CHAPTER 1 A View of Life and Thought

Every school kid now learns that genes are made of DNA and that DNA codes for proteins. And everyone knows that human language comprehension and production typically involve coded sequences--strings of speech sounds, written sounds or syllables or words, sign language hand and face movements, or lines of Braille dot patterns. Therefore, the suggestion that cellular life and human thought might have something in common is not immediately surprising. Nor does it seem that profound.

I think there are actually very deep, exciting, and fruitful conceptual similarities between these two very different systems. More precisely, there are exactly two naturally-occurring examples (and as yet, no artificial example) of this particular kind of symbol-using representational system on the Earth. The first, single-celled version of such a system arose from a chemical prebiotic substrate at the origin of life, initiating Darwinian evolution. Subsequently, multi-cellular organisms evolved and they developed elaborate humoral and neural control mechanisms; but a similar, autonomous symbol-using system did not reemerge on any intermediate level until the origin of thought and language from the substrate of prelinguistic neural activity patterns in the brains of early hominids.

The similarities between these two systems are fundamental and specific enough that they can help us to come up with new kinds of theories to explain the still poorly understood features of the neural activity that underlie peculiarly human cognition. Also, an abstract analysis of how these systems work may be able to provide us with a guide for making artificial copies. Those are several entirely unphilosophical motivations for going forward with a project that at first glance might look merely like philosophy of mind or, in a positive sense, science fiction.

At this point, the reader may wonder how, or even if, such an fundamental likeness could possibly have been missed amid the prodigious yearly output of papers and computer programs by millions of biologists, cognitive scientists, engineers, linguists, and philosophers on the topics of molecular biology, evolution, neuroscience, and language. A number of factors are responsible, not the least being the fearful ramification of late twentieth century scientific knowledge just mentioned set against the backdrop of a reading speed probably unchanged since the origin of writing. An equally important reason for the lack of attention, though, is the non-obvious starting point for the analogy, which was found only after a basic difference between the two systems was recognized. The alignment of the two systems described in this book differs radically from three other mappings between biological and cultural/linguistic evolution that have occupied previous writers. A second main distraction has to do with complementary nature of the experiments that can currently be done on the two systems; but this is certain to change. A final diversion, worthy of detailed consideration, is the fundamentally different way in which code strings are used in present-day computers.

In this view of cellular life and human thought, the momentous events that came after the origin of each of these symbol-using representational systems were due to a new kind of intentional control that each began to exert on their worlds. In both cases, these systems made a new kind of evolution possible, though its details differ significantly in the two systems. Given the palpable evidence of this power, it will finally be worth pondering why systems of this kind have arisen so rarely.

A thick carpet of DNA

Life arose soon after the Earth was cool enough for water on its surface to be a liquid. Cellular life has had an enormous impact on the rocky, watery geology of the Earth, thickly coating its entire surface of land and water. Initially, there was no oxygen in the atmosphere (and there remain to this day organisms for which oxygen is a poison). Cellular life is responsible for virtually all of the oxygen in the planet's atmosphere--the result of a billion years of photosynthesis. The chemical factories in cells are also the sole source on Earth of strong molecular polymers like polynucleotides, polysaccharides, and polypeptides. One normally thinks of rocks as ’strong’; and they certainly are when kept dry. Dissolved in water, however, the crystalline minerals in rocks typically fall completely apart into small pieces consisting of only a few atoms each. By contrast, biological polymers--even millions of segments long--are created and can remain strongly bonded together in water, a requirement for flexible, multifarious interaction.

It is easy to forget the utter bizarreness of the picture that has been delivered to us by molecular biology. Virtually every square millimeter of the Earth (pavement included) is carpeted with a fantastically dense mat of biological information. Each individual eukaryotic cell contains hundreds of megabytes of DNA code. The tangled stream bank is much more odd than even Darwin imagined! There are terabytes of information jammed into every visible speck of leaf and dirt--all of it carefully copied, constantly scanned and error-corrected, and much of it accessed every day in the life of each of those uncountably numerous cells. The individual digital segments of information are already molecule-sized; it is hard to imagine how the information and interpreting apparatus could be made much smaller. It is as if life has turned the entire surface of the Earth into an unimaginably large, maximum density disk drive.

