

Review

Top-Down Causation and the Rise of Information in the Emergence of Life

Sara Imari Walker ^{1,2,3}

¹ Beyond Center for Fundamental Concepts in Science, Arizona State University, Tempe, AZ, USA;
E-Mail: sara.i.walker@asu.edu

² School of Earth and Space Exploration, Arizona State University, Tempe AZ, USA

³ Blue Marble Space Institute of Science, Seattle, WA, USA

Received: 4 March 2014; in revised form: 23 June 2014 / Accepted: 14 July 2014 /

Published: 21 July 2014

Abstract: Biological systems represent a unique class of physical systems in how they process and manage information. This suggests that changes in the flow and distribution of information played a prominent role in the origin of life. Here I review and expand on an emerging conceptual framework suggesting that the origin of life may be identified as a transition in causal structure and information flow, and detail some of the implications for understanding the early stages chemical evolution.

Keywords: origin of life; emergence; top-down causation

1. Introduction

Living systems are the congruence of two core functions: metabolism (reproduction) and genetic inheritance (replication). This has led many researchers to assign prominence to the appearance of one or the other as the crucial first step leading to the emergence of life. Thus, origin of life research has traditionally been divided into the respective frameworks of “metabolism-first” and “genetics-first” scenarios. For the former, emphasis is placed on the importance of energy flows through material substrates required to sustain living matter, while for the latter emphasis is placed on the propagation of genetic information. More succinctly, metabolism-first approaches are cast in term of physics (e.g., matter, energy, forces), and genetics-first cast in terms of the copying of information. These physical and informational perspectives are traditionally treated as two parallel, but distinct, narratives for describing living systems. A challenge in unifying these approaches is that the concept of information—which

relies on the existence of counterfactuals—is not easily reconciled with describing physical systems, whose dynamics are typically described by an initial state and deterministic laws. The counterfactual nature of information is evident when one considers that for a message to carry information, a different message must also be possible. Deterministic laws only allow one outcome for a given initial state and therefore do not allow more than one possibility. Thus, as physics is traditionally cast, there appears to be no room for information to physically influence the world. However, we know information does affect the world: *information is physical*, it must both be physically instantiated and—perhaps most importantly—it can influence the dynamics of physical systems [1]. The latter is particularly important in the context of understanding the nature of life: living systems do not merely passively accumulate and store information, they also actively process it [2,3]. Many of the most fundamental biochemical processes involve the transfer and processing of information [4]: DNA sequences provide a set of algorithmic instructions for the execution of specific chemical processes within cells (in addition to their role as a digital information storage repository) [5]; protein-based “circuits” determine the behavior of cells [6]; gene regulation is tightly controlled by a variety of feedback mechanisms, including switch-like behavior [7]; and even metabolism has an underlying logic [8,9]. Accommodating the active—or causal—role of information within physical theory is one of the great challenges associated with explaining life’s origins.

The causal role of information is often discussed under the conceptual framework of “top-down” causation, where higher levels of organization in structural hierarchies constrain the dynamics of lower levels of organization [10]. An example is how information encoded in the spatiotemporal patterning of electromagnetic fields dictates the patterns of a developing embryo by influencing local gene expression [11]. Top-down causation by information encoded in the current state is one of the most distinctive features of living physical systems [12–14]. In light of this, we previously suggested that the emergence of life may be associated with a transition in the causal and informational architecture of matter [15,16]. Within this proposed framework, the debate between genetics-first and metabolism-first scenarios takes on a new dimension: both may be unified under a common information-based descriptive paradigm, where genetics may be thought of in terms of digital information processing and metabolism roughly as a form of analog information processing [16]. Thus the debate on which came first—genetics or metabolism—may be recast as a debate about the informational hardware of the first living systems. A complimentary perspective to this hardware-based view must focus on the “software”, *i.e.*, the way information is managed in nascent living systems [17]. From the perspective of information flow and processing, genetics and metabolism are two core sub-functions of a common cellular operating system (COS) [14]. It is the emergence of the COS system that is the critical question to be addressed when asking how life first emerged. Cast in this light, the emergence of life is a question for physics: specifically, it is a question of how information acquires a causally efficacious role in physical systems. Here I review insights into the potential physics underlying the emergence of life based on this perspective, which suggest that the transition from non-life to life may be identified as a transition in causation and information flow. I also detail some of the implications of this approach for our understanding of the early stages chemical evolution.

2. The Limits of Chemistry

In traditional information-based—or “genetics-first”—theories for the origin of life, much of biological evolution is reduced to the chemistry of nucleic acids (such as RNA) [18]. To accomplish this reduction, two important assumptions are made: (1) genes are the fundamental units of biological information, and (2) evolution is primarily concerned with the accumulation of (genetic) information through the processes of replication and selection. These two features can be captured by appealing to chemical models for the catalyzed replication of genetic polymers, such as RNA. This forms the foundation of the RNA world hypothesis [19], which suggests that the early Earth was populated by “riboorganisms” who utilized RNA as both genetic material and the only genetically encoded biochemical catalyst.

Collapsing genetic information and catalytic function to the same molecule introduces a paradox. If genotype (genetic information) and phenotype (catalytic function) are encoded in the same molecule, there is an inherent chemical trade-off between genotypic (selection for replication) and phenotypic (selection for function) selection [20]. For example, well-folded RNA sequences are likely to be poor templates for copying, conversely poorly folded RNA sequences are unlikely to be good ribozymes (catalytic RNAs) [21]. This trade-off imposes a *physical* limit on the information content of primitive genomes. This physical constraint imposed on information content should be viewed in contrast with the more widely discussed *informational* limits of early replicators, defined by Eigen’s “error threshold” as the maximum amount of genomic information that can be reliably propagated from generation to generation for a fixed mutation rate [22]. Both the informational limit(s) on physical systems and the physical limit(s) on informational systems are important factors shaping biological evolution.

