

Measurement of Evolutionary Activity, Teleology, and Life

Mark A. Bedau

Philosophy Department, Dartmouth College
Hanover, NH 03755
mark.bedau@dartmouth.edu

Norman H. Packard

Physics Department and
Center for Complex Systems Research – Beckman Institute
University of Illinois at Urbana Champaign
405 N. Mathews Street, Urbana, IL 61801
n@complex.ccsr.uiuc.edu

October 7, 1996

Abstract

We consider how to discern whether or not evolution is taking place in an observed system. Evolution will be characterized in terms of a particular macroscopic behavior that emerges from microscopic organismic interaction. We define evolutionary activity as the rate at which useful genetic innovations are absorbed into the population. After measuring evolutionary activity in a simple model biosphere, we discuss applications to other systems. We argue that evolutionary activity provides an objective, quantitative interpretation of the intuitive idea of biological teleology. We also propose using evolutionary activity in a test for life.

Keywords: evolution, teleology, life, Turing test, adaptation, artificial life.

1 What is Evolution?

Our paradigm of an evolving system is the biosphere. Emerging somehow from inorganic origins, it has produced a myriad succession of marvelously adapted organisms, from the very simple to the quite complex. And the process is ongoing still.

Although we all have some common-sense grasp of the process, it is difficult to say precisely what evolution is. Evolution is change, but not all change is evolution. What distinguishes evolution from other kinds of change? Some despair of the prospect for answering this question. Consider Dobzhansky [?]:

There is no satisfactory general definition of evolution. “Sustained change” comes probably as close as possible at present. In the special case of biological evolution this may be amended to become “sustained change over a succession of generations,” to differentiate the evolutionary development (phylogeny) from the development of an individual (ontogeny).

Evolution is clearly more than sustained change, more even than sustained *complex* change. A turbulent fluid is continually undergoing complex changes but it is evidently not evolving in anything like the way that the biosphere is. Biological evolution is also not just complex change that propagates through successive generations. Large amounts of complex genetic deadwood—junk DNA—can accumulate in a gene pool with no real evolutionary effect.

Intuitively, the distinctive mark of evolution is the spontaneous generation of innovative functional structures. Implicitly designed and continually modified by the evolutionary process (on an evolutionary time scale), the structures persist because they prove sufficiently adaptive. The growth of adaptations causes the biosphere to increase in complexity, thus providing an arrow of time not implied by mere complex change, even if sustained through many generations. But how can the idea of a system continuously and spontaneously generating adaptations be expressed quantitatively? How can it be measured in a model or in the real world?

Our approach to answering these questions is to quantify the degree to which a system exhibits the continual spontaneous generation of adaptive forms. Specifically, we measure the degree to which new genetic combinations are persistently used in a population. This quantity is an objective, empirical measure of the level of evolutionary activity in an artificial or natural system, but it has important broader implications, as well. For one thing, it provides a natural, quantitative interpretation for the controversial but intuitively compelling view that a biosphere inevitably exhibits teleology (purposive or goal-directed behavior). In addition, it suggests a new approach to an empirical, quantitative understanding of life.

The evolving biosphere is a complex web of organisms interacting with each other and with their environment. Following the tradition of statistical mechanics, we regard each organism as a microscopic element of the biosphere, and we regard evolution as a macroscopic property that emerges as a consequence of the interactions among all organisms and their environment. Evolution is a macroscopic, long-term property of a population of interacting organisms.

The macroscopic state of a thermodynamic system is characterized by thermodynamic variables such as temperature, pressure, and specific heat. We would like to define analogous variables that characterize the macrostates of an evolving biosphere. Thermodynamic macroscopic variables are typically static quantities that characterize a time-independent equilibrium. By contrast, an evolving biosphere changes constantly, so its macroscopic characterization must inevitably include dynamic properties, especially in the long run. Evolutionary dynamics seem to possess metastable states, which leads us to define macroscopic quantities that are averaged over a short time scale (and possibly averaged over all or part of the population), and examine how they change over a longer time scale.

Evolution is driven by genetic changes. We will adopt the simplifying assumption that genetic changes are random changes in a genome. Each organism interacts with other organisms and its environment by means of an information processing mechanism that takes sensory data as input and yields an action as output. The organism's genome encodes the mechanism that causes the behavior, though often behavior is not encoded directly but arises through the complex interaction of other directly encoded mechanisms. Selection occurs as a result of the behavior; more efficacious behavior increases an organism's probability of survival, on average. Typically selection takes

Figure 1: The part of the figure enclosed in the ellipses represents how individual genomes change through a lineage as a result of changes in the genomes g^t (taken to be random events), which cause a change in the resulting phenotypes Π^t . Selection acts on the phenotypes in the context of an environment consisting partly of a population of competing phenotypes (the series of ellipses). If an individual phenotype Π^t survives the selection process, its genotype g^t (possibly mutated) is transmitted in reproduction. The gray arrows represent how the evolving distribution of different genomes in the population, $P^t(g)$, emerges out of the combined effects of selection of individuals and genetic alterations.

place over a longer time scale than reproduction. The indirect nature of the effect of genetic change on the dynamics of the population is illustrated in figure 1.

We have described the process of genetic change, and the resulting selection, on an organismic level. The macroscopic state, however, is a function of the distribution of the entire population over the space of possible genomes. The net result of the genetic change and selection for the entire population is an evolution of the population distribution over the space of genomes, as illustrated in the lower right of figure 1. The statistical characterization of evolution concerns the dynamics of this population distribution.

In the final analysis, the classical adaptationist perspective embodied in figure 1 must be qualified to allow for neutral evolution [?]. It must also be modified due to the presence of nonadaptive evolutionary forces, such as those due to developmental constraints [?]. We will concentrate, however, on the adaptive forces of evolution precisely because it is this aspect that we seek to quantify.

Spontaneous generation of innovative change epitomizes the dynamical nature of adaptive evolution. Thus, evolutionary change gives the appearance of having a direction as successful lineages become progressively better adapted to their environments. In the long run, however, evolution has no specific predetermined global goal. Evolutionary change might show overall statistical tendencies, such as those captured by the statistical measures defined below, but the details of a biosphere's global state are in constant flux. No Master Plan explicitly specifies the biosphere's form. At any given instant in a given local biological context for a given species, some specific

evolutionary changes would be better and others worse. So, each species has certain temporary local optimization criteria, which generally differ for different species. These local optimization problems generally admit many possible approximate solutions. The particular evolutionary path a lineage follows results from many random genetic changes that survive as solutions or partial solutions to the local optimization problems.

The process of selection acting on the products of these genetic changes causes complex functional changes in the organisms in the long run. These changes emerge *a posteriori* out of the contingencies of the evolutionary dynamics. Evolutionary learning apparently takes place, but without an *a priori* specification of what needs to be learned. A feature that is good in one context might be bad in another; it depends on the surrounding population, which constantly shifts as the population evolves. In effect, the population is adapting to itself and to the environment, so the specific optimization criteria change implicitly during evolution. Whereas *extrinsic adaptation*, typical of artificial learning systems, has an explicitly specified goal, adaptation without an explicitly defined goal is *intrinsic adaptation*. Intrinsic adaptation is exhibited by a variety of models [?, ?, ?, ?, ?, ?, ?, ?], perhaps the simplest of which is [?].

A fundamental question in the general science of evolution is to determine what macroscopic variables characterize the behavior of an evolving system. Below we define a variable that reflects the degree to which new genetic combinations are continually produced and persistently used in the population. Initially sacrificing some amount of generality in favor of the clarity provided by an explicit example, we first explain this statistic in the context of a simple model biosphere. After describing the model, we measure this quantity in the model. Then we discuss how to make analogous measurements in a variety of other systems.

We then argue that evolutionary activity may reflect the occurrence of teleological or goal-directed activity. Though teleological behavior has often been identified on strong intuitive grounds as a hallmark of living, evolving systems, the concept has not been formalized in scientific theory. Much of the scientific literature is inimical to the idea that teleology is even compatible with formal scientific theory. Below, we argue the opposite—the intuitive notion of teleology is indeed quantifiable, by means our measure of evolutionary activity. Just as recent studies of chaos have shown that deterministic systems could be unpredictable, we claim that deterministic systems may be

teleological. Similarly, just as a central question in the study of chaos has been how to quantify unpredictability, in our study of teleological systems we must face the analogous question of how to quantify teleological behavior.

