

The Web and Life

The story so far: In the beginning the Universe was created. This has made a lot of people very angry and has been widely regarded as a bad move.

—Douglas Adams, *The Restaurant at the End of the Universe*

INTRODUCTION

In some corner of a bland, ambiguous office somewhere in the world today, there is a high likelihood that the following scenario will play out. A young and rather attractive clerk will wander over to the office's token computer technician and ask, "So . . . Why do they call you a geek?" most likely as a bet from some fellow worker. Our technician will be unflustered by this question, having encountered it several times before in their unremarkable career. Nevertheless, they will still raise their head from their coffee-stained keyboard, smile and reply sincerely, "Because I talk to my computer as if it's my best friend." The clerk will never speak to the technician again.

This scene may appear to be an unrealistic and cruel caricature, but for many of those involved in the Information Technology industry it holds much truth. Scores of individuals have followed a career in computing inspired entirely by the expectation that, perhaps one day, they might be able to interact with machines as if they were alive. Whole sections of the industry openly aspire to it, with Artificial Intelligence gurus rushing from their closets to freely foretell a brave new world just over the horizon. But this world has never appeared and we are still faced with the reality of computer systems as lifeless and unnatural creations. Or are we?

A few facts are beyond question. Today, one particular example of combined computational power has emerged on a scale and power far greater than any individual or organized collective could have ever hoped to ascertain or understand in the past. This is the World Wide Web, or "Web" for short, and its unparalleled power is growing steadily by the second. Its growth is almost scary. Widespread use of the Web did not really begin until around 1995, when studies accounted for around 16 million users. By 2001 this figure had grown to over 400 million, and some estimates now predict that it should have topped 1 billion by 2005 and will surpass 2 billion by 2010 [80]—around one-third of the world's population by most common accounts. Couple the fact that the Web is well on its way to absorbing significant

portions of mankind's joint knowledge with the raw processing power that is inherent to its technical infrastructure and a social machine the likes of which we have never experienced before is plain. As Gustavo Cardoso, professor of information and communication sciences at ISCTE, Lisbon, said in 1998, "We are in the presence of a new notion of space where physical and virtual influence each other, laying the ground for the emergence of new forms of socialisation, new lifestyles and new forms of social organisation." Some have even referred to the Web as a higher level of human consciousness, a post-human [40] existence with its own independent cognitive capabilities and conscience, a "Metaman" [1] if you will, emerging from unapparent macroevolutionary processes.

A sizable and controversial claim without a doubt. Surely no form of man-made technology could, or should, ever be considered in the same vein as life itself. Perhaps not, but there are certainly a number of apparent similarities between the type of highly complex computer systems we see today and both the development and composition of many real-world systems that natural science has chosen to classify under the heading of "life." In short, and this point is key, a number of reoccurring patterns and themes appear prevalent across both real and virtual worlds. "When a pattern recurs in different systems which bear no obvious relationship to one another, we must suspect a common causative principle, one which can be understood in the most general terms without reference to the specifics of this or that case [80]." To take one brief example, binary characteristics linked to a significant number of macro and micro natural systems have been a matter of well-proven scientific fact for some time. In particular, they are recognized to play a key role in the state transitions surrounding many of the complex systems directly related to the idea of organism. Such characteristics are also fundamental to both macro and micro control in the digital systems world, but until relatively recently these have not been investigated under the same light as their evidently similar real-world twins. It is this duality of pervasion across real and digital worlds that makes binary systems both influential and enthralling. As a recurring theme that appears in the most unexpected of guises, there is much evidence in our Universe to point to the true power of two.

At this point in our understanding there is a need to be careful, however. It is important to emphasize the term *similarity* and not confuse it with the absolute understanding of proven equivalence. Even so, by indulging the creative license granted by established prior work, it should be remembered that mankind has made many great leaps by initially recognizing similarity alone. First comes imagination, inspiration, and recognition, then speculation, investigation, and ultimately proof or disproof. Some ideas are reeled into our minds wrapped up in facts, and some burst upon us naked without the slightest evidence that they could be true but with all the conviction they are. The ideas of the latter sort are the more difficult to displace [21]. Currently, we may well be somewhere between imagination and investigation when considering ideas of life through technologies like the Web, but if one accepts this position and allows a certain level of trust in the proven research of recent decades, then at the very least a highly compelling and provocative case can be presented for others to later validate properly. Many of the observations in the pages to

follow have already been substantiated, whereas others are somewhat more speculative. Nevertheless, it is hoped that where propositions are unproven, they are supported by sufficient evidence to make serious consideration at least credible, and it is in this spirit that the majority of this book is presented.

It is surely a truism to state that life has many interpretations for many good reasons. Try to encapsulate just one of these and an untold number will stand by laughing at you. In such respects, it is beyond doubt an essence—that is, an untouchable and complex collection of entangled interpretations, interactions, and dependencies across an unimaginable number of facets. But some of our latest discoveries have started to challenge many of our oldest beliefs about life. For example, today we know that life can survive in conditions far beyond those we would consider to be “normal.” There are places on our planet where no light and little oxygen can be found and where temperatures can soar to levels in excess of 400°C (750°F), yet unbelievably life still thrives. These are the hydrothermal vents of the deep sea floor that teem with a fascinating array of life such as tubeworms and huge clams. Furthermore, such vents provide just one example of a seemingly intolerable habitat that life has concurred. There are many more examples, so many in fact that biology has devised a term to describe the life forms that overrun them—Extremophiles.

But such habitats do not even start to push the extremes of life as we can now perceive it. For instance, as we shall see later, more abstract interpretations of life have now been conceived that do not require even the slightest morsel of material embodiment. To be blunt, in its purest guise, life is formless. Although the human race may well be a noble illustration of its many marvels, we are but one of many millions of different occurrences of nobleness in this world. For reasons like this, it is difficult to produce a condensed, coherent, and flowing description of life, or any matter seriously associated with it. Bits of the argument simply refuse to fit into the box, flopping out over the sides at the merest hint of classification.

Although every effort has been made to string together a compelling sequence of arguments in order to associate the Web we see today with the concepts of life, any number of routes could have been taken through the various discussions included, and any number of valid interpretations of life could have been given ultimate favor. Nevertheless, a primary aim of this work is to be as accurate and factual as possible in the material presented. Reading from start to finish will undoubtedly provide an erudite set of overall ideas, but the task of highlighting the most and least convincing evidence involved has been deliberately left to personal discretion. Perhaps several attempts will be needed to take in the subtle and sometimes organic associations entailed, but it is hoped that the effort involved will all be part of an enjoyable and stimulating read.

ALMOST NONE OF THIS IS NEW

Most who are even remotely familiar with the Web would concede that it is a concept inseparably linked with the modern day digital computer. Most would also acknowledge that, although such machines are becoming “cleverer” by the day, they

still do not possess one ounce of the same cleverness normally associated with any living being. The two are quite different; computers just do sums while we do something else. That's the way the world goes round: Computers simply do not think! Furthermore, computer engineers are not interested in the types of "stuff" that goes on in a living mind, are they? They are off in a land of space invaders, sort algorithms and slide rules. They can't even tie their own shoelaces let alone interact and appreciate the world to a level that could tackle the subtleties of a cognitive process. After all, why should they? They are making enough progress on their own without needing to be troubled by the challenges of such problems—they have the Web to play with now.

This is a fine parody for certain, but one that is nonetheless somewhat distant from the real truth. If you look as far back in the annals of modern computing as is practical, you will soon find that most, if not all, of the great pioneers in this area were fascinated by the concepts of life and mind. John von Neumann, for instance, one of the most influential individuals in the entire history of computing, was driven by such concepts, being fascinated by the concept of self-operating machines—often referred to as automata. In 1966 he published a book entitled *The Theory of Self-Replicating Automata*. Ted Codd, the father of relational database technology, also published in this area with his book *Cellular Automata* in 1968. So the trend went on and the search for lifelike ways of computing continued. Soon a number of new ideas and a whole collection of new fields of research began to sprout. Take, for instance, the groundbreaking work by Marvin Minsky, who started the Artificial Intelligence movement, John Holland's work on classifier systems, Friedrich Hayek and Donald Hebb's inspiring research into neural networks, and last, but not least, Chris Langton's outstanding contributions that culminated in the notion of genetic algorithms and artificial life.

In fact, until the mathematician Alan Turing first conceived the very notion of a universal programmable computing machine, the word "computer" typically referred not to an inanimate object but to a live human being. It was 1936, and people with the job of computing, in modern terms, crunched numbers. Turing's design for a machine that could do such work—one capable of computing any computable problem—set the stage for the theoretical study of computation and remains a foundation for all computer science. "But he never specified what materials should be used to build it" [83].

"Turing's conceptual machine had no electronic wires, transistors, or logic gates. Indeed he continued to imagine it as a person, one with an infinitely long piece of paper, a pencil, and a simple instruction book. His tireless computer would read a symbol, change the symbol, then move on to the next symbol, according to its programmed rules, and would keep doing so until no further rules applied. Thus the electronic computing machines made of metal and vacuum tubes that emerged in the 1940s and later evolved silicon parts may be the only 'species' of nonhuman computer most people have ever encountered. But theirs is not the only form that a computer can take" [83]. Living organisms also carry out complex physical processes under the direction of digital information. Biochemical reactions and ultimately an entire organism's operation are ruled by instructions stored in its genome, encoded in sequences of nucleic acids.

