

Autocatalytic confusion clarified

Wim Hordijk^{a,*}

^a*Konrad Lorenz Institute for Evolution and Cognition Research, Klosterneuburg, Austria*

Abstract

There is frequent confusion about the terms autocatalytic *reaction*, autocatalytic *cycle*, and autocatalytic *set*. As the use of the same adjective implies, these three systems do indeed share common properties, in particular their potential for exponential growth. However, the ways in which they achieve this potential are different, giving rise to different internal network structures and dynamics. Therefore, care should be taken which term is used in which context. Here, we explain and discuss the similarities and differences between the three systems in detail, in an effort to avoid any further confusion. We then also discuss the relevance of these autocatalytic systems for possible origin of life scenarios, with an emphasis on how autocatalytic sets may have played an important role in this.

Keywords: Autocatalysis, exponential growth, origin of life

1. Introduction

In chemistry, a *catalyst* is a compound (molecule) that speeds up the rate at which a chemical reaction happens, without being used up in that reaction. Catalysis is ubiquitous in living systems. The majority of biological reactions are catalyzed, and catalysts are essential in determining and regulating the functionality of the chemical reaction networks that support life (Szöke et al., 2003).

Not only do catalysts significantly *increase* the rates at which these reactions happen, they also *synchronize* these rates more closely. For example, Wolfenden and Snider (2001) measured the uncatalyzed and catalyzed rates of several biological reactions (see Figure 1). The rate increase from uncatalyzed to catalyzed reactions is many orders of magnitude. Furthermore, the

*Corresponding author. Email: wim@WorldWideWanderings.net

range of rates is decreased from roughly 15 orders of magnitude to only about three or four orders of magnitude. Both of these properties, increase in absolute rates and decrease in range of rates, are important for living systems to function properly.




Figure 1: Logarithmic scale of uncatalyzed (or spontaneous) and (enzyme) catalyzed reaction rates for several representative biological reactions. Reproduced from Wolfenden and Snider (2001).

A reaction, or set of reactions, where (some of) the reaction products catalyze their own production, directly or indirectly, is called *autocatalytic* (meaning “self-catalyzing”). However, such autocatalysis can happen in several different ways. In particular, the term “autocatalytic” is often used in the context of an autocatalytic *reaction*, an autocatalytic *cycle*, or an autocatalytic *set*.

Unfortunately, though, there is frequent confusion regarding these terms. For example, even the Wikipedia page on the topic of autocatalysis talks about these three systems as if they can all be thrown in the same pot

(Wikipedia, 2017). Of course the three systems do share common properties, in particular their potential for exponential growth. In fact, this is usually the property of interest when the term “autocatalytic” is used.

However, there are also important differences between the three systems. For example, in an autocatalytic reaction and an autocatalytic set each individual reaction is explicitly catalyzed by some molecule that is a product of the reaction or set itself. In contrast, in an autocatalytic cycle none of the individual reactions are necessarily catalyzed by anything, or at least not by its own products. On the other hand, an autocatalytic set does not necessarily contain a cycle, or if it does it often also consists of more than just a cycle.

Because of these differences, the specific internal structure and dynamics of these systems can also be different, besides the fact that they all give rise to (potentially) exponential growth of at least one of their products. So, care should be taken which term is used in which context, as it may lead to a misunderstanding of which chemical reaction network structure or internal dynamics is meant.

In this paper we explain the similarities and differences between the three systems in detail, in an effort to avoid any further, potentially exponentially growing, confusion. We also briefly touch upon the notion of hypercycles (an additional source of confusion), and discuss the relevance of all of these systems for possible origin of life scenarios.


A similar comparison was included in a overview by Szathmary (2000). However, at that time the detailed theory of autocatalytic sets (now known as RAF theory; see below) had not been developed yet, and there were no sophisticated experimental examples of real chemical autocatalytic sets. In fact, Szathmary (2000) referred to them as “hypothetical” networks. Here, an updated overview and comparison of autocatalytic systems is presented, with an emphasis on the latest developments in theory and experiments on autocatalytic sets. For brevity, in the remainder we will refer to autocatalytic reactions, cycles, and sets as AC reactions, AC cycles, and AC sets, respectively.

2. Autocatalytic systems

2.1. Autocatalytic reaction

An *autocatalytic reaction* is a single chemical reaction for which one of the products also catalyzes the reaction. An example of such a reaction is

the oxidation of oxalic acid by potassium permanganate (Issa et al., 1960):



Initially this reaction proceeds at very low rates. However, as more and more manganese(II) ions (Mn^{2+}) are produced, the reaction speeds up significantly. Alternatively, if some initial Mn^{2+} is already present, the reaction immediately proceeds at higher rates. In other words, this reaction is catalyzed by one of its own products, and is thus an AC reaction.

