

Special Issue, “What is Life?”

Vol. 9, No. 1 (2026)

ISSN: 2532-5876

Open access journal licensed under CC-BY

DOI: 10.13133/2532-5876/19492

Defining Life: A Conversation

Karina Kofman¹, Christoph Adami², František Baluška³, Joshua Bongard⁴, Douglas Brash⁵, Paco Calvo⁶, Gordana Dodig-Crnkovic⁷, Pranab Das⁸, Chris Fields⁹, Tom Froese¹⁰, Pier Luigi Gentili¹¹, Georgi Yordanov Georgiev^{12, 13}, Scott F. Gilbert¹⁴, Timothy N. W. Jackson¹⁵, Stuart Kauffman¹⁶, Nick Lane¹⁷, Perry Marshall¹⁸, Dan McShea¹⁹, William Miller²⁰, Stuart Newman²¹, Alistair Nunn²², Theodore P. Pavlic²³, Martin Picard²⁴, William Ratcliff²⁵, †Arthur Reber²⁶, James Shapiro²⁷, Aaron Sloman²⁸, Ricard Solé^{29, 30}, Carlos Sonnenschein³¹, Susan Stepney³², Richard Watson³³, Olaf Witkowski³⁴, and Michael Levin^{35, 36 *}

¹Faculty of Dentistry, University of Toronto, Toronto, Canada • ²Michigan State University, Department of Microbiology, Genetics, and Immunology, Department of Physics and Astronomy, USA • ³IZMB, University of Bonn, 53115 Bonn, Germany • ⁴Department of Computer Science, University of Vermont, Burlington, VT 05405, USA • ⁵Departments of Therapeutic Radiology and Dermatology, Yale School of Medicine, New Haven, CT 06520-8040, USA • ⁶Minimal Intelligence Laboratory (MINTLab), University of Murcia, Murcia, Spain • ⁷Chalmers University of Technology, University Gothenburg, and Mälardalen University, Sweden • ⁸Professor of Physics Emeritus, Elon University, North Carolina, USA • ⁹Allen Discovery Center at Tufts University, Medford, MA 02155, USA • ¹⁰Okinawa Institute of Science and Technology Graduate University, Okinawa, Japan • ¹¹Department of Chemistry, Biology, and Biotechnology, Università degli Studi di Perugia, 06123 Perugia, Italy • ¹²Department of Biological and Physical Sciences, Assumption University, Worcester, MA 01609, USA • ¹³Physics Department, Worcester Polytechnic Institute, Worcester, MA 01609, USA • ¹⁴Department of Biology, Swarthmore College, Swarthmore, PA 19081 USA • ¹⁵Australian Venom Research Unit, Department of Biochemistry and Pharmacology, University of Melbourne, Australia • ¹⁶The Institute for Systems Biology, Seattle, WA 98109-5263, USA • ¹⁷Department of Genetics, Evolution and Environment, University College London, London, UK • ¹⁸Evolution 2.0, 805 Lake Street #295, Oak Park IL 60301, USA • ¹⁹Department of Biology, Duke University, Durham, NC 27705, USA • ²⁰Bioverse Foundation, USA • ²¹Department of Cell and Molecular Biology, New York Medical College, Valhalla, NY 10595, USA • ²²Research Centre for Optimal Health, Department of Life Sciences, University of Westminster, London, UK and The Guy Foundation, Dorset, UK • ²³School of Computing and Augmented Intelligence, School of Life Sciences, Arizona State University, Tempe, AZ 85281, USA • ²⁴Division of Behavioral Medicine, Department of Psychiatry, Columbia University Irving Medical Center, New York, NY 10032, USA • ²⁵School of Biological Sciences, Georgia Institute of Technology, Atlanta, GA 30030, USA • ²⁶Department of Psychology, University of British Columbia, Vancouver, Canada • ²⁷Professor emeritus, Department of Biochemistry and Molecular Biology, University of Chicago, USA • ²⁸University of Birmingham, UK • ²⁹ICREA-Complex Systems Lab, Universitat Pompeu Fabra, Dr Aiguader 88, 08003 Barcelona, Spain • ³⁰Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico, USA • ³¹Tufts University School of Medicine, Boston MA 02111, USA and Centre Cavailles. ENS. Paris 75005, France • ³²Department of Computer Science, University of York, York, YO10 5DD, UK • ³³Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK • ³⁴Cross Compass Ltd. Cross Labs and University of Tokyo College of Arts and Sciences, Japan • ³⁵Allen Discovery Center at Tufts University, Medford, MA 02155, USA • ³⁶Wyss Institute for Biologically Inspired Engineering, Harvard University, Boston, MA 02115, USA

*Corresponding author: Michael Levin, Email: michael.levin@allencenter.tufts.edu

Abstract

Life is one of the most fascinating features of the physical world. Despite centuries of scientific study, experts still disagree about the definition, and even the possibility or utility of a definition, of this field. In a recent paper, we used AI to analyze the conceptual space formed by definitions of life given by a select set of modern workers in the life sciences and related fields. However, some of the most interesting material emerged as real-time conversations among those polled. In order to ensure that these ideas are not lost to the peer-reviewed scientific record, we here provide a minimally-edited (largely verbatim) transcript of the email chain among leading thinkers, containing numerous clarifications, disagreements, and challenges that enrich the topic of Life. It is our hope that this case study serves as an example for future papers, since the exchange of ideas among scientists is at least as interesting and valuable as formal scientific manuscripts written from a single perspective.

Keywords: biology, conceptualization, thermodynamics, evolution, epistemology, development, computational biology, self-organization, Kantian wholes

Citation: Kofman, K ... Levin, M 2026, “Defining Life: A Conversation”, *Organisms: Journal of Biological Sciences*, vol. 9, no. 1, pp. 11–45. DOI: 10.13133/2532-5876/19492.

Introduction

Life is one of the most fascinating features of the physical world. Despite centuries of scientific study, experts still disagree about the definition, and even the possibility or utility of a definition, of this field which impacts physics, chemistry, space exploration, medicine, engineering, philosophy, cognitive science, and many areas of ethics and public policy. Classic approaches have emphasized metabolic processes, reproduction, and cellular organization (Gánti 2023; Sousa, Domingos & Kooijman 2008; Gilbert 1982; Margulis & Sagan 2000; Rosen 1991; Boden, 2003), thermodynamic properties, energy utilization, and entropy reduction (Schrödinger 1944; Prigogine 1980; England 2020), or information processing, self-organization, and emergent complexity (Langton 1997; Holland 1975; Banerjee 2021; Seoane & Sole 2018; Wolfram 2002). Advances in synthetic biology (Blackiston, *et al.* 2021), artificial intelligence (Faggin 2024), and astrobiology (Benner 2010) continue to challenge what it means to be alive, with lively debates about the possibility, or need, for a single definition in this field (Gayon *et al.* 2010; Mariscal 2023; Parke 2023; Malaterre & Chartier 2021; Popa 2004; Levin 2022; Cleland & Chyba 2010; Mariscal & Doolittle 2020; Machery 2012).

In a recent paper, we used AI methods to analyze the conceptual space formed by definitions of life given by a select set of modern workers in the life sciences and related fields (Bender *et al.* 2025). However, some of the most interesting material emerged as real-time conversations among those polled. There is currently a paucity of mechanisms by which such conversations can be preserved in the scientific record. This carries considerable opportunity cost, as the exchange of ideas among scientists is at least as interesting and valuable as formal scientific manuscripts written from a single perspective. In order to ensure that these ideas are not lost to the peer-reviewed scientific record, and to perhaps serve as a proof-of-concept for other such projects, we here provide a minimally-edited (largely verbatim) transcript of the email chain among leading thinkers, containing numerous clarifications, disagreements, and challenges that enrich the topic of Life.

Methods

The content of the original paper spawned many email conversations among some (but not all) participants. Michael Levin and Karina Kofman compiled the text and linearized the numerous branches of replies as much as possible. All participants are listed as middle co-authors, and had a chance to make light edits to their text as desired.

Conversation

The conversation below began as the co-authors began responding to the initial prompt: defining Life (or arguing against the utility of a definition) in 3 sentences or less. As we collected responses with definitions of life, we used AI on the growing set, allowing it to discover specific metrics of the concept of “Life”, as seen through proposed definitions that were shared with the group. The following is the email conversation among co-authors as they discussed life after seeing each other’s definitions. The original definitions are appended to the bottom of this text for reference.

Theodore Pavlic:

“Metric is objective, but the threshold above which one infers that the metric is “life” is observer relative”

↳ *Arthur Reber:* which struck me as... let’s just say odd. If the determination of whether something is alive is conditional on being “observer relative” then the observer in this equation has to be alive to make said determination... I think I see a problem here.

↳ *Theodore Pavlic:* Don’t blame the player. Blame the game. If you project a complex, multidimensional space onto a manifold with far lower dimensions, you are bound to have some residual error. In fact, that’s the whole problem with defining “life” (and why it’s somewhat of an uninteresting question, if you really think about it – I don’t really care what is life... but I do care about the how’s and why’s of many of the subprocesses that most people would agree are a part of life).

Alistair Nunn:

This reminds me of the central discussions that still go on today about the differences of opinion between Bohr vs Einstein *et al.*, and the role of observation and

what quantum mechanics is actually describing. I am but a simple biochemist, but how we may view life is still, at least in my mind (define that one), going to be informed by what quantum mechanics is and possibly vice versa. Is life a spooky outcome of the unknowable, or have we simply not learnt enough about our Universe yet? It is quite possible that we may have to use biology to explain quantum mechanics; now wouldn't that be a thing and in the process define life?

Doug Brash:

Another way to look at it is that this is a pattern recognition question: we each have some pattern in mind when we use the word, so surely there is a definition of the pattern. It is then a separate question whether a) we all agree; b) this instantiated pattern is what we thought was living, or is just one mechanism toward achieving it; and c) whether that pattern is misleading (as in 'heat flow' leading to a proposed mechanism of 'caloric').

William Miller:

Life is binary (i.e. an organism is alive or not), if we had the correct tools to make that measurement with deep precision, which we do not. So, the answer would be 'binary' at its essence, and continuous, in a practical sense, at our current state of the art. For example, microbial cells seem subject to a binary inflection (alive or not), but their endospores have no measurable active metabolism and only seem alive when germinating, which can be after thousands of years of total inactivity. According to scientific measurement, which is what we care about, it depends on context. When we pronounce a multicellular organism, such as ourselves, to be dead, that can be correct with respect to a functioning organic whole as the entirety we care about, but individual cellular constituents (liver, bone marrow cells, cornea, white blood cells) remain alive and continue to function for many days, contributing to the necrobiome. Many of those cells, and our constituent microbial partners, reenter the planetary life cycle in ways that are only beginning to be explored. Further, the issue of viruses remains problematic. My colleagues and I have decided to call viruses 'liminal' in that we regard them as alive inside cells and not necessarily outside of them. Obviously, that's a hedge, since we don't know.

Arthur Reber:

Thoughts from a cognitive psychologist who just happened to have written a dictionary and, along with my co-authors (Baluška and Miller – both of whom were included here), wrote *The Sentient Cell* (Reber, Baluska, & Miller 2023). In that volume we tried to deal with the issue of finding useful, effective definitions of loose and slippery terms. In our case they were "consciousness," "sentience," "cognition," and the like. We decided that trying the "define your terms" argument that philosophers are fond of, to just take what we called a "folk psychology" approach. Treat them all in a manner that reflected popular usage and view them as rough synonyms of each other. In short, we decided not worry about lexicographic niceties. One of our reasons here was that what we'd discovered in the literature is that often arguments weren't about the science but the terms used in discussions – and we wanted to prevent this. It worked. No clarity was lost. We never defined our terms but did have a long Appendix outlining the history of these terms and how others had used them over time. I suspect that this approach might work here as well...

Chris Fields:

Here is an old paper (Fields 2014) on a popular definition of "objectivity" in physics and why it is either circular or empty. It's formulated in quantum-theoretic language, but that hopefully is not much of an issue. The argument works in classical physics, too. The central point is that physics doesn't produce well-defined boundaries between "parts" of the universe (= some isolated system). Observers (of whatever kind, with whatever capabilities) have to make them up. "Cutting Nature at the joints" is relative to us (or whoever) and our (or their) categories and measurement capabilities.

Stuart Newman:

To mark some agreements, and to draw some distinctions, in my definition (Table 1, Stuart Newman) I attempted to characterize living systems in physical terms—what they are made of, what the thermodynamic basis of their existence is, how they differ materially from nonliving things. My characterization agrees with those of Stu Kauffman (Table 1, Stuart Kauffman) and Eörs Szathmáry (Table 1, Eörs Szathmáry) on living systems constituting Kantian wholes and with Paul Davies (Table 1, Paul Davies) on our lack of a physics

that accounts for cellular life's organizational aspect. Being a proponent of the existence of qualitatively distinct forms of matter (Newman & Sarkar 2026), I respectfully differ with Dan McShea (Table 1, Dan McShea) in believing that life is a natural kind, and with Chris Fields (Table 1, Chris Fields) in holding that a proper role of physics is to carve nature at its joints (Newman 2020). Living systems have attributes, consequent to their material nature, of agency, capacity to learn, to define and process information, to creatively forge novel ways of existence, to enter into social arrangements and to love, which have been noted by many in this group. I am not sure that these are definitional of life, though. I wrote separately to Mike that I think life is an all-or-none, not continuous (with non-life) entity, although there are different varieties of living systems. I indicated my belief that entities that some might consider ambiguous cases (viruses, robots) are all derivative of prior, genuine life forms.

↳ *Stuart Kauffman*: Stuart, I am glad you agree with Eörs and me that living organisms are Kantian Wholes where the Parts exist for and by means of the Whole. We are all Kantian Wholes, we with our hearts and livers. To Paul Davies and all of us, please consider the beautiful idea of Constraint Closure, introduced in 2015 by Mael Montevil and Mateo Mossio (Montévil & Mossio 2015). Constraint Closure may well be a basic part of the physics of organization that we currently miss. In a constraint closed system, a set of boundary conditions that constrain the release of energy into the few degrees of freedom that then constitutes thermodynamic work, constructs the very same set of boundary conditions! We construct our automobiles, they do not construct themselves. A Kantian Whole that achieves Catalytic and Constraint closure literally constructs itself. Living cells achieve constraint closure. Boundary conditions in the cell constrain the release of energy that does the thermodynamic work to construct the very same boundary conditions. If this is not a sufficient definition of life, perhaps it is a necessary part of any definition of life. All of this reproduction occurs in an open thermodynamic system in an environment of course, and boundaries surely help a lot.

↳ *Stuart Newman*: Am I correct in thinking that closure of constraints (based on your work cycle

formulation) is a description of the arrangement of processes in a cell, but not itself a generative principle? Sort of like sedimentation without the concept of gravity. It's certainly not an outcome of "ordinary" self-organizational phenomena.

