Orig Life Evol Biosph (2013) 43:491–500
DOI 10.1007/s11084-013-9352-3

ASTROBIOLOGY

A New Criterion for Demarcating Life from Non-Life

J. H. van Hateren

Received: 1 November 2013 / Accepted: 13 December 2013 /
Published online: 9 January 2014

Abstract Criteria for demarcating life from non-life are important for deciding whether new candidate systems, either discovered extraterrestrially or constructed in the laboratory, are genuinely alive or not. They are also important for understanding the origin of life and its evolution. Current criteria are either too restrictive or too extensive. The new criterion proposed here poses that a system is living when it is capable of utilizing active causation, at evolutionary or behavioural timescales. Active causation is produced when the organism uses an estimate of its own Darwinian fitness to modulate the variance of stochasticity that drives hereditary or behavioural changes. The changes are subsequently fed back to the fitness estimate and used in the next cycle of a feedback loop. The ability to use a self-estimated fitness in this way is an evolved property of the organism, and the way in which fitness is estimated is therefore controlled and stabilized by Darwinian evolution. The hereditary and behavioural trajectories resulting from this mechanism combine predictability with unpredictability, and the mechanism produces a form of self-directed agency in living organisms that is absent from non-living systems.

Keywords Active causation · Agency · Darwinian evolution · Definition of life · Spontaneity

Introduction

General criteria for demarcating life from non-life are important for several reasons (Cleland and Chyba 2002). First, they may help to recognize life if it is discovered on other planets than Earth and if its composition and function differ strongly from those of life on Earth. Second, they may help efforts to produce artificial life in the laboratory (Stano and Luisi 2010), and may be used to evaluate to what extent those efforts are successful. And third, they may assist studies that aim to understand the origin and evolution of life on Earth.

There is currently no consensus on what would constitute sufficient and necessary criteria for establishing that a system lives (Bedau 2007; Tsokolov 2009; Benner 2010). It is clear that various properties are important, specifically material and physical requirements, requirements with respect to heredity and information, and requirements with respect to system integrity and autonomy. An example of a physical requirement is that some form of metabolism is needed

J. H. van Hateren (✉)
Johann Bernouilli Institute for Mathematics and Computer Science, University of Groningen, P.O. Box 407, 9700 AK Groningen, The Netherlands
e-mail: j.h.van.hateren@rug.nl

such that free energy can be harnessed from the environment to build and sustain life's structures and processes (Lineweaver and Egan 2008). Hereditary requirements are, first, that some form of structural memory (like RNA or DNA) is present such that replication and reproduction is possible (Pross 2004), and, second, that heredity can change such that life-forms can adapt to changing circumstances (Darwin 1859). Heredity and physiological structure are closely related to information, which suggests that the particular ways by which living systems accumulate and use information can be used to define life (Walker and Davies 2012; Michel 2013). Living systems are furthermore characterized by integrity and the ability of self-maintenance and autonomous functioning (Varela et al. 1974; Ruiz-Mirazo et al. 2004; Di Paolo 2005). Kauffman (2000, 2003) has argued that the agency of autonomous systems, their ability to act on their own behalf in an environment, appears to be the defining characteristic of life (see further the Discussion).

Not all criteria mentioned above have a clear proposed implementation, and no single criterion appears to be sufficient or even necessary, as there are exceptions to most of them. In this article I propose a criterion for demarcating life from non-life that has, to my knowledge, not been discussed before in this form. It proposes that the transition from non-life to life is accompanied by a transition of causality, from the passive forms of causation of non-living physicochemical systems to a form of causation that is active. It therefore conforms to the views of Kauffman (2000, 2003) and Di Paolo (2005) that agency is a defining characteristic of life. However, it reaches this conclusion not by taking the autonomy of organisms as a starting point, but by starting from Darwinian evolution by natural selection, and showing that a slight extension of Darwin's scheme leads to agency. The resulting criterion may be unique in its ability to demarcate all currently known life-forms from all currently known non-living systems.

