Beyond Digital Naturalism

Walter Fontana[a], Günter Wagner[b] and Leo W. Buss[b,c]

[a]Santa Fe Institute
1660 Old Pecos Trail
Santa Fe, NM 87501 USA

[b]Department of Biology
and
[c]Department of Geology and Geophysics
Yale University
New Haven, CT 06511 USA

walter@sfi.santafe.edu
gpwag@yalevm.ycc.yale.edu
Leo_Buss@quickmail.cis.yale.edu

Artificial Life, **1** & **2**, 1994, in press

Abstract

The success of Artificial Life depends on whether it will help solving the conceptual problems of biology. Biology may be viewed as the science of the transformation of organizations. And, yet, biology lacks a theory of organization. We use this as an example of the challenge that Artificial Life must meet.

> "If - as I believe - physics and chemistry are conceptually inadequate as a theoretical framework for biology, it is because they lack the concept of function, and hence that of organization. [...] [P]erhaps, therefore, we should give the [...] computer scientists more of a say in the formulation of Theoretical Biology."
>
> – *Christopher Longuet-Higgins*, 1969 [29]

1 Life and the organization problem in biology

There are two readings of "life": "life" as an embodied phenomenon and "life" as a concept. Foucault [20] points out that up to the end of the eighteenth century life does not exist: only living beings. Living beings are but a class in the series of all things in the world. To speak of life is to speak only in the taxonomic sense of the word. Natural history dominated the Classical age and is foremost a naming exercise. "The naturalist is the man concerned with the structure of the visible world and its denomination according to characters. Not with life." (p.161). There is a concurrent interest in how things work, but it is an interest that remains disconnected and in tension with the naturalist tradition.

In the early nineteenth century natural history makes a decisive step towards a biology when the notion of character becomes subordinate to the notion of function, when classification becomes comparative anatomy. Life is conceptualized as something functionally organized, and organization is foreign to the domain of the visible. A character is weighted according to the importance of the function(s) it is linked to. In contrast to the Classical age, characters are seen as signs of an invisible deep structure. Causal argumentation is reversed: a character is not important because it occurs frequently, but rather it occurs frequently because it is functionally important. Life as a concept becomes manifest as organization, organic structure, that is: organism. In the Classical period living beings were perceived largely as points in a coordinate system of names. Now they require an additional "space of organizational structure". This sets the stage for considering the problem of the modification and transformation of organization.

Darwin posited evolution as an effect of what basically amounts to be a force [41]: natural selection. Natural selection is a statement about kinetics: in a population, those variants of organisms will accumulate which are better able to survive and reproduce than others. If there is on-going variation and if variation is (at least partially) heritable, then the continuous operation of selection kinetics will lead to the modification of living organizations. One would like, however, to understand how organization arises in the first place. Darwin's theory is not intended to answer this. Indeed, this is apparent upon inspection of the *formal* structure of

the theory. Neo-Darwinism is about the dynamics of alleles within populations, as determined by mutation, selection and drift. A theory based on the dynamics of alleles, individuals, and populations must necessarily assume the prior existence of these entities. Selection cannot set in until there are entities to select. Selection has no generative power; it merely dispenses with the "unfit", thus identifying the *kinetic* aspect of an evolutionary process. The principle problem in evolution is one of *construction*: to understand how the organizations upon which the process of natural selection is based arise, and to understand how mutation can give rise to organizational, that is: phenotypic, novelty. A solution to this problem will allow one to distinguish between those features of organizations that are necessary and those that are coincidental. Such an endeavour requires a theory of organization. And, yet, biology lacks a theory of organization. The need for a conceptual framework for the study of organization lies at the heart of unsolved problems in both ontogeny and phylogeny. Can Artificial Life (ALife) illuminate biology?

2 Replicator equations without replicators

One way of viewing Darwin's theory is to consider "fitness" (short for "the ability to survive and reproduce") to be an undefined term, in analogy to an axiomatic structure [50, 51]. The principle is applicable whenever its conditions are met: autocatalytic growth kinetics, variation, heritability. At the same time it leaves open what the subject entities are. They may be molecules, genes, cells, organisms, but also populations, strategies, or even artifacts - it depends on the question one is asking. What is required, however, is a coherent development of "fitness" at the chosen level of description. While Darwin clearly had in mind the individual organism, it is not difficult to see how to abstract from Darwin's theory a structure capable of different behaviors that some might even object are "Darwinian".

Darwin's kinetic theory allows for multiple models in which fitness and its referent are specified. This prompted a debate about whether there is a fundamental referent and about what it might be: either the gene, or the individual, or the group, or the species, you-name-it. There are advocates of a single unit of selection, and others who claim that a description in terms of multiple units is essential [7].

Among the clearest and sharpest proponents of a single unit view is Richard Dawkins with his notion of *replicator selection*. The "fundamental level of selection", Dawkins maintains, is "among replicators - single genes or fragments of genetic material which behave like long-lived units in the gene pool" [11]. More generally, a replicator is defined "as any entity in the universe which interacts with its world, including other replicators, in such a way that *copies of itself are made*." [11] (Our emphasis.) We next exemplify how entities can behave kinet-

ically like replicators without being replicators, leaving room for more than one "fundamental" level of selection.

Consider a generic balance equation for the concentration n_i of an object i in an unconstrained population of n object species: $\mathrm{d}n_i/\mathrm{d}t \equiv \dot{n}_i = \Gamma_i,\ i = 1, \ldots, n$, where Γ_i describes the net growth of object species i. It is convenient to switch to internal coordinates or relative frequencies, $x_i = n_i / \sum_j n_j,\ 0 \leq x_i \leq 1$ and $\sum_i x_i = 1$, in which the balance equation becomes:

$$\dot{x}_i = \Gamma_i - x_i \sum_j \Gamma_j, \quad i = 1, \ldots, n. \tag{1}$$

This can also be viewed as the equation of a flow reactor where objects i are produced and a proportional dilution flow compensates for the excess production in the system at any time. If all Γ_i are just constants, $\Gamma_i = a_i > 0$, then the stationary state of the system will simply contain all object species *sorted* according to their relative magnitude of growth: $\bar{x}_i = a_i / \sum_j a_j$.

