

Is Life Fundamental?

Sara Imari Walker

*BEYOND: Center for Fundamental Concepts in Science, Arizona State University
Blue Marble Space Institute of Science*

“One can best feel in dealing with living things how primitive physics still is.”

Albert Einstein

Although there has been remarkable progress in understanding some pieces of the puzzle, the emergence of life is still a mystery, presenting what is arguably one of the greatest unsolved questions in science. For the physicist or biologist, this may seem a problem for chemistry, and that the difficulty is simply that we don't have the know-how to engineer chemical networks quiet as complex as life, at least not yet. However, current challenges and limitations in chemical synthesis and the design of complex chemical networks may be only part of the story. The central challenge is that we don't know whether life is 'just' very complex chemistry¹, or if there is something fundamentally distinct about living matter. Until this issue is resolved, real progress in understanding how life emerges is likely to be limited.

What's at stake here is not merely an issue of complexification; the question of whether life is fully reducible to just the rules chemistry and physics (albeit in a very complicated manner) or is perhaps something very different forces us to assess precisely what it is that we mean by the very nature of the question of the emergence of life [1]. Stated most acutely, if a fully reductionist account is sufficient, and life is completely describable as the nothing other than very complicated sets of chemical reactions, what then can we say originated? Taken to the extreme, the “all life is just chemistry” viewpoint advocates in a very real sense that life does not exist and as such that there is no transition to be defined. Cast in these terms, even the avid reductionist might be unwilling, or at least hesitant, to accept such an extreme viewpoint. At the very least, although it is an open question whether this viewpoint is fundamentally correct, it is counterproductive to think in such terms – without a well-defined distinction between the two, there is no constructive mode of inquiry into understanding the transition from nonliving to living matter. As much as (or perhaps more than) any other area of science, the study of the emergence of life forces us to challenge our basic physical assumptions that a fully reductionist account is adequate to explain the nature of reality.

¹ This is not to imply that life is any less remarkable if a full account of biological organization turns out to indeed reduce to nothing more than the underlying rules of chemistry and physics subject to the appropriate boundary conditions and no additional principles are needed.

An illustrative example may be in order. It is widely appreciated that the known laws of physics and chemistry do not necessitate that life should exist. Nor do they appear to explain it [2]. Therefore in lieu of being able to start from scratch, and reconstruct 'life' from the rules of the underlying physics and chemistry, most are happy to avert the issue nearly entirely. We do so by applying the Darwinian criterion and assuming that if we can build a simple chemical system capable of Darwinian evolution the rest will follow suit and the question of the origin of life will be solved [3]. Accordingly, the problem of the origin of life has effectively been reduced to solving the conceptually simpler problem of identifying the origin of Darwinian evolution. Although this methodology has been successful in addressing specific aspects of the puzzle, it is unsatisfactory in resolving the central issue at hand by quiet stealthily avoiding addressing when and how the physical transition from nonlife to life occurs. Therefore, although few are likely to be willing to accept a simple molecular self-replicator as living, the assumption goes that Darwinian evolution will invariably lead to something anyone would agree is "alive". The problem is that the Darwinian criteria is simply too general, applying to *any* system (alive or not) capable of replication, selection, and heritage (*e.g.* memes, software programs, multicellular life, non-enzymatic template replicators, *etc.*). It therefore provides no means for distinguishing complex from simple, let alone life from non-life. In the example above, the Darwinian paradigm applies to both the precursor of life (*i.e.* a molecular self-replicator) and the living system it is assumed to evolve into, yet most might be hesitant to identify the former as living. It is easy to see why Darwin himself was trepidatious in applying his theory to explain the emergence of life². If we are satisfied to stick with our current picture decreeing that "all life is chemistry" (subject to Darwinian evolution), we must be prepared to accept that we may never have a satisfactory answer to the question of the origin of life and in fact that the question itself may not well-posed.