“When the workings of biomolecular machines inside cells that process deoxyribonucleic acid (DNA) and ribonucleic acid (RNA)¹ are compared to Turing’s machine, striking similarities emerge: both systems process information stored in a string of symbols taken from a fixed alphabet and both operate by moving step by step along these strings, modifying or adding symbols according to a given set of rules [83].”

But Turing is not the only significant pioneer in computing. Many others have also contributed to important computational themes indirectly. Stuart Kauffman, Brian Arthur, and even the two great physicists Richard Feynman and Murray Gell-Mann, hugely influenced the advances made. Therefore, hardly any of the material presented here is either novel or speculative. Rather, it is calculatingly borrowed from some of the most well established computational works of theory of the past century. The only thing that may be new, however, is the context in which it is applied, namely the Web context. All the evidence and arguments have been applied in similar contexts before, with similar problem spaces all dealing with the same core set of parameters: large-scale dynamics, complexity and adaptation, to name a few.

WHERE TO BEGIN?—WEB MISCONCEPTIONS AND FOLKLORE

Before going any further, and certainly before probing the depths of any discussion on whether the Web may be alive or not, a number of common misconceptions must be dispensed with. The Web is not, for example, the Internet, although it is closely dependent upon it. The two are sometimes perceived as synonymous, but they are not [76]. For this reason, any use of the colloquialized term “the Net” in reference to the Web can only serve to confuse and is hence frowned upon here. The Internet is a communications network, a global framework of wires, routing devices and computers on which the Web rests, and to think of the Web just in terms

¹Structurally, RNA is indistinguishable from DNA except for the critical presence of a hydroxyl group attached to the pentose ring in the 2’ position (DNA has a hydrogen atom rather than a hydroxyl group). This hydroxyl group makes RNA less stable than DNA because it makes hydrolysis of the phosphosugar backbone easier [22].

Double-stranded RNA (or dsRNA) is RNA with two complementary strands, similar to the DNA found in all “higher” cells. dsRNA forms the genetic material of some viruses. In eukaryotes, it may play a role in the process of RNA interference and in microRNAs.

The RNA world hypothesis proposes that the universal ancestor to all life relied on RNA both to carry genetic information like DNA and to catalyze biochemical reactions like an enzyme. In effect, RNA was, before the emergence of the first cell, the dominant, and probably the only, form of life. This hypothesis is inspired by the fact that retroviruses use RNA as their sole genetic material and perform information-storing tasks. RNA can also act like a catalyst, a task mainly done by proteins today. There are several ribozymes, catalytic RNAs, that have been discovered, and peptide bond formation in the ribosome is carried out by an RNA-derived ribozyme. From this perspective, retroviruses and ribozymes are remnants, or molecular fossils, left over from that RNA world. Assuming that DNA is better suited for storage of genetic information and proteins are better suited for the catalytic needs of cells, one would expect reduced use of RNA in cells, along with greater use of DNA and proteins.

of electronics and silicon would be wrong.² Other terms like “Information Super Highway” may also be easily misconstrued as characterising the Web but don’t really quite get there. They do not convey the truly global, pervasive nature of its vast information bank,³ instead perhaps conjuring up unnecessarily artificial images, heavily dependent upon silicon-laden machines. The Web is not like that, it is something quite different, as we shall see.

The Web is not as young as one might first think. Computing pioneer Vannevar Bush⁴ outlined the Web’s core idea, hyperlinked pages, in 1945, making it a veritable pensioner of a concept on the timescale of modern computing. The word “hypertext” was also first coined by Ted Nelson⁵ in 1963 and can be found in print in a college newspaper article about a lecture he gave called “Computers, Creativity, and the Nature of the Written Word” in January 1965. At that time, Nelson also tried to implement a version of Bush’s original vision, but had little success connecting digital bits on a useful scale. Hence his efforts were known only to an isolated group of disciples.

Few of the hackers writing the code for the emerging Web in the 1990s knew about Nelson or his hyperlinked dream machine [35], but it is nonetheless appropriate to give credit where credit is due.

The origins of the Web as we would recognize it today eventually materialized in 1980, when Tim Berners-Lee and Robert Cailliau built a system called ENQUIRE—referring to *Enquire Within Upon Everything*, a book that Berners-Lee recalled from his youth. While it was rather different from the Web we see today, it contained many of the same core ideas.

It was not until March 1989, however, that Berners-Lee wrote *Information Management: A Proposal*, while working at CERN,⁶ which referenced ENQUIRE and described a more elaborate information management system. He published a more formal proposal for the actual World Wide Web on November 12, 1990, and implementation accordingly began on November 13, 1990 when Berners-Lee wrote the first Web page. During the Christmas holiday of that year, Berners-Lee built all the tools necessary for a working Web: the first Web browser, which was a Web editor as well, and the first Web server.

To move the timeline on, in August 1991 he posted a short summary of the World Wide Web project on the *alt.hypertext* newsgroup. This date also marked the

²This entails a form of “dualism,” the philosophical viewpoint advocated by the highly influential seventeenth-century philosopher and mathematician René Descartes. This asserts that there are two separate kinds of substance: “mindful stuff” and ordinary matter.

³It has been estimated that at the end of 2002 there were around 3 billion documents available on the Web, with several million more being added every day.

⁴Vannevar Bush was an American engineer, inventor, and politician, known for his political role in the development of the atomic bomb and for his idea of the memex—seen as a pioneering concept for the World Wide Web. He was allegedly a member of the secret committee Majestic 12 investigating UFO activities [90].

⁵Theodor Holm Nelson (born circa 1939) invented the term “hypertext” in 1965 and is a pioneer of information technology. He also coined the words hypermedia, transclusion, virtuality, intertwingularity and teledildonics [91].

⁶CERN is the European Organization for Nuclear Research (Organisation Européenne pour la Recherche Nucléaire) and is the world’s largest particle physics laboratory [92].

debut of the Web as a publicly available service on the Internet. In April 1993, CERN announced that the World Wide Web would be free to anyone, with no fees due.

The Web finally gained critical mass with the 1993 release of the graphical Mosaic Web browser by the National Centre for Supercomputing Applications developed by Marc Andreessen and Eric Bina. Prior to the release of Mosaic, the Web was text-based and its popularity was less than that of older protocols in use over the Internet, such as Gopher and WAIS. Mosaic's graphical user interface allowed the Web to become by far the most popular Internet protocol.

OUR UNDERSTANDINGS OF THE WEB

In one way or another, everyone is involved in systems. From the moment we are born until the time we die, our lives are governed and controlled by systems. We are born into a particular religious system, reared under some political system, and educated in a school system. We may be looked after by a health system, warmed by a heating system, and subjected to a draft system. There are systems of rivers, systems of organizations, and systems of systems. But what actually is a system?

The variety of systems can make a general definition difficult, but to be as broad as possible, a system may be defined as a set of components, connected in some fashion, directed to some purpose [33]. In such a way it is an entity in itself which maintains its existence through the mutual interaction of its parts. The key emphasis here is one of mutual interaction, in that something is occurring between the parts over time which maintains the system. A system is, therefore, different from merely a heap or collection of things.

Needless to say, the Web is one such system and categorizing it in this way provides a significant and important distinction from other more common definitions. Systems are also processes that consume, transform, and produce information. In other words, they change information in some noticeable way. This conflicts with the popular vision of the Web solely as a source of data en masse and implies that certain elements of its being are often overlooked. Certainly this is true when one remembers that systems can further incorporate other systems, like us as individuals. And here another important point is raised, for it is us, humans, who are the driving force behind the Web. Without us the Web would not have been created and would most certainly not have evolved to its current position today. We are intrinsic to its very being and are becoming increasingly dependent upon it in turn to maintain the stability of our various modern societies, cultures, and even core beliefs. We have truly become assimilated into its very fabric, thereby producing a hybrid hyperspatial system of natural flesh and blood combined with digital information tissue—a machine of sorts, a “social machine” the likes of which we have never seen before, to quote Tim Berners-Lee. So from at least that perspective the Web can currently be viewed as being partly alive in the truest sense, with society constituting its essential living element.