Theory

In this section I will first introduce several basic forms of causation, subsequently explain how the shift from passive to active causation is made in hereditary systems undergoing Darwinian evolution, and finally extend that result to behavioural systems. The explanation is qualitative, for a more quantitative analysis see several of the computational models discussed in van Hateren (2013).

Deterministic, Stochastic, and Modulated Stochastic Causation

Figure 1 illustrates several elementary forms of causation. The diagrams show only idealized forms, in reality mixed forms can occur, and usually along more dimensions than shown. The traces in the figure, called "signal" below, symbolize the time course of a variable, state, or property of a system, and the arrows a causal relationship with earlier or later signals (roughly "caused by" and "causing"). Figure 1a shows the work-horse of science, deterministic causation. The signal operates like a cog in a clockwork, driven by and driving other signals. The time course of the signal is in principle fully predictable from the signal driving it, and it does not originate new chains of causation but merely continues the current chain. In practice, the signal of deterministic causation is often accompanied by some noise, but the result (noisy deterministic causation) is still primarily deterministic in nature, with the noise doing little or no causal work.

This is different for stochastic causation (Fig. 1b), which occurs spontaneously, and is the start of new causal chains. Only the individual fluctuations can be causally effective, because

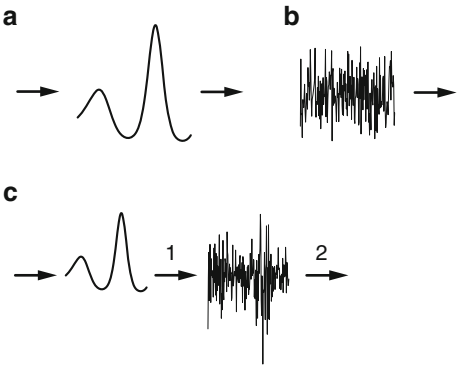

Fig. 1 Deterministic, stochastic, and modulated stochastic causation. Traces ("signals") symbolize the time course of system variables or properties, *arrows* symbolize causal relationships with earlier or later signals. In deterministic causation (**a**) the signal is in the middle of a chain of signals, caused by and causing the adjacent signals. Stochastic causation (**b**) forms the (random) start of new chains of causation. Modulated stochastic causation (**c**) lets a non-negative deterministic signal (*left*) modulate the variance of a stochastic signal (*right*). The *arrows marked 1 and 2* correspond to *arrows* in Fig. 2

the statistical properties of the signal (such as mean and variance) do no change over time, and therefore cannot contribute to causation. Physical origins of stochastic causation include thermal noise, quantum noise (e.g. radioactive decay), and fundamental uncertainty caused by nonlinear interactions that amplify disturbances, no matter how small, to significant effects (e.g. Laskar and Gastineau 2009).

Finally, Fig. 1c shows modulated stochastic causation, a crucial component in the models explained below. A deterministic signal, scaled to be non-negative, drives the variance of a subsequent stochastic signal. This form of causation is rather special, because it is exactly midway between deterministic and stochastic causation. Its mean remains steady, and therefore cannot participate in causation. Causation is neither fully predictable, because of the stochastic fluctuations, nor fully unpredictable, because of the deterministic time course of the variance. The arrows marked by 1 and 2 correspond to similarly marked arrows in Fig. 2 to be explained next.

Passive and Active Causation in Hereditary Systems

Figure 2a shows the basic Darwinian evolution scheme. The argument here stays close to the original one of Darwin (1859), and, for simplicity, neither replication and reproduction is distinguished, nor genotype and phenotype. A population of organisms consists of different types, which all reproduce in a common environment that is not fixed but varies unpredictably over time. Depending on the population size and the environment, each type has a fitness, which is defined here as the type's expected number of offspring (on average per lifetime). Because not all types have equal fitness in a particular environment (natural selection), the composition of the population gradually changes because of differential reproduction. The composition would eventually reach a steady state, were it not for two variable external inputs to the reproductive loop *R*. Apart from the time-varying environment already mentioned, there is a stochastic input that produces variations in heredity (mutations). New mutations produce new types in the population, and some of those types may have good fitness in a changed environment. The speed of environmental change determines what mutation rate is optimal, for the population as well as for the probability of long-term survival of individual lineages. Too

