

Autocatalysis Before Enzymes: The Emergence of Prebiotic Chain Reactions

Nathaniel Virgo, Takashi Ikegami

Ikegami Laboratory, University of Tokyo
nathanielvirgo@gmail.com

Abstract

How could complex, enzyme- or ribozyme-like molecules first have arisen on planet Earth? Several authors have suggested *autocatalytic cycles* as a partial answer to this question, since such reactions exhibit the life-like property of exponential growth while being composed of relatively simple molecules. However, a question remains as to the likelihood of an autocatalytic cycle forming spontaneously in the absence of highly specific catalysts. Here we show that such cycles form readily in a very simple model that includes no direct catalysis reactions. Catalytic effects nevertheless emerge as properties of the reaction network. This suggests that the conditions for the formation of such cycles are not difficult to achieve. The resulting cycles solve the problem of specificity not by being small and simple but by being large and complicated, suggesting that early prebiotic metabolisms could have been extremely complex. We predict that this phenomenon can be reproduced in wet chemistry. We discuss the challenges involved in this, as well as the implications for how we view the origins of life.

Introduction

A necessary requirement for biological metabolism is *autocatalytic kinetics*, i.e. the ability of a set of chemical species to increase its own rate of production. Without the ability to positively influence the production of its own chemical components, the prebiotic equivalent of a living organism would be able neither to reproduce nor to maintain its own composition over time. In this paper we investigate the possibility that the earliest proto-metabolisms achieved this through a mechanism known as an autocatalytic cycle.

In this paper we present a highly simplified model of a simple organic polymer chemistry operating away from thermodynamic equilibrium. This model is extremely simple, consisting only of basic synthesis and decomposition reactions, with no catalytic kinetics assumed *a priori*. We find that autocatalytic cycles form readily in such a system, suggesting that the chemistry in which the first steps toward metabolism took place could have been much simpler than generally supposed. The networks that emerge in our model are complex, consisting of many interlinked catalytic and autocatalytic cycles. The highly interconnected nature of

these autocatalytic subnetworks means that a reaction involving one of the intermediates is likely to produce another intermediate, thus overcoming the much-discussed problem of specificity in autocatalytic cycles. This suggests that complex autocatalytic reaction networks formed from simple molecules can be produced much more easily than simple networks composed of complex “replicator” molecules.

Because the requirements for this phenomenon are so easy to meet, it should be possible to observe it experimentally, in prebiotic chemistry experiments along the lines of the Miller-Urey experiment or HCN polymerisation. To achieve this one would need to change the conditions so that the breakdown of polymers via hydrolysis or oxidation occurs in the same system as their synthesis, at a comparable rate. This simultaneous build-up and break-down of polymers is analogous to anabolism and catabolism in biology. We comment on the potential implications of such a result, and the challenges that would be involved in attaining it.

It is worth pointing out a major difference between our model and one of the predominant existing approaches to explaining the origin of biological autocatalysis. As discussed below, there are many studies that model the emergence of autocatalysis in networks of reactions between peptide or RNA-like molecules, via a mechanism known variously as reflexive autocatalysis, autocatalytic sets or RAF sets. This work has shown that autocatalysis is easy to achieve via this mechanism even if the reaction networks are chosen at random rather than having autocatalysis “designed” into them (Kauffman, 1986); and that such autocatalytic sets are capable of evolution by natural selection via an attractor-based heredity mechanism, even in the absence of specific information-carrying molecules (Vasas et al., 2012).

However, this definition of autocatalysis presupposes the existence of single-step catalysis reactions, and therefore entails an assumption that the molecules involved are complex enough to behave as enzymes. Because our aim is to explain the emergence of such complex molecules from simpler reactants, we focus instead on a different mechanism: the *autocatalytic cycle* or *branching chain reaction* (King, 1978).

For our purposes, a branching chain reaction may be de-




Figure 1: (a) Schematic of the reductive citric acid cycle, redrawn from Morowitz et al. (2000). The branching step is the splitting up of citrate into oxaloacetate and acetyl CoA, which is then transformed into a second oxaloacetate, so that its concentration doubles on every turn of the cycle. (b) The mechanism of the formose reaction, as proposed by Breslow (1959). The branching step is the decomposition of an aldotetrose into two molecules of glycoaldehyde. The formose reaction has been observed experimentally, without the use of biological catalysts.

defined as a net chemical reaction, at least one of whose products is also an intermediate. This allows the concentration of intermediates to build up over time, which under the right conditions can lead to exponential growth. Such reactions are not uncommon and are often the mechanism behind combustion and explosive reactions. A more formal definition of this type of autocatalysis is given by Andersen et al. (2012). In the classification of Plesson et al. (2011), this definition includes direct, indirect and autoinductive forms of autocatalysis.