But there are still more elements of the equation to consider. As a case in point, it is wrong to consider the Web as being exclusively about raw information or hu-

man-based processes. Today the Web embraces both data and automated functionality as part of its make-up. Web Services, a recently introduced Web-friendly set of technologies, allow back-end computer systems, programs, and processes to interact using the Web as their door to door courier of data. Furthermore, multiple layers of data now exist on the Web, and so it is not uncommon to find cases of data about data, data about that data, and so on. This is the concept behind the newest branch of the Web, known as the Semantic Web, which points to a self-reflective world of information that freely talks about itself across multiple levels and dimensions. These dimensions not only add volume and weight to the descriptiveness of the Web's data, but also open doors to far greater levels of computerised automation. In some quarters the Web will soon be able to dispense with its human hosts altogether and start to do real work for itself, discovering its own weaknesses and strengths along the way, extending and compensating accordingly. It will become truly autonomous, changing into a place where Web service will find Web service, and those best suited will feed off each other's peer circle of information and capability. And the more powerful and connected such clusters of capabilities become, the greater the likelihood that they will produce even more superior capabilities. So an upward-spiraling current of self-feeding, self-organizing, decentralized computational power will have been started, if it has not been started already. The engine to a much more powerful instrument than the global entity we see today will be kick-started, controlled only by society's collective consciousness. As we shall see later, there is strong evidence to support the case for the eventual transition of this into something far greater, something far different from the Web of today. Will we recognize this transition when it happens? Has it happened already? Will we be able to harness or even comprehend the outcome of any higher Web state? Perhaps the philosophers amongst us might like to guess at the answer, but the vastness of the Web will surely work against them. Just like a single ant climbing the trunk of a great redwood, they may no more understand any future state of the Web than the ant understands the art of bonsai.

POWER OF THE PEOPLE

The great power plants of our planet might well provide the energy that invigorates the Internet, but they are not the driving force behind the Web; "that job lies entirely within the realms of human motivation and our dispositions toward action" [77]. Moreover, this is a force to be reckoned with, having led to both to the highest and lowest points of man's achievement. So, if one accepts that, as a global community, we are now joined with the Web into one great seamless being, it would not be too strong a conclusion to suggest that the underlying mechanisms of the Web are actually driven by the very same system of wants and needs that control us all as human beings.

Through us the Web feels its most basic survival instinct. We not only provide it with its principal motivations, but also provide its own set of aspirations and personal goals. What the Web wants for itself, the public presence of global humanity usually gets, and by that token the Web is incapable of hiding even its most secret

desires. Be they good, bad, pure or debase, every one of the Web's visible aspects simply represents a manifestation of mankind's averaged cravings, the need for useful information being just one of those. In such respects the Web is the best mirror we have ever had for our collective mindset. What it shows, we have asked for without question; morality holds little sway in the process of technological advancement.

But just what does this driving force look like? How can it be described? And how can it explain the many faces of the Web?

Maslow's hierarchy of needs is a theory in psychology proposed by Abraham Maslow in 1943, a theory of human motivation which he subsequently extended and which has been constantly updated by others ever since (Figure 1.1). Even so, the core content of his theory is still considered to be a key reference in modern psychological practice, contending that as we humans meet our basic necessities, we seek to satisfy successively higher needs that occupy a set hierarchy.

This model is often depicted as a pyramid consisting of five levels: The four lower levels are grouped together as *deficiency* needs, while the top level is termed as a *being* need. While our deficiency needs must be met, our being needs to continually shape our behavior. The basic concept is that the higher needs in this hierarchy only come into focus once all the needs that are lower down in the pyramid are mainly or entirely satisfied. Growth forces create upward movement in the hierarchy, whereas regressive forces push prepotent needs further down the hierarchy.

As such, Maslow's hierarchy can be categorized as follows:

- **Physiological Needs:** The first need for the body is to achieve a consistent physical and mental state capable of maintaining daily life. This is obtained through the consumption of food, drink and air, achieving adequate sleep, a comfortable temperature, and so on. When such requirements are unmet, a human's physiological needs take the highest priority by default. For instance, if someone were to simultaneously experience the desire for love and a hunger for food, they would rather eat than engage in a sexual act. As a re-




Figure 1.1. Maslow's hierarchy of needs.

sult of the potency of such physiological needs, we tend to treat all other desires and capacities with less urgency.

- **Safety Needs:** When all physiological needs are fulfilled, the need for safety becomes dominant and only in circumstances of extreme threat to physical condition will it overtake the need below it.
- **Love/Belonging Needs:** Once a person's physiological and safety needs are largely met, the third layer of human needs starts to become apparent. This involves emotionally based relationships in general, which includes the perceived need for companionship—both sexual and nonsexual—and/or having a family. There is an inherent sense of community or affiliation in our social makeup; in other words, humans want to belong to groups, whether they be clubs, work groups, religious groups, family, gangs, and so on. We need to feel wanted by others and to be accepted by them.
- **Esteem Needs:** There are two types of esteem need: (1) the need for the respect of, and recognition, by others and (2) the need for self-respect. Such needs can be seen as being motivations or drivers of behavior, as in the instinctual need of a human to make the most of their unique abilities.
- **Self-Actualization:** Self-actualization is the instinctual need of a human to make the most of their unique abilities. Maslow described it as follows: “A musician must make music, the artist must paint, a poet must write, if he is to be ultimately at peace with himself. What a man can be, he must be. This need we may call self-actualization.” Although Maslow tentatively placed self-actualization at the top of his hierarchy, this element has been discounted by most modern psychologists.

What is obvious and apparent about the relationship between the Web and Maslow's famous model is that the Web can assist in satisfying needs at every level in its hierarchy. It may not provide heat, nutrition, and shelter directly, agreed, but it can make the search for such things much easier. Regardless, the Web really comes into its own in the layers above purely physiological associations, and in such respects it is much more directed toward mental, rather than physical, fulfillment. For example, it can easily hypnotize us into a sense of belonging that is falsely secure. It can easily become a surrogate companion to many of those who are somewhat less than socially adept by either character or their own making. It provides an easy option, as it were, a quick fix for large regions of our needs' landscape, and in a culture where living life in the fast lane is becoming the norm rather than the exception, such quick fixes are not only becoming socially acceptable, they also are starting to take the front seat in our wider sociological systems.

Indeed academics have proposed that Maslow's hierarchy can be used to describe the kinds of information that individuals seek at different motivational levels, whether that be via the Web or not. For example, individuals at the lowest level seek coping information in order to meet their basic needs. Information that is not directly connected to helping a person meet his or her needs in a very short time span is simply left unattended. Individuals at Maslow's safety level need helping information. They look to be assisted in finding out how they can be safe and secure.

Enlightening information is sought by those seeking to meet their belongingness needs. Quite often this can be found in materials on relationship development. Empowering information is sought by people at the esteem level. They are looking for information on how their ego can be developed. Finally, people in the growth levels of cognitive, aesthetic, and self-actualization yearn for edifying information and material on how to connect to something beyond themselves [77]. The Web perfectly fits this higher space, openly promoting an image of communal connectedness and belonging. So as the various sectors of society become more comfortable within Maslow's framework, the Web provides a natural accompaniment to their upper motivational processes.

It should also be recognized that the psychological and sociological forces behind the Web have significant power, along with the fact that the source of this power is growing at an alarming rate. In the four centuries from 1500 to 1900, the human population of this planet grew by an average of around three million people per annum. However, in the century from 1900 to 2000, the average yearly increase was close to 44 million—in truth virtually 15 times the previous rate (Figure 1.2). So by such reckoning, some estimates conclude that more people are walking this planet today than all the humans who have ever lived—a profound conclusion if true.

THE DARK SIDE OF THE FORCE

There is, however, one notable and potent exception to the Web's preference for mental, as opposed to physical, stimulation amongst its human hosts. As intelligent beings, one of mankind's earliest discoveries was undoubtedly its capability to achieve sexual stimulation with relative ease, be that as part of the normal process of mating or through other acts of self-gratification. Over the centuries, our various




Figure 1.2. Global population growth.

cultures have controlled such activities through various rules and codes of conduct, often demoting them to a position of unacceptable public behavior. But still they are very much part of a normal human's instinctive behavior, forming an essential component of our drive to procreate the species. Sexual stimulation, after all, is nature's way of thanking us for thinking about its next generation of product, and, as such, it is a reward for which Mother Nature has deliberately created an inbuilt inclination. This is the predominant reason why pornography is so prevalent on the Web. In fact, it is so prevalent that if the Web could be visualized as a being in its own right, large parts of its skin would be colored pink through black to match our own range of exterior pigmentation.

Not wanting to be drawn into any debate about morals, from a purely technological standpoint, pornography has to be considered as a veritable plus for the Web, having forced the pace of change for many a valuable technology in the past. Throughout the history of new media, from vernacular speech to movable type, to photography, to paperback books, to videotape, to cable and pay TV, to CD-ROMs and laser discs, pornography has shown technology the way [36]. It may represent the darker side of the Web's public acceptability, but it is certainly playing a key role in its expansion and further evolution. Pornography draws curiosity seekers who stay to see what else the new media can do. And there is a convenient dovetailing in the audience for computers and pornography: Young Western males dominate both markets. Gadget-playing, girl-crazy young men will stay longer at a terminal that supplies both girls and gadgets—a sad indictment of the Web success, perhaps, but a hard reality nonetheless. Furthermore, it is from this pool of similar social demographics that most practicing technologists originate. Those who have and will make a difference in the advancement of the Web's technologies in the foreseeable future come from the predominantly caucasian economic states of the Western world. This is not a matter of racial prejudice or preference on the Web's part, but is simply an issue of macroeconomics. The Web goes where the funding is, it's as simple as that, and this in itself is causing a number of growing pains.