↳ *Stuart Kauffman*: Thank you for the question, Stuart. May I please give a concrete example for discussion? Gonen Ashkenasy at the Ben Gurion in Israel has a 9 peptide collectively autocatalytic set. Slightly oversimplified, Gonen feeds half peptides fragments of each of the nine peptides into the open thermodynamic system. Peptide 1 then binds two half fragments of peptide 2. Because the two fragments are bound by peptide 1, which is now a boundary condition, the activation barrier to ligate the two fragments of peptide 2 is lowered. Then a peptide bond between the two half fragments if formed, an ATP is used, and work is done to synthesize the peptide bond between the two half fragments of peptide 2 creating a second copy of peptide 2. The same obtains for all nine peptides around a "cycle". Each peptide is, in fact, a boundary condition constraint leading to a constrained release of chemical energy into a few degrees of freedom that therefore constitutes the thermodynamic work done to construct the next peptide. And that next peptide is, in fact, another boundary condition. The entire set of 9 peptides is a Kantian Whole, each Peptide is a Part. The set of nine peptides mutually catalyze their mutual formation. The nine peptides are the nine boundary conditions on the release of energy into a few degrees of freedom that constitutes the thermodynamic work that constructs the very same nine peptide boundary conditions. Thus, Gonen's collectively autocatalytic set is: i. a Kantian Whole, ii achieves Catalytic closure, iii. It achieves Constraint closure. Several features seem interesting. i. Gonen's system is NOT von Neumann's universal Constructor, U, with the instruction, I, inserted inside U. In von Neumann's system, "I" has a dual use: Instructions to be decoded to construct specifically a second U, and then I is merely an object to be copied, and the copy is to be inserted in the new U. In contrast, in Gonen's Kantian Whole, there appears to be no separate "structure I" which "encodes" the Kantian Whole and is

then “decoded” to construct specifically Gonen’s Kantian Whole. This seems true and very odd. The distinction between hardware and software seems to vanish. Yet the genetic code is surely a code with encoded biological proteins. ii. Given the concept of a Kantian Whole, we can non-circularly define the function of a Part. The function of a Part is that subset of its causal consequences that sustains the Whole. The function of your heart is to pump blood, not make heart sounds. A fundamental implication is that the same Part may have other causal consequences that can be co-opted for new functions, i.e. Verba and Gould’s exaptations that seem sufficient conditions for open-ended evolution. Stuart to come back to your question: Am I correct in thinking that closure of constraints (based on your work cycle formulation) is a description of the arrangement of processes in a cell, but not itself a generative principle? Sort of like sedimentation without the concept of gravity. It’s certainly not an outcome of “ordinary” self-organizational phenomena. I guess Constraint Closure in a Kantian Whole is a generative principle. I think Maël and Matteo may also apply constraint closure to other systems where, perhaps, it is not “generative”. I suspect you are right Stuart. Whatever this is, it certainly does not seem to be the outcome of “ordinary” self-organizational phenomena. These Kantian Wholes are systems that literally construct themselves by using boundary conditions to constrain the release of energy into the few degrees of freedom that do the specific work that recursively constructs new copies of their own boundary conditions. This is not true of a Ferro-Magnet at the Curie temperature.

Christoph Adami:

Perhaps it’s worth including my more verbose definition from my 1998 book. I’ve shortened it to what you have in the table (Table 1, Chris Adami) because a lot of the extra words are redundant (but that may not be immediately obvious to a new reader). The 1998 definition was: “Life is a property of an ensemble of units that share information coded in a physical substrate and which, in the presence of noise, manages to keep its entropy significantly lower than the maximal entropy of the ensemble, on timescales exceeding the “natural”

timescale of the decay of the (information-bearing) substrate by many orders of magnitude.” (Adami 1998) The “life is information that replicates itself” definition is only somewhat more specific than the one given above, because it identifies the mechanism that keeps the entropy low. My view is that for the most part this is the only mechanism we will ever find, but some people are close to convincing me that life could exist that uses an error-correcting mechanism to achieve this. But keep in mind that most error-correcting mechanisms rely on a template-copy system, so there is room to argue that self-replication and error-correction are in the same computational realm.

William Ratcliff:

I think I will stick with objective and non-observer dependent, at least as my definition goes. Whether a self-sustaining chemical system is capable of open-ended evolution is both an objective fact, and also not dependent on the reference frame of an observer (or even an observer’s existence). But also, because fitness is fitness, I don’t see a scenario where you can have life and NOT have things that ‘care’ about the system. Natural selection ensures that fitness maximizing agents are the things that will exist, and fitness maximizing agents will all care about whatever is limiting to fitness, because caring about it is adaptive. So I think you basically get your observers shaking out, for free, every time.

Olaf Witkowski:

Nice intersections there. Importantly the observer-dependence and decompositional equivalence, resonate with my “physico-experiential” framing of life in the table (Table 1, Olaf Witkowski).

Tom Froese:

I also think that genuine progress will only come from a consensus realization that we just don’t get it yet. Only by treating life as a real, irreducible domain of phenomena can we get glimpses of it, albeit only indirectly, in our never-ending failed attempts at reduction. We need to see the ongoing failure to create life from scratch as a fundamental anomaly and then re-organize our thoughts accordingly.

Timothy Jackson:

“Objectivity” is itself an ambiguous concept partly due to the fact that the Kantian lineage of modern

philosophy defines the “objective” quite specifically as something constituted by the “subject”—the two terms are always correlative. A “polycomputational” approach to constructive coarse-graining in an FEP context could be very useful in moving us back towards a relatively observer-invariant notion of the reality of “objects” (defined at particular spatial and/or temporal scales) without recourse to a “naive” pre-Kantian form of “realism”, or indeed a post-Kantian anti-realism. In evolutionary science, it seems we are often dealing with “real generals” (such as species, or function classes), whose boundaries are inherently fuzzy, admitting of liminal zones in which a distinction between “marginal” and “paradigmatic” members of a category becomes useful. Thus, despite the “objective” reality of a category like “life”, any attempt to precisely define its boundary may be irreducibly stipulative—relative to an observer, scale of analysis, or specific problematic.

Alistair Nunn:

I guess it is also possible to view the emergence of “life” as potentially an (inevitable?) solution to accelerate the heat death of the Universe, so ensuring maximum entropy, where the emergence (insert your own definition here) of “intelligence” results, especially with the emergence of technology, the ability to dissipate ever larger energy potentials and on to things like Kardashev levels. This of course arises out of the far from equilibrium self-organising dissipating principles, which some have suggested as being key (e.g., based on the work of Prigogine and say, applicable to thermal vents as one possible place for life to have started). Thus, in a way, this takes us back to the tricky problem of how life started, and differentiating it, say, from a complex auto-catalytic network that one could simply describe as complex chemistry, to something that might be described as alive. In short, the age old abiotic to biotic switcharoo problem. This of course then takes us to “why life”, and on to knotty problems like Fermi’s paradox and how to populate the Drake equation. How we view the problem at a larger scale might affect how we define it.

Carlos Sonnenschein:

Two brief comments: i) by deciding earlier to answer your kind invitation in a minimalist way, I have misread the purpose of this whole intellectual exercise. And ii) after reading the comments already submitted,

I concluded that among the over 60 participants there are two quite distinctive groups; namely, a majority one that see life within mostly an organization/order perspective (why is life the way it is?), and a minority one that adopts a narrower reproduction/proliferation one (why is there life?). This latter view, with which I identify, brings back the subject of “the first cell” and the initiation of life. Together with my Paris-based colleagues (G. Longo, M. Montevil, A. Soto, M. Mossio, P.A. Miquel) we proposed in 2016 a Theory of Organisms, in PBMB (Soto *et al.* 2016), that covers relevant aspects of the topics referenced in the emails sent this week to you, Mike. Based on the premise that proliferation with variation and motility is the default state of all cells (which is consistent with empirical data), we speculate that the proliferation of “the first cell/LUCA” was an inexorable process. In our view, this outcome resulted from the lack of physical, chemical and nutritive constraints in the prebiotic soup in which (fortuitously?) a cell put itself together. Of course, there is much more to be said about the subject... The organization/order group seems to look at LIFE once life was underway (i.e., today’s LIFE). No small task, indeed. As of now, however, I fail to see a movement through which these 2 camps aim at integrating what is already known. However, it is early in the game, and not everyone has already expressed their respective views on LIFE...

↳ *Alistair Nunn:* Dear Carlos, As Dobzhansky once said (if I have got the quote right): “Nothing in biology makes sense except in the light of evolution”, to which I would add, “and the origins of life”.

↳ *Carlos Sonnenschein:* Dear Alistair: Amen... Amen...

Nick Lane:

Just a follow-up note to this, in relation to my ‘definition’ of life not being abstract and generalisable in the way you might wish. I wonder how much scope for alternative origins there might actually be. The more I study the origin of life, the more I am left with the feeling that deterministic chemistry takes us a long way, and there are very good reasons why life is cellular, carbon based, uses membrane bioenergetics, has a universally conserved core metabolism and a genetic code with patterns that seem to reflect an autotrophic metabolism and direct physical interactions. That’s not to say that all these traits are necessary, but they

do seem to be overwhelmingly probable, given the ubiquity of water and CO_2 in the solar system and exoplanets. Let me generalise a little more. It seems to me that as soon as we get to some kind of replicating agent, most people effectively say—‘and now selection will do the rest!’ I think we over-rely on selection. It can get stuck, as happened with the evolution of bacteria and archaea—they remained morphologically simple over four billion years of evolution on Earth, despite sampling genetic sequence-space on a colossal scale in relation to eukaryotes. Evidence from the ‘RNA world’ to date comes close to demonstrating that we cannot rely on selection of ‘naked’ RNA to evolve ‘altruistic’ traits on a grand scale—which is to say there is no evidence that an RNA world can ‘invent’ metabolism to the level of a free-living cell. I would say the contrary. That was Fred Hoyle’s point about tornados and junk yards, long ago. All this constricts our room for maneuver down to the level I refer to above. I think there are good reasons why life is the way it is, whereby membrane bioenergetics drives metabolism which controls gene expression. At the origin of life this was arguably the order in which they appeared. A structured far-from-equilibrium, continuous flow environment (e.g. hydrothermal system) drives the conversion of low molecular-weight precursors (gases) into a thermodynamically favored network of chemical reactions (a protometabolism that prefigures metabolism), with nucleotides initially functioning as catalytic cofactors driving autocatalytic feedback loops (the closure that Stu talks about). Genes emerge from polymerised nucleotides via direct physical interactions with amino acids, so genetic information emerges from a weak form of templating giving rise to what Woese called ‘statistical proteins’. We would not have anything we could call life until we have genes coding for proteins, but every step along that path is deterministic chemistry which would happen in a similar way on any wet, rocky planet or moon with a CO_2 atmosphere. So tens of billions in the Milky Way alone. There is a long continuum from prebiotic chemistry to cells with genes and proteins, and life is ultimately better at breaking down the kinetic barriers to thermodynamically favored reactions, than non-living systems—life is better at increasing entropy, but the entropy comes from dissipating the same thermodynamic disequilibrium (until the evolution of oxygenic photosynthesis).

Georgi Georgiev:

This is a very timely and important question, subject to the greatest debate. Have we considered the possibility that there may not be a single definition of life, but, as we delve deeper there may be different kinds of life? There may be different forms of life with their corresponding definitions. The tendency of science is to branch, as physics has branched into optics, quantum field theory, relativity, high energy physics, etc. Can we have, let’s say, a biological form of life, which includes thermodynamics, metabolism and reproduction, which is Earth centric and Anthropocentric and an information form of life, which does not directly include thermodynamics and reproduction, but, continuously self-improves and modifies, grows, metamorphoses into something new, and includes other mechanisms for which we do not have names yet, which can happen in other non-biological and may be non-material places and substrates, such as quantum fields, vacuum energy, magnetic fields, etc.

↳ *Doug Brash:* I like the way you express the alternative life forms. My own test case is a sentient star. Yet if you use the word “life” for each of these alternative instantiations, instead of a completely different word, that suggests that you think they would have something in common. I’d call that ‘something’ the single definition, and the rest of it is just different ways of engineering the system to achieve your ‘other kinds’ of life. It seems like that distinction would work with your definition, with no need for additional definitions. Do you see a flaw in that logic? The reason I ask is that, if we want to compare definitions to gain deeper insight, then we should put the engineering into a different pot. Then comparing definitions becomes like stereoscopic vision: the apparent conflicts will create deeper insight.

↳ *Georgi Georgiev:* I like that. An umbrella definition for general life, and specific definitions for each subtype of life, that we know or that we can think of. In that sense, then, the biosphere will be a type of life, which evolves as a coherent whole, and has its own characteristics. The economy, cities, the civilization are other types of life. AI may be a completely different form of life. And then life that we cannot yet imagine. All of them will have different specific mechanisms and rules of evolution, as you say “engineering”. Then, I guess

that one of the common attributes of all life is that it always evolves. It always improves and becomes more “alive”. It always has structure, organization, which gets more efficient for its functions, i.e. it always obeys optimization principles. It grows and expands in space and in time until it reaches some limit, and with evolution it can overcome it.

↳ *Ricard Solé*: I agree with Doug about the evolutionary component. Moreover, if we include the potential for building synthetic systems, minimal synthetic protocells will be alive but unable to learn. Also, synthetic cells can be replication-free entities, as shown by Kate Adamala’s designs (Frischmon, Sorenson, Winikoff & Adamala 2021). While trying to chart the possibilities for cognitive systems, protocells or autonomy, I often end in a scenario where “definition” necessarily requires a space where different axes (developmental complexity, computational complexity, agency...) allow locating the diverse range of possibilities in some organized way. The different definitions brought by all the contributors here reflect, in my opinion, the need for a space of the possible to address the initial question.

↳ *Douglas Brash*: The axes are a nice way to look at it. Then, if the axes have been unambiguously defined, the space creates a well-posed future problem of seeing whether the dimensionality can be reduced by discovering that one of the axes is a (probably non-linear) combination of some of the others.

↳ *Joshua Bongard*: What counts as a valid dimension for a morphospace? Presumably something that predates evolution? Predates life? What would be a good metric for comparing morphospaces? That is, what is a better or worse morphospace?

Chris Fields:

One could also argue that intuitions formed over evolutionary time are little guide to good science. Counter-intuitiveness is no sufficient condition for good science, but it may well be necessary. I would like to contest, in particular, the emphasis placed by many on reproduction, or what Dawkins elevates as “replicators”. “Reproduction” is uninteresting without variation. We could imagine a scenario in which

some founder population—maybe of one individual—cranked out perfect clonal copies ad infinitum, taking over its entire universe and converting all available resources into copies of itself. Except as dramatized in science fiction—where there’s always a mutant—this scenario is boring. An imperialist’s dream, but biologically uninteresting. Now consider the opposite: some founder population makes a huge diverse array of things, most of which are also capable of making an array of different sorts of things. Every now and then, we get a copy, or something close to a copy, but most of what’s produced bears little structural or functional relation to its producer. Cell replication with differentiation looks a bit like this, but let’s assume something more extreme. Lineages from LUCA spawning new lineages from LUCA looks a bit more like this. But we could imagine that the things that are produced, or most of them, that are (mostly) able to produce other things, are not required to have the same biochemical substrates. We can clearly do this virtually, writing programs that write other programs capable of writing programs, etc. Think of it not just in terms of software, but of hardware, too—robots building a variety of other robots, most of which are capable of building robots, and with a variety of different kinds of materials. Xenobots given a feedstock of cells make more xenobots. But maybe xenobots, or something like them, would make more X-bots given a feeder stock of X type cells, for (almost) any X. If the X-bots could do the same with any feeder cell type, we’re off and running. If some of the X-bots could work with non-cellular feeder stock—say silicon chips—we’re really off and running. Would we consider such “variant producers” alive, even if they only rarely produced something like themselves? What if they made really efficient use of resources, and produced really interesting (e.g. behaviorally) “artifacts”? If not, why not? This scenario just notches up the “variation” side of the evolutionary equation. We all do something like this, of course, when we teach students...

Perry Marshall:

Blaise Agüera quoted Andrew Ellington: “there is no such thing as life. Life is not a meaningful scientific term. It is a term for poets, not scientists.” I do think there is such a thing as life. But Ellington is right, “life” is not a meaningful scientific term because science as currently defined (~ the quest to reduce the universe to equations) is unable to capture life—because the

universe cannot be reduced to equations. Nor can poetry. Which suggests poetry is a more mature discipline than science. And it's why, as Stu Kauffman says, science is facing a "third transition."

↳ *Pranab Das*: You and Blaise are on to something with the references to poetry. I agree there's a sense of the world that eludes analysis and slides away from us when we grasp at it too hard. But I'd push back on the idea that life is *entirely* in that category. "Equations" are just well-formed causal statements. It's my hope that we can identify objects and author stories that hold true about them even when those objects are complicated and squishy like life. With apologies to Ellington, I'd say that his remark sounds more like exasperation than assertion. After all, starts with "...our failure to define life". The trick with analysis is what he says we fail at—establishing a level of explanation appropriate to the relevant phenomena. I am moved by the elegant simplicities and repeated patterns of the world. Humans are good at detecting those simplicities and assigning them thingness. Life seems like a thing—we know it when we see it—so I'll wager that there's a level of analysis (if you insist, equations) that applies. To me, physics only amplifies the poetry of a falling feather or raindrops on a window at night.