The situation changes with autocatalysis, $\Gamma_i = b_i x_i$, as it occurs when objects are replicated (asexually), that is: are copied,

$$\dot{x}_i = x_i \left(b_i - \sum_j b_j x_j \right) = x_i \left(b_i - \langle b \rangle \right), \quad i = 1, \ldots, n. \tag{2}$$

This is one description of *selection*. The effect of (2) is competition, as can be seen from the stationary state which consists of only the object i with the largest b_i. b_i is the "fitness" of object i, and the evaluation of fitness which is implict in autocatalysis is nicely expressed by (2): at any time the growth term b_i is compared against the average $\langle b \rangle$. If it is below (above), the net effect will be a negative (positive) growth rate of species i. As a result the average will shift towards higher values until it matches the maximum b_i: the survivor. This is a choice mechanism which is quite different from sorting.

The story gets an additional twist when fitness is frequency dependent, that is: when the (asexual) reproduction of i depends on the composition of the population at any time: $\Gamma_i = x_i \sum_j c_{ij} x_j$. (1) then becomes:

$$\dot{x}_i = x_i \left(\sum_j c_{ij} x_j - \sum_{r,s} c_{rs} x_r x_s \right) = x_i \left(\sum_j c_{ij} x_j - \langle c \rangle \right), \quad i = 1, \ldots, n. \tag{3}$$

The major difference is that (3) can lead to cooperation. Several mutually dependent species may coexist indefinitely and exhibit complicated dynamical behaviors. Selection need not be a naive optimization device.

Variants of both equations, (2) and (3), have been widely studied in the context of, for example, imperfect reproduction (mutation), or genetics with Mendelian

as well as non-Mendelian transmission [16, 10, 12, 13, 44] (for an overview see [22]). Equation (3) has been termed the *replicator equation* [40], and represents the essence of replicator selection: autocatalysis induced by replication (reproduction, copying) causes the composition of a population to shift based on an endogenous comparison against the population average. Of course, what is being chosen in this process is not easy to say when the mutual interdependencies are intricate.

We resume our theme of organization with a simple observation: autocatalytic kinetics can be induced without replication. Suppose that the interaction between two object species j and k does not result in the replication of either j or k, but in the production of a different object species i, as is the case in a chemical reaction. Suppose further (for the sake of simplicity) that objects j and k are *effectively* not used up in the reaction, that is: each time they react we are given back one instance of each (call it "food"). The overall scheme of this "stylized reaction", then, is:

$$j \; + \; k \; \longrightarrow \; i \; + \; j \; + \; k. \tag{4}$$

Notice that this is not autocatalytic in j or k, since they appear on both sides of the reaction equation with the same stoichiometric coefficients. Let the rate constant of this reaction be $d^i_{j,k}$ (possibly zero). Equation (1) then becomes [43]:

$$\dot{x}_i = \sum_j \sum_k d^i_{j,k} x_j x_k - x_i \sum_{r,s,t} d^t_{r,s} x_r x_s, \quad i = 1, 2, \ldots, n. \tag{5}$$

Consider now a set K of object species such that for each $i \in K$ there exists a pair $j, k \in K$ which produces i. Such a set maintains itself, *but it does not copy itself - it only makes more of itself.* Consider now a number of disjoint self-maintaining sets $J, K, L, \ldots$ which interact with one another in such a way that for each reactive pair $k \in K, j \in J$ the product is in K or J (but distinct from k and j). Now we simply rearrange (5) by collecting all individual object species $i \in K$ into a set with relative frequency $x_K = \sum_{k \in K} x_{k \in K}$. Let the frequency of i in its own set be $y_{i \in K} = x_{i \in K} / x_K$. We obtain:

$$\dot{x}_K = x_K \left[\sum_J \mathcal{C}_{KJ}(t) x_J - \langle \mathcal{C}(t) \rangle \right], \tag{6}$$

with the coefficients

$$\mathcal{C}_{KJ}(t) = \sum_{i \in K} \sum_{j \in J} \sum_{k \in K} d^{k \in K}_{i \in K, j \in J} y_{i \in K} y_{j \in J}.$$

Equation (6) has the form of a replicator equation (3). However, the sets K do not replicate, nor do any of their members (by definition). These sets only grow,

and they do so by reflexive catalysis or self-maintenance at the set level. This induces a selection kinetics identical to the replicator case (a similar observation is mentioned in [12]). However, the relevant units are not replicators. Hence we cannot refer to it as "Darwinian selection", since Darwinian selection rests on reproducing entities. It is a generalization of it and we will simply call it *selection*. Hence we have *sorting* when the growth kinetics of entities is not autocatalytic, and *selection* when the growth kinetics is autocatalytic. In the special case where entities are replicators selection becomes *Darwinian selection*. In the case where entities are not replicators, selection can still occur at the level of aggregate entities, but cannot be Darwinian.

This is our first checkpoint towards a firmer notion of organization. An organization is a set of entities that continuously regenerates itself by transformation pathways (the $d^i_{j,k}$) internal to the set [30, 31]. All that is required is sufficient connectivity: the matrix $\tilde{d}_{ik} = \sum_j (d^i_{j,k} + d^i_{k,j})$ must be irreducible. Many researchers noticed the possibility and the significance of self-maintaining sets of chemicals ([8, 12, 39, 24, 25, 36], and probably many more). The disagreement, however, is over the likelihood of such sets given certain kinds of molecules and their importance in shaping the (early) history of life. Simple self-maintaining ensembles have recently been obtained in the laboratory of Pier-Luigi Luisi [1].