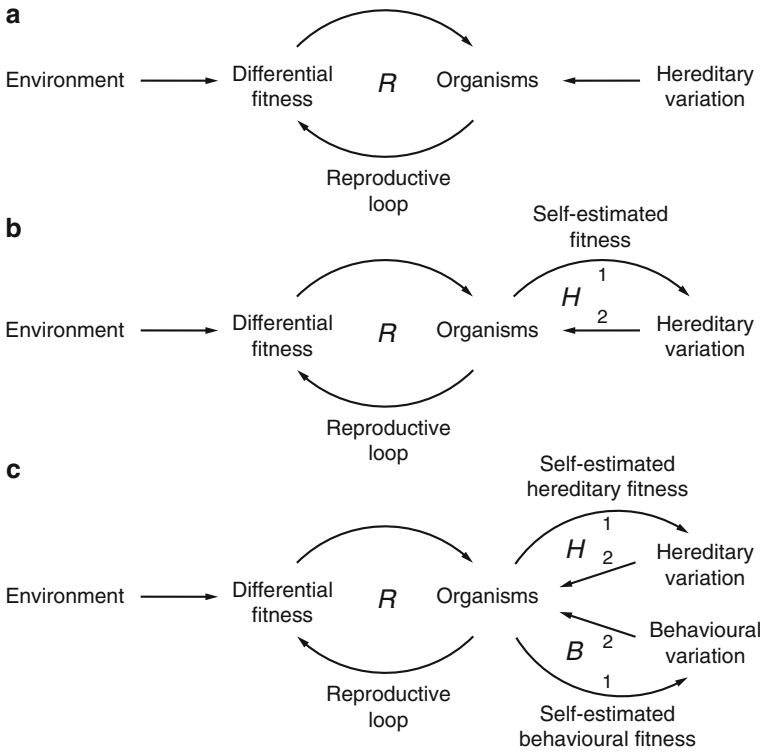

Fig. 2 Basic and extended schemes of Darwinian evolution. In the basic scheme (**a**) a population of organisms reproduces and changes because each type within the population has its own fitness (expected number of offspring per lifetime). Fitness depends on population size, a time-varying environment, and hereditary variation (mutations). All causation is passive (either internally deterministic or externally driven) from the point of view of the organism. In the extended Darwinian scheme (**b**) each organism implicitly estimates its own fitness, and uses that to modulate the mutation rate. The result is that hereditary trajectories (lineages) become partly driven by the self-estimated fitness, and therefore that the organism's lineage is partly actively caused. Active causation can be extended to behavioural timescales as well (**c**), leading to agency in the organism, where randomly generated behaviour that is estimated to imply low fitness is more quickly modified than high-fitness behaviour. *Arrows marked 1 and 2* correspond to those in Fig. 1c; *R* reproductive loop; *H, B* hereditary (behavioural) loop producing actively caused hereditary (behavioural) trajectories

low mutation rates make it impossible to follow fast changes in the environment, leading to extinction of the population. Too high mutation rates produce so many nonviable mutations that extinction is likely as well, unless it is compensated by a particularly high fitness of the viable mutations. If the speed by which environments vary can change unpredictably, there is no single mutation rate that works best.

Recently it has become clear that the mutation rate of a biological cell is indeed not a constant, but is under control of the cell itself, through a wide range of mechanisms that apparently evolved on multiple occasions (Galhardo et al. 2007; Shee et al. 2011; Al Mamun et al. 2012; MacLean et al. 2013). When a cell encounters adverse conditions, for example when toxins or unfavourable temperatures are present, the mutation rate is increased as part of the cellular stress response. The primary interpretation of this control is that the disadvantage of a higher mutation rate (many nonviable offspring) is overruled by the advantage of an increased probability of a mutation that helps some of the offspring to survive the stress.

