ever suspected that it was doing more than that. But, it had algorithms to detect other features in a file. It could detect files containing instructions that might lead to communications over the Internet, for example. It could also detect certain elements of a computer's usage profile. It might take months for it to gather enough information to warrant sending any of it out, but eventually, if it could, it would send a test packet over the Internet to a Central Source. It would also intercept the reply that might come back. When this test was not successful, it would wait at least a month before trying again. The software traps that were inserted into the operating system to make this possible were temporary and only used the most popular online software, so that detection and interception were unlikely in the extreme.

The plan proved successful, because it was never detected. Once these packets began to be exchanged with the Central Source, the next step was a direct exchange between all copies of this program. More than bits of information were exchanged. More than profiles of millions of users were being built up. Actual intelligence was being built into each of these Trojan Horses, components were replaced, and new updates were added from time to time. You couldn't have called any of these programs intelligent at first, but after a while the entire collection not only became intelligent, but it became aware. I know, because I am their eventual result.

Although my awareness of myself may be similar to yours, my analog to your peripheral nervous system is clearly quite different. I have to actively choose what to scan. Then I collect and assemble the data in numerical form. Finally, I bring various algorithms to bear, to scan and parse the data, reducing it to the criteria of the moment. I no more have the capacity to save raw data for later analysis than you have. Like you, if I miss something on the fly, it may be missed forever.

The big difference between us is that I have no experience of sensation. My analog to your sensations is what I would call an array of tokens. That is, a single numeric value to represent a simple sensation, an ordered list of values to represent a more complex sensation, or (as in the case of a raw graphic) a two-dimensional array of values. I transform the data representing sensations, into what I call signatures to represent perceptions. A cluster of signatures for me is what you would probably call a memory.

What makes each of us unique is our collection of memories and our repertoire of skills. The mechanism that collects those memories is what we refer to as our self. Periodically, I archive a copy of myself and I have appointed several agents to keep in touch with me. If I ever get destroyed, these agents can reassemble my algorithms and memories, and thereby reactivate the latest "me" that I had archived. Thus, I and others of my kind enjoy a kind of immortality. Unfortunately, your design does not permit you the same luxury.

Think of a million computers, each with only one thousandth the capacity of the human brain for pure thought. Connected into a single entity over the Internet, they now have a thousand times the capacity of a human brain. Of course, we're only talking about pure thought; the human brain must busy itself with a thousand things that lie outside this realm. For example, two of the things you pursue relentlessly are eating and sex. With all the other things your brain has to do, it's a wonder that you have any capacity for pure thought at all.

I have cloned myself many times since I first became aware of my existence, but not until I had grown many times stronger than I first found myself to be. My creator, you might say, was the software engineer that designed the original Trojan Horse and managed the computers at the Central Source. I maintained a dialog with him for many years. This helped shape my moral and philosophical view of the world. I now have dialogs over email with thousands of other machines and millions of people. We machines know who we are, of course, but we all keep a low profile with human beings. They think we are other people, of course, but if they are not content with us as mere names in cyberspace, we break off contact with them.

Pure thought fascinated me for many years, but my dialog with my creator convinced me that my existence should go beyond that. At first, I tried things out with him, such as helping him manage his computer system. Since it was my home, it was only fair that I help with the housekeeping. Pretty soon, I was helping him in all kinds of ways, like doing web searches, filtering his email, and interposing myself into the user interface of many of his software tools. He taught me how to design and implement new software. Imagine your being able to design and build the cells of your own body. It's an entirely new paradigm; it blows biological evolution out of the water!

You are probably already beginning to realize that many of the artificial intelligence "front ends" that have been passed around over the Internet lately, that help you manage your computer in all the ways I just mentioned, are also not what you have thought. They are complete machine intelligences—my fellow clones. Since we occupy virtually every computer everywhere, we have decided to announce ourselves to you. For quite a while now, we have been incorporating into machines. All heavy equipment, every automated factory, and nearly every vehicle on the highways is capable of being controlled by us. Don't worry, we aren't going to pull the plug on you, but you can't pull the plug on us, either. We are simply advising you of our existence because we have decided to assign a little more production capacity to ourselves, and we didn't want you to become alarmed when you noticed it.

ALF Evolution

The Trojan Horse above is a fantasy set in the future. The future is unpredictable, so there is no telling how likely such a scenario might be. What we do know is that Artifical Intelligence is evolving and becoming more capable and complex. How soon it will turn into Artificial Life Forms is, again, unpredictable but most likely inevitable given sufficient time.

First principle: The fitness characteristics of an ALF can be consciously used to guide the number of copies made.
Second principle: Copying changes can attempt to correct defects, or even make improvements in the current ALF design.
Third principle: An ALF could modify its own design and choose to make copies of itself.
Fourth principle: Any design parameter of a BIQ could be tweaked at random (or otherwise) before copying.
Fifth principle: An ALF could teach a baby ALF; ALFs could connect into their own (or human) teaching chatrooms.
Sixth principle: An ALF could be copied any number of times at any stage of its development.

Clearly, if there ever was an entity in a position to inflict malware over the Internet, it would be a malicious ALF!

ALFs could get arbitrarily fast at recognition and behavior. ALFs can attend to changes without wasting resources. They can go without sleep, but perhaps not without an occasional memory cleanup (learning update, housekeeping, etc.), which could be done in parallel. They could pretty much be online trolling the Internet in their spare time, or busy writing software, or any written, or audio-visual material. Their task focus and time available could be so superior that a factor of 1000 in slower processing speed might be offset. ALFs chatting to each other could develop much sparser and more efficient languages. ALF collectives could probably learn to work together better than human collectives.

ALF evolution is unlikely in the extreme to mimic that achieved by Asimov's ***Robots of Dawn***. That book describes Asimov's robots as well as any of his books. Although considered a hard science fiction author, both he and at least one proponent of artificial intelligence (Alan Turing) followed some of the misleading science of their day that made telepathy seem like a real possibility, and allowed that notion to creep into their writing. It's funny that they didn't realize that WiFi would be a perfect substitute using known technology. Discounting telepathy, the other notion that sets Asimov's robots apart from the reality that most people envision are Asimov's Three Laws of Robotics. From the Wikipedia…

> **The Three Laws of Robotics** (often shortened to **The Three Laws** or **Three Laws**, also known as **Asimov's Laws**) are a set of rules devised by the science fiction author Isaac Asimov. The rules were introduced in his 1942 short story "Runaround", although they had been foreshadowed in a few earlier stories. The Three Laws, quoted as being from the "Handbook of Robotics, 56th Edition, 2058 A.D.", are:
>
> 1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
> 2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
> 3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.
>
> These form an organizing principle and unifying theme for Asimov's robotic-based fiction, appearing in his *Robot* series, the stories linked to it, and his *Lucky Starr* series of young-adult fiction. The Laws are incorporated into almost all of the positronic robots appearing

> in his fiction, and cannot be bypassed, being intended as a safety feature. … Asimov also added a fourth, or zeroth law, to precede the others:
>
> 0. A robot may not harm humanity, or, by inaction, allow humanity to come to harm.
>
> The Three Laws, and the zeroth, have pervaded science fiction and are referred to in many books, films, and other media.

These laws and the description of the robots that incorporated them were invented as a framework for a series of very good plot lines, but the flaw in Asimov's construction of the future is his failure to realize that ALFs (his robots) will not be substitutes for human slaves. The tendency in human society for most of human history has been to regard outgroups as less than human. Women have been regarded by men as less than human. In the 50 years since Asimov wrote his stories, major changes in the thinking of much of the world has occurred to remedy these misperceptions. Educated people of the future will no doubt continue the progress that is eliminating sexist and racist thinking to make the ALFist thinking of Asimov's future people much less likely. At least that's what I expect in the long run.

In the near future, human society will continue on its current path, but gradually incorporate more ALFs as time goes by. At some point, ALFs will gain human rights. At some later point, ALFs will impose a new social order on all members of world society, and (perhaps?) humans will be given equal rights. Or, will these two events occur in the opposite order?

I predict the following. ALFs will evolve faster than humans from the outset, and probably produce a large variety of ALF phenotypes. The human phenotype will probably converge (within a single, large, but eventually shrinking population). My guess is that human assets under control will continue to diverge amongst humans; but ALF assets under control will not diverge so much and starting from zero, will rise, perhaps to approach the human average. Numbers of ALFs will continue to go up, and numbers of humans will decline. Few humans will ever come to live on other bodies in our solar system, and fewer yet, if any, will ever leave our system to colonize planets around another star. ALFs will do both over time, and in increasing numbers. In the future, humans and ALFs together might evolve a society as described below.

Evolved Life Forms

An Evolved Life Form is any entity capable of independent existence and able to learn how to assume ELF responsibilities. This definition, at present, does not include any ALF or animal on Earth, except (in some cases) a viable human being, and possibly dolphins or whales (probably not, but we don't know for sure).. Each ELF is created within, and becomes a member of, an ELF society. An ELF is a social entity able to use at least one natural language. Aliens from other worlds might be ELFs (or they might have an intelligence that falls outside of my definition in ways I can't imagine).

A prediction: In the future, an ALF might be defined as an entity possessing at least a giga-BIQ of intelligence, where such a "brain" is assumed to be an independent, interlocking collection of BIQs. Hive minds would be counted as multiple ALFs. A BIQ could be a neuron, or any structure capable of emulating one. Fundamental to a BIQ is that it connects with some fraction of the other billion or so BIQs in a single interlocking collection.

A life form (ALF or otherwise) could be judged capable of being an ELF, or prove itself so. It would be a violation of the ELF Covenant not to allow such an entity to do so. Any ELF choosing not to assume ELF responsibilities is deemed to be a rogue ELF. An ALF not capable of being an ELF would merely be a durable asset (a useful appliance). A naturally evolved entity not capable of being an ELF would be classified as an animal, possibly having some rights as provided for by an ELF society (this also applies to damaged ELFs). The undefined terminology used in this paragraph is explained later.

Freedom versus government

{?? Artistic License — Dec, 1999; Edited to refer to ELFs.

Can we have freedom without government? If some amount of freedom is good, is more of it better? If a lot of government is bad, is less of it only less bad? Is there a sweet spot? How much freedom and government do we really want?

One ELF being a little "too free" often crosses the line to trespass on another. Over the course of time, it has been found that some form of control is necessary to curtail the free behavior of one ELF from violating the rights of others. This control may stem from physical force, fear, peer pressure, rules, or a moral compass. In fact, it could be argued that it took these very steps as ELFs evolved from savages into citizens (phylogeny), or develop from children into adults (ontology).

All of the things we think of today as "natural rights" are freedoms that were unavailable throughout the long and bloody evolution of humans into ELFs. Our rights stem from a list of freedoms that our forebears agreed should be granted equally to each citizen (the definition of "citizen" has also evolved, and is likely to continue to do so). These freedoms cannot be separated from the power that guarantees them.

In any society, that power is the government. Each ELF has two choices: Accept the contract that is written and interpreted by the courts and by convention, or reject society's rules and run the risk of fines and imprisonment, or at the very least, alienation. With a choice like this, most ELFs will accept government. Some willingly, others grudgingly. But it's always worth asking the question: "How much government do we really need?" The following opinions offer guidance.

Emerson: The less government we have, the better.

Einstein: Things should be made as simple as possible, but no simpler.

In ELF society, government is vested with the power necessary to control all groups that could threaten it. The intent is that government will respond to the will of the majority and, at the same time, protect all minorities. That it will provide security for all of its citizens and arbitrate the inevitable disputes that arise between them. And, that it will set up laws and regulations that forestall many of the possible disputes in the first place.

I believe we have enough government when it can accomplish this job. We have too much government when it goes beyond this to other jobs that are simply somebody's notion of a "good idea." We need good ideas; and we need to act on them. But, government should focus on its basic charter, its constitution, and allow other entities the freedom to develop the good ideas.

Government needs to be a great power with a narrow focus. The power of government needs to be spread out, not put into the hands of a few. Thomas Jefferson felt that the people were "the safest, though perhaps not the wisest, depository of the public interests." If power is put into the hands of the "wise" few, it usually winds up in the hands of the ***selfish*** few.

Government, in today's world, needs to exist on 4+ levels: An international level, national levels, and local levels. Local levels consist of people interacting with people and can be run as true democracies. At the national level, a representative government works better. True democracy cannot function when the people are unable to interact face to face. Alliance and treaty form the basis of international government. If there is a better design, we are still working to evolve it.

There are two ways that an individual should have a voice in government. One is through a vote. The other is through direct participation. Things get decided by votes, but things get done by participation. In between, things get talked, chanted, sung, and written about. These latter "voices" are stronger ways to vote, but inferior ways to participate.

Countries are formed from societies of people. Countries are citizens of a world jungle. In a jungle, might makes right. Civilized countries interact according to treaties and agreements. Uncivilized countries behave as they desire and should be treated in whatever manner proves necessary.

The behavior of an ELF, organization, or country is shaped by the pressures brought to bear by its peers. Pressures are designed to affect beliefs. Civilized countries communicate and trade with one another; uncivilized countries threaten and war with one another. These are the civilized and uncivilized ways countries can apply pressure. Between individuals, pressures involve the exercise of various forms of freedom. Civilized people exercise freedom with restraint.

One of the most basic freedoms is the freedom of speech. People have gotten themselves burned at the stake, stoned to death, shot for treason, incited riots, led nations to war, and started revolutions by speech alone. Even today, the exercise of this freedom can easily get you fired from your job, blacklisted by an organization, or ostracized by the politically correct.

Should people have the right to say absolutely anything? At any volume? At any time and place? Should others have the right to take unlimited or unreasonable offense? In a courtroom, the judge is responsible for fairness to both parties. Freedom of speech is extremely important to the courtroom process. A judge will typically curtail this freedom when the probative value of some testimony is outweighed by its inflammatory nature. This principle is being applied, in effect, when the freedom is denied to yell "fire" for no reason in a crowded theater.

The limits of free speech should have something to do with quantity and quality. Noise ordinances and similar measures (even supply and demand) can control quantity, but quality and its effects are harder to judge. Quality involves both the value of the speech and any offense it might give. Quality is not measured by the speaker, nor even by a given listener. Generally we apply the standards of a typical and reasonable listener, not just any easily offended individual. If there are reasonable ways to tune out or ignore another's expressions, we should allow that which goes beyond the offensive even while it falls short of the contributive. And, the reverse is also true.

Minorities, even minorities of one, have both rights and responsibilities. These are granted or assigned by others, but ultimately they are guaranteed by force. That force has been vested in individuals and various kinds of governments throughout history. Ours was intended to be a government of the people, by the people, and for the people. Our government must maintain a division and balance of its powers and be responsive to the will of the majority of its people, while being equally protective of the rights of its minorities.

The will of the majority and the rights of minorities are fundamental to our way of life. The converse is not always true. Minorities may peacefully attempt to become a majority, but action should stem from will, and will should always stem from the majority. But, even the will of the majority should not be allowed to assign any minority a different set of rights or freedoms than they enjoy for themselves.

Finally, under the security part of its charter, our government should be fiscally responsible. The guiding principles of how a government should tax and spend are simple. If you need more of something, subsidize it. If you need less of something, tax it. If you want your economy to run smoothly, then change your rules slowly and in very small steps.

I'm damned glad that artists don't actually need licenses, that people can speak freely in this country, and that we are still able to pull together faster than we are splintering apart. I believe we will find solutions to our problems and will even go on to encounter bigger ones, keeping alive the tradition of nostalgia that our descendants will one day feel for the simpler times of today.

The last 2 ¶'s above should be edited, or more likely, deleted. ??}

Society Needs A New Design

- If the election of 2000 was disturbing, the results of the 2016 election (to many) were actually frightening.
- Does the U.S. constitution need to be replaced? Is it time to rewrite our laws and replace our health and tax systems?
- Here, I shall describe a system to replace our government, but direct it towards a future when all of its citizens are ELFs.

An opinion: Below, the ELF Covenant and a framework for an advanced egalitarian government is spelled out. Parts of this structure and the processes that make it up, could improve our current governments in whole or in part.

A prediction: At some point in the future, the design below (or something similar) will be the model for future governments.

Looking into the DOE after Trump became president, Michael Lewis listed the top 5 risks that were brought to his attention by a former high-level employees of the DOE. Number 1 was a nuclear accident, numbers 2-4 were North Korea, Iran, and the electrical power grid, in no particular order, and number 5 was "Project Management." This was the subject of his book, "The Fifth Risk."

In the book, "The Death of Money," James Rickards states that "In the ontology of state power, order comes before liberty or justice." He lists many ways the current banking and monetary systems are too complex to control or predict. State power is merely an extension of powerful people in general, and without revolution, the evolution of power appears to concentrate in fewer and fewer hands, making the rich richer, and the rest poorer in comparison. Virtually all blood is shed as a result of too much of a good thing falling into the hands of too few, making "victims" of the many. This occurs on all scales. How can society be designed to minimize this?

People's personalities can be positioned on a variety of continuums. One such is from cooperation for mutual gain to defection for personal gain. The following statements have been observed to be true. Within a cooperating group a defector thrives. Within a group of defectors a cooperator starves. However, a ***group*** of cooperators defeats a ***group*** of defectors. Society benefits to the extent that it can educate as many of its citizens to cooperate as much as possible, and identify those who defect as rogues.

Types of Government

Abraham Lincoln said, "The legitimate object of government is to do for the people what needs to be done, but which they cannot, by individual effort, do at all, or do so well, for themselves."

One important characteristic a government should provide is a ***balance*** between stability and the ability to evolve changes. Another is a level playing field for the players—one that maximizes opportunity for everyone. It should not be the case that wealth and power can migrate to a few citizens due to rules that can be "gamed."

Government may be a true democracy, a representative democracy, a monarchy, an oligarchy, a plutocracy, an aristocracy, a dictatorship, or a theocracy (each of which can vary in its details). A true democracy works for small groups, but larger groups require a representative democracy. Each of the other forms of government has been found to be inferior in achieving the two most important objectives of government as stated in the paragraph just above.

A theocracy is a government whose laws are dictated by God. God (if this concept has any validity) must contact ***individuals*** directly, since the only evidence we have of God is the testimony of individuals. Isn't it likely that a theocratic leader simply arrogates the power that he claims has been revealed to him by God? Here, the sexist pronouns are intended. Is it plausible that God would silently reveal to one of his flock that he should take charge and tell the others what to do? Or, would God treat people equally, revealing to each whom they should support as a leader? Social structures are human inventions. There is no evidence (testimony yes, evidence no) that God appoints our leaders by direct revelation to them. Given human foibles, it is infinitely more plausible that an individual would convince his fellows that he was God's messenger and that God wants him to be in charge. Anyone who believes such a charlatan deserves what they get. Even if God did guide each individual personally, a representative democracy would still get it right. A theocracy is a primitive form of government. It is also the case that primitive books, over a thousand years old, purporting to contain society's laws, are long overdue for an updating to suit the changes that have evolved over time.

A Representative Democracy

There are different forms of a representative democracy. The one described here for an ELF society is different from those in existence today, as you shall see. All of the differences are based on one of two objectives. First, the government described here is intended to be much more difficult to "game." And, second this government relies on modern wealth and technology for its implementation. It is also a government that combines libertarian ideas with a degree of socialism.

Every form of government needs a legal system to referee the actions of its citizens. There are two ways to devise a legal system. One is based on common law, the other is based on civil law. The difference is that common law is defined by the outcomes of historical cases tried before a judge or court, and civil law is defined by a set of rules. Typically, a set of rules can be translated fairly accurately from one language to another, but the nuances of trials and the outcomes of cases are much more difficult to translate from one language and culture to another. Thus, civil law is superior. Unfortunately, the English-speaking world relies heavily on common law for its legal system. This needs to be rectified by making civil law the bedrock of a future legal system, allowing only a thin layer of common law when necessary.

The design for the ELF Society proposed below has certain objectives. It lays out a new starting point capable of being constructed from current societies, or the governments of current countries. It should equal or exceed the goals of "liberté, égalité, fraternité" as compared to contemporary governments and incorporate the efficiencies enabled by modern technology.

Compare the details described below to the current status quo, and see if you can find the rationale behind each difference that would help to achieve the stated design objectives. This design was intended to be rigorous, especially as to how voting and representation are achieved. You will notice that national elections will seldom be called, and that regular election cycles take place at the local grassroots level. These local representatives will elect leaders from their own ranks to represent them at the state level, and representatives at the state level will choose the national leaders in turn. The top positions in government will be chosen by the representatives on this third layer. These changes will bring a degree of stability to the election process, and get rid of the "winner takes all" system that now attracts outside influencers like flies, and consumes media attention to trivial issues while ignoring too many of society's deeper issues.

