

Chapter 4

The Origin of Life, Evolution, and Functional Organization

Wim Hordijk, Mike Steel and Stuart Kauffman

Abstract The idea that autocatalytic sets played an important role in the origin of life is not new. Neither is the idea that autocatalytic sets can tell us something about the evolution and functional organization of living systems. However, most of these ideas have, until recently, remained at a conceptual level, and very few concrete, mathematically sound, and practically applicable results had been achieved. In this chapter, we review and discuss recent results on a mathematical framework of autocatalytic sets that could take the idea out of the conceptual realm, and provide a formal and powerful way to investigate autocatalytic sets in the context of the origin of life, evolution, and functional organization.

4.1 Introduction

Origin of life research seems divided between two paradigms: *genetics-first* and *metabolism-first*. However, one element that both these paradigms have in common is that of *autocatalysis*: molecules catalyzing their own or each others replication, in a mutually beneficial way. The idea of such *autocatalytic sets* could form a bridge between the opposing paradigms.

Autocatalytic sets are generally defined as a set of molecules and catalyzed chemical reactions between them, where each molecule type is created by at least

W. Hordijk (✉)
SmartAnalytiX.com, Lausanne, Switzerland
e-mail: wim@WorldWideWanderings.com

M. Steel
Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand
e-mail: mike.steel@canterbury.ac.nz

S. Kauffman
University of Vermont, Burlington, VT, USA
e-mail: stukauffman@gmail.com

one reaction from this set, and each reaction is catalyzed by at least one molecule type from the set. Moreover, an autocatalytic set is usually required to be *self-sustaining* in the sense that every molecule in the set can be built up from some ambient ‘food’ source (molecules assumed to be available from the environment), by using reactions only from within the autocatalytic set itself.

This idea of collectively autocatalytic sets is not new. It was introduced more or less independently several times from the 1970s onward (Kauffman 1971, 1986; Eigen and Schuster 1977, 1979; Dyson 1982, 1985), and used in various origin of life models and scenarios (Wächterhäuser 1990, 2007; Gánti 1997; Rosen 1991; Letelier et al. 2006). Of course the idea has received criticism as well (Lifson 1997; Orgel 2008; Vasas et al. 2010), but recent experimental evidence shows the possibility and viability of such sets (Sievers and von Kiedrowski 1994; Lee et al. 1997; Ashkenasy et al. 2004; Hayden et al. 2008; Taran et al. 2010; Vaidya 2012).

In some of the original arguments to support the idea, it was claimed that in sufficiently complex chemical reaction systems (CRSs) (i.e., with a large enough diversity of molecule types), autocatalytic sets will arise almost inevitably (Kauffman 1971, 1986, 1993). This was subsequently disputed by pointing out that this argument actually requires an exponential growth rate in catalytic activity with increasing system size, something which is chemically highly unrealistic (Lifson 1997). So, the question of the probability of autocatalytic sets arising spontaneously in arbitrary CRSs remained open, with arguments both for and against mostly being based on informal qualitative arguments, and a thorough mathematical analysis was still lacking.

In more recent work on autocatalytic sets (Steel 2000; Hordijk and Steel 2004, 2012b, a; Mossel and Steel 2005; Hordijk et al. 2010, 2011, 2012), we have placed these ideas within a firm mathematical framework, and investigated the probability and structure of autocatalytic sets both mathematically and computationally. Not only has this led to a confirmation of the original claim (although with a more complicated mathematical proof), it has also provided many new and interesting insights into the conditions under which autocatalytic sets are likely to emerge, and their actual structure such as size and composition. Furthermore, our formal framework includes an efficient algorithm for actually detecting autocatalytic sets in arbitrary CRSs (an important tool that was not available before).

Given the recent experimental, theoretical, and computational progress, it seems more and more likely that autocatalytic sets played an important role, perhaps in more than one way, in the origin of life (Morowitz et al. 2000; Braakman and Smith 2012; Martin and Russel 2007; Penny 2005; Lane 2009; Hordijk et al. 2010; Kauffman 2011). Furthermore, we claim that autocatalytic sets might also play an important role in the evolution and functional organization of living systems in general. As such, we argue for a generalized theory of autocatalytic sets that could also be applied to living cells, entire ecologies, and perhaps even the economy.

