

An investigation into irreducible autocatalytic sets and power law distributed catalysis

Wim Hordijk · Leonard Hasenclever · Jie Gao · Dilyana Mincheva · Jotun Hein

Abstract RAF theory has been established as a useful and formal framework for studying the emergence and evolution of autocatalytic sets. Here, we present several new and additional results on RAF sets. In particular, we investigate in more detail the existence, expected sizes, and composition of the smallest possible, or irreducible, RAF sets. Furthermore, we study a more realistic variant of the well-known binary polymer model in which the catalysis events are assigned according to a power law distribution. Together, these results provide further insights into the existence and structure of autocatalytic sets in simple models of chemical reaction systems, with possible implications for theories on the origin of life.

1 Introduction

RAF theory has been established as a useful and formal framework for studying the emergence and evolution of autocatalytic sets (Steel, 2000; Hordijk and Steel, 2004; Mossel and Steel, 2005; Hordijk et al, 2011; Hordijk and Steel, 2012b; Hordijk et al, 2012; Hordijk and Steel, 2012a, 2013; Steel et al, 2013; Hordijk et al, 2014; Smith et al, 2014), motivated by questions in the context of the origin of life (Hordijk et al, 2010; Hordijk, 2013; Hordijk et al, 2013). Here, we present additional and more detailed results on RAF theory, in particular on various aspects of so-called irreducible RAF sets, and on a more realistic variant of the original binary polymer model for chemical reaction systems.

Wim Hordijk
SmartAnalytiX.com, E-mail: wim@WorldWideWanderings.net

Leonard Hasenclever
University of Cambridge, E-mail: leonard.hasenclever@gmail.com

Jie Gao
University of North Carolina at Chapel Hill, E-mail: neojgao@gmail.com

Dilyana Mincheva
University of Oxford, E-mail: dillymincheva@gmail.com

Jotun Hein
University of Oxford, E-mail: hein@stats.ox.ac.uk

In Steel et al (2013) several theoretical results were established on the (expected) sizes of the smallest, or irreducible, autocatalytic (RAF) sets and the intractability of actually finding them. Some preliminary computational results were included as well to support the theoretical claims. In this paper, we investigate irreducible RAF sets in more detail and present various additional computational results. This provides useful empirical insights into the existence, expected sizes, and composition of the smallest possible RAF sets.

These investigations are done using the well-known binary polymer model (Kauffman, 1986, 1993). This model is sometimes criticized for a lack of chemical realism. However, there exist experimental systems in which the spontaneous emergence of autocatalytic sets has been shown or postulated, and which are basically a “chemical implementation” of (some variant of) the polymer model (Taran et al, 2010; Vaidya et al, 2012). In fact, such experimental systems can be modeled very accurately and efficiently using RAF theory (Hordijk and Steel, 2013). Moreover, chemically more realistic assumption can be easily built into the binary polymer model, such as template-based catalysis (Hordijk et al, 2011; Hordijk and Steel, 2012b; Hordijk et al, 2014). In this paper, we introduce and investigate another chemically more realistic variant of the binary polymer model, one in which catalysis is assigned according to a power law distribution, a signature characteristic of “real-world” networks. We show that incorporating this assumption does not change the overall results of the standard model in terms of the existence of autocatalytic sets. In fact, in this more realistic model variant RAF sets may be even more likely to emerge.

This paper is organized as follows. The next section provides a brief review of RAF theory and the binary polymer model. Section 3 then presents the results of investigating irreducible RAFs in more detail, while section 4 presents the results of using the binary polymer model with power law distributed catalysis. Finally, section 5 summarizes the main conclusions and suggests possible consequences for the role autocatalytic sets may have had in the origin and early evolution of life.

2 RAF theory

RAF theory is based on the original notion of autocatalytic sets as introduced by Kauffman (1971, 1986, 1993), in an attempt to formalize the idea of life as a functionally closed and self-sustaining chemical reaction system. This notion of autocatalytic sets is closely related to other such models and ideas (Gánti, 2003; Eigen and Schuster, 1979; Maturana and Varela, 1980; Dyson, 1985), and they are believed to have played a crucial role in the origin of life (Kauffman, 1993; Hordijk et al, 2010; Hordijk, 2013; Hordijk et al, 2013).

First, we define a *chemical reaction system* (CRS) as a tuple $Q = \{X, \mathcal{R}, C\}$ consisting of a set of molecule types X , a set of (possible or allowed) chemical reactions \mathcal{R} , and a catalysis set C indicating which molecule types can catalyze which reactions. Next, a food set $F \subset X$ is defined as a subset of molecule types that are assumed to be freely available from the environment (i.e., they do not necessarily have to be produced by any of the reactions in \mathcal{R}). Finally, an *autocatalytic set* (or RAF set) is now informally 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.

The first (RA) part of this definition captures the functionally closed property mentioned above, and the second (F) part captures the self-sustaining property. A more formal definition of RAF sets is provided in Hordijk and Steel (2004); Hordijk et al (2011), including an efficient (polynomial-time) algorithm for finding such sets in any (arbitrary) CRS. This RAF algorithm returns the union of all RAF (sub)sets that exist within a given CRS, or the empty set if the CRS does not contain any RAF set.

As a simple model of a CRS, we use the binary polymer model (Kauffman, 1986, 1993). In this model, molecule types are represented by bit strings up to a certain length n , with the food set made up of bit strings up to a given small length t (e.g., $t = 2$). The possible reactions are ligation (concatenating two bit strings into one larger one) and cleavage (cutting a bit string into two smaller ones). Finally, the catalysis events are assigned at random, with a given probability p that a molecule $x \in X$ catalyzes a reaction $r \in R$. The idea behind modeling catalysis randomly is that, in general, little is known about it, and as with chemical reactions, predicting catalysis is a hard problem (Kayala et al, 2011). This model has been used in other, related, computational studies on autocatalytic sets as well (Filisetti et al, 2011; Vasas et al, 2012).

