

Epistolution: a new principle necessary to a learning-first theory of life

Charlie Munford 

Talking Octopus, Guilford, Vermont, USA

ABSTRACT

Biological theory assumes the organized appearance of life and the reliable recurrence of traits are due to inheritance. Natural selection acting on blind variations produces phenotypes with heritable traits, one of which may be natural learning. The aim of learning, then, is solving problems related to survival and reproduction. But what if these views confuse cause with effect? Perhaps a learning algorithm is required for any phenotype at all to arise. If so, evolution proceeds learning-first, with individuals pursuing another telos entirely. I argue that this aim may be epistemological, the drive to understand the world through an *umwelt*. By “understand” I mean neither association nor prediction but Karl Popper’s concept of explanation through conjecture and refutation. I propose that if only genetic materials are truly heritable, not traits, then testing a successful physical theory of life will depend on building abiotic machines which can perform natural learning without the presence of any inherited materials or conditions. I name this process “epistolution,” combining “epistemology” and “evolution,” to distinguish it from other concepts. Epistolution is an integral consequence of any learning-first view of life, such as the Cellular Basis of Consciousness theory. This type of theory suggests that in all cells during the history of life full-blown agency, involving beliefs, intentions, and desires, generated all the phenotypes that have then been winnowed by natural selection. Unlike in other versions, I posit that the aim of agential living systems is the explanation of reality rather than inductive prediction or survival/reproduction.

ARTICLE HISTORY

Received 9 April 2024
Revised 2 June 2024
Accepted 4 June 2024

KEYWORDS

Agency; cognition;
evolution; intelligence;
intention; Neo-Darwinism;
sentience

The purpose of this paper

Our current fundamental theories of biology do not intersect cleanly with the laws of physics. Genes-first explanations of life posit that naturally selected nucleotide sequences encode instructions for traits, but the physical rules by which these instructions may be decoded into individual organismic phenotypes are still unknown. As a result, the basic logic of physiology, and why biology appears to violate the second law of thermodynamics, is still unexplained. The goal of this paper is to outline the reasons I believe Karl Popper’s epistemology, a theory of *knowledge*, provides a path forward to unlocking this problem. A genes-first explanation of life assumes that we can derive a physical theory of biology, and therefore an explanation at the level of individual motivations, from the Shannon information encoded in the genome by working out the physical manner in which nucleotides control traits. I argue that if the assumption that nucleotides encode traits is false then will be impossible to develop a physical theory of life until this assumption is discarded, and a new point of view emerges.

Such a change would entail developing a genuinely testable theory of biological agency, a learning-first theory of life. This theory would entail an alternative telos for living beings. If all life aims at survival and reproduction, all roads will forever eventually lead back to the nucleotides. Scientific inquiry rests on a large foundation of philosophical assumptions. It is these assumptions that are my target in this paper. My view is that if we take the idea of all living organisms as subjective agents seriously, the consequences of this change in our perspective may enable the discovery of a universal natural learning algorithm that is present in all life forms. This universal biological algorithm, unlike evolved programs, would be substrate-independent and could be instantiated in artificial abiotic machines. The consequences of this view, if correct, would be that “intelligent” machines can be built that rather than mimicking humans by processing our existing human-created symbolic tokens, genuinely invent new knowledge in the way organisms do, by creative, agential attempts to understand the physical world.

CONTACT Charlie Munford  charlie.munford@gmail.com  Talking Octopus, 3976 Hinesburg Rd., Guilford, Vermont 05301, USA

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

What is knowledge?

Natural learning might be defined as an agent's increase in knowledge about a subject matter. This common-sense, vernacular definition would seem to be at odds with versions in the biological literature which almost universally avoid or downplay the use of the word "knowledge." Why is this? Despite the commonplace utility of the term, biological theory tends to exclude it. Some researchers might object to its use because it is a philosophical term, imprecise and difficult to test. Others would claim that knowledge is uniquely human, or that it is contained only in symbolic abstract representations. None of these claims are true, but significant difficulties arise because although the presence of knowledge has obvious consequences in the physical world, we currently do not have a good *physical* theory of knowledge. Testing any quality without a physical theory of its effects is problematic.