It is important to reiterate the necessity of making a *logical* distinction between reproduction and self-maintenance. There have been occasions where these two orthogonal concepts were conflated [15, 24], probably because self-maintenance was considered as an *alternative* to an "origin of life" through primitive replicatory elements. It should be clear so far that self-maintenance has little to do with making *two* individuals out of one through transformations internal to the system. Hence it cannot be an alternative to reproduction. Reproduction does not necessarily require organization either, as exemplified by simple self-replicating molecules [48, 46] or a viral RNA in the presence of its replicase in the test tube [42]. This assumes that we are willing to make a distinction between a molecule or a pair of molecules in complementary association and an organization. Self-maintenance isolates a different aspect of individuality than reproduction. The former requires organization, the latter requires means for its multiple instantiation, as, for example, by compartmentalization through spatial separation, or membrane enclosure, or plain chemical bonds. Clearly, our intuitive notion of life includes both: organization and reproduction. They are conceptually different, and the implementation of one need not be the implementation of the other.

The main point here was to show that non-replicatory but self-maintaining sets do exhibit the kinetics required for selection, though, by definition, not of a Darwinian kind. (Provided one accepts (3) as a proper dynamical formalization of selection.) The integration of replicators and primitive organizations leaves room for multiple units of selection (in the sense of [7]).

Of course, the modification of organizations is hardly identical to the evolution of replicators. What does it mean, if anything, for non-reproducing organizations to "vary"? This question cannot be adequately answered within a description where the micro-entities are atomic structureless units as is the case in the conventional dynamical systems described above. The reason is that the question draws attention to structure-function relationships.

3 Organizations must be constructed

An *extensional* description is roughly one in which the entire universe of relevant objects is given at once, *in extenso*. It is basically a look-up table which may even be so large that nothing can store it. The modern set-theoretic view of a function is of this kind: a function is a collection of ordered pairs (in, out), for example $x^2 = \{(0,0),\ (1,1),\ (2,4),\ (3,9), \ldots\}$. An extensional framework also characterizes traditional dynamical systems. Their definition requires an extensionally given network that specifies which variables couple with which other variables in what ways. For a particular kind of reasoning this is quite useful. It certainly is adequate in setting up the gravitational equations of motion for a system of a few planets, where the relevant knowledge can actually be tabulated.

This framework does not easily fit biology, because the objects denoted by the variables are typically of combinatorial complexity. If equation (2), suitably augmented with mutational terms [12], were to describe the replication, mutation and selection of RNA sequences of length 100, then we would have to specify 10^{60} equations with their corresponding coefficients. The problem is not so much that this situation forces a stochastic description, since only a vanishingly small fraction of all these possibilities can be realized. More fundamentally, such a description is still extensional as long as the relevant properties - the replication rates b_i as determined by the tertiary structure, for example - are not a *function* of the sequences. The major point, the fact, namely, that there *is* some inner logic that connects sequences and their replication rates, is lost. It is precisley this logic which makes the problem an *interesting* one. The question is how that logic structures a population under a given dynamics.

The same holds for ecological modelling in terms of Lotka-Volterra equations, or for game dynamics. In his recent work Lindgren [28] considers agents with an internal structure that determines the strategy they play in a given game. The structure of both opponents i and j, then, allows one to infer the coefficient c_{ij} in (3). As in the case of RNA sequences, a "strategy-grammar" opens up a space of combinatorial complexity and permits with finite means the *endogenous* specification of an infinity of interaction coefficients.

To summarize, in contrast to extensional models, *constructive* models are founded

on objects with a grammatical and, therefore, combinatory internal structure. Note, however, that in the previous examples the internal structure of agents or objects does not affect the *functional* character of their interactions. These always remain copy actions: $i\ (+\ j)\ \longrightarrow\ 2\,i\ (+\ j)$, as is evident from the factorization of (2) or (3). The internal structure only codifies the *strength* of an interaction which is kept fixed in *kind* for all agents. This is precisely what changes in going from (3) to (5). While the $d^i_{j,k}$ may still specify strengths, they in addition require an underlying *logic* that specifies which object i is implied by a pair of objects (j, k).

This suggests a distinction. In the previous examples concerning the replication of RNA sequences or strategies new entities enter a finite population through *mutation*. The cause of a mutation is a *chance* event, meaning that it stands in no relation to its effect. We refer to models in which new agents are constructed in an unspecific (essentially stochastic) fashion as *weakly constructive*. This is to be contrasted with a situation in which the encounter of two agents *implies* a *specific* third one, as in (5). Models of this kind will be termed *strongly constructive*. The prime example of a strongly constructive system is chemistry. A strongly constructive system that contains agent A must cope with the network of its implications. But, then, it also must cope with the implications of the implications. And so on. Organization is here a network which results from convergence to both *relational* (logical) and *kinetic* self-consistency. The logical component induces a structure that is absent in the weakly constructive case. Think of a "knowledge system" where the agents A, B, C, etc., stand for propositions, and where the deterministic construction of new propositions results from "interactions" which we may call rules of inference. The organizational analogy between "consistent systems of belief" and metabolisms is, in our opinion, not completely superficial.

We have reached the next checkpoint towards a firmer notion of organization. An organization is defined in terms of a strongly constructive model: a system of transformations [47, 31]. This distinguishes it from a weakly constructive version of a Lotka-Volterra or a replicator equation (3) describing an ecology of individuals that only replicate and mutate. There is no doubt that an ecological population can be highly organized. However, the next section will clarify that the concept of organization suggested by a strongly constructive system is different in kind. Clearly, in real life, weakly and strongly constructive aspects are entangled. Disentangling them will be an important step in understanding what is necessary and what is contingent in the history of life.

An extensional system of equations, like (5), is useful to capture some dynamical aspects, but useless to capture the constructive nature of organization. If we were to introduce a new object, say ν, into (5), we would have to specify its constructive interactions with the other objects arbitrarily (for example, at random [24, 25]). This, however, eliminates precisely what is interesting about

organization. The constructive aspect is essential for addressing both the origin problem of organizations and the problem of their variation.

4 Organization – *De Arte Combinatoria*[1]

Chemistry gave rise to biology. This is an elementary indication that strongly constructive interactions are fundamental to organization. Chemistry, therefore, informs our attempt of conceiving a formal and transparent model of organization [17, 19, 18].