Because high fitness can, under favourable conditions, lead to exponential growth in numbers, this is presumably an evolvable strategy (Ram and Hadany 2012). In particular, it is expected to be evolvable if the mutation rate is controlled by an estimate of the fitness (Fig. 2b, loop *H*; van Hateren 2013). This control must be such that high estimated fitness keeps the mutation rate low, while low estimated fitness produces a higher mutation rate.

The fitness estimate is made, implicitly, by the cell itself, as part of the control loops regulating its own mutation rate. The cell cannot know the true fitness, because that would require a highly complicated simulation of itself and of external factors that are effectively unpredictable. But it can use rules of thumb to make a fair estimate of the fitness. For example, toxins will interfere with the normal functioning of the cell, and this can be readily detected and used as an implicit indicator of low expected fitness. Many different indicators may need to be factored in and combined in order to obtain a good estimate. The better the cell estimates its own fitness, the more accurately it can modulate its mutation rate according to the true fitness. The cell's ability to self-estimate fitness is an evolved property under control of Darwinian evolution. It is therefore under evolutionary selection pressure to estimate the true fitness as faithfully as possible within the constraints of the means available to the organism.

Interestingly, the schemes of Fig. 2a and b have quite different causal signatures. The basic Darwinian scheme of Fig. 2a consists of a core comprising the reproductive loop *R* and how it is affected by fitness changes driven by population size and environment. This core process consists of components determined by deterministic or noisy deterministic causation. It unfolds mechanistically like any other feedback system one can encounter in the physical world. The environment is an external driver of the system. Even though mutations happen inside cells, the cause (primarily thermal noise) is ultimately environmental in nature, because the cells are not isolated systems. The randomness producing mutations is therefore also an external driver of the system. The organism is thus driven by internal deterministic processes and two external drivers. The internal deterministic processes just unfold, happen, like in a clockwork or machine, and the organism has no opt-out, no other choice than undergoing them. From the point of view of the organism, they are passive forms of causation. The same applies to the two external causes, the organism is driven by them, undergoing them in a passive way. The organism is thus a passive player in the scheme of Fig. 2a, it is pushed around, and exposed to forms of causation that are merely passive as seen from the organism.

The situation is quite different for the extended Darwinian scheme of Fig. 2b. In this scheme the organism produces, implicitly, a fitness estimate and uses that to control hereditary variability (loop *H*). Note, however, that this control is indirect and without foresight, because the fitness estimate does not drive mutations in a particular direction. It is blind to whether particular mutations might decrease or increase actual fitness, i.e., it is undirected. Rather than driving mutations directly, the fitness estimate drives the mutation rate, which is equivalent to driving the hereditary variability. It therefore conforms to the modulated stochastic causation of Fig. 1c (see the corresponding arrows 1 and 2). Although undirected, this way of driving heredity does in fact influence the probability of success of the lineage to which the organism belongs. If actual fitness is high, the estimated fitness is expected to be high as well, and the resulting low mutation rate keeps the organismal lineage close to the heredity of the successful type. On the other hand, if actual fitness is low, the estimated fitness is expected to be low as well, and the resulting high mutation rate produces new types, some of which might have higher fitness. Most lineages will die, but the ones with higher fitness can multiply quickly because the reproductive loop *R* has the potential for exponential growth.

