Santa Fe Institute Working Paper 04-06-020
arxiv.org e-print adap-org/0406058

Objects That Make Objects: The Population Dynamics of Structural Complexity

James P. Crutchfield[1, *] and Olof Görnerup[1, 2, †]
[1]Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501
[2]Department of Physical Resource Theory, Chalmers University of Technology & Göteborg University, 412 96 Göteborg, Sweden

(Dated: July 6, 2004)

To analyze the evolutionary emergence of structural complexity in physical processes we introduce a general, but tractable, model of objects that interact to produce new objects. Since the objects—ϵ-*machines*—have well defined structural properties, we demonstrate that complexity in the resulting population dynamical system emerges on several distinct organizational scales during evolution—from individuals to nested levels of mutually self-sustaining interaction. The evolution to increased organization is dominated by the spontaneous creation of structural hierarchies and this, in turn, is facilitated by the innovation and maintenance of relatively low-complexity, but general individuals.

PACS numbers: 87.23.Kg, 89.75.Fb, 05.65.+b, 87.23.Cc
Keywords: complexity; evolution; population dynamics; entropy; structure, hierarchy, ϵ-machine, self-organization; computational mechanics, autopoiesis, autocatalysis

Long before the distinction between genetic information and functional molecules—molecules encoded by, but manipulating that stored information—there were objects that simply interacted and mutually transformed each other. How did structured objects with the dual roles of information storage and transformation emerge in such initially disorganized environments? Here we introduce a class of models that allows us to explore this question, in a setting often referred to as *pre-biotic* evolution [1]. In contrast with prior work, we focus on the questions of what *levels* of structure and information processing can emerge and, specifically, what population-dynamical mechanisms drive the transition from pre-biotic to biotic organization.

One of the key puzzles in this is to understand how systems, on the one hand, produce structure that, on the other, becomes substrate for future functioning and innovation. The spirit of our approach to this puzzle follows that suggested by Schrödinger [2] and found in von Neumann's random self-assembly model [3]. However, our model is more physical than chemical in the sense that we do not assume the existence of sophisticated chemical entities, such as macromolecules, nor do we even use chemical metaphors, such as information being stored in one-dimensional arrays—in *aperiodic crystals*, as Schrödinger presciently proposed. While ultimately interested in pre-biotic organization and its emergence, our focus is on what one might call *pre-chemical* evolution. As such, the following provides a first step to directly address how structural complexity and evolutionary population dynamics interact [4].

Here we introduce a model of the emergence of organization and investigate it in a setting that, at one and the same time, provides a well defined and quantitative notion of structure and is mathematically tractable. The goal is to develop predictive theories of the population dynamics of interacting, structured individuals and their collective organization. With a well defined measure of structural complexity one can precisely state the question of whether or not complexity has genuinely emerged over time in pre-biotic and pre-chemical processes. Additionally, with a predictive theory of the population dynamics one can identify and analyze the (at some point evolutionary) mechanisms that lead to such increases (and to decreases) in structural complexity.

We create a well stirred population—the *finitary process soup*—of initially random, finite ϵ-machines [5] that interact and transform each other, making new ϵ-machines of differing structure and so of differing transformational properties. The initial random soup serves as a reference that, for reasons to become apparent below, has "null" structural complexity, both in the individuals (on average) and across the population. Here we consider the case of a *process gas*: objects, ϵ-machines T_A and T_B, are successively randomly paired (*pan-mixia*) and act on each other to create progeny: $T_B + T_A \rightarrow T_C$. No externally applied selection or variation is imposed.

An ϵ-*machine* $T = \{\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{T}}\}$ consists of a set of *causal states* $\boldsymbol{\mathcal{S}}$ and *transitions* $\boldsymbol{\mathcal{T}}$ between them: $T_{ij}^{(s)},\ s \in \mathcal{A}$. We interpret the symbols labeling the transitions in the alphabet $\mathcal{A}$ as consisting of two parts: an *input symbol* that determines which transition to take from a state and an *output symbol* which is emitted on taking that transition. ϵ-machines have several key properties [5]: All of their recurrent states form a single *strongly connected* component. Their transitions are *deterministic* in the specific sense that a causal state together with the next input symbol determines the successor state. And, $\boldsymbol{\mathcal{S}}$ is *minimal*: an ϵ-machine is the smallest causal representation of the transformation it implements.

