Darwinian demons, evolutionary complexity, and information maximization

David C. Krakauer
Santa Fe Institute, New Mexico 87501, USA

(Received 13 July 2011; accepted 4 September 2011; published online 30 September 2011)

Natural selection is shown to be an extended instance of a Maxwell's demon device. A demonic selection principle is introduced that states that organisms cannot exceed the complexity of their selective environment. Thermodynamic constraints on error repair impose a fundamental limit to the rate that information can be transferred from the environment (via the selective demon) to the genome. Evolved mechanisms of learning and inference can overcome this limitation, but remain subject to the same fundamental constraint, such that plastic behaviors cannot exceed the complexity of reward signals. A natural measure of evolutionary complexity is provided by mutual information, and niche construction activity—the organismal contribution to the construction of selection pressures—might in principle lead to its increase, bounded by thermodynamic free energy required for error correction. © *2011 American Institute of Physics*. [doi:10.1063/1.3643064]

Over the 4.6 billion year history of the earth, an incredibly diverse range of species have evolved capable of exploiting an equally diverse range of environments. A small subset of these species, emerging around 1 billion years ago, are the multicellular eukaryotes, ancestors of the animal lineages. These species have evolved networks of cells capable of extracting large amounts of information from their environments for inferring sources of free energy and predicting the behavior of competitors. I show that the evolution of inferential mechanisms can be captured using a complexity measure and seek to explain how informational bottlenecks at the genetic level are overcome by evolving mechanisms of learning and plasticity. I present a fundamental principle for all adaptive systems, the demonic selection principle, which states that an organism or its strategies cannot exceed the complexity of their selective environment. In this framework, genomes and brains are both information gathering devices for inferring the future from regular or predictable patterns in the past.

I. THE COMMON STRUCTURE OF COMPLEX SYSTEMS

Is it possible, or even desirable, to understand the history of life on earth in relation to the emergence across key biological lineages of increasing complexity? Darwin had much to tell us about diversity, but far less about complexity. Many prominent scientists consider the question of increasing evolutionary complexity nonscientific and at odds with the empirical distribution of life on earth.[17] Multicellular organisms are vastly outnumbered by microbes. I shall pursue a slightly different question, seeking to identify those conditions when organisms are selected to increase the amount of information they encode about their environments and other organisms, and the diversity of mechanisms that have evolved to accomplish these goals (see Refs. 20 and 9 for related projects of comparable scope). I shall focus on evolutionary dynamics, information, and mechanisms of learning, placing the simple adaptive capabilities of the genome and the flexible representational power of brains, on a complexity-continuum defined in terms of memory and action. The challenge will be to explain the key innovations that have led to significant increases in these capacities in a few lineages through time and the continued success of those lineages that have not evolved powerful inferential mechanisms. Success in answering this and related questions would amount to illuminating one of the great mysteries of life, the origin of complex information processing architectures of which the human brain is one outstanding example.

Research on the evolution of complexity typically proceeds by picking functional traits and plotting a time series to reveal systematic trends across lineages, for example, increases in cell size and diversity,[5,42] increases in genome size[32], changes in tissue form and pigmentation,[27] changes in shell morphology,[35] in the number of discrete taxonomic features[36] or the fewest number of phenotypic dimensions that capture variation across lineages.[33] There has been a parallel, neutral approach, to the evolution of complexity that views increasing numbers of components (genes or traits) as a statistical feature of basic genetic mechanisms of mutation and duplication,[8,26] or increasing genetic interactions through epistasis.[6,13] In both cases, there is awareness that complexity is often self-limiting, and that as the number of components increase, so do constraints on diversification.[31,44] And that whereas lineages can never fall below a minimum complexity to sustain life (minimal life—Refs. 28 and 25), there is no known upper bound to complexity, and so the maximum could be increasing simply through incremental sampling of possible variation over long stretches of time without any systematic drivers of complexity.

We are currently experiencing a comparative data-deluge (Refs. 2 and 3) with unprecedented amounts of genetic sequence, expression, and molecular proteomic data sampled from a very diverse set of biological lineages. As of 2010 around 4600 genomes, chromosomes and plasmids

have been sequenced. These sequences are deposited in publically available databases that include representatives from all of the five major kingdoms of life and the viruses. In GenBank alone, the NIH-supported genetic sequence database, there are currently 136 000 000 sequence records, spanning 380 000 organisms. Of perhaps, greater relevance to this project are databases such as the Kyoto Encyclopedia of Genes and Genomes (KEGG, Ref. 21), a reference knowledge base for linking genomes to function through t he process of PATHWAY mapping, whereby genomic or transcriptomic content is mapped to reference pathways in order to infer systemic behaviors of the cell or the organism. And the Gene Expression Database (GXD, Ref. 19) which stores information from a variety of gene expression assays. What this means for us is the possibility of taking a large number of different quantitative measurements of species, many of which can be placed in an ancestor-descendant relationship, and thereby track through time systematic variation in gene content, expression patterns, and critical sets of shared functional capabilities bearing on information acquisition, storage, and inference.

A. Outline

Here, I review and build on recent progress in the study of evolutionary dynamics,[29] information theory and inference,[12,18,38] evolutionary complexity,[1] the theory of reinforcement learning,[46,47] the study of entropy as it relates to information in the framework of Maxwell's demon,[23] and the thermodynamics of scaling.[45] It is striking that it is only recently that we have come to appreciate that many of these frameworks seek to answer a very similar question: how starting from simple initial conditions, or a position of maximum ignorance, a sequence of carefully chosen steps can increase the information that a system encodes using a limited source of free energy. This problem has important roots in classical problems of undecidability in mathematics formalized in terms of string rewriting systems where one wants to determine whether one can sort among a set of inputs to yield a desired output.[34] This problem has a correspondence with a fundamental problem in statistical physics describing the increase in entropy over time and the concomitant loss of information in a coarse grained closed system. This paradox is typically resolved in the framework of Maxwells demon (Maxwell 1871), where one seeks an automaton capable of producing order locally from disordered states. These (rewrite systems and demons) are in turn analogous to natural selection mechanisms that sort among alternative genomes selecting out adaptive variants (for a biological review of selection mechanisms see Ref. 4). I shall consider selective demons and learning mechanisms that lead to systematic increases in the information of a system capable of predicting states of the environment. Increasing effective inference from genomes to neurons corresponds to the evolution of nested selection mechanisms (recursive Darwinian demons) operating over finer spatial and temporal scales. All however need respect a fundamental constraint and that is that the environment establishes an upper bound on the stable coding capacity of the genome or brain. The argument that I shall present proceeds by making the following points:

(1) All selection mechanisms assume a similar mechanism of filtering, where the description length or complexity of a selective filter establishes an upper bound on the description length of a target. Maxwell's demon needs to be at least as far from equilibrium as the system that it sorts, and selection pressures need to be as information rich as the genomes that they propagate.
(2) The activity of a Darwinian demon can be shown to maximize genomic information acquisition over a spatially identifiable distribution of free energy resources. The Darwinian demon dissipates energy through morbidity.
(3) Free energy budgets dedicated to error correction (genetic repair) scale with genome size, with smaller genomes expending less energy per unit mass while still propagating genetic information reliably.
(4) A natural, evolutionary complexity measure defined over an ensemble is the mutual information between distributions of genomes and distributions of environments.
(5) Mechanisms of plasticity and learning can overcome the evolutionary constraint whereby information is fixed through differential death, by allowing that information be fixed through the elimination of false hypotheses with associated neural decay.
(6) The demonic selection principle states that all agent complexity must mirror that of the selective or demonic filter. How then can system complexity increase? I suggest that niche construction, the construction of selection pressures by organisms increases distinguishable states of the environment, revealing new sources of free energy for exploitation by mobilizing excess coding capacity.