In this chapter, we provide a non-technical overview of our formal framework for autocatalytic sets, some of the main results obtained with it so far, and our arguments for how it might be generalized and applied beyond the origin of life.

For full details, we refer to the original publications (Kauffman 1986, 1993; Steel 2000; Hordijk and Steel 2004, 2012a, b; Hordijk et al. 2010, 2011, 2012).

4.2 A Formal Framework of Autocatalytic Sets

4.2.1 Chemical Reaction Systems

In our framework, a *CRS* is defined as a set X of molecule types, a set \mathcal{R} of possible chemical reactions between these molecule types (transforming reactants into products), and a catalysis set C indicating which molecule types can catalyze which reactions. A catalyst speeds up the rate at which a reaction happens (sometimes by several orders of magnitude), without being “used up” in the reaction itself. Mathematically, a CRS is written as a collection $Q = \{X, \mathcal{R}, C\}$. We also include the notion of a *food set* F , which is a (small) subset of molecule types that can be assumed freely available in the environment.

When using mathematical models to describe real-world systems, it is desirable to start with the simplest type of model that still captures the essence of the behavior of interest in the real system. Here, we use a simple model of CRSs known as the *binary polymer model* (Kauffman 1986, 1993). In this model, the molecule types are represented by bit strings up to a certain length n . The food set F consists of a few small bit strings, for example all bit strings of length one and two (i.e., six food molecules in total). The possible reactions are simply *ligation* and *cleavage*: two bit strings are “glued” together into one longer one (ligation), or one bit string is “cut” into two smaller pieces (cleavage).

The catalysis set C in the binary polymer model is generated at random as follows. Assume a probability p that a given molecule $x \in X$ catalyzes a given reaction $r \in \mathcal{R}$. To generate an instance of the binary polymer model, each possible molecule-reaction pair $(x, r) \in (X, \mathcal{R})$ is included in the catalysis set C independently with probability p . The main idea behind modeling catalysis at random is that, overall, little is known about which molecules can catalyze which reactions [predicting catalysis is known to be a hard problem (Kayala et al. 2011)].

Together, the set X of bit strings up to length n , the set \mathcal{R} of possible ligation and cleavage reactions, and a catalysis set C generated at random with probability p , form a CRS $Q = \{X, \mathcal{R}, C\}$, thus providing a simple model of CRSs with tunable parameters n and p .

4.2.2 Autocatalytic Sets

Given a CRS $Q = \{X, \mathcal{R}, C\}$ and a food set $F \subset X$, an *autocatalytic set* (or RAF set, in our terminology) is now defined as a subset $\mathcal{R}' \subseteq \mathcal{R}$ of reactions (and associated molecule types) which is:

1. *reflexively autocatalytic* (RA): each reaction $r \in \mathcal{R}'$ is catalyzed by at least one molecule type involved in \mathcal{R}' , and
2. *food-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.

For a full mathematical definition of RAF sets, see Hordijk and Steel (2004); Hordijk et al. (2011), where we also introduced an efficient algorithm for detecting RAF sets in a general CRS \mathcal{Q} .

Figure 4.1 shows an example of an RAF set that was found by our algorithm in an instance of the binary polymer model with $n = 5$. Black dots indicate molecule types, and white boxes indicate reactions. Solid black arrows indicate reactants going into and products coming out of a reaction, and dashed gray arrows indicate catalysis.

The RAF sets that are found by the RAF algorithm are called *maximal* RAF sets (maxRAFs). However, as we will show below, a maxRAF can often be decomposed into several smaller subsets which themselves are RAF sets (subRAFs). If such a subRAF cannot be reduced any further without losing the RAF property, it is called an *irreducible* RAF (irrRAF).