Indirectly, the fitness estimate thus promotes high fitness and survival. It does so by a peculiar form of causality, that is, to my knowledge, unique to living systems. This form of causality consists of harnessing and modulating randomness (the stochastic processes

producing mutations) by a criterion (estimated fitness) that is controlled by a feedback loop (*H*) that contains the organisms and their hereditary structure that resulted from previous cycles of the same control loop. Although each cycle of the control loop is the result of random mutations, a trajectory of subsequent mutations (i.e., consisting of consecutive organisms belonging to a particular lineage) is not random at all. It is continually driven by a combination of actual fitness (*R* loop) and the criterion of fitness as estimated by each organism of the lineage (*H* loop). As a result, a hereditary lineage of organisms is neither completely predictable (because of the randomness of mutations) nor completely unpredictable (partly because of the self-estimated fitness produced by each organism). Because the randomness of each step is utilized as a source of spontaneity by each organism, this form of causality is properly called active. More formally, active causation is then defined here as the causation that arises from a feedback loop *H* (Fig. 2b) that uses a fitness estimate (itself under selection pressure to conform to the actual fitness driving the *R* loop) as the driver of the modulated stochastic causation that produces new heredity.

Passive and Active Causation in Behavioural Systems

The above discussion concerns organismal control of hereditary variation, for which there is ample empirical evidence. More speculative is an extension of this scheme to behaviour, here taken as any change within a particular organism, thus at timescales shorter than the lifetime of the organism. As before, behaviour is considered to be passively caused when it is purely due to internal (noisy) deterministic processes or to external drivers. Examples are behavioural responses that rely on mechanisms available through genetic memory (established by natural selection in the past of the organism's lineage) and through physiological or neural memory (established during the lifetime of the organism). Although the establishment of such mechanisms may involve active causation, once they are present they can unfold in a mechanistic, predictable way and therefore do not require active involvement of the organism.

In addition, behaviour may originate from active causation, with Fig. 2c showing an elementary implementation. The organism makes now also a behavioural fitness estimate (incorporated in loop *B*), which may partly correspond to the hereditary one but need not be identical. If the behavioural fitness estimate is low, the organism produces large (but random) changes in behaviour or behavioural disposition, whereas behaviours that have high estimated fitness are changed little and therefore tend to persist. The estimator evaluating behaviour is itself an evolved property of the organism, and is therefore driven by basic Darwinian evolution to conform to the true fitness (as operating in loop *R*) as closely as possible. The causal structure is again active: the organism harnesses randomness (again primarily from thermal noise of molecular processes in cells), and utilizes that as a source of spontaneity to drive behavioural trajectories. Such trajectories are again neither completely predictable nor completely unpredictable (van Hateren 2013).

An example of an active mechanism may be the run-and-tumble behaviour of *E. coli* (Macnab and Koshland 1972), where the frequency of changing direction is increased when low concentrations of chemical attractants are sensed. This results, on average, in moving closer to higher concentrations. In general, showing that behaviour is actively caused will require, first, determining the part of the behaviour that can be explained by passive causation (through predictable, deterministic processes), and, second, showing that the unexplained behavioural variability is high when the organism's estimated fitness is low, and low when this fitness is high. Although it is not known if active causation at the level of behaviour is generally present in all life-forms, it appears evolvable (van Hateren 2013) and seems plausible, and I assume here that it is indeed generally present.

Discussion

This article proposes that life distinguishes itself from non-life by its utilization of a peculiar, active form of causation that is not encountered in the non-living world. It enables living organisms to initiate new chains of causation in the following way. It starts with a random event driving a change in heredity or in behaviour, with the variance of the randomness controlled by the organism. Although any single change is undirected, it is part of a series of changes (a hereditary or behavioural trajectory) that is continually nudged by a feedback loop under control of the organism (through its self-estimated fitness that modulates the variance of subsequent changes). As a result, the trajectory is random in detail, but as a whole it is not random or undirected at all, but complies with the criteria used for the fitness estimate. Using estimated fitness is an evolved capability that is driven, on an evolutionary timescale, to comply with Darwinian fitness.

Active Causation as a Demarcation Criterion

Active causation is assumed to be not only a general property of the hereditary apparatus of living cells, but also of the molecular apparatus that drives changes within cells during their lifetime, and of the molecular or neural circuits that drive behaviour of multicellular organisms (van Hateren 2013). The presence or absence of the capacity for active causation can thus serve as a criterion for classifying a system as belonging to life or not, also systems that pose problems for other demarcation criteria. For example, mules would be classified as life, because they utilize active causation at a behavioural level, even if they do not reproduce and do not take part in continued Darwinian evolution. A similar conclusion holds for a living cell that has stopped reproducing, for example because it belongs to a multicellular organism. Entities like flames and growing crystals are classified as non-life, because they lack the capacity for active causation.

