

Primordial Evolution in the Finitary Process Soup

Olof Gornerup
James P. Crutchfield

SFI WORKING PAPER: 2007-05-008

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu



SANTA FE INSTITUTE

Primordial Evolution in the Finitary Process Soup

Olof Görnerup^{1,*} and James P. Crutchfield^{2,†}

¹*Complex Systems Group, Department of Energy and Environment,
Chalmers University of Technology, 412 96 Göteborg, Sweden*

²*Center for Computational Science & Engineering and Physics Department,
University of California, Davis, One Shields Avenue, Davis CA 95616, USA*

(Dated: April 27, 2007)

A general and basic model of primordial evolution—a soup of reacting finitary and discrete processes—is employed to identify and analyze fundamental mechanisms that generate and maintain complex structures in prebiotic systems. The processes— ϵ -machines as defined in computational mechanics—and their interaction networks both provide well defined notions of structure. This enables us to quantitatively demonstrate hierarchical self-organization in the soup in terms of complexity. We found that replicating processes evolve the strategy of successively building higher levels of organization by autocatalysis. Moreover, this is facilitated by local components that have low structural complexity, but high generality. In effect, the finitary process soup spontaneously evolves a selection pressure that favors such components. In light of the finitary process soup’s generality, these results suggest a fundamental law of hierarchical systems: global complexity requires local simplicity.

PACS numbers:

Keywords: structural complexity; entropy; information; computational mechanics; population dynamics; hierarchical dynamics; emergence; evolution; self-organization; autocatalysis; autopoiesis.

I. INTRODUCTION

The very earliest stages of evolution—or rather, pre-evolution—remain a mystery. How did structure emerge in a system of simple interacting objects, such as molecules? How was this structure commandeered as substrate for subsequent evolution—evolution that continued to transform the objects themselves? One wonders if this recursive interplay between structure and dynamics facilitated the emergence of complex and functional organizations. Since these questions concern the most fundamental properties of evolutionary systems, we explore them using principled and rigorous methods.

To build a suitable model a few basic ingredients are required. First, one needs some type of elementary objects that constitute the state of the system at its finest resolution. Second, one needs rules for how the objects interact. Third, one needs an environment in which the objects interact. Fourth, one needs quantitative and calculable notions of structure and organization. These requirements led us to the *finitary process soup* model of primordial evolution [1]. Simply stated, the soup’s ingredients are, in order, ϵ -machines, their functional composition, a flow reactor, and the structural complexity C_μ of ϵ -machines.

After explaining each of these ingredients, we will relate the model to classical replicator dynamics by reducing the soup to a special case. We then move on to contrast the limited case with the full-fledged finitary process

soup as a constructive, unrestricted dynamical system.

II. OBJECTS: ϵ -MACHINES

Here we employ a finite-memory process called an ϵ -machine [2–4], as our preferred representation of an evolving information-processing individual. Using a population of ϵ -machines is particularly appropriate in studying self-organization and evolution from an information-theoretic perspective as they allow quantitative measurements of storage capacity and randomness. Rather than using the abstraction of a formal language—an arbitrary finite set of finite length words—we consider a discrete-valued, discrete-time stationary *stochastic process* described by a bi-infinite sequence of random variables S_t over an alphabet \mathcal{A} :

$$\overleftrightarrow{S} = \dots S_1 S_0 S_1 \dots \quad (1)$$

A process stores information in a set of *causal states* that are equivalence classes of semi-infinite histories that condition the same probability distribution for future states. More formally, the causal states \mathcal{S} of a process are the members of the range of the map $\epsilon : \overleftarrow{\mathbf{S}} \mapsto 2^{\overleftarrow{\mathbf{S}}}$ from histories to sets of histories:

$$\epsilon(\overleftarrow{s}) = \{ \overleftarrow{s}' | P(\overrightarrow{S} | \overleftarrow{S} = \overleftarrow{s}) = P(\overrightarrow{S} | \overleftarrow{S} = \overleftarrow{s}') \}, \quad (2)$$

where $2^{\overleftarrow{\mathbf{S}}}$ denotes the power set of $\overleftarrow{\mathbf{S}}$. Further, let $\mathcal{S} \in \mathcal{S}$ be the current casual state, \mathcal{S}' its successor, and $\overleftarrow{S}^{\rightarrow 1}$ the next symbol in the sequence (1). The transition from one causal state \mathcal{S}_i to another \mathcal{S}_j that emits the symbol