If we shrink ourselves down to molecular size and look at what cells are doing, it becomes clearer that cells invented a new kind of molecular-level intentionality--a way to partly overcome the deterministic thermodynamic buffetings to which all matter is subject--that went far beyond the chemical dynamics of the landscape before there was life. This does not imply that cells create mysterious, irreducible holistic forces; in fact, we know quite a lot about how they work. But it is a natural way of characterizing what goes on in cells that distinguishes them from the prebiotic chemical cycles in clouds, rock piles, streams, beaches, and ocean floors. The prebiotic state is already a complex, highly interactive, energy-dissipating, deterministically evolving soup containing a number of different types of dynamically stable subunits and structures. Cells, however, found a way to encode, use, and reproduce information about how to cause thousands of different chemical reactions in this soup to happen. The tricky part is that the information, as well as all of the interpreting apparatus has to be in the soup; everything is still subject to the soup's deterministic buffetings. Some of the chemical reactions can already happen a little by themselves. The cellular system, however, speeds many reactions, slows or prevents others, invents many new ones that never used to happen at all before, and above all, orders and organizes the reactions. In short, code-using cells have taken over forceful control of chemical phenomena in local regions of the otherwise still prebiotic soup.

To get the system off the ground, a way had to be found to camouflage the coded information from the dissipative attack of the soup, but without hiding the information so thoroughly as to make it inaccessible. In a sense, the resulting system is still locally deterministic since no new chemical forces or low level rules of chemical interaction have been added. But at a slightly larger scale, there is another clear sense in which the system escapes determinism. By exploiting partially hidden, partially arbitrary information that the soup has trouble seeing and thus destroying, the system is able to evolve, by reproducing and adapting that information in a new, Darwinian manner, very far away from its initial prebiotic state into configurations that are exceedingly improbable from the soup's viewpoint. In this sense, cells are intentional--microscopic, but enormously willful and deliberate.

Origin of the human mind

Humans, too, have achieved a new level of intentional ability when compared to other animals. And it has resulted in a drastic reworking of the Earth similar to that brought about by the origin of life. It took a while for humans to get into high gear. For most of the evolutionary history of modern humans, we lived in small bands of gatherers and hunters, having more benign effects on our surroundings. It is difficult to avoid the conclusion, however, based mostly on artifacts like tools, burials, and artwork, that even the less numerous, somewhat more environmentally-friendly early humans had a new kind of control over what went on in their heads. In recent years, our collective effect on the geological and biological environment has begun to rival that of rare events in the Earth's history like asteroid impacts or giant (100,000 sq km) flood basalts.

Our increased intentional power compared to animals is correlated with language use. As with the case of cells and DNA compared to geology, it can be easy to forget how odd our behavior is with respect to other animals--in this case due to immersion in our own language rather than the unfamiliar microscopic scale of the cellular code. An animated, hour-long conversation consisting of a sequence of perhaps 30,000 closely connected speech symbol segments must appear pretty weird to a non-linguistic animal. A walk about a busy market in a country where you don't understand the language brings back some objectivity. The only examples of serial vocal behavior in present-day animals that even remotely resembles this in scale are from songbirds and whales; wrens, for instance, typically sing hundreds of distinct songs, each consisting of a group of syllables with a few sound segments per syllable. It is easy to tell the human and avian behaviors apart, though, since in humans, the ordering of the sound sequence at intermediate scales (as reflected in word and sentence order, spanning tens to thousands of segments) is essential in coding for a meaning, while the songbird's intermediate range ordering (as reflected in within-song or between-song order) does not appear to code for or mean anything. And birds don't talk directly to each other the same way we do.

As was the case with cells in relation to the prebiotic soup, prelinguistic animal cognition is already quite complex; moreover, it is supported by brains with a design and operating parameters very similar to our own brains. However, there is a clear sense in which animal cognition and brain activity is much more beholden to the deterministic buffetings of their physical and social environments. It is not that humans have outgrown feeding, fighting, fleeing, mating, and sleeping. But we have transformed these exercises almost beyond recognition (by an animal).

I think the main point of language is to allow humans to encode, use, and reproduce information about how to make thousands of *mental* reactions take place--to support a new kind of more intentional mental metabolism of neural activity patterns. As in the case of cells, the code and the interpreting apparatus cannot have climbed out of the soup of prelinguistic neural patterns, but must have somehow devised tricks for concealing arbitrary, coded information from the dissipative attack of the prelinguistic neural pattern 'soup' in order to get the language-level system off the ground.

There is a tendency to focus on the communicative functions of language. There is no question that language is useful for that purpose. But it is also possible to think of language as essentially a new brain operating system that turned out to also be quite useful for communication. It is much harder to see, manipulate, and understand internal neural activity patterns flitting across billions of neurons than it is to make theories about external code strings. Nevertheless, the specific aspect of the neural activity patterns in human brains that have granted us such a quantum increase in intentional power when compared to animals remains the real grail for the next century.

Code-using systems vs. Darwinian evolution

Many previous attempts to extend evolutionary ideas to language and culture have rather unadventurously hewn closely to a subset of the features of the biological/cell-based version. The results over the last century have generally been disappointing. I would like to focus instead on what I think are more fundamental similarities between the two systems in the architecture of code use.