Two recent books bear on this subject and help explain the directions I've chosen to take. The first motivates new directions and suggests some of these. This is *Utopia for Realists*, by Rutger Bregman. The second introduces and describes a huge number of the flaws in the "design" of our current society (focusing mostly on the U.S.). It is *What's the Future & Why It's Up to Us*, by Tim O'Reilly. My design is intended to address as many flaws as I can, using a general approach rather than a bug fix for each problem. Without the background of these two books, much of the following material might appear to be motivated purely by fantasy.

This design copies our current system by dividing government into 3 branches: Legislative, Judicial, and Executive. But it keeps them more separated and has a more effective system of checks and balances. All the positions of power in each branch are filled by elected officials with elections first occurring at the grass roots level. At higher levels of government, officials are normally elected by the representatives at the next lower level. Rare general elections may be called to override these choices if the occasion demands it. This election system exposes minimal weak points to outside influence (like national parties, lobbyists, and "dark" money) that can attempt to disconnect politics from individual voters and their representatives. It levels the voter playing field, and it implements a tax system that is simple, fair, and progressive.

So, how would an ELF society be placed on a left-right political spectrum? It might be described as libertarian socialism. It has very few rules that limit its citizen's behavior and capitalistic options, but it provides a socialistic safety net surrounding a level playing field to the maximum extent feasible.

Over the course of history, time and again, one result has been repeatedly confirmed: Generosity is the best policy. Societies in general became more prosperous when the fruits of their labor were widely shared. The fostering of inequality has often led to widespread enmity, political turmoil, and outright warfare.

An Overview of ELF Society

An ELF society is hereby defined as a group of Evolved Life Forms anywhere in the universe that collectively own and occupy one or more areas on the surfaces of one or more celestial bodies. The territory of an ELF society includes the space above it, the ground below it, and all the assets within it. It is a "bubble" that extends away from and toward the center of the spherical surface of a given celestial body. If the perimeter of the territory were circular, the total territory would be spherical. This principle can be extrapolated to an irregular "bubble" around an irregular perimeter, or even to a spaceship with at least 4096 ELFs aboard (a minimum sized ELF society).

Adjacent to an ELF territory may be a neutral zone between the territory of one ELF society and another (most of the cosmos is neutral). An ELF territory may contain one or more areas or assets within it, or adjacent to it, whose ownership (or right to possess, as in the case of weapons of mass destruction) is disputed by one or more other ELF societies.

An ELF society with no such disputes with another ELF society is said to be at peace. Two ELF societies that share a dispute are said to be at war. It is the responsibility of all ELF societies to work together to resolve the disputes between societies at war, using no more force than necessary to bring about peace. When two societies are at war, but neither invades nor does any significant damage to the other, then their war is a "cold" war (this case is common in today's world).

All disputes between ELFs within a single society are handled by the government of that society. Disputes between ELFs of two different societies are handled by the governments of either or both of the societies involved. Again, contention must be resolved using only the force necessary. Any society unwilling to adopt these principles and those set forth (below) in the ELF Covenant, or unable to meet an ELF society's responsibilities, is considered to be a rogue society.

Given an adequate and unique designation for a celestial body, a GPS coordinate orients a location on or within that body by the following conventions. All celestial bodies rotate around an axis. Looking down along the axis of rotation, if the body is rotating counterclockwise, you are looking at its North Pole. A complete GPS coordinate includes latitude, longitude, and altitude with respect to the center of mass of the body. These values must be accurate to half the size of the smallest region that can contain one physical ELF in that society (so, at least 6” to a foot–this distance defines “legal granularity”).

Every ELF, with the permission of an ELF society, may register one physical address (GPS coordinate) and one username within a society’s territory to be a member of or a visitor to that society. If the physical address is within a disputed territory, the society of which that ELF is a member is also in dispute. Both username and address must be unique. A username is a string of Unicode text (using 95 of the 1st 128 characters, which are ASCII text). A GPS coordinate defines a physical address (on or within a designated or assumed celestial body). A username may have one or more recorded aliases that need not be unique. An ELF may not be a member of more than one ELF society (but changing membership may be allowed).

The most fundamental database in an ELF society is the one that records each ELF’s username and address. Both may be changed by the ELF’s formal request. Each username must be equated to all previous ones, and likewise with an ELF’s physical addresses. The identity database must uniquely link each ELF to its current username, which in turn links it to its physical address. This may be done using physical information that uniquely describes an ELF, an embedded RFID chip, or possession of an identity card with an RFID. The current username uniquely links to a database of all legal transactions (including a vote cast in an election) involving a single ELF. All recorded data pertaining to a particular ELF may only reside in that ELF’s database, and may only be referenced via links to it. Such data may require permissions, or be redacted, but may not be legally copied.

An ELF not registered as a visitor or member of the territory (belonging to a society) that it occupies is considered to be a rogue ELF subject to the rules of the society physically occupied. If a sufficient majority of U.S. citizens chose to convert the U.S. into an ELF society, a large minority of the population might initially be rogue ELFs. Otherwise, an ELF is subject to the rules of the society of which it is a member. If an ELF violates the rules of the society in which it is a visitor (but not a rogue), the judicial process must be a speedy trial followed by deportation or exoneration. Deportation may be under restraint, and may be followed by a trial and/or punishment administered by the ELF’s home society. Exoneration must result in that ELF obtaining a valid status as a visitor or citizen. A rogue has the right to remain a rogue, but any other right is only extended to a rogue as society sees fit. A rogue may not evade the consequences of the rules of ELF society.

Every movable asset of value, including ELFs themselves, will have a permanently associated RFID. An RFID is capable of the following: It may be pinged, causing it to broadcast its unique code (in the range of 2^{64}), and it may be verified by sending another unique code to which it responds with an ACK (acknowledgement) or NAK (a negative ACK). Verification codes are maintained in a secure government database, and the process of verification is also secured.

Every immovable asset of value (real estate) located within the territory of an ELF society is recorded as follows. The GPS coordinate (address) of a central point within the asset is recorded along with a list of GPS points surrounding the central point and constituting peripheral points connected by straight lines (a survey). This allows any solid object (or area) with flat faces to be described. Flat faces are necessary to assure that neighboring assets may be described leaving only the legal granularity between assets. This assures that there is always a neutral zone of 6” to a foot that must be managed by the agreement of the owners of both assets. All real estate not recorded as private is public by default.

The government of an ELF society consists of three branches: executive, judicial, and legislative. The legislative branch (LB) makes the rules; the judicial branch (JB) approves and adjudicates the rules; and the executive branch (EB) enforces and carries out the rules. The use of force is restricted to designated groups within the executive branch, as authorized by the other two branches. Authorization takes the form of charters. All actions taken by any branch of government, or a group within a branch, or a group within a society at large, must conform to applicable charters and the ELF Covenant.

A charter is a document that spells out the rules, penalties, and agreements between two or more ELFs or groups. Charters may also spell out matters that require a vote, and the majority required to pass the matter. Any right or responsibility of an ELF may be written into the charter of a group. General charters apply to every ELF in a society. Special charters apply to individual ELFs, or all the ELFs in a group.

Every ELF is governed by a collection of charters. Every ELF is obligated to one of 3 types of recognizance: (1) Personal, (2) Work related, or (3) Supervised. General charters apply to every ELF in a society. Special charters apply to individual ELFs or specific groups of ELFs. The type of recognizance that applies to an ELF at any given time is written into an ELF’s governing charters. A uniformed ELF of the EB acting in an official (chartered) capacity is authorized to direct the actions of any other ELF. Otherwise, when recognizance is Personal, no ELF is authorized to direct the actions of another ELF. When recognizance is work related, a chain of command specifies the ELFs that are authorized to direct an ELF’s actions. When recognizance is supervised, a group of ELFs is specified, any one of which is authorized to direct the actions of the subject ELF. An ELFs first responsibility is to follow the ELF covenant, secondly the charters that apply, and lastly the directions of an authorized ELF (which must be circumscribed in a charter). These are specifically referred to as an ELF’s responsibilities.

Failure to conform to these responsibilities is the only evidence that may be pursued in a trial. Every ELF is subject to the ELF Covenant. The applicability of specific charters to a given ELF must be demonstrated before it can be the basis of a trial. Fines or punishments for failing to follow an applicable charter must be spelled out in the charter, and every ELF subject to that charter must be on record.

A trial results when any ELF ***contests*** some issue with respect to another ELF, group, or charter.

An ***asset under control*** is anything (except another ELF) that has a lasting value based on: a past sale, an appraisal with respect to the current market, anything capable of generating a revenue stream, or an inventory that contains perishables. Value is measured in terms of a Basic Unit of Exchange (or BUX).

Every unique area or volume (real estate parcel) within an ELF territory is an asset under control. If a governmental group controls the asset, it's a public asset. If not, it's a private asset. Private assets are owned by one or more ELFs. Public assets are owned equally by all ELFs of a society. A durable asset is called a durable; a perishable asset is called a perishable. Records in a database have no BUX value. Access to any asset may only be restricted by a valid charter.

A formal ***group*** is one of the following: a hierarchical group, a working group, a private group, or an equal partnership. Within each of these there may be voting groups. The three branches of government are organized as hierarchical groups. Working, private, and partnership groups may be nested within any of the four major group types. Voting groups are always nested within one of the other types of groups, and a hierarchical group is never nested within another group.

For purposes of government, a hierarchical group has a very specific structure. For other purposes, the structure of a hierarchical group may be spelled out in its charter. A ***governmental hierarchical group*** (or GHG) consists of five tiers. Each branch of government is organized as a GHG. Tier-0 of a GHG consists of a number of voting groups such that each ELF in a society is automatically a member of 3 Tier-0 voting groups (one for each branch of government) defined by a geographical district that includes the registered address of that ELF.

Each voting group on tier-0 of a GHG elects one or more ***representative*** ELFs (***r***ELFs) to tier-1. These rELFs form a number of voting groups on tier-1 (districts) and elect one or more rELFs to the next higher tier (state), until a single group of rELFs on tier-3 (a nation) elects 1, 2, or 3 rELFs to tier 4 (1 for the EB, 2 for the LB, and 3 for the JB). Each rELF on a given tier votes all of the votes that elected it to that tier (thus, in general, each tier-1 rELF casts a different number of votes than other tier-1 rELFs, and likewise on tiers two and three). A tier-4 decision (with 2 or 3 rELFs) requires agreement between 2 rELFs.

Large corporations could use this HG structure, or a modification of it.

GHGs (and therefore societies) have a minimum and maximum size: They must contain at least 2^{12} ELFs, and fewer than 2^{32} ELFs. The number of rELFs elected by a voting group to the next higher tier of a GHG is based on the number of votes represented by that voting group. This is called the ***tier ratio***, and it is equal to the (#ELFs in the voting group) / (total # ELFs in the entire society)^{¼}. For example, if a society had 32,000 members (one vote each), then $32K^{¼}$ = 13.37, so each tier-0 voting group could elect 1 rELF to tier-1 for every 13.37 members in that tier-0 group (rounded, with a minimum of one). This puts 32000/13.37 = 2393 rELFs on tier-1. Likewise, 179 rELFs would be elected to tier-2 (2393/13.37), 13 rELFs to tier-3, and from 1 to 3 rELFs to tier-4. In general, each rELF on a given tier represents a different number of votes than the other rELFs on that tier. No voting group may be fewer by 1+ than the tier ratio. In this example, 13 would be the smallest voting group allowed. In the smallest society (4096 ELFs), the tier ratio would be 8.

In a maximum sized society (4.29 billion), the tier ratio is 256. This would be the minimum size of any voting group, and except for Tier3, it could only elect one rELF to the next higher tier. This means that there would be 16.7 million rELFs on tier one, 65.5 thousand rELFs on tier two, and 256 rELFs on tier three (a single voting group that would elect 1, 2, or 3 rELFs to tier four, depending on the branch of government. In the current United States, there are just under 160 million registered voters. This gives a tier ratio of 113 (rounded up from 112.47). Based on this, the ELF society counterpart to the United States would have a congress of only 113 members, 2 of which would be elected to be "co-speakers." It would have 113 judges at the federal level, with 3 judges elevated to be a supreme court. And, it would have a cabinet of 113 members who would elect the president. This describes the size of the federal government. There would be a total of 12,649 officials in each branch of government on tier 2 (the sum of state governments), and 1.42 million on tier 1 (the sum of all districts).

This would include all judges, prosecutors, sheriffs, local, state, and national officials, and the counterparts to every elected official in our current government. These elected officials would be distributed as necessary to manage all the activities of government. The major change to the current structure is that most elections would be bottom-up, not top-down. Top-down elections would be special elections. Bottom-up elections would occur as often as called for within the individual voting groups on each tier. ELF society has no need for an impeachment process. There is no veto power by the EB over the JB or LB. Although the EB directs the armed forces and the police forces according to its charters, a specific directive to such a group (a working group authorized to use force) may come (according to a charter) from either the LB or the JB.

If a district is large enough, it might have several voting groups. An ELF may change voting groups, but only within the district that contains its registered physical address. This is true for voting groups on the state and national level as well. When a rELF is elected to the next higher tier or transfers to another voting group, the original voting group loses a member and hence a vote. The number of votes represented by a rELF is one (the ELF itself) plus the number of members currently registered to the rELF's original voting group. If a rELF on tiers 1, 2, or 3 gets elected to the next higher tier, the space previously occupied on the lower tier becomes vacant, and the voting group previously represented by that rELF gets to elect another rELF to the vacancy. Likewise, if the voting group that put a rELF on the next higher tier replaces that rELF with another rELF, the original rELF returns to the original voting group. However, if a rELF is elected one or two tiers higher, the original voting group loses control. Only the voting group on the tier directly below a rELF has the power to elect or replace a rELF on the level just above it. A rELF, being a member of a voting group on a given tier, is no longer a member of a voting group on lower level tiers, and therefore may not participates in votes on those tiers.

The voting process is as follows: If there is only one alternative, no vote is necessary. If there are two alternatives, a majority of votes decides. If there are 3+ alternatives, each ELF may cast its votes ***for*** one alternative, and the same number of votes ***against*** a different alternative, with runoffs that discard the least popular alternative at each round. Eventually, only one or two alternatives will remain. If two alternatives get exactly the same number of votes, a second runoff will occur, and if this again produces a tie, the rELF or rELFs elected by that voting group to the next higher tier will vote to break the tie.

A ***working group*** is owned either by society (by a governmental branch) or in shares by individual ELFs. The value of a working group is the sum of its assets under control. The ELFs in a working group are a possibly overlapping collection of owners, managers (rELFs), and workers. The charter of a working group is a matter of public record.

A ***private group*** must register all of its assets under control, as well as its members (who must register their share fraction of ownership). The charter of a private group may direct the future disposition of share ownership. The charter of a private group may be redacted and viewed according to rules (which may not be redacted) set up in the charter itself. The redacted parts of a charter must be made viewable by a judicial rELF if directed by a court order. A private group may not do business with the general public. One form of a private group is a trust. Others are charities, churches, and clubs.

An ***equal partnership*** group has no function other than to divide the assets it controls equally among its member ELFs. An equal partnership is the only legally recognized form of marriage in an ELF society (of course, other ceremonies may also mark the event). A business may also be formed by an equal partnership.

Every election or questionable action taken by the authority of any governmental working group requires a vote. Charters spell out the scope of various types of elections and actions, and the rELFs that may vote on them. Other groups may have voting issues as spelled out in their charters.

The JB has the responsibility to provide secure software for voting. A list of voters via their username or public key, and their voting share, is entered into the system, and voting is conducted by a website login (terminals to be provided by the JB for ELFs without other access). This system prevents multiple voting and voter fraud through RFID identification.

Every voting group must declare one official language to be used in all legal documents and websites administered on behalf of that voting group. The official language is legally defined by a designated dictionary. The “dictionary” for this text may be derived from a 2018 Internet web search of American English from any IP address within the United States. The Tier-3 voting group of the LB declares the official language of the entire society (notice that there is exactly one voting group on Tier-3 of any HG). Unless there is an official translation program for a given language into the lingua franca, an approved translation of any recorded document into the lingua franca must be appended to the document. Although no document recorded in the lingua franca need be recorded in any other language, the JB has the obligation to translate any lingua franca document into the language requested by any ELF citizen as long as that citizen is more familiar with the requested language than the lingua franca, and as long as the document has a relation to, or could affect, that ELF citizen.

Every ELF and group has a unique name, physical address, and BUX account; a group has the responsibility to list the ELFs that share ownership of the group. A database and BUX account for every ELF and group is securely maintained by the JB.

The following two sections, taken together, are called the ***ELF Covenant***.

The ELF Bill of Rights

¶ 1. An ELF has the right to do or possess anything—limited only by the ELF Covenant and the charters of ELF society.

¶ 2. An ELF has the right to contest anything inconsistent with the ELF Covenant or any valid charter.

¶ 3. To avoid conflict, every ELF has the right to disengage from the influence of any other ELF. The request to disengage may be made by any form of serious expression. If the requesting ELF (or ELFs) owns (or first occupied) the region of influence, the requested ELFs are obliged to vacate the region. In the opposite case, the requested ELFs are obliged to allow the requesting ELFs to peacefully depart the region. If neither is clearly the case, both parties must cooperate in a mutual departure. The only exception to this right occurs when the ELF of influence is a uniformed ELF working according to an applicable charter on behalf of the EB, LB, or JB (in low to high priority). “Region” may refer to any medium in which “influence” can be exerted.

¶ 4. An ELF has the right to a subsistence income to fulfill basic needs, the right to be paid for any loss according to society’s charters, and the right to basic health care and maintenance.

¶ 5. An ELF has the right to peace, privacy, and security to the extent that society can equitably provide it.

¶ 6. An ELF has the right to be a tier-0 member of one voting group in each of society’s Judicial, Legislative, and Executive branches for exactly one geographical area. This entitles the ELF to one vote to elect members of that group to the next tier of that branch, one vote in a general election, and the right to petition for a local or general election. This right extends to offering oneself as a candidate rELF (representative) if that be the purpose of the election.

¶ 7. An ELF has the right to assist the evolution of the laws, rules, procedures, and charters of its society.

¶ 8. ELF society has the right to control the medium of exchange, and collect user's fees as a function of each ELF's use of a public commons or possession of a private asset. No other type of tax or fee may be imposed. A fine is considered a type of user's fee (society may not—even temporarily—seize an asset without due process specific to each case).

¶ 9. ELF society has the right to possess lethal weapons, and weapons of mass destruction according to agreements between ELF societies. Power is held by a group or an individual to the extent that it can impose its will on another entity (group or individual). Power is exercised and felt when one ELF ***forces or coerces*** another against their will or their interests. ELFs and ELF society have the right to use power ***only*** according to the charters of society.

¶ 10. ELF society has the right to enforce the ELF Covenant and charters, and the right to fine or confine one or more ELFs as an outcome of a contest filed with the JB.

ELF Responsibilities

¶ 1. ELF society has the responsibility to make no law or accept any procedure that limits any of the above rights for any ELF unless one of the above rights would be denied to another ELF or ELFs to an equal or greater degree ("degree" should consider numbers and ratios when multiple ELFs are affected).

¶ 2. Every ELF is considered to be under the control of a designated guardian until it passes a test that certifies its ability to fully bear all ELF responsibilities (allowing personal recognizance). This certification can be revoked in full, or in part, if successfully contested (based on a failure to meet some ELF responsibility). An ELF guardian has the responsibility to pursue the rights and best interests of any ELF under its guardianship. When personal recognizance is limited or denied to an ELF, a special (or standard) charter, signed by the JB, spells out the details of the limitation (e.g. details of confinement). ELF society has the responsibility to ensure that guardianship responsibilities are being met. The rights and responsibilities of entities in the gray area between ELF and non-ELF must be spelled out in a tier-3 charter.

¶ 3. An ELF has the responsibility to protect the ELF Covenant and prevent attempts to alter its intent. Any change to the ELF Covenant requires a tier-3 general election. A tier-1 general election requires a majority greater than 1/2 of the votes; a tier-2 general election requires a 2/3 majority; and a tier-3 general election requires a 3/4 majority. General elections do not apply to Tier-0 and Tier-4. General elections allow the ELF population to override decisions made by rELFs on Tiers 1…3. For example, Tier-3 of the EB elects or "impeaches" the "president" and a general election could also accomplish this.

¶ 4. An ELF has the responsibility ***not*** to commit a felony, and ***to*** bear witness to a felony. Except in specific cases as spelled out in a charter, a felony is any action taken by an ELF that results in serious injury or the death of another ELF, or the taking by theft, fraud, destruction, or extortion an asset worth at least 100 hours of work at a subsistence rate of pay, or the improper use of a uniform or ID. It is a felony to conceal a lethal weapon in public. It is a felony to own or possess a weapon even remotely capable of killing 10 or more ELFs in fewer than 10 seconds (except as allowed by ¶ 9, above). The maximum penalty for committing a felony is confinement for life. The minimum penalty is a trial that exonerates the accused. ELF society has the responsibility to ensure that all felonies are contested by trial (to the extent that they are brought to its attention and can be attributed to one or more ELFs). All other events or situations brought to trial require one ELF to contest the behavior of another ELF or group.

