EVOLUTION WITHOUT VARIATION AND SELECTION

LIANE GABORA AND MIKE STEEL

ABSTRACT. A central tenet of evolutionary theory is that it requires variation upon which selection can act. We describe a means of attaining cumulative, adaptive, open-ended change that requires neither variation nor selective exclusion, and that can occur in the absence of generations (i.e., no explicit birth or death). This second evolutionary process occurs through the assimilation, restructuring, and extrusion of products into the environment by identical, interacting Reflexively Autocatalytic and Food set-generated (RAF) networks. We refer to this more primitive process evolutionary process as Self–Other Reorganisation because it involves internal self-organising and self-maintaining processes within entities, as well as interaction between entities. Since there is no self-assembly code, it is more haphazard than natural selection, and there is no discarding of acquired traits (a signature characteristic of natural selection). In the extreme, it can work with just one entity but it differs from learning because it can operate in groups of entities, and produce adaptive change across generations. We suggest that this more primitive process is operative during the initial stage of an evolutionary process, and that it is responsible for both the origin and early evolution of both organic life, and human culture. In cultural evolution, this 'evolution without variation' process can increase homogeneity amongst members of a group and thereby foster group identity and cohesion.

Keywords: autocatalytic network, cultural evolution, evolution, origin of life, social learning

Addresses for Correspondence:

Liane Gabora
Department of Psychology, University of British Columbia, Kelowna BC, Canada
liane.gabora@ubc.ca

Mike Steel
Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand
mike.steel@canterbury.ac.nz

1. Introduction

A central tenet of evolutionary theory is that evolution requires variation upon which selection can act. This follows from Darwin's theory of natural selection, which holds that evolution is due to differential replication of randomly generated heritable variations in a population over generations, such that some traits become more prevalent than others. Darwin's theory of evolution by natural selection is the foundation of modern biology and provides a unifying principle for describing species, well after the origin of life. Darwin's theory arose in response to the paradox of how organisms accumulate adaptive change despite that traits acquired over a lifetime are eliminated at the end of each generation. His solution was to come up with a population-level explanation: although acquired traits are discarded, inherited traits are retained, so evolution can be explained in terms of preferential selection for those inherited traits that confer fitness benefits on their bearers. However, neither origin of life research nor cultural evolution research is plagued by the problem that acquired changes are extinguished at the end of a generation; therefore, in these domains, the paradox that Darwin's theory was designed to solve does not exist [17]. Accordingly, it would seem appropriate to pay particular attention in these domains to the role of epigenetic and non-Darwinian (e.g., Lamarckian) processes. However, it is widely assumed that any kind of evolutionary process requires variation and selection. Even in research on cultural evolution—which is sometimes (though not universally[1]) regarded as Lamarckian—it is generally assumed that variation and selection are what enables evolution to take place (e.g., [3, 6, 54]).

Indeed, although the term 'evolution' is sometimes assumed to be synonymous with variation and selection, it is increasingly more widely accepted that non-Darwinian processes play a role in evolution [16, 43, 44, 71, 77]. This paper goes a step further by providing an existence proof that evolution is possible in the *absence* of variation and selection. Although this has

---

[1]Some maintain that Lamarckism requires genetic transmission to biological offspring, a view held by early biologists (e.g., [63]), though current scholars (e.g., [54]) sometimes take a more equivocal view.

been suggested before [18], here, we formulate the argument in a more rigorous fashion, by mathematically describing what kind of structure an entity must have to engage in this kind of evolutionary process. This is important, because although entities of many kinds change due to acquired traits—e.g., a rock tumbling down a stream will acquire rounded edges—this change may not be adaptive, or to contribute to the entity's efficiency, survival, or replication. Indeed, in the case of the rock, the process that rounds its edges culminates in its disintegration. In contrast to the rock, in the example we will provide, the change is adaptive.

In origin of life research, there is a sizeable literature on particular models for the evolvability of primitive metabolic networks prior to DNA (see for example [48, 62, 64, 69, 78], and the references therein). Here, in order to describe evolutionary processes as seemingly disparate as the origin of life and the origin of culture within a common framework, we necessarily take a higher level and more general approach, using a different mathematical modelling technique from these earlier studies.

The paper begins with a list of definitions of terms and a table of the acronyms used. Next, we introduce the mathematical framework that will be used to develop the argument. We then model evolution in the absence of variation and selection with this framework (with further mathematical details provided in the Appendix). Next, we discuss what sort of structure is able to evolve without variation and selection. We conclude with a discussion of the implications.

2. Preliminaries

We begin by alphabetically listing and defining key terms used in this paper. We believe this is necessary in order to maintain clarity, since some miscommunication has arisen due to inconsistencies in how evolutionary concepts are applied in the sciences and social sciences.

- **Acquired trait:** a trait obtained during the lifetime of its bearer (e.g., a scar, a tattoo, or a memory of a song) and transmitted horizontally (i.e., laterally).

- **Culture:** extrasomatic adaptations—including behavior and artifacts—that are socially rather than sexually transmitted.

- **Darwinian (or 'selectionist') process:** an evolutionary process—i.e., a process that exhibits cumulative, adaptive, open-ended change—occurring through natural or artificial selection.

- **Darwinian threshold:** transition from non-Darwinian to Darwinian evolutionary process [77, 71].

- **Generation:** a single transition period from the internalized to the externalized form of a trait.[2]

- **Horizontal transmission:** non-germ-line transmission of an acquired trait from one entity to another. Thus, social transmission is horizontal because the information is not inherited, i.e., it is not passed from parent to offspring by way of the genes.

- **Selection:** differential replication of randomly generated heritable variation in a population over generations such that some traits become more prevalent than others. Selection may be natural (due to non-human elements of the environment) or artificial (due to human efforts such as selective breeding), and it can occur at multiple levels, e.g., genes, individuals, or groups [50].

- **Self-assembly code:** a set of coded instructions that is: (i) *actively interpreted* through a developmental process to generate a soma, and (ii) *passively copied without interpretation* through a reproduction process to generate self-copies, i.e., new sets of self-assembly instructions that are in turn used in these two distinct ways.

---

[2]Note that, with respect to biological evolution, a new generation generally (though not in horizontal gene transfer) begins with the birth of one or more organism(s). With respect to cultural evolution, a new 'generation' begins with the expression of an idea (which is considered one transmission event). Thus, over the course of a single discussion, an idea (a cultural trait) may undergo multiple generations.

- **Self–Other Reorganisation (SOR):** a theory of how both culture, and early life, evolve through interacting, self-organizing networks, based on theory and findings that have come out of the post-modern synthesis in biology.

- **Vertical transmission:** Germ-line inheritance of a trait from one generation to the next. (In other words, transmission occurs by way of a self-assembly code, such as DNA.)

It is important to emphasize that we are using the term 'selection' in its technical, scientific sense. The word 'selection' also has an 'everyday' sense in which it is synonymous with 'choosing' or 'picking out'. One could say that selection—in the everyday sense of the term—occurs in a competitive marketplace through the winnowing out of inferior products. However, the discussion here concerns selection in the scientific sense of the term.

Our definition of 'self-assembly code' is consistent with Von Neumann's work on relf-replicating automata [72], and subsequent insights by Holland [32], and Langton [49]. In a truly Darwinian or selectionist process, traits acquired through interaction with the soma (the culmination of the *active interpretation* of the code) do not affect traits vertically transmitted through reproduction (the culmination of the *passive uninterpreted* use of the code). As Holland [32] put it, the self-assembly code passed down through the germ-line is *sequestered*, i.e., protected from the influence of acquired traits. Thus, it is the division of labour between these two ways of using the self-assembly code that is responsible for the signature characteristic of a selectionist process: lack of transmission of acquired traits.

Note that vertical and horizontal transmission must be defined with respect to the relevant evolutionary process. A common error is to refer to the transmission of cultural information from parent to offspring as vertical transmission (e.g., [6]). Although the individuals in question are in a parent–child relationship with respect to their status as biologically evolving organisms, with respect to their status as participants in cultural evolution, this need not be the case.