II. DARWINIAN DEMONS AND THE NATURE OF ADAPTATION

The origins of systematic, quantitative research into complex, ordered, states begins with James Clerk Maxwell in 1871 seeking to explore the limitations of the second law of thermodynamics. The law states that it is impossible to create an inequality of temperature or pressure in a closed system without the expenditure of work, or that entropy tends to increase towards thermodynamic equilibrium. Maxwell proposed a molecular intelligence (Demon) capable of discerning variation in molecular velocities and sorting them according to velocity into two compartments of a vessel isolated from the environment. In this way a non-uniform, out of equilibrium, configuration can be achieved. Maxwell also observed that if there are few molecules in the vessel, then disequilibrium might emerge as a simple fluctuation without a demon. William Thompson (Lord Kelvin) described the demon as capable of operating selectively on individual atoms reversing the natural dissipation of energy captured in average descriptions of populations. Thus, the law is statistical and can be violated, unlike the first law describing the conservation of energy.

Szilard in 1929 (see extensive discussion in Ref. 23) identified dissipation in an intelligent demon with increasing order in the vessel inaugurating a series of deep ideas on the nature of information, the erasure of memory, and the theory of computation, leading to the whole area of the

thermodynamics of computation, and more recently, quantum computation.

When Darwin proposed his theory of natural selection most fully in the origin of species in 1859—he could have had no idea (and most biologists and physicists still have little idea today) that the structure of his theory would closely resembles the demonic thought experiment proposed by Maxwell. For Darwin, natural selection was an environmentally situated filter (more recently called the selective sieve), which allows well-adapted variants to populate the next generation, and poorly adapted variants to be eliminated. Natural selection is a demon (in the Maxwell sense) that possesses sufficient intelligence (in the Kelvin and Szilard sense), to be able to detect, memorize, and act upon variation in one generation, in order to induce an adaptive distribution of genotypes and phenotypes (organisms) in the next generation.

What distinguishes thermodynamic demons from selective demons is primarily the degree of intelligence or inferential capability with which each is charged. Selective demons are temporally, non-stationary filters of adaptive variability and comprised of numerous biotic and abiotic elements in the environment sensitive to continuous variability. Thermodynamic demons are typically capable of only simple binary recognition and action. Second, selective demons operate iteratively at each generational bottleneck and lead to changes in the composition of the selected agents. This significantly expands the scope of the selective filter beyond the physical demons working with unchanging particles. In adaptive systems, particles evolve "agency" acquiring inferential attributes that preempt those of the selective demon. Third, selected agents are capable of modifying properties of the demon (in part because they contribute to the filtering dynamics through frequency dependence) during niche construction. Physical demons are not modified by their decisions. Fourth, evolution builds new selective and adaptive filters/demons that take the place of natural selection, introducing learning and plasticity into biology over short periods of time. These four factors imply that non-equilibrium statistical physics is vastly underspecified to serve as a theory for biological evolution. However, some form of nested, demonic dynamics serves as a powerful platform upon which to build selective and adaptive mechanisms in order to generate a more general theory of selection and adaptation. And we argue it is this theory that is required if we are going to understand the nature of biological complexity and sophisticated forms of cognition.

III. EVOLUTIONARY DYNAMICS

Darwin's explanatory contribution to science was to provide a mechanism for changes in the architecture of organisms by allowing for the differential accumulation of modifications across generations. Hence, differential survival is promoted or demoted through "environmental" feedback (Darwinian demon) acting on population level variation in function. This is illustrated in a very simple case in Figure 1 as a directed graph, where traits j of generation i are indicated as g_{ij} with function f_j. This network represents an elementary Darwinian process by virtue of assigning negative (inhibitory or competitive) values to the diagonal f_j terms. The Darwinian framework describes the changing features of populations of organisms over time and explains these in terms of both inherited features of ancestors and the differential success of features across a range of temporally separated environments.

For a population of N different genotypes (classes of organism), this scheme can be translated directly into a fundamental equation of evolutionary dynamics, the replicator equation (RE), which in continuous time drops the generational index preserving only the functional index, and imposes a constraint on total population size through density dependent regulation (negative self-interaction).

$$\dot{g}_i = g_i\left(f_i - \sum_j f_j g_j\right), \tag{1}$$

$$= g_i(f_i - \langle f \rangle). \tag{2}$$

By introducing mutation rates q_{ij} where type $j \to i$, we allow for one to many mappings across generations and drop the unrealistic requirement that all diversity be present as an initial condition. This generates an expanded system of equations described as the replicator-mutator equation or quasispecies equation.[29]

$$\dot{g}_i = \sum_{j=1}^{N} g_j f_j q_{ij} - g_i \sum_i^N f_i g_i. \tag{3}$$

This framework describes the temporal variation in the composition of a population of density-limited replicators. According to one popular short text in population genetics, evolution is the change in frequencies of genotypes through time, perhaps due to their differences in fitness.[15] The "perhaps" in the

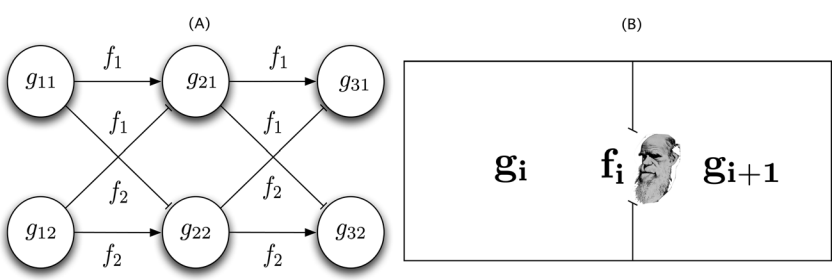

FIG. 1. Evolution as a directed, acyclic graph (A) and as a Maxwell's demon box (B). Genomes in successive generations encode functions that enhance their own proliferation at the expense of competitors. Evolution by natural selection describes a heterogenous flow of genotypes forward in time. In demonic dynamics, the complex environmental interactions are aggregated into the figure of the demon capable of computing adaptive mappings between generations. There dynamics of natural selection are captured by cognition by the demon.

sentence refers to the role of drift, or demonic indifference, promoting neutral variability in the population.