The central argument of this is essay is that we should not be satisfied with this fully reductionist picture. If we are going to treat the origin of life as a solvable scientific inquiry (which we certainly can and should), we must assume, at least on phenomenological grounds, that life is nontrivially different from nonlife. The challenge at hand, and I believe this is a challenge for the physicist, is therefore to determine what, *if any thing*, is truly distinctive about life. This is a tall order. As Anderson put it in his essay *More is Different*, "The ability to reduce everything to simple fundamental laws does not imply the ability to start from those laws and reconstruct the universe" [5]. From this perspective, although an explanation of the physics and chemistry underlying the components of living systems is fully reducible to known physics, for all practical purposes we just can't work in the other direction and expect to really nail the issue down. If we can't work from the bottom-up, then we must work from the top-down by identifying the most distinctive features of the organizational and logical structure of known living systems, which

² Darwin is famously quoted as stating, "It is mere rubbish thinking, at present, of the origin of life; one might as well think of the origin of matter" [4].

set them apart from their nonliving counterparts. We must therefore assume, right at the outset, that the “all life is chemistry” picture is inadequate to address the question at hand. We must ask, if life is not just complex chemistry, then what is it?

Despite the notorious difficulty in identifying precisely what it is that makes life seem so unique and remarkable, there is a growing consensus that its informational aspect is one key property, and perhaps the key property. If life is more than just complex chemistry, its unique informational aspects may therefore be the crucial indicator of this distinction. The remainder of this essay focuses on an illustrative example of how treating the unique informational narrative of living systems as more than ‘just’ chemistry may open up new avenues for research in investigations of the origin of life. I conclude with a discussion of the potential implications of such a phenomenological framework – if successful in elucidating the emergence of life as a well-defined transition – on our interpretation of life as a fundamental natural phenomenon.

“It from Bit from It”

Wheeler is quite famously quoted as suggesting that all of reality derives its existence from information, captured cleverly by the aphorism “it from bit” [6]. If Wheeler’s aphorism applies anywhere in physics, it certainly applies to life, albeit in a very different context than what Wheeler had originally intended. Over the past several decades the concept of information has gained a prominent role in many areas of biology. We routinely use terminology such as “signaling”, “quorum sensing” and “reading” and “writing” genetic information, while genes are described as being “transcribed”, “translated”, and “edited”, all implying that the informational narrative is aptly applied in the biological context. The manner in which information flows through and between cells and sub-cellular structures is quiet unlike anything else we observe in the natural world.

As we now learn it in school, the central dogma of molecular biology states that information flows from DNA → RNA → protein. In reality the situation is much more complicated than this simple picture suggests. The central dogma captures only the bit-by-bit transfer of Shannon (sequential) information; however, biology seems to employ a richer and more problematic concept than that tackled by Shannon, to the point that it is hotly debated as to what is even meant by “biological information”. Consider as an example DNA, which acts as a digital storage repository for the cell. The human genome, for instance, contains roughly 3.2 billion base pairs, corresponding to roughly 800 MB of stored data. Compare this to rare Japanese plant *Paris Japonica*, with a genomic size of a whopping 150 billion base pairs or 37.5 GB of data – one of the largest genomes known [7]. *Paris Japonica* therefore vastly outstrips humans in terms of its genome’s Shannon information content. Does this somehow imply that this slow-growing mountain flower is more complex (*i.e.* processes more information) than a human? Of course the answer is no. Across the tree of life, genome size does not appear to readily correlate with organismal complexity. This is because the genome is only a small part of the story: DNA is not a

blueprint for an organism³, but instead provides a database for transcribing RNA, some (but by no means all) of which is then translated to make proteins.

The crucial point here is the action is not in the DNA, no information is actively processed in the DNA itself [8]. A genome provides a (mostly) passive access on demand database of special DNA sequences called genes, which contribute biologically meaningful information by being read-out to produce functional (non-coding) RNAs and proteins. The biologically relevant information stored in DNA therefore has nothing to do with the chemical structure of DNA (beyond the fact that it is a digitized linear polymer). The genetic material can just as easily be another variety of nucleic acid and accomplish the same task [9]. What is important is the *functionality* of the expressed RNAs and proteins. Functional information is a very strange beast, being dictated by in part by the global context rather than just the local physics [10]. For example, the functionality of expressed RNA and proteins is context-dependent, and is meaningful only in the larger biochemical network of a cell including other expressed proteins, RNAs, the spatial distribution of metabolites, *etc.* Sometimes very different local biochemical structures will fill the same exact functional role – a phenomenon known as functional equivalence (familiar from cases of convergent evolution) where sets of operations perform the same functional outcome [11]. Only small subsets of all possible RNA and protein sequences are biologically functional. A priori, it is not possible to determine which will be functional in a cell based purely on local structure and sequence information alone (although some algorithms are becoming efficient at predicting structure, functionality is still determined by insertion in a cell, or inferred by comparison to known structures). Biologically functional information is therefore not an additional quality, like electric charge, painted onto matter and passed on like a token. It is only definable in a relational sense, and thus must be defined only within a wider context.