Modern life has overcome the chemical limitations imposed on template-based replication through the decoupling of selection for replication from selection for function. In effect, this amounts to physically separating information propagation (replication) from the processing of that information (function) [3]. Life achieves this decoupling by utilizing two species of biopolymer that play different biochemical roles: DNA stores genetic information and proteins execute the majority of catalytic and structural functions (with RNA playing a prominent dual role as an informational and functional intermediate). In modern life, replication of genetic information is not constrained by the physics or chemistry of particular DNA sequences. This is not to say that chemistry does not play a role, as it clearly must: thus, for example, some polymerases are more efficient and less error-prone for certain reading frames than others. However, in general the efficiency of copying sequences is not as widely varied as it is for non-enzymatic replicators. The utility of genes is that they carry *instructional* information that can be read-out by the translation machinery of the cell to execute specific functions: selection is ultimately operating on the informational output, or function, of a gene and not strictly on the gene itself.

The above considerations have led many researchers to distinguish replicators from reproducers, going as far back as the early work of von Neumann [23]. An interesting side note, highlighting the power of logic-based views of living systems, is that von Neumann conceptually formalized a distinction between trivial and non-trivial replicators based solely on logical considerations, before the inner workings of cellular self-reproduction, such as the structure of DNA, were elucidated. von Neumann identified two classes of self-replicators: trivial and non-trivial. *Trivial replicators* are constrained by strict physical

limits imposed on their capacity to generate copies, such as the chemical trade-off between templating ability and catalytic function imposed on an RNA genome encoding RNA catalysts. In contrast, *non-trivial replicators*—also commonly referred to as reproducers—have in some sense transcended many of these constraints due to the explicit use of programs, symbolic logic and codes to mediate self-reproduction. Non-trivial replicators require constructors, defined as entities that can cause some change in a physical system while retaining the ability to cause it again. Examples of constructors include heat engines and chemical catalysts. von Neumann was particularly interested in the idea of a *universal constructor*, a machine capable of taking materials from its host environment to build any possible physical structure (consistent with the available resources and the laws of physics), including itself. Universal constructors are therefore, by definition, capable of self-reproduction. A key feature of a universal constructor is that it must be programmed with an explicit set of instructions to construct itself (these systems are therefore self-referencing, another distinctive feature of the living state [24–26]).

Living systems are constructors in the sense proposed by von Neumann. A caveat is that known life may not be truly universal—it is an open question whether living systems possess the capacity to construct *any* possible object consistent with available resources; however, technological civilizations may approach this ideal. While constructors certainly exist outside of living systems (e.g., the examples of the heat engine or catalyst cited above), life represents a unique class of physical systems in that the relevant constructors are *virtual*. That is to say, in living systems the constructors are programs. For a living cell, it is the program that must be preserved through the process of construction and not necessarily the matter it is instantiated in (see, e.g., discussion on functional equivalence in Section 5). *The emergence of life is therefore describable in physical terms as emergence of causally efficacious abstractions, or virtual constructors.* As I discuss below, the emergence of virtual constructors is intimately tied to the notion of causal emergence as it requires the appearance of new “higher-level” causes that are necessarily abstract (e.g., programs). Identifying when such systems emerge therefore requires rigorous quantification of the notion of causal emergence, including “top-down” causation, which we discuss in detail in Section 3.

Identifying the role of (virtual) constructors in the early evolution of life highlights the need for an information-based understanding of how life first emerged by identifying the limits of chemistry (hardware)—only origin of life scenarios. This is perhaps most strikingly evident when one considers the relevant complexity classes for describing replicative systems, which we turn to next.

Szathmáry and Maynard Smith, for example, have offered a classification scheme that distinguishes the complexity of replicators based on hereditary potential [27]. They identify *limited heredity* replicators as those where the number of types is smaller than, or roughly equal to, the number of individuals in a realistic physical system (*i.e.*, short oligonucleotides). Alternatively, *unlimited heredity* replicators are identified as those where the number of types far exceeds the number of individuals that could possibly exist in any realistic system (*i.e.*, genes and genomes of extant life). By definition, only unlimited heredity replicators allow the potential for unlimited growth in complexity: unlimited heredity replicators occupy a vast state-space (much larger than anything physically achievable) that is mostly unoccupied by physical systems where “complex” states could, at least in principle, be populated through Darwinian evolution. The distinction between limited and unlimited heredity is therefore in the size of the relevant state space. Life is an example of the more “complex” unlimited heredity replicator, because the state

space of hypothetically possible genomes is much larger than could ever be realized by summing over all genomes on Earth.