Physics is about mechanisms. So is computation theory. But the latter has a twist which physics lacks: it is about mechanisms in which things build other things. Such "things" are *processes* and *functions*. As opposed to the clockwork or the steam engine, computation is inherently constructive. Computation need not only be about calculations that are of interest to a mathematically inclined person. What we emphasize here is the aspect of computation as a formal system which enables symbolic structures to build further symbolic structures in a consistent way. A first grip on organization can be obtained by studying the collective phenomena occuring in a dynamical system consisting of many such interacting symbolic structures. We briefly review such a platform. Details can be found in [18].

I: Constructive Part

(I.1) calculus: Our entities are literally *functions* expressed in a canonical syntactical framework known as the λ-calculus [9, 4] in which they can be *applied* to one another yielding new functions.

The grammar of λ-expressions, E, is

$$E \ ::= \ x \mid \lambda x.E \mid (E)E, \tag{7}$$

where x is a variable. Thus, a variable is an atomic expression. There are two expression forming schemes - "combinators" - which define syntax: One, $\lambda x.E$ (termed abstraction), binds a variable in an expression E making it the equivalent of a formal parameter in a procedure, that is: E is intended as a function in x. The other, $(E)E$, (termed application) expresses the notion of a function being applied to an argument - except that there is no syntactical distinction between function and argument. While application and abstraction are purely syntactical operations, they are given an operational meaning through *substitution*:

$$(\lambda x.A)B \longrightarrow A[B/x] \tag{8}$$

[1] Gottfried Wilhelm von Leibniz (1646–1716)

where $A[B/x]$ denotes the textual substitution of all occurrences of x in A with B. (We assume unique names for bound variables, distinct from names of free variables.) The arrow means that the expression on the l.h.s. can be rewritten as the expression on the r.h.s., *thereby only replacing equals for equals*. The process of carrying out all possible substitutions within an expression is termed *reduction*, and the final stable form - if there is one - is unique and is called a *normal form*.

(I.2) normal form: In this model universe every expression is reduced to normal form within preset computational limits. If no normal form is attained, the expression is not allowed.

II: Dynamical Part

λ-calculus is now put in the context of a (stochastic) dynamical system which mimics a constrained flow reactor containing a finite number of "expression-particles".

(II.1) initialize: A system is initialized with N particles. These are randomly generated (usually unique) λ-expressions.

(II.2) interact: Two expressions, A and B, are chosen at random (in this order). Denote this choice by $[A, B]$ and denote the normal form of $(A)B$ by C. Then, the following "reaction"-scheme applies:

$$[A, B] \longrightarrow C + A + B \tag{9}$$

(II.3) boundary conditions: Apply syntactical or functional boundary conditions to C to determine whether C is allowed to enter the system.

(II.4) constant size: If C is added to the system, choose one expression-particle at random and remove it:

$$D \longrightarrow \emptyset \tag{10}$$

This keeps the system constrained at N particles at any time.

(II.5) iterate. Continue with (II.2).

The reader will recognize that this is just the description of a discrete stochastic analogue to equation (5), where the possible $d^C_{A,B}$ are implicitly given by a calculus, in this case λ-calculus:

$$d^C_{A,B} = \begin{cases} 1 & \text{if } (A)B = C \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

In section 5 we discuss the motivation for this approach and its basic assumptions. First, we briefly review some results. The model provides a formalization of our intuitive notion of organization. It generates organizational levels that can be described without reference to the micro processes that give rise to them. Organizational levels beyond a "molecular ecology" (Level 0, see below) emerge even in the absence of Darwinian selection. Details can be found in [18].

Level 0

Level 0 arises with no specialized boundary conditions (II.3). The system becomes dominated by either single self-copying functions or ensembles of mostly hypercyclically [12, 14] coupled copying functions (i.e., functions f with $(f)g = g$ or f, for all g in the system). Thus, equation (5) reduces to the situation described by the replicator framework, equation (3): the reproduction of an object species, i, depends on itself and one (or more) other object species, j. From a purely functional point of view i is a "fixed point" of its interaction with j. Under perturbation, i.e., the introduction of random expressions, Level 0 ensembles frequently reduce to a single self-copying function, i.e., a function f with $(f)f = f$.

Level 1

Level 1 arises under a variety of conditions, all of which involve a restriction on copy actions, up to their complete elimination. The reason for the importance of such an extreme boundary condition is the elimination of Darwinian selection, thereby allowing to assess the necessity of Darwinian selection in the generation of organization. The following features, therefore, need not be the result of *Darwinian* selection.

Under no-copy conditions the set of objects in the system changes until it becomes confined to a subspace of the space of all λ-expressions. This (infinite) set is *invariant* under applicative action and is characterized by three properties:

(1) Grammatical structure. The objects of the subspace possess a specific syntax described by a grammar (beyond just conforming with (7)). Sometimes the objects are grouped into several "families" with distinct syntactical structures. Grammatical closure here defines membership: an object can be unambiguously assigned to a given organization on the basis of its grammatical specification.

(2) Algebraic structure. All relationships of action between objects of the subspace are described by a (small) set of equations. The system frequently admits "coarse-grained" identities that emphasize symmetries and particular roles of objects (like inverse elements, successor functions, neutral elements, etc.). It is important to emphasize that neither the formulation nor the discovery of the laws which define an organization require knowledge of the underlying λ-calculus. An organization has a level of description that

can be considered independently.

(3) Self-maintenance and kinetic persistence. The flow-reactor contains a finite number of objects. Hence only a small subset "carries" the organization under this dynamical system. Typically this subset maintains itself in the sense that every object is produced by at least one interaction involving other objects of the same set. Notice that self-maintenance is not a statement about kinetics, but rather about the constructive relationships within a set of objects. We also have observed borderline cases where the objects change constantly (while remaining confined to their invariant subspace). In all cases, that which is kinetically persistent is the organization as expressed by its grammatical and algebraic description.

These characteristics endow Level 1 organizations with some interesting properties.

Center. An organization has many *generators*, that is: subsets of objects capable of spawning a given organization, if the reactor is intialized only with them. Such sets are typically small. So far Level 1 organizations were observed to have a unique smallest and self-maintaining generator set that constructs the organization *ab initio*. We call it the *center*. The center is typically present in the reactor.