Some known examples of autocatalysis via branching chain reactions are shown in Fig. 1 and 2. This definition is similar in spirit to that of an autocatalytic set, but in our case the catalysis mechanism emerges from the system's dynamics, rather than being a property of individual molecules.

Autocatalytic cycles have been hypothesised as playing an important role in the origins of life. Wächtershäuser (1988), and later Morowitz et al. (2000) proposed the *reductive citric acid cycle* (Figure 1a) as a possible means by which molecules such as sugars, lipids and amino acids could have been generated on the early Earth. The citric acid cycle is important in modern biology but its intermediate steps are catalysed by enzymes. Wächtershäuser's argument was that inorganic surface catalysts might have been able to play the same role on the early Earth. Morowitz et al. argued that the reductive citric acid cycle might be unique, in the sense of being the only autocatalytic cycle that could lead to the complexity of modern life on an Earth-like planet.




Figure 2: Some other known examples of autocatalysis via chain reactions. (a) A few of the most important reaction steps in the early stages of the combustion of H₂, demonstrating autocatalysis via a more complex network than a single cycle. H₂ and O₂ can be mixed without reacting, but due to this mechanism they will react very rapidly after an initial spark produces small quantities of H and O. (b) Template replication is a special case of chain reaction autocatalysis. Here, an AB dimer catalyses the formation of another AB dimer through complementary base pairing. Figure 2b is taken from Virgo et al. (2012), in which a physical instantiation of template replication was demonstrated using macroscopic “monomers” floating above an air table.

These ideas have been criticised on the grounds that it would be difficult to find mineral catalysts that would catalyse every step in this relatively complex cycle (Orgel, 2000) without also catalysing side-reactions that would reduce the replicator's specificity to a non-viable level (Szathmáry, 2000). This latter problem must be solved by any approach to the origins of life. In any autocatalytic chemical system there will be reactions that contribute to the autocatalytic network (branching reactions and propagating reactions) and reactions that deplete its constituents (terminating reactions). If the latter dominate then growth will not occur.

In this paper we offer solutions to these problems. King (1982) gave a heuristic argument that the formation of autocatalytic cycles is very likely in systems that are driven by a flow of energy across their boundary but closed to matter flow. This is because the products of any reaction will eventually be recycled, and this recycling process has a high probability of forming part of an autocatalytic cycle. Our model confirms that this phenomenon can occur in very simple driven systems, even if the system is not completely closed to matter flow. This suggests that there may be a great number of simpler autocatalytic systems that could have preceded the reductive citric acid cycle, perhaps ultimately leading to the production of complex organic molecules that could play the role of enzymes.

In the same paper, King argued that autocatalytic cycles with many intermediate species are statistically unlikely to be viable, in the sense of being able to grow exponentially. This is because every step in an autocatalytic cycle is vulnerable to side reactions. Every reaction step may be assigned a number between 0 and 1 representing its specificity, and it can be shown that the cycle is only viable if the product of the specificities passes a threshold. Hence, all else being equal, a cycle with many steps is less likely to be viable than one with only a few. However, in our model we observe fairly large autocatalytic systems that are composed not of a single cycle but of many intersecting catalytic and autocatalytic cycles. It seems that such systems avoid the need for specificity simply by including such a large number of species that the production of molecule that *isn't* part of the autocatalytic network is comparatively low.

Our model shows that branching chain reactions occur rather easily under certain conditions. Essentially all that is needed is the simultaneous presence of synthesis reactions (such as polymerisation) and decomposition reactions (such as oxidation or hydrolysis), as well as a source of free energy that causes some reactions to be favoured over others. These processes are closely analogous to anabolism and catabolism in a living cell. On the early Earth there were a wide variety of potential energy sources (Deamer and Weber, 2010) as well as, presumably, a wide variety of environments of varying temperatures, pressures, pH values, redox conditions etc., making it fairly likely for such conditions to be satisfied somewhere on the planet. Such conditions should also be relatively easy to achieve experimentally.

Below we survey the two main existing approaches to the emergence of autocatalysis within the field of ALife, before presenting our own model and its results. This is followed by an extended discussion of how this phenomenon fits into our picture of the early Earth, as well as the challenges involved in demonstrating it in wet chemistry experiments.