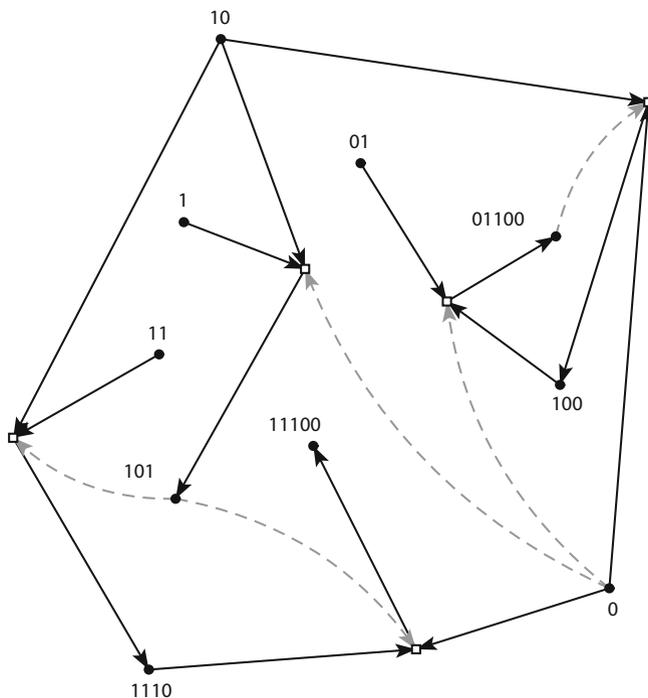


Fig. 4.1 An example of an RAF set that was found in an instance of the binary polymer model with $n = 5$, and food set $F = \{0, 1, 00, 01, 10, 11\}$

4.3 The Emergence and Evolution of Autocatalytic Sets

To investigate the probability, structure, and dynamics of autocatalytic sets in CRSs, we used the RAF framework and the binary polymer model to perform and analyze computational simulations and derive mathematical properties and theorems. Here, we present a brief review of the main results and insights obtained so far.

4.3.1 *The Probability of Autocatalytic Sets*

In Hordijk and Steel (2004) we showed computationally, and then proved mathematically in Mossel and Steel (2005), that at most a *linear* growth rate in the level of catalysis (with increasing maximum molecule length n) is sufficient for autocatalytic sets to exist with high probability. The level of catalysis is defined here as the average number of reactions catalyzed by any one molecule. Note that this result is in stark contrast with the required exponential growth rate in the original argument (Kauffman 1986, 1993). Furthermore, each molecule only needs to catalyze (on average) between one and two reactions to get a high probability of autocatalytic sets (for n up to 50, at least) (Hordijk and Steel 2004; Hordijk et al. 2010). Thus, it seems that autocatalytic sets are highly likely, even for very moderate levels of catalysis.

One might argue that the binary polymer model is perhaps too simple to say much about reach chemical systems. However, in recent experiments, an almost literal chemical implementation of the binary polymer model was used, which has the potential to form autocatalytic sets (Taran et al. 2010). Furthermore, the theoretical linear upper bound on the growth rate in the required level of catalysis also holds for alphabets other than binary, e.g., polymers consisting of four building blocks (such as nucleotides), or 20 (such as amino acids), or any number. Finally, we also investigated several more realistic model extensions, such as *template-based catalysis*. The idea here is that a molecule must match a certain number of bits around the reaction site to be considered a candidate catalyst (similar, for example, to base-pair bonding in RNA molecules). In this case, the linear bound on the required growth rate in the level of catalysis still holds (Hordijk et al. 2011). Moreover, the required level of catalysis in the template-based case can be predicted mathematically from the purely random case (Hordijk and Steel 2012b).

4.3.2 *The Structure of Autocatalytic Sets*

Having established the high likelihood of autocatalytic sets, we next investigated the structure of such sets. In particular, it turns out that RAF sets can often be decomposed into smaller RAF subsets (subRAFs) (Hordijk et al. 2012). For