*Electronic address: olofgo@chalmers.se

†Electronic address: chaos@cse.ucdavis.edu




FIG. 1: Three examples of ϵ -machines. T_A represents the identity function and has the causal state A . T_B has two causal states (A and B), accepts the input string $1010\dots$ or $0101\dots$, and operates by flipping 0s to 1s and vice versa. T_C has the same domain and range as T_B , but maps input strings onto themselves.

$s \in \mathcal{A}$ is given by a set of labeled transition matrices: $\mathcal{T} = \{T_{ij}^{(s)} : s \in \mathcal{A}\}$, where

$$T_{ij}^{(s)} \equiv P(S' = S_j, \overset{-1}{S} = s | S = S_i). \quad (3)$$

The ϵ -machine of a process is the ordered pair $\{\mathcal{S}, \mathcal{T}\}$. One can show that it is the minimal, maximally predictive causal representation of the process [4]. Unlike a general probabilistic ϵ -machine, for simplicity, here we take causal-state transitions to have equal probabilities. The finitary ϵ -machines can be thought of as finite-state machines with a certain properties [4]: (1) All states are start states and accepting states; (2) All recurrent states form a single strongly connected component; (3) All transitions are deterministic: A causal state together with the next value observed from the process determines a unique next causal state; And (4) the set of causal states is minimal. Here we use an alphabet of input and output pairs over a binary alphabet: $\mathcal{A} = \{0|0, 0|1, 1|0, 1|1\}$. This implies that the ϵ -machines work as mappings between sets of strings. In other words, they are transducers [5].

In contrast to prior models of pre-biotic evolution, ϵ -machines are simply finitely-specified mappings. More to the point, they do not have two separate modes of representation (information storage) or functioning (transformation). The advantage is that there is no assumed distinction between gene and protein [6, 7] or between data and program [8–11]. Instead, one recovers the dichotomy by projecting onto (i) the sets that an ϵ -machine recognizes and generates and (ii) the mapping between these sets. Examples of ϵ -machines are shown in Figure 1.

III. INTERACTION: FUNCTIONAL COMPOSITION

The basic pairwise interaction we use in the finitary process soup is functional composition. Two machines interact and produce a third machine—their composition. Composition is not a symmetric operation. Machine T_A composed with another T_B does not necessarily result in the same machine as T_B composed with T_A : $T_B \circ T_A \neq T_A \circ T_B$.

The upper bound on the number of states of the composition is the product of the number of states of the parents: $|T_B \circ T_A| \leq |T_B| \times |T_A|$. Hence, there is the possibility of exponential growth of states and machine complexity if machines are iteratively composed. A composition, though, may also result in a machine with lower complexity than those of its parents.

A. Interaction networks

We represent the interactions among a set of ϵ -machines with an *interaction network* \mathcal{G} which is a graph whose nodes correspond to ϵ -machines and whose transitions correspond to interactions. If $T_k = T_j \circ T_i$ occurs in the soup, then the edge from T_i to T_k is labeled T_j . One may represent \mathcal{G} with the binary matrices:

$$\mathcal{G}_{ij}^{(k)} = \begin{cases} 1 & \text{if } T_k = T_j \circ T_i \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Consider the ϵ -machines in Fig. 1, for example. They are related via composition according to the interaction graph shown in Fig. 2, which is given by the matrices

$$\mathcal{G}^{(A)} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad (5)$$

$$\mathcal{G}^{(B)} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad (6)$$

and

$$\mathcal{G}^{(C)} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}. \quad (7)$$