Due to these properties, one can quantify an ϵ-machine's structural complexity. To do this we need the probability distribution over the states in $\boldsymbol{\mathcal{S}}$, how often they are visited, and this is given by the normalized left eigenvector associated with eigenvalue 1 of the stochastic matrix $\mathbf{T} \equiv \sum_{s \in \mathcal{A}} T^{(s)}$. Denote this eigenvector, normalized in probability, by $p_{\boldsymbol{\mathcal{S}}}$. An ϵ-machine's *structural complexity* is the amount of stored information: $C_\mu(T) \equiv -\sum_{v \in \boldsymbol{\mathcal{S}}} p_{\boldsymbol{\mathcal{S}}}^{(v)} \log_2 p_{\boldsymbol{\mathcal{S}}}^{(v)}$. When C_μ is finite, we say the ϵ-machine (or, more properly, the transformation it describes) is *finitary.*

Thus, unlike previous models—such as λ-expressions [6], machine instruction codes [7], tags [8], and cellular automata [9]—ϵ-machines allow one to readily measure the structural complexity and disorder of the transformations they specify. It is well known that algorithms do not even exist to measure these quantities for machine-language programs and λ-expressions, for example, since these are computation-universal models [10]. As our results demonstrate, these tractable aspects of ϵ-machines give important quantitative, interpretive, and theoretical advantages over prior work on the pre-biotic evolutionary emergence of structural complexity. The finitary assumption is also consistent with the recent proposal that gene expression is implemented with finite-memory computational steps [11].

We should emphasize that, in these finitary process soups and in contrast with prior work, ϵ-machines do not have two distinct modes of representation or functioning. The objects are only functions, in the prosaic mathematical sense. Thus, one benefit of this model of pre-biotic evolution is that there is no assumed distinction between gene and protein [2, 3], between data and program [18].

Finitary process soups allow one to quantitatively analyze not only the structural complexity of individuals, but also the interaction between individual structure and population organization and dynamics in terms of how they store and process information and the *causal architecture* that supports these. Since this view of a system *intrinsically computing* applies both to individuals and to the population as a whole, we can identify the locus of a population's structural complexity. Is it largely the sum of the individuals' or largely embodied in the transformative relationships between individuals? Perhaps it derives from some irreducible interaction between these levels.

The finitary process soup differs from early investigations in which finite-state machines were evolved using an explicit mutational operator [12]. Here, novelty derives directly from how the objects themselves are structured, since this determines how they transform each other. Equally important, survivability is determined by an individual's rate of reproduction—the original biological notion of *fitness*; there is no externally imposed fitness function. In this, the process soup is similar to the molecular evolution models of Eigen and Schuster [13].

A *population* P is a set of N *individuals*, each of which is an ϵ-machine. More compactly, one can also describe the population as a distribution of ϵ-machine types: $\mathbf{f} = (a_1/N, a_2/N, \ldots, a_n/N)$, where n is the number of possible ϵ-machine types and a_i is the number of individuals of type T_i. A single replication is determined through compositions and replacements in a two-step sequence: First, construct ϵ-machine T_C by:

1. With probability $1 - \Phi_{\text{in}}$, forming the composition $T_C = T_B \circ T_A$ from T_A and T_B randomly selected from the population and minimizing [10].

2. With probability Φ_{in}, generating a random T_C.

Second, replace a randomly selected ϵ-machine, T_D, with T_C. Φ_{in} is the rate of influx of new (random) ϵ-machines. When $\Phi_{\text{in}} = 0$, the soup is a closed system. When $\Phi_{\text{in}} = 1$, the soup is open, but consists of entirely random ϵ-machines and so is unstructured. The initial population P_0—with $\mathbf{f} = (1/n, \ldots, 1/n)$—is similarly unstructured.

As a first step to detect population structure we define the *interaction network* $\mathcal{G}$ as the ϵ-machine compositions that have occurred in the population. For a population with n different types, $\mathcal{G}$ is described by an $n \times n$ matrix the entries of which are the machine types returned by the compositions $T_i \circ T_j,\ T_i, T_j \in P$. We represent $\mathcal{G}$ as a graph whose nodes are the machine types in the population and whose directed edges connect one node, say T_A, to another, T_C, when $T_C = T_B \circ T_A$. The edges are labeled with the transforming machine T_B. We also represent $\mathcal{G}$ as a transition matrix $\mathcal{G}_{ij}^{(k)} = \mathrm{P}(T_k | T_i, T_j)$, when $T_k = T_j \circ T_i$.

