

Biological Information and Natural Selection

Abstract. What is ‘biological information’? And how is it related to ‘complexity’? Is the ‘complexity’ of a cell the same thing as the ‘complexity’ of a hurricane, or are there two kinds of complexity in nature, only one of which involves the transmission of ‘information’? An account of biological information as a set of simplifying conventions used to coordinate the division of labor within cells is offered, and the question of why natural selection might tend to produce such systems of conventions is considered. The somewhat paradoxical role of mutations as occasionally-informative noise is considered, and a very simple formal model of the relationship between natural selection, complexity and information is developed and discussed in relation to adaptive landscapes. The question of whether we should expect evolution to slow down or speed up as the things that are evolving get more complicated is raised and dealt with. Finally, an epistemological moral is drawn, and briefly applied to the case of cancer.

In some ways, a cell is like a large, noisy, neural net.[1], [2] Molecules, or segments of molecules – genes, transcription factors, and the rest of the cell’s regulatory apparatus – act as nodes in the network. Interactions between molecules can be thought of as connections between the nodes. The propensity of a pair of molecules to interact can be identified with the weight of a connection. The pattern of response of a gene to the binding of a transcription factor seems analogous to the response function of a node. Either thing could in principle be modeled, somewhat imperfectly, with a huge, interconnected system of stochastic differential equations.

This mathematician’s way of looking at cells – a direct intellectual descendant of Stuart Kauffman’s original ‘random Boolean network’ model of gene regulation[3] – makes them seem rather like complex analog computers. And yet many biologists – including Kauffman himself – are skeptical of this further step. Are cells really just ‘processing information’? Are computers really a good model for cells? Isn’t life something rather different from computation? When we make this sort of idealized mathematical model, aren’t we greatly oversimplifying what is, in reality, a very complex physical system? Is the information, as opposed to the molecules we mentally associate it with, really there at all?

On the other hand, though, if we *can’t* legitimately talk about biological information, what sort of information is there left for us to talk about? We ourselves are living things, and all our information-processing devices are arguably just parts of our own extended phenotype. If we decide that it’s meaningless or superfluous to speak of genomes as ‘encoding information’, or of gene-regulatory networks as ‘processing information’, then shouldn’t the same arguments apply to neurons, and systems of neurons, and therefore to human thoughts, speech, and symbolic communication? Or is this taking skepticism too far, does this kind of biochemical nominalism throw the baby out with the bath water?

Some of these questions sound like philosophical ones. Perhaps that means that it's a waste of time to even think about them, but there really might be some utility in an attempt to at least state the obvious, to say what everyone already more or less knows, by now, about the nature of biological information. If this turns out to be easy to do, then we'll know that our ideas on the subject are in perfectly good order, and we can go back to doing science untroubled by any murky philosophical concerns. If it turns out to be difficult, we might end up having to do some thinking about philosophical problems, whether we want to or not. This seems like an elementary precaution, one that could hardly do us any harm.

1 A Criterion

Where do we start? If we want to know when it's legitimate to speak of encoded information in describing an event in nature, presumably we need to provide an account of what distinguishes cases in which it is legitimate to speak in that way from cases in which it is not. So one thing we are probably looking for is a criterion or rule that will allow us to decide which category a particular case belongs to.

A chromosome and a hurricane are both very complex things. But we tend to want to say that a DNA molecule encodes information in its sequence. On the other hand, we don't exactly ever want to say that a hurricane encodes much of anything. It's just a storm, it's random. Is there a principled distinction here, or merely an anthropomorphic habit of thought with respect to one sort of object but not the other?

There actually does seem to be a real difference behind this distinction. The information-encoding sequence of a DNA molecule, as opposed to the exact physical state of its constituent atoms at a particular instant, has a physically meaningful existence to the extent that the other machinery of the cell, RNA polymerases and transcription factors and things like that, would behave in the same way if presented with another molecule having that same sequence but whose individual atoms were in a slightly different exact physical state.

We can think of the sequence space associated with a biological macromolecule as a partition of the space of its exact physical states. Each sequence is associated with a very large number of slightly different exact physical states, all of which would be interpreted by the cell's machinery as representing the same sequence. It's meaningful to speak of this partition as something that really exists in nature because natural selection has optimized cells to act as if it does. Of course, if we had an omniscient, god's-eye view of the cell, we could still ignore the existence of genomic sequences as such and do a brute-force, atom-by-atom predictive calculation of its trajectory through state space, just as we could with the hurricane, but that's not what the cell itself is set up to do, that's not how it regulates itself. It uses much coarser categories.

Within certain limits, DNA molecules having the same sequence will be transcribed into RNA molecules having the same sequence, regardless of the exact physical state the parent DNA molecule is in. Of course, if it is folded in the wrong way, or methylated, or otherwise distorted or modified beyond those limits, a

molecule may evoke a different response, but as long as it is within them, the same sequence will get the same transcription. Thus, it is the sequence of the parent molecule, and only the sequence of the parent molecule (not its various other physical features) that constrains the gross physical state of the daughter molecule. In that sense, it really is physically meaningful all by itself, because the rest of the machinery of the cell is optimized to treat genes with the same sequence as if they are the same thing. There just isn't anything like that going on in a hurricane, even though a cell and a hurricane are both 'complex systems'.