B. Meta-machines

For a machine to survive in its environment somehow it needs to produce copies of itself. This can be done directly by self-reproduction, e.g. $T_A \circ T_A = T_A$, or it can be done indirectly in cooperation with other machines: e.g., T_A facilitates the production of T_B , which facilitates the production of T_C , which, in turn, closes the loop by facilitating the production of T_A . In other words, there can be sets of machines that interact with each other in such a way that they *collectively* self-reinforce the overall production of the set. This leads to the notion of an autonomous and self-replicating entity, which we call a *meta-machine*. Inspired by Maturana and Varela's *autopoietic set* [12], Eigen and Schuster's *hypercycle* [13], and Fontana and Buss' *organization* [14], we define a meta-machine Ω to be a connected set of ϵ -machines




FIG. 2: Interaction network of the ϵ -machines in Fig. 1. There is a transition, for example, that is labeled T_C from the node T_A to the node T_C , since T_A composed with T_C results in T_C (in fact, each transition from T_A has the same label as the label of its respective sink node since T_A is the identity function).

whose interaction matrix consists of all and only the members of the set. That is, a set Ω is a meta-machine if and only if (1) the composition of two ϵ -machines from the set is always itself a member of the set:

$$T_j \circ T_i \in \Omega, \quad \forall T_i, T_j \in \Omega; \quad (8)$$

(2) all ϵ -machines in the set can arise from the composition of two machines in the set:

$$\exists T_i, T_j \in \Omega, \quad T_k = T_j \circ T_i, \quad \forall T_k \in \Omega; \quad (9)$$

and (3) there is a nondirected path between every pair of nodes in Ω 's interaction network \mathcal{G}_Ω . The third property ensures that there is no subset of Ω that is isolated from the rest of Ω under composition. Consider, for example, the union of two self-replicators, T_A and T_B , for which $T_B \circ T_A = T_A \circ T_B = T_\emptyset$. According to property (3), they are not a meta-machine.

IV. COMPLEXITY MEASURES: C_μ

In previous computational pre-biotic models, the objects have been represented by, for example, assembly language codes [8–11], tags [15, 16], λ -expressions [17] and cellular automata [18]. We employ ϵ -machines instead mainly for one reason: there is a well developed theory (computational mechanics) of their structural properties. Assembly language programs and λ -expressions, for instance, are computational universal representations and so one knows that it is not possible to calculate their complexity [5].

For finitary ϵ -machines, in contrast, complexity can be readily defined and analytically calculated in closed form. Define the stochastic connection matrix of an ϵ -machine $M = \{\mathcal{S}, \mathcal{T}\}$ as $\mathbf{T} \equiv \sum_{s \in \mathcal{A}} T^{(s)}$. The probability distribution $p_{\mathcal{S}}$ over the states in \mathcal{S} —how often they are visited—is given by the normalized left eigenvector of \mathbf{T} associated with eigenvalue 1.




FIG. 3: A schematic illustration of the finitary process soup. Two ϵ -machines, T_A and T_B , are composed and produce a third machine T_C , or a random machine T_R is introduced to the soup. In either case, another randomly selected machine T_D is removed to maintain a fixed population size. Note that this is a well stirred setting, and so there is no spatial relationship in the population.

The *structural complexity* C_μ of M is the Shannon entropy of the distribution given by $p_{\mathcal{S}}$,

$$C_\mu(M) \equiv \sum_{v \in \mathcal{S}} p_{\mathcal{S}}^{(v)} \log_2 p_{\mathcal{S}}^{(v)}. \quad (10)$$

The structural complexity of an ϵ -machine is the amount of information stored in the distribution over \mathcal{S} , which is the minimum average amount of memory needed to optimally predict future configurations [4].