Szathmáry and Maynard Smith’s classification cannot distinguish trivial and non-trivial replicators. It is based solely on identifying the hypothetically possible state space for the chemical hardware and does not take into account the existence of counterfactuals. However, in determining the complexity of physical systems (including life) a distinction must be made between *possible states* and *physically accessible states*. Distinguishing between possible and impossible transformations is the basis of constructor theory, currently being developed by Deutsch and Marletto [1,28]. Their proposal is to recast all fundamental laws of nature entirely in terms of statements of which tasks (*i.e.*, classes of physical transformations or constructions) are possible and which are impossible, and why. Constructor theory therefore proposes to place counterfactual statements as first-class descriptors of physical systems, a position of prominence normally only awarded to laws of motion and initial conditions [1]. Taking their proposal seriously, the subset of physically accessible states among all possible states is defined by the constructors present in a given system. This yields a new perspective on emergent phenomena—emergence can occur only when new constructors arise that open possible, but previously inaccessible, physical states (this is somewhat akin to the notion of the “adjacent possible” as popularized by Kauffman [29]). In the context of the current discussion, the complexity of trivial and non-trivial replicators is distinguished by the size of the constructible state space. The size of the constructible state space, although difficult to quantify in practice, may be a natural complexity measure for the biosphere: humans are arguably more complex than other organisms because of the space of all possible constructions we can build (contrast with the length of our genome, which does not accurately capture our “complexity” relative to other organisms).

Table 1. Intersection of trivial/non-trivial (software) and limited/unlimited heredity (hardware) replicator classification schemes. The limited/unlimited heredity classification scheme distinguishes replicators based on the total number of all possible states of the chemical substrate, whereas the trivial/non-trivial classification scheme distinguishes replicators based on the number of physically accessible states due to the presence of constructors. Living systems are members of the class of non-trivial replicators with virtual constructors, and are capable of unlimited heredity.

	Trivial	Non-trivial
Limited heredity	short oligonucleotides such as non-enzymatic template replicators, crystals	physical constructors such as catalysts and autocatalytic sets
Unlimited heredity	monomolecular genes and genomes such as would exist in RNA-world riboorganisms	virtual constructors such as the cellular operating system of extant life and von Neumann automata

The intersection of these two schema is shown in Table 1. A trivial replicator such as an RNA replicase could, for example, be a limited or unlimited heredity replicator depending on its length.

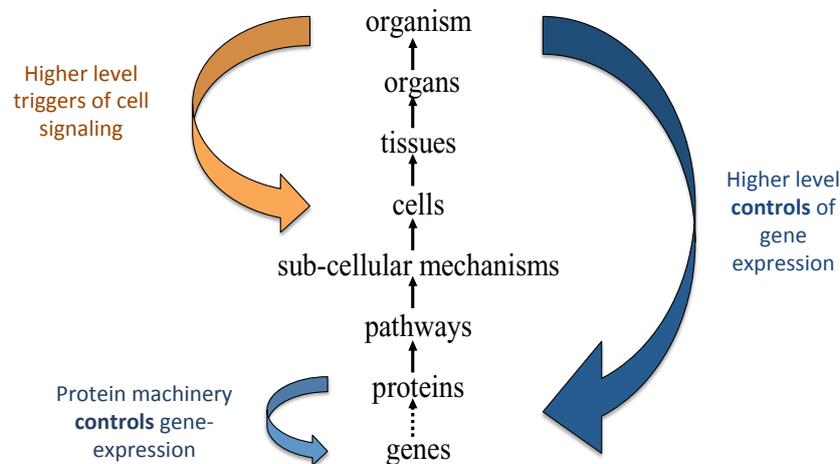
A trivial replicator with unlimited heredity may therefore appear to have open-ended evolutionary potential due to the large size of its hypothetically possible state space (e.g., consider the size of state space of a 50 nucleotide RNA—it would take more carbon than is available in the Earth’s crust to synthesize one picomole of all possible RNA sequences in this space [30]). Genetics-first RNA world scenarios are based on the premise that limited heredity replicators will evolve into unlimited heredity replicators through the process of Darwinian evolution. However, as noted above, much of this state space is physically inaccessible due to chemical trade-offs. Conversely, metabolism-first scenarios that often involve autocatalytic sets (sets of mutually catalytic species that collectively reproduce) provided examples of non-trivial replicators in that they explicitly rely on the presence of constructors (catalysts). However, they are limited in their potential for open-ended growth in complexity (limited heredity) by the physical and chemical limitations imposed on the physical constructors: in short, they are not programmable. Life arises through evolutionary processes converging on the lower left-hand box of Table 1. To understand the transition to non-trivial replicators with unlimited heredity—and thus the emergence of life—therefore requires explicit acknowledgement of the causal role of information in physical systems.

3. Top-Down Causation in Biological Systems

The forgoing discussion suggests that life first arose through the emergence of virtual constructors in chemical systems. The emergence of virtual constructors is intimately tied to the notion of causal emergence as it requires the appearance of new “higher-level” causes that are necessarily abstract (e.g., implemented through the use of symbolic logic and programs). Identifying when such systems emerge therefore requires rigorous quantification of the physics underlying causal emergence, and in particular “top-down” causation. We first discuss top-down causation as a conceptual framework for understanding living matter, and then turn to how quantifying causal emergence could provide new insights into how life first emerged.

In living systems, causal influences run both up and down organizational hierarchies [10] (see e.g., Figure 1). Characterizing living processes therefore requires both bottom-up causation—such as when a gene is read-out to make a protein that affects cellular behavior—and top-down causation—as occurs when changes in the environment initiate an organismal response that permeates all the way down to the level of individual genes [5]. Bottom-up causation is the status quo in modern physics, whereas top-down causation is less familiar and much more challenging to quantify. Top-down (or downward) causation occurs when a higher hierarchical level exerts causal influence over a lower level by setting a context (for example, by changing some physical constraints) within which the lower level actions take place [31,32]. A relatively simple example of top-down causation is the three-dimensional folding of an RNA molecule. The global constraint imposed by the network of structural interactions between residues causes the molecule to acquire a specific three-dimensional conformation [13]. RNA (or protein) folding therefore presents an example of whole-part causation, albeit localized to a single contiguous molecule.