2 Two Types of Complexity

This certainly seems like an important difference. We can't really have much hope of answering our opening questions about how we should think of cells until we can explain it. Perhaps both things are complex systems, but, on an informal level, we habitually treat cells as being complex in a rather different way from hurricanes. For one thing, we never speak of optimization when talking about something like a hurricane. Optimization only comes in to a discussion when it somehow involves natural selection, not something that hurricanes are usually thought of as being subject to. Are there, then, two entirely different kinds of 'complexity' in nature, one naturally selected kind that involves 'information' and optimization, and another, non-self-replicating kind that involves neither? Why are they usually considered one phenomenon? Do we designate two entirely different things with one English word, 'complexity', or are the differences between the two types inessential?

Well, what could we actually mean by the word 'complexity', in the case of the hurricane? The general idea seems to be that because it's so big and contains so many different molecules moving in so many different directions at so many different speeds, it has quite a few degrees of freedom, that there are rather a lot of different physical states the system could be in. Turbulent flows can vastly magnify the effects of events that initially involve only a few molecules; so any and all of these degrees of freedom could affect the gross behavior of the system over time with indifferent probability. Predictions of the detailed behavior of the system far in the future will be inexact unless they are all taken into consideration. The absolute complexity of a hurricane might perhaps be equated with the volume of its statistical-mechanical phase space, a space with six dimensions for each molecule in the storm (three for position, and three for momentum along each of the three spatial axes) in which any instantaneous state of the whole system can be represented as a single point.

If that's what we mean by the word, though, we seem to have a bit of a problem. By that measure, any reasonably large collection of atoms is extremely complex, and cells are no more complex than drops of water of similar size. Both things are made of atoms, in approximately equal numbers, so the phase space for each should have a similar number of dimensions. A similar number of dimensions implies a roughly similar volume. Actually, though, we usually seem to want to treat cells as containing another layer of complexity beyond this obvious physical one, as being complex in a way that mere drops of water are not, as somehow having additional features. These

additional features are precisely the kinds of things that seem to differentiate ‘complex’ entities like cells from ‘complex’ entities like hurricanes, so apparently it would be easier to keep our philosophical story straight if we were in a position to maintain that they actually exist in some physical way.

3 Simple and Complex Simplifications

The idea that these additional informational features must have some sort of genuine physical existence isn’t exactly wrong, but we’re going to have to adjust it a bit to make it coherent. From a physical point of view, all these ‘additional’ details cells supposedly have, though they’re very complex as far as we’re concerned, are actually the details of a set of ways in which they are *simpler* than oily drops of water. It’s just that they’re simpler in a very complicated way, unlike ice crystals, which are simpler than drops of liquid water in a very simple and uniform way. In both cases the actual physical entropy of the ordered object is lower than it would be if its constituent small molecules were all scrambled up and broken apart into a randomized liquid. Biological order, like classical crystalline order, is a simplification and regularization of a complex, irregular molecular chaos, one achieved, in the biological case, at the expense of creating even more disorder elsewhere. Repetition – for example, the same kinds of amino acids and nucleotides repeatedly appearing everywhere in the cell – and physical confinement (to compartments, or within the structures of macromolecules) are both conditions that, all other things being equal, decrease entropy. What increases, during evolution, is the complexity of the physical simplifications life imposes on itself and the world around it, not absolute physical complexity *per se*.

This rather obvious fact should help us make some sense of the idea, mentioned above, that biological information exists just when physically slightly different molecules with the same sequence are treated by the cell as being the same thing. To the extent that the machinery of a cell indifferently treats two different RNA molecules with the same sequence as if they were identical, it interacts with itself in a way that involves simplifying reality, *throwing away* some of the available details, responding in almost exactly the same way to many slightly different states of affairs. This is a necessary part of the process by which life continually re-imposes its own relative simplicity on the complex world around it. It’s an energetically expensive imposition of sameness, an endothermic filtering-out of perturbations and suppression of irregularities. Some theoretical biologists call this general sort of process ‘canalization’,^[4] though that term normally has a much narrower meaning.

Life responds to the members of the collection of possible states of affairs in the world in a way that is much simpler than that collection. In this sense, ‘information’, in a cell, isn’t something that is *added* to the bare physical particulars of its constituent molecules, it’s those bare particulars themselves with something expensively and repeatedly *subtracted* from them, a sort of topological invariant that can be conserved through all distortions and translations precisely because it is not very dependent on exact details. Biological information is filtered complexity,

complexity that becomes information when some of its features are systematically and routinely discarded or suppressed as noise, when it is consistently responded to, by the cell's own machinery, in a way that only depends on some selected subset of its features.

(We might wonder, then, whether the idea of *transmission over a noisy channel*, so central to mathematical information theory, should really also be central to our theory of 'biological information', when *filtration* might seem like a better description of what's going on. Here, however, we may actually have succeeded in making a distinction without a real difference behind it, because the filtration of information is just the systematic ignoring or removal of noise, and that's what's required in transmission over a noisy channel as well. In the first few steps of visual perception, the conversion of light impinging on the eye to action potentials in neurons, the chaos of the external world is filtered down and transformed into a stereotyped, coded signal – but this is also what a ribosome gets, via an mRNA, from the complex chaos of an actual, physical chromosome, a stereotyped, coded signal. There is no deep conceptual difference between the transmission of information over a noisy channel and its filtration.)

4 Codes as Conventions

The privileged subset of a complex molecule's features that matters typically consists of ones which have conventional mechanical relationships to other collections of features, elsewhere in the cell's machinery. A tRNA encodes a convention that relates a selected subset of the features of any mRNA to a selected subset of the features of some amino acid. In this way a codon 'represents' a particular type of amino acid. Only that type of amino acid will be allowed to randomly bounce into a situation where it will form a peptide bond with a growing protein chain, if that codon is currently the one being translated by the ribosome that is making the protein. The tRNA seems to encode a 'meaning'. Should this bother us, philosophically, in any way?