In the final section, we propose using evolutionary activity as a test for life. Where the necessary experimental measurements can be made, the test should be applicable to evolving systems in the real world. In the realm of artificial life models, the experimental measurements are nearly always possible, at least in principle. Thus, the test is a quantifiable analog to Turing's test in artificial intelligence.

2 A Model Biosphere

Every model of living processes must model phenomena on some specific level. A model of a chemical soup might attempt to show how life could originate. Here, we model a population of organisms in an environment to show how evolution emerges. Following the tradition of statistical mechanics, the organisms and the environment in our model are highly idealized. Just as the Ising model from statistical mechanics represents a magnetic solid merely as a lattice filled with zeros and ones, the model presented here abstracts from a host of details about organisms and their environments.

2.1 Tropic Bugs

The model used here is a modification of an earlier model designed to be simple yet able to capture the essential features of an evolutionary process [?]. The earlier model consisted of organisms in a two dimensional world. The only thing that exists in the world besides the organisms is food. Represented as a field of values, food is put into the world in heaps that are concentrated at particular locations, approximating a continuum field with a gradient away from a central location. Food is refreshed periodically in time and randomly in space. The frequency and size of the heaps are variable parameters in the simulation.

The food represents energy for the organisms. They interact with the food field by eating it, decrementing the food value at their location, and incrementing their internal food supply. The organisms in this earlier model are endowed with enough innate intelligence to follow the gradient; they

survey their local neighborhood and move in the direction of most food. Since the behavior of these organisms is simply the tropism of ineluctably climbing the food gradient, we call them *tropic bugs*. Movement expends energy, so each step taken exacts a tax on the internal food supply. If this internal food supply drops to zero, the organism dies and disappears from the world. On the other hand, an organism can remain alive indefinitely if it can continue to find enough food.

When an organism accumulates enough food, it produces some number of offspring. Reproduction by tropic bugs is controlled by two quantities regarded as “genes” for the organisms: the number of offspring, g_{off} , and the threshold for reproduction, g_{th} . These genes are changed during reproduction by random amounts, analogous to point mutations.

Evolution will occur in a population only if the environment stresses the population, so that some of its members can be better adapted to coping with the stress than others. The only stress faced by the tropic bugs is to find enough food to remain alive. Any evolution that occurs in the model is the effect of this one environmental imperative.

2.2 Strategic Bugs

The model biosphere used here differs from the tropic bug model in two respects. The primary modification is to allow organisms to follow individually different strategies for finding food, so we call these organisms *strategic bugs*. The other modification provides a second source of genetic novelty.

The behavioral disposition of both tropic and strategic bugs is genetically hardwired. But whereas every tropic bug follows the same rule for climbing the food gradient, each strategic bug follows its own individual food-finding strategy. A behavioral strategy is simply a map taking sensory data from a local neighborhood (the five member von Neumann neighborhood) to a vector indicating a magnitude and direction for movement:

$$S : (s_1, \dots, s_5) \rightarrow \vec{v} = (r, \theta). \quad (1)$$

A strategic bug’s sensory data has two bits of resolution for each site (least food, somewhat more food, much more food, most food). Its behavioral repertoire is also finite, with four bits of resolution for magnitude r (zero, one, ..., fifteen steps), and four bits of resolution for direction θ (north, north-northeast, northeast, east-northeast, ...). As with tropic bugs, strategic bugs

pay a movement tax proportional to the distance traveled. A tax is also levied just for living, so strategic bugs must continue finding food to survive.

The graph of the strategy map S may be thought of as a look-up table with 2^{10} entries, each entry taking one of 2^8 possible values. This look-up table represents an organism’s overall behavioral strategy. The entries are input-output pairs that link a specific behavior (output) with each sensory state (input) that an organism could possibly encounter. Whereas tropic bugs have only the two genes g_{off} and g_{th} , strategic bugs have 2^{10} additional genes, one for each entry in their strategy look-up table. Although still finite, the space of genes in the strategic bug model is greater than in the tropic bug model by three orders of magnitude. This allows evolution in a much larger space of genetic possibilities, which better approximates a biological world with an infinite number of possibilities.

Evolution requires a source of random variation. Just as the two genes carried over from the tropic bug model can change during reproduction by random amounts, analogous point mutations of the strategy genes can change the output values of entries in the strategy look-up table. A parameter regulating the mutation rate, i.e., the “strength” of mutations, determines what fraction of the table mutates during reproduction.

The genome in the strategic bug model is large enough that it becomes reasonable to allow sexual reproduction, or at least a simple version with haploid crossover. The second respect in which strategic bugs differ from tropic bugs is that, whereas both tropic and strategic bugs reproduce asexually, strategic bugs can also reproduce sexually. A strategic bug can tell when it is next to another bug. If two healthy bugs (i.e., bugs with sufficient internal food) are adjacent, they flip a coin to decide whether to produce offspring. Analogous to the exchange of genetic material during crossover, each child contains a mix of genetic material randomly chosen from the two parents. There is no distinction between “female” and “male,” so sexual reproduction here simply means offspring produced with a mixture of parental genetic material.

The strategic bug model illustrates a novel form of evolving dynamical system that has recently been developed for models of immune networks and autocatalytic networks [?, ?, ?, ?, ?, ?] as well as for models of parallel computation in machine learning [?, ?, ?, ?]. These systems consist of a state space that changes with time, with a meta-dynamic specifying the state-space evolution. In the strategic bug model, the momentary dynamical rule

includes all the individual strategies for each organism present in the population, as well as the rule that governs the input of new food. The strategic bug metadynamic specifies how new elements come into the population and what rules govern them. Since new genes interact with the world and with each other just as old genes do, the metadynamical element in the strategic bug model is simply the generation of new genes by mutation and crossover, which creates new elements of the population. More complex forms of interaction between organisms (besides simply reproducing) would necessitate a more complex metadynamic.

There many free parameters in the strategic bug world which must be set at the beginning of a simulation; we now describe them all. The parameters are named as variables that might appear in a computer program; the values listed in square brackets are those we actually used for the measurements presented in the following section. To interpret the parameters it is useful to realize that the time scale is basically set by the amount the organisms are taxed for their activity, and the rate that food is coming into the world. The world was a 128×128 lattice, with an integer food value between zero and 255 at each site.

Bug initialization: The world begins with an initial population of size `initial population size` [50]. Some of the parameters that specify the bugs are the same for all bugs. One of these is `mouthful` [50], the maximum amount of food a bug can eat in one gulp (if there is less than this amount at the bug's current location, all food is eaten). Other parameters that are the same for all bugs are the taxes, which subtract food from the internal food supply each elementary time step. One is `move tax` [10], the amount of food used per unit distance moved. Another is `reproduction tax` [0]. A third is `overall tax` [10], a metabolic tax for survival every time step.

Another group of parameters are needed to specify the initial values of genes that may be changed during reproduction. Each bug has a `reproduction thresh` [1000] for the amount of food needed to reproduce, and a `sex threshold` [800] for the amount of food needed to have sex (provided another healthy bug is in the neighborhood). Asexual reproduction yields `offspring num` [2] children. The strategy look up table is initialized with a fraction `strategy density` [.25]

of the entries set to random output actions, with all other entries set to the output action “do nothing.”

World initialization: The environment is specified by parameters that govern the input of food. The first such parameter is `food time` [2], the number of elementary time steps between each food input. Each food input consists of a lump of food placed at a random location. The lump has a maximum food value `food max` [250] at the center of the lump, and extends spatially, diminishing linearly to zero at a radius `food width` [0.2], given in units of the size of the entire world.

Evolution: Parameters that specify the evolution are primarily the mutation levels. First there is the mutation level `tropic mut` [0] of the tropic genes, `reproduction threshold` and `offspring num`. Then there is also the mutation level `strategy mut` [0.1] of the strategy table, given in terms of the fraction of total number of entries. The final evolutionary parameter is `crossover fraction` [0.4], the fraction of strategy table entries that are exchanged during sexual reproduction.

3 Evolutionary Activity for Strategic Bugs

An actively evolving system is continuously and spontaneously generating adaptive change. Its gene pool is continually shifting, absorbing new genes and rearranging existing ones, but genetic changes can persist in the long run only if the organisms with the new genes thrive. Any change that lessens an organism’s ability to survive will, on average, be unlikely to appear in later generations. In other words, those changes that *do* get absorbed into the gene pool must, on average, either enhance survival or, at worst, be neutral. Thus, the continual retention of new useful genetic material indicates that the population is continually enhancing its gene pool.