The second step is to introduce a natural notion of organization that encompasses interaction and dynamic stability, we define a *meta-machine* as a set of ϵ-machines that is both closed and self-maintained under composition. That is, $\Omega \subset P$ is a meta-machine if and only if (i) $T_i \circ T_j \in \Omega$, for all $T_i, T_j \in \Omega$ and (ii) for all $T_k \in \Omega$, there exists $T_i, T_j \in \Omega$, such that $T_k = T_i \circ T_j$. This definition of self-maintenance captures Maturana et al's *autopoiesis* [14], Eigen and Schuster's *hypercycles* [13], and *autocatalytic sets* [6, 15]. In a process soup, awash in fluctuations and change, a meta-machine is a type of organization that can be regarded as an autonomous and self-replicating entity. Note that, in this sense, the initial random soup P_0 is not organized. To the extent that interaction networks persist, they are meta-machines.

To measure the diversity of interactions in a population we define the *interaction network complexity* $C_\mu(\mathcal{G}) = -\sum_{f_i, f_j, f_k > 0} v_{ij}^k \log_2 v_{ij}^k$, where

$$v_{ij}^k = \begin{cases} f_i f_j / \sum v_{ij}^k, & T_k = T_i \circ T_j \text{ has occurred,} \\ 0, & \text{otherwise.} \end{cases} \qquad (1)$$

Finally, a machine type's frequency changes at each generation according to its interactions and is given by

$$\mathbf{f}_t^{(k)} = \mathbf{f}_{t-1} \cdot \mathcal{G}_{ij}^{(k)} \cdot \mathbf{f}_{t-1} \Big/ \sum_{k=1}^{n} \mathbf{f}_{t-1} \cdot \mathcal{G}_{ij}^{(k)} \cdot \mathbf{f}_{t-1} \; . \qquad (2)$$

Let us now explore a base case: the population dynamics of one simple subset of ϵ-machines, those consisting of only a single state. This class is especially instructive since it is closed under composition: the composition of two single-state machines is itself a single-state machine. There are 15 single-state ϵ-machines; excluding the null machine. As a consequence, there is a finite number of possible interactions and this, in turn, greatly facilitates an initial analysis. Although a seemingly trivial case, a population of these machines exhibits nontrivial dynamics and leads to several insights about unrestricted populations.

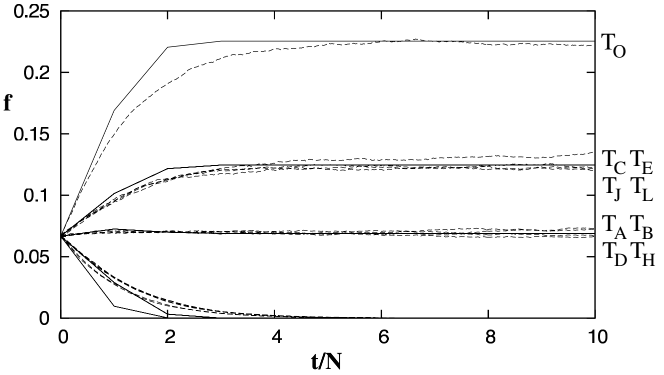

FIG. 1: Population dynamics starting with $N = 10^5$ randomly sampled single-state ϵ-machines: Fraction $\mathbf{f}_t$ of ϵ-machine types as a function of time t (number of replications). Simulations (dashed lines) and theory Eq. (2) (solid lines).

The first obvious evolutionary pressure driving the system is that governed by trivial self-reproduction (copying). ϵ-Machines with the ability to copy themselves (directly or indirectly) are favored, possibly at other ϵ-machines' expense or in symbiosis with other ϵ-machines. The number of such self-reproducing machines grows in relation to the whole population, further increasing the probability of self-reproduction. Interaction networks that sustain this will emerge, consisting of cycles of co-operatively reproducing ϵ-machines. These are chains of composed mappings that form closed loops.

Figure 1 shows the dynamics of a population sampled from $\mathbf{f}_0 = (1/15, \ldots, 1/15)$ and in a closed system ($\Phi_{\text{in}} = 0$) of $N = 10^5$ ϵ-machines. The figure shows that Eq. (2) predicts the simulations quite well. Out of the interactions between all possible ϵ-machines the population settles down to a steady-state interaction network of nine ϵ-machines. Figure 2 shows this meta-machine. Note that the structural complexity of individual ϵ-machines is always zero: $C_\mu(T) = 0$ for single-state machines. Thus, the population's structural complexity is due solely to that coming from the network of interactions.