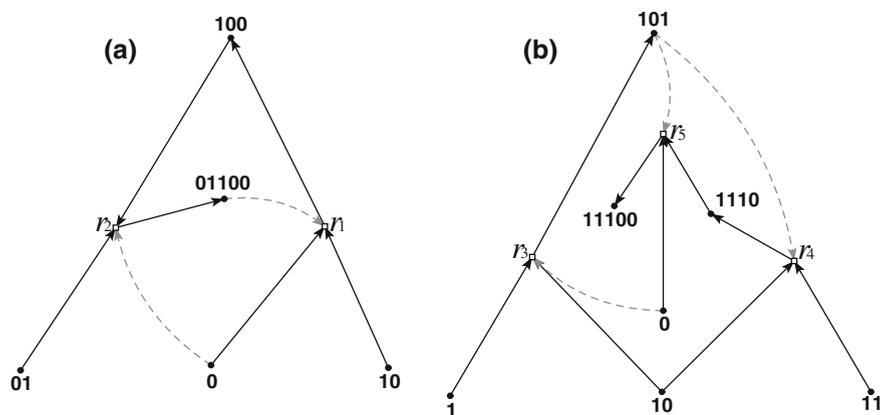


Fig. 4.2 The two independent subRAFs \mathcal{R}_1 (left) and \mathcal{R}_2 (right) of which the maxRAF shown in Fig. 4.1 is composed

example, consider the maxRAF shown in Fig. 4.1. This RAF set, consisting of five reactions, can be decomposed into two smaller and independent subRAFs \mathcal{R}_1 and \mathcal{R}_2 of size two and three, respectively. These subRAFs are shown in Fig. 4.2.

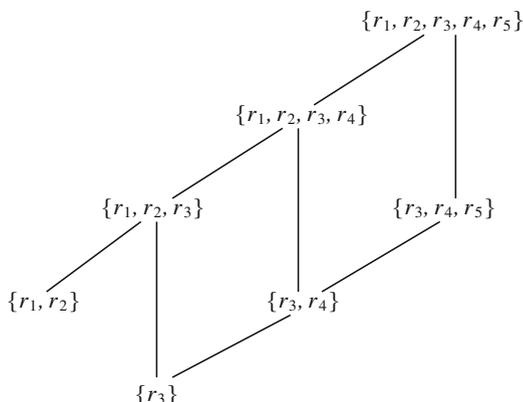
Furthermore, the subRAF \mathcal{R}_2 itself contains even smaller subRAFs. In fact, any given RAF set can (potentially) contain an exponentially large number of subRAFs. All these subRAFs together form a partially ordered set, which can be visualized in a so-called *Hasse diagram*. Such a diagram of the subRAFs of the RAF set in Fig. 4.1 is shown in Fig. 4.3.

This decomposability of RAF sets provides a possible mechanism for evolution to happen in autocatalytic sets (Vasas et al. 2012; Hordijk et al. 2012). The process of combining and recombining different subRAFs into various (compartmentalized) reaction sets, giving rise to inheritance with variation, selection, and thus *evolvability*, was convincingly shown to work in Vasas et al. (2012). Furthermore, in Hordijk et al. (2012) we also argued how such a process can give rise to the emergence of higher-level structures (autocatalytic sets of autocatalytic sets). In that sense, the Hasse diagram shows the possible “trajectories” by which autocatalytic sets can emerge and evolve.

4.3.3 The Dynamics of Autocatalytic Sets

To confirm these ideas about the emergence and evolution of autocatalytic sets, we also performed actual molecular flow simulations on RAF sets. Given a (max)RAF set \mathcal{R} and starting with only molecules from the food set F , we used the well-known Gillespie algorithm (Gillespie 1976, 1977) to simulate the flow of molecules on the given reaction network. What these simulations show, is that indeed

Fig. 4.3 The Hasse diagram of all subRAFs contained within the RAF set shown in Fig. 4.1



the reaction system generally starts out with smaller (irreducible) RAF sets, which over time grow and combine into larger ones (Hordijk and Steel 2012a). Moreover, the existence of some (sub)RAFs can give rise to the coming into existence of other (sub)RAFs. For example, a particular subRAF \mathcal{R}_B might be dependent on a molecule produced by another subRAF \mathcal{R}_A , and so \mathcal{R}_B can only come into existence once \mathcal{R}_A exists.

Figures 4.4 and 4.5 show a simple example [taken from Hordijk and Steel (2012a)]. Figure 4.4 shows another (maximal) RAF set that was found by the RAF algorithm in an instance of the binary polymer model. The colored boxes show the various subRAFs that exist within this maxRAF. Note that some boxes (subRAFs) are nested (such as the purple box within the red box), and some boxes depend on or form possible extensions of yet other boxes (such as the green box forming an extension of the blue box).