Let's consider the replicator dynamics in a little more detail (see Refs. 39, 37, and 29 for a synoptic treatment of this equation and its properties). Consider a 1-dimensional, discrete lattice of free energy (meaning thermodynamic internal free energy availability for work)—food, partners, shelter. This space represents a preference function of the Darwinian demon by providing a selective gradient that is maximum at some lattice site. This is achieved by making available to each genotype a quantity of resource f_i. In Figure 1(A), we described these as weights mapping genotypes onto themselves and inhibiting competitors. Populations distribute themselves over this functional resource space generating in a sufficiently large population a distribution over the entire space, g_i. So g_i is an investment in the function f_i. Selection in a population of fixed size can be described deterministically by the RE,

$$\dot{g}_i = g_i\left(f_i - \sum_j f_j g_j\right).$$

The states of this equation are confined to the unit simplex s_N in n-dimensional Euclidian space,

$$s_N = \left\{\mathbf{g} \in \Re^N : g_i \geq 0 \forall_i \in N, \sum_i g_i = 1\right\},$$

This allows us to interpret g_i as a probability distribution and evolution in terms of Shannon information. We define the amount of information gathered (IG) about the demon's preferences at time t as the difference between the maximum possible certainty over the choice of function and the observed information. This can be written in terms of standard Shannon entropies

$$IG(t) = \sum_i^N g_i(t)\log g_i(t) - \log(1/N).$$

The solution to the RE is unique and depends only on the relative values of f, such that,

$$g_i = 1, \quad i = max_i(f_i),$$
$$g_j = 0, \quad \forall_j \in n \setminus i.$$

This demonstrates a fundamental Darwinian or selective principle: the amount of information that is acquired for performing a preferred function is directly proportional to the loss of life. The ability to find the adaptive peak requires that all resources are channeled into a single genotype. This can be stated slightly more formally,

$$IG(t \to \infty) = \log(N). \tag{4}$$

Hence, around N variants must be lost to obtain a log(N) quantity of information about the environment. This is the first and most fundamental thermodynamic constraint on the evolution of biological systems, as the ability to propagate a functionally superior genotype is proportional the diversity of genotypes and their dissipation from the population. For the Darwinian demon, information bearing degrees of freedom are payed for with the loss of information about the remaining states of the environment.

A related complexity measure that could take into account population size, as well as population diversity (N), has been suggested to me by Eric Smith building on the concept of effective complexity of Gell-Mann and Lloyd.[14] If we now consider a population of M individuals, then we might measure the total information gain as the relative entropy, or Kullback-Leibler divergence, between the evolving distribution and the uniform prior distribution at time $t=0$. For the sake of brevity, let us define the vector $x_i = g_i(t=0)$. Then, this information gain is given by,

$$IG_{kl}(t) = M\sum_i g_i(t)\log\left(\frac{g_i(t)}{x_i}\right). \tag{5}$$

At equilibrium, this would provide a measure of evolutionary information that is extensive in the total population size,

$$IG_{kl}(t \to \infty) \approx M\log(N). \tag{6}$$

The population size variable M increases the total information acquired by expanding the repertoire of distinctions available to selection. The population size parameter M acts rather like effective population size in population genetics, which modulates the relative contribution of selection and drift. Higher values of M correspond to a reduction in sampling noise.

For a fuller exploration of the replicator dynamic as an inferential or informational process, Harper[18] and Frank[12] have recently discussed how the replicator equation and its extensions such as the Price equation are equivalent to gradient flow of the Fisher information metric. Also note that it is perfectly appropriate to use the phrases "selection against" and "selection for" some trait. Selective demons select against non complementary genomes and select for complementary genomes. This is just as a porous sieve selects against retaining small particles and for retaining large particles. These are mechanically dual operations.

IV. ERROR-CORRECTING FREE ENERGY BUDGETS AND EVOLUTIONARY MAXIMA

In Sec. III, I considered information at the level of populations and neglected to mention how much information could be carried by the winning genome. For each of the f_i functions, there must be some mechanism generating an appropriate phenotype. In this section, I explore the idea that the amount of information present in the winning genome is the maximum quantity that can be propagated reliably given a budget for metabolic free-energy apportioned to genetic repair.

To explore information at the individual level, we need to consider the expanded quasi-species equation, which adds to the replicator equation provisions for individual informational properties in terms of variation in genome length. Recall the quasi-species equation,

$$\dot{g}_i = \sum_{j=1}^N g_j f_j q_{ij} - g_i \sum_i^N f_i g_i. \tag{7}$$

Define the probabilities of mutating among genomes by considering a simple bit-wise operation, in which the probability of any one element of the genome of length L mutating is given by a probability p. Hence,

$$q_{ij} = p^{H_{ij}}(1-p)^{L-H_{ij}}. \tag{8}$$

The value H_{ij} is the hamming distance between genomes i and j, and the binomial distribution of genomes arises because a transitions between any two genomes become exponentially more rare the further they are apart in sequence space.

A. Genomic allometry

One of the defining characteristics of biology is the pervasiveness of scaling laws which relate macroscopic features of living systems back to the constraints on efficiently distributing energy rich compounds for growth and repair. The best known of these are the so called allometric scaling laws relating brain size and body size, or basal metabolic rate to body mass.[45] These relationships are frequently described in terms of quarter power exponents.

Of immediate interest in this paper is that the repair rate of genetic damage is limited by body mass through the basal metabolic rate.[16] Empirical data from prokaryotes reveals relationship of the form,

$$p = km^{-1/4}$$

and

$$L = k'm^{1/4}.$$

As the mass increases, the per site probability of a mutation declines as one quarter power, whereas the genome length increases as mass to the one quarter. I shall assume the first of these laws as a means of deriving the second and conjecture that selection seeks to maximize the information transmitted by genomes. We shall return to the informational grounds for this conjecture in Sec. IV B.

First, substitute the first scaling law into the genome wide mutation rate in order to obtain error rates proportional to mass,

$$q_{ij} = (km^{-1/4})^{H_{ij}}(1-km^{-1/4})^{L-H_{ij}}, \tag{9}$$

Define the per bit fidelity as $q = 1-p$ and the probability of the entire genome replicating without error as

$$Q = (1-km^{-1/4})^{L}.$$

Assume (see Ref. 29) that the wild type genome (with no mutations) has fitness $f_w = f_1 = 1+s$, and all mutants are equally less fit $f_m = f_i = 1, \forall i \setminus 1$. This allows us to consider a two dimensional system, rather than the 2^L dimensional space in the case of a binary genomes, or 4^L for a nucleic acid genomes. As every mutation to a wildtype generates a mutant, whereas only 1 in 2^L generates a wild-type from a binary mutant, ignore back-mutation. Note that this imposes the strongest possible selective penalty on mutants, and any reduction of severity (for example allowing for a continuous fitness function) will only make matters worse for the wild-type which will face stronger competitors. Continuous landscapes are important when it comes to finding the best genome, but are not useful for calculating upper bounds on survival. The quasi-species equations become,

$$\dot{g}_w = f_w g_w Q - g_w(f_w g_w + f_m g_m), \tag{10}$$

$$\dot{g}_m = f_m g_m + f_w g_w(1-Q) - g_m(f_w g_w + f_m g_m). \tag{11}$$

We seek to determine the stability of the fixed points of this system rather than the full time-dependent solution, and this allows us to drop through symmetry the non-linear density dependent terms of the equations, yielding

$$\dot{g}_w = f_w g_w Q, \tag{12}$$

$$\dot{g}_m = f_m g_m + f_w g_w(1-Q). \tag{13}$$

The ratio of the wildtype to the mutant converges to,

$$r = \frac{(1+s)Q-1}{(1+s)(1-Q)}.$$

For the wildtype to persist in equilibrium, we require that $Q > 1/(1+s)$, hence

$$q^L > 1/(1+s),$$

which provides us with a threshold per bit fidelity above which the wildtype persists

$$q > 1/(1+s)^{1/L}.$$

Take the log of both sides assuming that $s \approx 10$ (around an order of magnitude above the mutant), and obtain,

$$L < \frac{1}{-\log(q)},$$

and recognizing that $-\log(q) \approx 1-q$ for values of q close to one, we find that,

$$L < \frac{1}{p}.$$

Substituting $p = km^{-1/4}$, we find the threshold value above which L leads to the loss of genomic information is given by,

$$L = k'm^{1/4}.$$

This fundamental error threshold[10] allows the derivation of the second empirical scaling law relating genome size to mass. Hence, the conjecture of a maximum rate of evolutionary change subject to preserving the per bit fidelity with an apportioned free-energy budget, is justified by reproducing the scaling constraint placed on the maximum coding capacity of a genome. Note that it is the scaling relationship between L and q that we care about, and all more inclusive errors rates (coincident modifications of more than one site), will not modify this bound.