In its most simple form, an AC reaction can be written as follows:



where molecule type A catalyzes its own production from some “food source” f (this could be a single molecule type, or represent several types). The variable k above the arrow indicates the *rate constant* of the reaction, which usually depends on the activation energy to make the reaction happen, and environmental conditions such as temperature and acidity. Figure 2 shows a graphical representation of this reaction, where black dots represent molecule types and the white box represents a chemical reaction, with solid arrows indicating reactants going into and products coming out of the reaction, and the dashed arrow representing catalysis.




Figure 2: A simple AC reaction. Black dots represent molecule types and the white box represents a chemical reaction, with solid arrows indicating reactants going into and products coming out of the reaction, and the dashed arrow representing catalysis.

The change in concentration of molecule type A over time can now be written as an ordinary differential equation (ODE):

$$\frac{d}{dt}[A] = k[f][A] \tag{2}$$

where the square brackets (as in $[A]$) indicate molar concentration (i.e., mole per unit of volume). Assume, for simplicity, that there is a constant concentration c of food molecules, i.e., $[f] = c$. In other words, as soon as one food

molecule f is converted into A , it is replaced by a new food molecule f . This could, for example, represent a *buffered* food supply. Equation (2) can then be rewritten as:

$$\frac{d}{dt}[A] = kc[A] \quad (3)$$

which has as solution:

$$[A](t) = C_0 e^{kct} \quad (4)$$

where C_0 is the initial concentration of molecule type A , i.e., $C_0 = [A](0)$.

Equation (4) immediately makes it clear that the concentration of A will grow exponentially over time, given an unlimited supply of food molecules. If there is a limited food supply, this exponential growth will eventually slow down and level off. Figure 3 (black line) shows this (initial) exponential growth over time for a rate constant $k = 1.0$, a constant food concentration $c = 0.05$ (e.g., 5 mole per 100 μL), and an initial concentration of molecule type A of $C_0 = 0.01$.

However, this ODE solution represents an idealized situation, which forms a good approximation for systems with high molecular concentrations in a well-stirred reaction vessel. When concentrations are very low, though, especially during an initial transient phase, the actual behavior of chemical systems can deviate from their idealized ODE representation. In these cases, stochastic simulations using the well-known Gillespie algorithm (Gillespie, 1976, 1977) often form a better tool for studying the system’s behavior. The Gillespie algorithm keeps track of actual molecular counts (i.e., number of molecules of a given type), simulating the “execution” of individual reactions, one by one. Figure 3 (red line) shows the result of one such simulation of reaction (1) with similar parameter values, but appropriately rescaled to match the ODE solution.

Note that the mathematical analysis presented here represents the (ideal) theoretical case, where the *reaction order* p is equal to 1, i.e., the (initial) rate of product formation is proportional to the initial concentration C_0 . In practice, however, autocatalytic reactions may have a reaction order $p < 1$, i.e. they may not show full exponential growth.

In reality, the single reaction presented in Equation (1) proceeds in three





Figure 3: The (initial) exponential growth in concentration of the product of an auto-catalytic reaction, resulting from a solution to the ODE representation (black line) or a Gillespie simulation (red line).

separate steps:




where $f \bullet A$ and $A \bullet A$ represent molecule compounds. In other words, first the catalyst binds to the reactant(s) (remember f could represent multiple molecules), which is a reversible reaction. Next, the reactants are transformed into the product (while still bound to the catalyst), which is (often) a non-reversible reaction. And finally, the original catalyst and the new product dissociate, which again is a reversible reaction.

Therefore, reactants (f) may be temporarily bound into a compound ($f \bullet A$) but not give rise to a product if the compound dissociates again before the second reaction step happens. Similarly, a newly formed product (A) may be temporarily unavailable as a catalyst as long as it stays bound in a compound ($A \bullet A$). As a consequence, especially if the dissociation of $A \bullet A$ has a relatively low rate, in many chemical experiments a reaction order of $p = 1/2$ has been observed (or some value $1/2 < p < 1$), i.e., the (initial) rate of autocatalytic product formation is proportional to $C_0^{1/2} = \sqrt{C_0}$. This is known as the *square-root law of autocatalysis*, which gives rise to hyperbolic (e.g., quadratic) growth rather than exponential (von Kiedrowski, 1993; Paul and Joyce, 2004). A similar phenomenon applies to the other autocatalytic systems, as discussed next, as well.

2.2. Autocatalytic cycle

An *autocatalytic cycle* is a sequence of reactions that, once completed, results in two (or more) copies of the molecule type that was started with. An example is the *formose reaction*, which consists of a sequence of reactions that produces two glycolaldehyde molecules ($CH_2OH \cdot CHO$) starting from one glycolaldehyde and adding two formaldehyde molecules (CH_2O) along the way (Orgel, 2000). This reaction sequence can be written in short-hand notation as:



In other words, the formose reaction sequence can be represented by a single AC reaction, and is therefore considered an AC cycle.