WHAT WOULD A WEB BE WITHOUT THE HOLES?

Although the Web certainly has the capability to be globally pervasive through the technologies and infrastructure coverage offered by the Internet, it is not truly this way yet, and significant gaps in its potential to reach a full global audience still exist. Web pages cannot be created using many character sets with regional specificity, for instance. For this reason, you will not find any native Web content including text in Ethiopic—often referred to as *Ge'ez* or *Classic Ge'ez*—one of the ancient languages of Ethiopia which is now mainly used in Ethiopian Orthodox Church as a liturgical language [37]. This is an issue predicted by the very fact that at present there is no universally recognized Ethiopic character set.⁷

⁷There are concerted ongoing efforts by organizations like the World Wide Web Consortium (W3C) to address such internationalization problems. For more information, please refer to <http://www.w3.org/International/>

This particular deficiency is not because the demand for such material is low, quite the contrary, it is purely because, up to this point, there has not been sufficient funding or knowledgeable interest to make the development to make such capability a reality—a symptom common to most, if not all, of the Web’s current technological concerns.

STRUCTURE ABOUNDS

There are those who would argue that the Web is a completely unstructured place. One need only take a quick look at the NASDAQ Most Active list to support this fact, they will say, adding that Web portals and the search engines exist primarily because the Web is a tremendously disorganized space, a system where the disorder grows right alongside the overall volume. Yahoo and Google, for example, function (in a way) as man-made antidotes to the Web’s natural chaos, being engineered attempts to restore structure to a system that is incapable of generating structure on its own. This is an often-noticed paradox of the Web: “The more information that flows into its reservoirs, the harder it becomes to find any single piece of information” [42].

Without a doubt, it is easy to see why the less experienced Web user might form such an opinion. The Web is certainly an impossible thing to digest in one sitting, and it isn’t hard to become overwhelmed by its dominance without the help of some pretty clever tools by your side. Furthermore, it is certainly true that it is precisely for such reasons that search engine technologies have become so popular, providing a relief map of the Web based on information hewn straight from its rock face. Even so, an unstructured view of the Web’s reality is still somewhat a naive interpretation of its form, being, in truth, almost as far from the actual reality as one can get.

The Web is not without structure, it simply has too much of it. It is our fault that we do not understand that vast amounts of information become a complex “interwinding,” regardless of whatever medium it is stored in. This is due to a fundamental flaw in our mental capability. People just keep pretending that they can make complex things deeply hierarchical, categorizable, and sequential when they can’t accomplish this directly; that is a task for raw complexity to achieve itself. All information becomes deeply intertwined beyond a certain point, and then such classification types begin to drift off into higher levels of framework by natural tendency. It’s not that they disappear beyond recognition, it’s just that there are lots of them, overlapping, merging, and generally making things difficult to comprehend from a distance. In reality, it’s simply a question of not seeing the forest for the trees.

Take any two pieces of Web-friendly content, a couple of simple text-based Web pages say, and have one point to the other via a hyperlink. By this very act, structure is automatically introduced. At a finer level of detail, even the very text of each page has its own structure without the need for hyperlinks, presumably following the grammatical rules of the natural language in which it is written. In such ways the Web exudes structure to the point where it can be extremely counterproductive, often confusing its use or definition for any singular purpose. Actually that’s a little

unfair, because the structure we are discussing here is not just any run-of-the-mill architecture. Rather it is a composite of multiple types, layers, and hierarchies of more simplistic configurations spreading out in every direction. One could never consider the Web to be a singularly hierarchical framework, for example, although it does undoubtedly embody many thousands, if not millions, of items amassed by such configurations. In a similar manner, it is not solely a peer-to-peer, hub-and-spoke, distributed-parallel, topological, or any other type of organized system; it is all of them rolled into one complex, all encompassing whole, for many reasons and across many dimensions. But this is not just one almighty and unruly mess, because each structure carries its own purpose and identity. Furthermore, many fields of science and philosophy have a name for such arrangements. The Web's structure is in many senses a collection of "ontologies," a systematic account of existence in its own right.

In information engineering, such ontologies are the product of an attempt to formulate a conceptual description about a topic or area of interest. They are typically structures containing all the relevant entities and their relationships and rules within that area. Even so, the technological usage of the term "ontology" is derived from the much older usage of the term in philosophy, where it is also used to refer to the study of existence. The purpose of a computational ontology is therefore not to specify what does or does not 'exist', but rather to create a corpus of details containing concepts that refer to entities of interest to the examiner and that will be useful in performing certain types of computerized investigations. For this reason the rationales used by philosophical ontologists can be helpful in recognizing and avoiding potential logical ambiguities.

An ontology that is not tied to one particular area of interest, but attempts instead to generally describe things and the relationships between them, is known as a 'foundation' or 'upper' ontology, and it is from such a base that the general setup of the Web's organization is formed. This is a large open framework within which many more specialized, domain-specific ontologies are enclosed to make the artifacts presented more useful for real-world purposes. These smaller structures are hence deliberately tied to a specific subject, audience, or other kind of distribution, like the websites constituting the eBay organization for instance or the various pages on any company's intranet.

Ontologies can be referred to as being explicit specifications of particular conceptualizations, and it is across such specifications that the various search engine spiders and bots roam to collate their own interpreted indexes of the Web. In doing so, however, they only pick up nuances that suit their own particular purposes, unless, that is, they are specifically redirected. For this reason it is dangerous practice to consider the results returned by any particular search engine to be an absolute reference of the Web's true character. Google may be the Web's best ally from an approachability point of view, but it might not be considered the best candidate to act as its biographer, because it may take many more generations of new Web technologies before such a role can be fulfilled adequately. Even so, when the Web's story does eventually mature into an account worthy of accurate representation, ontological structure will almost certainly be one of the materials from which that story will be spun.

One new set of technologies that might help establish an accurate interpretation of the Web's true form in the long run are those that belong to its latest offspring, the Semantic Web. Here much more formal types of ontology are the norm, allowing far greater and more accurate levels of descriptiveness amongst the Web's content. The overall intent behind the Semantic Web is not only to make the general Web more accurate for human consumption, but to increase its capability for direct understanding by computers themselves. If successful, it will be like the Web folding back upon itself, making it capable of seeing its likeness, warts and all.

FOUR DIMENSIONS ARE NOT ENOUGH

A surreal, but nonetheless true, characteristic of the Web, as all conceptual ontologies, is the fact that it is quite literally not of this world. It is materially just a collective made up of an unimaginable number of magnetic fluctuations on the surfaces of the world's various disk drives and memory chips. In truth, in its native form it has little practical physical presence at all; you can't see it, touch it, smell it, or even taste it, and in such respects it is a truly ethereal thing. To purposefully misquote Grady Booch,⁸ the famous Software Engineer, it is "invisible to most of the world. Although individuals, organizations, and nations rely on a multitude of Web-intensive systems every day, most Web content lives in the interstitial spaces of society, hidden from view except insofar as it does something tangible or useful" [57]. Furthermore, because of its physical absence the Web does not even necessarily conform in usual ways to the dimensions of space and time. For example, the fact that one Web page could theoretically link to another physically located in either the next room or the next galaxy makes no conceptual difference at all to the Web's constitution. It may take much longer for the page furthest away to render in a browser, but this is more a concern for the physical scaffolding of the Internet that underlies the Web, rather than the Web itself. This also points to that notion that the Web works on a different timescale to our own. We may see Web pages appear in front of us in the blink of an eye, but this is merely a microscopic detail, again more applicable to the Internet rather than the Web as a whole. In the Web's sphere of existence, "life" plods along at a much slower speed, synchronized to the averaged rhythm of change coming from humanity's collective stumbling en masse.

None of this means to say that the Web does not embody or aspire to comply with any of the Universe's physical laws, as we shall see from Chapter 3 onwards. It is simply a quantum leap in our joint technological capability which we do not fully understand yet. As such, it may take us some time to refocus on just what exactly is taking place inside this new technology's metabolism. The term "quantum leap" is not intended to be used as cliché here, because it perfectly describes the ideas involved. The change in mindset required to understand and harness the real power of the Web is similar to the difference in perspective needed to appreciate both Newtonian and quantum mechanics. Furthermore, this comparison has been

⁸Grady Booch is best known for developing the Unified Modelling Language with Ivar Jacobson and James Rumbaugh.

investigated seriously on a number of occasions already. In one example, a fascinating result was found by Jon Kleinberg, a computer scientist at Cornell University who discovered that, when the matrix of the Web is analyzed like a quantum mechanical system, stable energy states correspond to concepts under discussion [53,54].

OUR UNDERSTANDINGS OF LIFE

So what about the Web as a living thing? How can we name it so? In fact, how can we even identify and distinguish life itself, even in its most obvious of forms?