According to the adherents to the ideas of the twentieth century philosopher of science Karl Popper, a group called the "critical rationalists," knowledge consists of good explanations. In their view the aim of natural learning in humans is neither association of one event with another nor inductive prediction of the future, but the explanation of the past. Knowledge about the past allows for anticipation of events in the future that are quite unlike anything that has ever occurred before. This is very different from statistical association and prediction.

Nothing about the past logically entails that the future will resemble it. This is called the "problem of induction," and it was discovered by the philosopher David Hume [1]. Hume argued that it is illogical to assume that events in the past allow prediction of the future. His famous example was that by examining the swans of Europe, one might have developed the notion that all swans were white. Nothing within this experience would have prepared Europeans for the discovery of black swans in New Zealand. Indeed, if generalizing from past examples was the real basis of their knowledge of birds, it would have been impossible for them to recognize that a black bird could even *be* a swan. This problem was a major obstacle in the philosophy of science until Karl Popper developed his theory of knowledge in the 20th century [2]. Popper suggested that knowledge is formed not by logical induction but by a process of conjecture and refutation. The truth can never be finally established, but good explanations can be conjectured and then subjected to reasonable discussion and empirical testing so as to be logically *eliminated*. In this way science can approach the truth gradually, without ever arriving at a final destination.

A good explanation is parsimonious, logical, informative, and testable, a set of requirements that the physicist David Deutsch has summarized as "hard to vary" [3].

Notice that despite Popper's success in clarifying the scientific process, our current paradigm of artificial intelligence and machine learning is at odds with his theory. In these fields, learning is considered to be a matter of logical processing of data garnered from past events according to a fitness function, in other words, *induction*. This mistake is extremely confusing because in machine learning both salient information and precise instructions for its processing are supplied to the algorithm by the programmers. It is not a mistake to think of this as a technology that is highly useful and transformative, or that *humans* might gain knowledge by using it. It is a mistake to think that the algorithm *itself* is doing anything resembling natural learning.

In natural learning, and consequently in scientific inquiry, it is the operation of the algorithm itself that determines the salience of information. For a natural learner not all events are meaningful, but the few that do take on meaning result very quickly in a capable intelligence with only a sparse training set. That's the power of a good explanation; it "understands" things in the Popperian sense of a conjecture. The algorithm works efficiently because it only has to refute one of two rival explanations, and for this purpose even the smallest relevant difference can sometimes suffice. Facts become relevant only by virtue of their role within an explanation. Every fact in the world always supports a false theory except the few outstanding examples that refute it. If another theory agrees with *all* the facts, it wins. For a machine learning algorithm, on the other hand, information has no salience at all. It does not understand any of the data that it processes. Nothing about the data forms an explanation of the world from the point of view of the algorithm. It has no more point of view than an abacus, a slide rule or any other manmade tool for doing computations. Consequently, it requires enormous datasets to train these machine learning algorithms and they never really learn anything new, anything that wasn't implicitly present in the selection of the training data and the fitness function.

The technologist's habit of metaphorically naming part of a machine learning program an "agent" is highly misleading because biological agents like humans have individual perspectives formed by their explanations of the world. These agential perspectives allow them not only to process predefined information in a controlled environment but to find or create salience from

differences they encounter in their stochastic natural environments. The concept of a subjective point of view itself is intimately bound up with the idea of natural learning and of explanatory knowledge. Knowledge must not only be “about something,” but it must be “from somewhere.”

Popper’s theory is currently the best epistemology available, but it is still not a *physical* theory. Why is this? Because currently all conjectures come from organisms. New explanations arise only through biology. Life and knowledge are inextricably linked. Wherever one is present, the other is present as well. As far as we can tell, subjectivity only occurs in biological systems. Subjectivity, including agency, is an essential ingredient for the growth of knowledge, but we can’t yet link it to the known laws of physics.

Knowledge is not necessarily coded information

Knowledge is, in my view, Bateson information that has been assimilated. Gregory Bateson calls information “a difference that makes a difference” [4]. It’s hard to imagine that heritable materials like nucleotide sequences, whole genomes, cytoplasmic inheritance, or karyotype could make no difference in the development of cells, so by that test they all certainly carry Bateson information. But any influence at all carries this sort of information; by definition an influence makes a difference. Temperature, salinity, pH, gravity, the properties of water, and so forth are all influential, but they are not encoded.