In one of its simplest forms, an AC cycle consists of the following reaction sequence:



where f again represents some food source (either a single molecule type, or a set). This reaction sequence starts with a single molecule of type A , uses the food source to create two intermediates B and C , with C then splitting up into two molecules of type A . Note that each reaction could, in principle, create additional products (“waste”), but this is not represented here for




Figure 4: A simple AC cycle consisting of three reactions, creating two molecules of type A starting from one A and using two food molecules along the way.


simplicity. Figure 4 shows a graphical representation of this cyclic sequence of reactions.

This sequence of reactions can be written in short-hand notation as:



Note that this reaction is indeed similar to the AC reaction given in equation (1). Therefore, its overall dynamics will also be similar, in particular the (potential) exponential growth of molecule type A . However, in this case the corresponding rate constant is constrained by the lowest rate constant of the individual reactions in the cycle, and the precise internal dynamics of the cycle may be more complicated.

A more complex example of an AC cycle is the *reverse citric acid cycle*. The core cycle consists of a sequence of eleven reactions, using CO_2 as a food source to produce two citric acid molecules ($C_6H_8O_7$) starting from one (Morowitz et al., 2000). The overall reaction can be represented by:



This is again similar to the AC reaction in equation (1), and thus also represents an AC cycle. Note that the square-root law of autocatalysis holds here as well, possibly giving rise to sub-exponential growth in practice.

2.3. Autocatalytic set

An *autocatalytic set* is a set of reactions and molecules such that these molecules mutually catalyze each other’s formation from a basic food source f , using only reactions from the set. More formally, an AC set (or *RAF set*) is a set of reactions \mathcal{R} that is:

1. *Reflexively Autocatalytic* (RA): each reaction $r \in \mathcal{R}$ is catalyzed by at least one molecule type that is either a product of \mathcal{R} or is present in the food set f ; and
2. *F-generated* (F): all reactants in \mathcal{R} can be created from the food set f by using a series of reactions only from \mathcal{R} itself.

A mathematically rigorous definition of RAF sets is provided in Hordijk and Steel (2004); Hordijk et al. (2011). A RAF set forms a catalytically closed (RA) and self-sustaining (F) reaction network. The concept of an AC set was first introduced by Kauffman (1971, 1986, 1993), and more recently studied extensively both mathematically and computationally as RAF theory (Hordijk and Steel, 2017). The terms AC set and RAF set will be used interchangeably here, with AC set usually referring to the general concept, and RAF set to the formal mathematical framework.

There are several chemical examples of AC sets. These examples were constructed by generalizing the idea of an *autocatalytic* reaction to a pair (or even larger set) of *cross-catalytic* reactions (Dadon et al., 2008). The first such example that was created experimentally consists of two oligonucleotides (two hexamers that form each other’s base-pair complement) that mutually catalyze each other’s ligation from a food source consisting of trimers (Sievers and von Kiedrowski, 1994). Later, a similar AC set of much longer ($>70\text{nt}$) RNA catalysts, or ribozymes, was constructed by Kim and Joyce (2004), and AC sets of up to 16 even longer ($\sim 200\text{nt}$) ribozymes were produced by Vaidya et al. (2012), with these ribozymes mutually catalyzing each other’s formation from shorter RNA fragments. However, not only nucleic acid polymers have been used, as a nine-member AC set of peptides (32aa) has also been constructed in the lab (Ashkenasy et al., 2004).

The basic AC set of Sievers and von Kiedrowski (1994) can be represented by the following set of reactions \mathcal{R} :



where f is again the food source (in this case representing a set of nucleic acid trimers), k the rate constant (assumed to be the same for both reactions here, for simplicity), and A and B mutually catalyze each other's formation from the food source. Figure 5 shows a graphical representation of this reaction network.




Figure 5: A simple AC set consisting of two reactions, with each product catalyzing the other's formation from a food source.

The changes in concentration of the molecule types A and B can now again be written as a set of (coupled) ODEs:

$$\frac{d}{dt}[A] = k[f][B] \quad (11)$$

$$\frac{d}{dt}[B] = k[f][A] \quad (12)$$

Assuming, as before, a constant (buffered) concentration of the food source ($[f] = c$), equation (11) can be rewritten as:

$$[B] = \frac{d}{dt} \frac{[A]}{kc} \quad (13)$$

Substituting equation (13) into equation (12) then gives:

$$\frac{d}{dt} \left(\frac{d}{dt} \frac{[A]}{kc} \right) = kc[A] \quad (14)$$

which can be rewritten as:

$$\frac{d^2}{dt^2}[A] - (kc)^2[A] = 0 \quad (15)$$

Equation (15) is a second-order linear differential equation in the variable $[A]$. Since $(kc)^2 > 0$, this equation has as solution:

$$[A](t) = re^{kct} + se^{-kct} \quad (16)$$

Finally, substituting (16) into (13) and taking derivatives gives:

$$[B](t) = re^{kct} - se^{-kct} \quad (17)$$

The variables r and s in equations (16) and (17) can be determined from the initial conditions $[A](0)$ and $[B](0)$. If they are equal, i.e., $[A](0) = [B](0) = C_0$, then from (16) and (17) it follows that:

$$C_0 = r + s \quad (18)$$

$$C_0 = r - s \quad (19)$$

which results in $r = C_0$ and $s = 0$. In this case, equations (16) and (17) simplify to:

$$[A](t) = [B](t) = C_0e^{kct} \quad (20)$$

So, if there is complete symmetry in rate constants and initial concentrations, both $[A]$ and $[B]$ behave as if they were independently autocatalytic, as in equation (4).

However, when the initial concentrations are different, e.g., $[A](0) = C_1$ and $[B](0) = C_2$, then there is also a difference in their dynamical behavior. From equations (16) and (17) it now follows that:

$$C_1 = r + s \quad (21)$$

$$C_2 = r - s \quad (22)$$

which has as solutions:

$$r = (C_1 + C_2)/2 \quad (23)$$

$$s = (C_1 - C_2)/2 \quad (24)$$

This then results in

$$[A](t) = \frac{C_1 + C_2}{2}e^{kct} + \frac{C_1 - C_2}{2}e^{-kct} \quad (25)$$

and

$$[B](t) = \frac{C_1 + C_2}{2} e^{kct} - \frac{C_1 - C_2}{2} e^{-kct} \quad (26)$$

So, in this case the dynamics of $[A]$ and $[B]$ are mutually dependent, but they both still have the ability to grow exponentially. Figure 6 shows this (initial) exponential growth over time for $[A]$ (solid black line) and $[B]$ (dashed black line) according to equations (25) and (26), with a rate constant $k = 1.0$, a constant food concentration $c = 0.05$, and initial concentrations $C_1 = 0.05$ and $C_2 = 0.00$. In other words, there are no B molecules present initially, but there are some A molecules, which will start catalyzing the formation of B , which then in turn catalyze the formation of more A , and so on. The red lines in Figure 6 show the result of a Gillespie simulation of the same system with similar parameter values, and again rescaled to match the ODE solution.

With this simple 2-reaction AC set, where both products are directly produced from the food set f , it is still possible to write down and solve a set of ODEs. However, for more complicated AC sets, where some products are multiple reaction-steps away from the food set, this becomes difficult or even impossible to do. However, the Gillespie algorithm can still be used, even for large AC sets. Figure 7 shows an example of such a more complicated RAF set that was found in a simple model of chemical reaction networks where molecules are represented by bit strings that can be ligated into longer bit strings or cut into shorter ones, and where catalysis is assigned randomly. In this example, the food set consists of the monomers and dimers (i.e., bit strings of lengths one and two), and the longer bit strings are built up from the food set by a sequence of one or more reactions, each of which is catalyzed by one of the molecules from the RAF set itself.

Note that in this case not all reaction products will grow at an exponential rate. For example, molecule 100 is produced by one reaction, but consumed again by another reaction. However, the “end products”, such as molecules 00100 or 11100, which are not used as reactants, can (potentially) still grow exponentially in concentration (provided an unlimited supply of food molecules), due to the “collectively” autocatalytic nature of the reaction network.

Furthermore, it turns out that RAF sets often consist of a hierarchy of smaller and smaller RAF subsets (Hordijk et al., 2012). This is indicated in Figure 7 by the colored outlines. For example, the yellow subset will always exist as long as there are food molecules around, as the reactions




Figure 6: The (initial) exponential growth in concentration of the products A (solid lines) and B (dashed lines) of the autocatalytic set shown in Figure 5, resulting from a solution to the ODE representation (black lines) or a Gillespie simulation (red lines).

within that subset only use food molecules as their reactants and catalysts. However, both the red and the blue subset need at least one reaction to happen spontaneously (i.e., uncatalyzed) for them to come into existence (as with a single AC reaction when the catalyst is not initially present). Once the blue subset exists it can be extended with the green subset, which also requires a spontaneous reaction.