Yes and no. It shouldn't actually surprise us, but it should make us cautious. An analogy with a meaning in a human language does exist, but it's not a precise one. In this sort of discussion, we can't just pretend it isn't there, but on the other hand it can easily become misleading if we don't pay careful attention to the differences. That means that we need to dwell on it here, to make sure we understand its subtle pitfalls.

What, in general, is a 'meaning'? The philosopher David Lewis argued persuasively that we can often think of meanings, even in human languages, as conventions for solving what the game theorist Thomas Schelling dubbed 'coordination problems'. [5], [6] Lewis argued that they can be thought of as public agreements committing us all, for the sake of convenience, to treat x's as valid proxies for y's under some circumstances, to treat the word 'cow' as a valid proxy for a cow, allowing us to collectively and individually reap the gains available from better coordination of our joint activities. Better-informed dairy farmers, who can speak to each other concerning cows and their affairs, mean that there is more butter

for everyone than there would be if we couldn't discuss such things, and it is the traditional semantic convention that makes the noise a proxy for the animal that puts some of that butter on the table. tRNA molecules seem to function as 'traditional' semantic conventions in more or less this sense, since they facilitate the widespread use of particular codons as proxies for particular amino acids.

5 Selection, Coordination and Conventions

Why should natural selection, as a mechanism, tend to produce this sort of complex and rather artificial-seeming convention? What's already been said about Lewis's theory of meanings as conventions suggests a very general answer to the question. Somehow, to persist, grow, and reproduce, life must need to solve a lot of coordination problems, in a way not equally necessary to crystals and storms. But why is life, in particular, so beset with this kind of problem? Why don't crystals or hurricanes face a similar need for coordination across time and space, how is it that they can they grow and persist without transmitting coded information from point to point and time to time if cells can't?

That's a big question, but to make any progress in clarifying the analogies and dis-analogies between cells and computers, we must have an answer to it. To get one, we're going to have to back up a bit. Though it may seem like a digression, before we can really think clearly about how natural selection produces biological information and the conventional codes that carry it, we must first think a bit about the character of natural selection itself, as a process, and what sort of things it actually acts on.

Continual selective change in gene frequencies has to go on in a population over long periods of time for very much evolution by natural selection to occur. This can only really happen in a population that has both heritable variation and Malthusian dynamics. We tend to focus on heritable variation as the really essential thing about life. This way of thinking won't help us with our present inquiry, however, because the way inheritance works in actual modern organisms is already completely dependent on the existence of an elaborate form of encoded information. If we want to understand why living things create and process something recognizable to us as encoded information in the first place, we have to think carefully, instead, about the other side of Darwin's great idea, about the Malthusian dynamics of biological populations, and where those dynamics come from.

John Maynard Smith and Eörs Szathmáry have pointed out[7] that one feature of the molecular tool-kit on which life as we know it is based that seems to be essential is the fact that the basic building blocks – amino and nucleic acids – can be put together into an infinite variety of different shapes. This makes them a bit like a Lego kit, a set of standardized blocks with conventional shapes and affinities which allow them, collectively, to be assembled compactly into almost any form at all. On a molecular scale, 'almost any form at all' means about the same thing as 'catalyzing any reaction at all'. Because macromolecules made out of the same basic constituents can catalyze an enormous variety of different reactions, a self-replacing collection of

such macro-molecules can divide the labor of self-replacement up into an arbitrarily large number of distinct reactive tasks, each of which will be performed far more efficiently, by the specialized catalyst that has evolved to take care of that aspect of the overall task, than any competing abiotic reaction.

In contrast, there's really only one kind of hurricane. The raw ingredients for a hurricane can only go together in one general way. The spiral can be larger or smaller, and can differ in various other insignificant ways, but the possibilities are very limited. A similar thing could be said of classical crystals – at a given temperature and pressure, there tends to be only one form the crystalline lattice can take. Complex organic molecules are a much more promising substrate for evolution from this point of view, since a few types of small molecules can be assembled into an infinite variety of large ones.

Geometrically, this difference between classical crystals and large organic molecules might have something to do with the fact that carbon atoms are effectively tetrahedral, and therefore don't pack into three-dimensional space in any compact, periodic way unless the bond angles are distorted as they are in diamond. This fundamental geometric frustration, a three-dimensional analog of the five-fold symmetries that frustrate periodic tilings with Penrose tiles, is an intuitively appealing candidate for the role of culprit in life's complexity. Be that as it may, for our present purposes, it's enough to just register that the difference exists.

This whole process of endlessly dividing up the self-replicative task faced by cells sounds very similar to something that Darwin talked about in the chapter on 'divergence of character' in *Origin of Species*.^[8] There he describes the evolution of both ecosystems and systems of organs as an endless process of efficiency increasing through an ever-finer division of labor, which allows narrower and narrower tasks to be performed by more and more specialized structures. His arguments work just as well on the bio-molecular scale, inside of cells, as they did on the scale of whole ecosystems. The virtue of a flexible, Lego-like kit of tools which can evolve and conserve designs of arbitrary complexity is that it allows this division of self-replicative labor to go to its logical extreme.