Using the binary polymer model, it was shown that RAF sets are highly likely to exist in general CRSs, even for very moderate and chemically plausible levels of catalysis (Hordijk and Steel, 2004; Mossel and Steel, 2005; Hordijk et al, 2010). Furthermore, this result still holds when (i) a more realistic “template-based” form of catalysis is used where potential catalysts have to match a certain number of bits around the reaction site (Hordijk et al, 2011; Hordijk and Steel, 2012b), (ii) only the longest polymers can act as catalysts, also in combination with the template constraint (Hordijk et al, 2014), and (iii) a “partitioned” polymer set is used, where polymers can undergo only reactions within their own partition, but catalysis can be both within and between partitions, such as in an RNA/protein world (Smith et al, 2014).

The RAF sets that are found by the RAF algorithm are called *maximal* RAF sets (maxRAFs). However, it was shown that a maxRAF can often be decomposed into multiple smaller subsets which themselves are RAF sets (subRAFs) (Hordijk et al, 2012). If such a subRAF cannot be reduced any further without losing the RAF property, it is referred to as an *irreducible* RAF (irrRAF). The existence of multiple autocatalytic subsets can actually give rise to an evolutionary process (Vasas et al, 2012), and the emergence of larger and larger autocatalytic sets over time (Hordijk et al, 2012).

Finally, RAF sets are not just a theoretical construct, but have been shown to exist in real chemical systems (Sievers and von Kiedrowski, 1994; Ashkenasy et al, 2004; Lincoln and Joyce, 2009; Vaidya et al, 2012). In fact, RAF theory can be applied directly and successfully to model such real chemical reaction systems (Hordijk and Steel, 2013), providing more insight into their structure and properties. Moreover, RAF sets were recently found to exist in an actual bacterial metabolic network (Sousa et al, 2014).

3 Irreducible RAF sets

In Hordijk et al (2012) it was shown that, in principle, there can be exponentially many irreducible RAF sets within a maximal RAF. So, in general it is not possible to efficiently enumerate all irrRAFs that exist within a given CRS. Furthermore, in Steel et al (2013) it was shown that even finding a *smallest* irrRAF is an NP-complete problem.

Irreducible RAF sets are important in at least two ways. First, they represent the smallest possible autocatalytic sets that can exist within a chemical reaction system, and are thus likely to emerge first in a dynamical sense. Second, irrRAFs are the equivalent of the “viable cores” of Vasas et al (2012), where it was shown that the existence of multiple such sets can give rise to an evolutionary process (provided some other conditions are met as well). In this sense, the possible existence of exponentially many irrRAFs is actually a desirable property.

So, despite their computational intractability, it would still be useful to have more insight into the (expected) sizes and structure of irrRAFs. In Steel et al (2013), a search algorithm was introduced to randomly sample irrRAFs within a given RAF set \mathcal{R}' . Briefly, this algorithm works as follows [see Steel et al (2013) for details]:

irrRAF sampling algorithm

1. Randomly reorder the reactions r_i in the given RAF set \mathcal{R}' .
2. For each next reaction $r_i \in \mathcal{R}'$, remove r_i and apply the RAF algorithm to $\mathcal{R}' - \{r_i\}$, resulting in a subRAF $\mathcal{R}'' \subset \mathcal{R}'$.
 - (a) If $\mathcal{R}'' = \emptyset$, return r_i to \mathcal{R}' .
 - (b) Otherwise, replace \mathcal{R}' with \mathcal{R}'' .
3. Go back to step 2, until all reactions have been considered for removal.
4. Return the resulting irreducible RAF set \mathcal{R}' .

Note that the particular irrRAF that is returned by this algorithm depends on the order in which the reactions in \mathcal{R}' are considered for removal. Therefore, the reactions in the given RAF \mathcal{R}' are randomly reordered each time the algorithm is applied, so a (possibly) different irrRAF can be found.

Using the above randomized algorithm, Steel et al (2013) presented some preliminary results on the average sizes of irrRAFs. In the current paper, we perform a more detailed and thorough investigation into several properties of irrRAFs. For this, separate and independent implementations of both the irrRAF sampling algorithm and the binary polymer model were created (in Python 3.3). Next, the algorithm was applied to many instances of the binary polymer model for different values of the parameters n and p (with $t = 2$), and various statistics on irrRAFs were collected.

3.1 Average irrRAF sizes

First, to verify the correctness of both the new implementation and the previous (preliminary) results, Figure 1 (left) reproduces the results of Steel et al (2013) on the average sizes (in number of reactions) of maxRAFs and irrRAFs for increasing levels of catalysis. The *level of catalysis* f is defined as the average number of reactions catalyzed per molecule, which is $f = p|\mathcal{R}|$ in the binary polymer model.

A maximum molecule length of $n = 10$ is used, and each data point is averaged over 39 model instances. Note that for $n = 10$, RAF sets start showing up (with low probability) for a level of catalysis around $f = 1.2$, exist in about 50% of model instances for $f = 1.3$, and are very common for values of $f = 1.4$ and larger.