So, even though all these different RAF subsets are already present in the underlying reaction network, dynamically they could exist in different combinations in different simulation runs or spatially separated locations, depending on which of the required spontaneous reactions happen first, if at all. A short movie showing this, resulting from simulating the RAF set in Figure 7 using the Gillespie algorithm, can be found in Hordijk (2016a). A




Figure 7: An example of a more complicated RAF set where the molecules are represented by bit strings, with the food set consisting of monomers and dimers. All reactions are catalyzed by one of the molecules from the set, and all longer molecules (bit strings) can be built up from the food set through a sequence of reactions from the set itself. The colored outlines indicate RAF subsets within the larger RAF set.

general and more complete overview of RAF theory and its main results is presented in Hordijk and Steel (2017).

2.4. Hypercycles

We will briefly mention the concept of hypercycles, which combine properties of AC reactions and AC cycles. A *hypercycle* is a collection of self-replicating macromolecules (i.e., they each catalyze their own formation from a food source), where each molecule, in addition, also catalyzes the replication of the next molecule, in a closed cyclic manner (Eigen, 1971). The idea behind this is that if, for example due to mutations, one macromolecule loses its ability to self-replicate, it can still be formed through a reaction catalyzed by the preceding molecule in the cycle, thereby maintaining the integrity of the cycle as a whole. Indeed, the concept of a hypercycle was introduced as a possible way to overcome the *error threshold* (Eigen, 1971).

We will return to hypercycles later on, after discussing the similarities and differences between the various autocatalytic systems.

3. Similarities and differences

As the mathematical (ODE) analyses showed, AC reactions, cycles, and sets all have the potential for exponential growth of one or more of their products. In fact, this is usually the property of interest in these systems. However, there are also clear differences between these systems, in particular in the way they achieve this exponential growth. They should therefore not be confused with each other, as is, unfortunately, too often done.

The most obvious difference is that an AC reaction is just one single reaction, whereas an AC cycle and an AC set generally consist of multiple reactions. In principle an AC set could also consist of just a single reaction, though, and an AC reaction for which all reactants are in the food set f is an example of the simplest possible AC set. However, an AC set is, in general, not simply a set of AC reactions (a subtle but important difference). Moreover, an AC reaction for which at least one of the reactants is not in the food set does not form an AC set (as it is not F-generated). An example is shown in Figure 7, where the green subset consists of a single AC reaction (the product 01111 catalyzes its own formation). However, one of the reactants is 111, which is not in the food set (i.e., it is not a monomer or dimer). So, by itself this AC reaction is not an AC set, but it can form an extension to the blue subset, which is an AC set and which produces molecule type 111.

Another important difference between the three systems is that, even though an AC cycle can be written in short-hand as a single AC reaction, none of the actual reactions in an AC cycle need necessarily be catalyzed. For example, none of the reactions in the formose reaction cycle is explicitly catalyzed by anything. However, without any catalysts this reaction cycle does not proceed at an appreciable rate (Orgel, 2000), although it does run more efficiently with the addition of various mineral catalysts (Schwartz and de Graaf, 1993). Similarly, even though the reverse citric acid cycle is catalyzed by highly evolved enzymes in modern-day organisms, there is evidence that several of its reactions can also be catalyzed by certain minerals (Zhang and Martin, 2006).

Note, though, that in these AC cycles, even if the individual reactions are catalyzed, they are not catalyzed by products of the cycle itself. This is in stark contrast with AC reactions and AC sets. In an AC reaction, one of the products catalyzes its own formation, and in an AC set the reaction products mutually catalyze each other's formation (although in an AC set some of the catalysts may also come from the food set, such as certain minerals). So,

whereas an AC cycle is cyclic in the sense that it ends up with the same molecule type that it started with, an AC set is cyclic in the sense of being *catalytically closed*.

More precisely, an AC set (or RAF set) usually contains at least one closed catalytic loop (Contreras et al., 2011; Steel et al., 2013). An example is shown in Figure 8, which shows the so-called *catalysis graph* of the simple 2-reaction RAF set in Figure 5. In such a graph the nodes represent molecule types, and there is an arrow from node i to node j if molecule type i catalyzes a reaction that produces molecule type j . In the simple 2-reaction RAF set, molecule types A and B mutually catalyze each other’s formation, so there is an arrow from A to B and vice versa, forming a closed catalytic loop. Such a closed loop in the catalysis graph is what is called a *viable core* in Vasas et al. (2012).



Figure 8: The catalysis graph of the simple 2-reaction RAF set in Figure 5. The two molecule types A and B mutually catalyze each other’s formation, creating a closed loop in the catalysis graph.

However, RAF sets usually consist of more than catalytic loops alone. In particular, a closed catalytic loop within a RAF set can support (provide catalysis for) reaction paths branching out from the closed loop. Such “branches” are called the *periphery* in Vasas et al. (2012). Figure 7 shows an example of this, where the formation of molecule type 00100 (produced within the red subset) is catalyzed by another molecule from the set, but it does not catalyze anything in return. However, it is an “end product” of the RAF set, and can (in principle) grow at an exponential rate.