B. Evolutionary light speed

This result provides an intriguing insight into the maximum rate of sustainable, evolutionary change via selective dissipation. In Figure 2, the implications of these evolutionary scaling laws are illustrated in terms of maximum rates of genetic information accumulation in analogy with a better known fundamental limit on rates of change, the maximum velocity of light in a vacuum. For organisms, the unit of time is the generation and the units of space, the number of mutations per genome. For electromagnetic phenomena, the units of time are seconds and the units of space, light seconds. In both cases, the diagonal represents the maximum rate of propagation—for evolution one mutations per genome per generation and for light, one light second per second. All points within the light and life cones are accessible from an arbitrary point in the causal present, and are accessible from arbitrary points within the causal past. Hence, the scaling law would seem to reflect an allocation of free energy to error repair that maximizes the rate of evolutionary diversification while remaining within the bounds of the error threshold. It should be emphasized that these data are for prokaryotic organisms that evolve very fast. For eukaryotes and multicellular organisms that live well within the bounds of the life cone, there will be excess coding capacity not expended in adaptive search.

In the following sections, I shall argue that the increase in biological complexity consists in the evolution of mechanisms that permit organisms to achieve rates of change that move them outside of the causal future of their life cones by identifying alternative mechanisms of transmission that do not violate the thermodynamic constrains of Darwin's demon: that the process of selective information acquisition is limited by the rate of error-free generational turnover. If for some reason, the vector f_i fluctuates on a time scale shorter than the fixation time of a genome g_i (the mapping is not stable), then adaptation by means of natural selection becomes impossible. Differential death would no longer be sufficient to acquire environmental information and some new information acquisition and propagation mechanisms need to evolve.

This fundamental constraint on evolutionary dynamics allows us to relate the dissipation of information at the level of population measurement—the number of distinct states of the environment, $\log(N)$, to the genomic maximum of information within individuals—the number of functionally distinct states of the genome, L

$$L = \log(N).$$

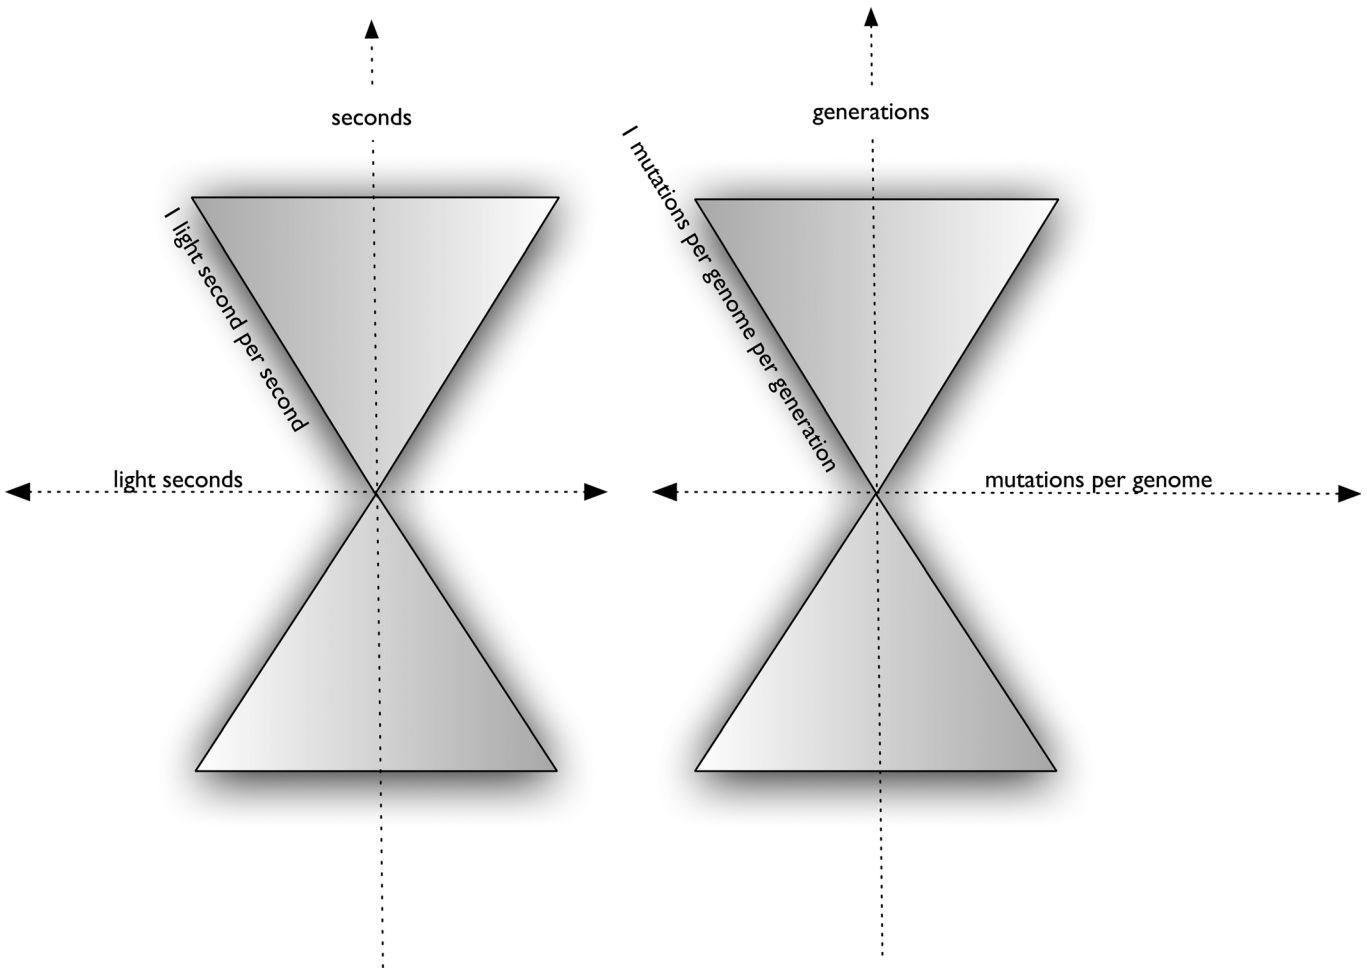

FIG. 2. Light cones and life cones illustrating two maxima in rates of information propagation. The horizontal line indicates the present. The diagonal in both cases sets the upper limit on the rate of change and anything outside of the cones are unreachable from the causal present. The diagonal of the life cone is the locally stable maximum L solution to the quasi-species equation, $L = 1/p$. Similarly, all events in the past influencing the present must lie within the cones.

This equality arises because each adaptive function (each discernibly different input to the Darwinian demon) must be encoded by a different state of a genome. Because the information that can be accumulated by selection is extensive in genotypic diversity, the increase in genome coding capacity must be extensive at the sequence level.