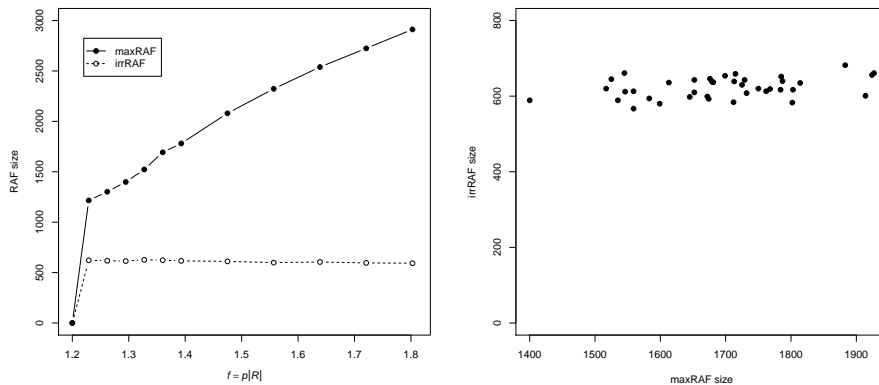


Fig. 1 Left: The average sizes of maxRAFs (black dots) and irrRAFs (white dots) in the binary polymer model with $n = 10$ and various values of p . Right: The average size of irrRAFs vs. the size of the maxRAF they were found in, for $n = 10$ and $p = 0.000083$.

As Figure 1 (left) shows, the average size of maxRAFs increases linearly with increasing levels of catalysis. In other words, when maxRAFs become more common, they also become larger. However, somewhat surprisingly, the size of irrRAFs remains constant with increasing values of f , and thus seems to be independent of the level of catalysis. In fact, the size of an irrRAF even seems to be independent of the size of the maxRAF it is a subset of. Figure 1 (right) shows the average size of irrRAFs against the size of the maxRAF they were sampled from (50 irrRAFs per maxRAF), for $n = 10$ and $p = 0.000083$ ($f = 1.36$). Even though the sizes of the maxRAFs range from about 1400 to 1950 reactions, the (average) size of their irrRAFs has a very small variance around about 600 reactions.

In the limit of very high levels of catalysis (most molecules catalyzing most reactions), one would expect to find small irrRAFs. So an interesting question is how large the level of catalysis needs to be before the average irrRAF size starts to decrease and become very small. Figure 2 shows this kind of data for three different values of n (averaged over 16 model instances for each data point). Again somewhat surprisingly, it requires rather high levels of catalysis to start seeing small irrRAFs. Note the log-scale on the horizontal axis, with levels of catalysis one to two orders of magnitude larger compared to Figure 1.

In conclusion, in the range of realistic values of the level of catalysis, the (average) size of irrRAFs seems to depend mostly on the value of n (or, similarly, on the size of the full reaction network), but is independent of the value of p (probability of catalysis) or the size of the maxRAF they are part of. So, whether they are rare or common, the sizes of irrRAFs can be expected to be about the same for similarly sized CRSs, and cannot be expected to be very small. In Steel et al

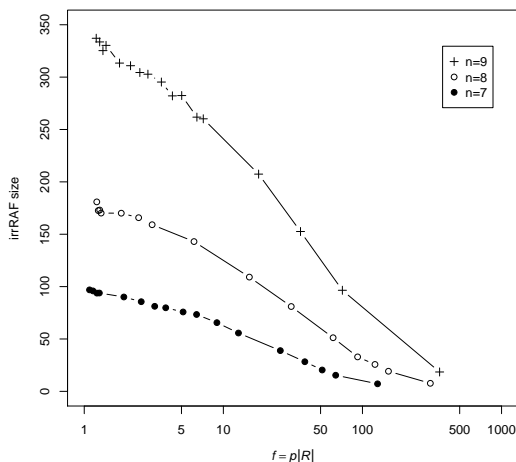


Fig. 2 The average irrRAF sizes for high levels of catalysis (f) for various values of n .

(2013) it was already proved that it is unlikely to find very small irrRAFs at a level of catalysis where RAF sets are just starting to show up (with low probability). However, the current results seem to indicate that this fact still holds over a much larger range of f values.

3.2 Smallest irrRAF sizes

Even though finding a *smallest* irrRAF set is, in general, an NP-complete problem (Steel et al, 2013), with the irrRAF sampling algorithm it is still possible to get an idea of the overall range and distribution of irrRAF sizes. Figure 3 shows two histograms of irrRAF sizes sampled from one particular maxRAF in an instance of the binary polymer model with $n = 10$. The histogram on the left is for a sample size of 250 irrRAFs, while the one on the right is for a sample size of 1000. The main difference between the two sample sizes is that for the larger sample the histogram approaches a bell-shaped distribution more closely. A striking observation, however, is that all irrRAFs in the sample turn out to be unique. In other words, none of the 1000 irrRAFs sampled consist of the exact same subset of reactions, even though many of them have the same size. In principle there can be exponentially many irrRAFs within a maxRAF (Hordijk et al, 2012), and these sampling results indicate that indeed, in general, many different irrRAFs exist within a given maxRAF (although they may overlap in at least some reactions). As already mentioned, from an evolutionary point of view this is a useful feature.