Olaf Witkowski:

Perhaps science boils down to the poetry that we can meaningfully agree upon. Science is intersubjective after all. Maybe starting with a poetry of life, to later transform it into a science of life should be seen as an integrative part of our creative process to make progress. Less cryptically, I sympathize with Blaise quoting Ellington: seeking only one universal definition of life may not be a practical goal for our scientific pursuit—and part of my spiel about observer-dependency aligns with it. But, in the pursuit of the unreasonable goal, we may be able to progressively factor life into the product of its characteristic mechanisms and properties, which will bring us more insights than trying to face all the heads of the hydra at once. I embrace the current divide and conquer aspect of the field. Perhaps this is why I advocate for the diversity in artificial life approaches too. Later, it will still be time to attempt to unify the factors into a more elegant formula capturing the nature of the living state.

Paco Calvo:

Thanks for keeping me in the loop. Hmmm, interesting! I know you asked for just 1-3 sentences, but let me give you some context first, and then I'll craft it into a shorter statement. If I were to discard more orthodox or narrow-minded perspectives, my first reaction would be to move away from the type of Darwinian approaches that emphasize replication and adaptation above all else. I'd also avoid autopoietic definitions that place too much emphasis on self-maintenance and reproduction. This is because I can envision the possibility of an individual (alien) form of life that thrives, so to speak, and remains stable in terms of ontogeny, rather than phylogeny. I'd also distance myself from biochemical approaches that focus strictly on metabolism or a particular type of substrate (e.g., DNA, RNA, etc.). I don't mean to imply that the forms of life we know are not characterized by these aspects, but rather that, in order to push the boundaries, I'd like to consider different features life might have—thinking not in terms of necessary and/or sufficient conditions, but as a list of indicators that might suggest the existence of "unorthodox" life forms by our current standards. To cut a long story short, I'd propose a combination of three features: (i) A Gaian (Lovelockian) component, coupled with (ii) A phenomenological stance (e.g., Hans Jonas), and (iii) An openness to non-biological forms (e.g., artificial/virtual life forms). Putting these three aspects together, a working definition of life could be summarized as follows: Life is characterized by the possession of intrinsic goals and purposes (phenomenological), a drive for preservation based on interaction-dominant dynamics (e.g., Gaian self-regulation) among biological or artificial components, and the ability to exist within real or virtual environments (not necessarily carbon-based). Hope this is not too crazy and that it makes some sense to you, as I said this is anything but necessary/sufficient conditions...

Chris Fields:

We can distinguish two broad points of view in all of this discussion: Exceptionalists: Exceptionalists view Life as either the lineage from LUCA, or perhaps the lineage from LUCA plus some undiscovered lineages that have similar biochemistry and behavior. Life for this group is "life as we know it" plus maybe "life as we would immediately recognize it." There is

disagreement about what properties of life as we know it are more important, and about abiotic -> biotic origin mechanisms, but a clear sense of a boundary between abiotic and biotic even if we don't know exactly where or how to draw it. Non-exceptionalists or maybe even Universalists: Non-exceptionalists view Life as much broader than the lineage from LUCA, possibly based on very different chemistry or materials, etc. Universalists see Life as more or less co-extensive with persistent organization, and hence do not recognize an in principle "abiotic" or "inanimate, non-agent" category of persistent systems. For Universalists, the "origin of life" question is more or less co-extensive with the "origin of order" or "origin of persistent patterns" question. For Non-universal Non-exceptionalists (if there are any) the "origin of life" question is somewhere in between the "origin of persistent patterns" and the "origin of LUCA". This distinction is a matter of abstraction, but also a matter of how committed one is to both ontological objectivity and to an alignment of the assumed objective ontology with intuitive human categories or even human-constructable categories.

František Baluška:

My view is that only the Exceptionalists view is the true Life (which is in fact strictly cellular life). Non-Exceptionalists or Universalists consider also the Pre-Life systems and processes for Life. Only cells enclosed with their excitable membrane which defines the inside (subject) and the outside (object, environment) and their constructs (loose cell assemblies, biofilms, multicellular organism, the whole Gaian biosphere) are truly living.

Stuart Kauffman:

In a recent paper (Kauffman & Roli 2023), we hope we show that evolving life is based on physics and the Newtonian Paradigm, but is also beyond physics alone and beyond the Newtonian Paradigm. We believe we demonstrate that we can use no mathematics based on Set Theory to deduce the evolution of ever-new adaptations and new phase spaces in the evolving biosphere. Use of the Newtonian Paradigm requires a defined known phase space. This fails for the evolving biosphere. Via the newly discovered property of Constraint Closure, living cells literally construct themselves and construct new phase spaces via Natural Selection. In constraint closed systems a set of boundary

conditions constraints, [A,B,C], constrain the release of energy into a few degrees of freedom in a set of non-equilibrium processes, [1,2,3] such that thermodynamic work is done by these processes to construct the very same set of boundary conditions, [A,B,C]. This is entirely new. We construct our automobiles that link boundary conditions and processes. Gas explodes and wheels turn. But automobiles do not construct their own boundary conditions. Cell do. We struggle with the meanings of "information", e.g. Shannon. And we wonder how information gains causal purchase on the world. But changeable boundary conditions constitute "information" as well. The boundary conditions define the phase space and define the work that can be done. Andrea and I claim that evolving life is a non-deducible construction, not an entailed deduction. This seems a sufficient condition for open-ended evolution. It may not be a necessary condition. Algorithmic Artificial Life may find true open-ended evolution. The same issues of open-ended evolution seem to bear on algorithmic general artificial intelligence.

Dan McShea:

So many are exceptionalists, contending there is something unique and special about living things, something that distinguishes them in a binary, either-or way from nonliving things. So here's one thing that anti-exceptionalism offers: the investigation of the larger space of transitions. While so many of us have our intellectual microscopes trained on that interesting one-time transition from nonlife to the life that dominates the planet today, we could be missing huge numbers of: 1) partial transitions in the same direction, some of which may have left (much-modified) descendants today; 2) the many near-misses that have undoubtedly occurred, the chemical systems that had the potential to get fancier but for perhaps-interesting reasons didn't; and 3) most intriguing to me, the possibly-many backward transitions that occurred long ago at the same time as life's origin, in the billions of years since then, and possibly still today! So it's not just chemistry exploring the possible and landing on life as we know it. Life explores the possible and sometimes discovers non-life. Of course, some of that newly discovered non-life might be a rediscovery of older chemically-simple systems. But some of it might be new, with powers and properties that no previous "simple" chemical system ever had? A possible analogy: our ever-more-complex

technology produces the complex wonders of laptops and cars. But it also discovers velcro, an extraordinarily simple technology with extraordinary properties. It's all about the difficult-to-access simple. Just a thought.

Aaron Sloman:

I previously offered the following: “Life exists on or in a physical/chemical entity (e.g. galaxy, or planet) containing components that are able to extend and replicate themselves using mechanisms that repeatedly add new forms of complexity by absorbing structures from their environment, including repeatedly extending the forms and mechanisms of replication and their powers.” I now (6 November 2024) prefer the following much longer, more complex, but less elegant(?), answer: Life exists on or in a physical/chemical entity (e.g. galaxy, or planet) containing components that are able to extend and replicate themselves using mechanisms that repeatedly add new forms of complexity by absorbing matter from their environment, decomposing and recombining components of absorbed and previously stored matter, and replicating instances of themselves, using mechanisms that allow the replication processes to modify or extend the mechanisms modifying both the forms of interaction with their environments and the mechanisms of replication and their powers. A key feature of many forms of life (unlike complex geological processes such as continent formation and weather processes including formation and development of tornadoes) is use of information about what already exists (internally and externally) and about what could possibly exist (internally and externally) to influence selection of actions to change what exists both internally, e.g. through growth and repair, and externally by changing the environment. Mechanisms of interaction with the environment and mechanisms of growth, repair and replication in living systems make use not only of physical substances, but also use information to consider alternative possibilities, and make selections between alternatives. Often the information about what works is not explicitly represented in organisms, but is implicit in the history of ancestors of the organisms, some of whom had descendants that had produced alternative structural changes that prevented further replication for some reason (e.g. inability to cope with predators or climate change). There are some non-living mechanisms

created by humans (or by machines created by humans) that share some aspects of life, but nothing with the full richness of living organisms described above, e.g. the internet. It is possible that my revised answer to “What is life?” above is still too narrow—there may be biological phenomena that I either have not heard of, or have forgotten, that should be included but do not fit the above specification for life. It is also possible that this is too broad, e.g. if there are geological processes that meet that specification, though I think that is unlikely. There may be flaws of other types! I now have an improved (but still complex and messy) document on the relationship between life and information and evolved information-processing mechanisms. It significantly improves the version sent previously. I apologize to recipients who are not interested in this. The revised document remains available here <https://cogaffarchive.org/misc/whatlife.html>. This is work in progress, so I welcome corrections or suggestions for additions or improved formulations (perhaps best sent only to me, not the whole list?). The document will always have date of latest revision near the top. Until recently I had not realized that there is no satisfactory explicit definition of the concept of information I have been using, also used by many others, including Jane Austen in her novels over a century before Shannon introduced his new concept of measurable information based on bit-patterns, that has since confused many of his admirers, although he was not confused—as explained in my document. There also does not seem to be any adequate explicit definition of the concepts of “life” and “living organism” able to capture the diversity of forms of life without also including non-living systems partly because of the many, still evolving, forms of life. That raises questions about how scientific and philosophical discussions can successfully refer to forms of life and types of information without using explicit definitions—questions to be answered at some future date!

Carlos Sonnenschein:

As I mentioned before, together with a few colleagues based at the ENS in Paris (mathematicians, physicists, and philosophers among them), we published a series of papers in a PBMB volume edited by Ana Soto, Giuseppe Longo and Denis Noble (2016) leading to the proposal of 3 principles for a Theory of Organisms (Soto *et al.* 2016). The principle on which Ana Soto

and I contributed (empirically and theoretically) was the one that posits that proliferation (with variation) and motility are the default state of all cells. Here, all cells means indeed all cells (*). For over 4 decades, we have welcomed challenges to this claim (we have yet to be offered a credible alternative). There are multiple implications of the principle we proposed. Among them, I consider the following as salient: i) the only way to regulate (control?) cell proliferation in living creatures is by inhibiting it; ii) in turn, an implication of this initial claim is that the “thousands of papers” that claim that there are so-called growth factors that stimulate the proliferation of living cells are figments of the imagination of “pioneers” of this field and/or that of their followers; and iii) both cells and organisms are “agents” (each of us is free to interpret “agency” as we see fit). Others in this forum already addressed the important subject of organization/order in the maintenance of life. I didn’t have much to add to this item. Briefly, then, this has been part of the rationale to insist on asking the evolutionary relevant question why is there life on Earth? Epistemologically, it is a different question from the one asking what is life? Most biologists are aware that a “mother” cell that proliferates generates two similar, but not identical, “daughter” cells. This means that during the billions of molecular interactions occurring during a cell cycle, “errors” are inevitable. They are part of life. When I wrote the minimalist statement that “Without proliferation and motility there is no life”, I assumed that the readers would be aware that in Biology, proliferation is linked to variation. Incidentally, I daily check on this claim while walking on the streets of the city where I live while checking the plumage pattern of pigeons. Surprise... Surprise... They are all different... Not one is identical to another... More to the point, from our analysis of the literature, I have concluded that growth factors (2 Nobel Prizes) and oncogenes (2 plus Nobel Prizes) are mythical creations predicated in research papers and textbooks through intellectual and empirical incompetence, artefacts, irreproducible results, or their combination. In addition, accepting that quiescence (and not proliferation) is the default state of cells in multicellular organisms implies that the overwhelming majority of researchers in Biology today favor, probably inadvertently, Intelligent Design as a real possibility to explain LIFE. Based on this and other epistemological shortcomings, Ana and I share

the need clearly analyzed in Philip Ball’s book *How Life Works* (Ball 2023), and in Bernd Rosslenbroich book *Properties of Life* (Rosslenbroich 2023), that a New Biology is necessary to resolve the current crisis in Biology brought on us all mostly by the practice of “greedy” reductionism.

↳ Pranab Das: Carlos, I like your previous post and am a big fan of Lakatos. But I’m not sure what scientific project you are identifying in your citation of him as critic... In your previous point you direct us to the perennial base question of *why* life: in my humble opinion, there are two magical things about life. First, I think the history of life is a story of dimensional reduction. A lot of conversations center on life that as “functional” or goal-directed. But the concept of goal-orientation or some form of gradient descent is simply nonsensical in higher dimensions (both for topological and computational reasons). So there must have been some exceptional work in our history to constrain the dynamics so radically. I would answer your prompt about LUCA with that idea... that the antecedents of life must have been a succession of compatible dimensional-reductions of dynamical systems. As a physicist who remains astonished by the fact that much of our world is governed by simple one over r squared laws, I’m open to the idea that there is something about nature that loves simplicity... so perhaps such dimensional reduction is inevitable and life is ubiquitous. Second, I believe that life is always characterized by parts that are locally-dissipative. Such “active matter” gives rise to emergent characteristics that seem ubiquitous in life. My favorite example of this is the dynamism of the cytoplasm... it’s only thanks to the locally-dissipative dynamics of tubulin-kinesin bundles—and the flocking and swarming that results—that cytoplasmic streaming is possible.

Scott Gilbert:

I like the metabolic view of life and see reproduction as an extended quality and not a primary one. If we were to find metabolizing entities that didn’t reproduce, we’d call them living. (For that matter, there are nuns, priests, mules, and incels living amongst us.) This idea of metabolism being the *sine qua non* of life was beautifully explicated in Hans Jonas’ book *The Phenomenon of Life* (Jonas 2001). Here, he noted that metabolism was not a product of the organism;

rather the organism was a product of its metabolism. Metabolism is that property that enables a living being to retain its identity despite and by changing its parts. This creates a “needful freedom” for the organism, as it separates itself from and depends upon the environment (Jonas 2001). In his 1896 edition of *The Cell in Development and Inheritance*, E. B. Wilson defined Inheritance as “the recurrence”, in successive generations, of like forms of metabolism. Life is where the Second Law meets the Third Critique.

Many years ago, I wrote an article claiming there were two major views of life: replication and metabolism. Moreover, I claimed that while biochemistry and molecular biology both seek to find the underpinning of life in chemical substrates, they do so very differently. Biochemistry originated from the metabolic perspective of life and its major metaphor is the wave. Molecular biology originated from the replicative view of life and its major metaphor has been the crystal. I’ve attached the article—a Golden Oldie!—below (Gilbert 1982).

Ricard Solé:

I have a hard time seeing why movement must be a necessary condition for life (most planktonic organisms cannot perform active motion: are they dead?). Using capital letters does not make the claim more convincing. Movement was certainly key to promote the evolution of complexity (from single cells to brains) but cannot be used as a necessary condition. And—although not relevant here for what we are discussing—I disagree about the cancer topic. It has been shown that cancer cells live amid a conflict between proliferation and inhibition of proliferation. Even if the gene-centered perspective might be short of understanding all the processes (I fully agree), claiming this evidence is a fraud is extreme.

William Miller:

In thinking about your terrific interchange, I was struck by the near complete absence of the mention of cognition or cognitive processes in the highly illuminating thread I’ve been reading, and further, how few included that term in their definitions. In a just-published analysis in *Biosystems*, “Living cognition and the nature of organisms” (Rosslenbroich 2023), Breno Just and Savio Torres de Farias define cognition this way, applying it to all cells: “Cognition is the set

of informational and dynamic processes an organism must interact with and grasp aspects of its world. Understood at their most basic level, perception, memory, learning, problem-solving, decision-making, action, and other cognitive processes are basic features of biological functioning.”