Dormant life-forms such as spores and dehydrated eggs still have the capacity for active causation, be it in the future, and are therefore classified as life, though not alive (living at the very moment). Viruses must presumably be classified as not belonging to life, not so much because they lack a metabolic system, but because they may not be able to use their host to modulate viral genetic variability as driven by an estimate of viral fitness (as opposed to an estimate of the host's fitness). Viruses therefore probably do not have the capacity for active causation, although this cannot be ruled out for all viruses and is ultimately an empirical question. A colony of social insects might or might not be classified as a living entity. That would depend on whether it is possible to define a proper fitness for the colony as a whole, whether such fitness is estimated by the colony, and whether such an estimate drives either the colony's hereditary variability or its collective behavioural variability. In either case, hereditary or behavioural, the mechanism must be under control of Darwinian evolution, and therefore must increase fitness and have a hereditary component.

Active causation fits fairly well with naive intuitions as to what makes a system living. A primary phenomenological property of living systems is that such systems have agency, in the sense that they have some level of autonomy and act on their own behalf. They are unpredictable to some extent, and generally appear to be goal-directed and self-serving. Finally, living systems can physically die, indeed an essential requirement for the applicability of fitness and Darwinian evolution. The feedback scheme of active causation gives these naive intuitions a solid basis that is accessible to scientific analysis in terms of the underlying control loops and their physicochemical realizations.

Relationship with Approaches Emphasizing Autonomy, Replication, and Information

The current approach is related to several long-standing traditions that attempt to characterize the nature and origin of life, of which I will discuss here autonomy, replication, and information. Central to the present approach is the concept of fitness, defined here as the expected number of offspring per lifetime, extended to the concept of self-estimated fitness as modulating stochastic causation in a feedback loop. High fitness requires autonomy, faithful replication, and the acquisition of new information, as discussed below.

Firstly, autonomy is required by high fitness, because it provides the stability and time needed for effective reproduction. Autonomy in the sense of self-maintenance and homeostasis is central to the idea of autopoiesis ("self-production", Varela et al. 1974; see also Thompson 2007), and it was recently extended with adaptivity and agency (adaptive self-regulation, Di Paolo 2005). The concept of active causation (AC) as proposed here resembles, but is not identical to adaptive self-regulation (ASR). There are systems that have AC but no ASR (such as the *H* loop in Fig. 2b, where organisms modify their offspring, not themselves) and systems that have ASR but no AC (such as an adaptive extension of a conventional autopoietic system that is, by default, purely deterministic; such systems would lack genuine agency and would not be alive according to the criterion presented here). Nevertheless, it would be fairly straightforward to incorporate active causation and its fitness-modulated stochasticity into the framework of adaptive self-regulation, and thus provide it with what I view as genuine agency and goal-directedness. Kauffman (2000, 2003) defines an autonomous agent as a system that can act on its own behalf in an environment, but he can only explain this agency in a definitional sense (by invoking thermodynamic work cycles). Agency as explained here (exercising active causation, as in the *H* loop and especially the *B* loop of Fig. 2) solves this problem: it is a highly specific mechanism that directly explains the causal freedom of agents in terms of underlying physical processes (a feedback loop with stochasticity driven by self-estimated fitness driven by actual fitness).

Secondly, faithful replication is required by high fitness, because otherwise fitness-promoting properties previously acquired in evolution would quickly deteriorate (Eigen 1971; Szathmáry and Maynard Smith 1995). However, too faithful replication would hamper the rate of adaptation to a time-varying environment. It is proposed here that a controlled variation of the mutation rate, although presumably selected in basic Darwinian evolution for its survival value (Galhardo et al. 2007), has produced agency (in the form of active causation) as a spin-off (the *H* loops in Fig. 2).