Indeed, although childbirth entails one mother and one father, there is no limit to the number of 'parental influences' on the 'birth' of an idea. A related error is to say that in cultural evolution, there is a third form of transmission, namely, oblique transmission, in which traits are transmitted from non-kin members of the parental generation (e.g., [6]), because as far as cultural evolution is concerned it is irrelevant whether the information comes from biological kin or non-kin. In a similar vein, although dual inheritance theorists speak of culture as a second form of inheritance [30, 60, 75, 56], the distinguishing feature of an inherited trait is that it is transmitted vertically (e.g., from parent to offspring) by way of a self-assembly code (e.g., DNA), and therefore not obliterated at the end of a generation. This is not the case with respect to cultural traits [17], nor is it even the case for all biological traits.

The acronyms used in this paper are listed alphabetically in Table 1.

Acronym	Meaning
CRS	Catalytic Reaction System
RAF	Reflexively Autocatalytic and Food set-generated (F-generated)
OOL	Origin of Life
OOC	Origin of Culture
MR	Mental Representation
CCP	Cognitive Catalytic Process
SOR	Self–Other Reorganisation

TABLE 1. Abbreviations used throughout this paper.

3. AUTOCATALYTIC NETWORKS

The type of evolution without variation and selection process that we will describe involves entities that (1) have an organisation that is self-maintaining (i.e., the entities have a structure,

and a mechanism for maintaining that structure), and (2) they interact with each other by way of their environment. Autocatalytic networks provide a way to model such entities in origin of life (OOL) and origin of culture (OOC) settings.

The theory of autocatalytic networks grew out of studies of the statistical properties of *random graphs* consisting of nodes randomly connected by edges [10]. As the ratio of edges to nodes increases, the size of the largest cluster increases, and the probability of a phase transition resulting in a single giant connected cluster also increases. The recognition that connected graphs exhibit phase transitions led to their application to efforts to develop a formal model of the OOL, namely, of how abiogenic catalytic molecules crossed the threshold to the kind of collectively self-sustaining, self-replicating, evolving structure we call 'alive' [43, 42]. In the application of graph theory to the OOL, the nodes represent catalytic molecules, and the edges represent reactions. It is exceedingly improbable that any catalytic molecule present in the primordial soup of Earth's early atmosphere catalysed its own formation. However, reactions generate new molecules that catalyse new reactions, and as the variety of molecules increases, the variety of reactions increases faster. As the ratio of reactions to molecules increases, the probability increases that the system will undergo a phase transition. When, for each molecule, there is a catalytic pathway to its formation, they are said to be collectively *autocatalytic*, and the process by which this state is achieved has been referred to as *autocatalytic closure* [43]. The molecules thereby become a self-sustaining, self-replicating structure (such as a living protocell [38]). Thus, the theory of autocatalytic networks has provided a promising avenue for modelling the OOL and thereby understanding how biological evolution began [78].

Autocatalytic networks have been developed mathematically in the theory of Reflexively Autocatalytic and Food set-generated (RAF) networks [39, 67]. The term *reflexively* is used in its mathematical sense, meaning that every element is related to the whole. The term *food set* refers to the reactants that are initially present, as opposed to those that are the products of catalytic reactions. Elements of a 'foodset' are not the same thing as 'food';

the foodset is simply the raw building blocks available for catalyzed reactions. Autocatalytic networks such as RAFs have proven useful for modelling both the OOL and the onset of biological evolution [39, 67, 69, 78], and the origin of complex cognition and the onset of cultural evolution [12, 13, 22, 24, 23].[3] In the OOL context, the reaction network could be bounded by a semipermeable lipid membrane. In the OOC context, the network is a semantic network, which supervenes upon a brain within a human body, which thus constitutes a boundary on the cognitive RAF. In this case, the products and reactants are not catalytic molecules, but culturally transmittable *mental representations* [4] (MRs) of experiences, ideas, and chunks of knowledge, as well as more complex mental structures such as schemas and scripts. There can be multiple RAFs in the semantic network of a single individual. These RAFs may be subsumed by a larger RAF that connects them all; alternatively, they may be disconnected (as when inconsistent views are held by the same individual). An example that illustrates how a RAF can contain multiple RAFs within it is provided in Fig. 1.

Table 2 summarises how RAF theory terms apply in biological and cultural/cognitive settings. The fact that RAFs have proven useful in both these domains suggests that RAF theory may provide a broad conceptual framework that is applicable to the origins and early stages of diverse evolutionary processes.

We now summarise the key concepts of RAF theory. A *catalytic reaction system* (CRS) is a tuple $\mathcal{Q} = (X, \mathcal{R}, C, F)$ consisting of a set X of molecule types, a set $\mathcal{R}$ of reactions, a catalysis set C indicating which molecule types catalyse which reactions, and a subset F of X called the food set. A *Reflexively Autocatalytic and F-generated* set (i.e., a RAF) is a non-empty subset $\mathcal{R}' \subseteq \mathcal{R}$ of reactions that satisfies the following two properties:

---

[3]For related approaches, see [1, 5, 57].

[4]Although we use the term 'mental representation,' our model is consistent with the view (common amongst ecological psychologists and in cognition and quantum cognition communities) that what we call mental representations do not 'represent' but instead act as contextually elicited bridges between mind and world.

[5]Conceptual closure is the focus of another paper (in progress).

FIGURE 1. A RAF can contain a vast array of RAFs within it. The analysis of this 7-reaction system in [36] (which arose in an earlier RNA system constructed in a laboratory-based experimental model of early life) revealed that more than half of all the $2^7 = 128$ subsets of this RAF are themselves RAFs ('subRAFs'). In this figure, the reactions (r_i) are labelled by square nodes, and round nodes indicate molecule types (which consist of a single food-set element f, along with derived molecule types x_i). Catalysis is indicated by dashed arrows. Figure produced using the RAF program *CatlyNet*.

RAF Theory	Origin of Life (OOL)	Origin of Culture (OOC)
node	catalytic molecule	mental representation (MR)
edge	reaction pathway	association
cluster	molecules connected via reactions	MRs connected via associations
connected graph	autocatalytic closure [42, 43]	conceptual closure[5] [12]

TABLE 2. Application of graph theoretic concepts to the origin of life and origin of culture.

(1) *Reflexively autocatalytic (RA)*: each reaction $r \in \mathcal{R}'$ is catalysed by at least one molecule

type that is either produced by $\mathcal{R}'$ or is present in the food set F; and

(2) *F-generated*: all reactants in $\mathcal{R}'$ can be generated from the food set F by using a series
    of reactions only from $\mathcal{R}'$ itself.

A set of reactions that forms a RAF is simultaneously self-sustaining (by the F-generated condition) and (collectively) autocatalytic (by the RA condition) because each of its reactions is catalysed by a molecule associated with the RAF. A CRS need not have a RAF but when it does, there is a unique maximal one (the maxRAF). Moreover, a CRS may contain many possible RAFs, and it is this feature that allows RAFs to evolve, as demonstrated both in theory and in simulation studies, through selective proliferation and drift acting on the RAFs that are subsets of the maxRAF[39, 69].

In the OOL context, a RAF emerges in systems of polymers (molecules consisting of repeated units called monomers) when the complexity of these polymers (as measured by their maximum length) reaches a certain threshold [43, 55]. The phase transition from no RAF to a RAF incorporating most or all of the molecules depends on (1) the probability of any one polymer catalyzing the reaction by which a given other polymer was formed, and (2) the maximum length (number of monomers) of polymers in the system. This transition has been formalised and analysed mathematically, and by using simulations. Moreover, RAF theory has been applied to real biochemical systems [34, 35, 37, 39, 55], ecology [7] and cognition [22, 23, 24]; the theory has also proven useful for identifying how phase transitions might occur and at what parameter values. The evolution and reproduction of RAFs has been studied in an OOL context [55] and a OOC [24]; see also [1].