In order to increase organismal information about regular features of the environment beyond this maximum, new mechanisms need to be acquired that satisfy the metabolic constraints on genome size, but transcend these constraints at another level. I shall argue that these mechanisms are in a class of functions that we recognize as cognitive (sensing, learning, and memory mechanisms), and that increasing cognitive capability is synonymous with biological complexity at the organismal level, by which I mean increases in information storage capacity, information processing ability, strategic or behavioral diversity, and information transmission bandwidth.

V. AN EVOLUTIONARY COMPLEXITY MEASURE

Continued evolution requires that the environment is constantly changing. The Darwinian demon is changing its "mind" over its preference function (the position of the lattice site with maximum free energy available for metabolism etc). One reason for this is that the demon is itself experiencing dissipation alongside the populations that it is selecting. Another is that the demon can be directly modified by populations under selection (niche construction). Leaving aside these complexities for the moment, we shall consider the case where the demon periodically samples its preference function from some known probability distribution. In other words, a free-energy lattice site f_i on one dimensional lattice of dimension N is now a continuous random variable, where $\Pr[0 \leq f_i \leq F] = \int_0^F \theta(f_i) df_i$. At an interval of several generations (this must be greater than the number of generations required to fix a genotype), the lattice is periodically reconfigured sampling from the distribution $\theta(f)$.

We previously established that evolutionary dynamics maximizes the information content of the genome subject to budget of available free energy apportioned to error correction. Hence, selection will maximize the entropy $H(L)$ subject to $L \leq 1/p$. We also know that at equilibrium, for $\log(N) = L$, there is little or no information in the environment that is not present in the genome. Over evolutionary time, L is itself evolving and as with the environment there will be some probability distribution over the 4^L states of the genome, call this $\hat{L}$. The genome records the lattice position favored by the demon. This implies that selection seeks to minimize equivocation or conditional entropy $H(\hat{L}|\theta)$, capturing all environmental information not present in the genome. Let us define the complexity of our system as the maximum of the difference between the information in the genome and the equivocation,

$$C = H(\hat{L}) - H(\hat{L}|\theta), \tag{14}$$

$$= H(\hat{L}) + H(\theta) - H(\hat{L}, \theta). \tag{15}$$

This quantity is known as mutual information and measures the maximum amount of information that a selective demon can transfer into a successful genome from its own internal memory. Adami first proposed this quantity as a natural measure of evolutionary complexity[1] focusing on the universality of the genome as a natural basis for comparison. We note that there can be no information in the genome that is not already present in the demon. Any state of the genome that is not under scrutiny will simply drift towards some maximum entropy configuration. We find that as the information content of the demon/environment increases, so does the genomic complexity. We also understand that whereas $H(\hat{L})$ is bounded by free-energy constraints on L, we are not aware of similar bounds on $H(\theta)$. The joint entropy $H(\hat{L}, \theta)$ ensures that organismal complexity cannot grow beyond the limits imposed by the maximum L, as for any preferred state of an environment where $\log(N) > L$, there will be a maximum of ignorance over which state the demon prefers. Similarly, increasing the rate of error (e.g., mutation) will increase $H(\hat{L}, \theta)$ by randomly assigning genotypes to states of the environment. This measure of complexity is ensemble based and requires that we take measurements over many genomes in a variety of environments.

Mutual information is a natural measure of evolutionary complexity as we have shown it is maximized by evolutionary dynamics. If this is so, why are all lineages not equally complex? Why do some lineages become "living fossils" and preserve a fixed configuration for very long spans of time. And why under certain conditions does complexity reduce through time? Let us consider each in turn.

(1) *The diversity of complexity:* The most obvious explanation for simplicity is that the discernible environmental variability (the number of states of the demon) remains low. Not all organisms experience the same environment such that even "simple," low $H(\hat{L})$, lineages could be at the maximum of complexity given their environments. There is simply no further information available to be extracted. If we allow that environments differ in their complexity so should organisms.

(2) *The constancy of complexity:* There is no need to increase organismal information content when equivocation has been minimized and regulatory information maximized. If an environment is unchanging for long stretches of time, then the same should be true for the organism.

(3) *The reduction of complexity:* There are two scenarios that should lead to a reduction of complexity. The first is a reduction in the probable states of the environment. If N is reduced, then any marginal cost on increasing L should lead to a reduction in $H(\hat{L})$. This is beautifully exemplified in parasite evolution, where the free energy probabilities of a "lattice site" (trait value) become zero or one. Hence, if an enzyme required to complete the virus life cycle is predictably encoded by the host, or more generally if some feature of the virus becomes redundant, there is no need for the virus to encode it. This will lead to a reduction in $H(\hat{L})$. The second scenario corresponds to an increase in noise or a reduction in free energy available for error correction. This will tend to increase $H(\hat{L}, \theta)$ and reduce organismal complexity through increased dissipation of information.

A. From ENIAC to CMOS: Chimerical structural complexity

The technological evolution of turing complete digital computers provides a useful analogy for revealing a few perils of evolutionary complexity measures. From the electronic numerical integrator and computer (ENIAC) designed in 1943 and built by 1946, to the complementary metal oxide semiconductor (CMOS) designed just over twenty years later in 1967, we have witnessed non-linear returns to scale in size, energy efficiency, and processing power. An ENIAC weighing 30 tons filled a room of around 3 by 30 m, consumed 150 kW of power, and calculated at a rate of around 0.05 million instructions per second (MIPS). An Intel core i7-908X operates at just under 150 000 MIPS, takes up around 250 mm^2 and can be purchased at most electronic retailers by mail order and delivered in an envelope for around $1000.0.

From the perspective of mechanical complexity, the ENIAC has far and away more moving parts than an integrated circuit, and the assembly of an ENIAC was fraught with engineering challenges. The manufacture of i7 chips has reached a level of extraordinary precision, and once built at an incredibly expensive and energy-demanding facility, they experience a relatively low rate of failure. The manufacturing process is not however visible in the final product.

The question is how might we measure the complexity of each of these devices if we did not have access to their input output behavior, or the sequence of operations through which they were built? At first glance, the ENIAC might be assumed to be more complex as it has so many more moving parts, each of which requires such huge amounts of energy to be powered. A closer inspection reveals that the i7 is a marvel of miniaturization and contains extraordinary structure at the microscopic scale. But then so do the individual vacuum tubes of the ENIAC inspected at the atomic scale which possess equivalent atomic behavior. It is possible that the microarchitecture of the *i*7 would be interpreted erroneously as a regular crystal structure with impurities. The point is that without knowledge of behavior and armed only knowledge of static structure, even a principled complexity measurement is liable to produce complexity chimera. This obvious fact from the life of machines has largely escaped those interested in real life. There is still a desire and practice of comparing structure with structure as a measure of complexity, whether this is a simple comparison of genome size, or ratio of brain to body size. Without prior knowledge of process, these measurements are likely to be misleading.

The ENIAC and the integrated circuit suggest that a sensible complexity measure should endeavor to evaluate the computational capability of each device, subject to constraints on free energy.

VI. TRANSCENDING SELECTIVE CONSTRAINTS, RECURSIVE DEMONS AND INFERENCE

A strong implication of what has been said is that solutions to continuous survival questions posed by a selective demon need to be solved through differential mortality. It is evident that there are many bits of information acquired by evolved species that have not enforced this constraint. Members of our own species can store all kinds of information without the need for a huge army of clones, each of whom differs by only a few bits, and only a few of whom survive after something has been learnt. This is however what natural selection requires, and it is an extortionate price that most lineages have long since stopped paying. The way that this has been achieved has been to internalize our Darwinian demons by evolving mechanisms of cognitive inference and dissipating energy not through mortality but during neural processing.