Figure 4.5 shows the results of a molecular flow simulation on the RAF network in Fig. 4.4. Starting with only food molecules, at first only the concentrations of the product molecules of the smallest (irreducible) subRAFs start to increase (the yellow and purple lines; line colors in Fig. 4.5 correspond to the colors of the boxes in Fig. 4.4). The blue subRAF actually needs to be “seeded”, i.e., one of its three reactions needs to happen “spontaneously” (uncatalyzed) at least once for it to come into existence. Such an uncatalyzed reaction is always possible, but at a much lower rate than a catalyzed reaction. However, around time point 0.3 one such spontaneous reaction happens, and the blue subRAF comes into existence. After a while (around time point 0.5), the green extension of the blue subRAF (on which it depends) also comes into existence. Finally, the purple subRAF grows into the red subRAF after a last necessary spontaneous reaction happens around time point 0.55.

Furthermore, additional simulations on simple (collaborative) autocatalytic sets and equivalent but (selfish) non-autocatalytic sets of reactions confirm that autocatalytic sets do indeed have an advantage in terms of outperforming their non-autocatalytic competitors and in being more robust against perturbations.

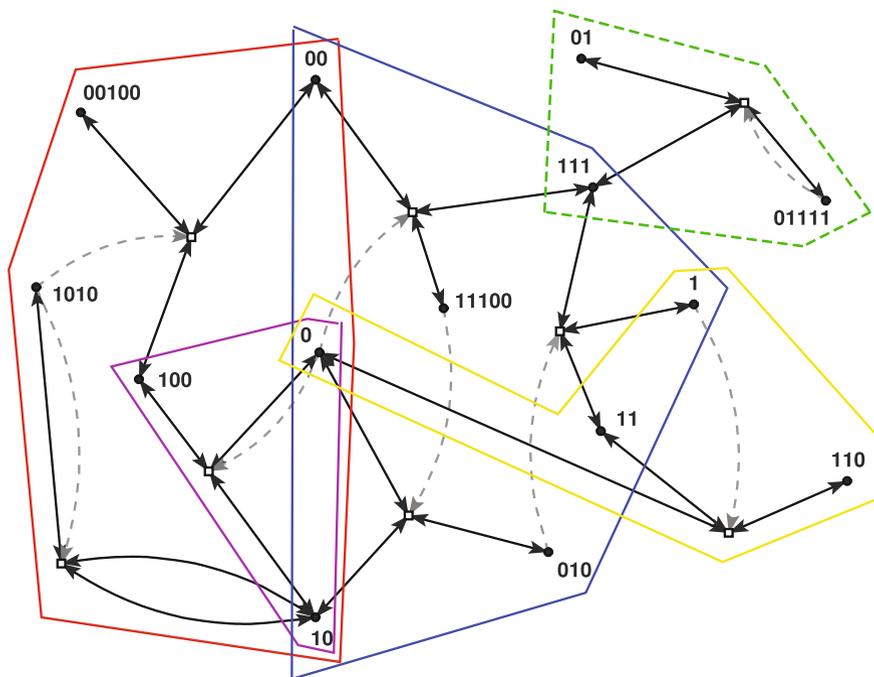


Fig. 4.4 Another example of a (maximal) RAF set found by the RAF algorithm in an instance of the binary polymer model. The *colored boxes* indicate the various RAF subsets that are contained within the maxRAF

4.3.4 The Importance of Autocatalytic Sets

In conclusion, using the formal RAF framework we have shown that autocatalytic sets are highly likely to emerge in CRSs, that they have the ability to evolve, and that they can outcompete, and are more robust than, equivalent non-autocatalytic sets. Note, however, that the examples shown here are deliberately kept small, so they can be visually inspected and understood. As a consequence, though, they are limited in their ability to show the full extent of possible structures, inter-dependencies, and dynamics that can (and do) occur in larger instances. But we trust that the simple examples provided here are sufficient to convey the main results and implications, at least conceptually.