4. A RAF MODEL OF EVOLUTION WITHOUT VARIATION AND SELECTION

We now demonstrate a more primitive non-Darwinian form of evolution mathematically using RAF networks. We begin by describing a simple process involving a group $\mathcal{G}$ of indistinguishable entities as illustrated in Fig. 2. These entities can be described as identical RAFs. They may be CRSs such as those associated with very early life [2, 8, 16, 27, 40, 66, 71], artificial

neural networks in a computational model of cultural evolution (e.g., [11, 25]), human semantic networks or some other structure we have never encountered.

FIGURE 2. Identical entities exhibit cumulative adaptive change over time, as indicated by transition from light to dark, as a result of cumulative environmental changes, indicated by p, p' .... The first product p, generated by i due to a stimulus s catalysing a reaction in the RAF structure of entity i. The particular structure illustrated on the left is a schematic example of a simple RAF (with molecule types $\alpha - \epsilon$ and s, reactions indicated by squares, and catalysis shown with dashed arrows; food set here is $F = \{\alpha, \delta, s\}$, and ϵ gives rise to p). The product p in turn catalyses a reaction sequence in j at time t_2, leading to p', and so on. In the simple scenario outlined here, stimulation of i by s occurs only once, but that is sufficient to result in the generation of p, which is sufficient to set in motion the sequence of adaptive changes.

One entity in $\mathcal{G}$, which we call entity i, encounters a stimulus from the environment, which we refer to as s. In the case of a primitive biochemical reaction network, s could be molecule that crosses the lipid membrane that encases the reaction network. In the case of culture, s could be a visual or auditory stimulus. What makes s a novel stimulus is that at t_0 it was not part of the reaction network of any entity in $\mathcal{G}$. Stimulus s is nonetheless part of the foodset,

because it is available as an initial reactant or catalyst, as opposed to being directly or indirectly constructed from the set of initial reactants.[6]

Stimulus s triggers one or more catalytic reaction(s) in i. In a OOL scenario the reactions would be chemical reactions, while in a OOC scenario, they would be cognitive catalytic processes (CCPs) [22, 24]. This reaction (or chain of reactions) culminates in the generation of a product p into the environment. Thus, p was not part of the foodset; it is the end result of the series of reactions (or interactions) triggered by s. Note that s can only trigger the formation of p if the necessary sequence of reactions is able to take place, and it is the RAF that enables this reaction sequence to take place.

Product p confers some adaptive benefit on i. Exposure to p may improve the efficiency of i's reaction network (or semantic network) by reducing the number of steps required for a particular reaction sequence (or problem-solving process). It may have a protective function (for example, by making the environment less inviting for a predator) or (in a cultural context) it may serve as a source of pleasure that enhances wellbeing (for example, a work of art). We can describe the result of this triggering event as a sequence of one or more reactions within the RAF network of entity i that generates p, and for which the initial reaction in the sequence is catalysed by s. For example, in the OOC context, s catalyses the formation of a transient cognitive catalytic process that is part of the RAF within the mind of individual i, as illustrated in Fig. 2(ii) of [25]. Having generated p, i returns to its initial state except that it now experiences the adaptive benefit of p.

Once p is released into into the environment it can be encountered by other entities in $\mathcal{G}$. In an OOL context, p may be transmitted because it passes through a lipid membrane or protocell wall. In an OOC context, p may be transmitted via social learning, or by way of an artifact. A

---

[6]It is a matter of time before one of the entities encountered s; the more abundant s is, the less time it is expected to take for this chance encounter to occur. It is possible that the encounter with s results from the discovery of a previously uncharted part of the environment that is rich in s, but it could also be that s entered the environment from elsewhere, or that s was always there but by chance had not previously interacted with any of the entities.

recipient of p could potentially have had the capacity to generate p on its own, but by chance had never done so. In a OOL context, for example, a particular catalyst may have never come in contact with its reactant. In a cultural context, an individual may had two concepts but never thought of combining them. Alternatively, it is possible that the recipient did not previously possess the capacity to generate the product.

Since the other entities are identical to i, and exist in the same environment as i, the adaptive benefit of p percolates throughout $\mathcal{G}$.[7] The generation of p makes it possible to generate a new product, p', which confers even greater adaptive benefit than p. An entity j assimilates p, which triggers a sequence of one or more reactions in j, resulting in p'. Entity j benefits from the enhanced adaptive quality of p' over p. The ability to produce p' and enjoy its adaptive benefit percolates throughout the group $\mathcal{G}$. In this way, each version of p pave the way for its successor.

The entities in $\mathcal{G}$ continue to generate increasingly beneficial versions of the original p (e.g., $p', p'', p''', \ldots$). They collectively come to possess the ability to generate and benefit from this family of products because of the catalysis of the internal RAF structure of entities by the products of other entities. Thus we have cumulative adaptive change over time. Notice that in this model neither the birth nor the death of entities occurs; it is the same set of entities at the beginning as at the end, and there is no competitive exclusion and no selection (as defined above) of entities.

---

[7]The word 'percolates' is used here to refer to the 'communal exchange' (between-entities) component of the process, as opposed to the 'self-organization' (within-entities) component of the process. If one entity encounters p before another, that first entity may temporarily function more efficiently. We note also that the scenario could unfold such that each entity is affected solely by its own products, and thus after the impact of the initial generation of p by i, there could be little interaction, direct or indirect, amongst them. In other words, it is possible that j is affected by the p' it generated itself (as when someone is affected by looking at their own art). However, it is also possible that j is affected by the p' generated by another entity in $\mathcal{G}$ (as in when someone is affected by looking at another entity's art). Finally, we note that possession of a reactant (which in the cognitive scenario is a MR of a particular product) does not guarantee catalysis of that reactant. For example, someone may have a memory of an artwork, but while the initial experience of it yielded little cognitive restructuring, the second experience of it triggers a productive stream of thought.

4.1. **Modelling this process.** We consider a simple model, based on two processes (percolation and generation). First, let ρ denotes the rate of percolation of products (or knowledge or 'traces' of these products) through the group $\mathcal{G}$, with neighbourhood structure represented by a directed graph $D_{\mathcal{G}}$ (the nodes of this graph are the entities in $\mathcal{G}$, and the arcs indicate directed percolation links). For the second (generation) process, we let λ denote the rate (per entity) at which a new product is generated. Provided that the ratio ρ/λ is not too small, the entities evolve via cumulative, adaptive change, with the only variation being between those in $\mathcal{G}$ that have not yet switched to a new product and those that have[8]. Moreover, for large values of ρ/λ, each percolation step will be complete before the next new product is generated and so there will be no variation amongst the entities.

Conversely, as $\rho/\lambda \to 0$, the entities diverge from each other with respect to the products they generate. Their complexity (which may be estimated through analysis of the maximal number of improvements of any product at any given time for the entities in $\mathcal{G}$) is expected to be lower than in the previous scenario of shared cumulative adaptive change. We formalise this model more precisely, and provide a justification for these statements above, based on mathematical modelling, in the Appendix.

5. WHY EVOLUTION WITHOUT VARIATION AND SELECTION IS POSSIBLE

We have shown how entities that possess an abstract structure mathematically described by RAFs can exhibit evolution without variation and selection. We refer to this kind of evolutionary process as Self–Other Reorganisation (SOR) because it consists of (i) self-organization of autocatalytic networks, and (ii) interactions amongst such networks that alter their potential for future configurations [18, 19, 21, 73]. SOR involves not *competition and survival of some at the expense of others*, but *transformation of all*.[9] Like natural selection, SOR has mechanisms

---

[8]This might be likely in an cultural setting when the entity generating a new product is equidistant from all the others or (more particularly in a social media context) where everyone sees something at the same time.

[9]SOR is an evolutionary process, and therefore by definition adaptive, but we note that communal interaction in and of itself is not necessarily beneficial to the recipient. Transmission of useful plasmids through horizontal

for preserving continuity and introducing novelty, but unlike natural selection, reproduction is a low-fidelity process, because it is the culmination of haphazard catalyzed interactions, as opposed to the accuracy afforded by copying from a code, as in a Darwinian process. The distinction between Darwinian evolution and SOR is summarised in Table 3 and illustrated in Fig. 3.