It would be foolish to pretend that we are anywhere near to understanding the true essence of life. There might be the odd fleeting moment when science unlocks another of nature's secrets, but life hides its secrets well and we soon realize that we still have an additional million locks to go. In truth we only really understand a small number of life's characteristics, a handful of its mechanics, a few of the words spoken in its native tongue as it were. But still we try, we aspire to understand, and slowly, ever so slowly, we are learning. In doing so, however, some of our most fundamental beliefs about life are being challenged.

In recent years, some have gone so far as to suggest that life may not be uniquely confined to a physical existence, for example contemplating the possibilities of consciousness and intelligence originating from purely virtual environments. The Web is most certainly one such environment, and this raises a number of profound questions. For a start, what observable properties would be required to know for certain if life existed on the Web or if the Web itself were alive? Furthermore, and perhaps more profoundly relevant, if the Web is not alive yet, but the expectation is that it might soon well be, how close to producing life are its relevant properties at the moment? For sure, these are complex and difficult questions that cannot easily be answered in one go, but there is encouraging evidence to suggest that we are not totally unequipped to tackle such conundrums. Ever since our wisest realized that they were different from the inanimate rocks that form our planet, we have tried to comprehend the very essence of our own vitality, collecting a wealth of knowledge along the way. Today we have definitive classifications for what is alive and what is not, but these mostly relate to our classical interpretations of being. Certainly they provide strong points of reference, but perhaps the real issues at hand today have now moved on. Perhaps now we should be less interested in the bipolar definitions of old and more interested in blurred edges of reality where new life forms might conceivably begin and end.

LIFE'S PLAYGROUND—A UNIVERSE OF INFINITE POSSIBILITIES

From early childhood we are urged to achieve; “eat all your greens,” “be a straight ‘A’ student,” “do the best you can” should all be familiar chants from one of society's most common mantras. Our lives are filled with goals, each one specific to a

certain property of our own being. But even though we as individuals might strive to excel at anything we do, ask even the world's most successful to define "what it takes to be the best" and heads will start to scratch. "Five years training and complete dedication", an Olympic gold medalist might volunteer, "a lifetime of serenity and prayer" could easily be the answer from a devout Buddhist, while "courage, complete conviction and the willingness to go that extra mile" might well be the retort from a world leader. All are good answers, but all are nonetheless equally misdirected from nature's point of view.

If Mother Nature were privy to such testament, she would no doubt raise a wry smile, for she has indeed learned the true answer to this question through countless generations of hard-earned experience. The solution to becoming the best is no real secret and, in fact, hides a deceptively simple strategy: Just find out who or what is the best and make sure that you are better. That's all there is to it. And it doesn't make the slightest bit of difference how hard you try along the way or how long it takes you to get there. You could turn out to be a million times more superior and become so in the blink of an eye, but still, in the natural world, just being "good enough" has always been, and always will be, good enough. There are no grades handed out, no fixed pass marks to point to, and few second chances. Evolution has no long-term goal other than to make all life forms better than they were a second ago. "There is no long-distance target, no final perfection to serve as a criterion for selection, although human vanity cherishes the absurd notion that our species is the final goal of evolution" [51]. Thus there are no definite targets, just complex gray regions of competition and influence leading to black or white outcomes based on pure survival. Evolution does not care about what makes sense; it cares about what works.

But why should this be so? Why should the concept of targets be so irrelevant in the natural world? The real answer is that the Universe is an immensely large space of possibilities [67]. In fact, it is as close to being infinite in such respects as to make the very concept of boundaries practically irrelevant. Hence, in our material Universe, no matter how expansive, adaptive, or progressive any given system, it will never be constrained solely by its environment unless it attempts to transcend all but the most fundamental of its laws. Take for instance the cheetah, the fastest land animal known to exist. It could evolve and theoretically continue growing faster and faster over countless generations. On its own, it would only stop accelerating as a species when the most basic of universal laws—gravity, friction, and so on—prevented its muscle mass from propelling it any faster over its natural terrain. Moreover, without such restrictions, only the speed of light itself would stand in the way of its ever-increasing speed.

Look also at what happens with a simple species of seaweed, containing 1000 interconnected genes. To be sure of finding the highest level of adaptation to its environment, natural selection would have to examine every considerable combination of genes and appraise all conceivable mutations in order to determine the best possible variant of the species. When the total number of combinations of gene mutations is worked out, the answer isn't two multiplied by 1000, it's actually two multiplied by itself 1000 times. That's 2^{1000} , or roughly 10^{300} —a number so vast that it dwarfs even the number of potential moves in a chess game. Quite simply, the number is huge and evolution could not even begin to try out that many alternatives, not

even in the time taken for the entire Universe to develop thus far. And remember that's just for seaweed. Humans and other animals roughly have 100 times as many genes, and most of these genes come in many more than two varieties [67].

In short, most adaptive systems such as life explore their way through this immense space of possibilities with no realistic hope of ever finding the single 'best' place to be [67]. They do this by a simple process of trial and error, fumbling their way to find their most appropriate evolutionary path. All that evolution can do, in essence, is look for improvements, not perfection. Hence it either builds upon the successes of previous generations or abruptly stops at the outright dead end of extinction. Through such hierarchies of successes, building blocks upon building blocks of slow progress are constructed, revealing a deep secret—namely, that hierarchy utterly transforms a system's ability to learn, evolve, and adapt [67]. But there is a price to pay for evolutionary success and higher orders of environmental fit. Through the compounded advances piled generation upon generation also comes at least one side effect; complexity. In all but the most specific of cases, this increases as a direct result of evolutionary advancement, applying its own rules and constraints from the ground up in its wake.

So, nature's systems are complex by definition and, given the meager amount of time we as a race have had to get our act together so far, it is not surprising that the man-made systems in our grasp are mostly simple by comparison. But here again there are tradeoffs. By being simple, the fabrication of our systems can be precise and efficient, whereas in comparison the complex, adaptive machinery of nature mostly cannot. A truly complicated structure has many masters, and none of them can be served exclusively. Rather than strive for optimization of any one function, a complex system can only survive by pacifying a magnitude of goals. For instance, an adaptive system must participate in a tradeoff between exploiting a known path of success, thereby optimizing a current strategy, or diverting resources to exploring new paths, thereby wasting energy to try less efficient methods.

So vast are the blended drivers in any complex natural system that it is often impossible to unravel the actual causes of its continued success or survival. Survival is a journey with many different and sometimes contradictory endpoints. Hence most living organisms are so multifaceted that they are blunt variations that happen to work, rather than precise renditions of their ultimate purposes. It is one of nature's unwritten rules; in creating something from nothing, as is nature's origin, forget elegance; if it works, it's beautiful [21]. That said, evolution's overriding cause is obviously to seek out the best possible solutions to any given problem in all possible situations; it achieves its peak fitness as it were. But in so doing, nature is always running blind. It only knows and cares about the condition of its current world and all the imperfect notions of "best" within it, infinitely stumbling in the dark to find any new options capable of stealing pole position. It is such stumbling that is evolution's defining quality. If, for example, one were theoretically able to ask it for the answer to an apparently predictable question, say "how can you cross a stream if there are three boats of different size in front of you?", completely nonsensical answers should not be unexpected. "Ignore the boats and dig a tunnel under the stream" might be one solution suggested—a ridiculous proposition maybe, but nevertheless a plausible solution. This seemingly haphazard approach to problem-solv-

ing is one of evolution's hallmarks, explaining why life continues to enlarge its own being [21]. Nature is simply an ever-expanding library of possibilities positioned in an open Universe of endless possibilities.

At first sight our modern digital world might appear to be totally different. As a species we like to think that we have progressed far beyond the point where we had to fight for our existence on the Great Plains, with our subconscious instinct for sustainability driving the search for the security and protection of certainty in our surroundings. To know our world and control it through that knowledge has been a powerful ambition of society ever since we crafted our first stone axe. Each time we have created such a new technology, we have pushed it, honed it, until it far exceeds the "just good enough" requirements of the natural world around us. We control it with ever-greater levels of precision and so involuntarily strive to push our lead in evolution's natural pecking order. In fact we have become so good at this that today we are shrouded in a world so locally precise, so digital and powerful, that we could easily account for every second of every human life on this planet and still have plenty spare computational capability left over. Nevertheless, although ours might appear to be a world destined for precise digital servitude, this is actually a somewhat naive view of our real contemporary technical context. Take a step back for a moment, in fact take 10 steps while you are at it, and you will soon see the abundance of controlled accuracy start to mount into an ever-changing, interconnected pool of swirling grayness, a swirl of interconnected system upon interconnected system not dissimilar to that from which we once fled—the imperfect and complex swirl of life itself. And the ultimate example of such interconnected technical co-evolution? Enter the Web we see today.

ENQUIRIES INTO THE DEFINITION OF LIFE

In 1872, H. C Bastian wrote: "Amongst the carbon compounds we find in our world there is an abundance of evidence to prove the existence of internal tendencies or molecular properties which may and do lead to the evolution of more and more complex chemical compounds. And it is such synthetic processes, occurring amongst the molecules of colloidal and allied substances, which seem so often to engender or give 'origin' to a kind of matter possessing that subtle combination of properties to which we are accustomed to apply the epithet 'living'" [28].