The rate at which new genes are introduced does not reflect genuine evolutionary activity, for the new genes may well be useless. Likewise, the mere persistence of genetic innovation alone is insignificant, because a persistent gene may well be unused and irrelevant. *Persistent usage* of new genes is what signals genuine evolutionary activity. In the context of the strategic bug model it is simple to defend the appropriateness of measuring persistent usage. Since “using” a gene amounts to moving and movements exact a

tax, persistent usage of a gene necessarily indicates that the benefits brought about by the movement result in enough of a gain to offset the tax; otherwise the gene would have disappeared from the population.

In this section we define a statistic designed to measure the rate at which new useful genetic material is being incorporated into the gene pool. We initially implement this measure in the model for strategic bugs. In a later section we discuss how to apply this measure to a wide variety of other artificial and natural systems.

3.1 Usage Statistics

If we define evolutionary activity as the population’s continual absorption of new genetic combinations that come to be persistently used, we must have a way to measure how persistently genes are used. To do this for the strategic bugs, we first define quantities that measure the usage of the genes in a bug’s strategy look-up table. (For simplicity’s sake, we refrain from measuring the usage of the two genes, carried over from the tropic bug model, that govern reproductive threshold and number of offspring.) Letting i label the bug and j label the gene within the bug, consider every gene g_{ij} of every bug as having a “usage counter” u_{ij} attached to it, initialized to zero. Recall that each entry in a look-up table is an input-output pair. Every time a particular input situation is encountered and its paired output entry in the table is used, the corresponding usage counter is incremented. During asexual reproduction, the usage is reset to zero if the corresponding gene mutates, and otherwise the usage is carried with the gene to the offspring. Offspring produced by sexual reproduction inherit their parents genes with their corresponding usages intact. In this way, a given gene’s usage preserves information accumulated over many generations along the lineage through which the gene is inherited.

To record what percentage of the genes in the entire gene pool for the strategic bugs have given usage values at a given time, we define a *usage distribution function*, $N(t, u)$, by apportioning all of the genes in all of the bugs into “bins” for given usage values u at given times t , as follows:

$$N(t, u) = \frac{1}{N_g} \sum_{i,j} \delta(u - u_{ij}^t). \quad (2)$$

Here, u_{ij}^t is the usage that gene g_{ij} (the j^{th} gene of the i^{th} organism) has

Figure 2: A view of the usage distribution function $N(t, u)$ for three particular values of t . At a given time t , $N(t, u)$ is a usage histogram, with usage increasing to the right and the proportion of genes with a given usage increasing vertically. Initially, at $t = 0$, all genes have zero usage, so a usage histogram has a single peak at zero usage. Then some genes get used and acquire positive usage, so a usage histogram for $t = t_1$ shows a tail sliding off the peak at zero usage. Later, at t_2 , a gene or group of genes has been used repeatedly, so a usage histogram shows a clump of genes with positive usage. Later still, at t_3 , after those genes have seen even more use, the bump has moved toward higher usage values. In this way, activity waves propagate through the usage distribution $N(t, u)$.

accumulated by time t , and $\delta(u - u_{ij}^t)$ is the Dirac delta function, equal to one if $u = u_{ij}^t$, and zero otherwise. So, $\sum_{i,j} \delta(u - u_{ij}^t)$ simply counts the number of genes that have usage u . This sum is then normalized by dividing by the number of genes in the population, N_g . Thus, $N(t, u)$ is the total fraction of genes in the entire population having usage u at time t .

One can visualize a usage distribution function $N(t, u)$ as a three dimensional surface (landscape) over a two dimensional time/usage grid. The value (height) of any given location (t, u) on this surface simply reflects the proportion of genes that have usage u at time t . A significant peak in the surface around a location (t, u) would indicate that at times near t a significant proportion of genes had been used about u times.

3.2 Activity Waves

The usage distribution function for the strategic bug model turns out to have a complex and interesting structure. Initially, at $t = 0$, all genes have zero usage, so $N(0, u)$ has just one peak at $u = 0$. As time progresses and genes are used, the usage distribution function becomes positive for other values of u . If a beneficial gene or gene-cluster enters the population, it will come to be used persistently, and the time/usage surface $N(t, u)$ will show a certain structure. Specifically, at a certain time after the introduction of the beneficial genes, $N(t, u)$ will have a bump due to the genes' persistent usage. (See figure ??.) As long as the genes remain beneficial, they will persist and their usage will increase. As their usage increases, the bump will move in

Figure 3: Above: A diagram of $N(t, u)$ for the strategic bug model. Time increases to the right, and usage increases up. Darker shades of gray indicate higher values in the usage “bins.” The value at a given point (t, u) on the surface indicates what proportion of the genes in the gene pool have usage u at time t . Swaths of higher values (darker shades) moving up and to the right are activity waves. New activity waves start when genetic novelties prove to have some utility for the organisms. The slope of an activity wave reflects the frequency with which the genes contributing to it are used. Below: The corresponding graph of evolutionary activity $A(t)$, with the reference point u_0 chosen to be 75. Glancing at the distribution $N(t, u)$ shows that in general usage exceeds u_0 only for those genes that are in activity waves. Peaks in $A(t)$ correspond to an activity wave bursting past $u_0 = 75$.

time with a velocity proportional to the frequency with which the relevant situation is encountered.

Such moving features appear as waves over the time/usage plane. We call them *waves of evolutionary activity*, or simply *activity waves*. The upper parts of figures ?? and ?? illustrate the waves in a time/usage diagram for the strategic bug model. Each of the graphs in figure ?? corresponds to a vertical column in the time/usage diagram, with higher values in the usage “bins” indicated by darker shades of gray.

If the genes in a wave continue to be used by all organisms with roughly the same frequency, the wave propagates at a constant velocity and appears as a relatively straight line over the time/usage surface. The slope of a wave reflects the frequency with which the genes in it are used. If this frequency is changing, the wave curves. If usage of a group of genes stops, then the slope of the wave levels to zero for as long as the genes persist in the gene pool. As the genes in this flat wave are pushed out of the gene pool by mutations, the wave’s height drops to zero. If the genes start to be used again, the wave’s slope (but not its height) will increase.

A gene will continue to contribute to just one wave during its time in a lineage in the gene pool. New activity waves are created only when a newly created gene is first used in its “life” in the gene pool. Mutations can converge, creating a new gene that accidentally “copies” an existing gene; use of the accidental copy gene will initiate a *new* activity wave.

Activity waves emerge out of the local interactions between “microscopic”

Figure 4: $N(t, u)$ for the same simulation as figure 3, but for a longer time. New bursts of evolutionary activity continue to emerge.

organisms. The medium of the activity waves is actually the genetic material of the entire population. All genetic changes are perturbations that can potentially initiate a new activity wave. Waves will actually start when the new genes are beneficial in the local biological context. “High” activity waves reflect “large” clusters of useful genes.

An activity wave in a biological system reflects the persistence of a useful piece of information about the world (coded by a gene or gene cluster) within the information processing machinery of an organism. New genes are like guesses about what might be good to do in the world. When a guess is more or less correct, the information persists because it promotes the organism’s survival. In this way, the continual emergence of new activity waves indicates that (on average) the organisms are effectively improving their internal models for what the world is like and what behaviors are most efficacious within it.

3.3 Evolutionary Activity

The existence of activity waves suggests a method for measuring a system’s evolutionary activity. As long as activity waves continue to occur, the population is continually incorporating and repeatedly using new genetic material; in short, the system is evolving. Our measure of activity, then, is based on measuring the flow of usage into the gene pool, with a new burst of usage corresponding to a burst of evolutionary activity.

Before assigning any evolutionary significance to a gene’s usage, we must distinguish short-term usage from long-term usage. Glancing at $N(t, u)$ shows that a large fraction of the genes have usage slightly greater than zero, but this short-term usage means little. Eventually, after further use, unhelpful genes are weeded out of the population. The only genes that accumulate usage above a certain value u_0 are those that end up contributing to activity waves; these are the genes that have proven their usefulness through acquiring long-term usage. Thus, we will call a gene *persistent* at a given time if its usage at that time exceeds a certain reference point u_0 .