Codes are only one form Bateson information can take. Codes are templated, Shannon information, which is a matter of digitization [5]. Claude Shannon discovered that interpreting symbols in a sequential message as templates, as interchangeable units for which the number of possible positional combinations could be quantified, resulted in the ability to compare a message systematically with a copy of that message [6]. This digitization of information can, when executed by electrical signaling apparatus, nearly eliminate errors in transmission. Digitization is only one way that Bateson information (differences with meaning to a subject) can be represented. It is a very important way because it allows for error correction and fidelity, which is the key to faithful inheritance. This was the insight that led directly to the digital age, but its unintended consequence has been to confuse theorists. Many have now come to conflate all *knowledge* with coded information, a mistake that elides the vitally important role of a subject, and of subjective meaning-making, in producing codes. All messages must have

a sender and a receiver that understands them. It conflates two very different things to say that some type of code must be the only source of biological information. This mistaken insistence flows directly from the assumption that knowledge can only enter biology through inherited blind variation followed by natural selection, a view that avoids the concept of subjects who make meaning from information.

It is simply not true that all the informative influences that result in living beings are inherited. Two cloned humans, sharing all their heritable materials and any processes those materials are responsible for causing, are nevertheless able to learn, for example, different spoken languages. Many other organisms can also learn, although there is currently a robust debate about how far down the tree of life this ability extends. Either way, at least one species is able to encode significances that were not present in inherited package in the germline cell. More fundamentally, the heritable Shannon information in a germline cell is compatible with all the cell fates that arise in differentiated tissues, therefore it cannot be the source of that differentiation. If the influences of the environment were encountered with only an inherited knowledge base with which to deal with them, the divergence of twins would be random; twins could not learn two different languages. Indeed, they could only speak a common language encoded in their DNA.

To be faithfully inherited, information must be in its coded, templated, Shannon form. It caused a great bolt of inspiration for the Neo-Darwinist explanation of life when it was discovered that the process of error correction in DNA was extensive and followed a similar logic as Shannon’s procedure with messaging templates. DNA is arranged in a double helix such that it can be split and copied, and faithful replication of a very long sequence of nucleotides can proceed with great accuracy by correcting errors in transmission. But the environmental influences on cells, which no one can deny make a difference, are not represented in digital code, and also not directly inherited. If they were not put in order somehow, they should send the complex system of the developing organism into complete disarray. How, then, are these other, non-Shannon, non-encoded influences turned into orderly signals during development? I propose that to account for the derivation of salience like this from a stochastic *umwelt*, in order to account for knowledge, we have to posit a universal learning mechanism.

How does knowledge get into biology?

Our best current biological theories are genes-first; they assume that knowledge enters biology through

evolutionary descent. In them the organized appearance of life and the reliable recurrence of traits are due to inheritance. Natural selection removes all phenotypes except ones with heritable traits which promote survival and reproduction. These phenotypes may feature natural learning, but the aim of learning in this case is confined to solving problems related to survival and reproduction. In this general mode of Neo-Darwinian explanation, what is inherited is not only a set of materials and physical conditions, but a set of instructions. This is why traits are considered to be heritable. A trait is a result not only of inherited materials and a niche, but also of a set of instructions for assembling, operating, updating, and reproducing these materials in a niche [7].

When translated into physical terms, this view suggests that knowledge emerges from natural selection acting on alternate genetic programs. Blind errors are introduced into the programs by chance, and through a long gradual process of natural selection the most functional programs have simply not been eliminated. This is how the highly improbable chemical and physical facts of biology are currently explained, facts that flagrantly violate the second law of thermodynamics. If this is your explanation for the presence of all knowledge in the world, then there is no reason to suspect that biological machines like organisms will be capable of anything other than random behaviors that passed the fitness function of survival and reproduction. Shannon information is the only source of fitness, therefore the only possible goal toward which life could be inclined. Thus, current difficulties have arisen with the concept of individual agency.