The open-ended character of this project, its ability to accommodate contingencies with contingency plans and make lemons into lemonade, has made life a robust and common kind of self-organizing system, able to grow and persist in a wide variety of environments and survive a wide range of perturbations. The other kind of 'complex' (i.e., only slightly simplified) dissipative system in nature, hurricanes and crystals and things like that, depends on exactly the right conditions to grow, and will immediately begin to shrink if the environment deviates from those perfect conditions, but life can evolve to fit itself to almost any conditions it encounters. The difference is basically that between a toolkit that can only do one thing and a toolkit that can do anything, that can do many things, that can do several different interacting things at once, or several things one after the other.

Darwin's idea of an evolved, flexible division of labor is useful in a modern context because it lets us more precisely characterize the difference between the sort of self-replication a living thing engages in and the kind of 'self-replication' we see in crystallization, where the cellular structure of the crystalline lattice can create copies

of itself very rapidly under just the right conditions. The molecules that make up biological systems replicate themselves by processes that are much more *circuitous* than those involved in the growth of crystals.

Circuitousness and a fine-grained division of labor are actually more or less the same thing. A division of synthetic labor means things are produced in many small steps, which is to say circuitously. Because catalysts can very efficiently accelerate many very specific reactions, this form of self-replication both is more powerful and more flexible than crystallization, better at grabbing atoms away from other competing geochemical processes under a wide variety of circumstances. It is this superior efficiency that makes it possible for living things to quickly fill an empty environment to the point of saturation. Thus, it is the flexibility of the underlying toolkit that ultimately explains the Malthusian dynamics of biological populations.

This circuitousness creates scope for alternate ways of doing things, and the possibility of selection between accidental variations. The more steps there are to a process, the more different ways there are to change a few of those steps, and the more different ways there are to change it, the more likely it becomes that one of those changes will improve its efficiency. Circuitousness also creates a need for coordination across time and space. Things must be done over here in time to get them ready for their role in a process going on over there, and if that process stops for any reason, they may have to stop as well. It is to these sorts of facts that the existence of anything recognizable as biological information must ultimately be attributed. The continual flow of information carried to the genome of a cell by transcription factors causes mRNA's to be produced when and where they are needed, and to stop being produced when they are no longer necessary. Circuitous self-replicators that have coordinated their activities most efficiently in time and space have had the most descendants, and in this way a whole genomic apparatus for coordinating and conserving circuitous patterns of molecular self-replication has evolved.

6 Mutations as Informative Noise

The division of the labor required for efficient homeostasis and self-replication requires coordination in the face of perturbation, and this requires the filtering out of noise. If cells were little robots made by tiny gnomes, our analysis could stop here. But life as we know it is the product of evolution. It would be a mistake to end our discussion of the relationship between noise and information in living things without considering the rather central case of mutations, a form of 'noise' which also seems to be the ultimate source of the genome's 'signal'. How can these two faces of the phenomenon be reconciled?

The existence of all the elaborate genome-repair adaptations found in modern cells shows that from one point of view – that of individual cells – mutations really are just noise, to be suppressed and filtered out if possible, like any other random noise. Of course, the cell does not actually have a 'point of view'. We should have said 'what typically maximizes the individual cell's fitness is the filtering out of as many

mutations as possible.’ Still, the fact that it has a fitness to maximize in the first place puts a cell much *closer* to having what we might normally think of as a ‘point of view’ than the other ‘player’ in the mutation/biological information game, Nature, in her role of selector of the fit and unfit.

We’re sometimes tempted to personify this actor, too, at least on an unconscious level. Even Darwin used to talk this way occasionally, mostly for the deliberate Epicurean purpose of shocking us into seeing how different ‘Nature the selector’ really is from a mere person. Before he settled on the term ‘natural selection’, though, he considered the possibility of speaking, instead, in terms of ‘the war of Nature’. (‘Survival of the fittest’ would then presumably have become something like ‘continual repopulation by the victors’, which at least has the virtue of descriptive accuracy.)

‘Selection’ may sound vaguely like a kind of information processing, but ‘war’ really does not. In fact, one of the most interesting things about natural selection is just exactly how *unlike* any normal form of ‘information processing’ it is. Still, somehow, out of this inchoate struggle a coded message eventually emerges. Our analysis of ‘biological information’ will be incomplete unless it includes some non-anthropomorphic explanation of how this magic trick works.

7 Serial Syntax Meets Holistic Semantics

Much of the way we habitually think about information and information processing comes from the fact that we ourselves are living things, and as such, are not infinitely capable. We have to break computational tasks down into pieces, and deal with one piece at a time. Bandwidth is a real constraint in the informational dealings of mere creatures, and everything they do is arranged to get around that fact. It takes us time to read a book, just as it takes an RNA polymerase time to produce an error-free transcript of a gene. There are many small steps involved. Whenever we want to do anything complicated, we have to do it in small chunks. Messages must be arranged in a sequential way, one simple piece after another, so that that they can be transmitted and decoded a little bit at a time.

When we think about computational problems, we of course assume that it will take a certain amount of time or space to solve a problem of a given complexity, and we are very interested in classifying such problems with respect to this kind of difficulty. But nature, as selector, does not have to respect these classifications, because it is not an information-processing device solving problems of limited complexity sequentially in stages. The ‘war of Nature’ happens in the real world, so the whole idea of bandwidth is irrelevant. A predator may kill its prey in an arbitrarily complicated way. This doesn’t make it any harder, or any more time consuming, for the prey to die. The prey doesn’t have to use any computational resources, or do any work, to be affected by the predator’s arbitrarily complex strategy, and the event doesn’t have to take any particular amount of time. Indirect ecological interactions – acorns from masting oaks feeding mice who bring ticks with Lyme disease to deer – of arbitrary complexity can also affect fitness in serious ways.