In a cultural context, the external component of SOR may involve, not direct communal exchange with other individuals, but individual learning of information from the natural environment or from products others have put into the environment [23, 24]. In learning, the internal component of SOR involves accommodation, i.e., self-organization of one's semantic network to accommodate new information. We emphasize that both social learning and individual learning involve foodset elements because the newly-acquired information is available as initial reactants; it does not have to be constructed from foodset elements to exist. This brings us to a significant advantage of the approach: because it distinguishes new items that arise as a result of creative thinking (modeled as foodset-*derived* items) from items that were previously available (modeled as foodset items) the approach offers a straightforward means of identifying where novelty first arises in a lineage and tracking how it gets embellished or modified over time [23, 24].

exchange among bacteria or protists may be beneficial, but transmission of viruses may be damaging. Similarly, transmission of useful technologies may be beneficial to the recipient, but transmission of misinformation may be harmful.

FIGURE 3. Comparison of (a) evolution through natural selection and (b) evolution by Self–Other Reorganisation (SOR). An X indicates that the entity was selected against (i.e., did not contribute to the next generation). Individual entities are represented by circles, with variation among them represented by different patterns. (In the example shown in (b), there is variation but there need not be, as illustrated in Fig. 2). In both (a) and (b), the darker the color the better adapted the individual (regardless of pattern). In (a), the darker color of one of the individuals in the third generation is due to mutation. Both (a) and (b) result in cumulative, adaptive change over time. However, (a) is relatively slow because it works due to differential replication of randomly generated heritable variation in a population over generations, such that some traits become more prevalent than others, with acquired changes being discarded at the end of each generation. In contrast, change in (b) is entirely due to acquired change over time.

Feature	Variation–Selection	Self–Other Reorganisation (SOR)
Unit of self-replication[10]	Organism	RAF network
Preservation of continuity	Reproduction (vertical)	Communal exchange (horizontal)[11]
Generation of novelty	Mutation, recombination	Creativity, catalysis, transmission error
Self-assembly code	DNA or RNA	None
High fidelity	Yes	No
Transmission of acquired traits	No	Yes
Type	Selectionist	Lamarckian (by some standards; see Footnote 1)
Evolution processes explained	Biological	Early life, horizontal gene transfer, culture

TABLE 3. Comparison between evolution through selection and evolution through Self–Other Reorganisation.

SOR is distinctly different from a Darwinian or selectionist process. Recall that by 'self-assembly code' we mean a set of coded instructions that is: (i) *actively interpreted* through a developmental process to generate a soma, and (ii) *passively copied without interpretation* through a reproduction process to generate self-copies, i.e., new sets of self-assembly instructions that are in turn used in these two distinct ways. As mentioned above, in Darwinian evolution, interpretation of the self-assembly code in development is separate from the uninterpreted copying of the self-assembly code in reproduction. As a result, while acquired traits (those obtained during an organism's lifetime such as a tattoo, or knowledge of a recipe) may affect the soma, and potentially affect interpretation of the self-assembly instructions, they do not affect the *content* of the self-assembly instructions. Therefore, though inherited traits (e.g., blood type) get transmitted vertically from parent to offspring by way of genes, acquired traits do not. Thus, it is the division of labour between these two ways of using the self-assembly

[10]We do not use the term 'replicator' here since it is often assumed that replicator evolution necessarily involves inheritance of germ-line material.

[11]Note that communal exchange is not necessarily beneficial to the recipient. Transmission of useful plasmids through horizontal exchange among bacteria or protists may be beneficial, but transmission of viruses may be damaging. Similarly, transmission of useful technologies may be beneficial to the recipient, but transmission of misinformation may be harmful.

code that is responsible for the *sine qua non* of a Darwinian process: lack of transmission of acquired traits.

Now let us turn to the earliest structures that could be said to be alive, prior to the evolution of something as complex as a DNA- or RNA-based self-assembly code. Without a self-assembly code, there were no vertically inherited traits; all change was horizontally transmitted (i.e., acquired). Therefore, the evolution of early life was not due to differential replication of heritable variation in response to selection; it was non-Darwinian [71, 77, 16]. The situation is analogous for culture. Human culture does not (currently) involve a self-assembly code that gets used in these two distinct ways, and as such, it does not exhibit this signature trait of a selectionist process. No traits are inherited by way of a self-assembly code; all change is acquired, and acquired change is transmitted. We posit that both early life and culture evolve through SOR.

A Darwinian explanation works in biology to the extent that retention of acquired change is negligible compared with retention of selected change; otherwise, the first (which can operate instantaneously) can overwhelm the second (which takes generations). Moreover, the lengthy period we associate with biological generations, from birth through to reproductive maturity and parenthood, is in general significantly longer than the stretch of time between when an individual acquires a cultural trait (e.g., an idea) and then expresses (their own version of, or their own take on) that cultural trait. This can make the relative speed of Darwinian evolution slower still.

One might suggest that natural language is a cultural self-assembly code. However, (1) although natural language involves *encoding,* it is not a set of encoded instructions for the self-replication of natural languages, and (2) language evolution does not exhibit the signature characteristic of evolution by way of a self-assembly code: lack of transmission of acquired traits. Though some have argued that humans are biologically predisposed to achieve language, language itself is characterised by horizontal—not vertical—transmission. Therefore, language evolution is not due to the mechanism Darwin proposed: differential replication of heritable

variation in response to selection [14, 18]. Results from computation modelling suggest that to cross the Darwinian threshold from non-selectionist to selectionist evolution requires the emergence of a self-assembly code [71]. There is no evidence that language or any other component of culture has crossed this threshold, but it is possible that culture is moving toward a 'cultural Darwinian threshold'; in other words, it may exist in the state biological life was in before the last universal common ancestor [76].

6. IMPLICATIONS

The feasibility of evolution in the absence of variation and selection —i.e., evolution by SOR— and the fact that early life and cultural evolution are both promising candidates for this second form of evolution, imply that we must be cautious about applying concepts and methodologies developed in a Darwinian evolutionary context in these two domains. Since biological acquired traits are usually (though not always) discarded, and since a self-assembly code must stay intact to preserve its self-replication capacity, the joining of bifurcations in biological lineages is infrequent; thus, a phylogenetic tree correctly captures its branching structure. Speciation makes inter-lineage transfer of information relatively rare in biological organisms. By comparison, since cultural acquired traits are not discarded, and there is no cultural self-assembly code, the joining of bifurcations in cultural 'lineages' is commonplace, and thus cultural 'lineages' tend to be network-like rather than tree-like [15, 51, 68]. These distinctions become clear in phylogenetic studies. For multicellular species, phylogenetic trees calculated for different protein or RNA sequences tend to be conserved. However, phylogenetic trees constructed for different protein or RNA sequences across protists, bacteria or viruses reveal much more blending and greater network-like organisation [65]. Gene expression profiles exhibit cross-similarity and greater convergent evolution as well. Since cultural relatedness frequently arises through horizontal (inter-lineage) transmission, there is extensive blending of knowledge from

different sources. Extensive horizontal transmission gives a bushy, reticulated appearance to a phylogenetic tree, which is misleading, because it implies not just chronology but ancestry.