To measure the diversity of interactions in the soup we define the *interaction network complexity* $C_\mu(\mathcal{G})$ to be the Shannon entropy of the distribution of effective transition probabilities in the graph \mathcal{G} . We consider, in particular, only the transitions that have occurred between machine types that are present. That is,

$$C_\mu(\mathcal{G}) \equiv \sum_{p_{i,j,k} \neq 0} v_{ij}^k \log_2 v_{ij}^k, \quad (11)$$

where

$$v_{ij}^k = \begin{cases} p_i p_j / \sum v_{ij}^k & \text{if } T_k = T_j \circ T_i \text{ has occurred} \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

and p_i is the fraction of machines of type i in the population. To monitor the emergence of actual and functional reproduction paths, we consider only those interactions that occurred in the population.

V. FRAMEWORK: THE SOUP

The ϵ -machines interact in a well stirred reactor with the following iterated dynamics:

1. *Production and influx:*

- (a) With probability Φ_{in} generate a random ϵ -machine T_R .
 - (b) With probability $1 - \Phi_{in}$ (*reaction*):
 - i. Select T_A and T_B randomly.
 - ii. Form the composition $T_C = T_B \circ T_A$.
2. *Outflux*:

- (a) Select an ϵ -machine T_D randomly from the population.
- (b) Replace T_D with the ϵ -machine produced in the previous step—either T_C or T_R .

T_R is uniformly sampled from the set of all two-state ϵ -machines in our simulations (see below). This sampling is also used when initializing the population. The insertion of T_R corresponds to an influx while the removal of T_D corresponds to an outflux. The latter keeps the population size constant. See Fig. 3 for a schematic illustration. There is no spatial dependence in this version of the soup as ϵ -machines are sampled uniformly from the population for each replication and removal.

VI. CLOSED POPULATION DYNAMICS

To familiarize ourselves with the model we first examine a simple base case: a soup with no influx that is initialized with machines taken from a finite set which is closed under composition. This case is also intended to work as a bridge between classical population dynamics and the general, constructive dynamics of the finitary process soup. The closure with respect to composition enables us to describe the system's temporal dynamics of ϵ -machine concentrations by a coupled system of ordinary differential equations. In the limit of an infinite soup size, the rate equation of concentration p_k of machine type T_k is given by

$$\dot{p}_k = \psi_k - \Phi_{out} p_k, \quad k = 1, \dots, n, \quad (13)$$

where the conditional production rate ψ_k is the probability that T_k is produced given that two randomly sampled machines are paired:

$$\psi_k = \sum_{i,j=1}^n \alpha_{ij}^k p_i p_j, \quad (14)$$

and α_{ij}^k is a second-order reaction rate constant:

$$\alpha_{ij}^k = \begin{cases} 1 & \text{if } T_k = T_j \circ T_i \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

The outflux Φ_{out} equals the total production rate of the soup—i.e., the probability that a reaction occurs given that two ϵ -machines are paired. It keeps the size of the soup constant:

$$\Phi_{out}(t) = \sum_{i=1}^n \psi_i. \quad (16)$$




FIG. 4: A simple base case: Machine type frequencies of mono-machines as functions of time. $N (= 100,000)$ denotes the population size. Dashed lines: simulation; solid lines: Eq. (17).

Given a soup with no influx, $\Phi_{in} = 0$, that hosts machines which are members of a set that is closed under composition, the distribution dynamics can alternatively be predicted from by its interaction network:

$$\mathbf{p}_t^{(k)} = \mathbf{p}_{t-1} \cdot \mathcal{G}_{ij}^{(k)} \cdot \mathbf{p}_{t-1}^T Z^{-1}, \quad (17)$$

where $\mathbf{p}_t^{(k)}$ is the frequency of ϵ -machine type k at time t and Z^{-1} is a normalization factor. This approximates the soup's elements as updating synchronously.

We illustrate the closed case by initiating the soup with machines that consist of only a single state—*mono-machines*. There are 15 mono-machines, the null (transition-less) transducer is excluded, and they form a closed set under composition. See Fig. 4 for the temporal dynamics of their respective frequencies. Nine machine types remain in the population at equilibrium. They form a meta-machine \mathcal{M} with $C_\mu(\mathcal{M}) = 5.75$ bits. In this case, since $C_\mu(T_i) = 0$ for all mono-machines T_i , the population's structural complexity derives only from its interaction network.