Artificial Chemistry approaches to Autocatalysis

Our aim in this work is to apply an “artificial chemistry” methodology to the question of how autocatalytic cycles can arise in prebiotic chemistry. In this section we briefly survey previous work that has had similar aims. This previous work has two main starting points: the work of Kauffman (1986) and the work of Fontana and Buss (1994).

A central work in the metabolism-first school of the origins of life is the model of Kauffman (1986), who showed that, even if the reaction scheme of an artificial chemistry is chosen completely at random, the probability of a collectively autocatalytic set of protein-like polymers becomes high as the number of species present increases. This is an important idea, because it implies that under the right circumstances, the emergence of something akin to biological metabolism might be almost inevitable, even without the organising force of natural selection. With good reason, this

work has spawned a multitude of successors.

However, it must be stressed that, due to its origins in a theory of protein interactions, this body of work assumes a particular mechanism for autocatalysis, which can only occur in relatively complex chemistries. This mechanism relies on the idea that the molecules involved are each able to behave like enzymes, selectively catalysing some reactions but not others in a way that can be modelled as a single-step reaction. This would require the monomers to be of a certain level of complexity. Our aim is to show that similar phenomena can occur without assuming enzyme-like kinetics.

From a quite different direction, the work of Fontana and Buss (1994) looked for autocatalysis in chemistries where the molecules were represented as Lambda calculus expressions. The goal of this work was to investigate the generation of novelty through the formation of autocatalytic structures. This work also spawned a large number of successors, including the work of Ikegami and Hashimoto (1995), who looked for the emergence of autocatalysis in networks of Turing machines and tapes under a noisy environment.

Work in this sub-field tends not to include thermodynamic considerations, choosing instead to emphasise the structure of the reaction network itself. A secondary goal of our work is to investigate the impact of thermodynamic considerations on such “abstract chemistries”. In particular, in real chemistry, reactions may proceed in the forward or in the reverse direction, depending on the free energy difference between the reactants and the products. We will argue that giving the system the ability to “choose” the direction of reactions in this way is important for the emergence of autocatalysis.

A Simple Model

We are concerned with the question of whether autocatalytic cycles, or more complex branching chain reactions, can occur in simple (non-enzymatic) organic chemistry. To do so we use a model in which a reaction network is randomly generated by allowing or disallowing cleavage and ligation reactions between polymers. A key difference between our work and previous work is that in our model no polymer can directly catalyse any reaction, so any autocatalysis that occurs must be via cycles rather than enzyme-like catalysis.

In polymer models in artificial chemistry, molecules are usually considered to consist of a string of m different types of monomer. For the sake of simplicity, in this work we set $m = 1$, restricting ourselves to a single monomer type, denoted A. The possible species can therefore be written A_1, A_2, \dots, A_n , where n is a maximum allowed polymer length, imposed for reasons of computational tractability. These are intended to represent molecules based on simple carbon chains, rather than complex heteropolymers such as peptides or RNA strands.

All reactions must preserve the number of monomers. We consider only reactions of the form $A_i + A_j \rightleftharpoons A_k$, where $k = i + j$ is not greater than n and, to avoid duplicates,

$i \leq j$. To generate a reaction network we must decide, for every such reaction, whether to include it in the network or not. For simplicity, following Kauffman, we simply include each reaction in the network with a constant uniform probability p , independently of every other reaction. The double arrow indicates that the forward reaction $A_{i+j} \rightarrow A_i + A_j$ and the reverse reaction $A_i + A_j \rightarrow A_{i+j}$ are always either both included in the network or both not included. This is required for consistency with thermodynamics, and as we will see, it plays an important role in the emergence of autocatalytic networks in our model.

We assume that the rate constants of all the included reactions are equal. Setting the forward rate constant to 1 without loss of generality and letting k stand for the reverse rate constant, we let each reaction $A_i + A_j \rightleftharpoons A_{i+j}$ occur at a net rate $R_{ij} = a_i a_j - k a_{i+j}$, where a_i is the molar concentration of species i . R_{ij} may be positive or negative, representing a net synthesis or net decomposition reaction, depending on the concentrations of the three reactants. The justification for the $a_i a_j$ term is that the ends of two polymer molecules must meet in order for them to undergo a ligation reaction, and we assume that the polymer tips, rather than the polymers themselves, behave like point particles in a well mixed system. (Such “mass action” assumptions are common in models of polymerisation kinetics.) The $-k a_{i+j}$ term simply means that, for every decomposition reaction in the network, there is a constant probability per unit time that it will occur in a given molecule. We write $R_{ij} = 0$ if the reaction is not included in the network.