There is no increase in computational cost associated with passing through more complex versions of this sort of filter.

Using terms like 'select' and 'evaluate' conveys a somewhat misleading impression of the way natural selection works. Natural selection is not a cognitive or computational process at all. We face a temptation to imagine it as having some of the limitations of such a process, but it does not. Evolution has not made cells what they are today by breaking them down into pieces and evaluating or dealing with the pieces separately, one a time. Recombination does break repeatedly genomes down into pieces of variable size, but 'evaluation' – selection itself – mostly happens to entire cells and organisms, and to genes in the context of whole cells and organisms, embedded in whole, particular possible worlds. It takes a whole organism, in all its complexity, in a whole complex world, to live, or die, or raise a litter of kittens. (Will the rains fail this year?) Genghis Khan has a large number of descendants today because of who he was in all his individual complexity, and what 13th century Central Asia was like. His Y chromosome might not have done as well in modern Chicago.

Natural selection is not the kind of mechanism that evaluates things a little bit at a time, it is the sort of mechanism that evaluates them all at once. The information in cells is *processed* serially, but it is first *produced* by natural selection as a gestalt, a single tangled-up whole. (It normally takes many selection events to bring a gene to fixation in a population, but each of those events involves a whole individual's complex struggle to get through a whole life in a whole world with a whole genome, so this does not mean that the process has somehow been 'broken down into smaller pieces' in the relevant sense.)

Whether or not a mutation is new biological information or just more noise isn't even something that can generally be discovered just by inspecting the organism in which it occurs. What is adaptive depends on the environment that an organism finds itself in; the same alteration to the same genome can be fatal or providential depending on external circumstances. The parts of the genome coding for Darwinian preadaptations are noise until the environment changes in a way that makes them useful, at which point they become signal. This determination may depend on any feature of the environment, even one as tiny as a virus. Its full complexity is constantly available as a filter to turn mutational noise into biological information, and neither of these categories is particularly meaningful in the absence of that full complexity.

Biological information only counts as biological information when considered as part of an entire detailed world. This makes it look, to a philosopher, something like 'truth', a property sentences only have at particular possible worlds. Because we process the syntactic aspects of language serially, and because computers can be built that do the same thing in the same way, we can easily forget that semantic interpretations are holistic in character, in that they require whole worlds or collections of worlds in the background for sentences to be true or false at. But in logic as we now understand it, there is no just such thing as a semantic interpretation without a 'model' standing behind it. Sentences are only true or false in the context of whole worlds or whole theories, and similarly biological information is only really

distinct from meaningless noise in the context of a particular environment in all its fine grainy detail, down to the last pathogen.

8 An Oracle for Solving Decision Problems

Is there a more formal way of thinking of the process of natural selection? If we did want to try to describe it in the language of information processing, how would we do it? We can actually do it fairly well with a model of computation that is ridiculously strong, one usually considered too strong to represent any actual information-processing system, a model which tends to be presented as a limiting case showing what could happen if certain necessary assumptions were relaxed. From a formal point of view, natural selection is something like an oracle for solving a set of decision problems.

A decision problem is any logical problem that can be posed in a form that admits a yes-or-no answer. An example is Goldbach's conjecture, the assertion that every even number greater than two is the sum of two primes. The conjecture is either true, or false; which it is is a decision problem. An organism either contributes its genes to the next generation or it doesn't which makes fitness a yes-or-no proposition, or anyway lets us treat it that way as a first approximation.

An oracle is an imaginary computational device that, presented with a putative answer to any computational problem, can infallibly evaluate it as correct or incorrect in a single computational step. An oracle that could infallibly tell whether a yes-or-no answer to any decision problem was the correct one could resolve thorny mathematical questions like Goldbach's Conjecture in no time at all. We would just guess 'yes, every even number greater than two is the sum of two primes', and the oracle would tell us if we were right or wrong. Of course, we still wouldn't know *why* we were right, but evolution doesn't explain itself either. We have to reverse-engineer our way to the whys of things, as we would if an oracle told us that Goldbach's Conjecture was certainly true. We can think of the organism as the putative answer to the problem of how to deal with its environment, and successful reproduction as an evaluation as correct. There is no minimum number of steps this process must involve, and no dependence of the time required on the complexity of the decision problem being solved, so only an oracle will do as a formal model.

9 Oracles and Adaptive Landscapes

To tease out the full implications of this simple way of conceiving of natural selection, we need to think about a problem that came up in the 1960's about a particular type of evolution. At that time, some doubt was expressed that anything as complicated as a protein molecule could possibly have evolved in the amount of time available.

With twenty possible amino acids at each locus, a relatively modest chain of 100 amino acids has about 10^{130} possible sequences. Since the universe is only about 10^{18} seconds old, it seems as if an unreasonably large number of sequences would have to be tried out each second to find a particular sequence with a particular function in the amount of time available.

John Maynard Smith demonstrated, however, that this supposed difficulty is more imaginary than real.[9] He did it by making an analogy with a simple word game. Suppose we're playing a game in which we start with some four-letter English word – Maynard Smith illustrated his argument by starting with 'word' itself. The objective is to transform that word into some other actual English word through a chain of intermediates that are themselves all valid words. His example was 'word – wore – gore – gone – gene', but of course there are plenty of other examples we could give, one being 'love – lave – have – hate'.