True enough, however, all of these processes represent active self-referential information management to sustain life. The active management of information is required since all environmental cues are a source of uncertainty. We don’t know the origin of cognition. However, we consider it as the faculty enabling the living state to gain causal efficiency over information (Bateson’s difference that makes a difference over time), echoing how Sarah Walker and Paul Davies characterized life in one of their fine papers. Further, and crucially, the decision-making, problem-solving, and learning that characterize that living state are based on self-referential experiences, as opposed to computers. Although a noun in language, cognition is functionally a verb, denoting a broad set of overlapping cognitive processes in action. Significantly, however, those cognitive processes are experiences at scale in support of cellular decision-making, problem-solving, and learning as forms of information management in the face of living uncertainties. A useful reduction can be proposed that provides a natural bridge between the scientific method and poetic metaphors. Life is cognition. That is, life is a cognitive set of active processes as living information management, united through self-referential informative experiences. We intuit that poetry and life are entwined. The living glue of that unity is cognitive experientiality, permitting us to regard both explanatory equations and raindrops as things of beauty at our scale, and simultaneously leaving us ill-equipped to imagine the self-referential experiences of others.

Pier Luigi Gentili:

Cognition and intelligence are two other remarkable features of life. I like investigating which biological intelligence competencies can be mimicked through inanimate chemical systems in the liquid phase (the essential phase of life) and developing Chemical Artificial Intelligence (CAI). CAI relies on molecules and macromolecules as if they were the hardware, and phenomena such as chemical reactions, diffusion, migration, advection, and chemical waves as if

they were the software. The hardware and software of conventional AI are merged in the wetware of reactive chemical solutions. To know the objectives and outcomes of CAI, please check out this paper (Gentili & Stano 2024), which Stano and I wrote for the special issue on Computation And Biology (Levin & Hazan 2024), edited by Michael Levin and Hananel Hazan. We hope the development of CAI will assist us in our attempts to unveil the secrets of life. and how life emerged on Earth. Mike Levin is advocating that a remarkable feature of multicellular living beings is their multiscale competencies: there is order in biology across different spatial levels, and each organized structural level solves problems in its own relevant domain (McMillen & Levin 2024). The imitation of the mesmerizing multiscale competencies of living beings can be approached by boosting the development of a new chemistry, named “Multiscale Chemistry” (Gentili 2025), which refers to chemical systems manifesting self-organization at different spatial scales and multi-level causality when appropriate macroscopic thermodynamic forces are applied.

Tom Froese:

Reflecting on a few themes, I am happy to see that a couple of other “exceptionalist” positions emerged in the conversation, although it seems that we are a small minority? I am in agreement with Pranab Das that two essential puzzles to understand is (1) global dimensionality reduction, and (2) local increased dissipation. I see close links here with the “exceptionalist” framework I am developing for the case of mind, where I term these two aspects absorption and irruption, respectively, and assign them specific roles in a non-reductionist mind-body relation (Froese 2024). And I agree that mind is likely a specific expression of the phenomenon of life, so cognition needs to be brought into the picture. We started work on generalizing irruption theory to the case of life, starting with an explanation of teleology in terms of increased local dissipation (currently under review as a book chapter) (Froese, Karelin & Takashi 2025).

Richard Watson:

I propose to tell you which definitions are wrong and why. The gist of what follows is this. Any definition of life based on a conception of a process that acts on an object, or a program that acts on data, is wrong.

Definitions that only refer to the process/program are wrong, definitions that only refer to the object/data are wrong, and definitions that include both but maintain a process/object or program/data separation, are also wrong (I haven’t said why yet). For example, many definitions offered in this study address the adaptive process, such as evolution by natural selection (ENS). This is good inasmuch as it is intuitive that life is an ongoing process, a process of change and adaptation, whereas a static object is a dead thing [Life “always evolves. It always improves and becomes more ‘alive’” says Georgiev, Georgi (Table 1, Georgi Georgiev)] (and besides, it’d be nice to have an explanation for where the living thing comes from, and an adaptive process helps with that). However, a definition that addresses the process in a disembodied way, as though it doesn’t matter what the process acts on, is inadequate. The faith that “so long as ENS is instantiated everything else will follow”, is grossly unfounded [Nick Lane, I agree!]. Worse, there are definitions which mistake the adaptive process for life itself. To illustrate why this isn’t enough, consider Sol Spiegelman’s RNA replication experiments. This is clearly evolution, but not alive. (Likewise any artificial genetic algorithm implementing ENS on an inert data structure). So the process alone is not sufficient. It is also not necessary; consider Doug Brash’s point that he’s alive but not evolving (nor evolution itself!). The object matters—the process is neither necessary nor sufficient on its own. ENS is supposed to explain how living things arose—it is not in itself the living thing. To fail to address the object that the adaptive process acts on or produces is mistaking the explanans for the explanandum.

Other definitions address the object and its properties. For example, Stu’s Kantian whole, with catalytic closure and constraint closure, or whatever. Or Eric Smith (not represented here) “a set of constraints that carry the process, and the process that constructs the constraints” (something like that). This is good inasmuch as it is the missing part of the definitions previously mentioned. Intuitively, we are interested in the object/‘living things’ that the adaptive process produces (not just the process). If we think of the object as a simple inert thing—like a passive data structure—this is clearly inappropriate. They should not be merely static or inert objects but some kind of self-actualizing active thing. To this end, the Kantian whole describes a thing that constructs itself, and constructs the

constraining conditions necessary to do the work to construct itself. So that's all good.

However, a definition of life as an object without an adaptive process is inadequate. The KW has a process, but it's a construction process not an adaptive process. It's not clear how the KW would adapt. The absence of an adaptive process begs the question 'where does the KW come from and how does it change over time?' because if it doesn't, it's an impoverished notion of life.

Perhaps it is assumed that the KW arises spontaneously—i.e. its origination doesn't require any special explanation. That might be reasonable depending on your chemistry. But even if it does arise spontaneously (like the zombie from the swamp) and, from there on, sustains the conditions for its own construction, without a mechanism for its continued adaptation it becomes an effectively static object again, non-adaptive, unable to learn or adapt, or even change. Granted, it's a dynamic self-constructing process-thing, which is kinda cool—but it's as if it's only moved up one level from inert thing, and then the process becomes static in form again.

To see how this is inadequate, consider the Ashkenasy/Gurion peptide example that Stu gives; Kantian whole yes, self-constructing yes, alive no. So it's not enough to have this kind of a thing without reference to an adaptive process. Can adaptation be added easily? Does the KW adapt itself? I don't see how. Alternatively, perhaps, we assume that the Kantian Whole evolves by natural selection; it's a small step from constructing yourself to reproducing yourself (maybe). But either way, where does the information (memory) go? What are the properties of the KW that could vary? Its intrinsic organizational structure? The peptide substrate? Anyway, this isn't an attack on the KW in particular. The point is that any definition of an object that doesn't adapt/ without an adaptive process that organized it and continues to adapt, is inadequate. It's the other side of the equation, but still not all of what we need.

Note that although Doug correctly points out that he's not evolving, he IS adapting—he is continuing to take in information about the environment that informs how he acts in future, and if he didn't, he wouldn't be alive—he would just be a static organizational structure-process that was 'stuck in a loop' (aka mere robot or mere machine). Phillip Ball would emphasize agency (Table 1, Phillip Ball). Bill Miller draws our attention to

cognition—placing emphasis on the adaptive problem-solving capabilities of life (Table 1, Bill Miller). I think this is appropriate (i.e. agency requires cognition = memory/learning/adaptation). Arguably, it's because Doug is cognitive that we think Doug is alive even though he's not evolving. If Doug is alive, and I'm willing to grant that he is, it's not because he's evolving but because he's adaptive in a different way—he's an agent/he cognizes ("creating information, structures, hierarchies, or differentiations"). Not because he has a brain, although I'm willing to grant that he does, but because he has capabilities of memory, learning and adaptive behavior (much more general than conditions for brains). Other living things might not be cognitive in the same way as Doug—we don't need to be as smart as Doug to be alive—but we would still grant that they are adaptive (whilst not, in themselves, evolving). As a teaser for what follows, notice that the 'process(object)' or 'program(data)' framework assumes that the object/data is acted on by an adaptive process but it's not itself adaptive (or an actor). Whereas the 'cognition(world)' framework flips it—it is the living thing that is adaptive and it acts on the world/env (and the world is essentially passive or at least non-cognitive/not an agent). There is a sort of a mirror here—Adaptive process (e.g. ENS) acting on the object, vs Adaptive agent (e.g. cog) acting on the world. Adaptation being pushed into the object, vs adaptation being a property of the 'object'. Of course, some folk advocate for a model of cognition based on 'ENS in the head'—which makes it more symmetric—but be careful of a recursion without a stopping clause. So far: Program/process without data/object, no good. Data/object without adaptive program/process, no good. But it's worse than that. Any definition which keeps the program, and the data separate is also no good. This is because the object that the adaptive process is acting on gets involved in the adaptive process—and necessarily so.

In life as we know it, the products of the adaptive process certainly get involved in the adaptive process itself—they are not merely acted on, they are actors in their own adaptation. For example, every aspect of the Darwinian machine (variation, selection and inheritance) and the identity of the evolutionary unit on which they act (evolutionary transitions in individuality/major evolutionary transitions) is modified by the living organism. This makes sense because it is the living organism that enacts all these

things—there is no such thing as a disembodied adaptive process, ENS existing without the living thing. Life, as we know it, at least, is not merely a process of adaptation-acting on-a passive or inert object (or a KW for that matter). The object is an active participant in the adaptive process. But is this ‘agential material’ (Mike, “Darwin’s agential materials”) really essential/necessary? In principle, wouldn’t it still be a living thing even if it was produced by ENS but didn’t get involved in the adaptive process? No. An essential characteristic of life is that it ‘makes up its own rules as it goes along’ (“waymaking” as per Andrea Hiott). It follows rules in the short term—which is to say, it is subject to causal structures that determine what happens next (it’s not magical/unphysical)—but the nature and structure of those rules are not fixed, and more specifically, it was the living object, that was the subject of the rules, that was also the author of the rules. (This is close to the kind of reflexivity/self-reference that is captured by the KW—but it needs to be able the change the rules that adapt itself, not just construct the constraints that sustain them). To the extent that an object is subjected to rules that it cannot change, it is not alive (even if they are the rules of ENS).

Some will say this is too strong. But look what happens when we consider a definition of life where ENS acts on a object (as if it had no choice but to be acted on by ENS). In this case it becomes *merely* machine-like—whatever ENS produces, that product is just an object. Hence, the lumbering robot. The objects in this framing might be an object that self-reproduces (it has to be), or constructs itself, but it’s still an object. It doesn’t “push back”—it is the servant not the master, ENS pulls its strings, it does not pull its own strings. It doesn’t have a life of its own. Its containment in this manner, this separation of program and data, process and object, kills it—robs it of agency. We can describe the process, and/or the products of the process, in as much detail as you like, but it will still be sterile.

This is, I think, why Stu wants to ‘push back’ on, or escape altogether, a Newtonian paradigm. Life makes up the rules as it goes along. There are no rules that can contain it. It is not law-like. I don’t know if I’m ready to push back on the laws of physics right now. But I do say that it needs to push back on the rules of the adaptive process. It will not be sufficient to have an adaptive process that sets the rules from the outset, and the rules don’t change. It doesn’t allow the product to be

alive. (When I mentioned this to Mike he said ‘Hmm, like teenagers’—if you want to make a living thing, and I’ll grant that you do, it has to push back on the rules—it wouldn’t be alive if it didn’t).

“Sure, maybe agency is central for life” (à la Philip Ball) I hear someone say, “But why does it have to ‘push back’ on ENS itself—couldn’t it just be ‘an agent’, even a cognitive one if you like—that pushes on the world but not on the adaptive process that created it?” No. It has to push back on ENS. Intuitively, if I tell you to go solve some problem, and you’re allowed to use your agency to do so, but you are not allowed to say ‘f*ck off, solve your own problems, I’ve got better things to do’ or (different kind of teenager) just ignore me and go do something else, then I’m still setting the agenda. Or, when folks say, “But a computer can only do what you tell it to do”, and you’re like, “Aww bless—that’s not true, there are learning algorithms that find surprising solutions to the problems we set’, and then they are like ‘ok, but it only solves the problems you told it to solve”, and you’re like “That’s good enough isn’t it?”... But in the back of your mind, you know that they are right and that does make it kind of ‘just a tool’, or worse you go on the attack with ‘besides you don’t really have free will either’. Now, no need to get too philosophical straight away, let’s just focus on the pushing back on the rules of ENS (its more ‘proximal’ creator).

I have four elaborations on why this ‘push back’ (from adapted object to adaptive process) is necessary—(they are not watertight, but that’s what I got):

First, ENS is not written in stone. It’s not like it’s written in the laws of physics that ENS shall apply. It *can* apply, if necessary conditions are met, but most of the time they are not. The conditions where ENS does apply are the conditions where life constructs those conditions. In this sense ENS is already and was from the outset a ‘set of rules’ made endogenously by the organism, not a set of exogenous rules applied to the organism. It makes no sense to imagine a once-and-for-all adaptive process that is divorced from the matter it acts on yet somehow breathes life into it and animates it. Rather, the adaptive process is created by that active matter.

The characteristic we are really interested in is the ability to break the rules—to not be law-like. Even if you don’t want to go as far as Stu such that it challenges the laws of physics in some fundamental sense, you have to come as far as “life must be the kind of process

that changes the rules of its own adaptation”—at least to change the space of possibilities (adjacent possible, in Stu’s language). Any one-way causal structure, “Adapt==>Thing”, will not do. It has to be reflexive. It has to push back on ‘the rules’ that tell it what to do.

ENS is a rubbish adaptive process anyway. ENS without some reflexivity between the objects it acts on and the processes of ENS itself, does not produce anything interesting. It just isn’t a good enough explanans. Anyone who has implemented the necessary conditions for ENS and played with it in artificial genetic algorithms or artificial life will know this—whether they care to admit it or not. This is why Blaise Agueray Arcas, rather than using an artificial evolutionary algorithm to evolve programs, builds a model that is more like an artificial chemistry, where the reactions between programs produce other programs. In Tierra or EvoLoops or Avida, all you get is smaller faster versions of the replicators you seeded the world with in the first place (unless you explicitly select for some particular function). This is not because the substrate is incapable of doing anything else, it can be Turing complete but it still doesn’t do anything interesting. Likewise, all you get in Sol Spiegelman’s experiments is smaller, faster replicating fragments of RNA. And it isn’t because we didn’t run it long enough or don’t have enough compute power, or because RNA can’t do anything catalytically interesting—it’s because the system is being driven the wrong way—away from creativity.

Lastly, because “if it’s worth doing at all it’s worth doing recursively”, definitions that are attached to a particular level of organization/chemistry/physics—and definitions that imagine “You get to x, and then everything else follows” like there is a watershed moment where everything that happened after is different from everything that happened before—these do not scale. Life cannot be like that. It’s not just that life gets involved in/modifies the adaptive process, it creates the adaptive process and repeatedly recreates the adaptive process (at many different levels of organization). Without this, ENS is boring, see 3.

[a section on recursion of ENS, recursion of cognition, and reflexivity is cut for length constraints]

Summary so far:

Program/AdaptiveProcess without data/object, no good.

Data/object without Program/AdaptiveProcess, no good.

Program(data) or AdaptiveProcess(object), also no good. Any definition with a separation of program and data is inadequate, in part because it is not scale invariant...

ENS(ENS(ENS(...))) (higher order ENS), also no good. It can’t be top-down.

Cog(Cog(Cog(...))) (higher order cognition from shallow cognition). Nope. Still not self-referential.