The parameter u_0 is determined by the time scale on which useless genes

are replaced, which in turn is determined by the details of the organisms and how they interact with each other and the environment. In the case of the strategic bugs, u_0 is determined by a combination of the taxation rates and the rate and pattern by which food appears in the world. In any case, u_0 is to be set high enough so that most of the useless genes disappear before their usage reaches the value u_0 . In practice, it is easy enough to identify a plausible reference point u_0 by glancing at $N(t, u)$ and picking a usage value above the initial large fall-off in usage (clearly evident in figures 3 and 4) as useless genes are weeded out. It is evident that a gene's usage exceeds u_0 in general only if it ends up contributing to an activity wave. The exact value of u_0 is not crucial. As long as u_0 is large enough, patterns formed by activity waves passing u_0 will be similar for a wide range of u_0 values.

As a preliminary to quantifying activity waves, we want a measure of “bulk” usage over time in the gene pool that allows short-term usage to be distinguished from long-term usage. The *net persistence* $P(t, u)$ of a gene pool at time t relative to a given usage u is defined as the proportion of genes at t that have at least usage u :

$$P(t, u) = \sum_{u'=u}^{\infty} N(t, u'). \quad (3)$$

Long-term and short-term usage may be separated because the net persistence $P(t, u)$ explicitly depends on u . One can think of the net persistence $P(t, u)$ for given t and u as the “bulk” of the column of usage values stacked above (t, u) . Since the bulk of these usage columns decreases as u increases, $P(t, u)$ decreases monotonically with u .

If an activity wave passes a certain point in the $N(t, u)$ plane, the net persistence function $P(t, u)$ will be changing in the neighborhood around (t, u) . As time approaches t there will be a significant increase in $P(t, u)$, and as usage exceeds u there will be a significant decrease in $P(t, u)$. Thus, a passing activity wave can be quantified by the rate of change of $P(t, u)$ with respect to either t or u . To avoid noise introduced by fluctuations in the population, we focus on the rate of change with respect to u . Waves in $N(t, u)$ correspond to “cliffs” in $P(t, u)$. The height of an activity wave at a point (t, u) in $N(t, u)$ is reflected by the *steepness* with which $P(t, u)$ falls off at that point.

Thus, to quantify the passage of activity waves, we can simply measure the steepness of $P(t, u)$ at our reference point u_0 . That is, we define the

evolutionary activity $A(t)$ of a system as the rate at which net persistence is dropping at u_0 :

$$A(t) = - \left[\frac{\partial P(t, u)}{\partial u} \right]_{u=u_0} . \quad (4)$$

If new activity waves continue to be produced, then the population is continually acquiring new genetic material that is proving its usefulness through a significant amount of repeated use. In this case, the system is exhibiting significant evolutionary activity, and the measure of evolutionary activity $A(t)$ will be positive. If evolutionary activity is zero, then the gene pool is incorporating no new persistent genes, and the system is exhibiting no significant evolution whatsoever. Since $P(t, u)$ decreases monotonically with u , the activity $A(t)$ is never negative.

The usage distribution function $N(t, u)$ contains a wealth of information about the evolutionary process. Evolutionary activity $A(t)$ —the rate at which innovative genetic novelty is flowing into the system—is one especially fundamental aspect of $N(t, u)$, but $A(t)$ does not reflect all significant evolutionary events. Two examples can illustrate how other kinds of evolutionary events could be quantified from $N(t, u)$.

One kind of event with evolutionary significance is extinctions. The extinction of a species of organisms would appear in $N(t, u)$ as the diminution and eventual disappearance of an activity wave; a massive dying off such as the Cretaceous extinction would be an abrupt termination of a mass of waves. But since these waves would terminate above the reference point u_0 , extinctions would not be registered in evolutionary activity $A(t)$. To quantify the *net* change in a system's innovative genetic novelty, subtracting that portion lost from extinctions, one could simply take the time derivative of $P(t, u)$ at the reference point $u = u_0$.

Changes in genes' usage patterns are another significant kind of evolutionary event. These would be reflected in $N(t, u)$ as changes in the slope of activity waves. For example, a wave's slope will increase if a group of little-used genes starts to be used more frequently. Evolutionary activity $A(t)$ does not quantify changes in the complexity of the dynamics of the activity waves; if waves continue to be produced at a steady rate, then $A(t)$ remains constant even if the *pattern* of activity between the waves becomes dramatically more complex. However, if the usage distribution function $N(t, u)$ were separated into components consisting of waves travelling at given velocities, then the

complexity of the activity wave patterns could be quantified by correlating the strengths of the different components.

Our measure of evolutionary activity $A(t)$ is averaged over all individuals in the population. In addition, since usage is always passed to offspring except after a mutation, evolutionary activity is also averaged over generations. Thus, waves in a system’s evolutionary activity truly characterize the system’s global, long-term dynamics.

4 Evolutionary Activity in Other Systems

The measurements of evolutionary activity $A(t)$ reported above were all made in the strategic bug model. But evolutionary activity can be measured in a wide variety of other artificial and natural systems, as long as the system consists of a “macroscopic” population of “microscopic” entities for which usage can be defined clearly and appropriately. If usage at the system’s micro-level can be measured, then all the global macro-level quantities—the usage distribution function $N(t, u)$, net persistence $P(t)$, and evolutionary activity $A(t)$ —become well defined, just as before.

The key to applying our measure of evolutionary activity is a good definition of usage. In some evolving systems it is difficult to measure gene usage directly; in others, there are no genes *per se*. Nevertheless, it is often possible to develop other ways to measure usage. One of the merits of our approach to measuring evolutionary activity is this flexibility in the definition of usage. In this section, after indicating what constitutes an appropriate measure of usage, we illustrate how to measure evolutionary activity in a wide variety of other systems.

In the strategic bug model we counted the usage of genes in a gene pool, but it is possible to measure usage of other kinds of entities. Instead of genes, the micro-level entities could be *groups* of genes; they could even be the broad collection of genes and gene variations shared throughout a species. Since our aim is to measure the rate at which a system incorporates new functional units, the micro-entities can be *any* functional units, any units with adaptive significance. Our measure applies to any level at which natural selection operates. (Cf. the “units of selection” debate [?, ?, ?, ?].)

If usage counters are attached to inappropriate micro-level entities, our measure of evolutionary activity can register false positives and false neg-

atives. If the micro-level units fail to reflect some aspect of functionality, then some genuine evolutionary activity might be missed (a false negative). For example, assume that there is a genetic system in which *combinations* of genes can have adaptive significance over and above the adaptive significance of their individual component genes. In this case, crossovers could spark the spread of many quite beneficial new combinations of pre-existing genes, which would constitute significant evolutionary activity. But the proliferation of these adaptive genetic combinations would not generate activity waves if usage were counted only for single genes; the frequency of groups of genes could alter while the frequency of individual genes remained the same. To capture the occurrence of this sort of evolutionary activity, some way must be found to add usage counters to potentially functional gene combinations.

False positive readings of evolutionary activity can also occur if usage counters are attached to non-functional micro-level units. For example, only a small fraction of the eukaryotic chromosome has potential adaptive significance. Relatively short segments, exons, code for amino acids; these are the genes. The intervening segments, introns, are without adaptive significance. So, consider a genetic system that is undergoing no change in exons but rapid change in introns. In this case, the system would not be significantly evolving. However, if usage were counted at the level of individual base pairs in the nuclear DNA or at the level of the codon (triplets of DNA base pairs that code for individual amino acids), then our measure of activity $A(t)$ for this system could be positive. To prevent such false positives, one must count usage only at levels on which units have adaptive significance.

To some extent it is an open question exactly which changes occurring in an evolving biosphere are adaptive; some are the effect merely of random genetic drift or other non-adaptive processes [?, ?] Still, there is no real doubt that a significant proportion of the change *is* adaptive. And even though selection might be taking place on a variety of levels, there is no real doubt that a significant proportion of adaptive change is genetic. Furthermore, the degree of evolutionary activity measured at different levels should correspond at least roughly. Thus, a genetically grounded implementation of our statistical measure of evolutionary activity $A(t)$ should give a good first approximation of a system's overall evolutionary activity.