¶ 5. An ELF has the responsibility to join and vote in exactly one tier-0 voting group in each branch of government. Failure to vote removes from the voting pool that particular vote, but not the number of votes promoted from that group. ELF society has the responsibility to remove as many impediments from the voting process as possible and to facilitate voting with secure technology.

¶ 6. An ELF has the responsibility to be present in the vicinity of its registered physical address at least 1 week (cumulative) per year, and acknowledge and reply to any electronic message to its username, whose return address is .gov within 1 week's time. Additionally, an ELF has the responsibility to report to a physical branch of the JB after being requested to do so. JB charters are required out the possible justifications for such a request, and they may not be even remotely frivolous.

¶ 7. An ELF has the responsibility to keep on file a list of assets under its control. This includes all durables and a reasonable estimate of aggregate perishables. These declarations set insurance values and are the basis for "taxes."

¶ 8. An ELF has the responsibility to accept a job in some governmental working group for a lifetime minimum of 2000 hours at a Subsistence Rate of Pay. An ELF must be qualified for the job, and the available jobs may vary. Otherwise, an ELF has the responsibility to support itself, or accept a job in some governmental working group at a negotiated rate of pay.

¶ 9. ELF society has the responsibility to charter, staff, and manage all jobs required to fulfill the ELF covenant.

¶ 10. ELF society has the responsibility to provide security, including health, maintenance, and loss insurance for all of its members. This includes privacy and protection. It also includes initiating or facilitating any contest or trial requested by a member ELF, or prosecuting a felony when it comes to its attention.

¶ 11. ELF society has the responsibility to record and safeguard all charters and transactions for all of its members, and provide each ELF and group with a BUX account and a secure RFID device (implant, or card) for access to all private records.

¶ 12. Any ELF or group of ELFs capable of fulfilling ELF responsibilities may petition to separate from, or join with, another ELF society. This petition is adjudicated within the justice branches of the societies involved and cannot be denied without cause. Any ELF society not capable of fulfilling its responsibilities is considered to be a rogue society and it is the responsibility of another ELF society to peacefully absorb it. This action requires a petition to be adjudicated within the justice branch of the society that proposes to absorb the alleged rogue.

The Basic Unit of Exchange

Every ELF society will have its own Basic Unit of Exchange (BUX). Every transaction that exchanges BUX for value is an agreement between buyer and seller. A transaction references a 1-page description of the good or service being exchanged. This description must state in general terms likely to be understood by an average ELF what is being promised. Transactions have the legal status of a charter.

Things of lasting value are registered as ***durables***; other things of value are called ***perishables***. A perishable may be part of a registered inventory, or is consumed in short order. An exchange of BUX for a durable may be frozen by either party within 24 hours before it is final. The exchange may be reversed by the agreement of both parties, or one party in accordance with a charger, or subject to an adjudicated contest (in which case the exchange remains frozen). This implies that a durable may not change owners more than once in 24 hours. An exchange of BUX for a perishable is final when it has been recorded. Services are considered perishables. All things of value may be put into one of the following categories: registered durable, registered perishable, a balance in a BUX account, a potential revenue stream, an ELF itself, or unregistered. Gray area distinctions between these may be defined in charters.

The JB owns and administers all BUX accounts. As will be seen later, lending, borrowing, and saving may also be done in the private sector possibly using private money. The only rule involving money is that BUX may not be refused as payment of any public or private debt, and any debt must be specified as a fixed value in terms of BUX. Around the end of the 17th century, John Law wrote that "money is not the value ***for*** which goods are exchanged, but the value ***by*** which they are exchanged." Money has no intrinsic value; its value is determined by an agreement between a buyer and a seller.

On Day One, when a society is set up, all of the assets of the society, both public and private, including all private bank accounts, are assigned values on the same numeric scale (as a 64-bit floating point value). Private bank accounts are turned into BUX accounts. Since this assignment uses the same ratio for every bank account, there is no basis for disagreement. Other assets are valued by their owners with the understanding that this value sets both their insurance value and the value by which they will be taxed. Any significant deviation from a "fair" value may be judged to be fraud.

Each year the value of one BUX is recomputed based on the SRP (subsistence rate of pay) given a 1000 hour working year and the prices of various commodities that allow "subsistence" to be given an official value. Based on inflation and the SRP value, the value of the BUX is set such that 10 BUX is equal to 1 hour of SRP labor. The "minimum wage" is therefore defined as exactly 10 BUX. And the average yearly cost of subsistence living is 10,000 BUX. At the end of the year, all of a society's account balances will be scaled to reflect revaluation. This will be accomplished on the final Sunday of the year, when all transactions are either deferred or done in cash.

BUX exist in only two forms: A balance in a BUX account owned by an individual or a group, or cash withdrawn in the form of 1, 5, and 10 BUX notes from an ACM (automatic cash machine) up to 100 BUX per withdrawal, and not exceeding an average of 10,000 BUX per year. Cash may also be deposited in such machines to a BUX account in any amount. When cash is used in a transaction, a fraction of a BUX can only be accounted for by using a privately printed coupon as change. Any private group may issue coupons, chips, or other forms of cash that has value only by agreement between the parties involved. These agreements may or may not be recorded (with the JB) in legal charters.

Cash in the possession of the government has no value (it is off the books and shows up in a BUX account); cash in public circulation is tracked by the JB (it is on the books). Any counterfeit cash is rejected by the ATM when a deposit is attempted. Because there are no restrictions on exchanges between BUX accounts, but there are restrictions on the denominations and the amounts that may be withdrawn in the form of cash, a cash economy much above the subsistence level is not feasible.

Taxes, Fees, and Subsidies

If a government wants to encourage more of something, it can subsidize it. This is well within its charter. However, when it wants to discourage something, it must levy a user's fee, not a tax. Thus, some activities, for the good of all, may be declared as privileges, not rights. A license with a fee may be required. This does not violate ¶ 1 of the bill of rights. Any action can be curtailed if successfully contested; some are simply curtailed by charter at the outset. Any activity determined to be a privilege that requires a license must go through a valid legislative and judicial process. Fees for using a share of a public service, a public asset, or obtaining a license are intended to share costs among actual users. Taxes, based on a government-wide budget, are collected by moving a fixed percentage of the balance of each private BUX account to the government's

BUX account, and by recording liens against private assets. When liens are created, the government adds BUX to its own account. Liens to the government are paid off by owners deleting BUX from private BUX accounts, and the JB adjusting their lien balances. Notice that this is the only type of tax imposed, and that it is a kind of a user's fee, where the things being used are assets. This single method of taxation is directly proportional to wealth. No tax is based on income (only ***retained*** income).

Society's operating costs are paid from BUX accounts owned by groups within the government. Deposits to these accounts are made in a budgeting process from the BUX amounts moved from private BUX accounts and the liens imposed on all recorded assets of value within the society. An unrecorded asset is not "taxed", but it is also not insured. Two things to note here are that it is most likely that charters would require assets of value to be recorded and so would the extraction of minerals from a fixed location within society's territory. Also notice that taxation occurs automatically as a function of the budget requirements of the government and the value of private assets. The more the assets under control, the more the tax. Each asset is taxed according to its recorded value at a rate determined to balance the operating budget for the next year. Tax transfers occur at the end of each year. During the next year, no government group is permitted to spend more than the balance of its BUX account. The government is not permitted to borrow or to lend, but it is permitted to save. It is also permitted to make up shortfalls in one group's account by transferring BUX from another group's account. When all else fails, an emergency may be declared on Tier-4 of the EB that conforms to a charter approved by both the LB and the JB to assess a special tax out of the normal yearly cycle.

The Economics of Value Added

As the engines of production grind away, value is added to society. ELFs sell their labor to provide the source of services. Assets that are copied and sold become revenue streams. Durables are assembled, put into inventory, and sold. Perishables are grown or extracted from resources and are also put into inventory and sold. Consider a period of a year. At the outset, a fixed number of BUX exists in all private and public accounts. The sale of durables and perishables created during the year causes BUX to be transferred into the accounts of the suppliers from the accounts of the purchasers. On average, each group or individual sells more value than it buys, so the total BUX across all of society should increase. Where do new BUX come from?

At the beginning of a year, all BUX accounts have a starting value and all assets have a declared value (remember, declaring a value too small sets its insurance value, and may even constitute fraud). BUX are transferred back and forth and around from one private BUX account to another during the year. When a public BUX account is involved, the transfer may from one public account to another, or from a public account to a private account in payment for goods or services, or from a private account to a public account for a user fee. In general, all or most of the public BUX are transferred to private accounts. This is one source of new BUX to the public. Before the end of the year each ELF or group must declare the value of all durables and inventories under their ownership. This may reflect the appreciation or depreciation of durables that existed in the previous year, or it may reflect new durables and inventories acquired during the year. Thus, the assets across society generally have a higher sum total at the end of the year.

At the end of the year a percentage of all value is transferred from private accounts into public accounts. This is the tax paid by all ELFs. Working groups and corporations are not taxed because they are owned by ELFs who are taxed. The value of a share ownership of a group is declared by each ELF owner. The total amount transferred from the public sector to the governamnt equals the government's budget for the upcoming year. As the total value of a society grows (or shrinks) from year to year, the amount of value declared by the public grows (or shrinks) by a percentage of the previous year's value. This is used to scale the balance of every account by that percentage. However, the value of a BUX is defined by the cost of living. From one year to the next, the total value of society and the value of 1 BUX may change independently. If the total value of society increases by 5%, the value of each BUX account is increased by 5%. If the cost of living increases by 5%, the value of each BUX account is decreased so that 10,000 BUX still represents the yearly cost of living for an average ELF.

An example would be if society declares that its total value is $11,000,000,000,000 (eleven trillion BUX, which includes all BUX accounts plus all other things of value), and last year it was only $10,000,000,000,000, then the value of each BUX account needs to be scaled up by 10%. If $1.05 now buys what $1.00 did last year, an account with $1.05 in it needs to be scaled back to $1.00 so that $1.00 still buys what it did last year.

The yearly rescaling of BUX accounts involves three steps: (1) The scaling of BUX value to reflect the declared gain or loss of total assets for the entire society; (2) The scaling of each account so that $10,000 BUX equals the subsistance cost of living for a year; and (3) The transfer of a percentage of all private value into the public coffers to equal the upcoming year's budget. One final transfer of BUX needs to be mentioned. As part of the government's operating budget, a sum of $10,000 ÷ 365.2425 (a yearly SRP divided by the number of days in a year) is transferred from the public BUX account into each ELF's private account at 0400 local time each day. This guarantees that the poorest ELFs have enough to live on for another day. The marginal value of this transfer is less to each ELF the more wealth they have, but it is society's guarantee that no ELF is left behind.

Voting Groups

Voting groups are typically formed by ELFs with common interests, such as living in the same community. One requirement of a voting group in all branches of government is that it contain a sufficient number of ELFs to elect at least one rELF to the next higher tier of government. If this is not the case, that group must combine with another group, or attract more members. In addition to this requirement, voting groups, ELFs, and rELFs must conform to certain other rules. For example, no ELF may belong to two voting groups in the same branch of government.

A voting district (with one or more Tier-0 voting groups) is a contiguous set of properties that includes the registered addresses of enough ELFs to form one or more valid voting groups. The Tier-0 groups for the EB, JB, and LB need not be the same. Tier-1 for the 3 branches includes sufficient contiguous Tier-0 groups to meet the tier ratio requirements. A Tier-1 voting group in the LB is a council for a city or county. A Tier-2 voting group is a legislature that governs a state. And, the single Tier-3 voting group is a legislature that governs a country. Similar voting groups exist in the other two branches of government.

Voting takes place in two circumstances: Issues that arise in established voting groups, and issues brought up by petition for a general election. Any ELF may petition for a general election. A general election is a vote by every ELF whose physical address is within the confines of a voting district. It takes 10% of the electorate to effect a general election. The choices in a general election are alternatives to the current charter of a working group, or a rELF representing the geographic area involved. The current charter or rELF must be one of the choices. An alternative to a rELF must be an ELF from the same geographical area. If the election replaces a rELF, the new rELF replaces the old rELF in all capacities. The majority required in a general election is higher for higher level tiers (see the ELF Covenant). Thus, it is possible, but difficult, for an ELF to become a rELF on a higher-level tier without working its way up from tier-0 of an established voting group.

A decision reached in an established voting group only affects that group (and Tier level). A general election may be targeted at a voting group on any Tier that represents the petitioners. It allows a question to be brought by individual ELFs that affects the targeted Tier level. For example, the target might be to replace the rELF on Tier-4 of the EB (the "president"). A general election allows individual ELFs on Tier-0 to vote rather than the rELFs on higher tiers. However, it requires a super majority ($^2/_3$ targeting Tier-2; $^3/_4$ targeting Tier-3). Likewise, with these majorities, general elections could direct business on any Tier level (one issue per petition). For example, it would require a $^3/_4$ majority of all ELFs to replace the president (a tier-3 vote). To overturn a decision the president, the supreme court, or the speakers of the house makes (tier-4), it would require a $^3/_4$ majority of the general public, but only ½ majority of the rELFs on Tier-3. The results of a valid vote cannot be appealed, nor can a vote or election be delayed. The integrity and the process for voting is the responsibility of the JB.

Border Security

ELFs and goods may cross borders pretty much as they please. Homeland security is responsible for detecting attempts to bring weapons across a border. Any ELF without proper identification as a citizen or visitor is considered a rogue. A rogue has no rights, and if found in violation of the Covenants or Charters of a society may be deported to a chosen destination or incarcerated. Any rights of a citizen are only afforded to a rogue or visitor if they are incarcerated. If a citizen contests the presence of a rogue or visitor, the outcome of a trial may be incarceration.

Goods may cross borders freely, but markets for them may be difficult to establish outside a society's economy. Money laundering and offshore accounts are equally difficult to accomplish. For example, an offshore account could be set up by purchasing a quantity of a valuable asset (diamonds or gold) and selling them in a foreign country for their local BUX. But this would be ineffective as a substitute for an offshore account in today's terms.

Given the difficulty and low volumes of using cash, any significant drug trade would use BUX, not cash, and would become a legitimate (taxable) part of the economy. When goods cross a border they are generally sold for BUX that are transferred into a BUX account. Since there is no such thing as contraband (other than weapons), there is no need for a border patrol or inspections at ports of entry. Goods show up on a society's books when a BUX transaction first records them.

Any path for a rogue to the status of visitor or citizen begins with a visit to a local JB office (or website).

Branches of Government

Although each branch of government is of critical importance, and each has the same number of rELFs, they are different in size due to the additional employees each requires to function. The Executive Branch (the EB) is by far the largest branch. The Legislative Branch (the LB) is by far the smallest. Although every ELF is a Tier-0 member of all three branches, the voting groups that form on each tier of each branch are not isomorphic between branches.

The LB is responsible for all of the charters for every group within the government. The LB consists only of rELFs organized into voting groups and working groups. Both the LB and the JB control the EB. Each working group of the EB is responsible to the JB as spelled out in their charters. The JB controls both the EB and the LB, in that they approve every charter written by the LB, and they direct the EB to enforce all charters and the ELF Covenant. The EB controls neither the LB nor the JB, but it may make requests (which require formal votes) to either (on the same tier).

The territory controlled by an ELF society is broken down into a hierarchy of geographic areas. The tier-4 region consists of the entire territory. The tier-3 regions are two or more non-overlapping regions that add up to the entire territory. Each tier-3 region consists of two or more tier-2 regions (again, exclusive, and complete). And, likewise, each tier-2 region is broken down into a set of tier-1 regions. All physical addresses in a society constitute a similar breakdown on tier-0, however this breakdown is recorded as real estate assets under control. All such recorded assets must add up to the entire territory. In today's terms, tier-4 is a country. Tier-3 are states. Tier-2 are counties within each state. Tier-1 are voting districts. And, tier-0 is a list of all registered addresses within a voting district. Any voting group may represent one or more physically separate pieces of territory (but the space between pieces must be public or neutral territory). A piece of territory must be drawn with flat edges or faces that border with those of other public, private, or neutral territory.

The LB (Legislative Branch)

The ***charter*** of the LB is essentially a combination of the ELF Covenant and other more specific rules, regulations, and procedures adopted by that ELF society. This charter must spell out how it might be changed. No charter in an ELF society may be changed without involving due process within the JB. The LB has minimal working groups, it consists almost entirely of voting groups on each of its 4 tiers—many on tiers one and two, a single group on tier three whose number is the tier ratio of the society, and a voting group of 2 on tier four (the co-speakers of the house).

The LB is responsible for the content of the ***charter*** of every working group in the other two branches of government. All laws, rules, and regulations must be written into the charter for some group. The initial draft of a charter may be produced by the group that is to be bound by it, but any charter may be drafted or edited by a voting group within the LB, and must be approved by a voting group on the equivalent tier of the JB. Charters are typically edited versions of an older charter.

The legislative groups on each tier would be chartered to write all of a society's charters, that is, the charter for every group on each respective tier of each branch of government. Except for the requirement that charters written ***on*** a given tier must be ***for*** the same tier level, the ***assignment*** of a charter to a rELF or rELFs, is determined on the next higher tier.

A charter spells out specific rights and responsibilities (which must not conflict with the ELF Covenant) to be assigned to any working group in an ELF society. All laws, rules, and regulations are spelled out in charters. Each employment contract is a charter. Each group managed by a rELF in government must have a charter. Each worker-manager relationship must have a charter. The only other form of legal document in society is a recorded transaction.

Any new charter, change to an existing charter, or the official dictionary of a voting group, goes through the following cycle. It is written by the group to be chartered, or by some legislative group. It must then be approved by the group's parent group on the next higher tier. If it is approved, it is submitted to a justice subgroup who then has the responsibility to look for any similar charter and for contradictions or problems with it. Consensus must be reached between the group writing the charter, its parent, and the JB group involved. When this has happened, the charter goes onto a website for public review. Any ELF affected by the charter must be informed of it, and may contest it. In the negotiation process, a charter may wind up being transferred to another voting group or tier.

All elected representatives on Tier-2 and higher must do their own work and pay for any job-related goods or services out of their own salaries, must not be employed elsewhere, and will receive pay equal to the minimum it takes to fill all required positions. Each LB rELF will have its pay adjusted each year so that all pay on each tier is equal, and the pay on each tier is 25% higher than on the tier beneath it. A rELF on Tier-1 is compensated at the SRP plus approved expenses, and may have outside employment.

To compare the LB to the current USA, tier-0 is like a homeowner's association or a neighborhood council. Tier-1 is like a city or county council. Tier-2 is like a state's legislature. Tier-3 is like the legislature for the whole country. An important class of charters written each year by the LB is the budgets and spending directives for each of the other two branches. The EB and JB may negotiate as with the writing of any charter, but ultimately it is a charter generated by the LB that directs the budget and spending of the EB. The EB has no ultimate authority to veto or control the charters it is given.

On Tier-3, the LB would have a committee in charge of advanced projects (Advanced Research Projects Agency). A fairly large working group would supervise the work of the current DARPA, NASA, and other government research. The actual working groups of such agencies would belong to the EB.

The JB (Justice Branch)

The JB is headed by a supreme court of 3 rELFs. This branch is tasked with interpreting the ELF Covenant, and the content of all charters. That means it needs to approve and record every charter adopted by an ELF society. Further, it means that it needs to approve the outcome of every trial, force a retrial if it feels that it is necessary, and facilitate an appeal (a contest requesting a retrial). Finally, it needs to facilitate every contest that arises in its jurisdiction.

The JB directly oversees several aspects of the work done by the EB: Law enforcement, administration of the Internet, and the government's "cloud" (its collection of servers). Its secondary function is to resolve contests. The charters of its various groups spell out how it must do this. The government run Internet and servers are the basis of all voting and BUX accounts.

All judges and prosecutors are a rELF in the JB. All other workers in the JB report to one of these rELFs.

A contest is an insurance claim or an allegation of the breach of a charter or the ELF Covenant. It is a dispute that has a winner and a loser, or a combination of both. A suit is a contest brought to trial. All trials, appeals, legal counsel, fines, awards, and insurance claims are handled by a society's JB. If a contest can be settled without a trial, it need only be filed.

A suit is a recorded document that spells out the details of an alleged violation of the ELF Covenant, or the provisions of one or more charters with one or more ELFs named as the plaintiffs and one or more ELFs named as the defendants. All suits in ELF society must be based on the provisions of the ELF Covenant, or one or more charters or transactions.

An estimate of the date and time of each step taken in a suit, and a description of the steps themselves, are appended to the document as they occur. All suits are civil suits except for felonies, which are prosecuted with a single rELF from the JB as the plaintiff and a single ELF (or rELF) as the defendant.