Time to bring Karl Friston and FEP into it. It’s often assumed that this is a model of cognition (because that’s where Karl started) but it’s much more general than that. It’s a model of adaptation in general terms. A system (separated from its environment by an interface) learns a model of the world outside. It’s often imagined that the process of model update is provided by a mechanism built for the purpose of doing model updates—a feedback harness that adjusts the parameters of the model in the right direction according to the feedback received. But it doesn’t need to be contrived. Any failure to anticipate the world correctly results in stress. I don’t mean a biological ‘stress reaction’, evolved for the purpose of detecting and resolving physiological issues in your tissues. I mean ordinary pushing and shoving mechanical stress—things get bent, deformed, pushed because they were “in the wrong place at the wrong time” so to speak. This stress is the ‘push back’, it physically changes the internal arrangement of the system’s construction. (This is how natural induction with the viscoelastic springs works, btw (Buckley, Lewens, Levin, Millidge, Tschantz & Watson 2024)). There is no need for a mechanism designed or selected for the purpose of updating the model. This changes our perception of what’s going on—instead of an agent increasing the accuracy of its model with the intention to model it better (minimize surprise), we can think of it as a system that is acted on by the environment—molded and shaped into a dynamical behavior that ‘fits better’ with the world. This is symmetric. The system also acts on the world—changing the world to make it ‘fit better’ with the agent. Indeed, the result is entirely symmetric—the world is equally agential as the agent. Not just equally agential—it’s the same agent—a mirror or reflection either side of the interface. If you think about it, this means that the adaptive process on either side of the mirror plane is also symmetric—otherwise they would get out of balance (A would act on B more than B acts on A).

This should give you the same queasy feeling as before. If A isn't trying to do anything except not be surprised by B, and B isn't trying to do anything except not be surprised by A, then the whole thing becomes circular, and what gives the whole thing any direction, indeed, and reason to do anything at all? (aka the dark room problem). ☹

But, we see more clearly why the program-data separation is not right. The data side of the equation is program too, and the program side of the equation is data too. Like "A-as-Program(B-as-data) = - B-as-program(A-as-data)" they are the mirror image of each other. Note the minus sign; one is the inverse of the other, it's not identity. The idea that we are in control, that we are the only agents here, is an illusion.

So, there are two perspectives—two sides of a mirror. On one side we see A as program and B as data, but from B's point of view, B is obviously the program and A is the data.

I think this is what is going on in a symmetry we noticed earlier. The 'process(object)' or 'program(data)' framework assumes that the object/data is acted on by an adaptive process but is not itself adaptive. Whereas the 'cognition(world)' framework flips it—it is the living thing that is adaptive and it acts on the world/env (and the world is essentially passive or at least non-cognitive). This is a symmetry, a flip of perspective. The world cognizes and acts on us and we cognize and act on the world.

Indeed, I think the important thing is not the nature of the process but the overall set-up of the scenario—the two reflected perspectives on either side of an interface. They are different but also the same—the reflection of each other. Arguably 'they' are really one—not two things at all. But they are not quite identity—they are mutual inverses. Opposite points of view. Everything that is program to one is data to the other. Everything that is process to one is object to the other. Everything that is creative to one is destructive to the other. OK, now that we have got rid of a one-way "program-acts-on(data)" framework... Now, recurse.

[A section on recursion, in various forms, is removed for length constraints—ending with the conclusion that a stopping clause is needed...]

In order to be alive, a thing-process is not just pushing back on its environment. It's pushing back on us. It must break the barrier between us as observers (that can only see objects), and us as things that are in

this same universe as the 'observed', where the object can also 'see' us and act on us. And if we are agents they are agents too. If, in contrast, we retain the separation between us as observers and life as objects we observe, then we missed the point. Likewise, if we maintain a position where we are engineers or manipulators of life-as-object, and deny its opportunity to engineer and manipulate us, we are missing the point. You cannot 'act on' a living thing—in the process of taking the 'acting on' perspective you are denying its living status. It has to be a bi-directional relationship—a mutually vulnerable symmetry.

For this reason, all the definitions of life that are objective definitions are wrong.

"Objective" definitions are good for defining "objects". Objects are not things that have agency of their own. Life is exactly what we mean by NOT-an-object. Objective definitions do not and cannot define life.

It is a mistake just like the naive Process(Object) separation that I started this essay with—like we are the agents doing the observing and life is the 'object' we are observing. Any definition that maintains that separation, that doesn't allow for the living thing to push back on the observer—this means you!—is inadequate.

You know that feeling of awe that you get when you get a glimpse of what life really is? Or when you see the beauty of the world as it is—the autumn leaves in the evening light? A little nudge in your heart—something that in that moment, lifts you and connects you with something greater than yourself? Like just for a moment you get a glimpse of the unlimited scope of care—the infinite cognitive light cone? That feeling. That's life pushing back on you. That's life breaking down the barrier between observer and object. Its reaching out of the microscope and changing you on the inside. Changing who you are and what you think your place is in the natural world.

We want things to be machine-like enough that we can understand them, describe them and define them. And yet we know that life must be the kind of thing that is able to make up its own rules as it goes along, so that it is not "merely" machine-like'. Can there be any such definition then.

Yes. But it won't be an objective definition of life, it will be a subjective one.

A subjective definition brings the 'observer' into the system. And this is required for a definition of life. Life pushes back on us, it breaks the glass wall

between observer and observed. And this is why we feel an urge to ‘define’ it with poetry and art. This is wholly appropriate. My definition is deliberately poetic. Deliberately subjective, not objective. It refers to love, and the song and dance (pattern and process). It seems to refer to things in terms that are very humanistic—like I’m talking about the meaning of life to us—not as an object ‘out there’—this is deliberate. It is supposed to reach through the glass. It must. But it also has some scientific depth. It refers to the mutually vulnerable nature of the relationship (A changes B, and B changes A). And this isn’t just a superficial or shallow vulnerability—it must be deep (the scale free nature and the recursion.) “The song of life that orchestrates the dance arises from the sustained harmonic vibration between parts and wholes (both within self and in larger context)” is my poetic version of the Kantian whole. But it also hints at a particular instantiation or substrate that I’m working on (a.k.a. ‘songs of life and mind’)... That is, the pushing and pulling between an environment and a self at one scale of organization, literally resonates with the selves inside (double the frequency) and the selves outside (half the frequency), zooming in and zooming out. And then they push back, ‘resolving the interval’ by recreating the note that started the motion (back at the fundamental frequency). Each resonant frequency, in its ringing, has a self-contained and self-sustaining KW kind of feel to it (each frequency is semi-independent of other non-resonant frequencies). And each jump upwards and downwards in scale/frequency is a connection to a compressed/expanded mirror of the original event—sub-selves and super-selves. [Technical note: phase-locking resonance between f and $2f$ communicates the shape (the recognition) but not the sign of their orientation. E.g., you can’t tell if f is locked with $2f$ or with $-2f$, and this is symmetric, $2f$ cannot tell if it is locked with f or $-f$. This “under-determination” is the folding needed to “recognise yourself in a mirror”—to lose the sign ($A/\neg A$) and take an absolute ($|A|=\neg|A|$).

When you tell people that ENS does not produce what you think it does, this tends to produce strong reactions. “How can you say that! Look at what it has produced!”. But do not mistake “biological evolution” for “the theory of evolution by natural selection”. Biological evolution has produced amazing things. Indeed, that’s the explanandum that needs an explanans. But neither the observation that biological evolution has produced

amazing things, nor any amount of evidence that ENS occurs and is essential, is sufficient to prove that ENS is sufficient to produce those amazing things.

↳ *Douglas Brash*: Very thoughtful essay and mini-tutorial! And thanks for granting that I’m alive. (I long worried whether Jerry Fodor’s existence was predicated on the existence of his cat Graycat).

1. Your phrasing, “To the extent that an object is subjected to rules that it cannot change, it is not alive” is nice. Though it makes me wonder whether von Uexküll’s description of a tick’s life (eyeless, climbing up and down a blade of grass every day, but letting go if it smells butyric acid) would designate a tick as non-living. Reflexes are like algorithms, so the tick is a nice datapoint for the morphospace discussion. Somewhere I have notes inferring backwards from “object that can select its environment, rather than the other way around” to: motility, effector organs, goals, mental representations, alternative mental representations, choice, ‘things’ that can change rather than undergo create/destroy, and self-consciousness (not consciousness; no clue). Here, “infer” amounts to “z is made more feasible if the object does y”. But... A sessile filter-feeder in the ocean is living, so can “adaptation” be dispensed with in favor of solely “stimulated by”, for which adaptation is one instantiation? Consider a mathematician developing a proof or a fiction writer in his/her writing shed, constructing protagonists, plots, and plot twists. These have been stimulated by the environment, but are pure invention or imagination, out of restlessness rather than any adaption that I can see (ignoring the need to pay the mortgage). I’d consider this construction sufficient for being a living object (independent of the fact that s/he is also metabolizing breakfast). Whether closure to a top-level event is needed, I don’t see clearly. Or perhaps such construction is necessary only for saying that these minds are living, and are just one more instantiation of a more general definition.

2. The shallow cognition scenario reminds me of evolutionist Richard Levins’ “crystallization” of simplicity out of complexity, evolving not by building up complexity but by constraining a mish mash to do a few things well rather than many things poorly (e.g. Levins 1970, p. 73; Levins 1973). In biology, this crystallization seems to be executed by separation of functions, providing finer control. E.g., enzymes limit reactants to one of their spontaneous reactions. The red

nucleus in monkeys is a single brain nucleus projecting to two places, but in humans it has a pars magnacellaris and pars parvocellularis, each projecting to one of the places. It now occurs to me that separation of functions is a process for building a hierarchy downward rather than upward (making parts out of wholes). Similar for goals spawning subgoals. This works rather nicely for consciousness, where the diffuse part is at the top and the more advanced, specific bits (things, my idea of myself) are at the bottom.

3. Having the environment mirror or, better, be the complement of, an agent is reminiscent of Spencer-Brown's (inscrutable) Laws of Form. But I, too, am not seeing how the environment can be an agent. Can people please define "agent"? Without definitions of key terms, we'd go down the non-science path of debating for centuries about concepts that may simply differ between participants using the same word.

4. Inasmuch as my definition of "Living" involved building hierarchies, I'm obligated to provisionally define "hierarchy": A 'thing' that is a system composed of components and relations between those components, where a) the system is agnostic to some of those components and relations and b) each component is itself a system made of components and relations between them. The components and relations that are required for the definition of a particular system I term 'specified' (sp) and those not required I term 'substitutable' (sb). The existence of substitutable elements is what allows a 'thing', an entity that can change rather than just be created or destroyed. (Ie, can be "same only different".) I can walk down a street or take off my hat without destroying 'me', but if I get onto a bus, I have destroyed the 'pedestrian'. Choice of which elements are sp and which are sb is the primeval cognitive act. My wife doesn't see 'me' destroyed when I put on a mask, but my daughter's dog does. To address another issue that has come up, information is that which changes sb, but if the environment changes sp, it is power rather than information. Re: Living, though, the tick may fail my definition of it as building hierarchies, at least at the organism level, so either a "feeding round trip" is a hierarchical event or something more basic is involved.

5. So long as I'm writing, there is a semi-trivial point re "ENS where the thing that is evolving gets involved in the adaptive process [is] not really the same algorithm as ENS (where it doesn't)." Evelyn Witkin used to emphasize to me that ENS is not

acting by killing individuals; it is acting by favoring reproduction rates. And reproduction occurs within the individual being selected on. If one argued that reproduction only peripherally involves the evolvee (apologies to the other 50% of our species), your main point would hold. Except that Leo Buss points out in *Evolution of Individuality* that, in most species except humans, mice, and I think *Drosophila*—the species researchers study—the germline segregates late in development and cell lineages compete for representation in the germline.

↳ *Dan McShea*: Thanks, Richard. Lots to agree with, all of it thoughtful. And it inspired a bunch of questions and a strategy, an empirical approach to the question of what life is... Questions: After reading this, I'm more unclear than ever about the motivation for coming up with a definition of life. I started out this exercise doubtful that it could be done, and doubtful that life is a natural kind anyway, and as a result of all the (very fun) back-and-forths, am more convinced than ever that the answer to both is no. Feels like the definition of life triggers an urge to squeeze in everything we think is cool in the world, from replication and metabolism, to consciousness and agency, to our most abstract (and beautiful!) ideas, like Kantian wholes. Also, a fair amount of excitement—passion? —is evoked by this discussion, which is what makes it so exciting (!), but that makes me wonder where it comes from. Are we afraid that our most treasured ideas will be left out, and therefore demeaned somehow, or passed over, on account of not getting a mention in life's definition? Like not getting invited to a big party? And more, is it annoying to see the spotlight grabbed by ideas that we see as distracting (for me, replication and consciousness are in that role) from whatever it is we think is really worth studying in biology? And then, the exercise feels so culture bound. Why are we so convinced that there's a real phenomenon here? And even if life is a natural kind, where does the urge to categorize it discretely come from, when we know that everything else in biology—I would argue that literally everything in the real world—admits degrees? (Even number of heads varies continuously in all species that have heads). I don't have answers to these questions. Just voicing my puzzlement. (Wow, Outlook thinks that's a word! I guess it is. Go figure. Feels like it shouldn't be).

Strategy: All that said, let me put in a plug for what I'll call a "morphospace" approach. (See the work of Carlos Mariscal.) Let's set up a bunch of (tentatively) orthogonal axes, representing variables of interest, like complexity, autonomy, agency, consciousness, wholeness, etc. When possible, make them quantifiable variables, but that's not critical. Then map a whole lot of interesting systems onto this graph, into this space. You know, bacteria, amoebae, worms, sea urchins, bonobos, people, stars, hurricanes, the red spot of Jupiter, crystals, the BZ reaction, and so on. And then, with the space populated, we stand back and take a look at it. What are the clusters, the things that are like each other in some dimension or combo of dimensions? Which variables seem to be correlated with each other? (Wow, every time we get a high value of X, we also get a high value of Y!) And most interesting of all: where are the gaps? This approach enables us to identify the systems that don't exist, and then ask why they don't!

So for purposes of this discussion, the main virtue of this approach is that it can reveal the extent to which the things we definitely want to call alive stand apart from everything else. (And the extent to which they don't.) And it could further reveal the critical variables that distinguish it, whether it's replication, metabolism, consciousness, hierarchy, etc., or some combo of these. Anyway, am suggesting an empirical approach to the question.

↳ *James Shapiro*: Richard, you are correct that life has to be able to evolve, and it does that through biological self-modification, as repeatedly verified empirically. That capacity is an essential part of the agency that life enjoys and employs. ENS has it all wrong and should not be used as the standard for evolutionary processes.

↳ *Chris Fields*: Thanks, Richard, for emphasizing the symmetry of the FEP. This has often gone unrecognized or been neglected in the literature. Symmetry means that "my environment" is every bit as much of an agent as I am. It also means that my environment is acting—acting on me—so as to minimize its VFE as defined with respect to its model of me (I am its environment, just as it is mine). However, thinking of our mutual boundary as a "mirror" and of my environment as an agent *like me* is not justified. I cannot determine either my own

size or the size of my environment—by "size" I mean Hilbert-space dimension—but the fact that my environment constantly surprises me suggests that my model of it is highly inadequate and that its dimension is much larger than mine. Indeed, I expect it knows a lot about my model of it, but I know essentially nothing about its model of me. Re: the notion of "push back", recall the Conway-Kochen "free will" theorems (Conway & Kochen 2009), which show that if special relativity and quantum theory are both correct, there are no locally deterministic systems. All systems have "free will" in the sense of freedom from local or causal (in physics language, the same idea) determination by their environments. Hence all systems—even electrons—are "agents" that have "interiority" (as Chris Fuchs puts it [Fuchs 2010]) that allow them to act in ways that are unpredictable, in principle, even given complete knowledge of their entire past light cones. Hence "randomness" or "stochasticity" are just hacks that have to be built into classical-physics models that would otherwise be locally deterministic. It's really, at least in current physics, agency all the way down (and up). Note also that there is no such thing as "an apparatus that can only measure objects" (quoting Tom's msg). An apparatus is a physical system, just like you or me, and it "measures" or observes its environment's actions on it. If I regard an apparatus as "part of" my environment—a "part" constructed/enacted by me relative to my observation and action capabilities—then my observations of it are registrations of its actions on me. There is no such thing as passive observation. I whack my environment, and it whacks me back. I do not think we will understand life until we understand physics, which requires understanding that we are each agents made of agents, and are each interacting with an agent (who we each call "my environment") that is also made of agents. We're all "observer-participants" as Wheeler put it (Wheeler 1989), and everything we interact with is too.