4.1 Artificial Biospheres

In the strategic bug model the usage of a bug's genes increases every time the bug "uses" the gene, by instinctively following an entry in a strategy look-up table. Genes coding for any type of activity could be subject to the same kind of usage bookkeeping. This makes it straightforward to measure evolutionary activity in many other model biospheres.

For example, consider the tropic bug model outlined in section ???. Tropic bugs have only two genes, one controlling the number of offspring produced during reproduction, g_{off} , the other controlling the food threshold required for reproduction, g_{th} . Whereas a gene in a strategic bug's strategy look-up table is used only if and when the bug detects that gene's input condition, a tropic bug uses both of its genes each time it reproduces. Thus, if tropic gene usage is counted as for strategic genes, the "usage" of a tropic bug's gene would reflect simply the longevity of that bug's lineage. Thus, activity waves would correspond to the persistence of bug species, and the slope of a wave would reflect the average rate at which bugs in those species reproduce.

Evolution among tropic bugs is reflected by changes in the time-dependent population distribution function over the two-dimensional space of genes, $P^t(g_{off}, g_{th})$, which can be identified as an *a posteriori* fitness function. Simulations of this model show that the size of fluctuations of available food strongly determines the system's evolutionary dynamics [?], with the dynamics of $P^t(g_{off}, g_{th})$ reflecting the evolutionary development of the system. If there is a low level of evolutionary activity then $P^t(g_{off}, g_{th})$ goes in time to a fixed distribution showing that the bugs all fall into one broad cluster of species. This occurs typically when fluctuations in the food supply are small. In this case, the usage distribution function $N(t, u)$ would be dominated by one long-lived cluster of activity waves, overshadowing a carpet of short-lived waves reflecting new mutations that all quickly become extinct. On the other hand, if there is a high level of evolutionary activity, $P^t(g_{off}, g_{th})$ in time develops disjoint peaks that move about and eventually collapse. This phenomenon occurs typically when large quantities of food are put into the environment relatively infrequently. Plentiful food apparently causes a rapid proliferation and variation of organisms near a particular genome, followed by dying out of large fractions of the population, followed by another rapid proliferation, and so on. In this case, when distinct subsidiary species continually split off from the main population, $N(t, u)$ would show the back-

ground cluster of waves overlying a continual stream of relatively long-lasting secondary waves.

In general, usage statistics can be gathered for any computationally implemented model that contains explicit micro-level rules governing the behavior and structure of organisms in the population. This is true even when the micro-level rules are computationally more complex than the look up tables of the strategic bugs. Thus, the style of usage bookkeeping implemented for the strategic bugs can quantify evolutionary activity in virtually any artificial life model biosphere.

4.2 Natural Biospheres

If gene usage is defined as for strategic bugs, our measure of evolutionary activity $A(t)$ is well defined for actual biological populations, at least in principle. In practice, however, it is virtually impossible to get data about persistent gene usage from actual biological organisms. This makes it virtually impossible even to count gene usage in natural biospheres. Furthermore, in many actual biological populations it is difficult to tell exactly which strands of DNA have adaptive significance.

However, there are other levels at which one could count usage in actual biological populations. For example, one could choose *species* as the microscopic functional units. Then, general patterns in the evolutionary activity of actual biological populations could be rendered quantifiable using data from the fossil record or from living populations in the field or the laboratory.

When we counted usage for the strategic bugs, usage of a bug's gene was incremented every time the bug encountered a strategy gene's input condition and acted as prescribed in its output condition. A species does not correspond to any single given behavioral input-output rule; it is a rough "clump" of roughly "similar" complete strategies (with rules covering all possible input conditions). So, a new definition of the "usage" of a species is needed. The simplest approach would be to increment the "usage" of a given species for as long as it persists in the ecosystem, weighted by the proportion of individuals in the ecosystem that belong to that species. If $s_i(t)$ were the proportion of the organisms that were members of the i^{th} species at time t , the usage $u_i(t)$ of that species at that time would be the accumulated

proportion of the population constituted by the species:

$$u_i(t) = \int_0^t s_i(t) dt.$$

A quite long-lived species, like the shark, would show up as a long activity wave. The birth of a new species would generate a new activity wave which would terminate with the species' eventual extinction. The continual production of new activity waves would signal the continual generation of new species.

With usage defined at the level of species, it should be possible to obtain at least a qualitative picture of evolutionary activity from the fossil record. The evolutionary activity based on this measure of usage should swing dramatically at pivotal points in the evolution of the biosphere. The Cambrian explosion, for example, during which nearly all major groups of invertebrates with hard parts originated, would be seen as a period with large positive evolutionary activity.

It is instructive to compare our usage statistics with the gene frequency statistics standardly gathered by molecular geneticists [?]. There are similarities but also important differences; for the moment, focus on the usage of genes. A gene's *frequency* is defined as that proportion of the population that possesses the gene; if the whole population has a gene, it is said to be "fixed." If gene usage is incremented each time a gene is *used* (as in the strategic bugs), then usage reflects more than the mere *existence* reflected by a gene's frequency. Even if usage is defined by reference to a gene's persistence (as with species above), a gene's usage differs from its frequency; for in this case usage is a gene's *integrated* frequency. Thus, the molecular geneticists' data showing the route to fixation of one or two genes is qualitatively different from our data showing large clusters of persistent genes emerging out of the sea of all possible genes.

4.3 Chemical Systems

In some evolutionary models, the microscopic constituents are so simple that there is no genetic code, for example chemical soups [?, ?, ?] (and perhaps also populations of complex clay crystallites [?, ?]). Our measure of evolutionary activity could still be applied in these models, if "usage" were taken

to be simply the integrated concentration of the chemical species, on analogy with the equation above for biological species. The definition of usage could also include a weighting for the number of reactions in which a given chemical species participates; thus, if $R_i(t)$ were the number of reactions for the i^{th} chemical species at time t , and $c_i(t)$ were the concentration of the i^{th} chemical species at time t , then the usage $u_i(t)$ of that species would be:

$$u_i(t) = \int_0^t R_i(t)c_i(t)dt.$$

In the autocatalytic soup, positive evolutionary activity would correspond to the ongoing production and persistence of novel forms of chemicals. Again, as in the genetic case, the intricate nature of the interactions are not completely captured by our measure of activity, but could be reflected in activity patterns.

Chemical soups can represent not only interactions between polypeptide strings and RNA strands; they can also model interactions between antibodies and antigens in an immune system. Different kinds of antibodies may be regarded as different chemical species, with their usage calculated in the same way. The evolutionary time scale of immune systems is short enough to be readily observable. Thus, there is reasonable hope of actually measuring this type of evolutionary activity in real immune systems.

4.4 Computational Systems

Populations of information processing units are an especially interesting and powerful setting for evolution, one with extensive practical implications, in fact. The strategic bugs are one example of a population of simple computational units. Other types of computational populations, however, need not be embedded in a biologically motivated setting.

One computational model within which evolutionary activity can be easily measured is Fontana's Turing gas [?]. This is a model of interacting strings, similar to the interacting chemical strings mentioned above, except that here the strings are information processing elements, functions in a lisp dialect called AlChemY (for Algorithmic Chemistry). The Turing Gas model is extremely active because the micro-units interact in a way that is essentially computational. Interaction happens repeatedly between pairs of lisp functions chosen at random. It occurs as one function is evaluated with

another as its argument. Interactions can produce null programs (in which case the interaction is termed “elastic”) or a new program (in which case the interaction is termed “reactive”). Usage counters are attached to each function present in the population, and incremented with each reactive collision. These measurements will be reported in future work.

Another kind of computational evolution is produced by genetic learning algorithms [?], which operate on a population of “hypotheses” described by a set of genes. These hypotheses are assigned a fitness, which is typically their suitability for solving some problem. The genetic algorithm changes the population of hypotheses through a process of survival and genetically modified reproduction of the fittest. Usage bookkeeping can be easily implemented in this setting, often with usage counters for each genetic unit, just as for the strategic bug model. For learning algorithms with more machinery, it may be necessary to include other aspects in usage count. For example, Holland’s classifier system [?] associates with each classifier a strength variable, analogous to the concentration of a species or chemical. A classifier’s usage could be defined as its integrated strength, on analogy with the equations above for biological and chemical species. In this setting, one would expect to observe a flurry of evolutionary activity initially, as a variety of new hypotheses are tested, with activity dying down once optimal hypotheses have been identified.