One is left to conclude that the most important features of biological information (*i.e.* functionality) are nonlocal. Biological information is clearly not solely in the DNA, or any other biochemical structure taken in isolation, and therefore must somehow be stored in the current state of the system (*e.g.* the level of gene expression and the instantaneous biochemical interaction network). What's more, molecular biologists are continuing to uncover a huge variety of regulatory RNAs and proteins, which acting in concert with other cellular components, dictate the operating mode (*e.g.* phenotype) of a cell. Therefore, not only is the information specifying functional roles distributed, but information control is also a widely distributed and context-dependent feature of biological organization [12].

Superficially this may not seem to be anything particularly insightful or illuminating. One might argue that such distribution of information and control is an inevitable consequence of the complexity of biochemical networks. However, on closer

³ Here a blueprint is defined as providing a one-to-one correspondence between the symbolic representation and the actual object it describes.

inspection this state of affairs is really quite remarkable for a physical system and potentially hints at something fundamentally different about how living systems process information that separates them from their nonliving counterparts. Cutting straight to the point, in biology *information* appears to have causal efficacy [11, 13]. It is the information encoded in the current state that determines the dynamics and hence the future state(s) and *vice versa* [14].

Consider a simplified example; the case of the genome and proteome systems, where the current state of the system – *i.e.* the relative level of gene expression – depends on the composition of the proteome, environmental factors, *etc.* that in turn regulate the switching on and off of individual genes that then in turn dictate the future state of the system. The subsystems cannot function when disentangled. More colloquially, this dynamic is often referred to as a chicken-or-egg problem, where neither the genotype nor the phenotype can exist in isolation. Such a dynamic is well-known from the paradoxes of self-reference [15]; picture for example Escher's *Drawing Hands* where each of a pair of hands is drawing the other with no possibility of separating the two: it is unclear which hand is the cause and which the effect.

In biology, we cannot disentangle the genotype and phenotype because causation is distributed within the state of the system as a whole (including the relations among all of the subcomponents). Similar dynamics are at play throughout the informational hierarchies of biological organization, from the epigenome [16], to quorum sensing and inter-cellular signaling in biofilms [17], to the use of signaling and language to determine social group behavior [18]. In all of these cases where the informational narrative is utilized, we observe context (state) dependent causation, with result that the update rules change with time in a manner that is both a function of the current state and the history of the organism [14]. Here casting the problem in the context of an informational narrative is crucial - the foregoing discussion may be recast by simply stating that the algorithm describing the evolution of a biological system changes with the information encoded in the current state and *vice versa*. Contrast this with more traditional approaches to dynamics where the physical state of a system at time t_1 is mapped into the state at a later time t_2 in accordance with a *fixed* dynamical law and imposed boundary conditions. Thus, for example, Newtonian mechanics provides the algorithm that maps the state of the solar system today onto its state tomorrow by specifying a trajectory through phase space. The key distinction between this situation and that observed in biology is that information doesn't "push back" and actively influence the ensuing rules of dynamical evolution as it does in living systems. This feature of "dynamical laws changing with states" as far as we know, seems to be unique to biological organization and is a direct result of the peculiar nature of biological information (although speculative examples from cosmology have also been discussed, see e.g. [19]). It therefore serves as a contender for defining life in the transition from nonliving to living matter.

Wheeler's dictum, as applied to the biological realm should therefore read more as

*“it from bit from it”*⁴, where lower levels of matter dictate the informational state of a system which then in turn dictates its future evolution. In this picture, *life is a dynamical phenomenon that emerges when information gains causal efficacy over the matter it is instantiated* [20]. A situation made possible by the separation of information from its physical representation (*i.e.* through functional equivalence). Thus, in biology the informational narrative is freed up to be almost independent of the material one and we may sensibly discuss cell-cell signaling, or sense data flowing along nerves, without specific reference to the underlying activity of electrons, protons, atoms or molecules. Of course all information requires a material substrate, but the important point here is that life cannot be understood in substrate terms alone. Thus it is meaningless to say that any single atom in a strand of DNA is alive. Yet, it is meaningful to state that the organism as a whole is living. “Aliveness” is an emergent global property.