As Figure 3 clearly shows, the range of irrRAF sizes does not seem to depend on the sample size. In other words, using a larger sample size does not make it more likely to find even smaller or larger irrRAFs. Given that it is unlikely that very small irrRAFs exist, according to the above results, the smallest irrRAF found in this sample is likely to be close to the actual minimum size. However, the question

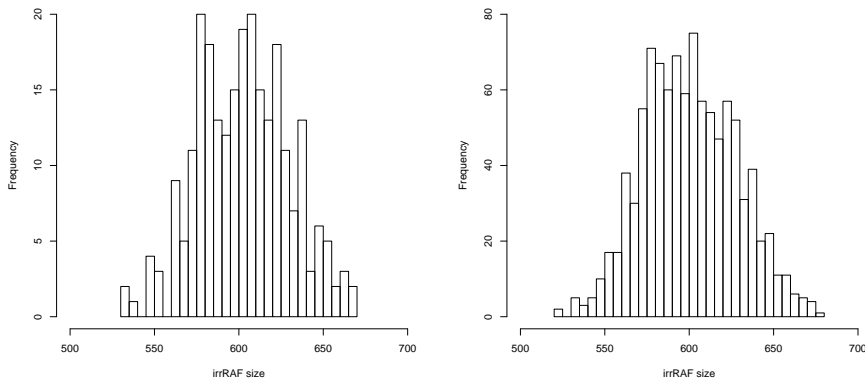


Fig. 3 The influence of the sample size on irrRAF sizes. Left: 250 samples. Right: 1000 samples. $n = 10$ and $p = 0.0000792$.

remains whether it is possible to get an analytical or efficiently computable bound on the size of the smallest irrRAFs.

In Steel et al (2013) it was suggested to look for directed cycles in the catalysis graph. The *catalysis graph* $G(\mathcal{R})$ of a CRS $Q = \{X, \mathcal{R}, C\}$ is a graph where each node corresponds to a reaction in \mathcal{R} , and with an arc from node r_i to node r_j if some product x_k of reaction r_i catalyzes reaction r_j , i.e., if $(x_k, r_j) \in C$. Assuming that no food molecules catalyze any reactions (a constraint that can be easily imposed on the binary polymer model without significantly changing any of the main results), an RAF set $\mathcal{R}' \subset \mathcal{R}$ generates a (directed) cycle in $G(\mathcal{R})$ (Contreras et al, 2011; Steel et al, 2013). So, if Q contains an RAF set then $G(\mathcal{R})$ contains at least one cycle, although, importantly, the reverse is not necessarily true. Also note that this cycle can be of any length, and does not necessarily include all reactions that are part of the RAF set (i.e, the length of the cycle can be much smaller than the size of the RAF set).

Using this idea, we constructed the catalysis graphs $G(\mathcal{R}')$ of several irreducible RAF sets \mathcal{R}' as found in instances of the binary polymer model with $n = 7$, and then used a function from the `NetworkX` Python package (NetworkX Developers, 2013) to detect the simple cycles in these graphs. However, it turns out that the lengths of such cycles is generally much smaller than the sizes of the irrRAFs themselves, as shown in Figure 4. The irrRAF sizes range from 70 to 125 reactions, whereas the simple cycle lengths never exceed 20. Obviously, these smaller cycles cannot correspond to RAF sets, as the catalysis graphs are explicitly constructed from irrRAFs, which (by definition) do not contain any smaller RAF subsets. The picture is the same when computing the strongly connected components in the catalysis graph. In this case, the largest component found is of size 24, still much less than the sizes of the irrRAF sets themselves. In other words, the lengths of cycles or sizes of strongly connected components in the catalysis graph do, unfortunately, not provide a useful bound on the size of the smallest irrRAFs.

However, these results do have important consequences for drawing conclusions about the existence of autocatalytic sets in a reaction networks \mathcal{R} by merely looking

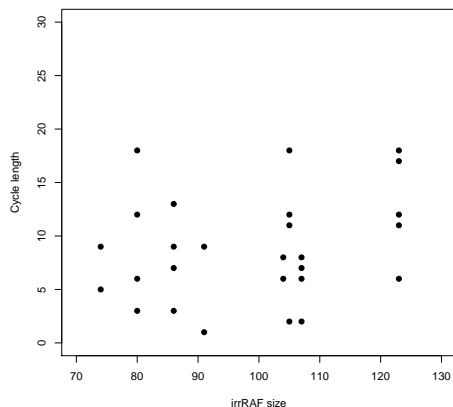


Fig. 4 Simple cycle lengths in the catalysis graph G vs. the size of the irrRAF that G was constructed from.

for cycles or strongly connected components in the corresponding catalysis graph $G(\mathcal{R})$, as for example done in Filisetti et al (2011). First, the presence of a cycle or connected component does not necessarily imply the existence of an RAF set. Second, as the current results (Figure 4) show, even if the presence of a cycle or component in the catalysis graph $G(\mathcal{R})$ does correspond to an actual RAF set, they still do not necessarily represent the full RAF set (at least not in the binary polymer model). These cycles or components do represent reflexively autocatalytic (RA) sets, but they are not necessarily food-generated (F), and are thus not self-sustaining.

3.3 Structural properties of (irr)RAFs

Finally, we looked at some structural properties of RAF sets. In particular, one can ask how important individual molecules and reactions are to maintain a given maxRAF or irrRAF. Figure 5 shows the (average) number of irrRAFs that molecules of different lengths are part of. Data points are averaged over 39 instances of the binary polymer model with $n = 10$ and $p = 0.000083$, with an irrRAF sample of size 50 for each of the 39 maxRAFs. Then for each molecule it was counted in how many of those $39 \times 50 = 1950$ irrRAFs it occurred, and those counts were averaged for each possible molecule length. As the figure shows, all molecules up to length four occur in all the irrRAFs sampled. Molecules of length five are part of most of the irrRAFs, but then the number drops off quickly to about one quarter of the irrRAFs for the longest molecules (length 10).