Figure 1. Information flows both up and down the hierarchy of structure and causation in biological systems. Figure adapted from Noble [33].



Ellis has identified at least five different classes of top-down causation [32]: algorithmic top-down causation, top-down causation by non-adaptive information control, top-down causation via adaptive selection, top-down causation via adaptive information control and intelligent top-down causation. Of these, top-down causation by information control and top-down causation via adaptive selection are essential to life [13]. In top-down causation by information control, a higher hierarchical level influences lower level entities to achieve a specific functional outcome (goal) through the use of feedback control loops. Feedback control depends on the flow of information through the system, where information on the system's current state is evaluated relative to a particular functional outcome. Top-down causation by information control thereby presents a framework for understanding how information—an abstraction—can play a causally efficacious role in biological systems. It is important to note that while information is abstract, in the sense that it involves one entity symbolically representing another, it is nonetheless physical and only exists when physically instantiated (it therefore holds similar ontological status to energy). There is thus no conflict with information playing a causal role in the dynamics of living systems.

Top-down causation is an important mechanism for adaptive evolution through natural selection, where the higher-level “goal” is survival [32,34]. The role of top-down causation in adaptive selection is particularly evident for cases of convergent evolution. A striking example is provided by the evolution of echolocation in dolphins and bats, where over 200 genes have independently changed in the same ways to confer both species with the ability to use sonar [35]. A common high-level selection pressure (e.g., for the ability to navigate) lead independently to the same specific mutations in the DNA of both species. Convergent evolution thus provides a clear example of top-down causation via adaptive selection, where causal influences run from macroscopic environmental context to microscopic biochemical structure.

Top-down causation operates through functional equivalence classes. Functional equivalence occurs when a given “higher-level” state leads to the same high-level outcome, independent of which “lower-level” state instantiates it [32]. Equivalence classes are defined in terms of their function, not their particular physical instantiation: operations are considered (functionally) equivalent (*i.e.*, in the same equivalence class) if they produce the same outcome for different lower-level mechanisms. Functional

equivalence classes therefore represent the physical manifestations of virtual constructors. Functional equivalence is evident, for example, in the case of convergent evolution presented above, where convergence occurs because natural selection optimizes a functionally equivalent outcome (in this case, echolocation). Jaeger and Calkins present a more compelling example of functional equivalence [14], citing an experiment where the RNA based RNase P in *Escherichia coli* bacteria is replaced by a purely proteinaceous RNase P derived from plant organelles, which is completely functional in the host *E.coli* cell [36]. In this example, the RNA and protein versions of RNase P are structurally unrelated, but they clearly fulfill the same functional outcome and are thus members of the same functional equivalence class. This kind of functional equivalency between RNA and protein likely played a prominent role in the evolution of the RNA-peptide world. Further studies of this kind will be important to understanding transitions in chemical hardware that maintain high-level function that would have been particularly important throughout the early evolution of life (discussed in Section 5).

Life may be defined as a self-reproducing system that functions via top-down causation by adaptive selection *and* information control [13,14,16]. Here information control is the critical factor distinguishing living systems. Using the notion of functional equivalence classes, Jaeger and Calkins define bacterial cell as the congruence of two master functional classes: replication (of the genome) and reproduction—or construction (of the cell machinery). Thus, their descriptive framework shares much in common with the logical structure of a von Neumann automata: the relevant functional equivalence classes are abstract (virtual) constructors of the state of the system. Here the “higher-level” is the information, or program, encoded in the two master functions. Top-down causation by information control operates via the flow of information among and between these two master functions, which acts to conserve their functions (though not necessarily the precise chemical substrate they are instantiated in). Thus, *a bacterial cell may be defined by its architecture of information flow*. This insight provides a new framework for quantifying the emergence of life that captures the significance of abstract constructors in defining the living state.

4. The Emergence of Life as a Transition in Informational and Causal Architecture

Living and non-living systems may be distinguished by how information and causation are distributed. For living systems, information—in the form of programs, symbolic logic, and codes—plays an active, and therefore causally efficacious, role. Based on this observation, we previously identified several hallmarks of living systems that could lead to new productive modes of inquiry into understanding how life first emerged [15,16]. These hallmarks are reproduced in Table 2. Many of these represent different ways of approaching the same key attribute(s) of life. For example, as discussed here, the logical structure of a universal constructor necessitates the presence of information as a causal agent, and in particular the presence of top-down causation, in that it requires “high-level” programs to direct the dynamics of its physical instantiation. Similarly, these concepts are intimately related to the notion of the co-evolution of laws and states (e.g., self-referential dynamics [24–26]), where the current state of a physical system encodes information that determines its future state(s) (thus highlighting the fact that counterfactuals can, and do, matter in setting the dynamics of certain classes of physical systems). As such, we see Table 2 as a framework for identifying what it is that distinguishes life as a unique state

of matter, by approaching the problem from multiple directions that each hold promise to eventually converge on identifying the most fundamental attribute(s) of life.

Table 2. The hallmarks of life. Table from Walker and Davies [16].