Mechanisms of plasticity and learning, supported by neural architecture, provide for the possibility of adaptation at the individual level without population selection. Population selection continues to play a role fixing mechanisms that enable adaptation and plasticity. In Sec. VI A, we consider three popular frameworks for the analysis of learning (imitation learning, reinforcement learning and Bayesian inference) and review evidence for a fundamental equivalence between these and the replicator equation of evolutionary dynamics. With this selective equivalence principle in hand, we shall see that the same problems of "demonic" selection thereby apply to brains and cognitive system generally, and this illustrates the need for nested selection mechanisms in support of all adaptive processes. Another way to say this is that learning does not overcome the fundamental computational constraints imposed by selective demons, and what is learnt cannot exceed the complexity of the reward signal.

A. Imitation

In this section, I follow[40] approach to imitation learning. Assume that each individual i varies in its strategy choices $\mathbf{s}_i$. Think of this as a vector specifying a mixed strategy defining how individuals should interact. Interactions between strategies $\mathbf{s}_i$ and $\mathbf{s}_j$ give rise to pairwise payoffs or rewards r_{ij}. Reward information is used to decide which strategy to adopt in subsequent interactions. In particular, differences in payoffs lead to the adoption of more successful strategies. Thus, the rate of imitation leading strategy j to imitate a strategy i, $\mathbf{s}_j \to \mathbf{s}_i$ can be written in matrix form,

$$f_{ij} = [(R\mathbf{g})_i - (R\mathbf{g})_j]_+, \tag{16}$$

where R is the matrix of rewards and g the vector of genotypic or strategic frequencies. Notice that this function is only non-zero in the positive half-space, such that lower payoff strategies are not imitated.

The population of players will evolve in time according to an imitation learning dynamics,

$$\dot{g}_i = g_i \sum_j (f_{ij} - f_{ji}) g_j, \tag{17}$$

$$= g_i \sum_j (R\mathbf{g})_i - (R\mathbf{g})_j] g_j, \tag{18}$$

$$= g_i[(R\mathbf{g})_i - \mathbf{g}R\mathbf{g}]. \tag{19}$$

The last of these is our familiar replicator equation once again. Imitation learning of successful strategies has an identical structure to evolutionary dynamics in a resource-limited lattice and requires an equivalently discerning demon.

Whereas in the evolutionary case individuals needed to die in order that the best adapted replicator fixed in the population, in this case, strategies are dying. Information is lost not at the level of the individual genome, but at the level of a plastic repertoire of learnt behaviors. In addition, the demon is now "frequency-dependent" and distributed over the population of learners.

B. Reinforcement or operant conditioning

Operant conditioning differs from imitation learning in that discrete behaviors are rewarded proportional to their probability in some environmental context. In other words, like the simple replicator equation and unlike imitation learning, operant conditioning is not frequency dependent. I follow the (Ref. 46) derivation for operant conditioning and replicator dynamics. Consider a set of behaviors X_i each associated with a probability x_i. For each behavior, there is a reward r_i. Mathematical learning theory typically assumes a "linear operator" model for updating behavioral probabilities based on rewards. Hence, the incremental change in the probability of any action, with a learning rate parameter α, can be written as,

$$\delta x_k = \alpha r_i (e_{ik} - x_k), \tag{20}$$

where the e_{ij} is the Kronecker δ function,

$$e_{ik} = \begin{cases} 1, & \text{if } k = i \\ 0, & \text{if } k \neq i. \end{cases}$$

This gives the updated probability of behavior assuming that a single action was performed. The average change of behavior takes into account the probability of an action x_i,

$$\Delta x_k = \sum_i x_i \delta x_k. \tag{21}$$

Combining these two equations, we deduce that

$$\Delta x_k = \sum_i x_i [\alpha r_i (e_{ik} - x_k)], \tag{22}$$

$$= \alpha \sum_i x_i r_i e_{ik} - \alpha x_k \sum_i x_i r_i, \tag{23}$$

$$= \alpha x_k r_k - \alpha x_k \sum_i x_i r_i. \tag{24}$$

The term $\sum_i x_i r_i$ is simply the mean reward, $\langle r \rangle$, and the change in behavior in continuous time is given by,

$$\dot{x}_k = x_k (r_k - \langle r \rangle). \tag{25}$$

Hence, behaviors that receive a reward that is greater than the average reward will be reinforced, and those obtaining a lower reward, extinguished.

C. Bayesian updating

Bayesian inference provides a convenient means of describing hypothetico-deductive reasoning. A limited set of hypotheses are subjected to scrutiny through experiment, and beliefs or posterior distributions over hypotheses are updated. It is fairly natural to think of organisms as reified hypotheses about the state of the environment, and selection as a natural experiment that either supports (through survival) or refutes (through morbidity) a given genotype and attendant phenotype. Prior beliefs are encoded in the variation of heritable states of the genome encoding regular features of past environments. The selective demon supports those phenotypes that provide the most compelling hypothesis about the whereabouts and mechanics for extraction of free energy. As Refs. 38 and 18 have demonstrated, this association can be made more concrete by connecting Bayesian updating directly to replicator dynamics. First, we need to assume that there exists some continuous distribution of hypotheses, $P(X)$, that we shall update according to a conditional likelihood function $L(X)$ capturing the filtering properties of selection. In discrete time,

$$P(X)_t = P(X)_{t-1} \frac{L(X)_{t-1}}{\langle L \rangle_{t-1}}, \tag{26}$$

and the change in the concentration of the probability mass around the highest likelihood values,

$$\Delta P(X)_t = P(X)_{t-1} \frac{L(X)_{t-1}}{\langle L \rangle_{t-1}} - P(X)_{t-1}, \tag{27}$$

$$= P(X)_{t-1} \left(\frac{L(X)_{t-1}}{\langle L \rangle_{t-1}} - 1 \right), \tag{28}$$

$$= \frac{1}{\langle L \rangle_{t-1}} P(X)_{t-1} (L(X_{t-1}) - \langle L \rangle_{t-1}). \tag{29}$$

And the evolution of the probability distribution in continuous time,

$$\dot{x} = \frac{1}{\langle L \rangle} x (L_x - \langle L \rangle). \tag{30}$$

The selective interpretation of Bayesian updating is that hypotheses that are better than the population average will tend to increase in frequency, reaching fixation at that environmental value which cannot be improved upon.

D. Isomorphism among dynamics of adaptive plasticity and the demon

All adaptive dynamics described in terms of probability distributions live in the real valued unit simplex. This geometric constraint imposes a fundamental equivalence among all these frameworks. All can be thought of as information maximization algorithms leading to a best strategy that is most effective at identifying peaks in free energy landscapes. The mechanisms underlying each of these inferential rules are distinct but they all respect similar constraints. Learning substitutes an ensemble of fixed replicators competing over many generations with an ensemble of cells or hypotheses competing during ontogeny. Whereas Darwinian demons dissipate through mortality, learning demons dissipate through the elimination of false beliefs.

Both are selective demons and possess the same elementary functional contraints.