The requirements for biological evolution are often summarized in the motto--’heritable variations in fitness’ (Lewontin, 1970). It has been suggested that this is a very general principle that could be applied to many things--namely, any case where there are populations of individuals that generate offspring that have features of both parents (but not the exact average), and where their inherited form is tested in a way that affects their rate of reproduction. The problem with this motto is that no one ever found a convincing naturally-occurring example of heritable variations in fitness that didn’t depend directly on DNA, proteins, and cells (or language and human brains).

This standard characterization of an evolutionary system picks out only part of the essence of what is necessary for the peculiarly biological (and cultural/linguistic) kind of evolution to occur--as distinguished from the physicist’s kind of "evolution of a dynamical system through time" applicable to cells but also to galaxies, streams, and beaches. The standard motto for biological evolution is rather like defining a bachelor as a someone who doesn’t wear a ring; a better definition would also mention marriage. The discussion of life and thought above focused first on how code is used to make devices for controlling reactions and establishing a new kind of control over the local environment. My argument is that this code-using core is required for the new kind of evolution to occur in a natural setting--that is, without an external intelligent agent for design and maintenance.

There is an interesting architectural difference between the two code-using systems that results in quite different boundary conditions on the kind of evolution that then occurs. Thus, I am suggesting that the real natural kind is an intentional code-using system, not a Darwinian evolutionary system; Darwinian evolution is one case of the historical accretion of adapted morphology with precise reproductive copying, no code production, and discrete generations. We will see that the dynamics of the historical accretion of adapted morphology at the level of the code-using system in human brains is quite different from Darwinian evolution.

Comparable discontinuities

Most anyone will recognize that the origin of life and the origin of human thought both constitute discontinuities of some kind. There are striking divergences of intuition, however, on how these two transitions compare in overall magnitude or importance. I say ’intuition’ since the topic is rarely explicitly broached. Rather, it involves background assumptions that usually congeal early in someone’s career, soon after they have broadly decided what field to settle into.

Many with a more humanistic bent would be taken aback by the suggestion that the origin of the mere chemicals of life could be placed alongside the origin of human language syntax, thought, religion, literature, art, music, architecture, fashion, cuisine, science, politics, and so on. Others with a more biological or developmental bent, see mere life as complex and subtle; they would emphasize the overwhelming continuities between the brains and cognitive abilities of animals (especially apes) and humans, and tend to regard language as important but much less fundamental, and a less discontinuous innovation than the origin of the universal cellular system that made the evolution of the language ability possible in the first place. The present project is conservative at least in the following sense--it stands firmly at the midpoint between these two extremes.

There are actually two discontinuities in both the evolution of life and thought. In each case, there was a long latent period after the initial origin of each code-using system. Most of the history of life, for example, is the history of single-celled life. Tightly coupled multi-cellular body plans with functionally specialized cells like neurons arose only much later and led to the explosive diversification of life near the beginning of the Paleozoic, the origin of the major phyla of modern organisms, and the rapid discovery, establishment, and occupation of the major ecological niches. This largely occurred without substantial alteration to the intracellular architecture of code-use. Similarly, though there were significant advances associated with the origin of modern humans and language in the Upper Paleolithic--including better stone tools, burials, painting, and figurines--the wide spectrum of innovations that followed the much more recent invention of agriculture and animal domestication--including writing, cities, and much more--constitute as abrupt an efflorescence as the one that followed the late pre-Cambrian origin of multicellularity. As with the Cambrian explosion, the internal architecture of symbol-use in food-producing humans seems not to have changed as much as the content and context of the symbol strings.

How computers fit in

Like cells and people, computers obviously also use code. However, computers are not naturally-occurring since they are dependent on humans for their programming, design, maintenance, and reproduction. And they use code in a fundamentally different way than cells and people do.

Computers are certainly not the first bit of human-designed machinery to be used as an analog for thought. Descartes, for example, envisioned the nervous system as a complex network of air-filled tubes (it has been suggested he was inspired by viewing an elaborate fountain). In his model, sensory stimuli are transduced into air puffs and carried to the brain by tube-like nerves. There the pineal acts like a pneumatic telephone exchange (controlled by the mind) that converts sensory air puff patterns to motor air puff patterns. These are then transferred by nerves back to the muscles, causing the muscles to inflate, shorten, and move limbs. This was an admirably explicit dynamical systems theory of the brain (except for the non-physical mind).

Because they use code, however, computers have often been proposed as a better model of the mind than other kinds of dynamical systems. There are two features of code use that computers share with cells. One is that code strings are read and written with digital effects. This means that a kind of categorization is performed on each symbol token--for example, 2.29 and 3.31 volts are both read to be the same as 3.3 volts. Something similar goes on in a cell accessing cellular symbols in DNA and RNA. For example, a first position ’G’ (guanine) in a messenger RNA will reliably be paired with a first position ’C’ in a transfer RNA despite the fact that thermal vibrations of the atoms in the ’G’ and the vagaries of the exact positioning of the messenger RNA strand in the ribosome will actually present a slightly different physical situation to the ’C’ of the transfer RNA during each individual recognition event.