There are 26 letters in the alphabet, so the space of four letter English words contains 26^4 or 456,976 points. However, no point is more than four changed letters away from any other. Whether or not it is possible to get from point a to point b in this space in a given amount of time is not a matter of how many points there are in the space in total, it is a matter of whether or not there is a valid bridge between the two words. Whether the space as a whole is easy to get around in depends on whether there is a percolating network of such bridges all through the space, which will take you from almost any point to almost any other in a few steps. There may be other valid destinations also reachable from the starting point by chains of valid intermediates, and there may be dead ends, and detours, all of which complicates the statistics a bit, but there is certainly never going to be any need to try out half a million words just to get from one point to another.

Why is the problem so much simpler than it initially seemed? It's easy to mistake this for a counting argument of some kind, to suppose that the person who was raising difficulties about the time it would take to search these sorts of combinatoric spaces was simply making a mistake about the numbers. But that isn't really what's going on here at all. What makes the word game easier than it otherwise might be is the fact that *we* all speak English, and therefore are able to instantly recognize that some of the four-letter combinations adjacent to a given word are themselves valid English words, while others are not. We aren't really blundering through the space at random. Each successive step involves a process in which we evaluate the available candidates to determine whether they are actually words. So the evolving word has to pass repeatedly through the filter of our brain, and pass a test for validity on the basis of knowledge stored there. Paths that terminate in dead ends don't get searched any further. We don't need to sample every point in the space, we just need to search down a tree.

Here, our knowledge of what counts as a valid English word is playing the same role that the environment plays in the process of natural selection. Both act as filters, weeding out unsuitable candidates. The word itself doesn't do any 'information processing' in this process. It is subjected to the decisions of an external oracle, which gives a yes-or-no answer to the question of whether or not a proposed next step is actually a meaningful English word in its own right. It is in this sense that this sort

of process can be thought of as the solving of a series of decision problems by an oracle. Similarly, the evolving protein itself doesn't need to 'know' whether an adjacent sequence represents an adaptive variant of the original design, or what it might be useful for. The environment will supply that information instantaneously and for free, by either killing a mutant cell or letting it prosper.

To appreciate the enormous creative power of natural selection, we really only need to ask one further question. What happens to the process of evolution, in this sort of combinatoric space, if the evolving objects get more complicated?

An obvious and intuitive way to extend the analogy would be to move from thinking about the space of four letter English words to the space of eight letter English words. As the examples of the reachability of 'love' from 'hate' and of 'gene' from 'word' illustrate, it is relatively easy to get from word to word in the space of four-letter words through a continuous chain of viable intermediates. However, there are also isolated islands like ALSO and ALTO, a pair of words that is not connected to any other word. Do these sorts of isolated islands become more or less common as the number of letters in each word goes up?

Take any eight-letter word – say, CONSISTS or CONTRARY. Does it have a meaningful single-mutation neighbor? Typically, no. It is usually impossible to move from eight-letter English word to eight-letter English word through a continuous chain of valid intermediates. So it might easily seem to us as if this sort of evolution ought to get more difficult, ought to slow down, as the evolving objects get more complicated. But this is precisely where the analogy between natural selection and any kind of cognition breaks down most badly.

The greater difficulty in navigating through the space of eight-letter words is not a consequence of some innate topological property of high-dimensional spaces. It is a consequence of the limits of the human mind, and the consequent simplicity of human languages. Natural human languages have, at most, on the order of 10^5 words. The space of eight-letter sequences contains on the order of 10^{10} sequences. Naturally, since less than $1/10^5$ of the sequences in the space are meaningful words, and since each word only has 208 neighbors, most meaningful eight-letter English words will not have even a single meaningful neighbor. But the evolution of proteins is completely different. Nature just doesn't face this sort of cognitive limitation. It's not as if she can only remember the designs of a few thousand proteins. She doesn't remember things at all, that's not how this particular oracle does its magic. Anything that actually happens to work in competition, that in some way, however improbable-seeming, leads to the production of adult offspring, simply works.

In the absence of human cognitive limitations, what actually happens to this sort of space as it acquires more dimensions? There is no a priori reason why viable protein designs must be vastly more sparsely distributed in a space of longer sequences than they are in a space of shorter ones. The change in their density depends on things like the degree of modularity in the molecule's design, and the commonness of neutral mutations. Suppose, as a limiting case, that there is a constant, unvarying probability per locus, p , that some mutation at that locus leads to an adjacent sequence which is as adaptive as or more adaptive than the original sequence. As the length of the evolving sequence increases, there are more loci on it. All other things being equal,

the probability that *some one* of these loci can be mutated in some adaptive way should go up as the number of loci increases. Call the length of the evolving protein, the number of loci, therefore the dimensionality of the sequence space, L . The probability that there is some adjacent sequence that can be moved to without loss of fitness is then just $1-(1-p)^L$. This quantity goes up with the length of the evolving chain. As L goes up, this probability asymptotically approaches 1.

As it becomes more likely that, for any given sequence in the space, *some one* of the adjacent sequences is as fit or fitter, it also becomes much more likely that the space as a whole is permeated by a percolating network of adjacent sequences each of which is just as fit as its neighbor. Thus, as the length of the evolving chain goes up, the space becomes easier and easier to get around in, not harder, as intuitions derived from the way we humans filter information might suggest.

On the basis of this model, it seems possible to argue that there may actually be a *minimum* complexity below which things cannot evolve efficiently. There may be countervailing effects that in fact cause p to fall in most cases as L goes up, but if the decline is less than proportional to the increase in L , percolation still becomes more probable as L goes up. No doubt some times it is, and some times it isn't; the sorts of evolved complexity we actually see in the world around us should be biased in favor of those cases in which p declines at a less than proportional rate.