Moreover, recent results show that the RAF framework can be applied directly and meaningfully to *real* CRSs exhibiting autocatalytic sets (Hordijk and Steel 2013) [in this case an RNA replicator system (Vaidya 2012)]. Not only does the model replicate many of the original experimental results, it also provides additional insights and predictions that would be difficult or impossible to obtain from experiments alone. This forms an important step toward merging experimental and theoretical lines of work on autocatalytic sets and their emergence and evolution.

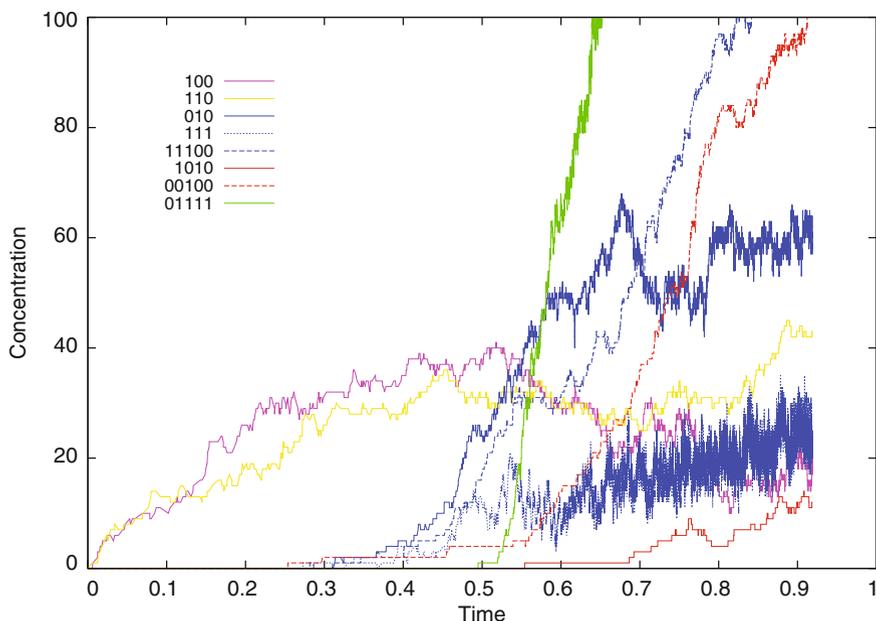


Fig. 4.5 The results of a molecular flow simulation on the maxRAF shown in Fig. 4.4. The *lines* (colors correspond to the boxes in Fig. 4.4) show the concentrations of the reaction products over time, starting from an initial concentration of food molecules only

As mentioned in the introduction, the general idea behind autocatalytic sets is not new, as it was already developed during the 1970s. In fact, this idea of catalytic (or functional) closure and self-sustainability as a necessary property for living systems is incorporated in several similar formalisms, such as autopoietic systems (Varela et al. 1974; Maturana and Varela 1980), (M,R) systems (Rosen 1991), the Chemoton model (Gánti 1997), and various others. However, many of these formalisms remained mostly at a conceptual level, and have not generated the same richness and depth of results as the RAF framework. As summarized in this brief chapter, with this new framework we have obtained clear and concrete insights into the emergence, evolution, and structure of autocatalytic sets, which has direct relevance for a possible mechanism for the origin of life.

4.4 A Generalized Theory of Autocatalytic Sets

The original RAF algorithm is specifically formulated in the context of CRSs. However, it is possible to state the algorithm in a more general form. This is useful for extending it in further directions, both within the context of chemical systems as well as for other applications.

In Hordijk and Steel (2012a), we presented a generalized RAF algorithm, which we refer to as the “*gf*–algorithm”. We then also showed that this algorithm can efficiently (i.e., in polynomial time) solve the HORN-SAT problem, which is a basic problem in propositional logic. HORN-SAT is of interest, as it is P-complete, which means that every problem in the computational complexity class P (i.e., every problem for which an efficient (polynomial-time) algorithm exists), can be reduced to HORN-SAT. Consequently, the *gf*–algorithm can efficiently solve any problem in the problem class P, and can thus be applied to a wide variety of problems other than finding autocatalytic sets in CRSs. One example was presented in Hordijk and Steel (2012a) where we described a simple toy problem in economics that can be solved with the *gf*–algorithm.