Hallmarks of Life
Global organization
Information as a causal agency
Top-down causation
Analog and digital information processing
Laws and states co-evolve
Logical structure of a universal constructor
Dual hardware and software roles of genetic material
Non-trivial replication
Physical separation of instructions (algorithms) from the mechanism that implements them

In particular, the unique role of top-down causation in biological organization potentially enables a methodology for identifying a non-trivial distinction between life and non-life, identifiable as a fundamental difference in causal and informational architecture [16]. Rigorously quantifying top-down causal information flow and causal emergence may therefore provide plausible candidate parameter(s) for tracking the transition from non-living to living matter. Furthermore, fundamental differences in the causal architecture of information flow may provide insights not only into differentiating non-life from life, but also in identifying intermediate states of the transitory regime that could include a gradation of states of “almost life”. This framework thus holds promise for identifying transitory phases in the evolution of chemical systems on the pathway to modern cells that have previously proved elusive to characterize.

We recently proposed, via a toy model, one such framework for identifying transitions in information flow, using Schreiber’s “transfer entropy” (TE) [37] to explicitly study the transfer of information from local (individual nodes) to global (the mean-field) and from global to local degrees of freedom in a lattice of coupled logistic maps [15]. TE was chosen for this study as it provides insight into the directionality of information flow that more traditional measures commonly used in the biological sciences such as mutual information (*i.e.*, static, symmetric measures) cannot capture. Non-trivial collective behavior was observed to emerge each time the dominant direction of TE transitioned from local-to-global (bottom-up) to global-to-local (top-down), indicating that the emergence of collective organization is correlated with a dominance of information transfer from global to local scales. Thus, collective behavior may be characterized by “top-down” information transfer for this toy model system. Similar dynamics have been observed applying TE to other coupled non-linear systems, including cases where individual nodes do not have direct access to global state information [38,39].

Implementing measures of information transfer such as TE, combined with the appropriate aggregate variables describing informational degrees of freedom, may provide a quantitative framework that accurately captures the causal and informational architecture of biological networks. However, TE is

a measure of predictability and thus may capture the notion of emergent computation but not true causal emergence. Identifying causal influences requires an interventionist approach [40], which is not captured by TE [41]. These types of analyses therefore need to be extended to investigate the *causal* flow of information. A measure of “information flow” (IF) that accurately captures causal effects has been proposed by Ay and Polani [42], where causality is detected by using a measure similar to TE that replaces observational conditioning on the probability distributions (as used in TE) with interventional conditioning [40]. However, thus far, IF has been applied to identifying local (micro)-level causation only and not to understanding global causal architecture.

A different measure of causal information that does capture global architecture is provided under the framework of integrated information theory (IIT). “Integrated information” (ϕ) was originally devised by Tononi as a possible measure of consciousness [43]. In addition to quantifying conscious experience, it may also be a contender for defining life. ϕ quantifies how much information is generated by a physical system when it enters a particular state through the causal interactions among its elements, above and beyond the information generated independently by its parts [44,45]. More specifically, ϕ is the effective information (*ei*) of the minimum partition determined by the system’s “wholeness” [44,46]. *ei* characterizes the information generated by the causal structure of a network, when it enters a particular state. The effective information for each physically realized state is the relative entropy of the *a posteriori repertoire* with respect to the *a priori repertoire* of the network. The *a priori repertoire* is calculated by intervening on the system and perturbing it into all states with equal likelihood; in other words, it is the maximum entropy distribution. The *a posteriori repertoire* is defined as the repertoire of possible states that could have led to the realized state through the causal mechanisms of the system. As such, *ei* measures how much the causal mechanisms of the system reduce the uncertainty about the possible causes for a given state. IIT thus potentially provides a framework for understanding how constructors, which define the causal mechanisms of physical systems, can generate physically manifest information by eliminating counterfactuals. *ei* can be calculated at any organizational level (e.g., by coarse-graining) and has been implemented by Hoel *et al.* as a measure of causal emergence by comparing the *ei* at microscopic and macroscopic levels of organization [47]. They showed that in simple networks with fixed microscopic causal mechanisms *ei* peaks at a macroscopic level in space and/or time rather than at a microscopic level for certain causal architectures. These measures hold promise for identifying the appropriate formalism for characterizing the causal structure of living systems. However, much work remains to be done.

5. Top-down Causation in Chemical Evolution

The prominent role of top-down causation in extant life suggests that many of the transitions in the complexity of chemical reaction networks leading to life’s emergence may be characterizable as transitions in informational and causal architecture. The hope is that these may eventually be identifiable using the results of approaches as outlined in Section 4. In addition to the potential to measure states of life, non-life and almost-life, there are many other insights to be gained by considering explicitly the causal role of information in studies of the emergence of life. Here we briefly discuss a few of

the insights that derive from this perspective, which may have the potential to shed new light into the processes driving prebiotic evolution and the origin of life.

Top-down causation through functional equivalency permits alternative chemistries for life to be viable functional alternatives to natural biochemical systems. These alternative biochemistries represent different physical instantiations that have not been explored by extant biology but can be studied through synthetic biology. This includes alternative nucleic acids (other than DNA or RNA), such as expanded genetic alphabets (*i.e.*, more than two base pairs) or nucleic acids with modified backbone sugars. A particularly striking example is the recent engineering of a bacterium to host an additional unnatural DNA base pair synthetically incorporated in its genome [48]. This is the first successful demonstration of an organism stably incorporating an expanded genetic alphabet in its genome over successive generations. This represents an example of an alternative instantiation of the replication master function [14,17]: both the genomes of natural life and of this synthetically engineered life are members of the same functional equivalence class for replication within the cellular operating system.