A more abstract illustration of evolutionary activity is found in the complex patterns produced by the temporal evolution of cellular automata. A cellular automaton is a population of automata filling the sites in a lattice. The dynamics of the automaton maps a configuration of symbols over the lattice to another configuration, using a local rule applied simultaneously at all sites. The local rule followed by a cellular automaton site is analogous to a strategic bug’s behavioral strategy but, whereas the bugs’ strategies vary between individuals and change over generations, the rule in a cellular automaton is the same for all sites and never changes. If the number of symbols is finite, and the local neighborhood is finite, a cellular automaton’s local map may be thought of as a look-up table whose inputs are site values over a neighborhood, and whose outputs are the value of a particular site at the following time step. Usage counters can be attached to each entry in this look-up table.

Langton has seen that for certain classes of cellular automata with complex dynamics, the number of neighborhood configurations visited slowly

grows with time [?]. As new local configurations are encountered, new usage activity would be stimulated and activity waves would form. If the set of local configurations visited stabilized to a fixed set, no new waves would be produced. Although this must eventually happen for any cellular automaton rule, it might take a very long time compared to the iteration of the rule. Langton's result thus indicates that cellular automata with complex dynamics might exhibit a primitive form of evolutionary activity that occurs in the absence of any "genetic" variability.

4.5 Mental Systems

Generating and repeatedly entertaining new ideas seems to be one of the hallmarks of an active mind. One can view an individual mind, whether real or simulated, as a macroscopic system, with individual ideas as its microscopic elements. Evolutionary activity could be measured by attaching usage counters to ideas. Although difficult for a real mind, in simulated minds it is usually quite easy to keep track of "ideas." Even the activity of real minds might even be able to be inferred from their products, at least in narrow realms. For instance, in the realm of mathematics, a crude measurement of evolutionary activity could be implemented by assigning usage counters to theorems, incrementing them each time a theorem is used to prove another theorem.

Moving beyond individual minds, evolutionary activity could occur in a communicating community of mental agents, such as the modern scientific community. Again, the pattern of mental activity could be inferred from patterns in its products. For example, evolutionary activity could be calculated from the Science Citation Index, with usage straightforwardly defined as the number of citations accumulated by a given article. An interesting feature of the population of "interacting" articles is that no member of the population ever "dies." The bulk of published articles would generate activity waves that move only a very short time; articles referred to quite often would generate high velocity activity waves, and articles referred to infrequently but over a long time period would generate low velocity activity waves. Concentrated periods of fruitful scientific activity would initiate many long activity waves. The flattening of most existing waves coinciding with the beginning of a welter of new waves would signal a scientific "revolution" [?].

Evolutionary activity could be applied on social and cultural levels, as

well, if only one could identify the elementary units of cultural transmission that Dawkins referred to as *memes*, on analogy to the genes that encode our physiological structure [?]. If memes could be identified, evolutionary activity in a whole culture could be measured with the same bookkeeping that we have used for genes.

5 Teleology in Evolving Systems

In this section we argue that evolutionary activity may be interpreted as a measure of the extent to which an evolving system’s behavior is teleological, i.e., goal-directed or purposeful. After briefly describing what teleology is, we argue that it is related to evolutionary activity in the strategic bug model. Any attempt to revive teleology in biology runs the risk of provoking controversy and criticism. As J. B. S. Haldane once quipped, “Teleology is like a mistress to a biologist: he cannot live without her but he is unwilling to be seen with her in public.” Thus, we make sure to explain why the teleology in activity waves is perfectly respectable.

5.1 The Nature of Teleology

Behavior that can be explained by reference to the utility of its effects we will call *teleological* (telic, goal-directed, purposive, for the sake of some end), and a *telic explanation* will explain something by reference to its beneficial effects [?, ?, ?]. Our’s is not the only approach to teleology; comparisons with the three most attractive alternatives are detailed elsewhere [?, ?].

In ordinary parlance, telic explanations are offered for a wide variety of things, such as the behavior and structure of biological organisms and their parts, the actions of conscious human agents, and the structure and behavior of artifacts designed and used by people. All of these can be given some variation of telic explanation [?]. In each case, an essential part of the explanation is a beneficial effect brought about by the thing being explained.

The benefit promoted by any form of telic activity can be identified with the activity’s purpose or goal. In the case of a human being consciously trying to produce some specific beneficial effect, the telic agent is consciously and explicitly aware of the goal. But some goal-directed activity is directed to goals that are not entertained consciously or explicitly; in these cases, the

“goal” is simply a beneficial effect that explains the activity [?].

Functionality is sometimes confused with teleology; the two are related, but they must be distinguished. Functional behavior is merely any behavior that is beneficial, that “serves a purpose,” regardless of its cause. Telic or goal-directed behavior, on the other hand, is not merely functional or beneficial; it does not merely serve a purpose. It must occur specifically *because* it is beneficial, *because* it serves a purpose. Telic behavior cannot occur merely accidentally or for some reason wholly unconnected with its utility.

For a given organism at a given time in a given local environment, there is a range of possible behaviors that would be more or less functional (beneficial). The “temporary local optimization criteria” set for a species by the evolutionary dynamics (recall the discussion in section ??) are the criteria for its maximal local functionality. If an organism contains a favorable new mutation, the new behavior caused by the mutation might immediately be functional. But that behavior will not be telic until its utility becomes a causal factor in its continual production. This can happen if the behavior persists through a lineage *because* of its utility.

5.2 Telic Activity Waves

The presence of activity waves in the strategic bug model reflects the occurrence of this sort of telic behavior. The usefulness of a gene is tested when and only when it is used. Unused genes exact no tax, so their “persistence” means little. But a well-tested gene persists in the gene pool only if the gene makes a significant contribution to the welfare of those organisms containing it. So, the presence of an activity wave shows not only that a significant number of genes are useful; it shows that a significant number of genes are present in the gene pool because they are continually verifying their usefulness. That is, these genes, and the behaviors they encode, persist because the behaviors are continually performed and continually benefit the organisms exhibiting them. Thus, activity waves reflect teleological behavior, and the continual production of new activity waves reflects the continual emergence of new teleological behavior. In this context, then, it is appropriate to speak of *telic activity waves*.

It might seem that not all genes that contribute to a telic activity wave need be beneficial. After all, an organism might use a harmful gene a number of times and still pass it on to offspring, provided the organism possessed

enough *other* genuinely beneficial genes to outweigh the harm produced by the use of the harmful gene. Thus, a harmful gene could acquire some positive usage. However, harmful genes make no significant contribution to telic activity waves. For one thing, a harmful gene is unlikely to persist long enough to contribute to net persistence $P(t)$. But more importantly, since activity waves occur on much longer time scales than generations, the continual persistence of well-tested genes cannot be attributed to happenstance. Telic activity waves reflect genes that persist because they contribute to strategies of proven usefulness. They are not merely useful; they persist *because* they are useful. Thus, the behaviors in these well-tested strategies are teleological (goal-directed), not merely functional.

The behavior of a strategic bug is genetically hardwired, “instinctive.” Its strategy allows for no flexibility of response; unable to deliberate consciously about what course of action to take, the bug cannot “freely choose” its actions. In a given local environment it has one and only one “option” for what to do: the behavioral output coded by the gene with that local environment as input condition. Nevertheless, its behavior can still be genuinely telic, goal-directed, or purposive, because it can be persisting due to its usefulness. Not every action produced by every gene in a bug’s behavioral strategy is telic; it depends on why the bug has the gene. If a particular gene is present in the bug’s genome because that gene has produced behavior that was beneficial for the bug’s ancestors, *then* the (instinctive) behavior produced by that gene will be telic. Instinctive behavior of this sort is the simplest kind of genuinely telic behavior. More than merely functional, it is a limiting case of teleology, located on the telic spectrum at the opposite end from behavior produced by open-ended conscious deliberation.

Are all evolutionary activity waves telic? Our argument that activity waves in the strategic bug model are telic depends on the premise that a persistent gene is valuable, in the sense that it benefits the organism by enabling it to gather more food. The activity waves that might occur in other systems discussed above do not necessarily have a similarly unambiguous value-based interpretation. An evolutionary activity wave is telic only if there is a value in the persistence reflected in the activity wave. At this point, there is no theory of value comprehensive enough to include all the systems discussed above, though for many of the systems a value-based interpretation is intuitively clear. For such systems, the activity waves are telic; for others, the question remains open.