Another key physical ingredient for a computer is the provision for long one-dimensional chains of individual symbols that do not directly interact with each other. The code recognition device (CPU) can count on finding a symbol where it was last put; it doesn’t have to worry about between-symbol interactions occurring while it’s ’not looking’. The code chain must not routinely ’fold up’ and interact with itself. Symbols only interact with each other indirectly via the code recognition device. Again, very similar considerations apply to DNA chains in cells.

However, cells put these recognized symbols to work in a remarkably different manner than computers do. Computers use code exclusively for *operating on other code*. This is powerful idea. By breaking any direct connection with the non-code world (a computer symbol has a completely arbitrary relation to something in the world) we paradoxically make it *easier* to program a digital computer to compute useful things about the world.

Cells, by contrast, use all their code to make proteins; starting from a protein, they can’t even write their code at all! The assembly of proteins relies fundamentally on non-arbitrary constraints in the (chemical) world. Proteins are constructed by simply bonding amino acids together into a 1-D chain that is parallel to the recognized symbols in the messenger RNA chain. But then this chain folds up into a precise 3-D structure with a precisely shaped surface according to an elaborate set of chemical constraints. These include strong and weak chemical bonds, interactions with water, and the precise structural details of the amino acids--a large set of hard-wired, non-arbitrary constraints that the cell harnesses, but cannot change.This would be a terrible way to make a general purpose computer; the effects of a change in one symbol are propagated throughout a cell in a tangled manner that is difficult to predict and ’debug’. But then the point of cells is not to compute useful things for humans, but to control chemical reactions using precisely shaped 3-D surfaces. The argument of this book is that language competence in humans brains is not a form of computation either, but rather the second example of cell-like code use--this time with code production.

A salutary blow

When a newcomer is introduced to a scientific field, there is often a reaction something like why (on earth!) are you studying that? And why in that way? Why not go after this obvious thing instead? More often then not, the answer the student gets is, yes, of course, we eventually want to study ’X’ but for various practical reasons (it’s too hard now, we don’t have the proper techniques, funding is not available yet), we can currently only study ’Y’ and ’Z’. In time, as a student learns more about the field, these practical constraints are internalized as an unconscious platform for everyday theory and experiment; they are crucial to getting useful work done. However, a brief look at the history of science suggests that sometimes, the scope of our understanding changes more abruptly and we actually do get to begin studying the formerly off-limits ’X’.

In molecular biology, there was a remarkable increase in understanding relative to classical genetics with the discovery of the structure of DNA and means by which it codes for proteins. Though we cannot yet make a living cell truly from scratch because it is much too hard to put all the necessary chemicals in the right place at the right time, it is fair to say that the main principles of life are now becoming well understood. Our ability to manipulate cells at a molecular level has been vastly increased as a result of this knowledge.

We currently also know an enormous amount about how the brains of animals and humans work. In contrast with cells, however, we still have only an extremely vague lower-level appreciation of what it is about human brain activity patterns that make them so powerful. There is a telling contrast between late nineteenth century chemical biology and late nineteenth century neuropsychology. The first field is now mainly of interest to historians of science and its texts would barely be recognizable by a modern molecular biologist; documents from the second field sound all too familiar to contemporary ears, even ones without historical training. The most detailed theories about human-specific abilities at present are phrased only in abstract, high level terms quite disconnected from activity patterns in neural circuits, often accompanied only by an embarrassed or arrogant proclamation that these theories must somehow be implemented in neurons. There is nothing wrong with higher level theories. But if the history of classical genetics and molecular biology is a guide, we will need to make more explicit connections between higher level theory and lower level implementation to gain a similarly deep understanding. None of this is to denigrate systems neuroscientists. Molecular biologists use the same local, reductionistic methods; and they aren’t any smarter people. Those methods just work a whole lot better once the basic operating principles of the target system are known.

This is where a large-scale analogy can help--especially one aligning, with a certain unavoidable brutality, systems very far apart in scale. It provides a healthy jolt to dislodge the mind from the well-worn tracks in which it runs in our everyday scientific lives. There are certainly a number of differences as well as parallels between the code-using systems in cells and humans. Some of these differences are easy to recognize; others will be difficult to see, and may generate misleading predictions. It is a chance that we will take to get ourselves into the right ball park.

One difficulty with a project of the present kind is that these two systems both have many parts, but also many levels of organization. It is difficult to keep them in mind all at once, it is easy to get lost, and many people have. Therefore, we shall have to spend a substantial amount of time setting the stage and closing off less productive--but very interesting--alleyways. Once we have our bearings, we will begin by slavishly going back and forth between the two systems. After some of the main correspondences have been sketched out, we will be able to adopt a less monotonous cadence. [~4,100]