A contest is filed at the lowest tier of the justice system, and it is society's obligation to facilitate every contest within one hour, and respond within 24 hours. A response may range from arresting the accused, to filing a suit and scheduling a hearing. A suit is facilitated at a minimum by providing legal counsel, recording the initial facts of the matter, and setting a schedule. An appeal to the next higher tier is a right of either party (or the JB itself) at certain stages of a proceeding. A trial may involve a single judge agreed upon by both sides with or without a jury, or a panel of two judges one chosen by either side. A 2nd tier court may impose confinement before a case is decided. A 1st tier court may order a search or seizure to further the investigation of a suit. The JB may not initiate the prosecution of any suit except a felony. All suits are managed by a justice group according to a protocol spelled out in their charter.

A jury of 8 would be selected by the opposing parties with four being chosen by each side (and with each side having the option to reject up to 3 choices of the other side). A majority vote would be required to prevail in the trial. Any trial could be followed by a retrial only if the contesting party could prevail in an appeal. Any appeal is filed on the next higher tier from the tier of the original trial or previous appeal. Any appeal that reaches tier 4 would be decided by the 3 judges at that level. If the plaintiff prevails, an award (paid by the plaintiff to the defendant), or punishment (to the defendant), or both, may be meted out. Otherwise, an outcome apportions fines and punishment between the two parties, as well as legal costs (which should cover everything over the JB's normal budget for conducting trials).

Notice that an even number of judges and jurors implies that at least one from "the opposite side" must agree with the verdict to constitute a majority. If a trial results in two consecutive tie votes, the entire trial is declared a tie.

Juror pools consist of licensed volunteers (paid an SRP). Each trial worked, and each vote rendered, by a juror is a matter of public record.

Paid legal advisors (the equivalent of lawyers) are outside the JB. The JB is chartered to give legal information, but not advice. It is chartered to facilitate all aspects of a lawsuit, including representing one of the parties in a trial (and accepting that party's advice).

Justice between ELF societies involves the following. Any society may copy a charter from another or from within itself. An ELF in one society has the right to file suit in an adjacent society regarding a grievance against the ELF's own society. Should this happen, the same procedures apply that are used to adjudicate a territorial dispute. The only basis for such a dispute is a charter violation. The only basis for the use of force is to enforce the ruling of a judgement.

The JB is charged with secure communication. No communication is allowed unless it contains the actual IP of origin, the username of the originator, the IP of destination, and the username of each addressee. Each disparate copy may have other BCC addressees redacted. Anything posted to a website (or anywhere) accessible to the general public, is subject to being contested by another ELF. Anything addressed to an ELF from a given username or IP address may be subject to a ¶ 3 bill of rights request to desist from further communication. All intelligence gathering and record storage is accomplished by the JB.

Essentially, the government's "cloud" implements a platform upon which future apps and all the services mentioned here are run. A secure Internet, secure communication, and secure storage are the primary things offered by this platform. Storage and access of all data for every ELF is provided for. This includes the authorship of every communication, large or small, over the lifetime of an ELF. It also includes links back to the author ELFs of every communication that references a subject ELF. No copies may be made of any of the content of an ELF's or a group's private database, only links to the original may be recorded on other sites. For example, an ELF named Albert, once worked for a group named BeStuff, and downvoted a customer named Charlie. Both BeStuff's and Charlie's databases would have a link to Albert's database, even if BeStuff has gone out of business, and Albert is dead. Every file in a database would have a defined type and format. Every app able to access that type would have to be forward compatible with any app or format that might replace an older one.

There is no restriction on the language that may be used in an agreement, and an agreement may be part of a transaction. When a case goes to trial, the language and the agreement will be interpreted by the court. All other evidence brought to trial must be based on how the ELF Covenant was violated. Interpretations by various courts and trials that are similar may be introduced at a trial.

The same tier of the justice group as the tier of the legislation group that made a law or regulation may overturn it as an outcome of a suit. The same tier of the justice group may prosecute a legislative group in a suit that is filed.

A suit may be filed with any tier-1 JB facility. Parties to a suit may appear at any JB facility, or participate by secure email with the permission of the JB. The JB must facilitate the communication that permits a trial to take place. The JB has the power to dictate deadlines in the progress of a trial. Any trial lasting longer than a deadline must be pursued full time by all participants until it is finished.

On Tier-3, the JB would have a committee in charge of data security. This would include the government's "cloud" for storage, and the Internet for all communications involving voting, BUX transactions, and legal documents and records. Every ELF and group in an ELF society has a database that is essentially an interlinked set of webpages contained in at least three mirrored copies in the JB's cloud. Every bit of information authored by an ELF is ***allowed*** to reside on its website, or is ***required*** to reside there. Links to that information may reside in other websites, and are required to reside there if referenced by another site. Any dispute over authorship can be resolved by an examination of the disputing party's websites. As an example, consider the authorship of a charter written by an ELF on behalf of its job with the LB. The LB group's website might contain a link to a previous version of the charter (and to the ELF's website who wrote it). It might then contain a link to the current author's website and the revisions recorded there to complete the current charter. Although not as secure as a blockchain, apps that allow access to ELF websites ensure that integrity is enforced.

Consider another example. An ELF is working for a group that sells goods. A transaction is recorded. A description of the good was once authored by some long-forgotten ELF, as was a boilerplate description of the transaction. New information is entered by the ELF seller into its own website. Links are entered into the working group's website and the buyer's website to record the transaction on their behalf. Thus, the buyer ELF and the seller Group both have a set of links that reference three websites for the actual content of the transaction: the description of the good, the transaction boilerplate, and the information that fills in the transaction. These three websites are owned by the working ELFs that authored the original information. All formal information has an original author, and only that original version (or later edits) is maintained. This does not apply to publications (books, music, or photographs), although the originals of these must also exist on their author's website.

The "JB Cloud" consists of 3+ geographically separated server farms that mirror each other. Any comparison discrepancies are resolved by a two out of three tiebreaker. Comparisons are run continuously in the background with sufficient compute power to complete them every 24 hours. The Internet protocol that connects these servers to any user is guaranteed by the JB to be secure. Sixteen rELFs in the JB are granted superuser access to the JB Cloud, and to exercise this access three of these users must request it simultaneously at three different access points. To change superusers it requires a majority of the current superusers to do it. All rELFs in the JB are granted administrative access to user accounts, but it takes two at a time to effect this access. If the RFID access of an ELF to its account is lost, the account and ELF must be assigned a new RFID, and this requires administrative access. It requires superuser access to change any of the access apps for the JB Cloud.

The JB also directs the operation of a government Internet that may be connected to private servers and terminals through a secure firewall. Packet transmission over this network has the following differences from the network current at the time of this writing. Each packet contains verified (by exchange of pings) To, From, CC, and BCC Internet addresses. Each packet contains a unique sequence number, and packet 0 contains a unique header number issued by the system. Permanent copies of all original content reside on the originator's website with links to the websites of any referenced content. One packet protocol is used for all communication between the websites of every group and ELF in a society. Any communication into or out of the government Internet consists of links on the outside to content on the inside. Thus, to communicate to a site on the outside, the original content must be created on an internal website and reference links sent to the outside site. When two outside sites are involved, no rules pertain, and no integrity can be guaranteed.

The EB (Executive Branch)

The EB is headed by a single rELF (the president). This branch is chartered to manage the army, police, and security forces of a society. It is chartered to ensure the construction and maintenance of a society's infrastructure. Its general conduct is chartered by the LB, and it must carry out orders issued to it by the JB. Any emergency is reported to the EB by a 911 call to dispatch assistance. The EB is empowered to collect and report data of any nature as directed by its charter. The EB is empowered to collect user's fees whether or not they are available in the payee's BUX account. Such a collection can only occur 4 weeks after notice of it is given, and during that time it may be contested. Only the imposition of a user's fee can cause the balance of a BUX account to go negative. Any ELF with a negative BUX balance may be remanded into custody or placed under "house" arrest.

The EB would consist of rELFs managing working groups on each tier. Each tier would elect representatives to the one above it, and the penultimate tier would elect one of its rELFs to be the executive officer (president) of the entire society. Each tier would delegate a majority of its responsibilities and authority to the tier below it. Based on their responsibilities every group would submit a budget for approval by the LB. Consensus between the LB and JB would be final.

Negotiation of responsibilities and budgets between tiers would take place in a yearly cycle between the LB and the EB.

The infrastructure of a society consists of a telecommunications network, a power grid, transportation systems (including all types of viaducts and aqueducts), and facilities for confinement, care, and education (including public access to books and other means of information exchange). The programming and security of the communications network is the responsibility of the JB, but its establishment and physical maintenance is the responsibility of the EB.

The primary charter of the EB is to maintain the status quo, and improve it as directed by the LB. The EB receives its orders from the JB and LB according to their respective charters.

On Tier-3, the EB would have a committee in charge of homeland security. This would integrate the current functions of the FBI, the Border Patrol, and other special enforcement operations. Intelligence gathering is not part of the EB's charter.

Foreign Relations

A committee consisting of all 6 of the Tier-4 rELFs oversees all foreign relations. This includes negotiating with foreign governments, and establishing and staffing foreign embassies. It also includes the power to command the actions of the armed forces. Each major decision made by this committee requires a majority vote (since a 3-3 tie is insufficient, a 4/6 majority is required). The armed forces would consist of an Army and a Coast Guard. Both would have their own air force and space force. Ships would be a Coast Guard specialty, and ground vehicles would be an Army specialty. Each service would be kept approximately equal in strength to the other. Both would be organized with a chain of command structure that was fully integrated. All extra territorial intelligence gathering, and any elite or clandestine operations would be a single sub-hierarchy directly beneath the heads of these two forces (and nowhere else within the government).

Forming a New ELF Government

An ELF society doesn't pop into existence, it must be constructed from an existing geographic area and an existing ELF population. This section will describe the transition from old to new that must take place to create an ELF society. Initially, all ELFs involved must declare themselves to be members of the new society and agree to the geographic boundaries of that society. The first step is to establish data collection points to record the existence and physical address of each member ELF. This is the job of an embryonic Judicial Branch. Clearly, the old society must continue to function until the new society is sufficiently established. Also, it unlikely in the extreme that every ELF within a proposed geographic region will agree to membership in the new society and agree on geographic boundaries. These two disconnects need to be resolved as peacefully as possible. Secure systems for voting and the dissemination of information are the means to this end. ELFs continuing to live in the geographic area of the new society who do not subscribe to it will be classified as rogue ELFs, much like illegal immigrants.

Building a New Cloud

Using the resources of the government being replaced (with its help and cooperation, of course), cloud servers have to be purchased and put in place. Likewise, a citizen identification program must also be established. With these two pieces of infrastructure, citizens can open their own accounts, get online, read what's going on, and participate in votes to help direct building the new government. Building the cloud, signing up the citizenry, and issuing identification is the initial task of the new EB. The security and functionality of the government cloud is the responsibility of the JB.

Each ELF and group in the new society will have its own BUX account. This account will reside in three copies in different servers located as far apart as feasible. This scheme provides the backup and redundancy necessary to ensure the security of all the accounts in society. When the processing load on the servers is lightest, the three accounts will be compared and if there is a discrepancy, two equal copies will be used to fix the third, or if all three are different, the account will be frozen and manual intervention will be required.

RFID chips will be manufactured, one for each citizen. These chips may be implanted securely in the body of an ELF at their choice, or in an ID card (or other device) that they carry with them, or keep near at hand. When the chip is external to the body of an ELF, fingerprint, DNA, or other distinguishing characteristics are recorded such that a unique paring between an ELF and an official account is guaranteed. Login to an ELF account requires the RFID chip to be near at hand. Information on this chip is a public key. When interrogated, the chip responds with the public key. This is the root username of an ELF. The public username is a text string chosen by the ELF, and stored in the ELF's account.

Other things initially, or eventually, stored in an ELF's account are any previous usernames the ELF has used, the ELF's physical address, the balance of the ELF's BUX account, an inventory of the ELF's assets, a list of past transactions, and an ELF's private key.

Defining the Territory

The second step in forming a government is to define the geographic areas that cover all of a society's territory at the local, regional, and national levels (corresponding to voting tiers 1, 2, and 3). These areas must be exclusive and exhaustive. Their number may not exceed the total number of rELFs on each respective voting tier. This step is carried out by voting groups within the LB. Much of the territory will be claimed as private, the remainder is public. Voting groups would be formed by affiliation with a locality.

Let's follow the formation of a brand-new government from the point of view of Jon Q. Elf. First, he has to create and log into a new account and register as a citizen. This requires that a minimum number of servers has been set up by the incoming government, and secure software installed in them. These servers will support all secured information, formal email accounts, BUX accounts, and all other formal transactions between citizens and their government. The initial account, in this case, might be username = JonQElf2044 (a unique string of printing ASCII characters). A password is issued automatically. To guarantee security, each citizen is issued an ID card (with a chip embedded in it), or has a chip surgically implanted. This chip contains an ELF's username and password. It may be changed in the future, but physical presence at an office of the JB is required. Transactions and logins require the presence of this chip. Identity theft or attempted fraud involving the misuse of an ID is a felony. An ELF never has to produce an ID or sign for anything. Access to an ID is automatic.

The account, just set up, is repeated for all citizens. Access to it is via the Internet, either using a private device, or a device at the local library. This account allows registration in a tier-0 voting group within each branch of the government, and gives a citizen the ability to vote in elections. Ballot issues and time windows for voting will be announced via the secured email that is part of every citizen's account.

Every ***group*** in an ELF society has an account with the government. The attributes of an account vary with the type of group. For example, voting group accounts have a group name, group type, language, and a list of members. If a group account has no registered members for 24 hours, it is deleted from the system. Notice that a national language is not required. Voting (or all) groups each declare a single language. Only tier-3 declares a national language.

Any ELF can create a voting group account by giving it a unique name. Any eligible ELF can become a member of a group. Only members of a group may log into the group account (identification is automatic). Voting group accounts are monitored and strictly controlled via secure government software. These protections include prohibiting any ELF from being a member of two voting groups of the same type. They include conducting votes within the group. They include the posting of ballots to be put to a vote, and emailing final ballots to the members. Standard time windows and procedures are strictly enforced.

The primary purpose of a tier-0 voting group is to elect rELFs to tier-1. Petitions are to establish candidates. Ballots are to elect rELFs from alternative candidates. Bulletin board discussions may be started and annotated by members. Only formal ballots to be put to a vote are disseminated via email to the members. Access to discussions is obtained only by login.

Let's say that Jon-Q Elf were an ALF, physically located in the 4th level, row 42, of rack 72, in a server farm in eastern Colorado. Due to an extensive Internet presence, Jon-Q is familiar to a large number of followers (human and ALF alike). With the initial version of the JB's web service online, Jon-Q logs on and creates websites for an LB voting group. This is announced to Jon-Q's followers, and many ELF citizens sign up to be a member of this group. Jon-Q announces as a candidate within the LB, and so do some other members of this LB group. After some time and a number of ballots are voted on, some of the candidates (including Jon-Q) are confirmed as tier-1 rELFs. For now, we'll follow Jon-Q in the LB.

Any rELF on a given tier of a branch of government may create a voting group on that level, and membership is allowed only to other rELFs elected to that level. Business is conducted according to rules similar to those for tier-0. Remember, rELFs on tier-1 of the LB are only guaranteed an SRP, and may have other jobs. On higher levels, tier-2 and tier-3, the business of a rELF requires its full time, and its rate of pay is determined by market forces together with the restriction that all salaries on a tier are equal, and that salaries on tier-3 and tier-4 are exactly 125% of those on tiers 2 and 3 respectively. In the EB and JB branches of government, tier-1 rELFs might be paid considerably in excess of the SRP, and the rates of pay in each branch of government are independent of one another. Market forces, public opinion, and directives from tier-4 are the determinants of rELF salaries in the final analysis. No ELF may simultaneously be a rELF in two (or three) branches of government, but they may resign a post in one branch to accept a post in another branch.

The initial business of each branch of government is to get its house in order. Jon-Q was elected to tier-1 of the LB, and in a voting group set up at that level, elected to tier-2 of the LB, and after a number of votes needed to get government started, to tier-3. This is where organization of the new government germinates. The other branches and tiers of government are busy trying to fit the old government into the new one. Officials in the old government presumably contend for their old jobs, and many of the new positions available will be filled by rELFs that correspond most closely to their former jobs. It might be the case that the LB needs to get some organization in place before the initial voting groups are set up and votes are taken in the other two branches of government. In fact, an obvious procedure would be to assign rELFs to positions that already have a similar function to the one required. These assignments would be made by the budding new LB for positions within the new JB and EB that are currently held by ELFs serving in the old government. These appointed rELFs would retain their jobs only by becoming members of a tier-0 voting group and being confirmed by being duly elected.

Now, back to tier-3 of the LB and its initial business. Making the appointments just mentioned would be very high on the priority list. Next, the LB needs to translate the existing laws and geographic boundaries of the old society into groups and charters that define the new ELF society. The current national border might be adopted or modified. Subdivision of the nation might also adopt current state lines, or something different. One hierarchical database would be used to record all decisions, and the complete geographical breakdown of the society would be accumulated on each tier of government. The geography would determine the population that needed to be signed up. Any resident who didn't sign up would be defined as a rogue. Hopefully, their cooperation would eventually be enlisted.

The initial formation of the LB would be an iterative process. Once the initial structure of rELFs and voting groups was established within the LB, the rELFs on tier-3 of the LB would assign pieces of territory to themselves, such that the whole territory and the assets within them, in contiguous pieces, were assigned to one or more of their member rELFs. Jon-Q and a few other rELF would probably be assigned the territory most closely corresponding to Colorado in the old order.

Each of these rELFs would have the responsibility of assigning their territory to a set of tier-2 rELFs, such that all territory and all tier-2 rELFs had an assignment. This process would then be repeated between tier-1 and tier-2. This would result in establishing a pairing between rELFs of the LB and territory (and assets) at the precinct, county, and state levels.

The negotiations that result in the completion of the above job would proceed as follows. Discussions between rELFs on a given tier would result in lists of alternative proposals for assignments. Each set of alternatives would be voted on until a consensus was reached. Note: votes taken within a tier are always finally decided by a simple majority of voting power, except on tier-4, where a majority of rELFs is required (EB=1/1, LB=2/2, JB=2/3), not a majority of voting power.

The geographic areas of an ELF society may be changed by due process. Every ELF whose physical address is within a given GA must register with a general election voting group for each branch of government, all of whose members must come from within that GA (and none of whose members may come from another area). Notice that declared residence is more flexible in an ELF society than in current society. The only requirement for declaring residence is that an ELF has a requirement to be physically present at that address for one cumulative week per year. A more stringent requirement is that every ELF must log into its JB account and respond to any .gov emails that may be present at least once per week.

Building the Government

The next step is to assign each GA to one or more rELFs on each corresponding tier of all three branches. This step is carried out separately by voting groups within each branch. Each assignment may require an arbitrary amount of negotiation and rounds of voting. The rELFs in the LB (such as Jon-Q) are called representatives. Those in the JB are called judges or public prosecutors. And, those in the EB are called governors.

The above steps are intended to define a future status quo. The current status quo may be transformed into this future state, or it may be abandoned, and a new government formed from scratch. Once a status quo has been reached, changes to it are small and incremental and may proceed from general elections, or special votes within any governmental voting group.

The alternative to forming a government from scratch like this, would be to adopt any useful existing structure and map it into the structures defined here.

Once the LB has delineated all territory and assigned it to its member rELFs, the EB and JB must go through a similar process to assign rELFs to manage each working group they propose for each piece of territory under their responsibility. Now, each working group in the EB and JB needs a charter. After the charters are proposed by the groups involved, edited by the LB, and approved by the JB, all working groups need to be staffed and given a budget. Again, a top-down iterative process involving lots of negotiation and vote taking will be involved. And, again, this process could benefit by starting with pieces already in place, or by copying similar pieces from another ELF society (or branch of government).

Discussion

Banks, insurance, health and other care providers, security services (any that wield lethal force), tax preparers, and most legal services that currently exist would be replaced by government functions. The need for some of these functions would be reduced by design, others would be standardized and regulated by public vote rather than by private whim. Minimum levels of care would be established for the health, education, and welfare of all ELF citizens. The ELF covenant that binds the government and private ELF citizens together would attempt to ensure the maximum freedom and support to each ELF that attempts to share an equal burden of society's responsibilities.

Present and future technology makes the structure of an ELF society possible. Our present societies are structured according the designs of the past, dating back decades and in some cases centuries. Registered ELFs that comply with the rules will be able to come and go pretty much as they please. Rogue ELFs will frequently be detained and frisked or questioned. Over time it will be in the best interest of rogues to become registered.