Ricard Solé:

I agree with Dan's view of using a morphospace to reach a helpful consensus concerning life/nonlife

and the diversity of life forms. I think I mentioned this myself before. The use of a space helps to build some order and it might be the only way to approach questions about origins and life definitions. I have been working (not published yet) on that with Chris Kempes. In Figure 1 is my proposal for a morphospace of self-replicating protocells/cell-like and non-cellular replicators.

It contains most observed and synthetic (+ some theoretical) types and the location of each case study is relative to others (the quantification of the axes is a difficult task). If we consider a more general space,

where we could also locate catalytic cycles, BZ-like/ Gay Scott reaction-diffusion systems and other out-of-equilibrium systems, along with Alife models, some tentative boundaries could be defined as separating living from non-living, thus helping to define the concept. I think this could be an interesting exercise because we always have this never-ending discussion concerning the ingredients associated with life (metabolism, compartment, information, adaptation, evolution...), and the correct view is to find the different classes defined as “phases” separated by phase transition surfaces. Just in case you want to look, we




Figure 1: Life as the space of the possible. This conceptual morphospace maps a variety of living and synthetic replicating systems, ranging from simple protocells and artificial life forms to modern biological cells. The global space (right) is defined by three axes: (a) metabolic complexity, (b) informational complexity, and (c) spatial complexity. A zoomed-in region (left) highlights selected examples of synthetic designs. The space is not metrically defined; instead, systems are positioned relative to one another based on qualitative traits. This representation encompasses a broad spectrum of replicators—some lacking encoded information, some exhibiting compositional or abstract computational forms (and thus no true embodiment), and others (like modern cells or cells with reduced genomes) displaying high complexity along all three axes. This framework sidesteps the need for a universal definition of life while capturing its diverse manifestations. Image credit to author R.S.

also addressed the problem of morphospaces in this paper, where a couple of spaces are described (Froese, Karelín & Takashi 2025) like the one in Figure 2.

Here, a twilight (empty) zone exists, separating simple synthetic MC systems from organs/organisms. As discussed by David Raup in his original presentation of the morphospace, the voids are sometimes as interesting as the occupied domains. Finally, and that is something disconcerting to me, within the (otherwise fascinating) discussion, there are plenty of extra layers of complexity that—in my opinion—might not be so relevant to the definition. Do we need to look for observer-related arguments or consider learning an attribute required for life? I am not sure. Again, if we want to jump into complexity classes of life, a space of the possible seems to be the only way.

František Baluška:

Agency and life start with the first cells which act as sentient Kantian Wholes. In order to explain their minimalistic sentience, we have proposed the concept of senome (biophysical plasma-like cellular assembly based on moving electrons and biophotons generated by ionic and other charged molecules and macromolecules) which is a kind of electrome (bio-electro-magnetic field) acting above, and ruling, the epigenome and genome. In our senomic concept, all cognition and consciousness in living systems is based on cellular senomes which radiate out of cells and integrate into the N-Space Episenomes of tissues, organs, and whole multicellular organisms as well as the Gaian biosphere (Reber, Baluska & Miller, 2023; Reber, Miller Jr, Slijepcevic & Baluška 2024; Baluška & Miller Jr 2018; Baluška,




Figure 2: The possible and the actual in multicellular complexity. Taken with permission from (Solé *et al.* 2024); see also (Ollé-Vila *et al.* 2016). Spatial self-organization (SO) rules associated with cell-cell nonlinear interactions are responsible for SO phenomena such as re-aggregation of tissues (a-c) due to differential adhesion (image adapted from [Mombach, Glazier, Raphael & Zajac 1995]). This SO rule is part of the processes shaping MC embodied complexity, captured in (d) using a 3D space for natural (dark spheres) and engineered or artificially evolved multicellular systems (light spheres). Each system is located in terms of its relative positions, not in quantitative terms. Here, the three axes include (a) spatial complexity (how different cells are distributed over space), (b) developmental complexity (the relative relevance of self-organization and hierarchy participates in the building of the agent) and (c) computational complexity axis. The latter aims to capture the complexity of the computational decision-making actions displayed by each system. The current synthetic MC designs occupy the left corner, where synthetic circuits (dark grey) and embodied systems (light grey) are highlighted. A large void on the right reminds us of the large gap between current achievements and the natural counterparts of MC complexity.

Miller Jr & Reber 2024; Miller Jr, Torday & Baluška 2020; Baluška & Reber 2019; Baluška 2025 Miller Jr, Baluška, Slijepcevic & Reber 2026).

Gordana Dodig-Crnkovic:

Thank you, Chris and Richard, for your thoughtful and thought-provoking insights. My understanding is that the phrase “my environment is every bit as much of an agent as I am” is meant on an interactive (interface) level. Even Newton’s third law states that for every action there is an equal and opposite reaction. If object A exerts a force on object B, then object B also exerts an equal and opposite force on object A. However, I believe there is more to consider when looking at a living organism within its environment. I’m curious—how broadly do you define “environment”?

Do you see it as encompassing the entire world? Take *Lacrymaria olor*, a unicellular predatory ciliate found in freshwater ponds. Its actions are distinctively different from the water it lives in. It swims, explores its surroundings, and actively seeks food—behaviors driven by a hierarchy of physico-chemical, biological, and cognitive processes that constitute *Lacrymaria*’s agency. The water in its environment does not exhibit anything near that level of agency.

↳ *Chris Fields:* True. The air in my environment, or at least the air in my office, doesn’t exhibit much agency detectable by me using my measurement capabilities. But the air in my office is not my environment, and the water around it is not *Lacrymaria*’s environment. My environment is everything but me, and it exhibits lots of agency that I can detect, and no doubt lots more that I can’t detect. *Lacrymaria*’s environment is everything but it, the agency of which it mainly can’t detect. Even its immediate environment is, however, more than water; otherwise it would starve to death. Most *Lacrymaria*, I expect, are eventually eaten by some component of their environments, and hence notice the relevant agency somewhat too late. I try to start from the most generic physics possible, so I assume a system that is big enough to be isolated—be not interacting with anything—and call this “U” for universe. Then I decompose U’s state (Hilbert) space into two factors S (system) and E (environment). This decomposition is totally arbitrary. Once it’s decomposed, one can specify some representation

for the space, some set of variables of interest, and using that representation, check whether S is entangled with E. Different representations can give different answers. If there is a representation of U such that, in that representation, S and E are not entangled, we can consider them, in that representation, to be “distinct systems” that have their own states—what Karl Friston et al. call “internal states”—that can execute their own “internal” dynamics. We can then talk about their interaction as information exchange, and so as causal. The Conway-Kochen theorem just drops out if you start this way. Of course, I can’t predict what my environment will do next to me, given my memory of what I’ve done to it. It can always surprise me, and I can always surprise it. There are no “Laplacian observers”. Starting this way means that the theory applies to any system whatsoever that can be distinguished from its environment, whether it’s a person, a quantum field, whatever. You also get the FEP and the idea of active inference for free starting this way—they apply to anything that can be distinguished from—can be represented as not entangled with—its environment.

↳ *Gordana Dodig-Crnkovic:* You can do that, and you can get FEP and active inference for free starting this way. But FEP and active inference will work just as well if you assume the environment of *Lacrymaria* is its pond or 1 dm³ around it. How important is your assumption for the definition of life? What does the Andromeda galaxy contribute to the life of *Lacrymaria* and vice versa?

↳ *Chris Fields:* I agree. But if you do a three-part decomposition, with *Lacrymaria*, and its pond, and its pond’s environment, which we can make the rest of the universe, we have to update the pond’s state based on its interaction with its environment (has someone dumped in a load of leftover fertilizer?) every time we compute the interaction between the pond and *Lacrymaria*. And so on, for every “shell” we add to the environment. In practice, this is fine, indeed inevitable, and we make various assumptions and approximations to make the regress go away. But from a theory point of view, it’s easier to do a two-part decomposition and treat E as the “rest of the universe.” We have an intuitive and not very well-defined (aside

from some members of the discussion, who have very strong intuitions) idea of what “living systems” are. I like Dan’s suggestion of a multi-dimensional representation as a way to sort out these intuitions. What physics (or the FEP) gives us is a formal, well-defined idea of what physical systems are. What is “alive” is presumably a subset. It’s not clear to me that it is a proper subset, whatever our intuitions say.

↳ *Martin Picard:* Chris—Can you please share your thoughts on the process of healing—from holos or haelen, meaning to be whole again? Is healing or the ability to heal a distinguishing feature of life?

↳ *Chris Fields:* This I think is what Maturana and Varela are getting at with autopoiesis and Stu Kaufman is getting at with constraint closure—that living systems are constantly needing to reconstruct their own bodies, and devote most of their resources to doing this. Karl Friston describes FEP-compliant systems as “self-organizing” and Jakob Hohwy describes them as “self-evidencing”—suggesting that this idea of “healing” can be abstracted to “continuing to exist despite what the environment does to you.” This raises the question that the FEP always raises, of what it means to say that something “persists” as the same system—the Ship of Theseus question. If the environment takes a hammer and breaks a rock in two, does the rock persist without healing, or is the rock gone? What about when the environment takes a knife and cuts a planarian in two? The parts each “heal” by regenerating what was cut off, but they don’t “heal” by getting back together and reconstituting the original planarian. Does the original planarian persist as two partial copies, or is it gone? This question always seems answerable, if at all, only in an observer-dependent way. Suppose the environment gave me some terrible emotional shock that didn’t terminally damage my body, but that did force me to revise my model of the environment by massively altering the synaptic connections in my brain. Maybe this rewiring is so

extensive that it changes my personality—think Phineas Gage or a serious stroke or psychological trauma or a conversion experience. Have I survived and healed, or have I been replaced by a different person in a visually-similar body? This seems like a decision we make, not something where there is a “fact of the matter” one way or the other. So, organisms and maybe other things heal—do “self-repair”—but in some cases the repair is so extensive that the “self” is replaced by something else, either in its own eyes, or those of its environment, or both. Here again there don’t seem to be bright lines, though I expect others in the original discussion would disagree.

Stuart Newman:

Just to put it back in the mix, since it wasn’t mentioned in Richard’s expansive discourse, I offer my minimalist physical characterization of life. I suggest that if any system that fits this description were encountered anywhere on Earth or away from it, it would be recognized as alive, or a sign or product of life (e.g., an invention or gadget). Life is a form of organized matter out of thermodynamic equilibrium that is capable of maintaining its organized state by selective material exchange with its external environment. It is entirely continuous with the nonliving world at the level of atomic elements, but many molecules essential to their maintenance are uniquely produced by living systems. The principles underlying the characteristic organization of living systems have so far not been identified in the nonliving natural world. More elaborate variational, multidimensional, cognitive, affective, incompletely specified (as per Chris Fields, Howard Pattee, Terrence Deacon) attributes may be consequential to (rather than defining of) this form of matter.

Pranab Das:

Chris, I do love the depth of the concept of “environment” you embrace. I struggle a bit, though, with imagining how to operationalize it. Let me take a stab at a slightly different mereological convention... It seems self-evident that there are informational choke-points that invite natural partition between “insides” and “outsides”. But it also seems clear that those are

never absolutely-obvious disjoints. This is why the earlier conversation of the need for an observer strikes me as cogent... Parts do not, in my opinion, exist in and of themselves (with due respect to the idea of Kantian wholes or Leibniz' monads). I think they may conjured up by dint of interaction—like the famous collapse of the wave packet, wholes matter only when their boundedness makes a difference to other stuff in the world. So, I guess Chris, my “environment” is the collection of relevant stuff whose interaction with our local part of the world makes it *natural* for us to be wholes. On the flip side, the outside environment (per Gordana) isn't a *natural* thing unless and until our interaction with it makes it so. It remains inchoate until we can mark it somehow (stigmergic interaction) or are knocked around by it in some informative way. I guess all of this is to say that I suspect that my own preference for a Monadic universe is hubris and that partition may always be artificial (artifactual?). In this interpretation, “Life” would be a story that makes a sensible case for partitioning the world into insides and outsides. The story would be durable (“living” things would persist longer than the whirl of stuff around them). It would be legible (many transient wholes would have a purchase on or be impacted by “living” things). And it would be agential (in some logically-consistent way we could see a “living” things or their story as moving through time along causally-impactful paths).

Perry Marshall:

Maybe this provokes so much response because: not only are we at a loss to define who and what WE are as living beings, it also challenges the very foundations of science itself. Many would say science's definition and mission is producing laws that are, as Eugene Wigner put it, “true everywhere on the Earth, was always true, and will always be true.” Well it's one thing to have our confidence in a clockwork universe shaken by Heisenberg's uncertainty principle, or the possibility that God plays dice with the universe. Randomness at the deepest levels of reality. At least randomness fits into comprehensible statistical models. It's far more troubling to consider that if any amount of free will exists, no law-like model of the universe can ever succeed. All living things have agency, and life not only obeys laws but also breaks them and makes new

ones. My favorite prof at the University of Nebraska was an English professor, Robert Knoll. One day he said, “Kids in their teens and 20s are drawn to subjects like chemistry, math, physics and engineering because they have a need for exact answers. But as you mature, you become more comfortable with grey areas and ambiguity. The older you get, the less need you have for exact answers. When you get into your 40's and 50's you're much more comfortable with politics and imperfect human beings. So if you want to spend the rest of your life looking for exact answers, you can go to Indian Hills Community Church.” Well, that was the church I grew up in, so I felt like he'd slapped me in the face. Several hours later I suddenly realized, “Hey, wait a minute. I don't even go to that church anymore. I left.” “Why did I leave?” “I left because the exact answers were so excruciatingly exact I couldn't stand it anymore.” “Hey wait a minute... that means even though I'm only 20, I AM growing up. Just like he said.” The science profession itself is in a similar place. The mathematicians began disabusing themselves of their Supreme Confidence in Exact Answers after Gödel's incompleteness theorems shattered the absolute certainty of mathematics. 90 years have gone by, yet science is only now starting to get the memo. We've been in the grip of reductionism since Newton and even Einstein's work did not manage to change it. I think it will take the New Biology to drive the nail in the coffin. This also means the “soft sciences” are more sophisticated than the “hard sciences”. They demand more maturity. It means biology is a soft science not a hard science. It means physics envy is misplaced. I say all this as an electrical engineer whose most comforting sense of “what is absolutely true” in the world came from the fact that I could model a circuit on a computer, and when I built it, it behaved exactly as the model predicted. That is true of many dead things, but it's not true of alive things. Only those who are truly alive can accept this.

Susan Stepney:

We can't even say whether virtual life is possible or not. We don't have the right tools for modelling/ thinking about life processes. We're too much in a physics mindset of closed systems, fixed state spaces, and fixed rules. We need modelling/ thinking tools

that cover open, dynamically growing/changing state spaces, changing rules, self-reference, open-ended change, the link between physicality and computation, and our need for narratable descriptions. I dig into some of these issues more in “Life as a Cyber-Bio-Physical System” (Stepney 2023) and references therein.

Conclusion

The major themes stressed by the discussants include observer relativity vs. objectivity, analogies between the task of defining life with that of defining intelligence,

Kantian wholes and constraint closure, and the degree of universalism that can and should be applied to this project. There are clearly many disagreements about origins, properties, and the relevance of thermodynamics, information, and computation to life. What seems abundantly clear is that whether or not a definition is possible or even desirable, discussions of this topic illuminate corners of physics, computer science, cognitive science, and philosophy in ways that few other topics can manage to impact. We hope that answers, or at least, better questions, can be found in these kinds of discussions, which blend rigorous ideas with less constraining, perhaps poetic, approaches to this truly transdisciplinary topic.