Another (related) problem that arises when methods developed for selectionist evolutionary processes are applied to culture is due to convergent evolution, in which similar forms arise independently because they are alternative solutions within similar design constraints.[12] Because biological organisms must solve many problems (reproduction, locomotion, digestion, etc.), the probability that a species will be mis-categorised because of convergent evolution (i.e., on the basis of how it solves any one problem) is low. Cultural artifacts, on the other hand, are generally constructed with a single use in mind (though artifacts developed for use in one context may be used to solve other problems; for example, a screwdriver may be used to open a can of paint.) Thus, for cultural outputs, the probability of mis-categorisation arising through the assumption that similarity reflects homology is significantly higher. Therefore, the cost of proceeding as if a Darwinian framework were applicable to culture when it is not is high. This problem does not arise with SOR because it does not assume that superficially similar artifacts are homologous. Indeed, a SOR take on culture suggests that what is evolving is the structures of the semantic networks that generate artifacts, not the artifacts themselves. Some have claimed that in practice this does not invalidate a phylogenetic approach to culture [28]. However, such conclusions come from analyses of datasets that involve little horizontal transmission (indeed, the creative blends that are the source of cultural novelty are often treated as 'outliers' and are intentionally discarded from analysis).

Such considerations have led some to develop network-based models of cultural evolution [4, 9, 11, 21, 20, 46, 51, 70]. This body of research suggests that horizontal transmission can significantly alter the pattern of relatedness. For example, a network-based analysis of Baltic

---

[12]Examples include (i) the body shape and structure similarity between the Tasmanian tiger and the fox, and (ii) wasp-imitating hover flies.

psaltery data that incorporated not just superficial physical attributes but also abstract conceptual attributes (such as markings indicative of sacred symbolic imagery), it was possible to resolve ambiguities arising from a phylogenetic analysis and generate a lineage more consistent with other historical data [70]. Horizontal cultural transmission may involve change in superficial features despite a preservation of deep structure, as occurs in metaphor [47], analogy [26, 33], and cross-domain transfer, in which a source from one domain (e.g., music) inspires or influences a creative work in another (e.g., painting) [59, 61]. This kind of complexity and hierarchical structure cannot be captured without taking a network approach to cultural evolution, which provides further support for the theory that culture evolves through SOR.

Interestingly, similar issues arise with the simplest life forms. Because of phenomena such as mutualism, lineage reticulation (due to horizontal gene transfer and allopolyploidy—the combining the genomes of different parental species), certain traits evolve with astonishing speed, thereby diminishing the continuity and distinctiveness of species [58, 77]. Indeed, the stability of genetic information is so compromised that sequencing three *Escherichia coli* genomes revealed that fewer than 40% of the genes were common to all three [74]. As a result, the boundaries between many prokaryote species are fuzzy, and exhibit reticulate patterns of evolution, thus calling into question the appropriateness of the notion of the "tree of life" [31, 41, 53]. The limitations of Darwinism as an explanation of the forms and dynamics of living things is increasingly recognised, while the role of epigenetic processes has become increasingly appreciated. Nevertheless, because such phenomena are much less present in biological evolution than cultural evolution, natural selection provides a reasonable approximation.

7. Concluding comments

Using RAF networks, this paper demonstrates by way of example, or 'existence proof' that evolution is possible in the absence of variation, selection, and competition. We refer to this kind of evolution as Self–Other Reorganisation (SOR) because it occurs not through competitive

exclusion such that only the fittest reproduce, but through the assimilation, restructuring, and exchange of components. SOR is a primitive form of evolution that can generate change quickly since it doesn't require the discarding of acquired traits, and it can take place in the absence of birth or death. Because it does not use precisely coded self-assembly instructions, it is more haphazard than a natural or artificial selectionist process, but sufficient for cumulative, adaptive, open-ended change.

RAFs have proven useful in two areas of research that might appear to be quite distinct, but that both involve networks and evolutionary processes. In the OOL, the underlying reaction system is a biochemical reaction network, and in the OOC it is a semantic network. Since cultural evolution lacks a self-assembly code, and lacks the signature characteristic of Darwinian evolution—discarding of acquired traits at the end of a generation—it seems reasonable to apply the RAF formalism and to posit that culture evolves through SOR. This is consistent with the proposal that the entities that evolve through culture are not discrete artifacts or units of behaviour such as songs, jokes or memes, but the minds that generate them [14], which have a structure that is self-organising and self-maintaining [29, 52], and therefore lends itself to the application of RAF formalisms [22, 24, 23]. We suggest that SOR may be operative during the early stage of any evolutionary process, that it was responsible for both the origin and early evolution of both organic life and human culture, and that RAFs provide a means of modeling evolution through SOR. We further speculate that human culture may be at a relatively early stage of its evolution.

In the extreme, SOR could work with just one entity; indeed, learning can be considered a kind of SOR. It occurs through (i) social transmission, or (ii) individual learning of information obtained from products others have put into the environment, or from elements of the natural environment [23, 24]. In learning, the internal component of SOR involves self-organization of one's semantic network to accommodate the new information. However, SOR is a broader concept than learning, as it can produce adaptive change in groups of entities, and it can

span generations. Competition may occur in SOR, but unlike selectionist evolution it is not essential. Since SOR does not require competition or variation, it may increase homogeneity among members of a culturally evolving group of individuals by increasing the amount of shared knowledge and experience among group members. It may thereby play an important role in fostering group identity, cohesion, and cooperation [73] (see also [1]).

8. Acknowledgements

We thank three reviewers for numerous helpful comments and suggestions on an earlier version of this manuscript.

References

[1] C. Andersson and P. Törnberg, *Toward a macroevolutionary theory of human evolution: The social protocell*, Biological Theory, 14 (2019), pp. 86–102, https://doi.org/10.1007/s13752-018-0313-y.

[2] D. Baum, *The origin and early evolution of life in chemical composition space*, Journal of Theoretical Biology, 456 (2018), pp. 295–304.

[3] R. Boyd and P. Richerson, *Culture and the evolutionary process*, University of Chicago Press, Chicago, IL, 1988.

[4] A. Buskell, M. Enquist, and F. Jansson, *A systems approach to cultural evolution*, Palgrave Communications, 5 (2019), pp. 1–15.

[5] K. R. Cabell and J. Valsiner, *The catalyzing mind: Beyond models of causality (Annals of Theoretical Psychology, Volume 11)*, Springer, Berlin, 2013, https://doi.org/10.1007/978-1-4614-8821-7.

[6] L. L. Cavalli-Sforza and M. W. Feldman, *Cultural transmission and evolution: A quantitative approach*, Princeton University Press, Princeton, NJ, 1981.

[7] R. Cazzolla Gatti, B. Fath, W. Hordijk, S. Kauffman, and R. Ulanowicz, *Niche emergence as an autocatalytic process in the evolution of ecosystems*, Journal of Theoretical Biology, 454 (2018), pp. 110–117, https://doi.org/10.1016/j.jtbi.2018.05.038.

[8] A. Cornish-Bowden and M. L. Cárdenas, *Life before LUCA*, Journal of Theoretical Biology, 434 (2017), pp. 68–74.

[9] M. Enquist, S. Ghirlanda, and K. Eriksson, *Modelling the evolution and diversity of cumulative culture*, Philosophical Transactions of the Royal Society B: Biological Sciences, 366 (2011), pp. 412–423, https://doi.org/10.1098/rstb.2010.0132.

[10] P. Erdös and A. Rényi, *On the evolution of random graphs*, Publication of the Mathematical Institute of the Hungarian Academy of Sciences, 5 (1960), pp. 17–61.

[11] L. Gabora, *Meme and variations: A computer model of cultural evolution*, in 1993 Lectures in Complex Systems, A. Goel, C. Seifert, and C. Freska, eds., Addison-Wesley, Boston MA, 1995, pp. 471–486.

[12] L. Gabora, *Autocatalytic closure in a cognitive system: A tentative scenario for the origin of culture*, Psycoloquy, 9 (1998), pp. [adap–org/9901002].

[13] L. Gabora, *Conceptual closure: How memories are woven into an interconnected worldview.*, in Closure: Emergent Organizations and their Dynamics, G. Van de Vijver and J. Chandler, eds., no. 901 in Annual Review Series, Annals of the New York Academy of Sciences, 2000, pp. 42–53, https://doi.org/10.1111/j.1749-6632.2000.tb06264.x.

[14] L. Gabora, *Ideas are not replicators but minds are*, Biology and Philosophy, 19 (2004), pp. 127–143.