On the other hand, the existence of a closed loop in the catalysis graph does not necessarily mean that there exists a RAF set in the corresponding reaction network (Steel et al., 2013). In particular, a closed catalytic loop may not be F-generated, and therefore not represent a proper RAF set. For example, closed loops (cycles) in the catalysis graph were observed in certain simulation studies, but were found to be unstable, as they were indeed not F-generated (Filisetti et al., 2011). This confusion between closed catalytic

loops and proper RAF sets was later resolved (Filisetti et al., 2014).

Finally, coming back to the notion of a hypercycle, such a system is actually a very specific instance of an AC set. In particular, it is an AC set where each molecule catalyzes exactly two reactions: its own formation and that of the next molecule in the cycle. However, this seems a rather tough requirement, which is probably why there are no known chemical examples of hypercycles. Unfortunately, though, there is also much confusion between strict hypercycles and the more general notion of AC sets. This “conceptual error” was already addressed in detail by Szathmáry (2013), so we will not elaborate on it here.

In short, there are important differences between AC reactions, cycles, and sets, and they should not be confused with each other. Even though they are all (potentially) capable of generating exponential growth, the ways in which they do so are different, giving rise to different internal network structures and dynamics.

4. Relevance for origin of life scenarios

Note that an autocatalytic system is a *self-replicating* system. Given an appropriate food source, a molecule (or set of molecules) that catalyzes its own formation, without being “used up” in that process, thus creates identical copies of itself. Since biological systems are self-replicators, chemical autocatalytic systems are obviously of great interest in the context of the *origin* of life, i.e., the transition from chemistry to biology.

Currently the dominant paradigm in origin of life research is that of an RNA world (Gilbert, 1986; Joyce, 2002), where the idea is that life started with one or more self-replicating RNA molecules. In other words, in such a world *AC reactions* would have played a dominant role (i.e., $\text{RNA} + \text{food} \rightarrow 2\text{RNA}$). Such self-replication would be *template-based*, i.e., it is the exact RNA sequence that is being replicated, nucleotide by nucleotide. However, despite experimental progress towards the spontaneous formation of RNA (Powner et al., 2009; Hud et al., 2013; Patel et al., 2015), the RNA world hypothesis still has significant problems (Benner et al., 2012; Szostak, 2012), and so far no one has been able to show that RNA can indeed catalyze its own template-directed replication.

An alternative to this “genetics first” view is that of a “metabolism first” view. For example, Morowitz et al. (2000) suggested that the (reverse) citric acid cycle may have emerged spontaneously from organic chemistry and

served as a central starting point for the origin of life. As already mentioned above, even though this metabolic cycle is catalyzed by complex enzymes in modern-day organisms, there is evidence that at least several of its reactions can be catalyzed by minerals (Zhang and Martin, 2006). So, in this alternative view *AC cycles* would have played a dominant role.

Furthermore, Gánti (1975, 2003) uses AC cycles as the essential subsystems in his chemoton model. Although not strictly a model for the *origin* of life, chemotons are models of *minimal* cellular life, relying on well-synchronized interactions between several coupled AC cycles.

However, what is missing in the chemoton model is explicit catalysis. Catalysts are ubiquitous in living systems, and life probably could not exist without them (Szöke et al., 2003). As mentioned above, simple AC cycles could (initially) be catalyzed by minerals, but an important property of living systems is that they produce their *own* catalysts and, moreover, these catalysts mutually catalyze *each other's* formation. This is exactly what allows living systems to evolve, diversify, and become more complex (Szöke et al., 2003; Ruiz-Mirazo et al., 2017).

So, what we suggest here is to replace the AC cycles in Ganti's chemoton model with AC sets, thus explicitly including catalysis while maintaining the potential for exponential growth. It has been shown that RAF sets have a high likelihood of existing in random chemistries, also for a chemically plausible level of catalysis (Hordijk and Steel, 2004, 2017). Furthermore, as mentioned earlier, there are several chemical examples of AC sets (Sievers and von Kiedrowski, 1994; Kim and Joyce, 2004; Ashkenasy et al., 2004; Vaidya et al., 2012), and it has even been shown that the metabolic network of *E. coli* forms a large RAF set (Sousa et al., 2015).

This, then, provides a further alternative for a possible origin of life scenario, where *AC sets* play a dominant role (Nghe et al., 2015). Small AC sets could have initially formed from prebiotic chemistry using inorganic elements (such as minerals and metals) as some of the original catalysts. As these initial AC sets evolve and become more complex (Hordijk et al., 2012; Vasas et al., 2012; Hordijk and Steel, 2014; Hordijk, 2016b), they produce more and increasingly efficient catalysts on their own, which over time will take over from the earlier less efficient catalysts, and so on. Support for such a scenario may exist in the fact that (i) many organic reactions can be catalyzed by inorganic elements (Schwartz and de Graaf, 1993; Zhang and Martin, 2006), (ii) many modern-day enzymes still use such inorganic elements as their cofactors (Christen and Mehta, 2001; Rees and Howard, 2003), and (iii) the metabolic

networks of living systems indeed seem to form AC sets (Kun et al., 2008; Sousa et al., 2015).