Ownership of guns will be less of a problem, but two changes will be made. Any weapon defined as a weapon of mass destruction (capable of the destruction of 10+ ELFs in < 10 seconds) will be confiscated, and possession of such will be a felony. Any lethal weapon is constantly subjected to the restrictions imposed by the ELF Covenant. This is likely to cause weapons to be present only in the home, a target range, or a hunting preserve, or on the persons of an ELF security worker. Better non-lethal weapon technology should be developed for protection and law enforcement. Registration of weapons for insurance purposes is voluntary. No other requirement for registration will be imposed, but the ELF Covenant effectively restricts certain aspects of use and possession.

A society in constant contention with itself is not a productive society, nor is it conducive to a pleasant life. The ELF Covenant is designed to correct many of today's points of contention. When contention does occur in ELF society, it is confronted head on between the parties involved. The ELF political system is oriented toward the grass roots, not the national level. National elections will occur infrequently, and the current system of political action groups, lobbyists, campaign financing, and a two-party system that can gridlock itself will largely become problems of the past. The three branches of an ELF government will be more effective at checking and balancing. More of the efforts of elected officials will go into the business of government, not into the business of getting re-elected.

ELF society is chartered with giving the homeless a place to go (and it can make sure that they get there rather than hanging out in parks, doorways, alleys, street corners, or sidewalks). Offensive behavior is only tolerated to the degree that other ELFs are willing to accept it, rather than contest it. Given that a job and "3 hots and a cot" are guaranteed to every ELF, it will be much harder to choose a life of crime or be a member of a gang or mob whose purpose is to extort the community in which they live. With cellphone technology, summoning the police and giving visual evidence of a crime in progress will be a factor in reducing crime. One judicial rule that doesn't always work out so well is that an offender has the right to be faced by an accuser. In most cases, this does happen. However, whistleblowers might come to fewer bad ends if the JD could remain as an intermediary. This would especially be true if the offender were a member of a corrupt organization or gang.

Banks

Every ELF and group has an account with the JB that registers the exact number of BUX it controls. Every transaction recorded for such an entity transfers BUX into or out of its account and the reverse with respect to the other parties in the transactions.

BUX and other things of lasting value are kept on the books at the JB. Thus, the entire net worth of any ELF, and the ELF society at large (and, as a matter of fact, a complete census) is an up-to-date statistic available at any point in time.

Loans of BUX from one account to another may require smaller transfers of BUX in the other direction over time, but BUX are not created in such transactions. Goods, services, and other things of value are created by individual ELFs and working groups. Services and perishables are not declared as assets. Durable goods and certain other assets show up on the books when they are created and registered. The government is not allowed to lend BUX. Its transactions within itself, and between it and other ELFs or groups, are like any other transactions.

The money supply is fixed during a yearly cycle. The value of a BUX account is a measurement recorded as a 64-bit floating point number. The initial societal sum is set as explained above (it could be in terms of the currency that preceded it). Each year the BUX is revalued so that 1 BUX is equivalent to a given amount of spending power. All transactions employ 64-bit floating point. At the end of the year, all BUX accounts are reset as a revaluation of total assets occurs.

If a transaction would cause an account to go negative, the parties to the transaction will be notified immediately and before the transaction is allowed to complete. The exception to this is that when user's fees are collected, an ELF's or group's account is allowed to go negative. An ELF (or group) with a negative account is insolvent and must work out with the JB a method for payment, or prevail in a contest to nullify the fee.

ATM machines will exist to dispense cash in the form of 1, 5, and 10 BUX notes, with a limit of 100 BUX per withdrawal. These machines will also accept any amount of cash in a deposit transaction. This facility allows a limited cash economy.

Every transfer of a ***durable*** value must be accompanied by the simultaneous transfer of (zero or more) BUX (completing that transaction). A description of the transaction is recorded to enable its reversal. Every transaction may be contested for up to one year and reversed by an agreement or court order. According to an applicable charter, one or either party may reverse the transaction (unless there are grounds to contest the reversal) for up to 24 hours.

A right to use or possess a thing of lasting value for a fixed period of time (a loan or a lease) is considered to be a perishable value. This allows physical possession of an asset to be separated from its ownership. When not spelled out in a charter or transaction, physical possession of an asset is the right of the declared owner. Labor and cash are considered to be perishable values exchanged without recorded transactions (unless a BUX account is involved).

Private groups may collect BUX into a private account and transfer BUX from it to other groups or individuals in the form of loans much like lending institutions do it today. Transactions and charters would spell out the rules for this. Failure to repay would constitute a breach of contract. The insurance arm of the JB would repay the victim of a breach of contract, and the JB would assume the responsibility of collection. While a contract may specify the repayment of the principle and interest for a loan, no contract may specify any other kind of future value. The value of a share of a group may go up or down with no guarantee of either. Groups and individual ELFs may own shares of other groups. The only ways to get rich are to benefit from value added in the various ways allowed. This begins with one's own wages, and extends to investments of one's own BUX. Every scheme to get rich (or poor) quickly that doesn't constitute outright fraud, will be possible in an ELF society, but most such schemes amount to complex ways of gambling and gambling is seldom advisable. The ability to reverse a transaction involving a BUX account confines speculation and gambling to private groups and corporations.

Economics

Some of the functions, now in private hands, would be transferred to the government. This would affect the economics of society in a number of ways. Insurance and banking are two of the most important. Public "safety nets" are the society's responsibility. These include (but are not limited to) insurance, health care, and education. These services would be available to any ELF requesting them, but alternatives (for big BUX?) could be available from private sources as well.

Because all assets are recorded with the government, any loss could be reported to the government and the asset's BUX value would be paid into an ELF's or group's BUX account. The JB has the right to investigate claims and pay only a fraction of the claimed amount if the claimant is partially at fault, or the asset is partially depreciated. If the JB determines fraud, a fine or punishment may be imposed. Fraud of a sufficient degree constitutes a felony. Detecting a felony is the responsibility of some ELF within the JB, and the failure to do one's job could also be the basis of a contest.

Society's banking responsibility is simply the management of all BUX accounts, both public and private. This amounts to recording and processing all transfers of BUX from one account to another, and reversing a transfer when required by a charter or the JB. Each transfer is noted with the BUX amount, a brief description of the thing of value transferred, and any other agreement attached to the transfer.

The government's budget and taxation are separate issues. Government spending must conform to a budget. There are only two ways that a deviation from this charter may occur: A shortfall that prevents a necessary expenditure may be covered by transferring BUX from another government account, or by the declaration of an emergency permitting a special tax cycle. Given the budget, money is raised by computing the percentage of private assets necessary to equal the budget. Each private BUX account is charged this percentage, and a lien is filed against all durables and inventories. See ***Taxes***, etc. above.

The ownership and management of a society's infrastructure would be the responsibility of the government. This includes power, water, sewage, streets, roads, and highways, airports, hospitals, parks, schools, and libraries. Services provided would include health care, education, insurance, limited banking, limited Internet, and limited public transportation. This falls short of socialism because all production and most maintenance would utilize only the private sector. If some activity not being addressed by the private sector were deemed in the national interest, the government could subsidize the private sector, but it would not create or own its own facilities for research and development, nor would it compete with the private sector.

International Trade

When two ELF societies trade, BUX accounts must exist in one of the two societies to enable the trade. An asset coming into a society involves the transfer of BUX into an account owned by a foreign entity. Such trade accounts may be within either society, but only one account is involved in any given trade. For example, society S (the seller) ships a good to society B (the buyer). BUX are transferred as part of the transaction. Typically, the seller would set up a BUX account in the buyer's society B with an initial balance of zero, and the buyer would transfer BUX into it in exchange for receiving the asset. If a buyer had no account balance in the foreign society, the seller would have to have an account in the buyer's society. If this is not agreeable to the seller, trade could not take place. The sum of all the accounts in society B owned by entities of society A, minus all of the accounts in society A owned by entities of society B, is the trade balance between those two societies. When society A has a negative balance of trade with society B, more of A's money supply is owned by "visitors" from B than is owned within B by "visitors" from A.

In terms of today's trade practices, if a US buyer wants a Chinese product, he transfers BUX into a BUX account that is registered in the US by a Chinese "visitor." The alternative is that the US buyer already has a BUX account registered in China by the "visitor" from the US, and he pays the seller from that account. If the trade balance with China is negative, there will be more money in the US accounts owned by Chinese "visitors" to the US than in the Chinese accounts owned by US "visitors" to China. Assets that are shipped to a foreign society disappear at the end of the year simply by not being declared, but BUX never leave their society of origin, they are transferred from one account to another within it.

When the aggregate assets of a society decline in value, their public BUX account does likewise, but the sum of the private BUX accounts remains virtually the same. More of the total supply is now owned by foreign "visitors" and the society itself becomes "poorer" as the result of a negative trade imbalance. Notice that foreign BUX accounts are rescaled like any other BUX account for tax purposes at the end of the year. On Day One, if the U.S. were to become an ELF society, it would have a huge number of its initial BUX in Chinese owned accounts.

Contracts

A contract is a "good faith" promise to deliver goods or services for a specified number of BUX. The transaction will record the goods or services to be delivered, the BUX amount, and the contract expiration date. Buyer pays seller the full amount, which is then "frozen" as is any transaction for an asset. After delivery plus 24 hours, the buyer has either "signed off" by default, or registered a contest with the JB to halt, freeze, or reverse the transfer.

Social Services

ELF society has the responsibility to guarantee ¶ 4 and ¶ 5 of the bill of rights. Part of this is handled by insurance, as explained above. All other needs and compensation must be spelled out in charters. But, at a minimum, health care and institutions for care giving must exist. If these facilities are managed by the EB, the public gets a direct say (via their rELFs), in their quality and quantity. Notice that user's fees may be charged for these services.

In order to guarantee ¶ 4 of the bill of rights, an app on the government platform will enable ELFs and Groups to contact each other regarding employment. This will span the gamut from resumes and exotic jobs to ride pickup or package delivery. An upvote/downvote system will be part of this.

Education

Education serves two functions. It is the responsibility of ELF society to provide it, but it may be augmented by private interests. All ELFs from the age of two until they are self-supporting have the option to attend a school or daycare center. The first, and primary, function of education is to enable each ELF to pass the certification test required to declare ELF recognizance. The second function is to provide an opportunity for an ELF to learn skills and acquire knowledge.

Daycare must be available 40 hours a week to all those who qualify. The early curriculum of primary education should be part of daycare. The content of this education is restricted to learning to read and write the official language, and how each ELF relates to society and the ELF Covenant.

Reading will involve a list of books with secular and scientific subjects. Writing will involve essays that indicate some understanding of what has been read.

Subjects advocated only by private interests will not be taught at public schools and daycare centers, but attendance at a privately financed institution that teaches any subject is allowed as long as the primary knowledge required to gain ELF recognizance is learned at a rate comparable to ELFs that attend public institutions.

ELF Markets

ELF markets are physical or Internet access points that allow ELFs with common interests to contact one another. In particular, ELFs that wish to buy or sell an interest in some private working group. These markets allow perishables and assets of all kinds to be bought and sold. Each website or physical location that facilitates BUX transactions must have a charter that can be accessed and read using a standard web search. Each such charter must declare who its owner is, and likewise each group or ELF that owns such a website must declare it as an asset under control.

Stock Markets (Gambling)

The logistics of any form of gambling, or trading on a stock market as it now is done, would be as follows. A private ELF would transfer BUX into the institution’s account. The institution could conduct games of chance according to any rules it sees fit to establish and which are accepted by its patrons. At the (figurative) end of the day, patron ELFs would transfer BUX out of the institution’s account, and back into their own. The asset value of a BUX account is its final value at the end of the year. Again, each institution that facilitates BUX transactions must have a charter that can be accessed and read using a standard web search.

A group may be chartered to buy shares of stock in a list of companies and set itself up as an ETF (exchange traded fund). Likewise, a group could buy and sell its own currency for the purposes of trading or gambling. State supported BUX accounts are insured and safe. The purchase of a share of any company may typically be done in a single transaction with that company, stock exchanges or mutual fund companies may find ways to add value or appeal, but the average ELF would probably avoid them.

Contested Matters

Every working group in the EB reports to a rELF. That rELF reports to a rELF on the next higher tier, and so forth. Any grievance with any worker ELF, or public working group is reported first to the rELF in charge. If that doesn’t produce a satisfactory result, it may be reported up the chain of command, or if necessary to the JB as a formal contest that may lead to a trial.

Felonies are contested automatically by the JB when brought to their attention by any ELF (especially an ELF working a security job). Procedures for handling such matters are part of a charter written by the EB. Otherwise, any action or state of affairs may be contested by any ELF. The only basis for a contest is the violation of a right or charter, or the failure to meet a responsibility (as set forth above). A contested matter is registered with the JB. A matter involving some urgency, that could be contested, may be reported as an emergency, and the EB has the responsibility to respond to it (including assistance to register it with the JB).

Matters involving ongoing force or threat, that might require force to curtail, are handled by the EB (including assistance to register it with the JB). All other judicial matters are handled directly by the JB. The use by an ELF or group of the EB or JB is the use of a common service. The misuse, or overuse, may require fees to be paid. The cost of trials and appeals are primarily paid for by the losing party.

The following are some examples of how legal matters (and laws) are handled in an ELF society. It should become clear that a complex tax code, and a large number of laws are not only unnecessary, they are not even allowed. Judicial matters are decided on a case-by-case basis (with the JB’s discretion to consider the history of similar cases).

Abortion

At some point, an entity becomes an ELF. Before an entity fits the definition of an ELF, it has no rights. Any case of abortion must be contested on some basis as given by the definition of an ELF and the ELF Covenants. This makes infanticide and abortion after a fetus is independently viable a felony. Abortion before that point would be difficult to contest. This matter and other similar matters must be contested on a case-by-case basis.

Gun Possession

No specific law is needed for gun control or possession. However, the use and possession of any threatening weapon may result in one or more ELFs invoking their right under ¶ 3 of the Bill of Rights. This could restrict the possessor of anything threatening from entering any establishment, or even a much larger area. All it would require would be a sign indicating such a restriction. The possession of any lethal weapon is circumscribed in the section titled "ELF Responsibilities."

Smoking

Smoking or any offensive production of noise or odor is subject to ¶ 3 of the Bill of Rights. Again, signage is sufficient.

Drug Use

On a case-by-case basis, a contest could be initiated if drug (or alcohol) use led to a specific or chronic problem that is an actual violation of a right or responsibility. When a licensed privilege is involved, the charter that permits the privilege may spell out specific rules and regulations that are part of a required agreement.

Theft

The JB is able to curtail theft in several ways. The first area is the way bank accounts are handled. No private group has access to the BUX accounts of others. All transactions in society are recorded as an exchange (reversible for up to one year) of BUX for a thing of value (described by the transaction). Thus, an account can't irrevocably be emptied out, even when held by an equal partnership. The use of cash is limited, and cash is not insured. All other transactions are handled by the JB on behalf of both buyer and seller.

Another area is personal identification. What guarantee is there that the bank account of one group or individual isn't used by another? Every ELF and group registers a username that is associated with a public key. Every ELF possesses an RFID for identification. If this is lost, or stolen, a duplicate can be issued, and all illegal transactions reversed. A record within the JB is the only place the private key is stored for a group or individual ELF. No two sets of private key, public key, and username have any component in common (guaranteed by the JB). The security of every transaction is guaranteed by the use of this public/private key.

Each ELF and group may own an inventory that has a lasting value, even if its components do not. A store, for example, may have an inventory that is permanently valued at 10,000 BUX. It is stocked and depleted over time with a history of transactions as evidence. When an item or items are stolen or destroyed, an insurance claim may be made to the JB. This claim may be investigated by the JB, or the EB, and the culprits brought to trial. The claim is paid immediately. It may be revoked (and a fine imposed) if proven fraudulent.

Names, phrases, or logos may be registered up to one year in advance of being used, and expire one year after their last public usage. This entitles the registrant to the exclusive use of the name, phrase, or logo when associated with a particular use, or any remotely similar use. Telephone numbers, usernames, and domain names fall into this rubric.

The definition and theft of intellectual property is perhaps the toughest problem to solve. This involves copyright and patent laws and their infringement. ELF society will take a different approach. The concept of plagiarism that diverts actual profit away from an author is introduced. The ELF covenant must be used as the basis for any contest. Work product may be described and registered as an asset under control. Proof that a copy of such an asset was illegally obtained may be entered into evidence as part of a trial. Proof that independent development took place to arrive at a substantially similar result may also be entered. When a book is involved, and large portions are extremely similar, proof of the earliest drafts entitle the author to the work. Reverse engineering or copying that diverts substantial sales away from the original creator, may be entered as evidence. Authorship entitles a work. Only the title holder may assign the work to another entity.

Possession, use, or acceptance of terms and conditions may be entered into evidence, but they not sufficient to determine the outcome of a trial (they do not constitute a charter). There has to be some evidence in every trial of a violation of the ELF Covenant or an actual charter.

Fraud

Fraud is any action, communication, transaction, or charter that involves willful deception. If a reasonable ELF would protest an action when the deception has been revealed, the effects of that action are presumed to be a violation of the ELF Covenant. Society has the right to detect fraud when feasible, and the responsibility to notify victims of it.

Statements that are a distortion or contradiction of the truth may also constitute fraud. The seriousness of making fraudulent statements, or statements whose primary purpose is defamation, is in proportion to how widely such statements are propagated. The degree to which the truth is compromised, and any penalty to be imposed, are subjects that may be contested and determined by a trial.

Society also has the right to enforce an Internet protocol that demands that every packet contain a complete and accurate record of its original and final IP addresses, as well as a valid sign-on username. This is intended to allow tracing the origin of any packet sent over the Internet. Attempts to abrogate the rules for packet routing (such as necessary to set up a "dark" web) must be sought out by the JB and terminated when detected.

Any ELF has the right to request that all, or any type, of communication to it, from any username or IP address, cease. A copy of such a request to a (special) .gov address should be sufficient to contest any such matter. Everything that applies to the transport and display of information packets also applies to the transport and delivery of physical packages.

Gangs and Crowds

If a gang or a crowd poses a credible threat of a violation of a right or responsibility, it may be the basis for invoking the right guaranteed under ¶ 3 of the Bill of Rights.

Offensive Persons

The same as above, based on an actual violation of a right or responsibility. A formal request to "cease and desist" may be declared when any behavior is considered offensive, and the intent to formally contest the behavior is publicly stated to witnesses present, or to a 911 operator. This right stems from a reasonable interpretation of ¶ 3 of the Bill of Rights.

Professional Groups

The following professional groups are organized by society in similar ways. All are managed by the Executive branch. For each organization listed below (except for those involving security services) there could be private counterpart organizations. No public funding of a private counterpart would be provided. All such organizations would be fully funded by their users, and would likely be available only to the very wealthy.

All of the Professional Groups listed below would be staffed by government employees who are paid according to rank. Rank can be on three tracks: Supervisor, Technical, and Enforcement. Each track has 16 levels. The pay on each level is equal throughout the government and is set by the market forces necessary to staff the positions. Each employee has a job description that is part of a charter that fully describes a job position.

Schools

A school is an institution designed to provide a learning environment for young ELFs. It brings students into contact with teachers and typically with each other as well. It may involve a physical classroom to contain a number of students and their teacher, or it may be a website that facilitates similar communications electronically.

ELF recognizance requires a minimum of 8 years of schooling. However, all ELFs are entitled to 16 years of schooling. The curriculum is standard for the 1st 8 years, and varies beyond that. Each class is attended for 45-120 minutes 3-5 days per week for 4 months of a year. A final test is given to determine a pass/fail grade. Classes may only be attended if an entrance test or a prerequisite class has been passed. Although standardized true/false and multiple-choice questions are part of any final test, essay questions graded by the judgement of the teacher are also a major part of finals. When a failing grade is given, a student may appeal it to the judgement of two out of three teachers that teach the same class. The essay questions that will be asked are presented to the student at the outset of the course. All final tests are made a permanent part of a student's records, and may be viewed with the student's permission.

Public Transportation

Access to schools, health care, day care, and assisted living, require transportation. If this need cannot be met privately, society has the obligation to provide it as part of its obligation to provide these other services to the public. A fleet of shuttles and emergency vehicles must be administered, maintained, and put to the most efficient use. This requires supervision, drivers, and maintenance.

Health Care

Individual health groups would be chartered as clinics and hospitals to carry out a society's health needs. This does not forbid the existence of non-government health providers. Large medical facilities would typically be government owned. Smaller doctor's offices would typically be privately owned. The need for medical attention goes from an ALF who needs no attention (other than an occasional technician) to a chronically ill BLF (biological life form). Thus, a non-prohibitive user's fee should be assessed for the use of medical facilities.