[15] L. Gabora, *The fate of evolutionary archaeology: Survival or extinction?*, World Archaeology, 38 (2006), pp. 690–696.

[16] L. Gabora, *Self-other organization: Why early life did not evolve through natural selection*, Journal of Theoretical Biology, 241 (2006), pp. 443–450.

[17] L. Gabora, *Five clarifications about cultural evolution*, Journal of Cognition and Culture, 11 (2011), pp. 61–83.

[18] L. Gabora, *An evolutionary framework for culture: Selectionism versus communal exchange*, Physics of Life Reviews, 10 (2013), pp. 117–145, https://doi.org/10.1016/j.plrev.2013.03.006.

[19] L. Gabora, *From deep learning to deep reflection: Toward an appreciation of the integrated nature of cognition and a viable theoretical framework for cultural evolution*, in Proceedings of the 2019 Annual Meeting of the Cognitive Science Society, L. Nadel and D. D. Stein, eds., Cognitive Science Society, Austin, TX, 2019, pp. 1801–1807.

[20] L. Gabora, S. Leijnen, T. Veloz, and C. Lipo, *A non-phylogenetic conceptual network architecture for organizing classes of material artifacts into cultural lineages*, in Proceedings of the 33rd annual meeting of the Cognitive Science Society, L. Carlson, C. Hőlscher, and T. F. Shipley, eds., Cognitive Science Society, Psychology Press, 2011, pp. 2923–2928.

[21] L. Gabora and C. Smith, *Two cognitive transitions underlying the capacity for cultural evolution*, Journal of Anthropological Science, 96 (2018), pp. 27–52, `https://doi.org/10.4436/jass.96008`.

[22] L. Gabora and M. Steel, *Autocatalytic networks in cognition and the origin of culture*, Journal of Theoretical Biology, 431 (2017), pp. 87–95, `https://doi.org/10.1016/j.jtbi.2017.07.022`.

[23] L. Gabora and M. Steel, *A model of the transition to behavioral and cognitive modernity using reflexively autocatalytic networks*, Proceedings of the Royal Society Interface, 17 (2020), p. 20200545, `https://doi.org/http://doi.org/10.1098/rsif.2020.0545`.

[24] L. Gabora and M. Steel, *Modeling a cognitive transition at the origin of cultural evolution using autocatalytic networks*, Cognitive Science, 44 (2020).

[25] L. Gabora and S. Tseng, *The social benefits of balancing creativity and imitation: Evidence from an agent-based model*, Psychology of Aesthetics, Creativity, and the Arts, 11 (2017), pp. 457–473.

[26] D. Gentner, *Structure-mapping: A theoretical framework for analogy*, Cognitive Science, 7 (1983), pp. 155–170, `https://doi.org/10.3389/fpsyg.2019.01426`.

[27] N. Goldenfeld, T. Biancalani, and F. Jafarpour, *Universal biology and the statistical mechanics of early life*, Philosophical Transactions of the Royal Society A, 375 (2017), p. 20160341, `https://doi.org/2017`.

[28] S. J. Greenhill, T. E. Currie, and R. D. Gray, *Does horizontal transmission invalidate cultural phylogenies?*, Proceedings of the Royal Society B: Biological Sciences, 276 (2009), pp. 2299–2306.

[29] A. G. Greenwald, M. R. Banaji, L. A. Rudman, S. D. Farnham, B. A. Nosek, and D. S. Mellott, *A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept*, Psychological Review, 109 (2002), pp. 3–25.

[30] J. Henrich and R. McElreath, *Dual-inheritance theory: The evolution of human cultural capacities and cultural evolution*, in Oxford Handbook of Evolutionary Psychology, Oxford University Press, 2007, pp. 741–750.

[31] E. Hilario and J. P. Gogarten, *Horizontal transfer of atpasegenes: The tree of life becomes a net of life*, Biosystems, 31 (1993), pp. 111–119.

[32] J. H. Holland, *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*, MIT Press, Urbana, 1992.

[33] K. J. Holyoak and P. Thagard, *Mental leaps: Analogy in creative thought*, MIT Press, Cambridge, MA, 1996.

[34] W. HORDIJK, J. HEIN, AND M. STEEL, *Autocatalytic sets and the origin of life*, Entropy, 12 (2010), pp. 1733–1742, `https://doi.org/10.3390/e12071733`.

[35] W. HORDIJK, S. A. KAUFFMAN, AND M. STEEL, *Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems*, International Journal of Molecular Science, 12 (2011), pp. 3085–3101, `https://doi.org/10.3390/ijms12053085`.

[36] W. HORDIJK AND S. M., *A formal model of autocatalytic sets emerging in an rna replicatior system*, Journal of the Systems Chemistry, 4:3 (2013), `https://doi.org/10.1186/1759-2208-4-3`.

[37] W. HORDIJK AND M. STEEL, *Detecting autocatalytic, self-sustaining sets in chemical reaction systems*, Journal of Theoretical Biology, 227 (2004), pp. 451–461, `https://doi.org/10.1016/j.jtbi.2003.11.020`.

[38] W. HORDIJK AND M. STEEL, *Autocatalytic sets and boundaries*, J. Syst. Chem., 6:1 (2015).

[39] W. HORDIJK AND M. STEEL, *Chasing the tail: The emergence of autocatalytic networks*, Biosystems, 152 (2016), pp. 1–10, `https://doi.org/10.1016/j.biosystems.2016.12.002`.

[40] W. HORDIJK, M. STEEL, AND P. DITTRICH, *Autocatalytic sets and chemical organizations: Modeling self-sustaining reaction networks at the origin of life*, New Journal of Physics, 20 (2018), p. 015011.

[41] A. L. HUGHES, *The origin of adaptive phenotypes*, Proceedings of the National Academy of Sciences USA, 105 (2008), pp. 13193–13194.

[42] S. A. KAUFFMAN, *Autocatalytic sets of proteins*, Journal of Theoretical Biology, 119 (1986), pp. 1–24, `https://doi.org/10.3390/ijms12053085`.

[43] S. A. KAUFFMAN, *The origins of order*, Oxford University Press, 1993.

[44] P. KILLEEN, *The non-Darwinian evolution of behavers and behaviors*, Behavioural Processes, 161 (2019), pp. 45–53.

[45] A. C. KIMBER, *A note on Poisson maxima*, Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete, 63 (1983), pp. 551–552.

[46] S. KIRBY, *Culture and biology in the origins of linguistic structure*, Psychonomic Bulletin & Review, 24 (2017), pp. 118–137.

[47] G. LAKOFF, *The contemporary theory of metaphor*, in Metaphor and thought (2nd ed), A. Ortony, ed., Cambridge University Press, Cambridge, UK, 1993, pp. 202–252.

[48] D. LANCET, R. ZIDOVETZKI, AND O. MARKOVITCH, *Coevolution of compositional protocells and their environment*, Journal of the Royal Society Interface, 15:20180159 (2018).

[49] C. G. LANGTON, *Artificial life*, in Lectures in Complex Systems, Vol. IV, L. Nadel and D. Stein, eds., MIT Press, Reading, MA, 1992, pp. 189–242.

[50] R. LEWONTIN, *The units of selection*, Annual Review of Ecology and Systematics, 61 (1970), pp. 1–18.

[51] C. P. LIPO, *The resolution of cultural phylogenies using graphs*, in Mapping our Ancestors, C. P. Lipo, M. J. O'Brien, M. Collard, and S. J. Shena, eds., vol. 5, Aldine de Gruyter, New York, 2005, pp. 89–107.

[52] H. MATURANA AND F. VARELA, *Autopoiesis and cognition: The realization of the living*, in Boston Studies in the Philosophy of Science, R. S. Cohen and M. W. Wartofsky, eds., vol. 42, Reidel, Dordecht, 1973.

[53] J. O. MCINERNEY, D. PISANI, E. BAPTESTE, AND M. J. O'CONNELL, *The public goods hypothesis for the evolution of life on earth*, Biology Direct, 6 (2011).