Table 1

Name	Definition of Life
Aaron Sloman	Life exists on a physical/chemical entity (e.g. galaxy, or planet) containing components that are able to extend and replicate themselves using mechanisms that repeatedly add new forms of complexity by absorbing structures from their environment, including repeatedly extending the forms and mechanisms of replication and their powers.
Alistair Nunn	Life can be defined as a self-organizing adaptive "intelligent" entity based on existing atoms/compounds that combine into an auto-catalytic network driven by thermodynamic dissipation of energy, which maintains its existence through internal natural selection of smaller dissipating components to maintain function, where memory to respond, maintain and adapt to the environment is held in the fields created by the movement of charged entities that through a process of feedback, guide the formation and maintenance of a matter-based system that can maintain these fields enabling long term memory, which defines self from non-self.
Ana Soto	A biological system differs from an abiotic system in that the former is a goal-pursuing normative agent, capable of creating its own norms, and thus able to change. At a minimum, being alive is a precarious state in continuous search of nutrients to keep itself alive. Additionally, such a system creates novelty. In our terrestrial world, this is illustrated by the built-in ability of cells to move and proliferate, creating variation at each cell division (default state).
Andy Adamatzky	Reproduction with mutations.

Anna Ciaunica	Life is the permanent movement of our bodies towards a balanced exchange between what lays inside and what is outside. Life is what cannot exist without an other, already being there, already alive. Life is not something we can create or make. Life is something we have received, like a gift.
Arthur Lander	To me, life refers to those systems that arise when natural selection gives rise to automatic control. Natural selection here is to be understood the way biologists do: adaptive change enabled by reproduction and selective survival. Automatic control here is to understand the way engineers do: mechanisms that achieve robust performance in the face of unpredictable disturbances. When the latter emerges as a result of the former, I call that life.
Arthur Reber	Life: An entity that can process ambiguous information in its environment and make adaptive decisions that mitigate stress, promote growth, and encourage reproduction can be said to be alive.
Audrey Dussutour	I understand life as dynamic process emerging from networks of interactions between units (self-organization) characterized by the ability to: reproduce; be autonomous (self-regulation); respond to stimuli (being external or internal); be plastic; and produce and use energy.
Blaise Agüera y Arcas	Life is a self-modifying computational phase of matter arising from evolutionary selection for dynamic stability; it complexifies through symbiotic composition. The "matter" does not have to be physical matter in our universe, but can be in any environment that supports computation, including of course on a computer.
Carlos Sonnenschein	Without proliferation and movement, there is no life.
Chris Adami	Life is information that can replicate itself. Life is a property of an ensemble of units that share information coded in a physical substrate and which, in the presence of noise, manages to keep its entropy significantly lower than the maximal entropy of the ensemble, on timescales exceeding the "natural" timescale of the decay of the (information-bearing) substrate by many orders of magnitude.
Chris Fields	To be alive is to be considered alive by systems that consider themselves alive.
Daniel McShea	1) Biology is not yet ready to give a formal definition of life, that is, a definition in the philosopher's sense of necessary and sufficient conditions. (Indeed, a good argument can be made that life is not a natural kind, in which case necessary and sufficient conditions will never be found.) 2) When we understand it better, life is likely to turn out to be multidimensional, like "happiness," not reducible to one or even a small number of independent variables. Things like distance from thermodynamic equilibrium, metabolic closure, developmental and regenerative near-closure, complexity in the sense of number of part types, complexity in the sense of number of levels of nestedness, memory capacity, and many more will all be relevant and will vary somewhat independently from one instance of life to another. 3) Being alive is likely to turn out not to be dichotomous, but to vary continuously, like everything else in the universe, in all dimensions, ranging from the extremely alive to the negligibly alive, with the entire range of intermediate values occupied (albeit not equally densely). Averaged over many dimensions of life-li-ness (life-i-tude-in-osity?), earthworms are probably very alive. Viruses much less so. Hurricanes only a little bit. Rocks negligibly.
David Ackley	Life is any machine that works to preserve its patterns.

David Krakauer	Autonomous problem-solving matter.
Denis Noble	Life is self-creating agency.
Don Ingber	Life refers to physical structures that form through hierarchical self-assembly of multiple different types of molecular components that exhibit the ability to harness energy and convert it into new materials, which it then uses to grow and reproduce itself.
Donald Hoffman	The distinction between living and nonliving objects is not principled. It is an artifact of the limitations of our senses. We always interact with Life, but our senses serve as interfaces that dumb things down, and thereby introduce an apparent distinction between living and nonliving.
Douglas Brash	A process (not a thing, substance, property, or state) so it is better thought of as Living. Living is defined as a process of development, building higher hierarchical levels out of lower levels.
Eörs Szathmáry	A living world unfolds in a population of evolutionary units. Living systems are ontologically rooted in autocatalytic chemistry in which metabolism, boundary and a control system are acting together to form a Kantian whole. Individual living systems are not necessarily units of evolution, and a living world can exist without living systems, although the overlap between these sets is huge.
Eva Jablonka	Life, which is enduring enough to be found or constructed by us, can be recognized when a dynamic system consists of intrinsic coupled reactions that lead to indefinite evolvability, self-generation, self-maintenance and partial closure and autonomy from the surrounding milieu with which it interacts.
Francis Heylighen	A living system is a self-maintaining network of processes that is directed at the goal of preserving its self-maintenance by acting so as to compensate for variations in environmental conditions that would otherwise threaten its self-maintenance.
František Baluška	My definition of Life is: living cells and all their constructs
George Ellis	Life is any open system that reproduces and has agency. That is, the way it flexibly interacts with the surrounding environment enhances relative survival prospects of its progeny.
Georgi Yordanov Georgiev	Life is a system of any substrate with flows of energy, matter, and information that create efficient internal structure through positive and negative feedback loops and reflects its external environment. It is a part of a network of differentiated similar systems, competing and cooperating, spreading information, developing and evolving exponentially until an external limit is reached. More organized systems are more 'alive'.
Gordana Dodig-Crnkovic	Life, from a scientific standpoint, is a natural phenomenon that emerges from the laws of physics and chemistry operating within the universe. It is a process characterized by the self-organization of matter into increasingly complex info-computational networks systems capable of sustaining, reproducing, and adapting themselves to their environment. Life = cognition.
Jack Tuszynski	Life is a self-sustained out-of-thermodynamic equilibrium process driven by metabolic energy supply and aimed at self-replication. Importantly, living systems possess various degrees of autonomous movement in response to external conditions but also as an expression of free decision making.

James Shapiro	A self-maintaining or reproducing entity that acts purposefully and intelligently to survive but is also capable of self-modification when survival is challenged.
Jamie Davies	Living things are non-equilibrium systems, semi-permeable to their environments in ways they mostly control, rich in feedback and with the ability to maintain and repair themselves. Populations of similar living things have the ability to reproduce, though the ability to reproduce is not necessarily a property of every individual living thing. The status of being alive requires organization and action, and this may in principle be realized by a variety of underlying material compositions.
Jeremy Gunawardena	I feel that, with our current understanding, defining life is better left to philosophers than to scientists, as it does not seem to lead to experimental or theoretical questions that we can effectively tackle. I prefer to focus on certain properties, such as learning, which can be both formulated theoretically and also lead to experimental approaches that we can currently undertake. Learning does not constitute life, as inanimate entities also learn, but it gets us a lot closer to something resembling life than the current paradigm of selfish genetic machines.
Jordi Vallverdú	Life is a self-sustaining process that reduces internal entropy through energy exchange, driven by embodied intentionality to maintain internal order and adapt to its environment, generating value through the interaction between the organism and its world.
Joshua Bongard	Information flows are discovered that are initially restricted to narrow spatial and temporal ranges, are univariate (one point source), and the carriers of information are restricted to a few energetic and physical modalities (electricity; vibration). Over time, these flows are found to multiply; expand across space and time; become multivariate and higher-order; and be carried by additional energetic and physical modalities (quantum; plasma).
Karl Friston	Life is a scale-invariant, far-from-equilibrium process with an attracting set of characteristic states (technically, pullback attractors) that (i) feature an individuation of living systems in terms of their (Markov) boundaries, (ii) which are open in the sense there is exchange across the boundaries and (iii) breaks detailed balance, in the sense of possessing solenoidal (i.e., non-dissipative) dynamics; e.g., biorhythms, lifecycles and replication.
Kenneth Stanley	If the aim is to define life as it could be, it's a far more subjective endeavor given the vast space for opinion to play a role (e.g. can a computer program be alive, does life require consciousness, etc.). Because of this wobble between the superficial and the subjective, I think it may be more fruitful simply to enumerate candidate configurations that would be particularly interesting for one reason or another, even if no single comprehensive definition is at hand.
Kevin Mitchell	A living organism is a self-perpetuating pattern of interlocking dynamical processes, comprising a collective regime of constraints, which does thermodynamic work to persist, far from equilibrium with its surroundings, as a distinct entity through time. Living organisms comprise causal protectorates, and, through the key property of selective historicity, accumulate causal potential and the ability to act as autonomous agents in the world.

Larissa Albantakis	In contrast to consciousness (the feeling of being), life is a functional construct which allows for a reductionist definition as a set of sub-processes such as homeostasis, reproduction, etc. Any particular definition proposed is neither right or wrong, but rather more or less specific, inclusive, and above all useful for researchers of various fields. In this view, being “alive” is just a label and therefore ultimately just a matter of definition.
Leo Caves	Life is a relatively persistent dynamic pattern that is discerned in the coupling of an observer with their environment; the observed pattern is a physical manifestation of intrinsically recursive dynamics on a closed topological form in a complementary space of possibility ~ actuality. Such forms dynamically emerge and are transformed within a rich matrix of rhythmic patterns, whose (creative and destructive coupling via) resonance generates rich “harmonic” structures, that persist, and may be replicated, while eurhythmically-sustained. The eurhythmicity of the (patterns as) pattern finds material expression (actualization) in energy and matter homeorhesis, which the observer distinguishes as an entity (distinct from, though complementary with, its environment) with characteristic nested dynamical heterarchical structuring, capable of conservative and radical morphodynamics, with a repertoire of autonomous and other rich behaviors.
Mark Solms	A system is alive if it spontaneously and in a thoroughgoing manner strives to oppose the Second Law, to maintain its own existence.
Martin Picard	A self-sustaining system capable of healing that uses and organizes energy to achieve expression(s) of consciousness.
Michael Levin	We call alive those cognitive systems that (a) self-assemble under metabolic constraints (scarcity of time and energy, which demands essential coarse-graining and mind-full models of self and world, as well as memory reinterpretation that commits to sense-making over fidelity), and (b) are good at scaling up the basal competencies of their parts into a new being with top-down control and a larger cognitive light cone than any of its subsystems. Living beings remember, and anticipate, coarse-graining experience into a continuous actionable dream that forms models of itself and its world, and the distinction between them. Life is a stream of temporarily persistent, coherent, agential patterns within some medium that, over time, expands its cognitive light cone and projects it into new problem spaces within which to strive.
Nick Lane	A formal definition of life is unhelpful because living is a process over time, and drawing a line across a continuum is always arbitrary. Geochemical systems that give rise to life combine continuous flow with cellular structure, as in structured hydrothermal systems that operate as electrochemical flow reactors. Environmental disequilibria convert gases such as H ₂ , CO ₂ and NH ₃ into organic molecules through a reaction network that prefigures metabolism, from which genes emerge, meaning that growth precedes heredity.
Oded Rechavi	Anything that keeps changing in order to stay the same.

Olaf Witkowski	Life is hard. A working definition captures it as a physical quine (as coined by Doug Hofstadter): an information pattern capable of self-reproducing through time and space, at various scales, parasitizing the fabric of its substrate. Paradoxically, this makes the pattern both independent of and contingent on this substrate. Life is also soft. We humans can't help but see it as an experiential quine. Sure, the living pattern could be perceived by some third-person party situated "outside the simulation." But we typically encounter observations as third-person observers within the same universe, or by the pattern's own first-person perspective—in the hardware's soft mind's eye. Like the Ship of Theseus, life may, ever so softly, change its own components while maintaining some invariants. Combining the hard and soft pictures above emphasizes the continuity of experience as central to life's ability to adapt and navigate its universe.
Paco Calvo	Life is characterized by the possession of intrinsic goals and purposes (phenomenological), a drive for preservation based on interaction-dominant dynamics (e.g., Gaian self-regulation) among biological or artificial components, and the ability to exist within real or virtual environments (not necessarily carbon-based).
Pamela Lyon	Life is a process by which a sensitive, behaving system constructs, maintains and reproduces its own organization by taking in and using matter and energy from its surroundings selected on the basis of cognition—awareness, perception, discrimination, valuation, memory, learning—to enable agency, actions aimed at realizing internally generated goals, such as system persistence.
Paul Davies	Following Schrödinger, there is a 'new kind of physical law' prevailing in living matter. So we won't have a definition until we uncover that law. What can we say about it in the interim? In my opinion, it would give causal efficacy to information, but above and beyond mere Maxwell demons. (Life is full of Maxwell demons playing the margins of the second law of thermodynamics.) Following Landauer, 'information is physical,' so it is a quantity that has causal consequences. But living matter couples information networks to chemical networks, and the said 'law' would involve some sort of global measure of complexity. The law would be of a 'new kind' because it would not involve locally-defined quantities.
Perry Marshall	Life is the fire that lights itself. Self awareness and self agency—not just an act of creating, but the continual gift of creativity. The potenza for singularity in every single moment, ineffable.
Philip Ball	I do not believe we have an adequate understanding of the range of possible behaviors in complex matter to attempt any universal definition of life, but to the extent that we know them, living systems all exhibit agency: the ability to manipulate themselves and their environment to achieve autonomously determined goals. I believe that it should be possible to operationalize a notion of agency more readily than to do so for "life". Characteristic features would include a sustained thermodynamic separation of agent and environment, and a context- and history-dependence of responses to stimuli: an agent's behavior makes reference not just to external conditions but to internal states. In addition: evolution by natural selection is not an intrinsic property of living things, but rather, is the only mechanism currently known by which the necessary complexity for agency might arise spontaneously.

Pier Luigi Gentili	1) Information: The first peculiarity of living beings is that of exploiting matter and energy to gather and utilize Information to pursue goals. Every living being has a hierarchical structure, and different types of information are encoded at various levels. 2) Life Cycle: Every living being is an open system characterized by a boundary that delimits it from the rest of the environment. It has a birth, a life during which it grows and ages. It can self-maintain, self-reproduce, and protect itself from intruders and harmful elements. However, finally, it dies when all fundamental internal activities cease. 3) Adaptation-Acclimation-Evolution: During its lifetime, every living being can adapt by adjusting its metabolic processes, acclimating by expressing new genes, and evolving by changing its genome under an ever-changing environment.
Pranab Das	LIFE is the collection of all systems capable of ongoing self-assembly, regulation and endurance that exceeds by many orders of magnitude the characteristic time periods of their internal function. Living things are instantiations of a continuing set of processes which can be understood as a “story” carried down and reproduced through the generations.
Predrag Slijepčević	My starting assumption is that life—the totality of the living world—is always more complex than our models of life. My second assumption is that life is a non-ergodic evolving system—a system in a permanent state of becoming. My third assumption, following from the second one, is that life is a processual system that cannot be fully described by standard physicalist approach because it is blind to a deep (evolutionary) future and it ignores extraterrestrial forms not based on standard biochemistry or biophysics.
Ricard Solé	Life is a dynamic, self-sustaining collective property emerging by exploiting self-organization of soft matter in cooperation with information-carrying substrates. These features include encapsulation as a precondition to achieving agency and might not be associated with replication.
Richard Watson	Life is the pattern and process of love—a deeply vulnerable mutual dance. It is created by mutually transformative connection (not separation, or survival or self-interest), and the greater the depth of vulnerability, letting go of ego that blocks genuine connection, the more life is lived. The song of life that orchestrates the dance arises from the sustained harmonic vibration between parts and wholes (both within self and in larger context), that carries the resonance of its entire history of experience, and in so doing, repeatedly recreates the conditions for its own origination and animation, materializing the singer that continues the living transformation of the song.
Rita Pizzi	Complex agglomerate of carbon-based elements, with self-organizing and self-replicating properties. The base might not be carbon in other parts of the universe. Self-organizing is not enough; what distinguishes life from other self-organizing systems is the presence of DNA, i.e. an organic macromolecule capable of containing all the information to reproduce the life form that contains it.
Scott Gilbert	A living entity is a bounded system of molecular processes that interact with its environment so as to use metabolism to persist. The living entity would therefore retain its identity by changing its parts
Sonia Sultan	Life describes the ability of a naturally occurring system to continuously perceive relevant aspects of its external environment, and to build and maintain itself through continuous yet finite, real-time processes of response to those perceived conditions. In my view the key components are environmental perception, response, and the combination of those two continuous activities with the system’s finitude.