Another way of looking at the importance of molecules is as follows. For each molecule x_i that is part of a maxRAF \mathcal{R}' , remove that molecule from \mathcal{R}' and re-apply the RAF algorithm, resulting in a (possibly smaller) RAF set \mathcal{R}'' . Then consider the difference in size between the original maxRAF \mathcal{R}' and the resulting RAF set \mathcal{R}'' . Figure 6 (left) shows the results of such an analysis for one particular maxRAF (with 1511 reactions) found in an instance of the binary polymer model

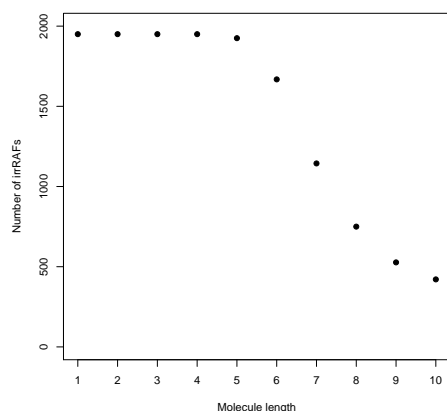


Fig. 5 The average number of irrRAF's (out of 1950) a molecule of a given length is part of, for $n = 10$ and $p = 0.000083$.

with $n = 10$ and $p = 0.0000792$. The food molecules were not included in this analysis, as it is obvious that they are essential.

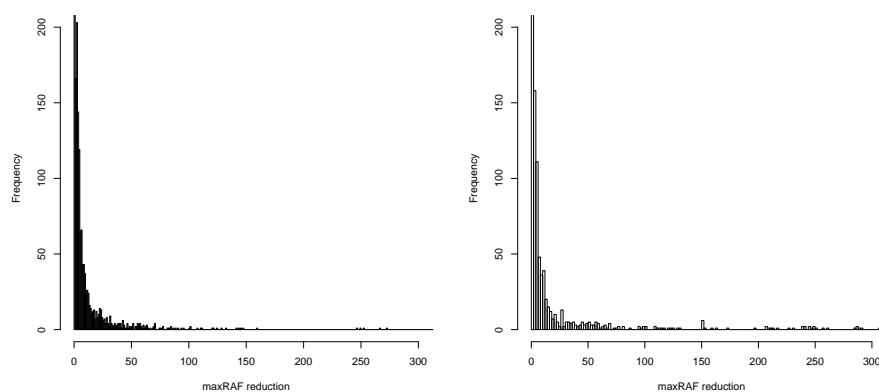


Fig. 6 Left: The influence of removing molecules from the RAF set. Right: The influence of removing reactions from the RAF set. Both diagrams are for $n = 10$ and $p = 0.0000972$.

As the histogram shows, most (non-food) molecules do not have any impact on the size of the maxRAF at all. In fact, about half of the molecules (769; off the scale of the histogram) have no impact. Most of the remaining molecules only have a very small impact, and just a few molecules have a large impact (a reduction in the maxRAF size of, say, more than 100 reactions). In this particular instance, there are actually four molecule types that, when removed, reduce the maxRAF to size zero, i.e., the maxRAF reduction is of size 1511 (also off the scale of the

histogram). These are what we call “essential” molecules: they are necessarily part of every possible irrRAF within the maxRAF.

A similar analysis can, of course, be performed by removing individual reactions instead of molecules. Figure 6 (right) shows a similar histogram, but for removing reactions from a maxRAF of size 2300. As expected, the overall pattern is similar: most reactions have no (1634 reactions; off the scale) or only a small impact on the maxRAF size, while only a few reactions have a large impact. In this particular instance, there are 55 essential reactions (off the scale), i.e., they will have to be part of every irrRAF.

Looking at other instances, the overall pattern is again similar, except that not all maxRAFs always have essential reactions. In fact, many maxRAFs do not have any essential reactions at all. They may have reactions that, when removed, have a large impact on the maxRAF size, but they do not necessarily reduce this size to zero. Figure 7 shows a histogram of the number of essential reactions in a sample of 400 maxRAFs (for $n = 10$ and $p = 0.000079$). About 130 of these maxRAFs do not have any essential reactions at all, and only a few instances have a large number of essential reactions.

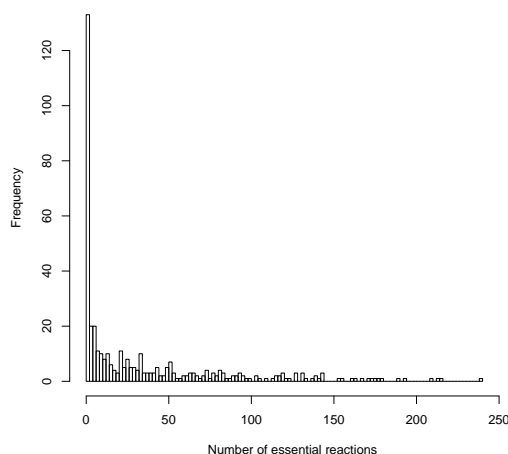


Fig. 7 The number of essential reactions in maxRAFs over a sample of 400 instances. $n = 10$ and $p = 0.000079$.

These results give some interesting insights into various structural properties of RAF sets. Since the binary polymer model is just an abstract model of a CRS, it is difficult to attach any biological interpretation to these results. However, in a separate study a similar “sensitivity” analysis is performed on RAF sets in the metabolic network of *E. coli*, in which case the results do have a clear biological interpretation (Sousa et al, 2014). Hence, the structural analysis as described here can indeed be quite useful.

4 Power law distributed catalysis

In Hordijk et al (2011); Hordijk and Steel (2012b), a more realistic version of the binary polymer model was investigated. In particular, a constraint was imposed on the catalysis events, so that a molecule can only be assigned as a catalyst to a reaction if that molecule matches a certain number of bits around the reaction site. This was inspired by, for example, base-pair complementarity in RNA molecules. In Hordijk et al (2014), the constraint that only the longest polymers can be catalysts was investigated, also in combination with the template-match constraint. Here, we investigate another variant of the binary polymer model in an effort to include more biological realism. In particular, a common property among “real-world” networks (including biological ones) seems to be a power law degree distribution (Dorogovtsev and Mendes, 2003; Newman, 2010). There are usually many nodes in such networks that have a very low connectivity, and a few nodes that have a very high connectivity.