Finally, it has long been recognized that information appears to play a crucial role in the origins and functioning of life, in particular when adapting to new conditions and thereby retaining or increasing fitness (Szathmáry and Maynard Smith 1995; Maynard Smith 2000; Nurse 2008; Walker and Davies 2012). However, information is a rather elusive concept (for an exhaustive overview of how differently it has been defined and used throughout science see Burgin 2010). There are two problematic aspects of information that need to be acknowledged when using it to explain life's properties. First, the type of information that has the most solid quantitative foundation, Shannon information, is value-free, i.e., it does not distinguish between meaningful (useful) and meaningless (useless) information. Second, a system that consists entirely of deterministic processes cannot generate new information. The second problem is resolved by letting stochastic processes produce new information, either with a fixed (Fig. 2a) or a modulated (Fig. 2b, c) mutation rate. The first problem is more difficult, but in principle resolved by the *H* loop of Fig. 2b, c. Because the *H* loop has an intrinsic goal-directedness (self-estimated fitness), it generates meaning (van Hateren 2013), with meaningful information effectively retained by the combined action of the *H* and *R* loops. A similar consideration applies to the *B* loop of Fig. 2c for information generated and stored during the

lifetime of an organism (where specific behaviours are not under control of the *R* loop, but hereditary behavioural strategies are). Note that this is different from the proposal of Maynard Smith (2000) of giving apparent agency to natural selection, which would, in my view, only lead to apparent, "as if" meaning.

Defining Life

Giving a definition of life may be a somewhat futile endeavour (Cleland 2012), primarily because the meaning of a single sentence is inevitably vague and open to different interpretations. I will nevertheless attempt to give one here, so that the present proposal can be readily identified in future discussions. It is a variant of NASA's working definition ("Life is a self-sustaining chemical system capable of undergoing Darwinian evolution"), reading: "Life is a material system capable of active causation." Active causation depends on estimating fitness and on the fact that modulating stochasticity can increase fitness itself. Because Darwinian evolution is the only mechanism currently known that will, in the long run, consistently promote high fitness, Darwinian evolution is presumably required for maintaining the long-term stability of active causation. The term "material system" is used to indicate that life is defined here as a phenomenon of the real world, and not of a purely symbolic system.

Taking active causation as the primary criterion for distinguishing life from non-life implies that any system that completely lacks active causation is classified as non-life. In particular, a self-replicating system that is subject to the basic evolution scheme of Fig. 2a, thus with a fixed mutation rate, is not considered to be life (in contrast to the NASA definition), unless it utilizes active causation at a shorter, behavioural timescale, for example in its metabolic pathways (such a system would correspond to Fig. 2c with a *B* loop but without an *H* loop). This does not pose a problem for defining current life, because fixed mutation rates do not seem to occur any more, and non-reproducing organisms have a *B* loop. However, the extended scheme of Fig. 2b has presumably evolved from the basic scheme, and the current proposal therefore implies that the transition from protolife to life happened somewhere between the schemes of Fig. 2a and b. Because this transition is bound to be gradual anyway, I believe this should not be taken as a major issue. Moreover, the scheme of Fig. 2a would presumably not last for long in a variable environment, because it is outperformed by schemes with variable mutation rates (Ram and Hadany 2012; van Hateren 2013). Finally, it may be argued that a system that evolves according to Fig. 2a, with all causal factors of a passive nature, is indeed closer to a mechanistic chemical system, merely maintaining and replicating itself, than to a living entity displaying some degree of agency.

Proving that a newly discovered system uses active causation would presumably require a detailed molecular analysis or an extensive analysis of its behaviour and evolution. However, several indicators for the presence of active causation might be easy to observe qualitatively: initiative, causal autonomy (i.e., partial independence of external causes), agency, goal-directedness, and self-interest.