VII. OPEN POPULATION DYNAMICS

We now move on to the general case of a soup with positive influx rate consisting of ϵ -machines of arbitrary size. The soup then constitutes a constructive dynamical system where there is a mutual dependence between its equations of motion and the individuals. Due to the openness, Eqs. (13)-(17) do not necessarily apply. We therefore turn to simulations.

In order to study dynamics that is ruled solely by compositional transformations we first set the influx rate to zero. A fine-grained description of the soup's history on the ϵ -machine level is given by a genealogy—a record of descent of machine types. By studying the example in




FIG. 5: Genealogy of ϵ -machine types in a soup with 100 machines. A solid line denotes that a machine type is present in the soup. Dashed lines (drawn from the parents to the child) denote composition. Note that almost the whole set of initial ϵ -machine types (with one exception) is replaced by the dynamics.

Fig. 5, a simulation with $N = 100$ individuals, one important observation is that nearly all the ϵ -machine types that are present in the soup's initial population are replaced over time. Thus, genuine novelty emerges, in contrast to the closed soup just described. Initially, there is a rapid innovation phase in which novel machines are introduced that displace the bulk of the initial machines. The degree of innovation flattens out, along with the diversity of the soup, and eventually vanishes as the population becomes increasingly closed under composition.

To monitor the soup's organization over time, we superimpose $C_\mu(\mathcal{G})$ time series from several runs in Fig. 6. One sees that plateaus are formed. These can be explained in terms of meta-machines. In addition to capturing the notion of self-replicating entities, meta-machines also describe an invariant set of the population dynamics. That is, formally,

$$\Omega = \mathcal{G} \circ \Omega, \quad (18)$$

where Ω is the set of ϵ -machines present in the population and \mathcal{G} is their interaction network. These invariant sets can be stable or unstable under the population dynamics.

Consider, for example, the meta-machine in Fig. 2. It is unstable, since T_A s are only produced by T_A s, and will decay over time to the meta-machine of Fig. 7. This also illustrates, by the way, how trivial self-replication is spontaneously attenuated in the soup.

The plateaus at $C_\mu(\mathcal{G}) = 4$ bits, $C_\mu(\mathcal{G}) = 2$ bits, and $C_\mu(\mathcal{G}) = 0$ bits correspond to the largest meta-machine that is present at that time. Since a meta-machine by definition is closed under composition, it itself does not produce novel machines; thus, one has the upper bound of $C_\mu(\mathcal{G})$. As a meta-machine is reduced due to an internal instability or sampling fluctuations by the outflux,




FIG. 6: Decomposition of meta-machines in a soup with no influx. Superimposed plots of $C_\mu(\mathcal{G})$ from 15 separate runs with $N = 500$. $C_\mu(\mathcal{G})$ is bounded by 4 bits while a 4-element meta-machine (shown), denoted Ω_4 , is the largest one in the soup. Ω_4 decays to Ω_2 , a 2-element meta-machine (shown) due to fluctuations, that in turn decays to Ω_1 , a single self-replicating ϵ -machine.




FIG. 7: The resulting meta-machine when the meta-machine in Fig. 2 decays under the population dynamics of Eq. (17).

the upper bound of $C_\mu(\mathcal{G})$ is lowered. This results in a stepwise and irreversible succession of meta-machine decompositions. Fig. 6 shows only three plateaus. In principle, however, there is one plateau for every meta-machine that at some point is the largest one in the population. The diagram in Fig. 8 summarizes our results from a more systematic survey of spontaneously generated meta-machine hierarchies in simulations of soups with 500 ϵ -machines.

We now examine the effects of influx by studying the population-averaged ϵ -machine complexity $\langle C_\mu(T) \rangle$ and the run-averaged interaction network complexity $\langle C_\mu(\mathcal{G}) \rangle$ as a function of t and Φ_{in} , see Fig. 9.