Natural selection is an explanation at the level of the trait and the population, not at the level of the individual. An individual can never experience nonsurvival or alternative rates of reproduction; these are events that happen in populations over many generations. There is no doubt that there is a form of knowledge embodied in the natural selection of nucleotides that are present in all organisms, but in between the nucleotide and the selection process stands the phenotype. If it were the case that knowledge could enter the phenotype and inform the trait, affecting the conditions of selection, then the whole genes-first explanation for life would be ruined. There would be an entirely unexplained, and vitally influential, process of learning happening at the individual level rather than at the population level. This is why the Weismann barrier was so important to the mid-twentieth century biological consensus called the Modern Synthesis [8]. The Weismann barrier insists that no information (Bateson)

may pass from the phenotype to the genotype, or else it spoils the entire epistemology of life [9,10].

Why is this possibility so problematic? Why can't selection simply favor phenotypes that develop gene-encoded traits that involve some learning, traits that allow them to improve their survival chances? Why can't selection for adaptation simply be recursive, allowing for phenotypes featuring learning to blindly emerge? This question has become a subject of great contention, but let me oversimplify the argument a bit by making a bold claim. Learning, by definition, is responding appropriately to something new, not responding in a blind, random fashion or responding to something that has been experienced many times before. Any response that is programmatically present in the genetic sequences is not a learned response. Any learning, in a coherent genes-first theory, is really only pseudo-learning.

Can knowledge enter through the phenotype?

Another possible type of explanation has only very recently been very tentatively advanced, a radically different sort of theory, a theory that includes individual agency [11–14]. This alternative asserts that knowledge enters biology through the phenotype, at the individual level. This form of explanation hints toward a learning-first theory of life which directly opposes the genes-first model that predominates today. In this learning-first alternative natural selection would still occur, but it would not select among genetic programs but rather among individuals. The heritable aspect affected by this selection would not include instructions but only vitally important nucleotide templates for making RNAs. Natural selection would merely be a selection of tools rather than instructions.

It is essential to recognize that this learning-first alternative does not neuter natural selection or genetic change. If this learning-first concept is true, knowledge that enters at the phenotype level could easily still get into the genes through natural selection. Imagine a giraffe who, learning that the leaves at the top of the tree are more nutritious, develops a habit of feeding only on high branches. This would create the selective conditions under which nucleotide sequences useful in developing long necks are favored. As a direct result genetic evolution would occur through blind mutation and selection, but it would incline toward a learned objective. Thus knowledge would enter biology through the phenotype but become cemented in part through genetic influence.

There is no easy way to eliminate this logical possibility by biological experiment. The road is blocked because of

the intertwining of nature and nurture, a well-known dilemma. The only aspects of organisms we can observe are either heritable materials or phenotypic traits. We cannot observe *umwelts*, the “world unto” an organism, the envelope of environmental conditions it exists within [15]. Presumably all observable traits must be a result of instructions, but all instructions arise from a combination of both genetic materials and an *umwelt*. Changing the genes affects the phenotype, but changing the environmental conditions also changes the phenotype. If environment is held constant, then genetic influences play a larger role in determining phenotype; if genes are held constant then environmental influences play a larger role. It is impossible to determine a precise ratio of the influence of genes vs. environment because although genes can be mapped, it is highly impractical to try to map all possible environments. So where are the instructions? Do they come from nature or from nurture? This debate has raged fruitlessly for decades because there is no conceivable *biological* experiment that can finally falsify either one. In biology, these two streams of causation are always combined. Only a *non-biological* experiment can decide the question because only abiotic materials are currently congruent with the known laws of physics.

The way to disprove a learning-first theory would be to develop a way to translate the messages that are purported to exist in the error-corrected Shannon information in the genome into physical instructions. This would be to develop a physical theory of biology. For this experiment we would have to build and operate artificial organisms from abiotic material that obey explicit coded instructions, instructions written by programmers. One interesting step in this direction has already been taken. Michael Levin’s lab at Tufts has begun to attempt writing bioelectric codes for heritable phenotypes in planaria, and they can control the inheritance of head type in this manner. They have yet to move this attempt beyond biotic materials [16]. Although Levin’s experiments are highly interesting, they in some ways prove the opposite of his code-oriented view of life [17]. Bioelectric coding of the sort he attempts is an example of environmental influence, not genetic influence, and so his results paradoxically tend to reinforce a learning-first view of life rather than demonstrate genetic programs. Levin is aware of this puzzle, but to my knowledge has not yet articulated a learning-first theory himself.