How fast p declines with rising L depends on the precise characteristics of the thing that is evolving. If mutations typically have very small effects on phenotype, p should be fairly large no matter how complex the evolving object is. Modular designs limit the extent to which p depends on L . The modularity of life on a molecular scale may have more to do with limiting the decline of p than anything else. Neutral percolation can continue to be easy no matter what L is, if the evolving molecules are set up in a way that means that many mutations have no great effect on fitness. None of this seems to put a dent in the counterintuitive conclusion that evolution should actually often get *easier*, not harder, as the evolving object becomes more complicated. Adding degrees of freedom makes discrete combinatoric spaces more densely interconnected, not less, so in the absence of human cognitive limitations or other countervailing effects, free evolutionary percolation should become more and more likely as the evolving entities get more complex.

This conclusion, while it seems to be inescapably implied by Maynard Smith's model, may strike some readers as repugnant to common sense. Perhaps it is, but on the other hand, it does seem to fit fairly well with the facts. As far as we can tell, the evolution of life on Earth is not particularly slowing down as the evolving organisms and ecosystems get more complicated. There has been rather a lot of complex evolution since the Precambrian, or in other words during the last ten percent of life's history on earth. Things seem to have moved much more slowly when the evolving organisms were simpler. The time needed to evolve a human's brain from a lemur's was actually much shorter than the time needed to evolve a lemur's brain from that of a fish.

A case could even be made that evolution has speeded up significantly as the evolving organisms have become more complex. An awful lot has happened to life since the relatively recent event of the rise of the flowering plants. The whole

evolutionary history of the apes and hominids has occurred in the last few percent of the Earth's history. The apparent implications of Maynard Smith's model and the data from the fossil record are actually in complete agreement. What the model seems to predict is what we actually observe. It's common sense, derived from our own human experience of designing things using our own limited cognitive capabilities, that makes the oracular power of the selecting environment, its ability to turn any problem no matter how complex into a decision problem which it can solve in no time at all, seem so unlikely.

Things get harder and harder for *us* to redesign as they get more complicated, because we're a little stupid. But there's no upper limit on the complexity of the systems natural selection can optimize, and no necessary dearth of extremely complex structures and ways of doing things in nature's infinite library of random designs. Complexity, as Kauffman first pointed out many years ago,[10] is available to nature for free. All the information represented by the human genome was obtained 'for free' from an oracle – but actually using it, and sending it from place to place in the cell, is energetically very expensive.

10 An Epistemological Moral

What use is it to know all this? Does all this philosophical reflection, all of this very explicit and careful restating of the obvious, actually help us in any way in, say, understanding cancer? In fact, there's an epistemological moral here for the sciences of complexity in general, one that is in fact particularly salient to the way we think about cancer. The moral of the story is that biological complexity is not at all simple, that it's actually really, really complicated.

The way we've done physics, since the time of Descartes and Newton, is by assuming that behind the apparent complexity of nature, there's actually a deep simplicity, that if we can read a set of rules off of the behavior of some physical system we're likely to have found a complete specification of its nature that is much simpler than the system seemed at first. We might think of this approach, which has worked so remarkably well in physics, as reflecting a sort of unconscious commitment to Platonism on the part of physicists, since the essence of Platonism is the idea that the apparent variety, changeability, and complexity of the world we experience is just the result of varying combinations of much simpler and more fundamental base-level entities which can not be perceived with the senses.

On the basis of this commitment, we really have been able to discover very simple rules – we call them 'fundamental physical laws of nature' – which can be stated, in their entirety, in mathematical language, and which actually do completely determine the character of the systems they govern. One electron can be treated, by us, as being pretty much the same as another because they really are all pretty much exactly the same and will both behave in exactly the same way in exactly the same situation. The most accurate description of nature and the simplest one often coincide in physics.

Because so much of the apparent complexity of nature has revealed itself to be the product of these sorts of simple physical laws, we now have the expectation, as

scientists, that behind complexity in general there is always an underlying simplicity, and that once the symbolically-expressible rules governing a system are read off of its behavior, we will know all about it. To think this way about cells, however, is to make too much of the analogy between the sort of rules that 'govern' the behavior of electrons and the sort of rules that 'govern' the expression of a gene.

Behind the apparent complexity of the electron's behavior is a simple set of rules. Once we have them, we know all about it. The expression of a certain gene may also seem complex, and yet we may discover that it, too, seems to follow certain simple rules. But, though it may well be true that behind the apparent complexity of the gene's behavior there is some set of simple underlying regularities, it's also certainly going to be true that behind these apparently simple underlying regularities there is actually even more physical complexity lying in wait. The simple conventional rules that govern the behavior of biologically meaningful categories of molecules in cells are not the ultimate physical 'foundations' on which the cell's dynamics are based, they are the complex biomechanical consequences of everything else going on in the cell.

The idea that a certain sequence is reliably transcribed in a certain way can only be pushed so far; at some point, physical differences in the molecule whose sequence it is – say, its folding state – start to matter, and a qualification must be added to the rule. But that qualification will not be a simple one; it must consist of an account of all the various things that can physically happen in the cell when the molecule is in all the various abnormal states that are physically possible for it. This is why the cellular mechanics of cancer are so difficult to pin down; once the cell's machinery is no longer functioning in the usual way, there are an unimaginably vast number of other things it could be doing. As it turns out, lots and lots of them are ways of being cancer.