5.3 Worries about Biological Teleology

Today, any form of teleology in biology tends to be viewed with suspicion and dismissed. The controversy stems partly from the differences among the many kinds of teleology. Technical terms like “teleonomic” have been introduced in the attempt to evade the controversy [?, ?]; ironically, these neologisms are used so divergently that they just add to the confusion. Whether it is fair to criticize biological teleology depends on the specific kind of teleology involved. We believe that the teleology reflected by telic activity waves is no cause for embarrassment; at least, it does not revive any of a quartet of familiar objections.

Anthropomorphism. One complaint is that teleology in biology anthropomorphizes nature (e.g., [?, ?]). This complaint takes two forms: Teleology might require either that each biological creature possesses sophisticated mental capacities analogous to those possessed by a person (mentalism), or that the diversity of well-adapted creatures is the result of the activities of a mental deity (the Designer supported by the notorious argument from design).

Worries about the argument from design would clearly be misdirected at telic activity waves. The teleology in telic activity waves presupposes no deity directing things behind the scenes. The more general worry that all teleology at bottom is mentalistic can also be deflected once it is realized that biological goals can be non-conscious. In the strategic bug model whether an organism remains alive is determined by whether it continues to find food, but the organisms are not “aware” that finding food benefits them. Thus, finding food should not be considered to be an organism’s *conscious* goal. Nevertheless, since finding food is what in fact determines whether an organism survives, it can be considered to be an organism’s *non-conscious* goal. As explained above, there is no requirement that all teleology be mentalistic and involve conscious goals; a non-conscious goal can be sufficient for teleology provided that it causes behavior that realizes the goal. In particular, the teleology in telic activity waves involves no mentalism. The behavior encoded by well-tested genes is teleological because it can be explained by its good consequences, but those good consequences are merely non-conscious goals.

The model could be enhanced in such a way that it could give rise to mental teleology. Organisms would need a more complex information processing

mechanism with the capacity to have explicit goals for which sub-goals can be formed in response to environmental contingencies, and with the capacity to gather information about what sub-goals are feasible in the current local environment. Whereas survival is measured on a time scale spanning generations, psychological value would be measured on the time scale within a single lifetime. So, whereas biological teleology takes place through a lineage existing over many generations, mental teleology would take place within one lifetime.

Predetermined Goals. Another reason for the jaundiced attitude towards teleology in biology is that teleology is thought to require that the evolutionary process itself has a predetermined goal (e.g., [?, ?]). Specifically, even if there is no Master Designer, teleology still must involve a Master Plan, a specific set of predetermined specifications for each species. The worry is *not* that the development of our model of biological teleology would show an average statistical trend (e.g., continual evolutionary activity); on the contrary, this would be desired in a model, since the actual biological world apparently exhibits the same statistical trends. The worry is that the genotype and phenotype of the organisms produced by the evolutionary process are determined *a priori*, independently of local environmental and ecological contingencies.

However, the sort of teleology signaled by telic activity waves involves no predetermined goals. Rather, the form to which organisms will evolve is determined by whatever happens to be sufficiently beneficial in the continually changing local biological context. (Recall the discussion of intrinsic adaptation in section ??.) Rather than being specified in advance, the organisms' forms depend on random, non-teleological genetic changes and the contingencies of the struggle for survival. So, the teleology in telic activity waves emerges *a posteriori*.

Future Causation. Teleology in biology is sometimes thought to require that events in the future (the realization of goals) have causal efficacy over present behaviors, a sort of "future causation" that seems patently absurd (e.g., [?]). Indeed, our view of teleology might appear to involve future causation, but this appearance evaporates under scrutiny.

Biological events can be viewed on either a micro or a macro level. The micro perspective involves events on a time scale within a generation. From this perspective, an individual organism's behavior has a telic explanation when it occurs because *in the past* the same kind of behavior helped the

organism or the organism’s ancestors to flourish. A behavior’s past beneficial effects cause the same kind of behavior in the present, which might well cause further manifestations of the same kind of behavior. The causation in this explanation is of the ordinary kind—past events causing present events. By contrast, the macro perspective involves a time scale spanning many generations. From a vantage abstracted from individual generations, one can say simply that a teleological behavior occurs because on average that kind of behavior promotes (in a tenseless sense) the survival of those organisms that exhibit the behavior. This kind of merely apparent future causation at the macro level is harmless because it is underwritten by ordinary causation at the micro level.

Value and Subjectivity. Our interpretation of teleology requires that good effects are causally efficacious. Some might judge that this reference to an effect’s goodness or value is inescapably subjective, possibly on the grounds that all value judgments are inherently subjective. However, we believe that an objective criterion of an organism’s welfare is its ability to survive and reproduce. In the strategic bug model, an organism’s welfare consists of no more and no less than this.

We ignore other possible components of an organism’s welfare, not because we believe that there could be none, but merely to simplify our model. A more complicated model could incorporate benefits that are unconnected with a creature’s survival, such as pleasures, the satisfaction of desires, and other “psychological” goods. In the strategic bug model, however, all telic phenomena are shaped by their value simply for survival and reproduction.

6 Vitality as a Test for Life

What is life? How can it be recognized? In an everyday context these questions seem tantalizingly clear—a cat is alive and a rock is not. But formalizing this distinction is difficult, especially if the formalization is to be used in empirical measurements.

Life is usually thought of as a property of individual organisms. We propose to make a gestalt switch and view life from a more global, statistical perspective. No single molecule of gas has a macroscopic property like temperature; temperature is meaningful only for large populations of molecules. Similarly, no single organism exhibits indefinitely ongoing life; in the long run

not even a lineage remains alive. Individual life is “here today, gone tomorrow” and, in fact, intuitively this transitory nature is one of its characteristic features. From a global perspective, only the complex web of interacting organisms—the entire biosphere—remains “alive” in the long run, through the continual cycle of birth and death of individual organisms. So, rather than try to define what it is for an individual “microscopic” organism to be alive, our concern is with what it is for a “macroscopic” system (population of organisms) to exhibit the property of indefinitely ongoing life.

Evolutionary activity is an especially salient global property of populations of living organisms; it seems that they cannot help but evolve, at least in the long run. In addition, minds and other non-biological systems with lots of evolutionary activity exhibit a kind of “liveliness.” Thus, we will say that a system is *vital* if and only if it has positive evolutionary activity.

We believe it is fruitful, theoretically and experimentally, to link the notions of an individual’s life to the vitality of the global system in which the individual lives. In fact, this link is already suggested by the common claim that an organism is alive only if it is a member of an actively evolving biosphere [?, ?]. Our measure of evolutionary activity sharpens this claim into the following *life-vitality hypothesis*: Vitality—positive evolutionary activity $A(t)$ —is a necessary condition of systems containing living individuals, and the measure of the vitality of a system of living individuals is the rate at which new telic activity waves are generated. Note that in the hypothesis vitality is only *necessary* for a system to contain living elements; it is not sufficient since the ideas entertained by a vital mind, for example, are not alive.

It is important to recognize what the life-vitality hypothesis is *not*. It is a contingent fact that the biosphere is the product of evolution; life might not have been linked to evolution in the way that it is. Organisms *could* have been designed and created by an omnipotent deity, and there *could* have been eternal, non-evolving forms of life, such as angels. Similarly, medical technology *could* improve to the point that individual organisms remain alive indefinitely and thus never evolve. These fanciful possibilities show that the life-vitality hypothesis is not a conceptual necessity; rather, it is the sort of contingent empirical claim that is characteristic of hypotheses in the natural sciences.

Certain facts about the biosphere might seem to contradict the life-vitality hypothesis. Infertile individual organisms live and die without affecting the

evolutionary activity of their population. Furthermore, the individuals in a lineage or a sub-population are all alive, but certain lineages such as the shark have persisted without evolving appreciably for a quite long time, and certain sub-population of infertile organisms such as mules are simply incapable of evolving. But these admissions do not contravene the vital systems hypothesis, for infertile individuals and infertile sub-populations are always transitory members of a global biosphere which is certainly evolving in the long run. You can't get a mule or a shark except from a vital biosphere.