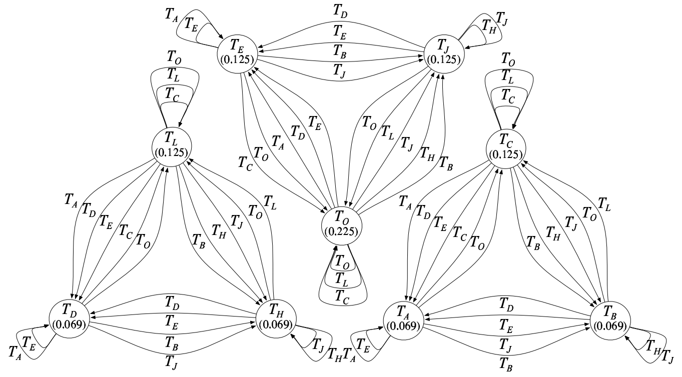

FIG. 2: Organization of the steady-state population—a meta-machine: The interaction graph $\mathcal{G}$ after 10^6 replications, at the end of Fig. 1. The graph shows only $T_A \xrightarrow{T_B} T_C$-denoted interactions. $C_\mu(\mathcal{G}) \approx 5.75$ bits.

With ways to predict the population dynamics and to detect the emergence of structural complexity in the soup, we now turn to the evolution of unrestricted populations. We summarize the results using 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 time and influx rate; see Fig. 3. One observes an initial rapid construction of increasingly complex individuals and interaction networks. In the closed system ($\Phi_{\text{in}} = 0$), both of these reach a maximum and then decline to less complex steady states within a small subspace of possible structures. In fact, both structural complexities effectively vanish at this extreme. The closed system specializes, ages, and eventually dies away. At the extreme of high influx ($\Phi_{\text{in}} = 1$), when the population looses the ability to store information, the network complexity vanishes and the individual complexity becomes that of a purely random sample of ϵ-machines.

Away from these extremes, the evolution of the open systems' network complexity is maximized at an intermediate influx rate $\Phi_{\text{in}} \approx 0.10$. Notably, the emergence of complex organizations occurs where individual ϵ-machine complexity is small. Survival, however, requires these individuals to participate in interaction networks and so to interact with a variety of other machines; they are generalists in this sense. At higher influx ($\Phi_{\text{in}} \approx 0.75$) large $\langle C_\mu(T) \rangle$ is correlated with markedly less complex networks. These more complex machines are specialized and do not support robust complex interaction networks.

It turns out that the maximum network complexity $\widehat{C_\mu}(\mathcal{G})$ grows slowly (linearly) with time. It is ultimately capped by the population size since there is only so much structure that can be built with a finite number of components. More extensive investigations show that it grows in an unbounded way—$\widehat{C_\mu}(\mathcal{G}) \propto \log N$—indicating the possibility of reaching highly structured populations at large sizes.

The finitary process soup demonstrates (i) that complexity of the entire system arises mainly from the trans-

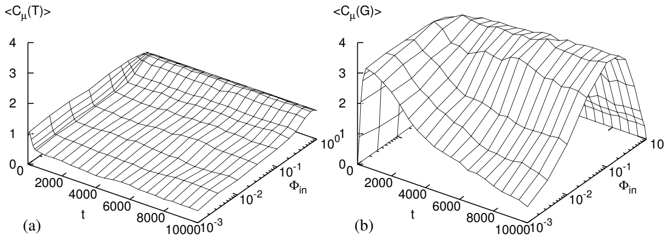

FIG. 3: (a) Population-averaged ϵ-machine complexity $\langle C_\mu(T) \rangle$ [bits] and (b) run-averaged interaction network complexity $\langle C_\mu(\mathcal{G}) \rangle$ [bits] versus time t and influx rate $\Phi_{\rm in}$ for a population of size $N = 100$ averaged over 50 runs at each $\Phi_{\rm in}$.

formative relationships between individuals and (ii) that those individuals tend to be noncomplex and to implement general rather than specialized local functions. Thus, the population dynamics makes a trade-off: simpler individuals facilitate the emergence of global structure. Conversely, for a system to become complex, it is not necessary to evolve complex individuals. The results strongly suggest that replicative processes will use this particular strategy to build successively higher levels of structural complexity from the compositional ("metabolic") network of interacting finitary components. In this way, the finitary process soup evolves higher computational representations, they are not built-in or accessible at the outset.