This leads to the following set of dynamical equations:

$$\dot{a}_i = \phi_i + \sum_{k=1}^{i-1} R_{k,i-k} - \sum_{j=1}^i R_{ij} \quad (1)$$

where the R_{ij} are as defined above, and ϕ_i represents the flux of A_i in or out of the system, as explained below. The two summation terms arise from the fact that each species A_i is involved in reactions of the form $A_k + A_{i-k} \rightleftharpoons A_i$ as well as $A_j + A_i \rightleftharpoons A_{i+j}$.

Thermodynamic properties

If we let the fluxes $\phi_i = 0$, the system will approach thermodynamic equilibrium. In such a state the reactions have the property of *detailed balance*, meaning that the forward and reverse rates are equal for every reaction. For a reaction $A_i + A_j \rightleftharpoons A_{i+j}$ this occurs when $a_i a_j = k a_{i+j}$, or $\log a_i + \log a_j = \log k + \log a_{i+j}$. We may therefore define the *chemical potential* μ_i of species i to be $(i-1) \log k + \log a_i$. This has the property that when the system is in thermodynamic equilibrium, $\mu_i + \mu_j = \mu_{i+j}$. (The usual thermodynamic definition of chemical potential would include a factor of RT , the gas constant times the temperature, which we have set to 1 for convenience.) From this we may define the Gibbs energy $G = \sum_i \mu_i a_i$. In accordance

with the second law, G cannot increase over time unless we allow some fluxes of matter in and out of the system. For a closed system, G is a Lyapunov function.

If we temporarily assume that every allowable reaction is included in the network ($p = 1$), we can see that the equilibrium concentrations must satisfy $a_i = k e^{C_i}$, for some constant C , in order for detailed balance to hold for every reaction. The value of the constant C depends on the initial conditions, which stems from the fact that the total number of monomers in the system, $\sum_i i a_i$, is conserved. Low initial concentrations will lead to decomposition reactions being favoured, and therefore low (negative) values for C , whereas high total monomer concentrations lead to synthesis reactions being favoured. High enough concentrations lead to positive values for C , meaning that the equilibrium conditions are dominated by the longest possible polymers rather than by short ones. This phenomenon is observed in real polymer chemistries. If $p < 1$ then it is possible for equilibrium situations to exist where this condition is not satisfied, because conservation laws arise that prevent some concentrations from becoming equilibrated with one another. However, in these cases higher concentrations still lead to longer products being favoured.

In order to observe the operation of autocatalytic cycles, the system must be held away from thermodynamic equilibrium. In real chemical systems this can be achieved in many ways. For example, by cycling the temperature or pH (both of which would effectively change k in our model), or through electrochemistry or photochemistry, which can drive reactions that would otherwise not be thermodynamically favourable. In the first set of results below we model the reactions as being held out of equilibrium by continually adding reactants and removing products, as in a flow reactor; in the second we simply start the system in an initial state far from equilibrium and observe the decrease in Gibbs energy over time.

Results

In this section we present the results from two different simulations based on the above model. The first serves as a useful demonstration of the formation of autocatalytic cycles in driven systems, but is somewhat contrived; the second shows that autocatalytic kinetics can arise in larger, randomly-generated systems.

In our first model we set $p = 1$, including every reaction in the network, but we limit the size of the largest polymer. There are many ways in which the system may be held out of equilibrium; in this illustrative example we do it by letting the fluxes ϕ_1 and ϕ_2 have nonzero values, with their rates chosen such that the concentrations a_0 and a_1 are held constant at 100 and 0.1 respectively. Conceptually, A_1 flows into the system, then undergoes a series of reactions until it is converted into A_2 , at which point it is removed. Boundary conditions of this type could be achieved experimentally




Figure 3: Time series of the concentrations in the model when $n = 6$ and $n = 7$. In both cases there is a period of exponential growth from $t = 0$ to about $t = 2.0$, indicating autocatalysis. This is followed by gradual saturation.

using membrane permeable only to small molecules.

Figure 3 shows the dynamics of this system when $n = 6$ and $n = 7$. In both cases there is a period of exponential growth followed by a period of saturation. Exponential growth is a key experimental sign of autocatalysis. (With $n \leq 5$ this effect does not occur.)