The capacity for synthetic biologists to engineer functionally equivalent chemistries for life opens the possibility to explore the alternative chemical hardware that early life may have utilized. Typically the functional equivalence class most frequently studied by prebiotic chemists is replication of genetic material (though this effort is not typically labeled as such). The transition from RNA-based to DNA-based life forms represents a “genetic takeover” that required maintaining higher level functions—such as replication of genetic information—through functional equivalency. The notion of genetic takeover was first popularized by Graham Cairns-Smith to describe how life could have started with genetic information encoded in clays, which eventually transitioned to being stored in more familiar organic genetic polymers through information transfer between the inorganic and organic substrates [49]. More generally, genetic takeover can apply to the transfer of information between any two genetic systems, including between different nucleic acid species [30]. Examples of possible alternative nucleic acids include TNA (α -L-threofuranosyl nucleic acid) and GNA (glycol nucleic acid), each of which can exchange genetic information with RNA [50]. This suggests that either TNA or GNA could have preceded RNA in the evolution of the genetic system of terrestrial life. TNA, GNA and RNA—although structurally different molecules—are therefore members of the same functional equivalence class. Additionally, six synthetic genetic polymers (including TNA) have been shown to be capable of information storage, retrieval (through transcription onto DNA) and functional folding [51], thus demonstrating the existence of a functional equivalence class of nucleic acids suitable for the replication master function in modern cells. Here, the members of the equivalence class do not need to have the capacity to exchange information with each other directly: information, in this case, is preserved and exchanged only through higher-level function.

The earliest stages of evolution may have exploited functional equivalency in a different manner than modern organisms do. Prebiotic synthesis is notorious for generating nonspecific products. An example is provided by the prebiotic synthesis of RNA, which typically yields 2'-5' backbone linkages, whereas in modern life RNA polymers have exclusively 3'-5' linkages. Early RNA would have therefore represented a mixed population of polymers with non-heritable information in the sequence of backbone linkages. RNA molecules with random compositions of mixed backbones of 2'-5' and 3'-5' linkages have been shown to fold and retain the capacity for molecular recognition [52]. This suggests that in

a prebiotic environment, low yield of nonspecific synthesis products may have provided sufficient fuel for evolution, with selection operating on classes of functional equivalent molecules, rather than on specific target structures. In this “top-down” view of prebiotic evolution, functional equivalence classes would have emerged first, as a result of selection on function, followed by refinement of the particular molecular instantiation as specificity evolved and the relevant functional equivalence class(es) became fixed in the population. This model for early evolution is advantageous in that it allows nascent life to take advantage of the mess of prebiotic chemistry, rather than viewing the nonspecific products of prebiotic synthesis solely as a hindrance to evolutionary progress [53]. It is also amenable to laboratory study, for example by performing *in vitro* selection experiments on nucleic acids of mixed composition and determining what compositions might enable robust selection for function over a large class of molecular compositions.

Selection for functional equivalence would have also operated at the level of entire reaction networks in addition to classes of molecular compounds within a network. The notion of functional equivalence is closely related to that of biological modularity, where a module is a separable functional unit within a biological network. Modules may be replaced by other members of their functional equivalence class without affecting the global network structure. Modularity has been demonstrated to emerge spontaneously in the formation of autocatalytic sets (collections of entities where production of each member is catalyzed by another member of the set) [54]. This may have been an important step in the evolution of the decoupling of information storage and processing: modules would permit evolutionary refinement of particular instantiations of functional equivalence classes in the top-down manner described above, thus also permitting storage of information about particular functional equivalence classes, rather than their precise chemical instantiation.

While evolution has been shaped by functional equivalency, it does not necessarily mean that one could construct early life by replacing the parts of modern cells with their ancient counterparts. For example, one might wonder if riboorganisms could be resurrected in the laboratory by starting with replacing all functions mediated by protein or DNA in modern cells (see e.g., [14]). In practice this would lead to several complications, the most significant of which is the potential to mix stored information with its execution. Riboorganisms populated a world prior to the origin of translation. Riboorganisms therefore represent a different functional equivalence class than modern life (constructor and replicator functions cannot be easily identified for riboorganisms, as they can for modern cells) and could therefore potentially represent a distinct instantiation of life from that of any known extant life.

6. Conclusions

In this review, we have identified life as a unique state of matter distinguishable from other physical states by its causal architecture. Under this view, the transition from non-living to living matter roughly maps to the transition from trivial to non-trivial replication and should therefore correspond to decoupling of “software” (information controlling the dynamics of the chemical system) and “hardware” (the chemical substrate). While we typically associate this decoupling with the origin of translation, more broadly it should also encompass those transitions in the complexity of chemical systems that would have separated information storage from processing, and enabled the emergence of new layers of control

architectures (feedback from state to dynamics or top-down causation from “software” to “hardware”). Transitions in causal architecture could have occurred through a series of “information takeover” events, which differ from the notion of “genetic takeover” in that they are not necessarily confined to the transfer of genetic information from one molecular species to another. Instead, information takeover refers to transitions in chemical complexity whereby advances in information processing and computation would have conferred nascent life with the logical architecture necessary to gain control over the chemical substrates in which it was instantiated. These transitions will likely be difficult to identify, however one key feature should be that the (informational) states of a nascent living systems played an increasingly prominent role in determining the future dynamics. This suggests that the processes leading to the emergence of life should be marked by a series of “information take-over” events corresponding to changes in causal architecture and information flow, which, in principle, is measurable.