The finitary process soup is a model of endogenous evolution. In particular, fitness is defined and measured in the biologically plausible way, as the rate of individual reproduction, and there is no externally imposed mutational operator. It is also a flexible model; we showed pan-mixia replication and will report on populations with spatial structure elsewhere. Moreover, it is extensible in a number of ways. Specifically, it is straightforward to couple in energetic and material costs of composition. This will allow one to analyze trade-offs between energetics, dynamics, and organization. Finally, one of the key determinants of evolutionary dynamics is the structure of selection-neutral genotype networks; see Ref. [4] and references therein. The neutral networks of ϵ-machines and of ϵ-machine networks can be directly probed.

We close with a few general comments. Given that organization in a population becomes hierarchical, we believe that powerful computational representations, when employed as the *basic* objects, are neither effectively used by nor necessary for natural evolutionary processes to produce complex organisms. We hypothesize that individuals with finitary computational capacity are appropriate models of molecular entities and transformations. From these, more sophisticated organizations and functions can be hierarchically assembled. This can only be tested by experiment, of course, but this will soon be possible.

It has been recently estimated that the genomes of higher species consist of a surprisingly small number of genes compared to the number found in lower species [16, 17], despite the higher species being markedly more complex and diverse in their behaviors. Moreover, many of those genes serve to maintain elementary functions and are shared across species. These observations accord with the evolutionary dynamics of the finitary process soup: global complexity is due to the emergence of higher level structures and this in turn is facilitated by the discovery and maintenance of relatively noncomplex, but general objects. In both the genomic and finitary soup cases, one concludes that an evolving system's sophistication, complexity, and functional diversity derive from its hierarchical organization.

This work is supported by Intel Corporation, core grants from the National Science and MacArthur Foundations, and DARPA Agreement F30602-00-2-0583. OG was partially supported by the International Masters Programme in Complex Adaptive Systems.

---

[*] Electronic address: `chaos@santafe.edu`
[†] Electronic address: `olof@santafe.edu`

[1] S. Rasmussen et al. *Science*, 303:963–965, 2004.
[2] E. Schrodinger. *What is Life? and Mind and Matter.* Cambridge University Press, Cambridge, UK, 1967.
[3] J. von Neumann. *Theory of Self-Reproducing Automata.* University of Illinois Press, Urbana, 1966.
[4] J. P. Crutchfield. In J. P. Crutchfield and P. K. Schuster, editors, *Evolutionary Dynamics*, pages 101–134. Oxford University Press, 2001.
[5] J. P. Crutchfield and K. Young. *Phys. Rev. Let.*, 63:105–108, 1989.
[6] W. Fontana. In C. Langton et al, editor, *Artificial Life II*, pages 159–209, Redwood City, 1991. Addison-Wesley.
[7] T. S. Ray. In C. Langton et al, editor, *Artificial Life II*, pages 371–408, Redwood City, 1991. Addison-Wesley.
[8] R. J. Bagley et al. *Biosystems*, 23:113–138, 1989.
[9] J. P. Crutchfield and M. Mitchell. *Proc. Natl. Acad. Sci.*, 92:10742–10746, 1995.
[10] J. G. Brookshear. *Theory of computation.* Benjamin/Cummings, Redwood City, California, 1989.
[11] Y. Benenson et al. *Nature*, 429:423–429, 2004.
[12] L. J. Fogel, A. J. Owens, and M. J. Walsh. *Artificial Intelligence through Simulated Evolution.* Wiley, New York, 1966.
[13] P. K. Schuster. *Naturwissenshaften*, 64:541–565, 1977.
[14] F. J. Varela, H. R. Maturana, and R. Uribe. *BioSystems*, 5(4):187–196, 1974.
[15] S. A. Kauffman. *J. Theo. Bio.*, 119:1–24, 1986.
[16] M. Lynch and J. S. Conery. *Science*, 302:1401–1404, 2003.
[17] Rat Genome Sequencing Project Consortium. *Nature*, 428:493–521, 2004.
[18] One recovers the dichotomy by projecting onto (i) the sets that an ϵ-machine recognizes and generates and (ii) the mapping between these sets [10].