The word "life" has probably been around ever since mankind began using language. It is a word of fundamental importance to all of us, and seldom do we make it through an entire day without putting it to use. We do so, however, with only a sketchy and subjective idea of what life actually means. This is because until recently, within the last century or so, it has been easy for people to distinguish between what they call living and what they call nonliving. There has been no need to define life precisely; its meaning has typically been intuitively understood [28].

The scientific revolution of the past few centuries has brought complications into this matter, and certain new fields of science are inherently concerned with systems that would be called living by some and nonliving by others. For example, advances in medicine have brought about the discovery of microscopic 'infectious agents'

such as viruses and plasmids. Are these alive? Many scientists are presently searching for extraterrestrial life. How will they know when they have found it, if life is not clearly defined [28]?

Scores of researchers have been, and still are, devoting their time to the study of biochemical evolution to unravel the mystery of the origin of life on Earth. Though even the ancients concerned themselves with this subject, not until recently have theories in this area depended on chemical systems whose status as living or nonliving is open to debate. Linus Pauling, one of the premier chemists of the twentieth century, once said “In connection with the origin of life, I should like to say that it is sometimes easier to study a subject than to define it.” Up to this point, most scientists have taken Pauling’s ‘easier’ route and avoided the issue of coming to a consensus on definition. As the *World Book Encyclopedia* has put it: “Rather than trying to define life precisely, biologists concentrate on deepening their understanding of life by studying living things.” So, how do they know what to study? What does the scientific community consider to be that subtle combination of properties needed to set all living things apart?

There are what might be considered to be classical properties of life, according to standard reference material. The 1984 *Random House College Dictionary* defines life as: “The condition that distinguishes animals and plants from inorganic objects and dead organisms, being manifested by growth through metabolism, reproduction, and the power of adaptation to environment through changes originating internally.” This latter property refers to the phenomenon of homeostasis, whereby an individual organism changes itself in response to a change in its surroundings. In other definitions, this is referred to as “response to stimuli” or just “responsiveness.” Homeostasis is not to be confused with response of the species to environmental changes through the process of natural selection. That is evolution, and it comes about through the transmission of random mutations in the organism to its offspring. This ability to transmit mutations during reproduction, and thus be subject to the processes of natural selection, is a criterion of life cited by many.

The *Encyclopaedia Britannica* concentrates on metabolism in its biochemical definition of life: “An open system of linked organic reactions catalyzed at low temperatures by specific enzymes which are themselves products of the system.” Some references also include movement against a force in addition to the other criteria. This may include locomotion, dynamics, or, in the case of most plants, growth against the force of gravity. The transfer of matter is another standard criterion listed. The consumption of raw materials and the excretion of waste materials are natural consequences of metabolism.

Hence, in general, life has traditionally been characterized in terms of growth, reproduction, metabolism, motion, and response through homeostasis and evolution. The phrase “in terms of” is used here because each of the criteria mentioned are exhibited by different systems to varying degrees. The caveats in such definitions are made evident by two sets of these systems:

- Systems that we consider to be alive but that don’t exhibit all of the classical properties
- Systems that we consider nonliving but that exhibit these properties

The classic example from the second set is fire. It grows, moves, metabolizes, consumes, transforms, and excretes matter, reproduces, and responds to stimuli—for example, wind. Crystals will grow in a saturated solution and reproduce more of their kind. The hydrosphere moves, as in flowing rivers, dissolves compounds and precipitates out different forms, reproduces from rain, and responds, by breaking a dam, for example. Similar analogies have even been carried as far as free radicals, a species of individual molecules that reproduce and that can grow into large polymers.

The nonreproducing mule is an often-used example from the first set. It does not take part in the process of natural selection, yet few would deny that it is alive. Various seeds, spores, and even insects can lay dormant for years without moving, growing, metabolizing, or reproducing, yet they are considered by most as part of what we call “life.” This dormant life is swept under a convenient carpet commonly referred to as “cryptobiosis,” meaning ‘hidden life’.

There is a third category whose anomalous systems defy classification into either set. One system in this category is the virus. On one hand, it can be said that “viruses are not living organisms because they are incapable of independent existence” and must use their host’s cell’s metabolic machinery in order to reproduce. On the other hand, it can also be argued that it seems unreasonable to deny that viruses are living just because they need help to do so. From the latter viewpoint, the virus is considered to be a parasite, a kind of microscopic leech. In many respects too, the Web fits into this category, because in its current state it is surely incapable of surviving without our support.

Certain man-made systems are claimed to be forms of “protolife” in a classical sense, implying that they may have been the first forms of life on Earth. One category of these includes the coascervates and microspheres. These are produced through combinations of various immiscible liquids—for example: Sidney Fox’s⁹ thermal proteins in water. These have been shown to form “cells,” also referred to as droplets, which grow and move about, bud and divide, and aggregate. Their membranes exhibit selective permeability and even catalytic activity.

Experiments have also been carried out with clay minerals. These spontaneously grow in layers. Various cations—an ion or group of ions having a positive charge and which characteristically move toward the negative electrode in electrolysis—when substituted into the silicate lattices of these clays, produce catalytically active sites. These active sites are ‘reproduced’ in subsequent layers that form on the clays. Thus they “are capable of replicative self-multiplication” and are subject to the process of natural selection.

⁹Arguably, Sidney Fox’s best-known research was conducted in the 1950s and 1960s, when he studied the spontaneous formation of protein structures. His early work demonstrated that under certain conditions amino acids could spontaneously form small polypeptides—the first step on the road to the assembly of large proteins. The result was significant because his experimental conditions duplicated conditions that might plausibly have existed early in Earth’s history.

Further work revealed that these amino acids and small peptides could be encouraged to form close spherical membranes, called microspheres. Fox has gone so far as to describe these formations as protocells, because he believed that they might be an important intermediate step in the origin of life. Microspheres might have served as a stepping stone between simple organic compounds and genuine living cells.

The systems mentioned here are just some of those that provide evidence that the classical definitions of life are not appropriate in many cases.

Some claim that there is a continuum of life complexity, with simple inorganic systems at one end and the highest life forms at the other, with the types of systems mentioned above lying somewhere in the middle of this continuum. It is therefore believed by some that there is no point along the continuum of existence from the simplest atom to the most complex animal, at which a line can be drawn separating life from nonlife. Consider the electromagnetic spectrum as an analogy: At a wavelength of 520 nm, light is green; at 470 nm, it is blue. But at what wavelength does it change from green to blue? There are an infinite number of blue-green shades in between the two extremes. The drawing of lines between green and blue, between living and nonliving, is said to be arbitrary, a matter of personal preference. Others take the stance that life is just an aspect of man's perception of matter, just as music is an aspect of his perception of sound. It is purely subjective which sequences of sounds will be perceived as music to an individual. This type of thinking has led authors such as Josephine Marquand to conclude that it is prudent to "avoid the use of the word 'life' or 'organism' in any discussion of borderline systems."

One way to draw the line between life and nonlife in such a continuum of complexity is to pick out a specific structural feature common to systems considered to be alive. The first such documented definition was that "all living systems are composed of cells." In his study on the origin of life in the 1920s, A. I. Oparin proposed that "only in discrete particles could the random chemical activity occurring in the waters of the earth be organized into harmoniously correlated chemical reactions." An effect of this definition would, obviously, be its converse: non-cellular systems are not living. But we cannot, however, deduce from this definition that all cellular systems are alive: the oil-vinegar emulsion in your salad dressing is composed of many cells, but is obviously at the nonlife end of the continuum.

Another feature of classical living systems is the ubiquity of the linear polymers of amino acids we call proteins. These make up much of the structure and catalytic machinery associated with life on Earth. Also ubiquitous are the polymers of nucleotides, the nucleic acids. DNA and RNA contain the genetic information that is passed on to the offspring of living organisms. Classical carbon-based life has been defined as that which makes use of or produces proteins and/or nucleic acids. More generally, living systems can be defined as those which exhibit optical activity and isotopic fractionation. The former refers to the rotation of polarized light by solutions of organic molecules. This is due to the chirality¹⁰, or handedness, of these molecules. Often, when a molecule has two possible mirror images, a left one and a right one, only one of these forms is found in nature. For example, only left-handed amino acids are used in proteins and only right-handed sugars are stored for energy. Isotopic fractionation implies that life forms selectively pick out certain isotopes of the elements. This results in relative concentrations of these isotopes within them that differ from the relative concentrations in other systems—this provides the basis for carbon-14 dating. These definitions are not foolproof; some inorganic and man-

¹⁰Of or relating to the structural characteristic of a molecule that makes it impossible to superimpose it on its mirror image.

made systems also display these properties. Crystals of one handedness are produced routinely by chemists, and phase changes such as evaporation can result in isotopic fractionation.

If you grew up in Canada, you might have been led to believe that only pine trees retain their foliage during the winter. If you then defined a pine tree as ‘a tree which stays green in the winter,’ you would run into trouble when traveling to California, where trees that are not even part of the conifer family stay green all year. In such a way, definitions of life here on Earth may be of a similar nature. They are based on the particular, most likely random, complex mechanism by which life arose in this locale. Nevertheless, life may have arisen elsewhere or earlier on Earth under different circumstances. It may then have lacked the familiar “life signs” we recognize today, yet still be considered life by other criteria, such as reproduction or metabolism.