[54] A. MESOUDI, A. WHITEN, AND L. K. N., *Is human cultural evolution Darwinian? Evidence reviewed from the perspective of the origin of species*, Evolution, 58 (2004), pp. 1–11.

[55] E. MOSSEL AND M. STEEL, *Random biochemical networks and the probability of self-sustaining autocatalysis*, Journal of Theoretical Biology, 233 (2005), pp. 327–336, `https://doi.org/10.1016/j.jtbi.2004.10.011`.

[56] G. B. MÜLLER, *Why an extended evolutionary synthesis is necessary*, Interface Focus, 7 (2017), p. 20170015.

[57] M. MUTHUKRISHNA, M. DOEBELI, M. CHUDEK, AND J. HENRICH, *The cultural brain hypothesis: How culture drives brain expansion, sociality, and life history*, PLoS Computational Biology, 14 (2018), p. e1006504, `https://doi.org/10.1371/journal.pcbi.1006504`.

[58] G. J. P. AND J. P. TOWNSEND, *Horizontal gene transfer, genome innovation and evolution*, Nature Reviews, 3 (2005), pp. 679–687.

[59] A. RANJAN, L. GABORA, AND B. O'CONNOR, *Evidence that cross-domain re-interpretations of creative ideas are recognizable*, in Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium (Creativity and Cognitive Development: A Perspective from Artificial Creativity, Developmental Artificial Intelligence, and Robotics), AAAI Press, Menlo Park, USA, 2013.

[60] P. RICHERSON AND R. BOYD, *A dual inheritance model of the human evolutionary process i: Basic postulates and a simple model*, Journal of Social and Biological Structures, 1 (1978), pp. 127–154.

[61] V. SCOTNEY, S. WEISSMEYER, N. CARBERT, AND L. GABORA, *The ubiquity of cross-domain thinking in the early phase of the creative process*, Frontiers in Psychology, 10 (2019), p. 1426, `https://doi.org/10.3389/fpsyg.2019.01426`.

[62] D. SEGRÉ, D. BEN-ELI, AND D. LANCET, *Compositional genomes: Prebiotic information transfer in mutally catalytic noncovalent assemblies*, Proceedings of the Royal Society B, 97 (200), pp. 4112–4117.

[63] R. SEMON, *The mneme*, Allen and Unwin, London, 1921.

[64] B. SHENHAV, A. OZ, AND D. LANCET, *Coevolution of compositional protocells and their environment*, Philosophical Transactions of the Royal Society B, 362 (2007), pp. 1813–1819.

[65] M. SHI, X. LIN, AND J. TIAN, *Redefining the invertebrate rna virosphere*, Nature, 540 (2016), pp. 539–543, `https://doi.org/10.1038/nature20167`.

[66] M. STEEL, *The emergence of a self-catalyzing structure in abstract origin-of-life models*, Applied Mathematics Letters, 13 (2000), pp. 91–95.

[67] M. STEEL, W. HORDIJK, AND J. C. XAVIER, *Autocatalytic networks in biology: Structural theory and algorithms*, Journal of the Royal Society Interface, 16 (2019), p. rsif.2018.0808, `https://doi.org/10.1098/rsif.2018.0808`.

[68] I. TEMKIN AND N. ELDREDGE, *Phylogenetics and material cultural evolution*, Current Anthropology, 48 (2007), pp. 146–153.

[69] V. VASAS, C. FERNANDO, M. SANTOS, S. KAUFFMAN, AND E. SZATHMÁRY, *Evolution before genes*, Biology Direct, 7 (2012).

[70] T. VELOZ, I. TEMPKIN, AND L. GABORA, *A conceptual network-based approach to inferring cultural phylogenies*, in Proceedings of the 34th annual meeting of the Cognitive Science Society, N. Miyake, D. Peebles, and R. P. Cooper, eds., Cognitive Science Society, Austin TX, 2012, pp. 2487–2492.

[71] K. VETSIGIAN, C. WOESE, AND N. GOLDENFELD, *Collective evolution and the genetic code*, Proceedings of the National Academy of Sciences USA, 103 (2006), pp. 10696–10701.

[72] J. VON NEUMANN, *Theory of Self-Replicating Automata*, University of Illinois Press, Urbana, 1966.

[73] B. VOORHEES, D. READ, AND L. GABORA, *Identity, kinship, and the evolution of cooperation*, Current Anthropology, 61 (2020), pp. 194–218.

[74] R. A. WELCH, *Extensive mosaic structure revealed by the complete genome sequence of uropathogenic Escherichia coli*, Proceedings of the National Academy of Sciences USA, 99 (2002), pp. 17020–17024.

[75] A. WHITEN, *A second inheritance system: the extension of biology through culture*, Interface Focus, 7 (2017), p. 20160142.

[76] C. WOESE, *The universal ancestor*, Proceedings of the National Academy of Sciences, 95 (1998), pp. 6854–6859.

[77] C. R. Woese, *On the evolution of cells*, Proceedings of the National Academy Science, 99 (2002), pp. 8742–8747.

[78] J. C. Xavier, W. Hordijk, S. Kauffman, M. Steel, and W. F. Martin, *Autocatalytic chemical networks at the origin of metabolism*, Proceedings of the Royal Society of London. Series B: Biological Sciences, 287 (2020), p. 20192377.

9. Appendix: Mathematical modelling and justification for the predictions of cumulative adaptive change vs individual-based change

To help formalise the model described in Section 4, we adopt the following terminology. Given the products $p, p', p'', \ldots$, we say that the entity that first gives rise to one of these products *generates* it. In the case where this product is produced from an earlier item in the series (rather than being the original product p in response to a stimulus), we say the entity *transforms* the existing item to the new one.

We model the increase in the adaptive value of products within $\mathcal{G}$ using two stochastically independent non-deterministic processes.

First, the generation of new products by entity i involves either generating a new product or transforming the most recent version of any product it has available; for simplicity, we assume that these are all equally probable. For example, if entity i has the products (p, p', q, r, r', r'') currently available, then it can either transform p', or q or r'', or generate a new product, with each of these four outcomes having the same probability, namely $1/4$. This process across entities is assumed to be described by independent exponential random variables with a fixed rate per entity of λ (making λ independent of the entity is consistent with the assumption that all entities are initially identical).

Second, each newly generated product then begins to percolate through the group by moving along the arcs of the directed graph $D_{\mathcal{G}}$ according to a continuous-time random walk on

this directed graph at rate ρ. For simplicity, we will treat the two processes (generation and percolation) as stochastically independent.

Let $N = |\mathcal{G}|$ (the number of entities in $\mathcal{G}$). We will assume that the directed graph $D_{\mathcal{G}}$ that describes the community interactions within $\mathcal{G}$ has the property that from each entity in $\mathcal{G}$, it is possible to reach any other entity in $\mathcal{G}$ by following some directed path in $D_{\mathcal{G}}$ (i.e. $D_{\mathcal{G}}$ is 'strongly connected'); however, no further assumptions are made regarding this graph. We now introduce some notation to keep track of the different versions of products that arise in the process described in Section 4. Suppose that product p first arises from entity i_1 and product p is then further modified by entity i_2 and so on, with the last modification (before time t) being made by entity i_k. In that case, we denote the sequence of products thus generated within $\mathcal{G}$ up to time t as: $p(i_1), p(i_1, i_2), p(i_1, i_2, i_3), \ldots$. More generally, we denote such a sequence by writing $(p(i_1), \ldots, p(i_1, i_2, \ldots, i_k) : k \geq 1)$ (thereby allowing the possibility that a product is generated but not transformed, in the case where $k = 1$). We refer to the number k of terms in this sequence as the *complexity* of the final product; thus, when an entity transforms a product, it increases its complexity by 1 (in particular, the first product $p(i_1)$ has complexity 1).