Learning and selective demons differ from Darwinian demons in that they require that selection pressures are internalized as reward systems. During Darwinian selection, individuals vary not only in their primary selective traits, but in their ability to represent the "preferences" of the environment. Learning can determine a best strategy without having to expend a large number of genomes in order to do so.

VII. THE CONSTRUCTION OF THE DEMON - THE ORIGIN OF SELECTIVE GRADIENTS

One virtue of Maxwell's thought experiment is that it allows the demon to be thought of as a coherent and localized computational entity. For evolutionary dynamics and learning, the demon is in reality a highly distributed system of agents and physical processes that exert an influence on the survival probabilities of different organisms or hypotheses. A model of the demon would imply a model of selection pressures in all of their ecological complexity. Evolutionary theory has until very recently ignored the origins of selection, focusing on the origins of traits under known selective constraints. If we are going to explain how evolutionary complexity increases, we need some means of explaining how and why a Darwinian demon becomes more selective, and over an increasing number of degrees of freedom, rather than dissipating into adaptive indifference over evolutionary time.

An analogous question exists in cosmology which seeks to explain sources of free energy available for life in terms of an "entropy gap," or the difference between the maximum entropy of the universe and the entropy at any given point in time after the big bang.[24] The conventional wisdom has it that early in the life of the universe entropy is at a minimum, and that low gravitational entropy allowed for the formation of gradients, nuclear and chemical, upon which all of known life depends. Life is assigned to the highest strata of a pyramid of free energy production and sits upon free energy created by the thermal and chemical disequilibrium of the earth, and the electromagnetic free energy from photons emanating from nuclear fusion in stars experiencing gravitational collapse. As the universe "evolves," the entropy gap systematically reduces, making free energy for useful work increasingly scarce. Life pursues the opposite path, having started as relatively high entropy hypercycles, the terrestrial ecosystem has evolved a variety of exquisitely coordinated regulatory networks with a rich hierarchical structure.

The question becomes, how can increasingly simple physical gradients select for complex adaptive mechanisms? In other words, what is it about chemical potentials that allows for a very significant increase in evolutionary coding capacity. This question has been pursued most rigorously by Smith.[41] One simple answer is that any chemical dynamics that confers upon a system elementary properties required for evolutionary dynamics (assumptions of the replicator equation) will have added to a system a new kind of demon sensitive to variation at scales neglected by physics and chemistry alone. For example, assume a simple autocatalytic network $A \rightarrow B \rightarrow C \rightarrow A$ that interacts at the level of rate constants with the coupled network $C \rightarrow D \rightarrow E \rightarrow C$. Modification of any element of the networks that increases the rate of flow away from the shared element C (either $C \rightarrow A$ or $C \rightarrow D$) will produce as a consequence a reduction in the concentration of the members of the unmodified network. The simple fact of coupling elements of minimal homeostatic or self-replicating systems thereby leads to accumulating complexity of interactions. Coupling leads to simple mechanisms of selection. The connection between simple prebiotic sets and the replicator equation is straightforward. Label distinct chemical species as X_i and their densities x_i. The hypercycle equation[11] describing neighboring interactions on the ring lattice is given by,

$$\dot{x} = x_i(x_{i-1}F_i(\mathbf{x}) - \langle F \rangle), \quad i = 1, \ldots, n, \tag{31}$$

where the indices are taken on modulo n and the functions F_i are positive rates of catalysis defined over the population of reactants. When F_i are constants this yields a first-order replicator equation.

$$\dot{x} = x_i(x_{i-1}r_i - \langle F \rangle), \quad i = 1, \ldots, n, \tag{32}$$

In this case, fitness is frequency dependent and it is the outcome of basic enzymatic interactions. Each element of the hypercycle encodes some part of the demonic preference function. As we move towards more elaborate networks, the locus of preference becomes more diffuse. Any complex trait will be composed of numerous, basic chemical interactions, and all them will be under scrutiny from their coupled interaction partners. A clue to increasing complexity would be some systematic expansion in the range of distributed selection pressures. I will touch on this briefly in Sec. VII A.

A. Niche construction as a complexity ratchet

Recently, there has been a renewed effort to provide a macroscopic theory for the origins of selection mechanisms. Much of evolutionary theory has posited selective values and determined the consequences of selection on the fixation and distributions of genotypes. This is what I have assumed in all of the models described in this paper. It would seem a natural next step to determine the origin, fixation, and distribution of selection pressures themselves. This effort goes under the name of niche construction.[30] One approach would be to determine the origin of all of those elementary biochemical processes indirectly in support of complex structures, recognizing that selection pressures are in reality composed of a multitude of physical and biological interactions. The theory operates at a higher level of abstraction, determining the origin of coarse-grained filters that responds to the average properties of a target genotype. Most importantly, the selective filter is partially constructed by the organism itself, and exerts an influence over a locally reproducing population—a deme. In this way, adaptations lift themselves up by their own bootstraps.

By the demonic selection principle, any niche construction algorithm sets an upper bound on the complexity of its target trait (following the same logic that a natural selective demon sets the upper bound on the degrees of freedom of an

organismal trait). But niche construction, by introducing a form of recursion into selection dynamics, might be able to generate traits of unbounded complexity. For example, a bird nest cannot stabilize through selection characteristics of nest-making birds that go beyond the code length required to specify nest-making behavior. But nest making behavior (the algorithmic specification of the behavior) is mutable and subject to rare elaboration. A bird b builds a nest n_0 that might select for an additional nest-related bird trait b_1. This trait could then favor a new type of nest n_1 that in turn favors traits b_2. This iteratively could select for new bird traits (b_i). There is a form of interaction, or in genetic terms "epistasis," between the trait and the niche that increases the resolution of selection. Organisms by modifying selection pressures generate a selective ratchet, achieving significant increases in adaptive complexity through positive feedback between niche construction and niche selection. This hinges of the idea that there is excess coding capacity (not yet at a maximum value of L) that can respond to the increased selective resolution. Obviously, this remains a rather speculative and non-rigorous conjecture that might be interesting to pursue further.

As we have suggested elsewhere,[22] the pinacle of environmental niche construction is the phenotype itself. The organismal phenotype exerts the most direct selection pressure on the genome, since gene expression products are first evaluated in the context of the cell and the tissues of the phenotype. The process of development, or the sequence of cell divisions and differentiation generating a mature multicellular organism, can be interpreted as a highly regulated instance of ecological niche construction. During development, the germ line gains largely unrestricted access to the diverse metabolic products of somatic cells and the benefits of their behavior. The genome encodes a regulatory algorithm for building phenotypes, in much the same way that some genotypes encode rules for building nests and spider webs.