Self-repair. Self-repair is a consequence of self-maintenance, kinetic persistence and the existence of a center. Organizations can tolerate vast amounts of destruction, while retaining the capability to reconstruct themselves.

Extensions. The model universe, in its present form, does not provide for "noisy" interactions (products are determined once the reactants are given). Without functional perturbations a Level 1 organization is a monolithic entity: once attained, all functional evolution stops. The easiest way out is to provide an exogenous source of noise by injecting random objects into an established organization. For a novel object to persist, it must create transformation pathways that sustain it in a kinetically effective manner within the established network of pathways. The modification of organizations is, therefore, highly constrained. However, when a novel object does become established, it imports new syntax elements, thus altering the grammatical structure and the algebraic laws which characterize the organization. This alteration occurs in a typical way: by extension. The unperturbed organization still persists as a "smaller" core organization to which the interactions with the new object have added another "layer". New algebraic laws are added to the existing ones, and the center of the extended organization is extended correspondingly. With constant reactor capacity, an organization

cannot be extended indefinitely, and upon several such extensions one observes a displacement of previous layers. Over several steps this can result in a substantial modification of the original organization.

Level 2

Level 1 organizations can be combined. At first one may expect a competitive situation. This need not be the case, since two organizations can generate novel objects, through cross-interactions, that are not members of either organization. This "cross-talk" consists in a set of objects that does not constitute a Level 1 organization (it is neither self-maintaining, nor grammatically closed under interaction). We refer to this set as the *glue*, since it knits both component organizations kinetically and algebraically together. The entire structure is a Level 2 organization: a meta-organization that contains the original organizations as self-maintaining subalgebras. The glue distinguishes this situation from a plain coexistence. (Indeed, a Level 2 organization is not described by (6) in terms of multiple component organizations - because of the glue.) A Level 2 organization is not as easily obtained as a Level 1 organization, since the kinetic requirements to the glue are severe. Nevertheless, the construction of Level 2 organizations from scratch has been observed: two Level 1 organizations form whose interactions integrate them into a Level 2 organization. Qualitatively the same properties as for the Level 1 case hold in the Level 2 case, although stability is not as pronounced.

Biology

The history of life is a history of the emergence of new organizational grades and their subsequent diversification [7]. A transition in organizational grade occurred when self-replicating molecules gave rise to (or became incorporated within) self-maintaining prokaryotic cells and the origin of procaryotes was itself followed by the emergence of a hierarchical nesting of different prokaryotic lineages to generate multi-genomic eukaryotic cells. The λ-universe mimics transitions seen in the history of life, i.e., the transition from self-replicating molecules to self-maintaining organizations to hierarchical combinations of such organizations. We refer to these organizational grades as Level 0, Level 1, and Level 2, respectively.

Organizations are very robust towards functional perturbations. Perturbing objects are frequently eliminated, otherwise they typically cause extensions to the grammatical and algebraic structure. Harold Morowitz [36] suggests that the core metabolism derives its architecture from the addition of several metabolic "shells" to a basic energy metabolism such as the glycolytic pathway. A shell attaches to another through only a few "entry points". These are pathways that introduce novel syntactical and functional elements, such as nitrogen (with amino acids among its consequences) and dinitrogen heterocycles (with nucleic acids among their consequences).

A Level 2 experiment in which two Level 1 organizations are brought into interaction can be seen as a massive perturbation of either organization. Frequently the interaction does not produce a glue that is sufficient to stably integrate both organizations while maintaining their autonomy. Rather, one organization loses autonomy and becomes a large extension to the other. This is reminiscent of the well-known pattern wherein an intracellular symbiont loses functions redundant with those possessed by the host.

The center of a Level 2 organization is just the sum of the corresponding Level 1 centers. Such a superposition does not hold for the grammatical and algebraic description of the resulting Level 2 organization, because of the extensions required to describe the glue. This recalls the general difference between the behavior of genotypes and phenotypes: genes combine independently, the organizations they spawn do not.

The transition from Level 0 to Level 1 emphasizes the tension between reproduction and organization. Replicatory elements are clearly kinetically favored over self-maintaining but non-reproducing organizations. The biological problem connected with the transition is one of restricting copy actions. The model shows how a new concept of organization arises – beyond the engineering of mutual dependencies between replicators (cf. hypercycle [12]) – when transformation operations are given the chance to attain closure.

Nevertheless, in all this we violate mass conservation, do not consider thermodynamics, assume all rate constants to be equal, do not use up reactants in reactions, have everything react with everything, and so on. *What does such a model really mean?*

5 A functional *perpetuum mobile*

The model explores the consequences of an extremely simple combination of a dynamical system with a calculus. With respect to biology the explanatory power of such a model comes from *defining a level of description*. Here it is given by the mathematical notion of function. The focus is on the expression and construction of functional relationships, rather than on their exact chemical implementation. λ-calculus is the canonical *language* to express such a consistent universe of construction. λ-calculus is clearly no more than a highly "stylized" chemistry. In the biological context it serves the purpose of a high level *specification language*, rather than a full fledged implementation language. The image of "organization" suggested by this abstraction is very simple: a *kinetically self-maintaining algebraic structure*. Universal algebra emphasizes constructions and their equivalences.

Take a set of two objects, A and B, such that $(A)A = B$ and $(B)B = A$ (assume

that cross-collision are not reactive). Clearly, the set $\{A, B\}$ is self-maintaining without containing copy actions. Is this an organization? Our definition suggests it is not, because the system does neither specify a non-trivial grammar, nor a non-trivial algebra. Trivial does not mean simple; it means that there is no compressed description of the set's composition and behavior short of listing it. Our definition suggests that an organization be over an infinite set of objects. Is that reasonable? Yes, because it is not required that this subspace be realized in its entirety. Only a subset of implicitly determined size is required to make the organization's specification kinetically persistent. If the size needed for supporting the organization is too large compared to the reactor's capacity, the organization disappears.