To include this property in the binary polymer model, we assign the catalysis events according to a power law distribution. In the standard (“uniform”) binary polymer model, each molecule-reaction pair (x, r) is included in the catalysis set C with independent probability p . In the power law variant of the model, instead of considering individual molecule-reaction pairs, we first independently draw a random number s_i from a power law distribution for each molecule type $x_i \in X$. Then, for each molecule type x_i , a total of s_i reactions from \mathcal{R} are chosen randomly (with equal probability and without replacement), to which x_i is assigned as a catalyst. This way, most molecules will catalyze only few reactions, while some may catalyze many reactions.

The power law distribution we use is the Zipf distribution with parameter a :

$$f_a(k) = \frac{k^{-a}}{\zeta(a)}$$

where ζ is the Riemann Zeta function and k ranges over the positive integers. This probability density function is named after the linguist Zipf, who noted that the frequency of any word in a sample of a language is inversely proportional to its rank in the frequency table (Zipf, 1932). The NumPy Python library provides a function for drawing random numbers from this distribution, from which we then subtract one to also allow for molecules catalyzing no reactions. The parameter a in the power law variant of the binary polymer model plays the same role as the parameter p in the standard model. However, comparing them directly is meaningless. Instead, we calculate the level of catalysis f (the average number of reactions catalyzed per molecule), which depends on p (standard model) or a (power law model), and can then compare the values of f between the two model variants.

Figure 8 shows the resulting distribution of the number of reactions catalyzed per molecule in one particular instance of the binary polymer model with this power law catalysis assignment for $n = 10$ and $a = 2.35$ (black bars). For comparison, the equivalent distribution for the standard model is shown as well for $n = 10$ and $p = 0.00008$ (gray bars). These distributions are indeed as expected: a Poisson distribution in the standard model (Newman, 2010) and a power law distribution in the model variant. In the instance of the standard model, the average number of reactions catalyzed per molecule is $f = 1.31$, while in the instance of the power

law model it is $f = 1.21$, i.e., the (average) level of catalysis is similar in both these instances.

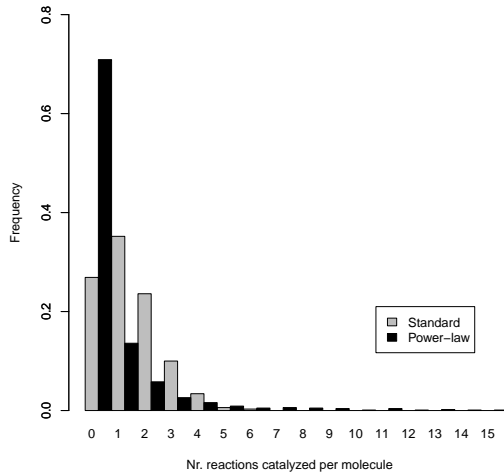


Fig. 8 Number of reactions catalyzed per molecule. Gray bars = standard model ($n = 10$ and $p = 0.000080$; Poisson distribution), black bars = power law distributed ($n = 10$ and $a = 2.35$).

Figure 9 shows the probability of finding RAF sets in instances of the power law model variant for various values of n (different curves) and (average) levels of catalysis f (on the x-axis), averaged over 1000 instances for each data point. In other words, each data point in the figure represent a particular combination of values for n and a (converted to the corresponding value for f), and for each such combination of parameter values, 1000 instances of the model were generated. The probability P_n of finding RAF sets (on the y-axis) is then the fraction of these 1000 instances that contain an RAF set (as determined by the RAF algorithm).

For comparison, similar data for $n = 10$ from the standard model is included as well in Figure 9 (the black dots). As in the standard model, in the power law case there is a fairly quick transition from not finding RAF sets ($P_n = 0$) to finding them most of the time (P_n close to one). However, the transition is less sharp, but RAF sets already start to show up at a lower level of catalysis (value of f) compared to the standard model.

A surprising feature in Figure 9 is that for the power law model there does not seem to be any increase in the level of catalysis f necessary to find autocatalytic sets for increasing values of n . In the standard model, a linear growth rate in the required level of catalysis f is observed for increasing n , i.e., the S-shaped curves (for each next larger value of n) move to the right at a linear rate (Hordijk and Steel, 2004). However, for the power law model, as Figure 9 shows, all the S-shaped curves intersect at a probability P_n between 0.5 and 0.6. Also, the transition from low to high P_n for a given n (with increasing level of catalysis f) becomes sharper for larger values of n . So, it seems that in this more realistic model variant, RAF sets are actually even easier to get (they already start showing up for smaller

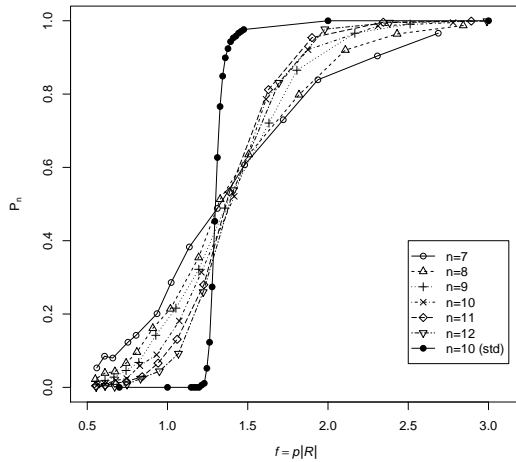


Fig. 9 The probability P_n of finding RAF sets for the power-law distributed catalysis model (various n and a). The black dots are for $n = 10$ (various p) in the standard model (std), for comparison.

levels of catalysis) than in the standard model, and do not require a larger level of catalysis for larger system sizes.