The straitjacket of inheritance

If the complex functions of organisms were shaped only according to blind chance followed by the eliminative process of natural selection, it would mean a sort of

inerrant obedience to fixed instructions. I call this paradox the “straitjacket of inheritance.” The evolutionary history of an organism is fixed, therefore all influences which derive from this history are also static. Each germline cell only contains one set of heritable materials, not one for every resulting cell fate and another for every possible cell behavior. The fact that heritable materials are utilized in the creation of these fates and behaviors does not necessarily mean they were coded for by them. Indeed, results from modern genetics research suggest they could *not* have been coded for. The human genome project was expected by genes-first theorists to explain how all human traits arose by decoding each of their corresponding genetic sequences. This project failed to show any such results. Genome wide association studies have now shown that the mapping from DNA sequences to traits is not at all a straightforward matter [18], there is both widespread epistasis and widespread pleiotropy. Many genes usually contribute to a trait and many traits are usually affected by a gene. This failure should have weakened the idea that the genome contains explicit digital codes for traits, and therefore undermined the genes-first theory of life.

This failure is more significant than many biologists may realize. The logical coherence of a genes-first theory requires that there be a map of physical rules by which a nucleotide sequence is translated directly into instructions for cells to follow. Without this map, biology and the laws of physics remain incompatible with one another. Today the genes-first alternative presupposes this map, but a mechanism has not been found that executes it. This means we still do not have a physical theory of biology. To explain not only the presence of genes but the presence of traits, we have to work out the physical rules that allow for the development of individuals. Heritable materials, of course, do influence the living phenotype somehow, this is not in dispute. But it is vital to recognize that the error-corrected material inheritance is in the form of templates for producing RNAs and proteins, not necessarily instructions for life. Traits are a result of the combination of heritable materials with instructions for their use. The question at hand is the origin of the instructions.

There is a basic intuitive problem with the genes-first conception of life – plasticity. If evolutionary inheritance were responsible for a definitive code for traits, this code would be subject to a continual narrowing of its scope of variation as complexity increased. Consider the fact that in our manmade devices the orderly operation of highly complex functions cannot tolerate even minute deviations. The very word

“machined” in common use means for an object to have its roughness and its plasticity removed to within very fine tolerances. In living systems, increasing complexity requires instead that the scope of variations becomes wider, encompassing more elaborate actions in more variable micro-conditions. Inheritance specifies rigidity, but complexity (in life) requires plasticity. In the logical operation of the CPU in a computer as it interprets instructions in a hard drive, even one bit out of place among billions may cause the program to crash. Life, on the other hand, takes place primarily in water. In liquid media, the stochasticity of Brownian motion ensures that a tremendous randomness is a feature of every complex set of molecular transformations. This randomness is fundamentally at odds with logical coding. If you were to give a computer program a simple set of “code” that was compatible with several trillion different behaviors and then turn it loose in a stochastic environment that triggered those behaviors unpredictably, why would you expect an orderly system to emerge? Such an expectation would be irrational. Computers and other complex machines work by virtue of detailed engineering that very carefully eliminates all this stochasticity.

If all the instructions for life were heritable, then life might have been doomed to remain a simple, invariant self-replicating system like an ice crystal, dependent on a constellation of very specific environmental conditions for its existence and, like such crystals, unable to project reproductions of its morphology into the future beyond these very narrow confines. If this were true then genetic material, instead of being a condensed structure confined to a particular region and deployed in a highly selective, flexible way during ontogeny, would be in effect the entire phenotype of the organism, again as it is for a crystal. Instead, life is quite unlike an inflexible crystalline solid. Life occurs largely in liquids, characterized by stochasticity and Brownian motion. The genome appears to be a specialized replicated library of templates for making RNA whose use is regulated, maintained, and reproduced by the entire cell [19].