An interesting facet of this perspective is that it characterizes life as logically and organizationally distinct from other kinds of dynamical systems, and thus life represents a novel, emergent state of matter. Our usual causal narrative, consisting of the bottom-up action of material entities only, could therefore be only a subset of a broader class of phenomena—including life—that admit immaterial causes through the action of virtual constructors. This viewpoint suggests new thinking as to how life might have arisen on lifeless planet, by shifting emphasis to the origins of computation, control and informational architecture, rather than focusing solely on the onset of Darwinian evolution for example, which does not rigorously define how or when life emerges in a physical system. This framework also permits a more universal view of life, where the same underlying principles would permit understanding of living systems instantiated in different substrates (either artificial or in alternative chemistries) anywhere in the universe.

Acknowledgments

The author thanks Paul C.W. Davies, George F.R. Ellis, Chiara Marletto and the emergence@ASU group for fruitful conversations. This publication was made possible through the support of a grant from the Templeton World Charity Foundation. The opinions expressed in this publication are those of the author and do not necessarily reflect the views of Templeton World Charity Foundation.

Conflicts of Interest

The author declares no conflict of interest.

References

1. Deutsch, D. Constructor theory. **2012**, arXiv:1210.7439.
2. Farnsworth, K.; Nelson, J.; Gershenson, C. Living is information processing: From molecules to global systems. *Acta Biotheor.* **2013**, *61*, 203–222.
3. Ruiz-Mirazo, K.; Umerez, J.; Moreno, A. Enabling conditions for open-ended evolution. *Biol. Philos.* **2008**, *10*, 67–85.

4. Walker, S.I.; Callahan, B.; Arya, G.; Barry, J.D.; Bhattacharya, T.; Grigoryev, S.; Pellegrini, M.; Rippe, K.; Rosenberg, S.M. Evolutionary dynamics and information hierarchies in biological systems. *Ann. N.Y. Acad. Sci.* **2013**, *1305*, 1–17.
5. Noble, D. Genes and causation. *Philos. Trans. R. Soc. A* **2008**, *366*, 3001–3015.
6. Bray, D. Protein molecules as computational elements in living cells. *Nature* **1995**, *376*, 307–312.
7. Ptashne, M. Principles of a switch. *Nature* **2011**, *7*, 484–487.
8. Danchin, A.; Sekowska, A. The logic of metabolism and its fuzzy consequences. *Environ. Microbiol.* **2014**, *16*, 19–28.
9. Braakman, R.; Smith, E. The compositional and evolutionary logic of metabolism. *Phys. Biol.* **2013**, *10*, 011001.
10. Campbell, D. Downward causation in hierarchically organised biological systems. In *Studies in the philosophy of biology: Reduction and related problems*; Ayala, F.J., Dobzhansky, T., Eds.; Macmillan: London, UK, 1974; pp. 179–186.
11. Levin, M. Reprogramming cells and tissue patterning via bioelectrical pathways: Molecular mechanisms and biomedical opportunities. *Syst. Biol. Med.* **2013**, *5*, 657–676.
12. Davies, P.C.W. The epigenome and top-down causation. *Interface Focus* **2012**, *2*, 42–48.
13. Auletta, G.; Ellis, G.; Jaeger, L. Top-down causation by information control: From a philosophical problem to a scientific research programme. *J. R. Soc. Interface* **2008**, *5*, 1159–1172.
14. Jaeger, L.; Calkins, E. Downward causation by information control in micro-organisms. *Interface Focus* **2012**, *2*, doi:10.1098/rsfs.2011.0045 .
15. Walker, S.I.; Cisneros, L.; Davies, P.C.W. Evolutionary transitions and top-down causation. In *Proceedings of Artificial Life XIII*, Michigan State University, East Lansing, MI, USA, 19–22 July 2012; Volume 13, pp. 283–290.
16. Walker, S.I.; Davies, P.C.W. The algorithmic origins of life. *J. R. Soc. Interface* **2013**, *10*, 20120869.
17. Danchin, A. Bacteria as computers making computers. *FEMS Microbiol. Rev.* **2009**, *33*, 3–26.
18. Joyce, G.F. Bit by Bit: The Darwinian Basis of Life. *PLoS Biol.* **2012**, *10*, e1001323.
19. Gilbert, W. Origin of Life: The RNA world. *Nature* **1986**, *319*, 618.
20. Szostak, J. The eightfold path to non-enzymatic RNA replication. *J. Syst. Chem.* **2012**, *3*, 2–14.
21. Ivica, N.; Obermayer, B.; Campbell, G.; Rajamani, S.; Gerland, U.; Chen, I. The paradox of dual roles in the RNA world: Resolving the conflict between stable folding and templating ability. *J. Mol. Evol.* **2013**, *77*, 55–63.
22. Eigen, M. Self-organization of matter and evolution of biological macromolecules. *Naturwissenschaften* **1971**, *58*, 465–523.
23. Von Neumann, J. *Theory of self-reproducing automata*; Illinois University Press: Chicago, IL, USA, 1966.
24. Goldenfeld, N.; Woese, C. Life is physics: Evolution as a collective phenomenon far from equilibrium. *Ann. Rev. Cond. Matt. Phys.* **2011**, *2*, 375–399.
25. Hofstadter, D.R. *Gödel, Escher, Bach*; Basic Books: New York, NY, USA, 1979.