The corresponding values for the parameter a in the Zipf distribution where this transition happens is roughly in the range $a \in [2.2; 2.7]$. RAF sets start to show up (with low probability) for values of a smaller than 2.7, exist in about half the instances for a value of a around 2.35, and are always present for a smaller than 2.2. This transition range corresponds very well with “real-world” networks, which tend to have a value for their power law parameter between two and three (Newman, 2010). Furthermore, as just observed, this range does not seem to depend on the actual value of n (at least not in the range of n values used here).

Note that the values on the horizontal axis (level of catalysis) in Figure 9 are averaged over many model instances. In the standard model, there is little variance in the average level of catalysis between different model instances for the same values of the parameter p (given a fixed value for n). However, in the power law case, this variance is much larger for a given value of a (the equivalent of p in the standard model). Figure 10 (left) shows this quite clearly, where a histogram of the average number of reactions catalyzed per molecule (f) is shown for $n = 10$ and $a = 2.35$ for 1000 instances of the power law model. The observed mean of this sample (1.395) is indeed very close to the expected mean (1.391), but the variance is quite large, with a few instances even having a level of catalysis of $f > 5$.

However, we prefer to present the results for the power law model in terms of the (average) level of catalysis (f), rather than the parameter a , since in this way they can be more easily interpreted and compared to the standard model. Also, this averaging, even though there is a large variance, does not seem to change the overall trends very much. Alternatively, the results can be presented unaveraged, as done in Figure 10 (right) for the sizes of maxRAFs and irrRAFs against the

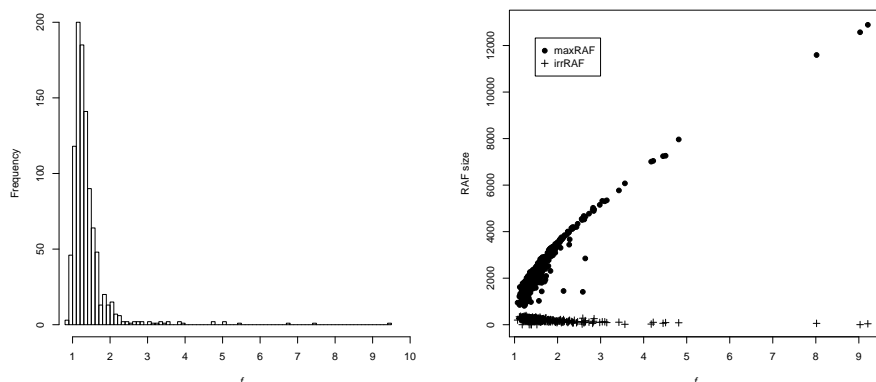


Fig. 10 Left: The distribution of observed average number of reactions catalyzed per molecule (f) for the power-law model. Right: The sizes of maxRAFs (dots) and irrRAFs (pluses) for the power-law case. Both diagrams are for $n = 10$ and $a = 2.35$.

value of f in close to 500 instances of the power law model with $n = 10$ and $a = 2.35$. This graph shows the same trend as can be observed in Figure 1 for the standard model: the size of maxRAFs increases linearly with increasing level of catalysis, but the size of irrRAFs is more or less constant. However, note that (for the same value of n), the maxRAFs are larger (ranging from about 2000 to 8000 reactions vs. 1200 to 3000), but the irrRAFs are smaller (around 200 reactions vs. around 600) in the power law case compared to the standard model.

In conclusion, by incorporating this more realistic way of assigning catalysis events, it appears that RAF sets already start showing up for smaller levels of catalysis, and the irrRAFs are smaller (and thus require fewer molecule types and reactions) compared to the standard model. Moreover, in the power law model there does not seem to be any increase in the level of catalysis necessary with increasing n to maintain the same probability of finding RAF sets. Other results, such as the (average) irrRAF sizes being independent of the level of catalysis or the maxRAF size (for a given n), the observed size range of irrRAFs being independent of the irrRAF sample size, and the impact of removing molecules or reactions from the RAF set, are very similar to those for the standard model.

5 Conclusions

We have investigated several aspects of RAF theory in more detail. First, a more thorough analysis of irreducible RAF sets was performed. Our results confirm that there are indeed many different such subsets within a given maxRAF. This has positive consequences for the evolvability of autocatalytic sets, which could have been crucial in the early stages of the origin and evolution of life. Rather than having to wait for one or a few specific chemical organizations to emerge, different (irreducible) RAF sets seem likely to have existed in large numbers. Their sizes do not seem to depend on the actual level of catalysis present, or the size of the maxRAF they are part of. In fact, even for low levels of catalysis, when RAF sets

are just starting to show up, irrRAFs are already quite large, thus immediately giving rise to a significant amount of chemical complexity.

Overall, the RAF sets we investigated seem to have common structural properties. They usually contain many molecules and reactions that have a low impact on maintaining the set as a whole, and a few that have a very high impact on the integrity of the system. Also, a significant fraction of RAF sets contain so-called “essential” molecules and reactions, the removal of which would break down the RAF set completely (and which, therefore, will be part of every possible irrRAF set within the larger RAF set). As expected, food molecules and other small molecules that are produced directly from food molecules tend to be essential.