The mechanics of λ-calculus teaches a lesson: *construction* alone, equation (7), is not sufficient; processes must be associated with construction that induce *equivalence*, equation (8). The basic combinators of the chemical grammar are not fundamentally different from those that build λ-expressions: a molecule is either an atom or a combination of molecules. That which makes such a universe capable of constructing networks, hence of organizing, is a consistent way of establishing that different combinations are effectively the "same". In chemistry, combinatory structures - molecules - combine into transition structures that undergo rearrangements. The nature of these rearrangements establishes which molecular combinations are effectively the same in regard to a particular product. Construction and equivalence are the essence of chemistry at the level of description set by the present model. The rest of chemistry is specific to the implementation of construction and equivalence with a given physics. Construction and equivalence are necessary for organization and, hence, for *any* life, artificial or natural. Consequently, the organizational features outlined in the previous section are not coincidental either, since they follow necessarily.

"Function" is frequently used in a colloquial way meaning different things. These meanings are disentangled at our level of description. First, there is the unambiguous mathematical concept of a function. In this sense every object in our model universe *is* a function. Then there is a more "semantic" aspect of function: the function of an organ, the function of a beta-blocker. This notion of function relates to specific roles within the context of a network. As a simple example consider a λ-expression, A, and certain elements B of its domain that are fixed points of A, i.e. $(A)B = B$. If B happens to be present in the reactor together with A, then A makes B a replicator: B assumes a specific kinetic role and A assumes a specific functional meaning for B in this particular context. Suppose A assumes this role for all C in the system. Then A is algebraically a "neutral element", but, again, conceivably only for those C which happen to be present, not for others which may perturb the system. If there is a neutral element, then this creates a "functional niche" for other objects to be "inverses" of one another,

and so on. Because objects *are* functions they construct networks within which they *have* functions. There is no need to assume the latter to be given *a priori*.

A reactor with many interacting λ-expressions is a system of concurrently communicating functions. Application of a λ-expression to another can be viewed as an act of communication - where a function is communicated to another function via a variable associated with a port named λ [6, 34, 5]. Despite equation (9), communication is symmetric. Here the "communication operator" is really the "+" symbol in the chemical reaction scheme. In fact, written as a chemical reaction the interaction scheme (II.2) just reads:

$$\begin{array}{ccc} & \nearrow & (A)B \ + \ A \ + \ B \\ A \ + \ B & & \\ & \searrow & (B)A \ + \ A \ + \ B \end{array} \tag{12}$$

The result of a communication between A and B is a non-deterministic choice between the process $(A)B$ and the process $(B)A$. This is manifest in the fact that for all random encounters between A and B half of the time A comes from the "right" (chosen second) and half of the time from the "left" (chosen first). Reaction, or communication, is clearly commutative in the model. When speaking (somewhat sloppily) of "interaction" in (II.2) we refer to the application event after a choice in (12) has been made[2].

When interpreting the reaction (or kinetic) scheme, it is imperative to remain consistent with the chosen level of abstraction. A kinetic scheme, like (II.2) and (II.4), serves two purposes. First and foremost, it defines a communication event between objects. Second, it summarizes the overall effect of many details whose full specification is inappropriate at the chosen level of abstraction.

The reason for not using up the reactants in (II.2) is, technically speaking, to induce a *recursion*. If no object were removed, the system would generate all possible iterated communications among the initial objects. The flux (II.4) eventually removes those communication products that are not sustained at a sufficient rate within the system. This focusses repeated communication to occur among a kinetically favored set of objects until structural consistency is attained as well (if it is attained). An organization is plainly the fixed point of such a recursion. It recently became apparent to us that in the process algebra literature a similar scheme - called "replication" (no biological connotations) - conveniently replaces recursion [35].

[2]Computer scientists should not confuse the "chemical communication operator" "+" with their choice operator denoted by the same symbol. If we symbolize communication by $\odot$, "choice" by the usual +, "interleaving" by | [6], then (12) reads as $A \odot B = (A)B|A|B \ + \ (B)A|A|B$.

A more "palatable" kinetic scheme would have been "recursion over food". That is: specify a certain number of (random) objects which are designated as "food", $\mathcal{F}$. Food is steadily supplied to the system, while reactants are used up. Hence (12) becomes

$$\begin{array}{ccc} & \nearrow & (A)B \ + \ \mathcal{F} \\ A \ + \ B & & \\ & \searrow & (B)A \ + \ \mathcal{F} \end{array} \tag{13}$$

where $\mathcal{F}$ stands for an object chosen randomly from the food set. Equation (10) would be turned off, since (13) preserves particle number. Under such circumstances the system is "anchored" to the food source, and any organization would have to contain pathways linking it to that source. Our original scheme (12) frees the system from having to maintain pathways to the initial objects. This is somewhat more abstract, but isolates *functional self-maintenance*, while (13) generates structures that need not be functionally autonomous, since they are exogenously maintained through food.

Of course, real organizations must be fed - but the feeding does not define an organization. The present system is clearly open and driven in many ways. In a sense, however, this is besides the point, since the ultimate nature of our model universe is "informational". (Information is not necessarily destroyed by its usage. Although its mechanics - no matter how remote - remains in a spooky analogy to chemistry.) The model suggests a view in which the world of thermodynamics gives rise to another world whose objects are functions (or processes). While there is no *perpetuum mobile* in the former, there is plenty of room for a *perpetuum mobile* in the latter: functional organization.

There is also plenty of room to improve on this model. For example, the model lacks a structured notion of communication which determines who communicates with whom and when. This suggests that formal processes may be more appropriate agents than functions. Formal processes do by and large not compute functions, but rather are characterized by their ability or inability to offer communications at various points in time [21]. This endows them with different notions of equivalence [33]. The particular abstractions from chemistry and the biologically motivated boundary conditions employed thus far informed our modelling platform and, therefore, our understanding of the emergence of functional organization of the Level 1 and Level 2 type. The proper abstractions and boundary conditions that may allow one to attain multicellular differentiating organizations, Level 3 say, or Level 4 (brains?) or ... remain to be captured. While replacing functions by processes may be part of the story, it is evidently an open question where this framework will lead – as is its applicability outside the biololgical domain. In any case,

> We shall first have to find in which way this theory of <processes> can be brought into relationship with <biology>, and what their common elements are. [...] It will then become apparent that there is not only nothing artificial in establishing this relationship but that on the contrary this theory of <communicating processes> is the proper instrument with which to develop a theory of <biological organization>. One would misunderstand the intent of our discussions by interpreting them as merely pointing out an analogy between these two spheres. We hope to establish satisfactorily, after developing a few plausible schematizations, that the typical problems of <biological organization> become strictly identical with the mathematical notions of suitable <communicating processes>.