Informational Efficacy and the Origin of Life

The liberation of the informational narrative from the material one potentially elicits a well-defined physical transition (even if currently not well-understood), which may be identifiable with the physical mechanism driving the emergence of life. In this picture, the origin of life effectively mediates the transition whereby *information* a ‘high-level’ phenomenon gains causal efficacy over matter in a top-down manner⁵ [20]. In physics we are used to the idea of ‘bottom-up’ causation, where all causation stems from the most fundamental underlying layers of material reality. In contrast, top-down-causation is characterized by a ‘higher’ level influencing a ‘lower’ level by setting a context (for example, by changing some physical constraints) by which the lower level actions take place, such that causation can also run in the opposite direction in organizational hierarchies [21, 22]. Thus, top-down causation opens up the possibility that high-level non-physical entities (*i.e.* information) may have causal efficacy in their own right [19, 23].

There is a vast literature suggesting that top-down causation as unifying mechanistic principle underlying emergence across the sciences, from quantum physics to computer science, to evolutionary biology, to physiology and the cognitive and social sciences (see *e.g.* [22]). In some areas of science, such as physiology, the existence of top-down causal effects is taken as self-evident and essential to making scientific progress. For example, it is not even a subject of debate that information control is widely distributed within a living organism (and thus that causation is also distributed). In other areas of science, such as chemistry and physics, which may be more familiar to the reader, top-down causation is not

⁴ Perhaps an even better dictum might be *“it from bit from it from bit ... ad infinitum”* to capture the self-referential nature of dynamical laws changing with states.

⁵ In practice, ‘top’ and ‘bottom’ levels are typically not easily identified in hierarchical systems. Conceptually one may view both top-down and bottom-up causal effects as inter-level phenomenon, occurring between neighboring levels in a hierarchy, a phenomenon referred to as ‘level-entanglement’ by Davies (not to be confused with entanglement in the quantum mechanics) [19].

nearly as widely accepted. In particular, its role in chemistry is not well understood at all [24]. Poised at the intersection of the domains of science where top-down causation is widely accepted (biology) and where it is as not readily apparent (chemistry and physics) sits the emergence of life, suggesting that some very interesting physics may be occurring at this transition, and it may have everything to do with the appearance of genuinely new 'high-level' causes.

Adopting this picture as constructive scientific inquiry into the emergence of life, an important question immediately presents itself: if a transition from bottom-up causation only (*e.g.* at the level of chemistry), to top-down (intermingled with bottom-up) causation may be identifiable with the emergence of life, what sets the origin of life apart from other areas of science where the role of top-down causation is clearly evident? As outlined by Ellis, there may in fact be several different mechanisms for top-down causation, which come into play at different hierarchical scales in nature [13]. In this regard, there may in fact be something unique to the emergence of life, and it has everything to do with the unique informational narrative of living systems as described in the previous section. Namely, biological systems (and other physical systems derivative of the biosphere such as computers and societies) seem to be unique in their implementation of top-down causation via information control [11, 13]. According to Auletta *et al.* who have rigorously defined this concept in the biological realm "Top-down causation by information control is the way a higher level instance exercises control of lower level causal interactions through feedback control loops, making use of functional equivalence classes of operations" [11]. The key distinction between the origin of life and other realms of science is therefore due to the onset of distributed information control, enabling context-dependent causation, and information thus effectively becomes a cause. Cast in the language of the previous section this is just another way of stating that the origin of life might be associated with the onset of dynamical laws changing with states [20].

In contrast to other quantities attempting to capture the role of information in living systems, such as functional or semantic information, or even 'dynamical laws changing with states' (*e.g.* self-referential dynamics), causality is readily definable, and in principle measurable (although often difficult in practice). This is a primary reason why top-down causation is widely heralded as one of the most productive formalisms for thinking about emergence [22]. This framework therefore potentially enables a methodology for identifying a non-trivial distinction between life and nonlife, delineated by a fundamental difference in how information is processed. For the later, information is passive, whereas for the former information plays an active role and is therefore causally efficacious. The catch is that one must be willing to accept (at the very least on phenomenological grounds) the causal role of information as a defining feature in the story of life right along side the substrate narrative of the underlying chemistry. This forces new thinking in how life might have arisen on lifeless planet, by shifting emphasis to the origins of information control, rather than the onset of Darwinian evolution or the appearance of autocatalytic sets (that lack control) for example, which do not rigorously define

how/when life emerges. It 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). It may also encourage new thinking about the emergence of the apparent arrow of time in the biosphere, trending in a direction of increasing complexity with time: dynamical evolution where laws change with states is likely to not be time-reversal invariant (although this remains to be rigorously demonstrated). Once life emerges, we might therefore expect it to complexify and diversify over time, particularly as information gains causal efficacy over increasingly higher-levels of organization through major evolutionary innovations [25].