Finally, note that the square-root law of autocatalysis, i.e., the experimentally observed sub-exponential growth of autocatalytic systems, may be important here. A strict Darwinian “survival of the fittest” scenario requires exponential growth. However, with hyperbolic growth the co-existence of two or more replicators is also possible, rather than one outcompeting all others (Szathmáry and Gladkih, 1989). Co-existence, in turn, may have been crucial for (initial) AC sets to emerge, function, and evolve into more complex ones.

Acknowledgments

The author thanks the Konrad Lorenz Institute, Klosterneuburg, for financial support in the form of a fellowship, an anonymous reviewer for helpful suggestions to improve the manuscript, and his autocatalytic network of collaborators for the many discussions that helped shape the ideas presented in this paper.

References

- Ashkenasy, G., Jegasia, R., Yadav, M., Ghadiri, M.R., 2004. Design of a directed molecular network. *PNAS* 101, 10872–10877.
- Benner, S.A., Kim, H.J., Yang, Z., 2012. Setting the stage: the history, chemistry, and geobiology behind RNA. *Cold Spring Harbour Perspectives in Biology* 4, a003541.
- Christen, P., Mehta, P.K., 2001. From cofactor to enzymes: The molecular evolution of pyridoxal-5'-phosphate-dependent enzymes. *The Chemical Record* 1, 436–447.
- Contreras, D.A., Pereira, U., Hernández, V., Reynaert, B., Letelier, J.C., 2011. A loop conjecture for metabolic closure, in: *Advances in Artificial Life, ECAL 2011: Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems*, MIT Press. pp. 176–183.
- Dadon, Z., Wagner, N., Ashkenasy, G., 2008. The road to non-enzymatic molecular networks. *Angewandte Chemie International Edition* 47, 6128–6136.

- Eigen, M., 1971. Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften* 58, 465–523.
- Filisetti, A., Graudenzi, A., Serra, R., Villani, M., De Lucrezia, D., Fuchslin, R.M., Kauffman, S.A., Packard, N., Poli, I., 2011. A stochastic model of the emergence of autocatalytic cycles. *Journal of Systems Chemistry* 2, 2.
- Filisetti, A., Villani, M., Damiani, C., Graudenzi, A., Roli, A., Hordijk, W., Serra, R., 2014. On RAF sets and autocatalytic cycles in random reaction networks. *Communications in Computer and Information Science* 445, 113–126.
- Gánti, T., 1975. Organization of chemical reactions into dividing and metabolizing units: the chemotons. *BioSystems* 7, 15–21.
- Gánti, T., 2003. *The Principles of Life*. Oxford University Press, Oxford.
- Gilbert, W., 1986. The RNA world. *Nature* 319, 618.
- Gillespie, D.T., 1976. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics* 22, 403–434.
- Gillespie, D.T., 1977. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry* 81, 2340–2361.
- Hordijk, W., 2016a. Autocatalytic sets arising. <http://www.WorldWideWanderings.net/blog/?p=561>.
- Hordijk, W., 2016b. Evolution of autocatalytic sets in computational models of chemical reaction networks. *Origins of Life and Evolution of Biospheres* 46, 233–245.
- Hordijk, W., Kauffman, S.A., Steel, M., 2011. Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems. *International Journal of Molecular Sciences* 12, 3085–3101.
- Hordijk, W., Steel, M., 2004. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology* 227, 451–461.

- Hordijk, W., Steel, M., 2014. Conditions for evolvability of autocatalytic sets: A formal example and analysis. *Origins of Life and Evolution of Biospheres* 44, 111–124.
- Hordijk, W., Steel, M., 2017. Chasing the tail: The emergence of autocatalytic networks. *BioSystems* 152, 1–10.
- Hordijk, W., Steel, M., Kauffman, S., 2012. The structure of autocatalytic sets: Evolvability, enablement, and emergence. *Acta Biotheoretica* 60, 379–392.
- Hud, N.V., Cafferty, B.J., Krishnamurthy, R., Williams, L.D., 2013. The origin of RNA and "my grandfather's axe". *Chemistry & Biology* 20, 466–474.
- Issa, I.M., El-Aasser, A.A., El-Merzibani, M.M., 1960. Low temperature oxidation of oxalic acid with potassium permanganate. *Analytical Chemistry* 178, 12–19.
- Joyce, G.F., 2002. The antiquity of RNA-based evolution. *Nature* 418, 214–221.
- Kauffman, S.A., 1971. Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. *Journal of Cybernetics* 1, 71–96.
- Kauffman, S.A., 1986. Autocatalytic sets of proteins. *Journal of Theoretical Biology* 119, 1–24.
- Kauffman, S.A., 1993. *The Origins of Order*. Oxford University Press.
- von Kiedrowski, G., 1993. Minimal replicator theory I: parabolic versus exponential growth, in: Berlin, D.H. (Ed.), *Bioorganic Chemistry*, Springer-Verlag. pp. 115–146.
- Kim, D.E., Joyce, G.F., 2004. Cross-catalytic replication of an RNA ligase ribozyme. *Chemistry & Biology* 11, 1505–1512.
- Kun, A., Papp, B., Szathmáry, E., 2008. Computational identification of obligatorily autocatalytic replicators embedded in metabolic networks. *Genome Biology* 9, R51.