The quote has been taken from John von Neumann's and Oskar Morgenstern's *Theory of Games and Economic Behavior* [49], except that we took the liberty to replace "theory of games of strategy" with <communicating processes> and "economic behavior" with <biological organization>. The quote characterizes the spirit of the work sketched here quite aptly. However, it is fair to add that von Neumann's and Morgenstern's intentions were not realized. But, then, their "theory of games of strategy" was not even weakly constructive.

6 ALife and Real Life

Our view of ALife relative to biology is perhaps best rendered by an analogy with geometry. For a long time Euclidean geometry was held to be *the* "science" of space. The development of non-Euclidean geometries led to the realization that terms like "line" or "angle" can be given different coherent specifications which result in different models of space. Geometry is not just one form of space; geometry is an axiomatic structure organized around *undefined terms* like "line" or "angle" [26]. In a particular model of space triangles have an inner angle sum of 180 degrees. If this fails to hold at large scales in our universe, it does not invalidate the former space as a model of geometry. Our picture of ALife is that of a variety of specifications and, perhaps, formalizations of different intuitive notions about what it means to be alive. A notion of "artificial life", however, makes sense only if there is an *implementation independent definition of life that informs biology*.

Succeeding in this vision means opposing the superficially informative. There is a risk of ALife becoming predominantly a community of digital Naturalists traversing their Classical age. Ultimately, the success of ALife will depend on the extent to which it succeeds in developing a concept of life that encompasses biology. But to inform biology means not simply imitating it, but understanding it. And

understanding it requires identification of the outstanding unsolved problems in biology and seeking their resolution.

Our work is ultimately motivated by a premise: that there exists a logical deep structure of which carbon chemistry based life is a manifestation. The problem is to discover what it is and what the appropriate mathematical devices are to express it. ALife started by emphasizing processes rather than that which implements them [27]. We believe this emphasis to be necessary, but not sufficient. If the practice of ALife does not contribute to the solution of the logical problem, ALife will fail. For how can we discern whether a construct is a manifestation of life or whether it is an imitation of life? We can't, unless the logical problem is solved. How do we know whether we have solved it? We know when the solution affects the research agenda of biology - the science of the instance that we have not constructed. Should the premise turn out to be wrong, then ALife will fail too. For then ALife becomes ARbitrary. ALife practicioners must avoid the presumption that bedeviled AI - that all the needed concepts were available to be imported from the other disciplines. ALife will find its first proving ground in real biology, in its capacity for catalysing the theoretical maturation of biology.

7 Sources

Similar ideas have certainly occurred to many people. Some contributions we have mentioned, others we acknowledge here.

A decisive step was undertaken by Bagley, Farmer and Kauffman who devised probably one of the first strongly constructive dynamical systems [3, 2]. Their model is based on simplified polynucleotide sequences which interact through complementary base pairing, thereby specifically cleaving and ligating one another. Rokshar, Anderson and Stein proposed a model somewhat similar in spirit [38]. Steen Rasmussen and his group studied self-organization in an ensemble of machine (assembler) code instructions [37]. Interestingly, in their setting self-maintenance is achieved not by construction of new instructions (that set being fixed), but by (weak and strong) construction of new pointers to the locations on which the instructions operate. George Kampis has been weaving an independent thread of thoughts which seems pertinent to the concerns addressed in this paper [23]. We were put on track by discussions with John McCaskill a few years ago. He suggested a model of interacting Turing machines where tapes, standing for stylized polymeric sequences, encode transition tables that read and write other tapes [32, 45]. It was still a long way to the present framework with its implied level of description and a formally more robust notion of organization.

Varela and Maturana were perhaps the first to think extensively about organization in a new way [30, 47, 31]. Their writings are at times not easily penetrable.

However, building a formal model made us understand that many of the issues raised here were foreshadowed by their thinking.

Acknowledgements

Thanks to Chris Langton for comments and careful reading of the manuscript. Thanks to Inga Hosp for suggesting the perpetuum mobile metaphor and to the "Stiftung Junge Südtiroler im Ausland" for financial support. This is communication number 10 of the Yale Center for Computational Ecology.

References

[1] P. A. Bachmann, Pier-Luigi Luisi, and J. Lang. Autocatalytic self-replicating micelles as models for prebiotic structures. *Nature*, 357:57–59, 1992.

[2] R. J. Bagley and J. D. Farmer. Spontaneous emergence of a metabolism. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, pages 93–141, Redwood City, 1992. Addison-Wesley.

[3] 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.

[4] H. G. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. Studies in Logic and the Foundations of Mathematics. North-Holland, Amsterdam, second edition, 1984. revised edition.

[5] G. Berry and G. Boudol. The chemical abstract machine. *Theoretical Computer Science*, 96:217–248, 1992.

[6] G. Boudol. Towards a lambda-calculus for concurrent and communicating systems. In G. Goos and J. Hartmanis, editors, *Lecture Notes in Computer Science*, number 351, pages 149–161. Springer-Verlag, Berlin, 1989.

[7] L. W. Buss. *The Evolution of Individuality*. Princeton University Press, Princeton, 1987.

[8] M. Calvin. *Chemical Evolution*. Clarendon Press, Oxford, 1969.

[9] A. Church. *The Calculi of Lambda Conversion*. Princeton University Press, Princeton, 1941.

[10] J. F. Crow and M. Kimura. *An Introduction to Population Genetics Theory.* Harper and Row, New York, 1970.