References

Al Mamun AA, Lombardo MJ, Shee C, Lisewski AM, Gonzalez C, Lin D, Nehring RB, Saint-Ruf C, Gibson JL, Frisch RL, Lichtarge O, Hastings PJ, Rosenberg SM (2012) Identity and function of a large gene network underlying mutagenic repair of DNA breaks. Science 338:1344–1348

Bedau MA (2007) What is life? In: Sarkar S, Plutynski A (eds) A companion to the philosophy of biology. Blackwell, New York, pp 455–471

Benner SA (2010) Defining life. Astrobiology 10:1021–1030
Burgin M (2010) Theory of information: fundamentality, diversity and unification. World Scientific, Singapore
Cleland CE (2012) Life without definitions. Synthese 185:125–144
Cleland CE, Chyba CF (2002) Defining 'life'. Orig Life Evol Biosph 32:387–393
Darwin C (1859) On the origin of species by means of natural selection. J. Murray, London
Di Paolo EA (2005) Autopoiesis, adaptivity, teleology, agency. Phenomenol Cogn Sci 4:429–452
Eigen M (1971) Selforganization of matter and evolution of biological macromolecules. Naturwissenschaften 58: 465–523
Galhardo RS, Hastings PJ, Rosenberg SM (2007) Mutation as a stress response and the regulation of evolvability. Crit Rev Biochem Mol Biol 42:399–435
Kauffman SA (2000) Investigations. Oxford University Press, New York
Kauffman SA (2003) Molecular autonomous agents. Phil Trans R Soc Lond A 361:1089–1099
Laskar J, Gastineau M (2009) Existence of collisional trajectories of Mercury, Mars and Venus with the Earth. Nature 459:817–819
Lineweaver CH, Egan CA (2008) Life, gravity and the second law of thermodynamics. Phys Life Rev 5:225–242
MacLean RC, Torres-Barceló C, Moxon R (2013) Evaluating evolutionary models of stress-induced mutagenesis in bacteria. Nat Rev Genet 14:221–227
Macnab RM, Koshland DE Jr (1972) The gradient-sensing mechanism in bacterial chemotaxis. Proc Natl Acad Sci U S A 69:2509–2512
Maynard Smith J (2000) The concept of information in biology. Philos Sci 67:177–194
Michel D (2013) Life is a self-organizing machine driven by the informational cycle of Brillouin. Orig Life Evol Biosph 43:137–150
Nurse P (2008) Life, logic and information. Nature 454:424–426
Pross A (2004) Causation and the origin of life. Metabolism or replication first? Orig Life Evol Biosph 34:307–321
Ram Y, Hadany L (2012) The evolution of stress-induced hypermutation in asexual populations. Evolution 66: 2315–2328
Ruiz-Mirazo K, Peretó J, Moreno A (2004) A universal definition of life: autonomy and open-ended evolution. Orig Life Evol Biosph 34:323–346
Shee C, Gibson JL, Darrow MC, Gonzalez C, Rosenberg SM (2011) Impact of a stress-inducible switch to mutagenic repair of DNA breaks on mutation in *Escherichia coli*. Proc Natl Acad Sci U S A 108:13659–13664
Stano P, Luisi PL (2010) Achievements and open questions in the self-reproduction of vesicles and synthetic minimal cells. Chem Commun 46:3639–3653
Szathmáry E, Maynard Smith J (1995) The major evolutionary transitions. Nature 374:227–232
Thompson E (2007) Mind in life: biology, phenomenology, and the sciences of mind. Belknap, Cambridge
Tsokolov SA (2009) Why is the definition of life so elusive? Epistemological considerations. Astrobiology 9: 401–412
van Hateren JH (2013) Active causation and the origin of meaning. arXiv:1310.2063 [q-bio.PE] http://arxiv.org/abs/1310.2063
Varela FG, Maturana HR, Uribe R (1974) Autopoiesis: the organization of living systems, its characterization and a model. BioSystems 5:187–196
Walker SI, Davies PCW (2012) The algorithmic origins of life. J R Soc Interface 10:20120869