This generalization makes it clear that the RAF algorithm can be applied in a wider context outside of chemistry and beyond the origin of life. In a similar way, we would like to propose that the entire RAF framework, i.e., the notion of an autocatalytic set as a functionally closed, self-sustaining set, can be generalized. Given that autocatalytic sets of molecules and chemical reactions can grow, (re)combine, evolve, and potentially give rise to the emergence of higher-level functionally closed structures, we argue that a living cell could perhaps also be considered an autocatalytic (super)set, consisting of several inter-dependent autocatalytic subsets. Taking this a step further, what about collections of inter-dependent living cells, such as bacterial colonies or multicellular organisms? And finally, to take the argument the whole way, perhaps even an entire ecology of organisms can be viewed in this way.

Moreover, if interdependent groups of cells or multicellular organisms can be viewed as collections of interacting autocatalytic sets, then the same might hold for interdependent groups of people or organizations, i.e., an economy. For example, a “production function” (such as converting raw materials into products) can be viewed as the economic equivalent of a chemical reaction. The equivalent of catalysts are then products that facilitate these production functions, but are not “used up” by them, and which themselves are the result of other production functions. As a simple example, consider several pieces of wood and some nails as “reactants”, hammering them together as a “reaction”, and the result, say a table, as the “product”, with a hammer being the catalyst, itself the product of some other reaction. In other words, an economy as a functionally closed, self-sustaining autocatalytic set, where the coming into existence of one subset (e.g., the Internet) enables the coming into existence of other, new subsets that could not have existed before (e.g., a company such as Google).

These arguments are still rather speculative, but as mentioned in the previous section, this idea of thinking about living (and perhaps even social) systems as functionally closed, self-sustaining entities is not new. However, with the RAF framework, we claim to have a concrete, mathematically well-founded, and widely applicable formalism to investigate these ideas in a more direct, meaningful, and detailed manner. If a system of interest, be it chemical, biological, or social, can be defined within the context of the RAF framework, the tools we have developed can be applied directly to gain more insight into the emergence, evolution, and

structure of autocatalytic sets within these systems. Indeed, we are actively discussing and exploring these ideas with ecologists, economists, and social and cognitive scientists who are interested in finding ways to apply the RAF framework in their own fields of study. We truly hope that these combined efforts will eventually lead to a generalized theory of autocatalytic sets, and a better understanding of functional organization and emergence, not only in the context of the origin of life, but well beyond.

Acknowledgments This work was partly funded by the *Allan Wilson Centre for Molecular Ecology and Evolution* (AWCMEE).

References

- Ashkenasy G, Jegasia R, Yadav M, Ghadiri MR (2004) Design of a directed molecular network. *PNAS* 101(30):10872–10877
- Braakman R, Smith E (2012) The emergence and early evolution of biological carbon-fixation. *PLoS Comput Biol* 8(4):e1002455
- Dyson FJ (1982) A model for the origin of life. *J Mol Evol* 18:344–350
- Dyson FJ (1985) *Origins of life*. Cambridge University Press, Cambridge
- Eigen M, Schuster P (1977) The hypercycle: a principle of natural self-organization. Part A: emergence of the hypercycle. *Naturwissenschaften* 64:541–565
- Eigen M, Schuster P (1979) *The hypercycle*. Springer, Berlin
- Gánti T (1997) Biogenesis itself. *J Theor Biol* 187:583–593
- Gillespie DT (1976) A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J Comput Phys* 22:403–434
- Gillespie DT (1977) Exact stochastic simulation of coupled chemical reactions. *J Phys Chem* 81(25):2340–2361
- Hayden EJ, von Kiedrowski G, Lehman N (2008) Systems chemistry on ribozyme self-construction: Evidence for anabolic autocatalysis in a recombination network. *Angewandte Chemie Int Ed* 120:8552–8556
- Hordijk W, Hein J, Steel M (2010) Autocatalytic sets and the origin of life. *Entropy* 12(7):1733–1742
- Hordijk W, Kauffman SA, Steel M (2011) Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems. *Int J Mol Sci* 12(5):3085–3101
- Hordijk W, Steel M (2004) Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *J Theor Biol* 227(4):451–461
- Hordijk W, Steel M (2012) Autocatalytic sets extended: dynamics, inhibition, and a generalization. *J Syst Chem* 3:5
- Hordijk W, Steel M (2012) Predicting template-based catalysis rates in a simple catalytic reaction model. *J Theor Biol* 295:132–138
- Hordijk W, Steel M (2013) A formal model of autocatalytic sets emerging in an RNA replicator system. *J Syst Chem* 4:3
- Hordijk W, Steel M, Kauffman S (2012) The structure of autocatalytic sets: Evolvability, enablement, and emergence. *Acta Biotheor* 60(4):379–392
- Kauffman SA (1971) Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. *J Cybern* 1(1):71–96
- Kauffman SA (1986) Autocatalytic sets of proteins. *J Theor Biol* 119:1–24
- Kauffman SA (1993) *The origins of order*. Oxford University Press, Oxford
- Kauffman SA (2011) Approaches to the origin of life on earth. *Life* 1:34–48