In practice, utilizing this framework as a productive paradigm for addressing the emergence of life will likely be very difficult. We currently don't have any good measures this transition. Although there is a vast literature in top-down causation, the role of a possible shift in informational efficacy (control) and thus causal structure as the key transition mediating the emergence of life has been absent in nearly all discussions of life's origins (see *e.g.* [20] for an exception relevant to this discussion). Part of the challenge is that we do not have the proper tools yet. Walker *et al* proposed one possible measure, applying transfer entropy to study the flow of information from local to global and from global to local scales in a lattice of coupled logistic maps [25]. Nontrivial collective behavior was observed each time the dominant direction of information flow shifted from bottom-up to top-down (meant to act a toy model for the transition from independent replicators to collective reproducers characteristic of many major evolutionary transitions). However, this measure falls far short of being satisfactory. In particular, it doesn't capture true emergence where the parts do not exist without the whole (*i.e.* the cells in your body cannot exist outside of the multicellular aggregate that is you). It also doesn't capture the causal relations among lower level entities and therefore is incapable of quantifying how the informational state of a system influences these lower level causal relations. In a very different context, a step in this direction may be provided by Tunoni's measure of integrated information ϕ , which has been proposed as a way to quantify consciousness by measuring causal architecture based on network topology [26]. This measure effectively captures the information generated by the causal interactions of the sub-elements of a system beyond that which is generated independently by its parts. It therefore provides a measure of distributed information generated by the network as a whole due to its causal architecture. A version of the theory whereby ϕ itself is treated as a dynamical variable that then may influence the underlying causal relations among sub-elements might provide a way of quantifying the causal efficacy of information in the context that has been discussed here. Additional formalisms will need to also account for reliable encodings, where the same high-level phenomenon is reliably produced. In biology we have the example of the genetic code, but are far from decoding more distributed aspects of algorithmic information in processing as occurs in the epigenome or the connectome.

It is an open question what will ultimately provide a useful phenomenological

formalism for understanding the emergence of life. At the minimum the framework presented here provides a non-trivial distinction between life and nonlife and thus formulates the origin of life as a well-defined scientific problem, a key requirement for rigorous inquiry into life's emergence as discussed in the introduction. Life may be identified as fundamentally distinct from 'just' complex chemistry due to its causal structure dictated by the causal efficacy of information. This immediately suggests several lines of inquiry into the emergence of life (which may or may not be practical at present). A top-down approach is to identify the causal architecture of known biochemical networks by applying measures (such as ϕ , or other measures of causal relationships [27]), for example by focusing on regulatory networks (information control networks) or on ancient biochemical pathways. A bottom-up approach is to determine how information control emerges *ab initio* from chemical kinetics as well as how control evolves once this "information takeover" has occurred. Some of these principles will likely be testable in simple laboratory systems. A third line of inquiry could focus on the fundamental aspects of the problem, *i.e.* the aspect of state-dependent dynamical laws, or the reproducibility of high-level outcomes via reliable encodings.

This is only a place to start, and it is entirely possible that additional and/or other novel physical principles will be required to pin-down what really happened in the emergence of life. Whatever proper formalism emerges, we should not shy away from treating life as a distinct and novel physical phenomenon when addressing its origins. If this line of inquiry provides a productive framework for addressing the origin of life, a question, which must eventually be asked, is: is life fundamental? For example, characterizing the emergence of life as a shift in causal structure due to information gaining causal efficacy over the matter it is instantiated would mark the origin of life as a unique transition in the physical realm. Life would therefore be interpreted as logically and organizationally distinct from other kinds of dynamical systems⁶, and thus be a novel state of matter emerging at higher levels of reality. Our usual causal narrative, consisting of the bottom-up action of material entities only, would therefore be only a subset of a broader class of phenomena – including life – which admit immaterial causes in addition to material ones and which are characterized by their causal structure. We would therefore have to consider that higher levels of reality admit the emergence of novel fundamental phenomena.