Figure 4 shows the reaction networks that arise once these systems have reached a steady state. (We were unable to find more than one attractor in these particular systems, although the existence of others cannot be ruled out.) The recycling structure of these networks can be seen as a response to the flux of Gibbs energy across the system's boundary, in accordance with Morowitz' (1966) cycling theorem. It is for this reason that we believe including thermodynamically realistic kinetics in such models is important for understanding the origins of autocatalytic cycles.

In both cases the mechanism behind the exponential growth is an autocatalytic cycle that produces two molecules of an intermediate for every molecule present initially; this exponential growth is countered by decay reactions once the concentrations become high. However, the two systems use different autocatalytic cycles. This is possible because the direction in which reactions occur is determined by the differences in the reactants' chemical potentials, and these depend upon the system's dynamics.

As n is increased further, more catalytic and autocatalytic cycles emerge (results not shown). However, it can be seen from Figure 3 that the concentrations of longer polymers are much higher than those of short ones; this trend continues as n is increased, leading to unrealistic results as n becomes large, since in reality a system composed mostly of long polymers will become viscous or solid, preventing further reactions by suppressing mixing.

However, this issue can be resolved by choosing different values for the parameters, so that shorter rather than longer ones are thermodynamically favoured. In addition to doing this we set n large enough that the longest polymer only




Figure 4: The reaction networks that form when $n = 6$ and $n = 7$, with $p = 1$ and the concentrations of A_1 and A_2 held constant. Propagating reactions are shown in black or grey, branching reactions in red, and terminating reactions in blue. Numbers represent the rate at which each reaction occurs once a steady state is reached, in multiples of 10^{-3} concentration units per time unit. The set of allowed reactions is predetermined, but the direction in which they proceed depends on the system's dynamics. Both networks contain several catalytic cycles, coupled to an autocatalytic cycle (highlighted in black). The reaction $A_6 \rightarrow 2A_3$ is the key branching step when $n = 6$, but when $n = 7$ it runs in the opposite direction, becoming a depleting reaction.

ever exists at a low concentration. The parameters we use are $K = 100$, $p = 0.2$ and $n = 40$. When such a system is driven toward a steady state, it produces very complex networks that are difficult to analyse. Andersen et al.'s (2012) algorithm could be used to detect autocatalytic subnetworks, but it cannot tell us how viable they are. Because of this, instead of driving the system we simply initialise it in a state with a high Gibbs energy and observe its return to equilibrium. This allows us to detect autocatalysis by observing exponential growth in the kinetics. We use the initial conditions $a_1 = 1000.0$, and $a_i = K^{1-i}$ for $i > 1$. This can be interpreted as a system that was initially in equilibrium, to which a large quantity of monomers has just been added.

Figure 5 shows the results of this simulation. We numerically integrated the dynamics of 50 randomly generated networks for 3 time units each. In 32 out of the 50 cases, no reactions occurred and the system remained in its initial state. In 14 out of the remaining 18 cases, there was at least one period of time in which a species' concentration increased with $d^2 a_i / dt^2 > 0$ and $da_i / dt > 0.01$. Such "accelerating" growth is an indication that there is a viable autocatalytic network within the system.

The behaviour of the system is quite sensitive to the choices of parameters, but the phenomenon of exponential growth appears to be fairly robust. Quantifying this is a task for future work.




Figure 5: (a) An example of the dynamics when $n = 40$, $p = 0.2$ and the system is closed but initially out of equilibrium. In this case the reaction $2A_1 \rightleftharpoons A_2$ rapidly goes to equilibrium and the system stays in this state for a while before a complex autocatalytic network arises and rapidly brings the system near to thermal equilibrium. (b) Superimposed results from 50 networks, showing the change in Gibbs energy over time. The initial conditions are identical for each network and have a Gibbs energy of about 6900, whereas the equilibrium state has a Gibbs energy of about 4300. The lines are coloured red when at least one species in the system is increasing with a positive second time derivative, indicating the presence of a viable autocatalytic network. The systems typically approach equilibrium more rapidly when an autocatalytic network is operating.

Discussion and Future Work

We have presented a model that couples a simple abstract chemistry with thermodynamically realistic kinetics, in order to show that autocatalysis via branching chain reactions can occur even in very simple chemical systems. The origins of life are often thought of in terms of a “self-replicating molecule” which, as Figure 2b shows, can be thought of as a small autocatalytic cycle composed of reactions between complex molecules. Our results suggest that it may be much easier to achieve the opposite: a large autocatalytic network composed of simple molecules.