The average ϵ -machine complexity $\langle C_\mu(T) \rangle$ increases rapidly initially before declining to a steady state. The average interaction network complexity $\langle C_\mu(\mathcal{G}) \rangle$ is relatively high where the average structural complexity of the ϵ -machines is low, and is maximized at $\Phi_{in} \approx 0.1$. Higher influx rates have a destructive effect on the populations' interaction network due to the new individuals' low reproduction rate. $\langle C_\mu(T) \rangle$ is, in contrast, maximized at a relatively high influx rate ($\Phi_{in} \approx 0.75$) at which $\langle C_\mu(\mathcal{G}) \rangle$ is relatively small. The maximum network complexity $\widehat{C}_\mu(\mathcal{G})$ of the population grows linearly at a positive rate of approximately $7.6 \cdot 10^{-4}$ bits per replication.




FIG. 8: Composition and decomposition hierarchy of meta-machines. Dots denote self-replicating ϵ -machines, solid lines denote $T_A \xrightarrow{T_B} T_C$ transitions and dashed lines denote equivalent $T_B \xrightarrow{T_A} T_C$ transitions. The label of the source node and the transition label are interchanged in the latter transition type. This results in a redundant representation of the interaction network, which is used to show how the meta-machines are related. The interaction networks are shown in a simplified way according to Ω_4 ; cf. Fig. 6.

VIII. DISCUSSION

We presented a conceptual model of pre-biotic evolution: a soup consisting of objects that make new objects [1]. The objects are ϵ -machines and they generate new ϵ -machines by functional composition. The soup constitutes a constructive dynamical system since the population dynamics is not fixed and may itself evolve along with the state space it operates on. Specifically, the dimension of the state space changes over time, which is reminiscent of the constructive population dynamics associated with punctuated equilibria [19].

In principle, this allows for open-ended evolution. The quantitative estimate quoted above for the linear growth of the interaction network complexity supports this in-

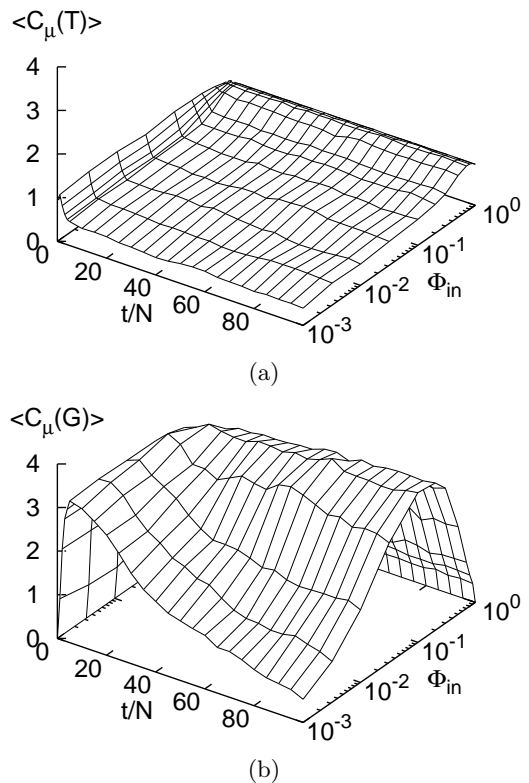


FIG. 9: (a) Population- and run-averaged ϵ -machine complexity $\langle C_\mu(T) \rangle$ and (b) run-averaged interaction network complexity $\langle C_\mu(G) \rangle$ as a function of time t and influx rate Φ_{in} for a population of $N = 100$ objects. (Reprinted with permission from [1]).

triguing possibility occurring in the open finitary process soup. In the case of no influx, though, the system reaches a steady state where the soup consists of only one self-replicator. Growth and maintenance of organizational complexity requires that the system is dissipative; i.e., that there is a small, but steady inflow of random ϵ -machines. Notably, in this case, the soup spontaneously evolves hierarchical organizations in the population—meta-machines that in turn are organized hierarchically.