- Morowitz, H.J., Kostelnik, J.D., Yang, J., Cody, G.D., 2000. The origin of intermediary metabolism. *PNAS* 97, 7704–7708.
- Nghe, P., Hordijk, W., Kauffman, S.A., Walker, S.I., Schmidt, F.J., Kemble, H., Yeates, J.A.M., Lehman, N., 2015. Prebiotic network evolution: Six key parameters. *Molecular BioSystems* 11, 3206–3217.
- Orgel, L.E., 2000. Self-organizing biochemical cycles. *PNAS* 97, 12503–12507.
- Patel, B.H., Percivalle, C., Ritson, D.J., Duffy, C.D., Sutherland, J.D., 2015. Common origins of RNA, protein and lipid precursors in a cyanosulfidic protometabolism. *Nature Chemistry* 7, 301–307.
- Paul, N., Joyce, G.F., 2004. Minimal self-replicating systems. *Current Opinion in Chemical Biology* 8, 634–639.
- Powner, M.W., Gerland, B., Sutherland, J.D., 2009. Synthesis of activated pyrimidine ribonucleotides in prebiotically plausible conditions. *Nature* 459, 239–242.
- Rees, D.C., Howard, J.B., 2003. The interface between the biological and inorganic worlds: Iron-sulfur metalloclusters. *Science* 300, 929–931.
- Ruiz-Mirazo, K., Briones, C., de la Escosura, A., 2017. Chemical roots of biological evolution: The origins of life as a process of development of autonomous functional systems. *Open Biology* 7, 170050.
- Schwartz, A.W., de Graaf, R.M., 1993. The prebiotic synthesis of carbohydrates: A reassessment. *Journal of Molecular Evolution* 36, 101–106.
- Sievers, D., von Kiedrowski, G., 1994. Self-replication of complementary nucleotide-based oligomers. *Nature* 369, 221–224.
- Sousa, F.L., Hordijk, W., Steel, M., Martin, W.F., 2015. Autocatalytic sets in *E. coli* metabolism. *Journal of Systems Chemistry* 6, 4.
- Steel, M., Hordijk, W., Smith, J., 2013. Minimal autocatalytic networks. *Journal of Theoretical Biology* 332, 96–107.
- Szathmáry, E., 2000. The evolution of replicators. *Philosophical Transactions of the Royal Society of London B* 355, 1669–1676.

- Szathmáry, E., 2013. On the propagation of a conceptual error concerning hypercycles and cooperation. *Journal of Systems Chemistry* 4, 1.
- Szathmáry, E., Gladkih, I., 1989. Sub-exponential growth and coexistence of non-enzymatically replicating templates. *Journal of Theoretical Biology* 138, 55–58.
- Szöke, A., Scott, W.G., Hajdu, J., 2003. Catalysis, evolution and life. *FEBS Letters* 553, 18–20.
- Szostak, J.W., 2012. The eightfold path to non-enzymatic RNA replication. *Journal of Systems Chemistry* 3, 2.
- Vaidya, N., Manapat, M.L., Chen, I.A., Xulvi-Brunet, R., Hayden, E.J., Lehman, N., 2012. Spontaneous network formation among cooperative RNA replicators. *Nature* 491, 72–77.
- Vasas, V., Fernando, C., Santos, M., Kauffman, S., Sathmáry, E., 2012. Evolution before genes. *Biology Direct* 7, 1.
- Wikipedia, 2017. Autocatalysis. <https://en.wikipedia.org/wiki/Autocatalysis>. Last accessed: July 2017.
- Wolfenden, R., Snider, M.J., 2001. The depth of chemical time and the power of enzymes as catalysts. *Accounts of Chemical Research* 34, 938–945.
- Zhang, X.V., Martin, S.T., 2006. Driving parts of Krebs cycle in reverse through mineral photochemistry. *Journal of the American Chemical Society* 128, 16032–16033.