In this huge wonderland, anything at all about the cell's state could end up mattering in almost any way. We are outside the narrow conventions of the normal human cell in cancer, and the volume of parameter space that lies available for exploitation just outside those boundaries is huge. *We* immediately get lost in these vast new spaces, but the oracle of natural selection already 'knows' all the paths through their complexities without actually needing to know anything at all. Since natural selection acts on the whole phenotype of the cell all at once, there is no particular reason why the adaptive effects of a genomic abnormality should be simple enough for us to easily understand. Every kind of cancer is its own vast universe of complexity, and each and every single tumor may end up being incomprehensibly complex. There is no upper bound on the complexity of the new adaptations that can be picked out and amplified by this ignorant but unerring process, because there are no cognitive limitations to create such a bound.

Probably the most surprising thing about metastasis is that it can evolve at all, that such a high fraction of the cancers that naturally occur in human beings are able to develop elaborate adaptations for spreading through the body and thriving in new locations over periods of a few months or years. Of course, already having a full human genome, cancer cells have a big head start on the project of evolving new adaptations for thriving inside the human body. But what must make the evolution of

metastasis possible is the overlay of a huge percolating network of equally viable, perhaps only neutrally different cellular phenotypes on human genome space. As creatures, with limited capacity for processing information, we seem to vastly underestimate the density of new, workable complex designs in the realm of unrealized forms, supposing that there are only a few, and fewer, proportionally, as complexity increases.

(In a habit that goes back millennia, we suppose that the things that *are* must be some reasonably large fraction of the things that *could be*, though in fact they are a vanishingly small and rapidly shrinking sub-set. Evolution explores a combinatoric space that gets easier to move around in in a way that depends more or less linearly on its dimensionality, while the number of points in that same space increases *exponentially* with its dimensionality, so an evolving life-form, as it becomes more complex, moves more and more freely through a vaster and vaster set of possibilities, exploring smaller and smaller fractions of them at greater and greater speeds.)

Because of this inappropriate, lingering Platonic assumption (in most kinds of Platonism, the Platonic forms are supposed to be fewer than the things in the world, though in *Parmenides* Plato admits that this supposition is problematic) we look for simple explanations of the adaptive effects of abnormalities, and we tend to think there must only be a few kinds of colon or lung cancer. But this is an illegitimate transfer of the sort of intuition we derive from our own experiences with things like eight-letter English words whose profusion is limited by our stupidity onto a universe that is not stupid (or clever) at all.

The rules we can deduce about how a cell will respond to a particular perturbation will only be simple if we are willing to allow a multitude of un-described possible exceptions. Any attempt to describe *all* of the possible exceptions to such a rule would involve describing all of a cell's possible physical states in exact detail. Then we would be no better off than we are in dealing with a hurricane. Exact prediction of future events is only possible when the simplifying conventions work as they're supposed to. The rules that govern cells are only simple when they are imperfectly stated; anything simple we say about them is only a sort of executive summary, an approximate characterization of what's 'usual' for a system whose real behavior can become arbitrarily different in particular cases. No matter what things are like in physics, the simplest description of a system and the most accurate one do *not* ever coincide in biology. A fully accurate description of a biological system is a description of that system in all its unique complexity. (Of course, such a description is impossible – but that just means biologists will never run out of things to do.)

This makes many of the regularities we discover in the living world very different from the sorts of laws studied by physicists, and the kind of knowledge we can have about living things very different from the kind of knowledge we can have of the fundamental physical laws. Those things are actual necessities of nature, descriptions of universal mechanisms that couldn't possibly work in any other way. The regularities we observe in cells really are just local conventions, which serve to allow coordination between large numbers of molecules, when they actually obtain, which is not all the time, or in all cells. In death, all these regularities decay and disappear.

As far as we know, there is never any such breakdown of the rules governing electrons.

So a large, noisy neural net might in fact be a useful model of information processing in cells, but only if we keep in mind the caveat that the system will act like one all the time except when it doesn't, when the underlying mechanical properties of the components come to matter in ways a neural net model doesn't reflect. It's very optimistic to suppose that we can ever arrive at a detailed simulation or model of cellular function that also covers all abnormal cases, that accurately predicts the system's response to any perturbation, because anything at all can matter in an abnormal case, and the suppression of idiosyncratic detail is the essence of simulation. The only fully accurate model of a cell as information processor is the cell itself in all its complexity, in its natural environment, where anything at all can happen, and information as such does not really exist.

References

1. A Ben-Hur, H Siegelmann, 'Computing with Gene Networks,' *Chaos*, January 2004.
2. H Siegelmann, *Neural Networks and Analog Computation: Beyond the Turing Limit*, Birkhauser, Boston, 1998.
3. S Kauffman, 'Metabolic stability and epigenesis in randomly constructed genetic nets', *Journal of Theoretical Biology*, Vol. 22, No. 3. (March 1969), pp. 437-467.
4. S Kauffman, 'The large scale structure and dynamics of gene control circuits: an ensemble approach', *Journal of Theoretical Biology*, Vol. 44 (March 1974) pp. 167-90.
5. D Lewis, *Convention: a Philosophical Study*, Blackwell, London, 2002.
6. T Schelling, *The Strategy of Conflict*, Harvard Press, Boston, 1960.
7. J Maynard Smith and E Sathmáry, *The Major Transitions in Evolution* (chapter 3) Oxford University Press, 1995.
8. C Darwin, *Origin of Species* (chapter 4) Charles Murray, London, 1859.
9. J Maynard Smith, 'Natural selection and the concept of a protein space', *Nature*, Vol. 225 (February 1970) 563 -564.
10. S Kauffman, *The Origins of Order*, Oxford University Press, 1993.