These hierarchies are assembled from noncomplex, general individual ϵ -machines. In this way, the soup's emergent complexity derives largely from a network of interactions, rather than from the unbounded increase in the structural complexity of individuals. It appears, therefore, that higher-order complex organization not only allows for simple local components but, in fact, requires them.

Acknowledgments

This work was supported at the Santa Fe Institute under the Networks Dynamics Program funded by the

Intel Corporation and under the Computation, Dynamics and Inference Program via SFI's core grants from the National Science and MacArthur Foundations. Direct support was provided by NSF grants DMR-9820816 and PHY-9910217 and DARPA Agreement F30602-00-2-0583. O.G. was partially funded by PACE (Pro-

grammable Artificial Cell Evolution), a European Integrated Project in the EU FP6-IST-FET Complex Systems Initiative, and by EMBIO (Emergent Organisation in Complex Biomolecular Systems), a European Project in the EU FP6-NEST Initiative.

-
- [1] J. P. Crutchfield and O. Görnerup. Objects that make objects: The population dynamics of structural complexity. *J. Roy. Soc. Interface*, 3:345–349, 2006.
- [2] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Lett.*, 63:105–108, 1989.
- [3] J. P. Crutchfield. The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11 – 54, 1994.
- [4] J. P. Crutchfield and C. R. Shalizi. Thermodynamic depth of causal states: Objective complexity via minimal representations. *Physical Review E*, 59(1):275–283, 1999.
- [5] J. G. Brookshear. *Theory of computation: formal languages, automata, and complexity*. Benjamin/Cummings, Redwood City, California, 1989.
- [6] E. Schrödinger. *What is Life? Mind and Matter*. Cambridge Univ. Press, Cambridge, United Kingdom, 1967.
- [7] J. von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, 1966.
- [8] S. Rasmussen, C. Knudsen, P. Feldberg, and M. Hindsholm. The Coreworld: Emergence and evolution of cooperative structures in a computational chemistry. In *Emergent Computation*, pages 111–134. North-Holland Publishing Co., 1990.
- [9] S. Rasmussen, C. Knudsen, and R. Feldberg. Dynamics of programmable matter. In *Artificial Life II: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems (Santa Fe Institute Studies in the Sciences of Complexity, Vol. 10)*. Addison-Wesley, 1992.
- [10] T. S. Ray. An approach to the synthesis of life. In C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, volume XI of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 371–408, Redwood City, California, 1991. Addison-Wesley.
- [11] C. Adami and C. T. Brown. Evolutionary learning in the 2D artificial life system ‘Avida’. In *Artificial Life 4*, pages 377–381. MIT Press, 1994.
- [12] F. J. Varela, H. R. Maturana, and R. Uribe. Autopoiesis: The organization of living systems. *BioSystems*, 5(4):187–196, 1974.
- [13] P. Schuster. The hypercycle: A principle of natural self-organization. *Naturwissenschaften*, 64:541–565, 1977.
- [14] W. Fontana and L. W. Buss. The barrier of objects: From dynamical systems to bounded organizations’. In S. Casti and A. Karlqvist, editors, *Boundaries and Barriers*, pages 56–116. Addison-Wesley, 1996.
- [15] J. D. Farmer, N. H. Packard, and A. S. Perelson. The immune system, adaptation, and machine learning. *Phys. D*, 2(1-3):187–204, 1986.
- [16] R. J. Bagley, J. D. Farmer, S. A. Kauffman, N. H. Packard, A. S. Perelson, and I. M. Stadnyk. Modeling adaptive biological systems. *Biosystems*, 23:113–138, 1989.
- [17] W. Fontana. Algorithmic chemistry. In C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, volume XI of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 159–209, Redwood City, California, 1991. Addison-Wesley.
- [18] J. P. Crutchfield and M. Mitchell. The evolution of emergent computation. *Proc. Natl. Acad. Sci.*, 92:10742–10746, 1995.
- [19] James P. Crutchfield. When evolution is revolution—origins of innovation. In *Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function*, Santa Fe Institute Series in the Sciences of Complexity, pages 101–133. Oxford University Press, 2001.