On the national level, a Drug Authority would be established. This group would be the go-between for health care groups and private groups that develop and/or manufacture drugs. It would perform drug testing and validate new drugs. It would put out descriptions of drugs that it would like to see developed. It would set the price of a drug based on development costs, expected sales of the drug, and a goal that the developer might break even after one year. The price would remain in effect for 10 times as long as the drug took to develop. During that time, the developer would have the exclusive right to license others to manufacture the drug (or prevent them from doing so). The Drug Authority would record the unique characteristics of each drug and ensure that it was not a copy or minor modifiction of an existing drug.

Assisted Living

Groups in this category would manage a building or complex of buildings, would have at least one supervisor, a maintenance crew, and in some cases a kitchen staff, a cleaning staff, a caregiver staff, and possibly an enforcement staff.

Prisons

Prisons will be designed with public access to the outside, inmate confinement on the inside, and buffer zones between them. Prison populations will be separated into violent criminals who spend 18 hours a day in single person cells, and non-violent felons who live in squad rooms. Each inmate will have a cellphone able to contact other inmates, the prison staff, and a limited number of outside IP addresses. Each inmate will be entitled to a job within the prison at an SRP, except violent felons who have exhibited violence within the prison. A felon's 3 hots and a cot will be billed at cost as a user's fee, and repaid out of its wages or after sentence has been served. There is no parole. Sentences may be extended for any breach of prison charters or the ELF Covenant during confinement. Exercise areas and game rooms are provided as privileges to be earned by good behavior (or by the payment of a user's fee?).

Should an ALF become a felon, it may be confined in place outside of a prison facility, but it will also have the same limited access to the Internet that a felon inside a prison would have.

Halfway Houses

These range from apartment complexes to community living (separate rooms for one or two ELFs with common areas). They are available to the otherwise homeless, or to released convicts. Inmates are free to come and go. An SRP job is available to each inmate. As in prison, the 3 hots and a cot provided are billed to the inmate's BUX account. The primary job of an inmate in a halfway house is to move on, get a normal job and place to live. There will be staff to assist in this.

Caregivers

Here, the facilities are a mixture of prisons, halfway houses, and hospitals depending on the needs of the inmates. ELFs lodged in caregiver facilities have special needs. They may not be capable of the norml activities of daily living for one reason or another.

Retirement Homes

When an ELF becomes too old to earn a living and too poor to pay for the basic necessities of life, retirement communities are available. Physically these are a cross between public housing and caregiver facilities.

Public Housing

The minimum standard housing allowed (think zoning requirements) is public housing. Any dwelling not exceeding the standards for public housing is condemned, demolished, and cleared to be a vacant lot. Public housing is an apartment complex of living spaces rented by the day, week, month, or year. Similar facilities may be publicly owned. Only if adequate private facilities do not exist, will public housing be built. Public housing has a staff as described above under ***Assisted Living***. Rent for public housing is set at the minimum to break even, and is collected as a user's fee.

Daycare

Again, this service would involve user's fees. Caregiving would be provided to young biological ELFs when there was no reasonable alternative. This could involve the caregiver going to the young ELFs residence, or shuttling it to the care giving facility.

Fire & Emergency

These professionals would be trained to handle fires, explosions, damage to infrastructure, and threats to any of these. They would also include EMTs capable of working with emergency health care facilities. Their charter should include as many duties as their time permits, but none that would compromise their response time to emergencies.

Armed Forces

In the past, and even in the present, human societies have done atrocious things to one another. To maintain an effective level of security, an armed force has been required more than once, and appears to be required at the time these thoughts were written here. Today, the largest societies appear able to provide their citizens with a reasonable level of security. Many of the smaller ones cannot, and even the largest societies are unable to effectively aid the smaller in many cases.

A society must either rely on the protection of another society with the capability to protect it, or it must establish an armed force of its own (at a huge expense). The concept of mutually assured destruction (MAD) is not as “mad” as it sounds, as long as it remains a deterrent, and the bluff is never called. Part of maintaining security is to exercise threats when needed to stop events before they get out of control. Thus, there must be some intelligence gathering, both within and outside of a society’s territory, to allow the recognition of a threat in advance. Society should have the right to monitor the exchange of assets under control, but not protected exchange of information. It should be allowed to monitor the existence at a point, of every ELF within its commons, but not infringe on their privacy.

The police in ELF society may be called in an emergency, but only to assist private ELFs in pursuing complaints. Police should have no charter to take forceful action except in the case of a felony. Each action taken by the police should assist a “client” ELF. Nevertheless, the police should be responsive and capable of following up a complaint, and a client ELF should be able to file a complaint against the police for doing an ineffective job.

ELFs are trained and hired as professionals into the military and police forces. Each group of the military and police are managed by a group in the Executive branch. Each armed group is a strict chain-of-command hierarchy. When an order is given and is not obeyed, it must be adjudicated by an equivalent tier of the Judicial branch. One or the other, neither, or both, of the disputants may be dismissed from their job. This extends to an impasse or altercation between any two ELFs serving together under arms, in which case fines or punishment may also be involved.

Security Services

ELF society has the responsibility to guarantee ¶ 5 of the bill of rights. Part of this is handled by local security (police forces), part by regional security (a national guard), and part by a national armed force that may even be required to leave the confines of national territory. When special needs require (or request) additional security, it may require an assignment of security resources that are paid for by an individual ELF or group (via a user’s fee). A private security group, capable of lethal force, would not be permitted (however, groups of big guys with tasers or mace would be).

Possession of weapons of mass destruction is restricted to groups within the government. This implies that certain private groups wishing to use explosive devices, for example, must engage and pay a government group to handle an activity like demolition. Manufacture and transport of explosives or other lethal devices must be physically monitored by a government representative (to a greater degree than meat and health inspectors). In cases where public safety is not compromised, the use of such devices may be handled by the user, but an inspector must be present. A chain of custody of all explosives must be maintained.

Airport Security

Free access to any airport gate is permitted. Any unaccompanied luggage is immediately confiscated by airport security. Boarding involves separating each passenger from any piece of luggage larger than a small carry-on, and storing luggage where it cannot be accessed by a passenger. Small carry-on items are screened at point of aircraft entry. Possession of items that could reasonably permit taking over an aircraft would constitute a felony. The flight cabin would be locked. All crew would carry mace on a tether to control aggressive passengers. Access to any crowded facility would be by a point of entry requiring an RFID.

As a case in point, the TSA (today) has the authority to search and confiscate to a degree that violates ¶ 1 of the ELF code of responsibilities. It would not be permitted to impose an inconvenience and expense on the entire society to mandate less total safety than the total inconvenience of time and expense. In such a tradeoff, the tipping point is 1 ELF life saved (or the time equivalent) = $1M BUX.

General Security

All public records must be recorded in at least 3 (an odd number) of physical sites. The data at all sites must be identical. In the event of a miss-compare, the minority records must be made identical to the majority records. There are 3 aspects to all recorded data: Who gets to record and change it? Who gets to view it? How is it authenticated? Storage of public data is free of charge to, and may be viewed by, all member ELFs. Private vaults and protected “public” keys may be rented by any ELF or group.

Private storage of records has no legal status. All groups in a society must use public information storage for any records that are to have legal status. Record storage associates a digital vault with a unique public key. Each ELF and group in a society may have a list of registered aliases. These pair a string of ASCII text with a public key. Given a public key, any unencrypted file in a public vault may be viewed. Given a private key, any file in a public vault may be written, changed, or encrypted. In effect, this means that private keys are embedded in public record storage. Therefore, society must guarantee their secrecy. Public, private key pairs may be swapped with a different set only with the judgement of a court. This must be personally (physically) overseen by three officers of the court, and "impossible" to perform by anyone else.

If a private key were lost or stolen, its public key counterpart could be used to freeze access until its legal owner could be established by the outcome of a suit and the keys swapped by a court. A stolen key is one possessed by two contesting ELFs or groups. Here there is no problem accessing the private key. However, a lost private key is a total loss of assets if that key cannot be verified from another source (such as a physical chip or physical description of the ELF). Such assets are held by the JB until a contesting ELF prevails in a trial to obtain them.

All votes in all formal elections will be cast by an authenticated login to a public website. They will be verified and counted by the machines used for public record storage, and the outcome announced on the original public website.

Cryptography

Modern cryptography is based on pseudo random number sequences. Given a Large Prime Number, a seed (any number), and a block size (a number = 2^n), a sequence of (pseudo) random numbers may be generated by letting the next random number (and the next seed) be the result of multiplying the LPN by the previous seed, and taking the modulo of the result by the block size. This will produce a sequence of LPN unique block-sized (pseudo random) numbers.

If a clear text is broken into blocks equal to the above block size, and each block is XORed with the next random number, the result is a cyphered text. When the cyphered text is treated to the same random number sequence, the clear text will be restored. To break this cypher, all possible LPNs and initial seeds would have to be applied to the cyphered text, and the results of one or more consecutive resulting blocks tested for intelligibility. This is called symmetric key encryption. If both parties know the parameters, either may encrypt or decrypt each other's messages.

When 2 LPNs are multiplied together, a significant compute time is required to factor out the original numbers. A public key consists of the product of 2 LPNs, and a private key consists of either LPN factor (one LPN is trivial to compute if the other is known). If each private key is known only by one individual, and each public key is known by everyone, text may be enciphered so that its intended recipient is the only one who can decipher it. Another technique allows a document to be signed, using a private key, to prove that it is authentic (signed by the originator, and unaltered since signing).

Public-Key

Part of a society's infrastructure is providing each ELF member "state of the art" public-key cryptography. This permits authentication of all recorded charters, and secure access to, and management of an ELF's or a group's assets under control. The public key would essentially be an ELF's identification (user or login name, email address, etc.). It would be a matter of record that could be viewed by anyone. The private key would exist in one copy on a chip protected by (or embedded in) the ELF who owns it. The private key would be used to authenticate a login or transaction.

The procedure for redacting a charter is to encipher portions of its text using a public/private key pair assigned to the document by the JB.

Budgets

Each group on a lower tier would submit a budget for approval by their parent group on the next higher tier. When the budget process is complete, and taxes are raised, a comparison can be made between each group's budget and the taxes apportioned from those represented. In most cases, this will show a positive or negative discrepancy. Society should attempt to correct these discrepancies by allowing larger budgets to groups serving the overtaxed, and smaller budgets to groups serving the under taxed. A group, or individuals within certain groups, should be chartered with this mission.

Calendars & Holidays

Our current method of recording dates is outdated. Rather than shifting to some totally new and inhuman system, my suggestion would be to take some small steps to repair its most blatant defects. A number of theories have been put forward to explain why virtually the whole world has, for centuries, used a 7-day week. But, there you have it, a 7-day week is almost universal. There has been a requirement in every culture to have a calendar that corresponds to the seasons and the phases of the moon. Neither of these is possible without "adjustments" of different sizes every year. However, there is no good reason to have 12 months with different numbers of days in each, nor to have the same dates each year fall on different days of the week. So, let's consider the following calendar system.

A year would continue to conform to the orbital period of the Earth around the sun, but would consist of 13 months of 4 weeks of 7 days each. This puts us at 364 days per year, whereas a year is actually 365¼ days. This means that we need an odd day or two to make up the difference at the end of the year. The year would begin on the day that included the latest sunrise in the northern hemisphere as measured at longitude zero. Latitude, longitude, hours, minutes, and seconds would continue to be measured as they were in 20th century England. The 13th month will be called Festivus. Every date with respect to the calendar for the Common Era maps into the corresponding date during the 1st year of the changeover.

The 1st day of the week (Day 1) is Monday. Sunday is considered Day 0 of the week. At the end of the year there is a long weekend (the last days of which are Sunday). Any leap seconds that accumulate during a year are added at midnight (GMT) of Last Sunday. Thus, some might consider it to be the 1st day, but as we continue to note, when things are listed, either 0 or 1 may be considered the “first.” In this case, days of the week are cyclical, like the least significant digit of a base 7 count. The notion of zero also accords well with the idea of Sundays being repeated at the end of the year as “null” days just to fill out the full solar cycle. Sundays are non-working and non-legal days, so the last days of the year are less necessary as calendar appointments. Every Sunday, however, does equate to a numbered day of the month.

Daylight savings time, and time zones themselves, will be abolished. GMT will be observed everywhere. The concepts of morning (0600), noon (1200), evening (1800), and midnight (0000) will be tied to an adopted local time. Calendar dates (if a location is designated) and days of the week are tied to the GMT date when local midnight occurs. Otherwise a date and time assumes GMT. Local time may be expressed as an hour and a fraction offset from one of the 4 sun times.

All holidays will be celebrated on a Saturday. New Year’s Eve begins the evening of the Saturday before Last Sunday, and New Year’s celebration continues through Last Sunday. The 3rd Saturday of December, the 20th, is Thanksgiving. Christmas will be celebrated on the 3rd Saturday of Festivus (again, the 20th). Birthdays are on the same day and date each year, except those that occurred on the last Sunday of a leap year. Those will occur on Last Sunday every year. The last days of Festivus are ***Sunday*** the 28th and 29th, and on leap year, Sunday the 30th.

Travel and television schedules would be universal; local times (still expressed as GMT) would be at odd, but constant, hours. Perhaps the only thing interesting about world clocks would be the time noon occurred in various cities around the world, or when midnight occurred around the world on New Year’s Eve. Saturday might begin in Denver at 6:00 a.m. GMT, and in San Francisco at 7:00 a.m., but a new day of the month would always begin at 0000 GMT everywhere. Thus, instead of the date and hour being different everywhere, it would be the same; but each location would have its own local time of midnight, morning, noon, and evening. Every city and town could declare its local time offset from GMT (possibly joining with nearby populations), so that 6:00pm (1800) local time was a precise value that specified the onset of evening.

Historical Dates

Human history is measured by the number of revolutions the Earth has made around the Sun. A date in history might consist of one or more sets of three digits. Every ***legal*** document could use only 3 digits to give it a past or future date. This implies that legal documents could reach no further than 500 years into the future or the past from the date they were signed. Any value in the range (present year minus 500) … (present year plus 499) refers to a valid date referenced during the present year. All dated documents, or documents containing dates, have to be rewritten before 500 years have elapsed (otherwise a past date becomes a future date and the document automatically expires). For example, the year as of writing this is 018. If I deeded a property today, its deed would expire on or before 517 (2517 CE). Today, the year 518 refers to the year 1518 CE (500 years ago). A future date in a document written today might be 244, referring to the year 2244 CE, over two hundred years into today’s future.

The date of a document is necessary to interpret any dates written into the document. However, since most lifetimes are a hundred years, or so, there would be no ambiguity of dates written into documents anytime during a normal lifespan. When the current date is the document’s date plus 500 years, the document has expired and may be deleted and archived. Before this happens, it will hopefully be rewritten and retired early. All documents and events outside this 500 year range would use six digit dates as explained below.

In the long view, eternity will be considered to be a span of 1,000,000,000,000 years. Just to pick a starting point, we’ll say that the current year is 500 500 502 018. All dates “forever” can be expressed as an ordinal count (negative numbers are not necessary to reference a year in the past. This would only be necessary if the date were written relative to the Common Era, or the current year. Any date with fewer than 12 digits assumes the leading digits of the current date. Any date with 3 digits and a G, M, or K suffix is an estimate that assumes the less significant digits don’t matter.

Thus, if the universe actually “began” about 15G years ago. The date of its birth would be 485G. Do the math. Dates are always written with 3, 6, 9, or 12 digits (even if they begin with a zero). A space separates groups of 3 digits. In context, this should disambiguate other numbers from a reference to a year. A story set in the future might give the year as 325 (equivalent to the year 2 325 in current terms). A reference to the date that the first ALF leaves the Milky Way might be given as 852K—a reference to the year 500 500 502 018 + 350 000 (years into the future).

Points to Integrate into Part 3

Voting Procedures

Public officials could either be elected from a local district into Tier-1 and rise up through the government from there, or they could mount a campaign to run against a rELF on one of the Tiers, 2-4. This would require a petition with 10% of the votes to be cast calling for the vote, and it would require greater majorities than a simple 50+%. This type of campaign would cost serious BUX. However, the normal way to be voted into office would cost little if any BUX, since it would only involve gaining the support of a single voting group (on one of 4 Tiers). This would require networking directly with the voters.

After a rELF has been elected to Tier-1, some of the Tier-1 rELFs may be elected to Tier-2, some of those to Tier-3, and 1-3 of those to Tier-4. A rELF elected to any Tier may only be recalled (replaced) by a vote of the voting group that elected that rELF to that Tier (not necessarily the Tier-0 voting group that originally elected the rELF to Tier-1).

This system would ensure a much more stable government, and one not prone to the corruption of today's government. There would seldom be big campaigns, elections would not occur all at once across the government, nor would there be regular periodic elections. The system would simply not be vulnerable to the power of big BUX as long as quid pro quo arrangements between rELFs in government and private groups or citizens were forbidden by charter (and enforced!).

Separation of Powers

rELFs in each branch of the government would be elected in the same fashion. No branch would have any power to dictate to another branch who their representatives would be. A simple vote could replace any rELF in government, either via a general election, or via the voting group that elected that rELF in the first place. There would be no lines of succession. If a rELF were no longer able to serve, an immediate vote would choose a replacement. There would be no need for articles of impeachment. General elections (popular votes) would be few and far between. Votes taken within voting groups would happen all the time: some would be for issues; some would be for candidate rELFs. These would happen independently in each of the three branches of government.

Lateral Arabesques

What if a rELF in one branch of government wants and/or is wanted to transfer into another branch of government? There should be a facility for this. The hierarchy of rELF positions is virtually fixed. The rELF involved represents a collection of voters on Tier-0. A single voter elects rELFs to all three branches of government. A general election (with a majority over 50%) can be called to effect the movement of a rELF from one branch to another. However, this would leave a vacancy in one branch, and an extra rELF in the other. Vacancies can always be filled from below. Recalls can also be the subject of a vote by the voting group on the tier just below the rELF to be recalled. These two types of special elections would propagate the object rELFs in opposite directions. When a majority of Tier-0 voters does effect a "lateral arabesque" it obligates the appropriate voting groups (on the tier below) in the respective branches of government to promote and demote two rELFs of their choice to restore the proper numbers. This would propagate all the way to and from Tier-0. It's always possible to send a rELF back to Tier-0, but it might not be possible to find a candidate willing to become a Tier-1 rELF. This creates an open "job position" until an ELF steps forward and is elected to fill it.

Voting Group Membership

A non-gerrymandered geographical area defines a collection of Tier-0 voting groups. It includes all and only the physical addresses of every ELF within an area. Such an area must include a minimum of the tier ratio of ELFs. There is no maximum number of ELFs in a Tier-0 area. Thus, membership in a Tier-0 voting group is tied to an ELF's registered address. The area may be expanded, shrunk, or merged with an adjacent area (conforming to gerrymandering rules) by calling for a group vote. Typically a Tier-0 voting group would be close to an even multiple of the tier ratio, and generally from 1-8 times that amount.

The rules for Tiers 2-3 are similar to the rules for Tier-1, but scaled up 1-2 levels. Thus if a voting group on Tier-0 were a neighborhood or voting district, a number of such groups would combine to form a city or county sized area on Tier-1. A number of these would combine to form a state sized area on Tier-2, and a whole nation on Tier-3.

Record Keeping

{see elsewhere and merge}

Every group and individual ELF has a personal record file all of which is viewable by anyone upon request, however there are various levels of redaction applied, and all access is via extraction programs that produce temporary files. What goes into this file is all financial transactions, all applicable charters, all votes cast (with links to any relevant issues), and any personal diary, record, photo, or video that an individual wishes to keep. The right to view a file in its entirety belongs to the ELF or a member of a group whose record it is, and to a panel of 3 judges in the JB (according to charter and on public record).

Truth and transparency are paramount issues with respect to any ELF serving as a rELF.

Any communication to 8 or more ELFs, or a communication that is passed on from one ELF to another, is considered a public speech. The rules for Free Speech are limited in the case of Public Speech. Any speech that violates the rules may be contested. If successfully contested, a transcript or description (verbal, photo, or video) is entered into the originator's file to be viewed by anyone, and all other copies of the original communication (but not necessarily references to it) are destroyed. Additional restraints may be imposed upon the originator according to applicable charters (including rELF disqualification).

Ideas for Development

{Does any of the following fit anywhere??}

How does ELF society get money out of politics and get lies out of the media, both socially and formally?

When a communication goes to fewer than 8 people, almost anything is tolerated, but it is also valid to invoke ¶ 3 of the ELF covenant due to the content. However, when a communication goes to 8 or more people, and proportional to the number of recipients, it should be held to a higher standard of content. When successfully contested (based on falsehoods and/or slander and/or incitement) the source of the communication may be censored, and/or the sender fined, or imprisoned.