Many definitions of life echo the belief that “living organisms are distinguished by their specified complexity.” In such definitions, the terms order, information, and complexity are frequently encountered. How are these concepts related? Complex structures and systems abound in the Universe, but most of them—the nonlife, by this definition—are random, that is, they are not specified. We could define the ‘information content’ of a structure as the minimum number of steps needed to specify that structure. To specify a random polypeptide—a sequence of amino acids—we need only state the proportions of the amino acids that go into making it. To specify a certain enzyme, however, we must state which acid occupies each position in the enzyme’s sequence. This takes many more steps for an enzyme of the same length as the random polypeptide. Thus enzymes have higher information content. “Order” has been defined as “situations that are unlikely to occur by random processes.” As the length of a specified protein goes up, the chance of producing that protein by hooking amino acids together randomly goes down exponentially. For a short protein of 10 amino acids, for instance, there is a probability of about 1 in 100 trillion that it would be formed randomly. Since proteins in organisms are usually substantially longer than this, we can hence state that they are highly ordered.

Life has also been defined not in terms of its order or information content, but in terms of its ability to transmit this information to its descendants—its ability to replicate in other words. The human chromosome has about 10 billion bits of information in it, and the half-set of these carried by a single sperm cell contains enough information to fill 500 volumes of a large book. Certain viroids reproduce by transferring an RNA string of less than 400 nucleotides to their host cell. While this is still a substantial amount of biochemical information, it is argued that viroids should not be considered living because they require not only the nutrients of their host, but also the information contained in the host cell’s reproductive machinery. Francis Crick included as a “basic requirement of life” the ability of the system to replicate both its own instructions and any machinery needed to execute them.

Erwin Schrödinger, the famous physicist, analyzed life from a statistical perspective. He noted that in all inorganic systems, it takes a statistically large number of molecules to produce a predictable result. In his “What Is Life?” lectures of 1943, he claimed that the characteristic of life is that it seems to defy the rules of

statistics. Statistically speaking, in a very small number of molecules, the genotype predictably governs the structure and function of a whole organism, the phenotype. He claimed that “this situation is unknown anywhere else except in living matter.”

The second law of thermodynamics states that, in any process, the total amount of entropy—apparently observed randomness—in the Universe must increase. This is a direct result of the natural tendency of the Universe toward equilibrium, a state of maximum disorder. Schrödinger wrote “It is by avoiding the rapid decay into the inert state of ‘equilibrium’ that an organism appears so enigmatic.” How, then, do we account for the processes of life, which seem to create order out of randomness? The answer lies in the fact that living things are open systems; that is, they exchange matter and energy with their surroundings. For every bit of order created within them, a greater amount of disorder is created in their surroundings. The process of building your body produces a great deal of heat, which causes the air around you to become more disordered. Thus, such processes stay within the bounds of the Second Law.

It is through the exchange of energy that life avoids the dreaded disorderly ‘equilibrium state’. Therefore, life’s “exquisite regulation of energy flow” has often been included in its definition. One such definition includes the flow of energy within the organism, as opposed to between it and its surroundings: “Life is a group of chemical systems in which free energy is released as a part of the reactions of one or more of the systems and in which some of this free energy is used in the reactions of one or more of the remaining systems.” The term “free energy” here refers to energy that can be put to use, as opposed to heat energy lost to the environment.

MORE CONTEMPORARY VIEWPOINTS ON LIFE

Certain groups, particularly those interested in the possible nature of extraterrestrial and artificial life, consider the preceding definitions for life to be too limiting. For example, they ask if it is relevant to differentiate between an individual organism and the entire biosphere, or between a particular website and the entire Web for that matter? A bacterium could easily mistake a person for a huge colony of one-celled organisms working in symbiosis, and, in the same way, a person could quite easily perceive the Web as a huge virtual colony of purposeful information-based technology. As such, one definition of life, from Feinberg and Shapiro’s *Life Beyond Earth*, is “the activity of a biosphere” [34], and they define a biosphere as “a highly ordered system of matter and energy characterized by complex cycles that maintain or gradually increase the order of the system through an exchange of energy with its environment.” The presence of the now familiar terms of order, complexity, and energy should be highlighted. Other classical notions associated with life, such as reproduction for example, are not required in this definition. So Feinberg and Shapiro propose that it might be more profitable for an organism to alter itself to adapt, rather than wait for randomly altered descendants to undergo the process of natural selection. In a perfect biosphere, with all elements in symbiosis, evolution of the parts tends to be detrimental.

Feinberg and Shapiro also allow room in their definition for the existence of

physical life, as opposed to chemical life. Examples they propose include plasma life, nuclear life, and radiant life. Plasma life would exist inside stars, where interactions between charged particles and magnetic fields would create self-sustaining, orderly systems. One example of nuclear life would inhabit a very cold planet. It would be composed mainly of solid hydrogen and liquid helium. The spins, or magnetic orientations of the hydrogen nuclei in the organism would be highly ordered. Magnetic fields caused by this organization of spins would induce further organization. Radiant life might inhabit interstellar nebulae, which are made of the dusty remnants of dead stars. This type of life is based on the properties of ordered radiation, using space dust as a tool for transforming the radiation. This can be viewed as an organized collection of self-stimulating lasers. It is interesting, however, that Feinberg and Shapiro do not include the possibility of computer-based life in their work. Perhaps this is just an accident of timing, given that they published in 1980, just before the use of computer simulations of life really took off and well before the Web materialized in its current form. Nevertheless, the 1980s and 1990s did indeed see a great upsurge of interest in the understanding of complex systems such as life, an upsurge that led to the front door of computing for many reasons.

The goals of creating artificial intelligence and artificial life can be traced back to the very beginning of the computer age. The earliest computer scientists—Alan Turing, John von Neumann, Norbert Wiener¹¹, and others—were motivated in large part by visions of instilling computer programs with intelligence, with the life-like ability to self-replicate, and with the adaptive capability to learn and control their environments. These early pioneers were as much interested in biology and psychology as in electronics and logic, and they looked to natural real-world systems as guiding metaphors for how to achieve their visions. Thus it should be no surprise that the earliest electronic computers were applied not only to calculating missile trajectories and deciphering military codes but also to modeling the brain, mimicking human learning, and simulating biological evolution. These biologically motivated initiatives have waxed and waned over the years, but since the early 1980s they have undergone a resurgence in the computation research community. This has led to great progress in such fields as neural networks, classifier systems, artificial intelligence, genetic algorithms, evolutionary computing, and ultimately the study of artificial life, a computational concept pioneered by Chris Langton.

Artificial life, says Langton, is essentially just the inverse of conventional biology. Instead of being an effort to understand life by analysis—dissecting living communities into species, organisms, organs, tissues, cells, organelles, membranes, and finally molecules—artificial life is an effort to understand life by synthesis: putting simple pieces together to generate lifelike behavior in man-made systems. This can be either intentional, as in the case of Langton's experiments, or completely unintentional, as is the case with systems like the Web. Its belief is that life is not a prop-

¹¹Norbert Wiener was a U.S. mathematician, known as the founder of cybernetics. He coined the term in his book *Cybernetics or Control and Communication in the Animal and the Machine* (MIT Press, 1948), widely recognized as one of the most important books of contemporary scientific thinking. He is also considered to be the first American-born-and-trained mathematician on an intellectual par with the traditional bastions of mathematical learning in Europe. He thus represents a watershed period in American mathematics [93].

erty of matter per se, but the organization of that matter. Its operating principle is that the laws of life must be laws of dynamical form, independent of the details of any particular carbon-based chemistry that happened to arise here on Earth four billion years ago. Its promise is that by exploring other possible biologies in a new medium—computers and perhaps robots for instance—artificial life researchers can achieve what space scientists have achieved by sending probes to other planets: a new understanding of our own world through a totally different perspective gained from other worlds. “Only when we are able to view life-as-we-know-it in the context of life-as-it-could-be will we really understand the nature of the beast,” Langton has declared [69].

The idea of viewing life in terms of abstract organization is perhaps the single most compelling vision to come out of early collaborative work on artificial life. And it’s no accident that this vision is closely associated with computers, given that they share the same intellectual roots. Human beings have been searching for the secret of automata—machines that can generate their own behavior—at least since the time of the Pharaohs, when Egyptian craftsmen created clocks based on the steady drip of water through a small hole. In the first century A.D. Hero of Alexandria produced his *Treatise Pneumatics*, in which he described, amongst other things, how pressurized air could generate simple movements in various gadgets shaped like animals and people. In Europe, during the great age of clockworks more than a thousand years later, medieval and Renaissance craftsmen devised increasingly elaborate figures known as ‘jacks’, which would emerge from the interior of a clock to strike the hours; some of their public clocks eventually grew to include large numbers that acted out entire plays. And during the Industrial Revolution the technology of the clockwork automata gave rise to the still more sophisticated technology of process control, in which factory machines were guided by intricate sets of rotating cams and interlinked mechanical arms. Moreover, by incorporating such refinements as movable cams, or rotating drums with movable pegs, nineteenth-century designers soon discovered controllers that could be adjusted to generate many sequences of action from the same machine. Along with the development of calculating machines in the early twentieth century, notes Langton, “the introduction of such programmable controllers was one of the primary developments on the road to general-purpose computers” [69].