An important property of our model is that both forward and reverse reactions are included, subject to thermodynamically realistic kinetics. As a result of this we can observe that an externally introduced source of energy drives cycling behaviour (Morowitz, 1966), and this recycling leads to autocatalytic kinetics (King, 1982). We therefore believe that adding reverse reactions and thermodynamic constraints to “abstract chemistry” models along the lines of (Fontana and Buss, 1994) could shed light on the process of novelty generation in general, as well as the origins of life in particular.

The main components of our model are (i) a system that is at least partially closed to matter flow, in which both synthesis and decomposition reactions can occur; and (ii) an energetic driving force, which causes some reactions to be

favoured over others. The simplicity of our model suggests that such conditions are essentially all that is required for autocatalytic networks to form. This makes it much more plausible that autocatalytic chemical systems could emerge on the early Earth, and the simplicity of the conditions makes the idea amenable to experimental testing in real chemistry.

It has been shown that Kauffman’s autocatalytic sets are capable of evolution by natural selection, even without the existence of specific information-carrying molecules (Vasas et al., 2012). Our hope is that something similar will be true of autocatalytic systems that occur via chain reactions rather than single-step enzyme-like catalysis. If this is the case then we may suggest that life did not start with the citric acid cycle but with a different autocatalytic system, perhaps composed of simpler molecules, but forming a much more intricate network of reactions. The catalysts required to produce the molecules of modern life via the reductive citric acid cycle could then have been arrived at by natural selection acting on the original autocatalytic system.

However, the models we have presented here seem not to exhibit the large number of attractors that would enable heredity in such a way. We must therefore discuss what additional conditions might need to be met in order for an evolvable system to arise.

Constraints on the Reaction Network

Our model obeys constraints imposed by mass conservation and the laws of thermodynamics, but beyond this we choose the permitted reactions at random. As we have seen, this results in autocatalytic networks that tend to include almost every possible species as part of their network. In order for the system to exhibit a large number attractors there would need to be multiple possible autocatalytic networks capable of out-competing each other.

Real chemical reaction networks are not random but are determined by the physics of molecular interactions. This imposes a number of constraints both on the topology of chemical reaction networks and on their kinetics, and such constraints might help to “partition” the network into multiple potential autocatalytic sub-networks. Perhaps the most obvious such constraint is imposed by stoichiometry: chemical reactions must conserve not only mass but also the number of nuclei of each chemical element, as well as electrons. Our system recycles monomers, but in biology (particularly at the ecosystem level) the recycling of specific nutrients such as nitrogen and phosphorous plays an important organising role. We therefore suspect that adding multiple conservation laws to our model will enable a richer range of behaviours than it currently exhibits. (However, this would greatly increase the number of possible molecular species in the model, requiring a change in modelling methodology from the simple ode integration that we have used here.)

Another important set of constraints are given by the shapes of molecules and the ways in which they interact

electromagnetically. This gives reaction networks the important property that similar molecules can undergo similar reactions. Modelling the relationship between the form of molecules and the resulting reaction network is of course rather difficult, but perhaps something like the approach of Fontana and Buss (1994), combined with the thermodynamic realism of the present model, would be a useful tool to investigate this question.

Finally, the existence of phase changes can also put constraints on the reaction network. King (1982) argued that this could enhance the formation of viable autocatalytic cycles. Adding phase separation to our model would allow us to investigate this idea.

The Importance of Spatial Structure and Compartmentalisation

In addition to phase separation, more complex spatial structuring may be important in going from simple to more complex autocatalytic networks. Many previous studies, including some by the present authors, have concluded that spatial self-organisation is important for avoiding “parasitic” side-reactions, i.e. sets of reactants that produce themselves autocatalytically, feeding not directly on the energy source but on the original autocatalyst (e.g. Boerlijst and Hogeweg, 1991; Froese et al., 2011, 2012, 2013). In (Froese et al., 2012; Virgo et al., 2013) we found that in a spatial context, parasitic reactions could become a positive benefit to the primary autocatalytic system, leading to evolvability. We expect that embedding a system along the lines of the present model in a spatial context will lead to richer dynamics.

Many hypotheses about the origins of life require “compartmentalisation”, the formation of a lipid vesicle, or similar small compartment, in which reactions take place. One reason for this arises from energetics: for complex biological polymers such as peptides to form, the monomers must be present in sufficient concentration, and since large monomers like amino acids or RNA bases are difficult to produce in such concentrations, a membrane is required in order to prevent them from diffusing into the environment.