As a word of warning, simply looking for cycles in the catalysis graph of a reaction system does most likely not result in actually finding autocatalytic sets. These cycles (especially the simple cycles) do represent RA sets, but they are not necessarily food-generated (F). However, using the RAF algorithm and the randomized irrRAF sampling algorithm, RAF (sub)sets can be found efficiently (although in the case of irrRAFs not necessarily exhaustively).

Finally, we have investigated a more realistic variant of the binary polymer model where the catalysis events are assigned according to a power law distribution. In this model variant, it seems that RAF sets are even more likely to show up (i.e., at lower levels of catalysis), and irreducible RAFs tend to be smaller (requiring fewer molecule types and reactions) compared to the standard model. Furthermore, the results suggest that no growth rate in the level of catalysis is necessary with increasing system sizes (the parameter n). It would be an interesting extension to combine this power law model variant with other more-realistic assumptions such as the previously investigated template-based catalysis constraint.

Acknowledgements

This paper resulted from a student project at the annual Oxford Summer School on Computational Biology (OSSCB13). We thank the Department of Statistics of the University of Oxford for funding. The complete project report, including additional and more detailed results, can be found online at:

www.stats.ox.ac.uk/__data/assets/pdf_file/0004/9859/Report_RAF_Theory.pdf

References

- Ashkenasy G, Jegasia R, Yadav M, Ghadiri MR (2004) Design of a directed molecular network. *PNAS* 101(30):10,872–10,877
- Contreras DA, Pereira U, Hernández V, Reynaert B, Letelier JC (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
- Dorogovtsev SN, Mendes JFF (2003) *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford University Press
- Dyson FJ (1985) *Origins of Life*. Cambridge University Press
- Eigen M, Schuster P (1979) *The Hypercycle*. Springer-Verlag

- Filisetti A, Graudenzi A, Serra R, Villani M, De Lucrezia D, Füchslin RM, Kauffman SA, Packard N, Poli I (2011) A stochastic model of the emergence of autocatalytic cycles. *Journal of Systems Chemistry* 2:2
- Gánti T (2003) *The Principles of Life*. Oxford University Press
- Hordijk W (2013) Autocatalytic sets: From the origin of life to the economy. *BioScience* 63(11):877–881
- Hordijk W, Steel M (2004) Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology* 227(4):451–461
- Hordijk W, Steel M (2012a) Autocatalytic sets extended: Dynamics, inhibition, and a generalization. *Journal of Systems Chemistry* 3:5
- Hordijk W, Steel M (2012b) Predicting template-based catalysis rates in a simple catalytic reaction model. *Journal of Theoretical Biology* 295:132–138
- Hordijk W, Steel M (2013) A formal model of autocatalytic sets emerging in an RNA replicator system. *Journal of Systems Chemistry* 4:3
- 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. *International Journal of Molecular Sciences* 12(5):3085–3101
- Hordijk W, Steel M, Kauffman S (2012) The structure of autocatalytic sets: Evolvability, enablement, and emergence. *Acta Biotheoretica* 60(4):379–392
- Hordijk W, Steel M, Kauffman S (2013) Autocatalytic sets: The origin of life, evolution, and functional organization. In: Pontarotti P (ed) *Evolutionary Biology: Exobiology and Evolutionary Mechanisms*, Springer, pp 49–60
- Hordijk W, Wills PR, Steel M (2014) Autocatalytic sets and biological specificity. *Bulletin of Mathematical Biology* 76(1):201–224
- Kauffman SA (1971) Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. *Journal of Cybernetics* 1(1):71–96
- Kauffman SA (1986) Autocatalytic sets of proteins. *Journal of Theoretical Biology* 119:1–24
- Kauffman SA (1993) *The Origins of Order*. Oxford University Press
- Kayala MA, Azencott CA, Chen JH, Baldi P (2011) Learning to predict chemical reactions. *Journal of Chemical Information and Modeling* 51:2209–2222
- Lincoln TA, Joyce GE (2009) Self-sustained replication of an RNA enzyme. *Science* 323:1229–1232
- Maturana H, Varela F (1980) *Autopoiesis and Cognition: The Realization of the Living*. Reidel
- Mossel E, Steel M (2005) Random biochemical networks: The probability of self-sustaining autocatalysis. *Journal of Theoretical Biology* 233(3):327–336
- NetworkX Developers (2013) NetworkX. [networkx.github.io](https://github.com/networkx/networkx)
- Newman MEJ (2010) *Networks: An Introduction*. Oxford University Press
- Sievers D, von Kiedrowski G (1994) Self-replication of complementary nucleotide-based oligomers. *Nature* 369:221–224
- Smith J, Steel M, Hordijk W (2014) Autocatalytic sets in a partitioned biochemical network. *Journal of Systems Chemistry* 5:2
- Sousa FL, Hordijk W, Steel M, Martin W (2014) Autocatalytic sets in the metabolic network of *E. coli*. *Journal of Systems Chemistry* (submitted)
- Steel M (2000) The emergence of a self-catalysing structure in abstract origin-of-life models. *Applied Mathematics Letters* 3:91–95

-
- Steel M, Hordijk W, Smith J (2013) Minimal autocatalytic networks. *Journal of Theoretical Biology* 332:96–107
- Taran O, Thoennesen O, Achilles K, von Kiedrowski G (2010) Synthesis of information-carrying polymers of mixed sequences from double stranded short deoxynucleotides. *Journal of Systems Chemistry* 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
- Vasas V, Fernando C, Santos M, Kauffman S, Sathmary E (2012) Evolution before genes. *Biology Direct* 7:1
- Zipf GK (1932) *Selected Studies of the Principle of Relative Frequency in Language*. Harvard University Press