References

1. Cleland CE, Chyba CF. Defining Life. *Orig Life Evol Biosph.* (2002) 32:387–393.
2. Davies PCW. *The Fifth Miracle: The Search for the Origin and Meaning of Life.* New York, USA: Simon and Schuster, 1999.
3. Joyce G. Bit by bit: The Darwinian Basis of Life. *PLoS Biol.* (2012) 418:214 – 221.
4. Darwin. C. Letter to J.D. Hooker, 29 March 1863. In F. Burkhardt and S. Smith (eds.) *The Correspondence of Charles Darwin 1863* (1999), vol. 11, 278.

⁶ Note this does not preclude that there may exist a gradation of states which are "almost" life with properties somewhere between completely passive and active informational dynamics, *i.e.* some parts might exist autonomously – an interesting question to consider in the context of astrobiology.

5. Anderson PW. More is Different. *Science*. (1972) 177:393–396.
6. Wheeler, JA. Sakharov revisited: “It from Bit”. *Proceedings of the First International A D Sakharov Memorial Conference on Physics, Moscow, USSR*. M Man’ko (ed) Nova Science Publishers, Commack, NY, 1991.
7. Pellicer J, Fay MF, and Leitch, IJ. The Largest Eukaryotic Genome of Them All? *Botanical Journal of the Linnean Society* (2010) 164 (1): 10.
8. Noble, D. Genes and Causation. *Phil. Trans. R. Soc. A* (2008) 366: 3001 – 3015.
9. Pinhero VB, *et al.* Synthetic Genetic Polymers Capable of Heredity and Evolution. *Science* (2012) 336: 341-344.
10. Auletta G. *Cognitive Biology: Dealing with Information from Bacteria to Minds*. Oxford, UK: Oxford University Press, 2011.
11. Auletta G, Ellis GFR, Jaeger L. Top-down Causation by Information Control: From a Philosophical Problem to a Scientific Research Programme. *J. R. Soc. Interface* (2008) 5:1159–72.
12. Alon, U. *An Introduction to Systems Biology: Design Principles of Biological Circuits*. CRC Press Taylor & Francis, 2006.
13. Ellis GFR. Top-down causation and emergence: some comments on mechanisms. *J. R. Soc. Interface* (2012) 2(1): 126–140.
14. Goldenfeld N, Woese C. Life is Physics: Evolution as a Collective Phenomenon Far From Equilibrium. *Ann. Rev. Cond. Matt. Phys.* (2011) 2(1): 375–399.
15. Hofstadter D. *Godel, Escher, Bach: An Eternal Golden Braid*. New York: Basic Books, Inc., 1979.
16. Davies PCW. The Epigenome and Top-Down Causation. *J. R. Soc. Interface* (2012) 2(1): 42–48.
17. Parsek MR and Greenberg EP, Sociomicrobiology: The Connections Between Quorum Sensing and Biofilms. *Trends in Microbiol.* (2005) 13: 27-33.
18. Flack JC, de Waal F. Context Modulates Signal Meaning in Primate Communication. *Proc Natl Acad Sci USA*. (2007) 104(5): 1581–1586.
19. Davies PCW. *The Physics of Downward Causation*. In: Clayton P, Davies PCW (eds). *The Re-emergence of Emergence*. Oxford, UK: Oxford University Press; 2006. p. 35–52.
20. Walker SI and Davies PCW. *The Algorithmic Origins of Life* (2012) arXiv preprint:1207.4803
21. Campbell DT. *Levels of Organization, Downward Causation, and the Selection-Theory Approach to Evolutionary Epistemology*. In: Greenber G, Tobach E (eds). *Theories of the Evolution of Knowing*. T.C. Schneirla Conference Series; 1990. p. 1–15.
22. Ellis GFR, Noble D, O’Connor T. Top-down causation: an integrating theme within and across the sciences? *J. R. Soc. Interface* (2011) 2: 1–3.
23. Ellis GFR. *On the Nature of Emergent Reality*. In: Clayton P, Davies PCW (eds). *The Re-emergence of Emergence*. Oxford, UK: Oxford University Press; 2006. p. 79-107.
24. Scerri ER. Top-down Causation Regarding the Chemistry–Physics Interface: A Sceptical View. *Interface Focus* (2012) 2:20–25.
25. Walker SI, Cisneros L, Davies PCW. Evolutionary Transitions and Top-Down Causation. *Proceedings of Artificial Life XIII*. (2012) p. 283–290.
26. Tononi G. An Information Integration Theory of Consciousness. *BMC Neuroscience* (2004) 5:42.
27. Pearl J. *Causality*. Cambridge, UK: Cambridge University Press, 2000.