With simpler molecules the energetics of polymerisation are less constrained, and simple monomers could more easily be produced abiotically. Compartmentalisation is therefore less critical for the kind of prebiotic system we consider in this paper, and one can therefore imagine such phenomena occurring in a relatively dilute “prebiotic soup”, or more accurately, a prebiotic flow reactor.

A second reason to require compartmentalisation is simply that there must be a population of multiple individuals in order for natural selection to occur. We suggest that simple spatial separation could have played this role originally, in a manner outlined in (Virgo, 2011; Froese et al., 2012), only later to be replaced by membrane-bound cell structures.

If autocatalysis can occur in solution, and if the autocatalytic network also produces lipid-like molecules, then

membrane-bound protocells may be able to form spontaneously (Ono and Ikegami, 2000; Madina et al., 2003). This neatly solves the chicken-and-egg problem of how membrane-bound autocatalysis could first have arisen.

Towards Empirical Verification

In our model, autocatalysis via branching chain reactions emerges in a system that contains only simple synthesis and decomposition reactions, together with a supply of free energy. This idea should be relatively easy to demonstrate experimentally. Previous experiments relevant to the origins of life, such as the Miller-Urey experiment Miller (1953) or the polymerisation of hydrogen cyanide (HCN) (see, e.g., Minard et al., 1998) have focused on the production of organic molecules through polymerisation. Both of these experiments produce a diverse mixture of products, including amino acids; however, these products form into a black, sticky “tar” called tholin that seems unlikely to self-organise into anything like a biological metabolism, despite the fact that tholin itself is thermodynamically unstable and can be used as an energy source by several common species of bacteria (Stoker et al., 1990).

Our results suggest that autocatalytic cycles may emerge in such experiments if the conditions are changed so that breakdown of polymers via hydrolysis or oxidation can occur simultaneously with the polymerisation, at a comparable rate. Since polymer molecules are continually built up and broken down, we would expect those that can produce themselves autocatalytically to persist at the expense of those that cannot. The kinetics and energetics of both polymerisation and depolymerisation are sensitive to environmental factors such as temperature, pH, monomer concentration and the presence of surfaces and inorganic catalysts. Achieving autocatalysis should simply be a case of setting the appropriate conditions for the reaction. We are currently working on demonstrating this in an HCN polymerisation experiment.

The challenge in such an experiment is in demonstrating that an autocatalytic cycle has indeed emerged. The sheer number of products means that the resulting mixture tends to have a continuous mass spectrum, making it difficult to identify which species are present. However, evidence for autocatalytic kinetics would be given by sudden changes in the mass spectrum, even if one cannot readily identify the species responsible.

The Prebiotic Ecosystem

Above we have mentioned several phenomena, such as nutrient cycling and parasitism, that one would normally associate with physical ecology than purely chemical systems. It is worth drawing an explicit conclusion from this: we believe that prebiotic systems should be thought of as resembling ecosystems, complete with food chains, nutrient cycling, energetic restrictions and all the rest — everything except for clearly differentiated living cells, which arose later.

We know that the early Earth was a very active world, with sources of chemical free energy from UV photochemistry in the atmosphere, shockwaves from asteroid impacts, radioactivity, lightning, volcanoes and geochemistry (Deamer and Weber, 2010), and matter cycling due to plate tectonics and the water cycle. In such a context, it is easy to imagine that such prebiotic ecosystems could have been a global phenomenon, leading to primordial equivalents of today's biogeochemical cycling of nitrogen, phosphorous and carbon. From this point of view a homeostatically self-regulating Earth system should be seen not as a *consequence* of the biosphere (Lovelock, 1987), but rather as the context in which it first arose.

Conclusion

We have set out to explain how autocatalysis could have emerged on the early Earth, before the existence of enzyme-like catalysts. We have shown, using a simple model, that autocatalytic cycles can emerge in chemical systems with only synthesis and decomposition reactions, without requiring the molecules to have special catalytic properties. The resulting autocatalytic networks solve the problem of specificity not by being small and simple but by being large and complicated. We conclude that the earliest origins of life may have lain not in a "minimal" autocatalytic system but in a "maximal" one.

The conditions required for this to occur are simple enough that we hope it can be demonstrated in wet chemistry experiments, and we have discussed how this could be achieved. Finally, based on our results, we have argued that the prebiotic Earth should be seen not as a soup but as an ecosystem.