26. Pavlic, T. P.; Adams, A. M.; Davies, P.C.W.; Walker, S.I. Self-referencing cellular automata: A model of the evolution of information control in biological systems. In Proceedings of Artificial Life XIV, New York City, NY, USA, 30 July–2 August 2014 ; Volume 14, pp. 522 - 529.
27. Szathmáry, E.; Maynard Smith, J. From replicators to reproducers: The first major transitions leading to life. *J. Theor. Biol.* **1997**, *185*, 555–571.
28. Deutsch, D.; Marletto, C. Reconstructing physics: The universe is information. *New Sci.* **2014**, *Issue 2970*.
29. Kauffman, S. *At Home in the Universe*; Oxford University Press: Oxford, UK, 1996.
30. Engelhart, A.E.; Hud, N.V. Primitive Genetic Polymers. *Cold Spring Harb. Perspect. Biol.* **2010**, *2*, doi:10.1101/cshperspect.a002196 .
31. Ellis, G.F.R. On the nature of emergent reality. In *The Re-emergence of Emergence*; Clayton, P., Davies, P.C.W., Eds.; Oxford University Press: Oxford, UK, 2006; pp. 79–107.
32. Ellis, G.F.R. Top-down causation and emergence: Some comments on mechanisms. *Interface Focus* **2011**, *2*, 126–140.
33. Noble, D. Claude Bernard, the first systems biologist, and the future of physiology. *Exp. Physiol.* **2012**, *93*, 1–16.
34. Okasha, S. Emergence, hierarchy and top-down causation in evolutionary biology. *Interface Focus* **2012**, *2*, 49–54.
35. Parker, J.; Tsagkogeorga, G.; Cotton, J.A.; Liu, Y.; Provero, P.; Stupka, E.; Rossiter, S.J. Genome-wide signatures of convergent evolution in echolocating mammals. *Nature* **2013**, *502*, 228–231.
36. Gobert, A.; Gutmann, B.; Taschner, A.; Gößringer, M.; Holzmann, J.; Hartmann, R.K.; Rossmanith, W.; Giegé, P. A single Arabidopsis organellar protein has RNase P activity. *Nat. Struct. Mol. Biol.* **2010**, *17*, 740–744.
37. Schreiber, T. Measuring information transfer. *Phys. Rev. Lett.* **2000**, *85*, 461–464.
38. Cisneros, L.; Jiménez, J.; Cosenza, M.G.; Parravano, A. Information transfer and nontrivial collective behavior in chaotic coupled map networks. *Phys. Rev. E* **2002**, *65*, 045204.
39. Paredes, G.; Alvarez-Llamoza, O.; Cosenza, M.G. Global interactions, information flow, and chaos synchronization. *Phys. Rev. E* **2013**, *88*, 042920.
40. Pearl, J. *Causality: Models, Reasoning and Inference*; Cambridge University Press: New York, NY, USA, 2000.
41. Lizier, J.T.; Prokopenko, M. Differentiating information transfer and causal effect. *Eur. Phys. J. B* **2010**, *73*, 605–615.
42. Ay, N.; Polani, D. Information flows in causal networks. *Adv. Complex Syst.* **2008**, *11*, 17–41.
43. Tononi, G. Consciousness as Integrated Information: A Provisional Manifesto. *Biol. Bull.* **2008**, *215*, 216–242.
44. Tononi, G.; Sporns, O. Measuring information integration. *BMC Neurosci.* **2003**, *4*, doi:10.1186/1471-2202-4-31.
45. Balduzzi, D.; Tononi, G. Integrated information in discrete dynamical systems: Motivation and theoretical framework. *PLoS Comput. Biol.* **2008**, *4*, e1000091.
46. Tononi, G. Information measures for conscious experience. *Arch. Ital. Biol.* **2001**, *139*, 367–371.

47. Hoel, E.P.; Albantakis, L.; Tononi, G. Quantifying causal emergence shows that macro can beat micro. *Proc. Natl. Acad. Sci. USA* **2013**, doi:10.1073/pnas.1314922110 .
48. Malyshev, D.A.; Dhimi, K.; Lavergne, T.; Chen, T.; Dai, N.; Foster, J.M.; Corrêa, I.R.; Romesberg, F.E. A semi-synthetic organism with an expanded genetic alphabet. *Nature* **2013**, *509*, 385–388.
49. Cairns-Smith, A.G. *Seven Clues to the Origin of Life: A Scientific Detective Story*; Cambridge University Press: Cambridge, UK, 1985.
50. Yang, Y.W.; Zhang, S.; McCullum, E.O.; Chaput, J.C. Experimental Evidence That GNA and TNA Were Not Sequential Polymers in the Prebiotic Evolution of RNA. *J. Mol. Evol.* **2007**, *65*, 289–295.
51. Pinheiro, V.B.; Taylor, A.I.; Cozens, C.; Abramov, M.; Renders, M.; Zhang, S.; Chaput, J.C.; Wengel, J.; Peak-Chew, S.Y.; McLaughlin, S.H.; Herdewijn, P.; Holliger, P. Synthetic Genetic Polymers Capable of Heredity and Evolution. *Science* **2012**, *336*, 341–344.
52. Englehart, A.; Powner, M.; Szostak, J. Functional RNAs exhibit tolerance for non-heritable 2'-5' versus 3'-5' backbone heterogeneity. *Nat. Chem.* **2013**, *5*, 390–394.
53. Walker, S.I.; Grover, M.A.; Hud, N.V. Universal Sequence Replication, Reversible Polymerization and Early Functional Biopolymers: A Model for the Initiation of Prebiotic Sequence Evolution. *PLoS One* **2012**, *7*, e34166.
54. Hordijk, W.; Steel, M.; Kauffman, S. The Structure of Autocatalytic Sets: Evolvability, Enablement, and Emergence. *Acta Biotheor.* **2012**, *60*, 379–392.

© 2014 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).