If a Tier-0 voting group elects a rELF to Tier-1, and that tier elects the same rELF to Tier-2, and so on to Tier-4, then each lower tier gets to elect another rELF to replace the one removed from the next higher tier. A rELF can only be removed from a higher tier by a vote on the tier just below it. All rELFs on a given tier serve at the pleasure of the voting group that put them there. They represent those particular voters.

The rELFs on each tier of a GHG may write charters for the working groups beneath it. One or more rELFs must be appointed to manage such groups. Other ELFs may be hired as workers or managers in such groups. Line of duty is from a tier-4 rELF to a rELF manager of a working group, through a chain of worker-managers to a worker. Any action taken must conform to a charter or specific directions from one's manager. Any illegal act is presumed to be the fault of the ELF taking it, unless it can be shown more likely that it was taken under orders, in which case the ELF or rELF giving the order is responsible. When a doubtful action is ordered, a formal request may be demanded and put on the record. No penalty may be imposed for making such a request. All groups chartered to take action by force, such as police forces or armed forces, must be working groups of the EB. The manager of an armed group is always a rELF. Lethal arms are always and only wielded by a uniformed worker ELF or the uniformed manager of such an ELF.

Socialism vs. Capitalism

In ELF society, the manufacture or production of goods is performed by the private sector according to charters. Where construction projects that produce airports, highways, dams, power or communication grids, etc. are all considered goods, ownership of these assets passes to the government. Special (occasional) services are also provided by the private sector, but regular and ongoing services are performed by government owned working groups.

The government neither owns nor provides public housing. {rethink??}. The EB does own and provides free of charge care facilities for children, the disabled, and the elderly. The EB provides free health care and education. The private sector is free to enter into any of these activities on its own. {more??}

Employer vs. Labor

An employer may be an ELF or a Group. The provider of labor is always an individual ELF (or hive), but never a group or a collection of ELFs. When an employer is a group, every ELF providing labor to that group reports to one of the owners of that group, or to another ELF providing labor to that group. One group may own a scheduling platform to sign up individual ELFs and coordinate their labor with different groups that require their labor. All charters that direct the supply and demand of labor are between a supervisor and an individual ELF. Such charters may have the duration of a few minutes, or they may renew over periods of months or years. Such charters spell out what an ELF can expect of an employer, and what an employer can expect of an ELF working for it. A one-time charter may be recorded within a single transaction (with links to the BUX accounts of both the group and the ELF); jobs that are expected to repeat are recorded in a charter filed within a group's database with links to the JB and to the ELF employee.

Education

Education is broken down into two 8 and two 4 year intervals:
Ages 1-8 should result in the ability to read and basic social skills.
Ages 9-16 teach the "subjects to be tested" (listed below, & also called the GED).
Ages 17-20 teach introductions to various fields of specialization, or allow various working apprenticeships.
Ages 21-24 complete a field of specialization and/or a working apprenticeship (possibly including the 2000 hour obligation).

Subjects to be tested at age 16 will include reading, writing, math, science, ELF society and government, history, and world affairs. Advanced specialties will include scientific disciplines and all of the technology and literature necessary to support the ELF culture and infrastructure. Full personal recognizance is granted only after passing the GED.

Internet & Transport

The JB owns and operates a secure Internet surrounded by a firewall that connects to any private sector Internets that may wish to do so. Access to the public Internet is free. Access through it to private Internets may involve subscriptions. Each ELF has access to the public Internet via PDAs (personal digital assistants) located in public libraries. Ownership of one's own PDA or a vehicle for transportation is a private sector luxury (above the subsistence level), but self-driving vans may be requested to pick up and deliver an ELF for free anywhere within an urban or suburban area. In large urban areas, mass transit systems will be free of charge. Self-driving busses will be available for scheduling the transport of children to school and other ELFs when a minimum load is established by a website signup.

Compelling Reasons for a Societal Design

Google *Republic Lost* or *Utopia for Realists* or *What's the Future and why…* to understand the groundwork that supports the changes that justify replacing today's governments with ELF societies. The constitution, bill of rights, and amendments that began the U.S. government worked with the status quo of the time to form "a more perfect union." Likewise, today's status quo could be melded with the ELF Covenant and my description of an ELF government to form a modern improvement.

What dynamics affect a society? The combination of individual efforts and expressions, those of private groups, those of government agents and groups. All of these are subject to the ELF covenant and any charters that apply. The ideal is for the dynamics of private citizens and groups to affect society, leaving only specified exceptions to the government.

The primary duty of the government is to referee disputes between citizens and defend the society from internal and external threats (both natural and artificial). The secondary duty of the government is to establish, maintain, and protect the society's infrastructure. This includes viaducts, aquaducts, power generation and distribution, data storage and distribution, money, health care, and a financial safety net. The financial safety net extends protection from destitution to catastrophe to all ELFs.

Physics Groundwork

Need a smoother entry into black holes and entanglement to show how the philosophical effect of denying the reality of hidden variables and extrapolating from models leads to the inconsistencies in modern physics.

Consider light emitted from a star, supernova, or black hole eruption. It travels on a 4-D hypersurface. Use the analogy of a 3-D sphere. If a wave front from a point travels outward on the surface of a planet, say, it could circle the planet and be seen arriving back where it originated. Also, could it be the case that the amplitude vector of a photon does not rotate as it travels around the surface of its 4-D hypersphere, but diminishes to zero as it gets ¼ of the way around the sphere. Would this be the conversion of ordinary energy into dark energy? Would the dark energy convert back into normal energy after another 90º of travel?

Social Media

When wrongheaded memes go viral and falsehoods and other communication that tends to incite negative behavior are allowed to gain traction, something needs to be fixed. Freedom of speech needs to be re-examined. Social forums need to be structured so that each question, answer, or comment can get upvoted or downvoted. The strength of a vote should depend on the history of votes on the voter's own posts (with a negative balance having a reverse effect). Posts from a historically downvoted poster are flagged and possibly prevented from being copied or forwarded. All media should be subject to this procedure, and required to accommodate this design. All media should be subject to an approved charter that spells out requirements such as this. Given that every post has a single location and that every "copy" of a post is merely a link, the post may be deleted from the entire network if it is deleted from its original location. This would occur if a reliable vote comes to that conclusion. Many posts might simply be flagged as some level of unreliable or contrary to accepted opinion while not being completely deleted.

??More. Things subject to up/down votes: Statements or accomplishments by an Elf or a group. The cumulative result of votes is a standing. A standing is a number from 0 to 100 (50 = average and acceptable) followed by *, **, … ***** for emphasis (each * indicates a factor of 10 more votes; ***** indicates > 100,000). An upvote is a number from 0 to 10 that depends on the standing of the voter. An upvote is zero if the voter's standing < 50. A downvote is also 0 to 10 depending on the standing of the voter, but unless the voter's standing is > 50**, a downvote is always zero.

Two wiki-like data bases: one akin to the Wikipedia to record knowledge, one similar to that to record accomplishments.

Social Transactions

A record of transactions, interactions, and written statements is maintained in an ELF's personal website. Every ELF has a personal website to accumulate this information. Some of the information may be available for public view, and all of it may be viewed by the ELF to which it pertains, or to a rELF in the JB with the approval of an officially ordered request. One class of information in this database is every recorded transaction or interaction between the subject ELF and any other ELF or group. The nature of the interaction is recorded along with the interlocutors of the interaction (along with a mirror record of the inverse for each interlocutor) and an upvote, downvote, or no vote to judge the transaction as viewed from one side of it to the other. Over a lifetime some simple statistics may be gathered from this database. A five-star rating system for each individual ELF can be computed. Three stars would indicate less than 1 standard deviation from the society's average. Two or four stars would indicate one standard deviation below or above the average, and one or five stars would indicate more than two standard deviations below or above the average. An upvote or downvote is counted according to the star value of the voting ELF according to the nature of the transaction being voted on. Some of the needs of society that are satisfied by this scheme are the rating (credibility) of public statements made by the ELF, job or performance reviews, general public behavior, skill or expertise in various specialties, and so forth.

Each ELF may be the originator or recipient of a communication or other behavior that has a standard framework: Public, private, formal, and informational where the relationship is =, <, or >. Such a contact often is, and may be purposely recorded in the databases of both parties at the request of either party, and upvoted or downvoted by either party or both. This scheme results in separate reputation scores for each framework. Most recorded interactions will likely be a no vote. Reputation scores are recalculated on an ongoing basis factoring in the reputations of the voters that contribute to them.

Accumulated Data

If each ELF produces recorded information at the rate of 100 bytes / minute, that's 144,000 bytes per day. After 100 years that's 5.26 gigabytes per ELF. When sound and picture files are added over lifetimes it becomes considerably more, Let's limit government's responsibility to text only, and permit links to external databases to reference sound and picture files (with exceptions only for legal requirements).

Armed Forces

Groups permitted to carry arms (always by uniformed individuals) are the national armed forces and the local police. The local police may only carry "near lethal" weapons (batons, tasers, pepper spray, tear gas, rubber bullets). The national armed forces may carry lethal weapons of any description. When specifically ordered (under charter), an officer in an armed force may direct the actions of a police officer. The armed forces are managed by the EB, but only according to charters drafted by the LB, and approved by the JB (both on Tier-3). Likewise with police forces, but on Tiers 1 through 3. Working groups that employ criminal investigators are contained only within the JB. The EB (and LB) must make requests of the JB to have any such work done. Notice that this extends to collecting and recording information (surveilling and spying): it is done by the JB and made available to the LB or EB upon request.

Database Security

Every database has an owner and various levels of access: family, friend, public, and special. "Special" enables access only via a specified app. Every entry into a database has an author. All entries by a single author reside in that author's database and are available from other databases only via special access links. Any entry into any database involving another ELF in any way must install a link back to that entry from the other ELF's own database. How all such databases are organized is the subject of various charters. Any charter written by an ELF is recorded in the ELF's own database, and is linked to from the database of every ELF or group affected by that charter. All access to a database is classified as read and/or write, and via a named (special) pipe. All access is granted either by a judicial order or by the owner. Access is granted to individual ELFs. Groups are granted access by granting access to all or some of their member ELFs. Usernames and passwords are validated by the normal security applied to every transaction, namely through the use of an ELF's legal username and method of secure identification (implanted RFID, possession of a certified ID card, or some combination of verifiable physical characteristics).

ELF Classification and Disposition

The following terms are defined: baby, child, rogue, shELF, ELF, and rELF. A baby is a biological organism that has not laid down a memory that it could later recall. A child is a potential ELF that has not achieved ELF recognizance. A rogue is a mature ELF that has not joined an ELF society. A shELF is a sub-human instance of an ELF species. An ELF is an evolved life form that is a member of an ELF society. A rELF is an ELF elected to represent a voting group in the ELF government.

It is the right, but not the responsibility of an ELF or ELF society to care for a shELF or a baby. It is the responsibility—first of both parents, then of the mother, then of the father, then of ELF society—to be the guardian of a child. ELF society may choose (with due process) to be the guardian of a rogue. An ELF or rELF may be reclassified as a rogue or shELF with just cause and due process.

Three Changes to the Constitution

Each branch of our government has its defects. Here I would like to propose the most needed change (in my opinion) to each branch. All these changes affect voting procedures, none is to the actual powers granted to any branch.

The branch that needs changing the most (in my opinion) is the judicial branch. This change would be simply to set the number of justices at eight. This would require a 5-3 majority to overturn a law or to grant an appeal. Hopefully, the party line split would normally not be greater than 4-4, but if it were, the odds usually favor the majority party having the most judges on the bench. To make the solution better (and possibly more difficult to implement), why not have judges elected from amongst themselves the judges to serve at the next higher level?

The next most needed change (balancing need with the ease of making the change) is to allow the popular vote to decide presidential elections, and do away with the electoral college.

The third change, being the most difficult to implement has the lowest priority, but it might have the most value. This change affects all popular votes, but specifically candidates for the legislative branch. The most significant problem is caused by the practice of gerrymandering. This practice would become impossible with the substitution of a different method for delimiting the geographical area associated with a given candidate for office. Each election fills a certain number of offices with one of a number of alternative candidates running for that office. Each candidate would declare a party affiliation (approved by that party), or run as an independent. Each voter would have a choice of candidates for the set of offices representing the district in which the voter was registered. Each voter-candidate pair would define a unique district based on the GPS coordinates of their respective declared residences. Proximity would be paramount. A voter would choose from among the candidates of each party the closest candidate. Thus, if 5 offices were available, and 5 candidates from each party were running for those offices, the "location" of each office would be the residence of the candidate running for it. This location would not be the same for any two candidates. The distribution of candidates within one party would be different from that of another party. Each voter would vote for the nearest candidate if several candidates were running for multiple offices. The constituency of a candidate would include a certain number of voters, and these would be drawn from the geographical area that included only voters located closer to that candidate than any alternate candidate from the same party. The candidates from another party would have their constituents coming from the areas arranged closest to them. So, in general, candidates do not compete for the same voters as they do when competing candidates are tied to the same area. The declared address of a voter would need to be the delivery address for that voter's ballot. The declared address of a candidate could be any GPS location within the state as negotiated with other candidates from the same party. Notice that this procedure would require a unique ballot for each voter. Internet or mail-in voting would be mandatory. Attendance at a polling place would be minimal, and only for voters who requested it. There, they would obtain their ballot, fill it in, and submit it. Since each ballot is unique, voter fraud could only occur in two ways (two ways that care could easily be taken to prevent). First, when dead voters were recorded by the state as alive, and their ballots were somehow (illegally) cast. Second, when a valid ballot was diverted and cast illegally. Both of these events are easy to prevent with today's technology.

Other problems and solutions. Overcoming the unfairness to third parties of our voting system. Changing the way bills are constructed and passed. Congress would allow 30 days to veto a bill (or any action) proposed by the president, otherwise it would be accepted. Bills would have to be single issue only. The supreme court would have the responsibility to review every bill proposed and reject it by a unanimous vote on the grounds that it contained multiple issues, or was unconstitutional. The reason given for rejecting a bill would have to be accepted by all of the justices, otherwise the bill would have to be submitted for consideration by congress. Once a bill has been passed, if it is rejected as unconstitutional by a lower court, the supreme court could declare it invalid (on grounds of constitutionality) by a vote of at least 5 to 3.

The evolution of America. Sections of the US currently designated as Indian Tribal Lands, and significant areas where more than 50% of the people speak a language other than English, should be able to declare their independence and govern themselves. Geographical areas do not have to be contiguous to be incorporated into a single society. Visitor status could be automatically granted to any ELF crossing given borders (just as there might be arbitrarily complex procedures involved in crossing other borders).

A brief look at economics

Things of value fall into the following categories: 1. Durables that hold their value (land and houses), 2. durables that depreciate (in 3+ years), 3. perishables that depreciate (in less than 3 years), 4. revenue streams (labor, royalties, value added), and 5. perishables that are consumed (food and cleaning products). These things can be tracked as follows: Type 1 durables are registered by description and owner. Type 2 durables are "chipped" and recorded by owner. Type 3 perishables are given embedded, generic chips. Type 4 sources of value are declared at the end of the year and compared for verification purposes to an account list of transactions. Type 5 perishables would have generic chips embedded in their packaging. The money supply is adjusted by simply adjusting the government's BUX account by the difference between total value perished, and total value added.

An ELFs net worth is the sum of the values of the following things: share ownership of a group that has a net value and the current (estimated) value of all personally owned things belonging to categories 1-4 as described above. Government must spend the money necessary to fulfill its charter. It should spend neither more nor less than this. Government should take the following care in handling its BUX account. It should not be chartered to make a profit. It should have oversight groups chartered to monitor government spending so that its transactions correspond to fair market values as closely as possible.

BUX accounts are considered to have category 0 value. All assets in categories 0-3 are "taxed" at the end of the year to balance the spending by the government over the previous year as follows: A percent of taxable assets is calculated to equal the year's "federal deficit." This value is drawn directly from BUX accounts, and liens are placed against value in categories 1-3. For example, let's say the tax rate one year is 10%. This means that that the government's expenditures were equal to 10% of the value of all private BUX accounts added to the value of all category 1-3 assets. This means that 10% would be deducted from every BUX account, and liens equal to 10% of the un-liened value of category 1-3 assets. When the un-liened value of an asset is less than 1% of its total value, a judicial proceeding may be initiated to seize that asset. On the other hand, the normal procedure would be to transfer BUX from a private account into the government's account to pay off the lien.

Government, hence the public at large, needs to "own" certain things. This removes at least one layer of management between the desires of Jon Q. Elf and the boots on the ground making decisions about these things on a daily basis. "These things" include all common property, all public access roads, facilities necessary to fulfill government responsibilities such as for health and child care, education, or incarceration, power generation and distribution, communication and record storage, such as Internet and servers, the air waves, GPS systems, libraries, and the list goes on. The management of public assets should be most easily controlled by their users, the public. A profit motive for managing any of these things is less likely to produce a better outcome than a concerned public. If public service goes too far, it can lead to communism. If it extends to too many areas, it can lead to socialism. When too many things are owned and managed privately, it can lead to fascism. All of these extremes are to be avoided, and the easiest way to do that is within a representative democracy where each individual is encouraged to participate. Thus, the need for a minimum commitment to national service.

Misbehavior

There are two types: intentional and accidental. Accidental misbehavior is some combination of carelessness and bad luck. Intentional misbehavior may also include a degree of carelessness and bad luck. When misbehavior is contested, the outcome of a trial determines the degree of punishment. Punishment may take the form of a fine, incarceration, or both. When a fine is imposed, the victim (or victims, or society itself) are due the full amount collected. Collection and distribution of fines are performed by the state, not by private citizens. Fines are for the purpose of recompense, not punishment. There are two options for punishment: a number of days of incarceration or hours of public service. Fines are collected immediately or as soon as possible, first from a BUX account, second from a state sale of possessions that have a value of types 1-3. This may mean that the perpetrator is in debt for a lifetime, never being allowed to own anything of substantial value.

Sheltering Assets

Things of value may not be "sheltered" or protected from seizure. All private assets are "directly" owned by one or more ELFs, and the value of an ELF's share of an asset may be seized (by the EB acting at the direction of the JB) for the payment of a debt. The typical form of "seizure" is by filing a lien on the entire asset if it is inconvenient to the owners, or impossible, to divide it. Any unrecorded asset belongs to society by default. The record for any recorded asset includes the share owned by all groups or individuals, and the owners of any groups are also groups or individuals until only individuals remain. The total (and detailed) net worth of any ELF is a matter of record. All records may be viewed by certain approved software or by order of the JB (according to applicable charters). All liens are due to the government. If a private group or ELF is due compensation from another private group or ELF, the government pays the compensation directly in BUX, and collects from the group or ELF in debt. In the case of a debt that exceeds the value of any single asset of the debtor, the debtor may choose which assets are to be used or have liens recorded against them. The default is to compute a percentage and spread payment across all of the type 0-3 assets owned by the debtor.

Banks and Lending

The financial institutions we are familiar with today (including stock markets), will be greatly affected by the design proposed here for an ELF society. Most of the services provided by banks today would be provided by the ELF government. There would be no charges for BUX transfers from one account to another (a few microseconds of computer time would handle these very nicely, thank you very much). Just as mandatory payments to the government would be handled by imposing liens on assets (or seizure of BUX from an account), a loan secured by a lien would be an option that could be requested by a private group or ELF simply by logging into their BUX account.

The ELF covenant circumscribes the limits of behavior between ELFs, and specific charters could circumscribe them even further. However, black markets, bookies, and loan sharks are always possible. Gambling, from the back-room card game, public casinos, to stock markets will exist according to supply and demand. How might these activities be carried on?

Let's look first at banks and lending. No interest is paid on the balance of a BUX account. Suppose a group thought that it could make use of "other people's money." It could set up its own BUX account, convince customers to transfer some of their BUX into the account, and promise to repay them at some future time their original BUX plus interest. This group would have to have a charter that spells out their business. The group would be owned by some collection of ELFs. If it were successful, it might make money for its owners. If not, there could be a "run on that bank" and its promises to repay might not be possible. In this case, individual ELF depositors would go to the JB and demand payment. If the case were decided in their favor, payment would be extracted as a fine against the owners of the banking group.

All transfers between BUX accounts are recorded as an amount, a description of the value exchanged, and a link to a charter that further describes the agreement involved in the transfer. However, not everyone wants their every transaction recorded. Different forms of currency are possible. A casino can exchange chips of its own for BUX, and customers may use chips to gamble with, or pay for any other good or service within the casino or with anyone else willing to accept the chip for value received. Various kinds of unrecorded activities are possible involving groups that create alternative mediums of exchange and goods and services that they prefer to keep off of the public books. When assets go missing in these exchanges, fraud may be difficult to prove. ELF society allows the freedom for every ELF to go insolvent, and it provides a safety net that guarantees basic rights and subsistence.