Note that under the assumptions of the model, the entities $i_1, \ldots, i_k$ are not assumed to necessarily be distinct (i.e., an entity may enhance a product more than once, either consecutively or later in the sequence). There may also be several such sequences generated within $\mathcal{G}$; for example, in addition to the previous sequence, one might also have $p(j_1), p(j_1, j_2), \ldots, p(j_1, j_2, \ldots, j_l)$, along with possibly other sequences generated over the time interval $[0, t]$.

Let $T_\rho(i)$ be the expected time for a product generated by entity i to percolate (within $D_{\mathcal{G}}$) to every entity in $\mathcal{G}$, and let $T_\rho = \max\{T_\rho(i) : i \in G\}$. For a wide range of standard percolation processes, the following properties then hold: (i) for $\rho > 0$, we have $\mathbb{E}[T_\rho] < \infty$; (ii) for all $\eta > 0$, $\lim_{\rho \to 0} \mathbb{P}(T_\rho(i) > \eta) = 1$, and (iii) $\lim_{\rho \to \infty} \mathbb{E}[T_\rho] = 0$. This last property implies that when ρ is large, items are highly likely to percolate throughout the entire group $\mathcal{G}$ in a short time.

If we start this process at time 0 with no products present, let $\tau_1, \tau_2, \ldots, \tau_k$ be the random variables that describe the time intervals between the generation of products across the collection of entities in $\mathcal{G}$. By the assumptions of the model, the τ_i variables are independent and exponentially distributed, with each variable having an expected value of $1/(N\lambda)$. Thus $\sum_{i=1}^{k} \tau_i$ is the time until k products have been generated (this has a gamma distribution with expected value $k/(N\lambda)$). Let $\mu = N\lambda$. Then, $\mathbb{P}\left(\bigcap_{i=1}^{k}\{\tau_i \geq \eta\}\right) = e^{-\mu k \eta}$ and $\mathbb{P}(T_\rho \leq \eta) \geq 1 - \mathbb{E}[T_\rho]/\eta$ (by Markov's inequality). Let $\mathcal{E}_k$ denote the following event: for each of the first k products generated, each product percolates to each entity in $\mathcal{G}$ before the next new product (in this collection of size k) is generated in $\mathcal{G}$. We then have:

$$(1) \qquad \mathbb{P}(\mathcal{E}_k) \geq e^{-\mu k \eta} \cdot (1 - \mathbb{E}[T_\rho]/\eta)^k = (e^{-\mu \eta}(1 - \mathbb{E}[T_\rho]/\eta))^k.$$

Setting $\eta = \sqrt{\mathbb{E}[T_\rho]}$ in (1) and applying Property (iii) above gives:

$$\lim_{\rho \to \infty} \mathbb{P}(\mathcal{E}_k) = 1.$$

Thus, as ρ becomes large, the entities evolve collectively, and any variation is transient and short-lived. We will refer to this limiting case as the *community-based model*. One can model this process by the following novel type of Pólya Urn model:

Consider an urn that initially has a single white ball. At each step (the timing of which follows a Poisson process at rate r), a ball is selected from the urn uniformly at random. If the selected ball is white, it is returned to the urn along with a ball of a new colour (not present already in the urn). If the selected ball is coloured, it is removed and replaced by a ball of the same colour but a darker shade.

To connect this urn process to the community-based model described above, note that selecting a white ball corresponds to the generation of a new product (which results in a ball of

a new colour being added to the urn), while selecting a ball of a darkest colour corresponds to the transformation of an existing product. Thus $r = N\lambda$.

We now compare the community-based model (corresponding to ρ large) to the opposite extreme, where ρ becomes small. In that case, the probability that there is percolation between any two entities in $\mathcal{G}$ over the interval $[0, t]$ tends to 0, and so products are only generated within entities but not shared between them. We will refer to this limiting case as the *individual-based model*. Note that in this individual-based model, entity i_j may possibly generate a new product $p(i_j)$, or it may generate $p(i_j)$ and then transform it (producing $p(i_j, i_j)$ and so on (or it might not generate any new products at all). Note that, in general, $p(i_j)$ may be different from $p(i_k)$ (for $k \neq j$) (i.e., different entities may produce different products, and similarly for transformed products).

For the individual-based model, we have N independent samples of the above Urn model but with $r = \lambda$. By contrast with the community-based model, we have a single sample of the above Urn model, but with $r = N\lambda$. Note that both models have the same expected number of generation events. However, they have quite different dynamics, as we now describe.

Firstly, in the community-based model there is only short-lived or transient variation among the entities, whereas in the individual-based model, the individuals diverge from each other in terms of the collections of products that are available to them. However, a more subtle difference is that in the community model, the complexity of items is significantly higher than in the individual model, in a sense that we now make precise.

Let $X = X_t$ be the number of balls in the urn at time t and let Y_i be the number of steps of the urn process that an urn that has just arrived at having $i \geq 2$ balls remains with i balls (until it moves to $i + 1$ balls). Then Y_i has a geometric distribution with expected value $\mathbb{E}[Y_i] = i$ for all $i \geq 2$.

Next, consider the random variable W_t that counts the number of times the first coloured ball is replaced by increasingly darker shades of balls of that same colour up until time t. The

expected value of W is given by:

$$\mathbb{E}[W_t] = \mathbb{E}\left[\frac{Y_2}{2} + \frac{Y_3}{3} + \cdots + \frac{Y_{X_t}}{X_t}\right] = \mathbb{E}\left[\frac{2}{2} + \frac{3}{3} + \cdots + \frac{X_t}{X_t}\right] = \mathbb{E}[X_t] - 1.$$

Thus, in the community-based model, one product has expected complexity of at least $\mathbb{E}[X_t] - 1$ at time t (and it is easily seen that there are products with expected complexity of $\mathbb{E}[X_t] - 2, \mathbb{E}[X_t] - 3$ and so on). Moreover, we claim that $\mathbb{E}[X_t]$ is bounded below by a term of order $\sqrt{rt}$ (a proof is provided at the end of this Appendix). Thus, since the community-based model has $r = N\lambda$, we arrive at the following conclusion regarding the influence of the size of $\mathcal{G}$ on complexity:

> Over a given period of time, some products in the community-based model have an expected complexity of order at least $\sqrt{N}$.

By contrast, for the individual-based model, we have $r = \lambda$ for each entity, and so we have N independent and identically distributed samples of a process where the maximum complexity of products across each group $\mathcal{G}$ will exhibit a lower (logarithmic) dependence on N (moreover, these complex products are likely to exist only in one or a few entities, rather than being shared across the group). To see this, note that the complexity of any product within an entity (up to time t) is bounded above by the number of generation steps for that entity, which has a Poisson distribution with mean λt, and the maximum of N independent and identically distributed Poisson random variables is known to be dominated asymptotically (with N) by a $\log(N)$ term [45].

Proof that $\mathbb{E}[X_t]$ is bounded below by a term of order $\sqrt{rt}$.

We have: $X_{t+\delta} = X_t + D_t$, where $D_t = 1$ with probability $r\delta/X_t + o(\delta)$ and $D_t = 0$ otherwise (note that the number of balls in the urn increases only if the white ball is selected from the X_t balls in the urn at time t). Thus, if we let $\mu(t) = \mathbb{E}[X_t]$, applying the law of total expectation

and taking the limit as $\delta \to 0$ gives:

$$\frac{d\mu(t)}{dt} = r\mathbb{E}\left[\frac{1}{X_t}\right]. \tag{2}$$

The initial condition here is $\mu(0) = 1$. Since the function $f(x) = 1/x$ is convex, Jensen's inequality gives:

$$\frac{d\mu(t)}{dt} \geq \frac{r}{\mu(t)}. \tag{3}$$

A straightforward analytic argument now shows that $\mu(t)$ is bounded below by the solution $\tilde{\mu}(t)$ to the differential equation obtained from (3) by replacing the inequality with an equality. This equation has a unique solution satisfying the same initial condition as $\mu(t)$, namely, $\tilde{\mu}(t) = \sqrt{1 + 2rt}$, which exhibits the claimed growth rate.