Several years prior to Langton’s work, the foundation of a general theory of computing had been laid by logicians who had tried to formalize the notion of a procedure, a sequence of logical steps directed to some useful purpose. This effort peaked in the early decades of the twentieth century with the works of Alonzo Church, Kurt Gödel, and Alan Turing—of whom we shall hear much more later—along with many others, who pointed out that the essence of mechanical process, the “thing” responsible for its behavior, is not necessarily a thing at all. It is a set of rules without regard to the material of which the machine is made and in such respect is truly ethereal. Indeed, comments Langton, this abstraction is what allows us to take a piece of software from one computer and run it on another: The “machine-ness” of the machine is in the software, not the hardware. Furthermore, once you have accepted this, then it is a very small step to say that the “aliveness” of an organism is also in the software—in the organization of molecules, not the actual molecules themselves [69].

Admittedly, that step doesn't always look so small, especially when you consider how fluid, spontaneous, and organic life can be and how controlled computers and other machines are. At first glance it seems ludicrous even to talk about living systems in those terms [69]. But the reality behind this truth lies in a second great insight, namely that living systems are machines—that is, machines with a kind of organization and structure different from those we are used to in the “every day” sense of the word, but machines nevertheless. Instead of being designed from the top down, using a standard reductionist approach, the way a human engineer might do it, living systems almost always seem to emerge from the bottom up, from a population of much simpler systems and parts, in such a way that a computer program is the summation of all its constituent instructions and the Web is the conclusion of all its contributing resources. In biological life a cell consists of proteins, DNA, and other biomolecules, a brain consists of many millions of neurons, an embryo consists of interacting cells, an ant colony consists of a multiplicity of ants, and for that matter an entire economy consists of nothing more than firms and individuals [69]. All are simply examples of purposeful aggregates structured from the ground up.

In fact, one of the most profound and surprising lessons learned over the past few decades, as we have begun to simulate evermore complex physical systems on computers, is that complex behavior need not have complex roots. As Chris Langton has pointed out many times, “tremendously interesting and beguiling complex behavior can emerge from collections of extremely simple components” [69]. Through their work, Langton and others have managed to show that one way to achieve lifelike behavior is to simulate populations of units instead of one big complex unit. Use local control instead of global control. Let the behavior emerge from the bottom up, instead of being specified and imposed from the top down. Moreover, while you're at it, focus on the ongoing behavior instead of the final result [67]. Final results are meaningless in lifelike scenarios because living systems provide pure examples of perpetual change, only settling down into a position of closure once all life has been extinguished and death ensues. And even then change continues as the components are assimilated back into their surroundings via natural processes such as decomposition.

By taking this bottom-up idea to its logical conclusion, it is possible to see it as a new and thoroughly scientific version of vitalism—the ancient idea that life involves some kind of energy, force, or spirit that transcends matter. The plain fact is that life does transcend mere matter, and according to Langton, not because living systems are animated by some vital essence operating outside the laws of physics and chemistry, but because a population of simple things following simple rules of interaction can behave in eternally surprising ways. Life in its classical sense may involve a kind of biochemical substance, but to make such a system sentient is not to bring life to a machine; rather it is to organize a population of machines in such a way that their interacting dynamics are “alive” [67].

To accept such a premise automatically leads to a final and striking conclusion about life—that is, that there is a distinct possibility that life isn't just like a computation, in the sense of being a property of its organization rather than its physical matter. Life literally is a computation [67], independent of any physical manifestation or linkages. To see why, start with conventional biological definitions of life. As biologists and others have been pointing out for more than a century, one of the

most striking characteristics of any living organism is the distinction between its genotype (the genetic blueprint encoded in its DNA) and its phenotype (the structure that is created from those instructions). In practice, of course, the actual operation of a living cell is incredibly complicated, with each gene serving as a design for a single type of protein molecule and with a myriad of proteins interacting in the body of a cell in a plethora of ways. But in effect you can think of the genotype as a collection of little computer programs all executing together, one program per gene. When activated, each of these programs enters into the complex logical maelstrom by competing and cooperating with all the other active programs in a finely tuned and self-organized overall balance. In unison, this entirety of interacting programs carries out an overall computation that is the phenotype: the structure that unfolds during the organism's development [67], thereby creating its own destiny and potentially solving a number of computational problems along the way. Indeed, Langton has managed to effectively capture a startling realism in his abstract computational framework. This is commonly referred to as artificial life, or A-Life for short, and is a model of computation that is closely analogous to higher-level computational systems such as the Web itself.

Next one can move from carbon-based biology to the more generalized biology of artificial life. The same notions apply, and to capture that fact Langton coined the term "generalized genotype," or GTYPE, to refer to any collection of low-level rules. He likewise coined the term "generalized phenotype," or PTYPE, to refer to the structure and resultant behaviour produced when those rules are activated in some specific context or environment. In a conventional computer program, for example, the GTYPE obviously represents the program code itself, and the PTYPE is the output produced by that running program as a result of the inputs by its users [67]. So, again by analogy, the Web can be viewed as being both a GTYPE and a PTYPE on the global technology scale.

Now, what is beautiful about all this is that once you have made the link between life and computation, you can bring an immense amount of theory to bear. For example, why is life quite literally full of surprises? Because, in the norm, it is impossible to start from a given set of GTYPE rules and predict what their PTYPE behavior will be—even in principle. This is the undecidability theorem, one of the most precious pearls of wisdom ever to have been discovered in the whole history of logic, and one that we will investigate in much more detail later. It states that unless a set of instructions—a computer program for want of a better term—is utterly trivial, the best way to find out what they will do is to actually execute them and see. There is no general-purpose procedure that can scan these instructions given a set of inputs and give a faster, or more accurate, answer than that. That's why traditionalists and those who do not understand the more esoteric capabilities of computing only see computers doing precisely what their programmers tell them. This is paradoxically both perfectly true and virtually irrelevant: Any piece of code that is complex enough to be interesting will always, but always, surprise its programmers, with its complexity literally stripping them of credible insight. That is why any decent software package has to be endlessly tested and reworked before it sees general release, and that's why the users always discover very quickly that the debugging involved was nearly always never perfect. Most important of all, for artificial life purposes,

that is why a living system can, for instance, be materialized as a biochemical machine that is completely under the control of a program, a GTYPE, and yet still have an unexpected, apparently spontaneous behavior in its PTYPE [67].

Conversely, there are other deep theorems in computer science stating that you can't go the other way either. Given the specification for a certain type of behavior, a PTYPE, there is no general procedure for finding a set of GTYPE rules that will produce it. That's why the overall behavior of an economy can very rarely be explained by the activities of a single company and likewise why the overall character of the Web is impossible to attribute to a single website. In practice, of course, these theorems don't stop human programmers from using well-tested algorithms and design methods to solve precisely specified problems in clearly defined environments. But in poorly defined, constantly changing environments, like the Web and those faced by living systems, there seems only one way to proceed: trial and error, also known as Darwinian natural selection. This process may seem terribly crude and wasteful, but in the long run it is actually highly effective. In essence, nature does its programming by building lots of different machines with a lot of randomly differing GTYPES and then smashing the ones that don't work well. In fact this messy wasteful process may be the best that nature can do [67]. But then unlike many modern-day programs of work, it has one distinct advantage on its side: There are no deadlines in the natural world, and evolution has all the time in the Universe to get its final code right!

We have now seen scientific definitions of life spanning an unexpectedly broad range of properties and criteria, and this is undoubtedly because we still do not fully understand the medium we are dealing with. That subtle combination of properties and criteria has included structural features, complexity, growth, reproduction, metabolism, motion, response to stimuli, evolvability, information content and transfer, and control of energy flow amongst others. Obviously, the word "life" has different connotations for different individuals within many different contexts, and this often leads to confusion and disagreement. For such reasons, Leslie Orgel¹² has coined the acronym CITROENS (complex information-transforming reproducing objects that evolve by natural selection) to describe certain borderline systems that appear to ascribe to life as a definition but for whom an absolute decision is not yet available. This is a veritable mouthful of an acronym without a doubt, but one that encompasses a great deal of up to the minute thinking nonetheless. It is also a fitting label for the Web as we move forward to examine more of its distinguishing features in depth in our search to establish its lifelike credibility.

¹²Leslie Eleazer Orgel is a chemist by profession. During the 1970s, Orgel suggested reconsidering the Panspermia hypothesis, according to which the earliest forms of life on Earth did not originate here, but arrived from outer space with meteorites. Together with Stanley Miller, Orgel also suggested that peptide nucleic acids, instead of ribonucleic acids, constituted the first pre-biotic systems capable of self-replication on early Earth.