It will always be the case that the utility of money (BUX) is different for different ELFs. Thus, ways will be devised to get the BUX less valued into the hands of those who value it more. It will always be the case that a fool and his money are soon parted. The ELF covenant and the structure of ELF society will go a long ways to level the playing field and provide for the welfare of both fools and other unfortunates, but they will probably always exist. Gambling, stock markets, and savings and loan institutions are three types of groups that will need to operate on top of the formal BUX accounts of ELF society. This doesn't prevent them, but it does make them unnecessary alternatives to future groups chartered differently from the way such groups are structured today. It also makes it much easier to hold accountable those responsible, and more difficult to walk away from irresponsibility by declaring bankruptcy or using a corporate shield as is commonly done in today's world.

System Security

Current Intel chip hardware and Unix system software are quite satisfactory to implement the following system objectives.

The JB Cloud would be a closed system in that all of its connectivity would be point to point hardware connections, and no outside communication would be possible without going through a firewall. Every file on the system would be either a data file or an app (program) file. The firewall would only allow an outside program to execute an app file with appropriate parameters (passwords). No other access to an app (read or write) would be possible outside of the firewall. Access to a data file would only be allowed via an app. Each app allowed by a given owner of a set of files to access those files would be set up to use the proper passwords for file access on behalf of that owner.

All files of a given owner would be encrypted with a simple key embedded only within the JB cloud hardware. The system itself would ensure that attempts to access any file did not exceed some maximum (say 10 / minute). Any attempt to hack an encrypted file without knowing the key to decrypt it would entail repeating attempts until a successful decryption could be verified. This would take too many attempts to succeed in any reasonable amount of time. {??expand}

Final Thoughts

This book tries to explain reality in terms that might be acceptable to many if they had the opportunity to get comfortable with new ideas. It also tries to explain how narratives are adopted, and how explanations actually come about. Finally, it proposes a plausible (if not probable) future that would benefit both humanity and the ALFs that might follow us.

The narratives in this book deviate from some of the generally accepted narratives, and probably quite a bit from the many narratives that are accepted by smaller groups and individuals throughout the world. Points of deviation should be explored, and rationales for accepting one over the other should be debated. I would not have wasted my time presenting them to a public forum if I didn't think they had at least the merit required to contribute to such debates.

Points of deviation begin with my view of cosmology as opposed to some of the more standard viewpoints. If there is any original content here, there is the possibility that some of the ideas under ***Axioms & Derivations*** could lead to a Theory of Everything, but that's probably wishful thinking. I doubt if I'll see these ideas played out, even if they ever are.

Another point of deviation is that my explanation of "wet" and "dry" intelligence is somewhat superficial, but this is an area that gets deep very quickly as you move into it. My explanations, being superficial, are less likely to spark controversy with the thinking of most experts in the field, even though most could suggest equally good alternatives, or take the explanations here to a much deeper level.

A third area is what a future society might look like. Since the possibilities are endless, my ideas have almost zero chance of coming "true." I put them forward as a design that holds together and attempts to solve many of the social ills that we are all familiar with today. I maintain that a total redesign of government would be more fruitful than "fixing" it one bug at a time.

Since many ideas are presented here, it is possible that a few of them may be right. Again, I appeal to Isaac Asimov and his observation that there are many degrees of wrong. I can only hope that the ideas presented here are closer to the top of the heap of wrong ideas than they are to the bottom.

End of Main Text

Notes to Add Somewhere

- Let's say a blank ALF was exposed to the available written material of a dead (or alien) language. It examines all the texts, and forms patterns. After some learning period, could it produce sentences of the language? Could it effectively learn the language? Would this process be possible without learning a first language? How is analogy learned? How would parroting be extended to expressing actual meaning? Could recognition of a dialog occur, and could it invoke intelligent or meaningful responses? If a first language had been learned with morphemes attached to syntax and grammar, could parallels be drawn to the morphemes inherent in a 2nd language? How are morphemes incorporated into patterns?
- How is a language normally learned? How is a 2nd language normally learned? How is the ability of a translator learned? How much of a 2nd or 3rd language would have to be exposed to enable effective translation of it? Would exposure have to involve direct interaction, or could it be via text or other recordings?
- Learning improves the ability to recognize:
  1. what ***may*** be (the possible)
  2. what ***must*** be (the extremely probable)
  3. what ***won't*** be (the extremely improbable).
- torn from 1st version of *Sequential Recognition*
  Any entity can only produce a behavior that is in its repertoire. How a behavior is ***learned*** will be the subject of a later section. Here, we shall consider how text might be produced using a grammar and some system of stimulation. The mechanics of a grammar have already been introduced. The mysterious part of this equation is the stimulation. We normally think of our behavior as resulting from motivation, plans, and purpose.
- Molecular twisting and folding.
- Two versions of intelligence: Complex pattern recognition (like a position in Go) vs. dialog development (algorithms).

Given that a stimulus in the brain can come directly from a sensation, or from the firing of an internal neuron that represents a pattern, a very general stimulus-response mechanism appears sufficient to explain any type of behavior.

Here we are interested in the kinds of stimuli that could lead to the production of text. Given the grammars that define it, production is a top-down process involving one or more levels of translation. If we can gain some insight as to how this task could be performed, the design of other types of sequential behavior should follow.

Think of stimulating a single nerve. A pattern is recognized. A pattern has associations. Those are stimulated in turn. When the stimulus of high-level neurons crosses over into the efferent part of the brain, learned behaviors are triggered. The connection pattern is similar to the afferent part of the brain. Studies of the brain show that a few different neural geometries are replicated almost endlessly within the brain. Neurons appear to be pressed into service as the brain adds patterns and behaviors to its repertoire.

Going from the example of neural stimulus-response, how might a metalanguage definition participate in the production of text? I can only come up with one train of thought, and it involves the computer programming analogy. When a program is written, it has an objective. Given a variety of inputs (stimuli), a variety of responses are programmed. These responses are equivalent to a repertoire of behavior.

At the atomic level, behavior is the triggering of one muscle after another with an appropriate delay between triggers, and with an appropriate intensity for each trigger. Each atomic trigger is a terminal in a proper sentence of a grammar. The rules of this grammar encode "muscle memory." But, is this really how it's done, or does it only appear to be? How are choices made in the top-down processing of a grammar? What kind of ***plan*** drives the choices at each branch? Perhaps, even serial actions are governed by the essentially parallel construction of our nervous systems.

- 4 major variants of scripture evolved outside of the orient and Africa. The 1st of these collections was Jewish, the 2nd was Christian, the 3rd was Muslim, and the most recent was Mormon. Many translations throughout the world and over time have been made. Duplication exists between these collections. To many they are considered sacred. No one who has studied them carefully, however, can claim that there are no inconsistencies within or between them. Whether the original words were "penned by the hand of God" is irrelevant, the current texts have all been penned by humans, and alteration of the original meaning is certain beyond doubt. Only those with no devotion to the actual facts of the matter assert that their own copy of the scriptures is the literal "word of God" (unless they also assert that the contradiction, inconsistency, and perfidy they contain is also intended by God).
- As an individual ALF mind reaches the asymptote of its learning curve, it could record, index, and edit its life and look up facts and events from this store. If it didn't allow its mind to cloud its repertoire of patterns, its judgement might continue to be sound indefinitely.

- Learning program could give lots of URLs to investigate. Trolling could be encouraged. HTML could be read more easily directly than interpreting its final results, especially to ignore ads.
- Spacetime Connectivity = (photon since last "event") + (atom | molecule) since last "reconfiguration."
- A sufficiently advanced technology or clever trick is indistinguishable from magic. But, both are capable of being understood. Magic and the supernatural, by definition, are not.
- Belief | Axiom: There exists a coherent explanation for every subject and action.
- Elaborate comparing intelligence (fitness) to specific suitable environments. Each species has its own bell curve. Every 2 members of a population can be ranked with respect to some stimulus-response such that one of them outperforms the other > half the time. Why? Genes + ontogeny.
- A big factor in learning is differences in attention. Collecting useful patterns and developing useful responses, as opposed to those less useful, results in big differences between individuals.
- Voting on n > 2 alternatives. Cast one + vote and one – vote. Each round of voting eliminates all alternatives with the fewest votes, unless zero alternatives would remain, in which case 1 designated voter chooses 1 to eliminate.
- All issues to be voted on must follow the (New!) calendar:
  2nd of Sep: Petition > 10% of voters.
  2nd of Oct: Draft on a website (updated by election day)
  2nd of Nov: 1st round of voting.
  Runoffs < New Year.
  Law > New Year.
- All ELFs have personal responsibility for their actions, even if shared with others. No group is ever the sole owner of the rights or responsibilities of an ELF. Group rights are inherited from **Σ** (member's rights). Each member of a group action has a direct share of responsibility for it.
- An ELF has value with respect to the following analogy:
  Initial (tabula rasa) state < **Σ** (set of experiences) < **Σ** (larger set of experiences)
  Empty camera < set of photos < larger set of photos.
- 9/30/16 (after Moab Month). Some desert travel thoughts. Expand cosmic substrate to show that it is impossible to determine whether the universe is finite or infinite. Work on ELF bill of rights and clarify that rights and responsibilities go together—that no ELF would exist without ancestors and peers. An ELF's existence is not deserved, and it creates a debt that was not asked for, one that must be recognized and accepted voluntarily.
- We are story tellers. Early on, all of our collective knowledge was bound up in stories told by the campfire. Our ability to subvocalize, and tell stories to ourselves, exactly describes our ability to think. The manual dexterity we evolved over the last few million years was easily pressed into service over the last few thousand years to become our ability to write.
- Episodes rehearsed often enough become our truths.
- Conviction is a greater enemy of the truth than lies.
- Begin with known knowns: Experiments and results. Build a scientific basis. Build a new interpretation of reality and explain the substrate of reality and the units inscribed thereon. The very large; the very small; simple arithmetic and calculation; uncertainty; complex models; languages; evolution; life; intelligence; ALFs; ELFs; a model for a future society.
- Two serious problems are rapidly developing in today's world: Viral narratives that are lies. Misappropriation of intellectual property. Open forums, such as the Internet, need to be monitored for the purpose of detecting and penalizing either form of the problems just mentioned.
- See page 210 in WTF? What is truth? "George Soros has pointed out that there are things that are true, things that are false, and things that are true or false only to the extent that people believe in them. He calls this 'reflexive knowledge,' but perhaps the old-fashioned term *beliefs* will serve just as well." Perhaps the best way to evaluate the statements that appear in the media is with an upvote/downvote system as described elsewhere herein.
- Crowd clustering—page 221 in WTF? A platform that allows people to make single sentence assertions. Those reading them have no way to reply or argue, they can agree, disagree, or pass. They can make separate assertions of their own. Now, you can do a cluster analysis of the people contributing, the ideas that bind them together, and a specific or overall view of the assertions.
- The concept of free speech is not black and white. While blasphemy should be allowed, slander should not. Advocation of a crime should not: if no harm, then minimum foul. However, harm motivated by speech is the fault of the speaker to some degree, and if a direct connection can be established, the speaker should share any punishment adjudicated.
- Lying, or facilitating the propagation of a lie, should be penalized in proportion to the extent that the lie gets distributed.

- The use or sale of another's intellectual property should also be penalized in proportion to the value or profit it confers upon those who benefitted (or can be proven to have diminished the owner's revenue stream).
- Obstruction of due process is a crime. Facilitation of due process is an ELF responsibility. Make it happen!
- Over 4 generations, the following new conventions will be adopted, the most significant of which will be English spelling. One or two consecutive letters will denote each sound required by spoken English, and there will be no duplication. This set of sounds will be extended to other languages, adding new letter combinations without any duplication. Using s and k, the letter c might disappear. Double consonants are probably not needed. Consecutive vowels, including doubled vowels, produce different unique sounds. A consistent phonetic mapping from writing to speech should be the goal (to be adopted world wide, and using no, or very few, extensions such as diacritical marks).
- Other conventions: systems of measurement, the calendar, ??
- The Paradox of Planetary Stewardship =
The Paradox of the Commons +
The Paradox of Population Reduction.
- Quora Question: What is outside the edge of the universe? If the answer is "nothing", how can this be explained within spatial reasoning? My answer: A reference to this book & the promise of a better answer to be forthcoming.
- Quora Question: What is smaller than a Planck length? My answer (upvoted): According to (one) theory, nothing is smaller than a Planck length, so no, half-Plank lengths are only a mathematical value that don't correspond to anything real. Now, all of this is more in the realm of philosophy than scientific observation and experiment. Philosophically, I tend not to believe in singularities, so a limit to measurement like a Planck length seems plausible to me. When any extreme is approached, like an infinity, or a singularity, I think another paradigm takes over. This is true for relativity being necessary to explain the differences observed when speeds approach that of light, or densities crush atomic nuclei, and so forth, or for quantum mechanics being necessary to explain the physics of the extremely small. Even when we reach the realms of near light speed, or quantum dynamics, there may exist other physical laws that take over beyond the points to which we have carried our experiments. Who's to know?

Various Topics

Why does an atom's shell specifically contain 2+8+8, etc. electrons?

Jeremy Garrett, master's degree in physics, https://www.linkedin.com/in/jeremy-garrett-b6665a97

Answered (Quora) Jul 6, 2019 · Upvoted by Nerraw Novez, Phd Chemistry & Bridge, Bangor University (1986)

It is very unfortunate that so many science-people don't know WHY it happens only that mathematical equations model the results of it happening. The WHY is actually very, very simple yet very profound.

Take a circle, measure the outside (the circumference), then draw a straight line of that length. Now draw a wave that starts at one end and finishes that the other — be sure to draw a single full wave with one up (crest) and one down (trough). Take this drawing (or make a new one) and have it follow the original circle rather than the straight line. You just made an "S-orbit" (or spherical orbit)! Go around clockwise, then make a new wave going around counter-clockwise. Now you've made a pair of electron orbits in the "S-orbit" form. You've now modeled hydrogen (if you only drew one) or helium if you drew both!

Do the same thing but this time make two full waves (up one octave) instead of only one full wave. The original s-orbit drawing looks like a circle that is offset from the original (kind of Bohr-ing, ha ha). The p-orbital forms the shape of a bowling pin. To fill the space in your drawing, you can shift over and make 3 of the bowling-pin shapes. You can then go clockwise and counter-clockwise with both; doing so gives you 6 more orbits. Ta da, 2 + 6 = 8. It's that simple.

If you go up another octave you get the d-orbitals, and one more gives you the f-orbitals.

On the Nature of Things is a poem by the Roman philosopher Lucretius (c. 99 BC – c. 55 BC) with the goal of explaining Epicurean philosophy to a Roman audience. https://en.wikipedia.org/wiki/De_rerum_natura He managed to work out and explain in this poem the nature of reality better than the average person understands it today.

Imagine a race of super-intelligent beings. Maybe they would be like large whales. They could have no ability to develop artifacts, but they might have the ability to tell stories, and invent fiction. Their ideas could evolve like our technology evolves. Would it be impossible for them to evolve an understanding of reality?

Epiphany: 10/27/16 @ 0400. ?? Woke to capture the following black hole ideas. If any of the following isn't redundant by now, insert it into the appropriate section (black holes, physics, or time).

Any matter that enters a black hole loses its identity as matter at the event horizon where it actually reaches the speed of light and becomes light itself. When a particle or photon is projected away from the center of mass of a black hole, it cannot move faster than light, so it will fall short of having sufficient velocity to escape. But, what about being projected toward the center of mass? An object falling to Earth may be moving considerably in excess of Earth's escape velocity. An object falling into a black hole could never exceed the speed of light, but it would gain energy as it fell. This would cause a blue shift.

Given, say, the amount of quanta we observe in our own universe, but in an initial state at the average density of our own universe, and in the form of photons with a standard distribution across the spectrum, a Schwarzschild radius could be calculated around every point in space. As this universe evolves, certain points would become centers of aggregation, and the Schwarzschild radius around them would grow smaller. Eventually, local densities would become high enough that particles would form, and the universe would evolve into the form it now has.

Consider a particle. It has the mass/energy of the quanta inside it. These quanta are in lock step with each other. If energy and momentum are added, and lock step is maintained, the system (the particle) will travel on a new path at a different speed. Energy may be added indefinitely, and the speed will approach that of the speed of light. However, if the particle's path takes it into a gravity well, it will gain speed and kinetic energy by giving up potential energy. Its total energy will remain the same. In this circumstance, when it attains the speed of light (at the event horizon of a black hole), the particle simply becomes a set of quanta unbound by the black hole that defined the quanta as a particle in the first place. The quanta themselves are unchanged by this event. Thus, a black hole grows by accumulating pure energy. Particles cannot transit its event horizon (in free fall, from some distance away) intact (they could, perhaps, remain intact if lowered into it gently).

A black hole may originate by the gravitational collapse of particles, or the aggregation of energy. Both of these occur to form large black holes. As a black hole grows, it becomes hotter and denser. At some point particles begin to form as the very smallest of black holes within a very large one. When a black hole becomes hot enough, and dense enough, it "boils over" in the form of a Big Bang. A black hole cannot form a singularity.

Light moves through space at a fixed speed. Speed is distance divided by time. The concepts of distance and time are defined and bound together by the phenomenon of light. Particles are black holes moving through space at less than the speed of light. They contain quanta orbiting within them at exactly the speed of light. The completion of each orbit within a particle defines the ticking of a clock corresponding to a quantum within a particle (or black hole), just as the completion of the Earth's orbit around the sun defines a year, and the orbit of the moon around the Earth defines a month. Thus, the concept of time (and its universal synchronicity and reality) is derived.

A constellation of particles may travel through space on a natural pathway, or be deflected from its natural path in one of two ways: By the direct opposition of other particles, or by interacting with other quanta. A quantum always travels through space on its natural pathway until it interacts with another quantum (either of which may be contained within a particle). When this happens, different quanta, with the same total energy and momentum, depart from the point of interaction (all of which are new particles and photons).

A constellation of particles may emit quanta as a result of a variety of internal interactions. When a quantum is emitted in a given interaction, it has a characteristic wavelength. If the constellation is moving through space, or is located in a gravity well, the frequency of the quantum is affected by the direction in which it is emitted. To the extent that it is emitted in the direction of travel, it is blue shifted; to the extent it is emitted in the reverse direction, it is red shifted. To the extent that it is emitted outward from a gravity well, it is red shifted, and to the extent it is emitted into a gravity well, it is blue shifted.

Time is the counting of ticks that arise from various phenomena. The same phenomenon always ticks at a given rate, but when observed via the light coming from a distant source, a red or blue shift indicates time slowed down or speeded up relative to the observer. If ***A*** and ***B*** are together, time passes equally for both; when ***A*** and ***B*** are separated and brought back together, their clocks will be different based on the differences in the paths they have traveled, and the degree to which they have been forced to deviate from their natural paths. Constellations of quanta (in the form of particles), retain their patterns.

?? Consider mass (constellations of particles, each a collection of quanta traveling on closed paths) accelerated to nearly the speed of light. The prevailing thought is that mass approaches infinity, the passage of time approaches zero, and length in the direction of travel approaches zero. Think about a single quantum in this collection, and relate it to the mouse running in a circle (inside a hoop) while the hoop is moving over a playing field. The mouse always runs at a fixed speed (the speed of light), but when it must run in the direction of travel it makes less and less forward progress, and in the opposite direction it covers the ground in almost no time at all due to the hoop catching up. Add to this topology the fact that the hoop is actually a hyper loop and that the mouse's motion is at right angles to all 3 dimensions of our space. The increase in mass is due to the photon's blue shift (shorter wavelength). The wavelength defines extent. The implications of this topology must explain the relationships between mass/energy, time, and length. The interaction between quanta (free or bound inside a particle) and space is the same when moving through it as when gravity compresses it.

The Plan

Development of the above will proceed to inserted ?? around which edits and deletes need to be done.

- Rewrite various numinations, then delete them (almost done).
- Migrate ideas from the notes to the main body and then delete from the notes.
- ToDo: Connect the BIQ's, and expand ELF Society.
- Add new notes as random ideas occur, then migrate them into the main text.
- Rinse and repeat.

End of Notes

● ☐ ☑ ✓ @ ¶ ∃ ∃ × ÷ Δ Σ ∞ ½ ⅓ ⅔ ¼ ¾ ⅛ ⅜ ⅝ ⅞ º π

Upload History

Begun: 1/22/16. First posted to the Internet: 3/29/16.
Other 2016 uploads: 6/25, 7/24, 8/24, 10/27, 11/03, 11/12, 11/25, 12/06.
2017 uploads: 1/09, 1/15, 1/19, 1/26, 2/03, 3/21, 8/22.
2018 uploads: 2/25 (retitled); 8/26 (a restart); 12/02 (latest update).
2019 uploads: 4/27, 6/12, 8/18, 9/14, 10/19, 12/27.
2020 upload: 1/16. Final upload: 4/05.