What I am suggesting is that genomes that can construct demons (construct their selective context) might be the key to systematic trends toward increasing mutual information between organism and environment $I(\theta;\hat{L})$. Not all lineages will possess this capability beyond the developmental niche construction required for survival. However, some lineages will have entered into positive feedback with their extended environment and initiated a trend of increasing inferential capability. This is very familiar in our own species as gene culture coevolution, and more recently, culture-machine coevolution, producing regularities such as the scaling relationship Moore's law. If niche construction is key to understanding increasing evolutionary complexity, then we shall need a more general formulation of the construction algorithms that build bodies and their extended environments. This intuition was in some measure articulated over fifty years ago by Waddington when he wrote,[43] "We need a hereditary system which does not merely contain information, but which acts as algorithm or program and thus leads to the production of a phenotype which takes its place between the genotype and environment." Waddington finished with the remark, "Possibly the people who are trying to discover how to set up a computer to learn to play good chess, or bridge, are among those most likely to make a major contribution to the fundamental theory of evolution.". At this point, I would only substitute the game of Go for the solved challenge of chess.[7]

VIII. CONCLUSION

I have sought to demonstrate that the Darwinian selection theory bears a structural similarity to Maxwell's thought experiment with a demon. In much the same way that the computational capability of the demon limits deviation from thermodynamic equilibrium, the selective demon limits the number of information bearing degrees of freedom in a genome. I have shown how evolutionary dynamics maximizes the information present in a genome, and that when environments vary, mutual information is a natural ensemble measure of evolutionary complexity. Variation in investment in error repair imposes an upper bound on the rate of evolutionary change and limits the coding capacity of a genome in a given environment. This limit can be overcome by evolving mechanisms of learning and plasticity which do not dissipate energy through mortality, but through the elimination of neurally encoded false hypotheses. However, these learning dynamics are themselves instances of selection processes, and hence subject to the same limits on information acquisition—learnt behaviors cannot exceed the complexity of the reward signal. Given that for both evolution and learning, the environment establishes a limit on complexity, we need to identify mechanism for enriching the environment. I suggest that this takes place through a process of niche construction. Organisms actively modify their environments, and thereby reshape selection pressures. Through positive feedback with mutability, niche construction in large communities of niche constructing organisms might be capable of generating a ratchet that iteratively increases mutual information between organism and environment by increasing access to fine-grained environmental resources. This implies that a meaningful measure of complexity will need to be ecological rather than structural, emphasizing statistical associations rather than individual traits, and build upon a general model of biological construction algorithms.

ACKNOWLEDGMENTS

Thanks to Eric Smith for many years of fruitful conversation and original insight into the statistical mechanics of adaptive processes and his suggestions for improving this manuscript. Thanks to Jessica Flack for many discussions on biological complexity and the nature of selection, and to Jon Machta for ideas and comments leading to the improvement of this manuscript. Thanks to Nihat Ay for his insights into information theory. Thanks to the Templeton Foundation for financial support.

[1]C. Adami, BioEssays **24**(12), 1085 (2002).
[2]P. Aldhous, Science **262**(5133), 502 (1993).
[3]A. Bassi and S. Denazis, Concurrency Comput.: Pract. Exper. **20**, 2061 (2008).
[4]G. Bell, *Selection: The Mechanism of Evolution* (Oxford University Press, Oxford, 2008).

[5]J. T. Bonner, *The Evolution of Complexity by Means of Natural Selection* (Princeton University Press, Princeton, 1988).
[6]E. Borenstein and D. C. Krakauer, *PLoS Comput. Biol.* **4**(10), e1000202 (2008).
[7]X. Cai and D. C. Wunsch, "Studies in computational intelligence," in *JanuszKacprzyk Studies in Computational Intelligence 1860-949X1860-9503* (Springer, Berlin, Heidelberg, 2007), Vol. 63.
[8]K. D. Crow, Mol. Biol. Evol. **23**(5), 887 (2006).
[9]T. W. Deacon, *The Symbolic Species: The Co-Evolution of Language and the Brain* (W.W. Norton and Company, New York, 1998).
[10]M. Eigen, Biophys. Chem. **85**, 101 (2000).
[11]M. Eigen and P. Schuster, Naturwiss. **64**(11), 541 (1977).
[12]S. A. Frank, J. Evol. Biol. **22**(2), 231 (2009).
[13]K. Frenken, *Innovation, Evolution and Complexity Theory*, J. Evol. Economics **17**, 107 (2006).
[14]Gell-Mann and S. Lloyd, Complexity **2**, 44 (1996).
[15]J. H. Gillespie, *Population Genetics: A Concise Guide* (Johns Hopkins University Press, Maryland, 2004).
[16]J. F. Gillooly, A. P. Allen, and G. B. West, "The rate of DNA evolution: Effects of body size and temperature on the molecular clock," in *Proceedings of the PNAS* (2005), Vol. 102, pp. 140–145.
[17]S. J. Gould, *The Structure of Evolutionary Theory* (Belknap, Cambridge, MA, 2002).
[18]M. Harper, The Replicator Equation as an Inference Dynamic, arXiv:0911.1763v3[math.DS].
[19]D. P. Hill, D. A. Begley, and J. H. Finger, Nucleic acids Res. **32**, D568 (2004).
[20]J. Jaynes, *The Origin of Consciousness in the Breakdown of the Bicameral Mind* (Mariner Books, Houghton, Mifflin, Harcourt, New York, 2000).
[21]M. Kanehisa, M. Araki, S. Goto, M. Hattori, M. Hirakawa, M. Itoh, T. Katayama, S. Kawashima, S. Okuda, T. Tokimatsu, and Y. Yamanishi, Nucleic acids Res. **36**, D480 (2007).
[22]D. C. Krakauer, K. M. Page, and D. H. Erwin, Am. Nat. **173**(1), 26 (2009).
[23]H. S. Leff and A. F. Rex, *Maxwell's Demon 2: Entropy, Classical and Quantum Information, Computing* (IOP Publishing, Philadelphia, 2002).
[24]C. H. Lineweaver, Phys. Life Rev. **5**, 225 (2008).
[25]P. L. Luisi, Anat. Rec. **268**, 208 (2002).
[26]D. W. McShea, Paleobiology **31**, 146 (2005).
[27]E. M. Mellgren, Trends Genet. **18**, 128 (2002).
[28]A. R. Mushegian and E. V. Koonin, "A minimal gene set for cellular life derived by comparison of complete bacterial genomes," in *Proceedings of the National PNAS* (1996), Vol. 93, pp. 10268–10273.
[29]M. A. Nowak, *Evolutionary Dynamics: Exploring the Equations of Life* (Belknap, Cambridge, USA, 2006).
[30]F. Odling-Smee, K. Laland, and M. Feldman, *Niche Construction: The Neglected Process in Evolution* (Princeton University Press, Princeton, 2003).
[31]H. A. Orr, Evolution **54**, 13 (2000).
[32]D. A. Petrov, Trends Genet. **17**, 23 (2001).
[33]M. R. Pie and J. S. Weitz, Am. Natur. **166**(1), E1 (2005).
[34]E. L. Post, J. Symb. Log. **12**, 1 (1947).
[35]W. B. Saunders and D. M. Work, Science **286**, 760 (1999).
[36]T. J. M. Schopf, D. M. Raup, S. J. Gould, and D. S. Simberloff, Paleobiology **1**, 63 (1975).
[37]P. Schuster, J. Theor. Biol. **100**, 533 (1983).
[38]C. R. Shalizi, Electron. J. Stat. **3**, 1039 (2009).
[39]K. Sigmund, Complexity **16**, 88 (1983).
[40]K. Sigmund, *The Calculus of Selfishness* (Princeton University Press, Princeton, 2010).
[41]E. Smith, J. Theor. Biol. **252**(2), 185 (2008).
[42]J. W. Valentine, Paleobiology **26**, 513 (2000).
[43]C. H. Waddington, Nature **218**(5141), 525 (1968).
[44]G. P. Wagner and L. Altenberg, Evolution **50**, 967 (1996).
[45]G. West, J. Brown, and B. Enquist, Science **276**, 122 (1997).
[46]J. Zhang, Neural Networks **22**(3), 220 (2009).
[47]J. Zhang, Neural Networks **22**(3), 229 (2009).