- Kayala MA, Azencott C-A, Chen JH, Baldi P (2011) Learning to predict chemical reactions. *J Chem Inf Model* 51:2209–2222
- Lane N (2009) *Life ascending: the ten great inventions of evolution*. Profile Books, Norton
- Lee DH, Severin K, Ghadiri MR (1997) Autocatalytic networks: the transition from molecular self-replication to molecular ecosystems. *Curr Opin Chem Biol* 1:491–496
- Letelier J-C, Soto-Andrade J, Abarzúa FG, Cornish-Bowden A, Cárdenas ML (2006) Organizational invariance and metabolic closure: analysis in terms of (M;R) systems. *J Theor Biol* 238:949–961
- Lifson S (1997) On the crucial stages in the origin of animate matter. *J Mol Evol* 44:1–8
- Martin W, Russel MJ (2007) On the origin of biochemistry at an alkaline hydrothermal vent. *Philos Trans R Soc B* 362:1887–1925
- Maturana HR, Varela FJ (1980) *Autopoiesis and cognition: the realization of the living*. Reidel Publishing Co., Dordrecht
- Morowitz HJ, Kostelnik JD, Yang J, Cody GD (2000) The origin of intermediary metabolism. *PNAS* 97(14):7704–7708
- Mossel E, Steel M (2005) Random biochemical networks: the probability of self-sustaining autocatalysis. *J Theor Biol* 233(3):327–336
- Orgel LE (2008) The implausibility of metabolic cycles on the prebiotic earth. *PLoS Biology* 6(1):5–13
- Penny D (2005) An interpretive review of the origin of life research. *Biol Philos* 20:633–671
- Rosen R (1991) *Life itself*. Columbia University Press, New York
- Sievers D, von Kiedrowski G (1994) Self-replication of complementary nucleotide-based oligomers. *Nature* 369:221–224
- Steel M (2000) The emergence of a self-catalysing structure in abstract origin-of-life models. *Appl Math Lett* 3:91–95
- Taran O, Thoennessen O, Achilles K, von Kiedrowski G (2010) Synthesis of information-carrying polymers of mixed sequences from double stranded short deoxynucleotides. *J Syst Chem* 1(9)
- Vaidya N, Manapat ML, Chen IA, Xulvi-Brunet R, Hayden EJ, Lehman N (2012) Spontaneous network formation among cooperative RNA replicators. *Nature* 491:72–77
- Varela FJ, Maturana HR, Uribe R (1974) Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems* 5(4):187–196
- Vasas V, Fernando C, Santos M, Kauffman S, Sathmáry E (2012) Evolution before genes. *Biol Direct* 7:1
- Vasas V, Szathmáry E, Santos M (2010) Lack of evolvability in self-sustaining autocatalytic networks constrains metabolism-first scenarios for the origin of life. *PNAS* 107(4):1470–1475
- Wächterhäuser G (1990) Evolution of the first metabolic cycles. *PNAS* 87:200–204
- Wächterhäuser G (2007) On the chemistry and evolution of the pioneer organism. *Chem Biodivers* 4:584–602