The possibility of an ecology that has reached a stable “climax” state and stopped evolving challenges the life-vitality hypothesis more directly. Although the existing genetic combinations in a climax biosphere would create continually propagating activity waves, the gene pool would no longer be absorbing new genetic variations. Thus, a climax biosphere would produce no new activity waves and its evolutionary activity $A(t)$ would consequently drop to zero. In this straightforward sense a static climax biosphere is “less vital” than one that is continuously evolving. Yet the organisms in the biosphere would still be merrily living. Thus, for climax biospheres the vital systems hypothesis breaks down. However, the extent of this breakdown vanishes in the long run, it seems, for real biospheres apparently do not remain indefinitely in a state of climax. On the contrary, in the long run biospheres seem to continue to evolve.

The field of artificial life is searching for a definition of life; even better would be a criterion of life—a public, empirical, repeatable, quantifiable test for whether a system (possibly artificial) is alive. The analogous situation in artificial intelligence is how to tell whether a system is thinking. Forty years ago Alan Turing proposed a criterion for thinking—the famous Turing test [?]. What the field of artificial life needs is an analogue of the Turing test.

Although public, empirical, repeatable, and quantifiable, the Turing test has two limitations: it detects thinking only *indirectly*, through its effects on overt behavior, and it evaluates this behavior through the *subjective* opinions of a panel of human jurors. It would be preferable to have tests that directly and objectively measure whether (and to what extent) a machine, or a model being implemented on a machine, exhibits intelligence or life.

Our quantity $A(t)$ is a direct and objective measure of a system's vitality or evolutionary activity. Thus, by invoking the life-vitality hypothesis, $A(t)$ provides a direct and objective measure of the degree to which a system exhibits life, yielding a Turing test for life.

As we noted in section ??, a biosphere’s vitality also is a measure of its teleology. This leads to another link between vitality and life. Purposeful behavior has often been cited as an especially characteristic sign of individual life [?, ?, ?, ?, ?]. But this proposal has been unhelpful until now, since no way to quantify and measure purposefulness was known. If this teleology-life relationship is recast at a global level, as we have done for the link between evolution and life, then telic activity waves allow teleology to be quantified. Purposeful behavior is signaled by telic activity waves, and the level of vitality measures the degree to which new purposeful behavior is continually emerging. Thus, vitality at one fell-swoop quantifies two intuitive signs of life: evolutionary activity and purposeful behavior.

Measuring vitality in models that include psychological factors could also provide a direct and objective substitute for the original Turing test itself. If the mind is viewed as a global, statistical system, positive activity would indicate a mind’s continual incorporation of new behavioral or psychological patterns of activity (*cf* section ??). Statistics such as our measure of evolutionary activity might provide a method for quantifying the purposefulness common to both life and mind.

Acknowledgements

We would like to acknowledge the hospitality of the Santa Fe Institute and the Complex Systems Group of the Theoretical Division at the Los Alamos National Laboratory, where this work began. We appreciate many helpful comments on the manuscript from C. Langton, J. Moor, J. Page, C. Voeller, and an anonymous reviewer, and helpful conversations with W. Fontana, J. Holland, S. Kauffman, S. Lloyd, T. Meyer, and P. Schuster. NP also thanks the “Complexity and Evolution” program at the Institute for Scientific Interchange, where part of this work was completed. This work was partially supported by National Science Foundation Grant number Phy86-58062.

References

- [1] F. J. Ayala, “Teleological Explanations in Evolutionary Biology,” *Philosophy of Science* **37** (1970): 1–15.

- [2] M. A. Bedau, “Against Mentalism in Teleology,” *American Philosophical Quarterly* **27** (1990): 61–70.
- [3] M. A. Bedau, “Where’s the Good in Teleology?” *Philosophy and Phenomenological Research*, forthcoming.
- [4] M. A. Bedau, “Naturalism and Teleology,” in S. Wagner and R. Warner, eds., *Beyond Naturalism and Physicalism*, forthcoming.
- [5] A. G. Cairns-Smith, *Genetic Takeover and the Mineral Origins of Life* (Cambridge University Press, 1982).
- [6] A. G. Cairns-Smith, *Seven Clues to the Origin of Life* (Cambridge University Press, 1985).
- [7] R. Dawkins, *The Selfish Gene* (Oxford University Press, 1976).
- [8] R. Dawkins, *The Blind Watchmaker: Why the evidence of evolution reveals a universe without design* (Norton, 1986).
- [9] J. Dupré, ed., *The Latest on the Best: Essays on Evolution and Optimality* (The MIT Press, 1987).
- [10] T. G. Dobzhansky, *The Biology of Ultimate Concern* (New American Library, 1967).
- [11] J. D. Farmer, S. A. Kauffman, and N. H. Packard, “Autocatalytic Replication of Polymers,” *Physica* **22D** (1986): 50–67.
- [12] J. D. Farmer, S. A. Kauffman, N. H. Packard, and A. S. Perelson, “Adaptive Dynamic Networks as Models for the Immune System and Autocatalytic Sets,” *Annals of the New York Academy of Science* **504** (1987): 118–130.
- [13] J. D. Farmer, N. H. Packard, and A. S. Perelson, “The Immune System, Adaptation, and Machine Learning,” *Physica* **22D** (1986): 187–204.
- [14] W. Fontana, “Algorithmic Chemistry: A Model for Functional Self-Organization,” this volume.
- [15] S. J. Gould, *Ever Since Darwin* (Norton, 1977).

- [16] S. J. Gould and R. C. Lewontin, “The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme,” *Proceedings of the Royal Society of London* **B 205** (1979): 581-598.
- [17] J. H. Holland, *Adaptation in Natural and Artificial Systems* (University of Michigan Press, 1975).
- [18] J. H. Holland, “Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems,” in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine Learning, An Artificial Intelligence Approach*, Vol. II (Morgan Kaufmann, 1986), pp. 593-623.
- [19] B. A. Huberman, ed., *The Ecology of Computation* (North-Holland, 1988).
- [20] S. A. Kauffman, “Autocatalytic Sets of Proteins,” *Journal of Theoretical Biology* **119** (1986): 1-24.
- [21] M. Kimura, *The Neutral Theory of Molecular Evolution* (Cambridge University Press, 1983).
- [22] C. G. Langton, *Computation at the Edge of Chaos: phase transitions and emergent computation*, Ph.D. dissertation, EECS Department, University of Michigan, 1990.
- [23] T. S. Kuhn, *The Structure of Scientific Revolutions*, Second edition (The University of Chicago Press, 1970).
- [24] R. Lewontin, “The Units of Selection,” *Annual Review of Ecology and Systematics* **1** (1970): 1-14.
- [25] J. Maynard Smith, *The Theory of Evolution*, Third edition (Penguin, 1975).
- [26] E. Mayr, “Teleological and Teleonomic, A New Analysis,” in R. S. Cohen and M. W. Wartofsky, eds., *Methodological and Historical Essays in the Natural and Social Sciences* (Reidel, 1974), pp. 91-117.
- [27] E. Mayr, *The Growth of Biological Thought* (Harvard University Press, 1982).

- [28] P. B. Medwar and J. S. Medwar, *Aristotle to Zoos* (Harvard University Press, 1983).
- [29] J. Monod, *Chance and Necessity*, translated by A. Wainhouse (Vintage, 1972).
- [30] E. Nagel, "Teleology Revisited," *The Journal of Philosophy* **74** (1977): 261–301.
- [31] M. H. Nitecki, ed., *Evolutionary Progress* (The University of Chicago Press, 1988).
- [32] N. H. Packard, "Intrinsic Adaptation in a Simple Model for Evolution," in C. Langton, ed., *Artificial Life* (Addison-Wesley, 1988), pp. 141–155.
- [33] O. Rössler, "Chemical Automata in Homeogeneous and Reaction-Diffusion Kinetics," *Lecture Notes in Biomathematics* **4** (1974): 399–418.
- [34] O. Rössler, "Deductive Prebiology," in K. Matsuno, K. Dose, K. Harada, and D. L. Rohlring, eds., *Molecular Evolution and the Prebiological Paradigm* (Plenum, 1984), pp. 375–385.
- [35] E. Sober, ed., *Conceptual Issues in Evolutionary Biology, An Anthology* (The MIT Press, 1984).
- [36] E. Sober, *The Nature of Selection* (The MIT Press, 1984).
- [37] A. M. Turing, "Can a Machine Think?" *Mind* **59** (1951): 433–460.
- [38] A. Woodfield, *Teleology* (Cambridge University Press, 1976).