References

- Andersen, J. L., Flamm, C., Merkle, D., and Stadler, P. (2012). Maximizing output and recognizing autocatalysis in chemical reaction networks is np-complete. *Journal of Systems Chemistry*, 3(1):1–9.
- Boerlijst, M. C. and Hogeweg, C. (1991). Spiral wave structure in pre-biotic evolution: Hypercycles stable against parasites. *Physica D*, 48:17–28.
- Breslow, R. (1959). On the mechanism of the formose reaction. *Tetrahedron Letters*, 21:22–26.
- Deamer, D. and Weber, A. (2010). Bioenergetics and life's origins. *Cold Spring Harbour Perspectives in Biology*, 2(a004929).
- Fontana, W. and Buss, L. W. (1994). What would be conserved if "the tape were played twice"? *Proceedings of the National Academy of Sciences*, 91:757–761.
- Froese, T., Ikegami, T., and Virgo, N. (2012). The behavior-based hypercycle: From parasitic reaction to symbiotic behavior. In Adami, C. et al., editors, *Proceedings of Artificial Life 13*, pages 457–464. MIT Press.
- Froese, T., Virgo, N., and Ikegami, T. (2011). Life as a process of open-ended becoming: Analysis of a minimal model. In Lenaerts, T. et al., editors, *Advances in Artificial Life, ECAL 2011*, pages 250–257, Cambridge, MA. MIT Press.
- Froese, T., Virgo, N., and Ikegami, T. (2013). Motility at the origin of life: Its characterization and a model. *Artificial Life*, Early Access publication February 1st.
- Ikegami, T. and Hashimoto, T. (1995). Active mutation in self-reproducing networks of machines and tapes. *Artificial Life*, 2(3):305–318.
- Kauffman, S. (1986). Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119:1–24.
- King, G. A. M. (1978). Autocatalysis. *Chemical Society Reviews*, 7(2):297–316.
- King, G. A. M. (1982). Recycling, reproduction and life's origins. *BioSystems*, 15:89–97.
- Lovelock, J. (1987). *Gaia: a New Look at Life on Earth*. Oxford University Press.
- Madina, D., Ono, N., and Ikegami, T. (2003). Cellular evolution in a 3d lattice artificial chemistry. In Banzhaf, W. et al., editors, *Advances in Artificial Life: 7th European Conference, ECAL 2003*, pages 59–68.
- Miller, S. L. (1953). Production of amino acids under possible primitive earth conditions. *Science*, 117:528–529.
- Minard, R. D., Hatcher, P. G., and Gourley, R. C. (1998). Structural investigation of hydrogen cyanide polymers: New results using TMAH thermochemolysis/GC-MS. *Origins of Life*, 28:461–473.
- Morowitz, H. (1966). Physical background of cycles in biological systems. *Journal of Theoretical Biology*, 13:60–62.
- Morowitz, H., Kostelnik, J., Yang, J., and Cody, G. (2000). The origin of intermediary metabolism. *Proceedings of the National Academy of Sciences*, 97(14):7704–7708.
- Ono, N. and Ikegami, T. (2000). Self-maintenance and self-reproduction in an abstract cell model. *Journal of Theoretical Biology*, 206:243–253.
- Orgel, L. (2000). Self-organising biochemical cycles. *PNAS*, 97(23):12503–12507.
- Plesson, R., Brandenburg, A., Jullien, L., and Bersini, H. (2011). Autocatalysis: At the root of self-replication. *Artificial Life*, 17:219–236.
- Stoker, C. R., Boston, P. J., Mancinelli, R. L., Segal, W., Khare, B. N., and Sagan, C. (1990). Microbial metabolism of tholin. *Icarus*, 85(1):241–256.
- Szathmáry, E. (2000). The evolution of replicators. *Philosophical Transactions of the Royal Society, Series B.*, 355:1669–1676.
- Vasas, V., Fernando, C., Santos, M., Kauffman, S., and Szathmáry, E. (2012). Evolution before genes. *Biology Direct*, 7(1).
- Virgo, N. (2011). *Thermodynamics and the Structure of Living Systems*. PhD thesis, University of Sussex, UK.
- Virgo, N., Fernando, C., Bigge, B., and Husbands, P. (2012). Evolvable physical self-replicators. *Artificial Life*, 18:129–142.
- Virgo, N., Froese, T., and Ikegami, T. (2013). The positive role of parasites in the origins of life. To appear in proceedings of IEEE SSCI 2013.
- Wächtershäuser, G. (1988). Before enzymes and templates: Theory of surface metabolism. *Microbiology and Molecular Biology Reviews*, 52(4):452–484.