[11] R. Dawkins. Replicator selection and the extended phenotype. *Zeitschrift für Tierpsychologie*, 47:61–76, 1978.

[12] M. Eigen. Self-organization of matter and the evolution of biological macromolecules. *Naturwissenschaften*, 58:465–526, 1971.

[13] M. Eigen, J. S. McCaskill, and P. Schuster. The molecular quasi-species. *Advances in Chem. Phys.*, 75:149–263, 1989.

[14] M. Eigen and P. Schuster. *The Hypercycle.* Springer Verlag, Berlin, 1979.

[15] J. D. Farmer, S. A. Kauffman, and N. H. Packard. Autocatalytic replication of polymers. *Physica D*, 22:50–67, 1982.

[16] R. A. Fisher. *The Genetical Theory of Natural Selection.* Clarendon Press, Oxford, 1930.

[17] W. Fontana. Algorithmic chemistry. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, pages 159–209, Redwood City, 1992. Addison-Wesley.

[18] W. Fontana and L. W. Buss. 'The arrival of the fittest': Toward a theory of biological organization. *Bull. Math. Biol.*, 1993. In press.

[19] W. Fontana and L. W. Buss. What would be conserved 'if the tape were played twice'. *Proc. Natl. Acad. Sci. USA*, 1993. In press.

[20] M. Foucault. *The Order of Things. An Archeology of the Human Sciences.* Vintage Books Edition, 1973. Originally published as *Les mots et les choses*, 1966, Editions Gallimard.

[21] M. Hennessy. *Algebraic Theory of Processes.* The MIT Press, Cambridge, Massachusetts, 1988.

[22] J. Hofbauer and K. Sigmund. *The Theory of Evolution and Dynamical Systems.* Cambridge University Press, Cambridge, 1988.

[23] G. Kampis. *Self-Modifying Systems: A New Framework for Dynamics, Information, and Complexity.* Pergamon Press, Oxford, 1991.

[24] S. A. Kauffman. Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. *J. Cybernetics*, 1:71–96, 1971.

[25] S. A. Kauffman. Autocatalytic sets of proteins. *J. Theor. Biol.*, 119:1–24, 1986.

[26] S. Mac Lane. *Mathematics, Form and Function.* Springer-Verlag, New York, 1986.

[27] C. G. Langton. Artificial Life. In C. G. Langton, editor, *Artificial Life*, Santa Fe Institute Studies in the Sciences of Complexity, pages 1–44, Redwood City, 1989. Addison-Wesley.

[28] K. Lindgren. Evolutionary phenomena in simple dynamics. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, pages 295–312, Redwood City, 1992. Addison-Wesley.

[29] C. Longuet-Higgins. What biology is about. In C. H. Waddington, editor, *Towards a Theoretical Biology. 2: Sketches*, pages 227–235, Chicago, 1969. Aldine Publishing Company.

[30] H. Maturana and F. J. Varela. *De Máquinas y Seres Vivos: Una teoría de la organización biológica.* Editorial Universitaria, Santiago de Chile, 1973. Reprinted in: H. Maturana and F. J. Varela, *Autopoiesis and Cognition: The Realization of the Living*, 1980.

[31] H. Maturana and F. J. Varela. *Autopoiesis and Cognition: The Realization of the Living.* D. Reidel, Boston, 1980.

[32] J. S. McCaskill. Polymer chemistry on tape. unpublished manuscript.

[33] R. Milner. *Communication and Concurrency.* Prentice Hall, New York, 1989.

[34] R. Milner. Functions as processes. In *Lecture Notes in Computer Science*, number 443, pages 167–180. Springer-Verlag, Berlin, 1990.

[35] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, I. *Information and Computation*, 100:1–40, 1992.

[36] H. J. Morowitz. *Beginnings of Cellular Life.* Yale University Press, New Haven, 1992.

[37] S. Rasmussen, C. Knudsen, and R. Feldberg. Dynamics of programmable matter. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, pages 211–254, Redwood City, 1992. Addison-Wesley.

[38] D. S. Rokshar, P. W. Anderson, and D. L. Stein. Self-organization in prebiological systems: Simulation of a model for the origin of genetic information. *J. Mol. Evol.*, 23:110, 1986.

[39] O. Rössler. Ein systemtheoretisches Modell zur Biogenese. *Zeitschrift für Naturforschung*, 26b:741–746, 1971.

[40] P. Schuster and K. Sigmund. Replicator dynamics. *J. Theor. Biol.*, 100:533–538, 1983.

[41] E. Sober. Darwin on natural selection: A philosophical perspective. In David Kohn, editor, *The Darwinian Heritage*, pages 867–899, Princeton, New Jersey, 1985. Princeton University Press.

[42] S. Spiegelman. An approach to experimental analysis of precellular evolution. *Quart. Rev. Biophys.*, 4:36, 1971.

[43] P. F. Stadler, W. Fontana, and J. H Miller. Random catalytic reaction networks. *Physica D*, 63:378–392, 1993.

[44] P. F. Stadler and P. Schuster. Mutation in autocatalytic reaction networks – an analysis based on perturbation theory. *J. Math. Biol.*, 30:597–631, 1992.

[45] M. Thürk. *Ein Modell zur Selbstorganisation von Automatenalgorithmen zum Studium molekularer Evolution.* PhD dissertation, Universität Jena, Germany, 1993.

[46] T. Tjiuikaua, P. Ballester, and J. Rebek Jr. A self-replicating system. *J. Am. Chem. Soc.*, 112:1249–1250, 1990.

[47] F. J. Varela. *Principles of Biological Autonomy.* North-Holland, New York, 1979.

[48] G. von Kiedrowski. A self-replicating hexadeoxynucleotide. *Angew. Chem.*, 98:932–934, 1986.

[49] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, Princeton, third edition, 1953 (1944).

[50] M. B. Williams. Deducing the consequences of evolution: A mathematical model. *J. theor. Biol.*, 29:343–385, 1970.

[51] M. B. Williams. The logical status of natural selection and other evolutionary controversies. In M. Bunge, editor, *The Methodological Unity of Science*, pages 84–102. Dordrecht, 1973.