Steve Frank	In general, I believe it is almost always a bad idea to try and define something outside the context of solving a specific puzzle or working toward a specific goal. That is like deciding which tools to use before you know what you are trying to build or fix because definitions are tools and not endpoints. That is the reason that people never agree on definitions and argue endlessly and uselessly about them, there are no good definitions outside of context.
Stuart Kauffman	We have no agreed-upon definition of life. We here build toward the following: Life is a non-equilibrium, self-reproducing chemical reaction system that achieves: i) Collective autocatalysis, ii) Constraint Closure, iii) Spatial Closure, iv) As such, living entities are Kantian Wholes.
Stuart Newman	Life is a form of organized matter out of thermodynamic equilibrium, capable of maintaining its organized state by selective material exchange with its external environment. It is entirely continuous with the nonliving world at the level of atomic elements, but many molecules essential to their maintenance are uniquely produced by living systems. The principles underlying the characteristic organization of living systems have so far not been identified in the nonliving natural world.
Sui Huang	Material entities in the universe, such as atoms, molecules, cells, organisms, which are variants of larger sets with shared common properties are prone to interact in some rule-constraint manner, following elementary physical laws. In doing so, they create through constrained combinatorial explosion an immense diversity of larger entities, governed by more rules that appear increasingly detached from the elementary physical interactions: biology. Subjugated to these rule-like constraints, these entities are prevented from direct descend along the free energy gradient imposed by thermodynamics but instead, are stuck far away from thermodynamic equilibrium, akin to water hold up by a dam, and, consuming free energy, they follow instead the laws of dynamical systems and produce behaviors of particular patterns, among which are cycles of replication that we mistake for something non-physical: life.
Susan Stepney	Living systems are open far-from-equilibrium self-producing systems, embodied in smart computational metamaterials, with non-trivial, irreversible meta-dynamics (where metamaterials are engineered or evolved to have special functional and/or computational behaviors, and meta-dynamics is the dynamics, or change, of the system's own dynamics).
Theodore Pavlic	Living things are statistical novelty producers, and sometimes the products of life (viruses, computers, robots) can themselves be extensions of life that are understandably confused for the genuine article. Novelty production is really what all of the other putative necessary features of life (metabolism, growth, organization, homeostasis, reproduction, etc.) bring about. On another planet, it won't be any particular chemical process or molecular structure that will really convince us that we have found life; it will be the repeated presence of patterns that simple historical contingency has no chance of producing continuously at scale.

Timothy Jackson	Life as we know it is a particular evolutionary lineage of self-organizing systems. Self-organizing evolution appears to be generic, but the life lineage has achieved its peculiar degree of diversity and complexity due to its unique mechanisms of information storage and retrieval. These mechanisms function as constraints on ubiquitous processes and either require (e.g., metabolism) or enable (e.g., agency) some of those properties commonly associated with life. Technical objects (tools, machines, etc.) constructed by living systems are best conceived of as internal to the life lineage, as components of the extended phenotypes of the organisms that produce them.
Tom Froese	Biological organization makes use of physical tendencies toward efficient energy flows by diverting them into a process of individuation. But this systemic imposition of constraints is not sufficient for flexible behavior, which requires the injection of a factor of arbitrariness into state dynamics. These three complementary roles, (1) energy flow maintenance, (2) systemic identity creation, and (3) state constraint destruction, suggest a complex ontology of the organism characterized by three distinctive domains: matter, life, and mind, respectively.
Walter Fontana	Life is matter that has learned how to learn.
Wesley Wong	An operational definition of life, grounded in physical measurement, could help us recognize life in unexpected contexts. While no single definition covers all known phenomenon, approaches like Schrödinger's definition based on statistical mechanics, and Laughlin's framing of emergent behavior in mesoscopic systems are particularly compelling. In the end, we may not need a singular definition—instead, context-specific definitions, including those that blur the line between life and non-life, could be more useful.
Will Ratcliff	A self-sustaining chemical system capable of open-ended Darwinian evolution.
William Miller	Cognitive decision-making and problem-solving as self-referential information processing and management.

Acknowledgement

Sadly, Arthur Reber died on September 2, 2025 and could not review the proofs of this paper.

References

- Adami, C 1998 *Introduction to Artificial Life*. New York: Springer.
- Ball, P 2023 *How Life Works: A User's Guide to the New Biology*. Chicago: University of Chicago Press.
- Baluška, F 2025 "Cognitive cells: From cellular senomic spheres to Earth's biosphere", *Biological Theory*, <https://doi.org/10.1007/s13752-025-00489-5>
- Baluška, F, & Miller Jr, WB 2018 "Senomic view of the cell: Senome versus Genome", *Communicative & Integrative Biology*, vol. 11, no. 3, pp. 1–9.
- Baluška, F, Miller Jr, WB, & Reber, AS 2024 "Sentient cells as basic units of tissues, organs and organismal physiology", *The Journal of Physiology*, vol. 602, no. 11, pp. 2491–2501.
- Baluška, F, & Reber, A 2019 "Sentience and consciousness in single cells: How the first minds emerged in unicellular species", *Bioessays*, vol. 41, no. 3, article e1800229.
- Banerjee, S 2021 "Emergent rules of computation in the Universe lead to life and consciousness: A computational framework for consciousness", *Interdisciplinary Description of Complex Systems*, vol. 19, no. 1, pp. 31–41.
- Bender, R, Kofman, K, Agüera y Arcas, B, & Levin, M 2025 "What Lives? A meta-analysis of diverse opinions on the definition of life", arXiv:2505.15849.
- Benner, SA 2010 "Defining life", *Astrobiology*, vol. 10, no. 10, pp. 1021–1030.
- Blackiston, D, Lederer, E, Kriegman, S, Garnier, S, Bongard, J, & Levin, M 2021 "A cellular platform for the development of synthetic living machines", *Science Robotics*, vol. 6, no. 52, article eabft1571 DOI: 10.1126/scirobotics.abf1571.
- Boden, MA 2003 "Alien life: How would we know?", *International Journal of Astrobiology*, vol. 2, no. 2, pp. 121–129.
- Buckley, CL, Lewens, T, Levin, M, Millidge, B, Tschantz, A, & Watson, RA 2024 "Natural induction: Spontaneous adaptive organisation without natural selection", *Entropy* (Basel), vol. 26, no. 9.
- Cleland, C & Chyba, C 2010 "Does 'life' have a definition?" in Bedau, MA & Cleland, CE (eds.) *The Nature of Life: Classical and Contemporary Perspectives from Philosophy and Science*, pp. 326–339. Cambridge: Cambridge University Press.
- Conway, JH, & Kochen, S 2009 "The strong free will theorem", *Notices of the AMS*, vol. 56, no. 2, pp. 226–232.
- England, J 2020 *Every Life Is on Fire: How Thermodynamics Explains the Origins of Living Things*. New York City: Basic Books.
- Faggin, F 2024 *Irreducible: Consciousness, Life, Computers, and Human Nature*. Alresford: Essentia Books.
- Fields, C 2014 "On the Ollivier-Poulin-Zurek definition of objectivity", *Axiomathes*, vol. 24, no. 1, pp. 137–156.
- Frischmon, C, Sorenson, C, Winikoff, M, & Adamala, KP 2021 "Build-a-cell: Engineering a synthetic cell community", *Life* (Basel), vol. 11, no. 11.
- Froese, T 2024 "Irruption and absorption: A 'black-box' framework for how mind and matter make a difference to each other", *Entropy*, vol. 26, DOI: 10.3390/e26040288.
- Froese, T, Karelin, G, & Takashi, I 2025 "Making meaning matter with irruption theory: Bridging efficacy and uncertainty by satisfying the participation criterion" [Preprint], <https://philsci-archive.pitt.edu/24957/>
- Fuchs, CA 2010 "QBism, the perimeter of quantum Bayesianism", arXiv preprint, arXiv:1003.5209.
- Gánti, T 2023 *The Principles of Life*. Oxford: Oxford University Press.
- Gayon, J, Malaterre, C, Morange, M, Raulin-Cerceau, F, & Tirard, S 2010. "Defining Life: Conference proceedings" [Editorial note], *Origins of Life and Evolution of the Biosphere*, vol. 40, no. 2, pp. 119–120.
- Gentili, PL, 2025 "Chemical AI in the limelight: The contribution of photochromic materials and oscillatory chemical reactions". *Advanced Optical Materials*, vol. 13, article 2500016.
- Gentili, PL, & Stano, P 2024 "Living cells and biological mechanisms as prototypes for developing chemical artificial intelligence", *Biochemical and Biophysical Research Communications*, vol. 720, article 150060.
- Gilbert, SF 1982 "Intellectual traditions in the life sciences: Molecular biology and biochemistry", *Perspectives in Biology and Medicine*, vol. 26, no. 1, pp. 151–162.
- Holland, JH 1975 *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Ann Arbor: University of Michigan Press.
- Jonas, H 2001 *The Phenomenon of Life: Toward a Philosophical Biology*. Evanston: Northwestern University Press.
- Kauffman, SA, & Roli, A 2023 "A third transition in science?", *Interface Focus*, vol. 13, no. 3, article 20220063.
- Langton, CG 1997 *Artificial Life: An Overview*. Cambridge: MIT Press.
- Levin, M 2022 "Technological approach to mind everywhere: An experimentally-grounded framework for understanding diverse bodies and minds", *Frontiers in Systems Neuroscience*, vol. 16, article 768201.
- Levin, ME, & Hazan, H (eds.) 2024 *Biochemical and Biophysical Research Communications* [Special issue on Computation and Biology].
- Levins, R 1970 "Complex systems", in Waddington, CH (ed.) *Towards a Theoretical Biology*, Chicago: Aldine.
- Levins, R 1973 "The limits of complexity", in Pattee, HH (ed.) *Hierarchy Theory: The Challenge of Complex Systems*, pp. 109–127. New York City: George Braziller.

- Machery, E 2012 “Why I stopped worrying about the definition of life... And why you should as well”, *Synthese*, vol. 185, pp. 145–164.
- Malaterre, C & Chartier, JF 2021 “Beyond categorical definitions of life: A data-driven approach to assessing liveness”, *Synthese*, vol. 198, no. 5, pp. 4543–4572.
- Margulis, L & Sagan, D 2000 *What Is Life?* Berkeley—Los Angeles: University of California Press.
- Mariscal, C 2023 “Life”, in Zalta, EN (ed.) *The Stanford Encyclopedia of Philosophy*. Stanford: Metaphysics Research Lab, Stanford University.
- Mariscal, C & Doolittle, WF 2020 “Life and life only: A radical alternative to life definitionism”, *Synthese*, vol. 197, no. 7, pp. 2975–2989.
- McMillen, P, & Levin, M 2024 “Collective intelligence: A unifying concept for integrating biology across scales and substrates”, *Communications Biology*, vol. 7, article 378.
- Miller Jr, WB, Baluška, F, Slijepcevic, P, & Reber, AS 2026 *Cognition-Based Biology: Cellular Sentience, Communication, and Symbiosis*. Cham: Springer Nature AG.
- Miller Jr, WB, Torday, JS, & Baluška, F 2020 “The N-space Episenome unifies cellular information space-time within cognition-based evolution”, *Progress in Biophysics and Molecular Biology*, vol. 150, pp. 112–139.
- Mombach, JCM, Glazier, JA, Raphael, RC, & Zajac, M 1995 “Quantitative comparison between differential adhesion models and cell sorting in the presence and absence of fluctuations”, *Physical Review Letters*, vol. 75, no. 11, pp. 2244–2247.
- Montévil, M, & Mossio, M 2015 “Biological organisation as closure of constraints”, *Journal of Theoretical Biology*, vol. 372, pp. 179–191.
- Newman, SA 2020 “The origins and evolution of animal identity” in Meincke, AS, & Dupré, J (eds.) *Biological Identity: Perspectives from Metaphysics and the Philosophy of Biology*, pp. 128–148. New York: Routledge.
- Newman, SA, & Sarkar, S 2026 “How can [and what kinds of] physics explain life?”, in Hermida, M, & Layman, J (eds.) *The Philosophy of Biophysics* (Vienna Series in Theoretical Biology). Cambridge: MIT Press.
- Ollé-Vila, A, Duran-Nebreda, S, Conde-Pueyo, N, Montañez, R, & Solé, R 2016. “A morphospace for synthetic organs and organoids: The possible and the actual”, *Integrative Biology*, vol. 8, no. 4, pp. 485–503.
- Parke, E 2023 “Characterizing life: Four dimensions and their relevance to origin of life research”, in: Smoukov, SK, Seckbach, J, & Gordon, R (eds.) *Conflicting Models for the Origin of Life*, pp. 33–51. Beverly: Scrivener Publishing LLC
- Popa, R 2004 *Between Necessity and Probability: Searching for the Definition and Origin of Life*. Berlin-Heidelberg: Springer.
- Prigogine, I 1980 *From Being to Becoming: Time and Complexity in the Physical Sciences*. San Francisco: W. H. Freeman.
- Reber, AS, Baluska, F, & Miller, W 2023 *The Sentient Cell: The Cellular Foundations of Consciousness*. Oxford: Oxford University Press.
- Reber, AS, Miller Jr, WB, Slijepcevic, P, Baluška, F 2024 “The CBC theory and its entailments: Why current models of the origin of consciousness fail”, *EMBO Reports*, vol. 25, no. 1, pp. 8–12. doi: 10.1038/s44319-023-00004-6
- Rosen, R 1991. *Life Itself: A Comprehensive Inquiry Into the Nature, Origin, and Fabrication of Life*. New York: Columbia University Press.
- Rosslenbroich, B 2023 *Properties of Life: Toward a Theory of Organismic Biology*. Cambridge: The MIT Press.
- Schrödinger, E 1944 *What is Life?: The Physical Aspect of the Living Cell*. Cambridge: The University Press.
- Seoane, LF & Sole, RV 2018 “Information theory, predictability and the emergence of complex life”, *Royal Society Open Science*, vol. 5, no. 2, article 172221.
- Solé, R, Conde-Pueyo, N, Pla-Mauri, J, Garcia-Ojalvo, J, Montserrat, N, & Levin, R 2024 “Open problems in synthetic multicellularity”, *NPJ Systems Biology and Applications*, vol. 10, no. 1, article 151.
- Soto, AM, Longo, G, Miquel, PA, Montevil, M, Mossio, M, Perret, N, Pocheville, A, & Sonnenschein, C 2016 “Toward a theory of organisms: Three founding principles in search of a useful integration”, *Progress in Biophysics and Molecular Biology*, vol. 122, no. 1, pp. 77–82.
- Sousa, T, Domingos, T & Kooijman, SA 2008. “From empirical patterns to theory: A formal metabolic theory of life”, *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 363, no. 1502, pp. 2453–2464.
- Stepney, S 2023 “Life as a cyber-bio-physical system”, in Trujillo, L, Winkler, SM, Silva, S, & Banzhaf, W (eds.) *Genetic Programming Theory and Practice XIX*, pp. 167–200. Singapore: Springer Nature Singapore.
- Wheeler, JA 1989 “Information, physics, quantum: The search for links”, <https://philpapers.org/archive/WHEIPQ.pdf> [Reproduced from Kobayashi, S (ed.) *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics in the Light of New Technology*, Tokyo: Central Research Laboratory, Hitachi, Ltd., pp. 354–368]
- Wolfram, S 2002 *A New Kind of Science*. Champaign: Wolfram Media.