The extended synthesis view makes confusion worse

Few biologists are still satisfied with the straightforward genes-first view presented in Richard Dawkins’ *The Selfish Gene*, the version where individuals have no agency but are rather “lumbering robots” controlled “body and mind” by their nucleotide sequences [7]. There have been recent attempts to harmonize learning and epigenetic inheritance with Neo-Darwinism. The

Extended Evolutionary Synthesis [20] is a thesis that acknowledges that organisms adaptively adjust their gene expression to fit their environmental conditions, that organisms influence their own conditions for selection, and even that direct learning is an alternate pathway for the inheritance of traits. But the EES version of evolution, although it includes observations from biology that conflict with Dawkins’ simple genetic determinism, does so at the expense of logical coherence. The EES does not go far enough, it is not a learning-first theory at all. It still presents living coherence as a consequence of inheritance. If life is a consequence only of inheritance and blind chance, it doesn’t solve the either the straitjacket of inheritance problem or the problem of induction.

In Dawkins’ view, individual-level behavior has no physical consequences; in that sense it doesn’t really exist. All we are seeing when we look at an organism is the physical effects of swarms of nucleotides. The EES would still hold that individual learning does occur, but it is in part a result of evolved genetic programs. In other words, the EES introduces individual-level causal programs but fails to discard gene-level programs. In the EES explanation, the individual appears to be a causal agent that affects physical results but only *sometimes*. This is not an acceptable claim in a physical theory of life because it makes the thesis untestable.

In my view, this EES version violates common sense intuitions about the nature of instructions. We cannot have it both ways. The real challenge in this question is explaining the eradication of stochasticity, of entropy, and the creation of harmonious living forms and functions that do not conflict with themselves. Life must follow orderly development and ontogeny in order to live. For this orderly flow of energy and materials to be sustained it requires a coherent set of instructions. It cannot be the case that life is ultimately controlled both by inherited genetic programs and also by the process of learning that creates another controlling set of programs. These sets of instructions would be in fundamental conflict. One set of instructions would be continually disrupting the coherence of the effects of the other set. Both sources of influence can be present, but one process must be the master and the other the slave. Either the phenomenon of learning is only an illusion and what we are actually seeing is preprogrammed genetic effects, or a learned program itself is truly what is shaping ontogeny, using inherited templates as mere tools for constructing phenotypes.

Once we look at it this way, we realize that there is a grand asymmetry between the two options. The genes-first alternative is hampered by its complete

insensitivity to local conditions. Natural selection is blind. By definition this means there is no contact between the subject matter learned by a phenotype and the evolutionary selective process. If there were such contact, then you could not even term what is happening “learning.” Explaining all aspects of life by reference to evolution by natural selection mirrors the process of logical induction that Hume refuted. A novel function cannot be produced by generalizing from past examples, this could only result in a modification of past functions in the direction of a more recent, but still past, example. Blind chance as the only source of novelty presents the problem of induction; no new forms could be fit for anything but past experience except by chance, just as the black swans couldn’t be identified as swans because they were not white. If phenotypes were modified by the local environment in a random chance type of manner this would solidify the genes-first theory, but they do not seem to be like this at all. It seems instead that they seek out function, even sometimes recovering, for example, from amputations.

Subjectivity and its role in a complete physical theory

The largest obstacle for the acceptance of a learning-first theory of life is its requirement that all living cells be capable of learning. If life really is organized by a universal natural learning algorithm, then it must be present in all life forms. This consequence of the theory runs counter to our traditional understanding of biology, where intelligence exists only in the highly complex brains of animals. It boggles the imagination to think that bacteria, for example, could be learners. This would require individual-level agential qualities like memory, beliefs, intentions, and desires to be present not only in other branches of the tree of life, but in all cells. On the other hand the more we look for these qualities, the more we seem to discover them. Not only corvids, cetaceans, pachyderms, and primates exhibit complex intelligence, but also cephalopods, which are mollusks, and honeybees, which are arthropods. Even plants and fungi have now been shown to have communicative abilities that surprise our intuitions [21,22]. Even further, the functions of prokaryotes appear, under deeper scrutiny, to exhibit much more complexity than previously thought, as the emerging field of basal cognition suggests [23]. The authors of the Cellular Basis of Consciousness theory have posited that consciousness is coterminous with life. They have argued, on the basis of extensive, detailed analysis of

cellular phenomena, that all cells display not only learning but consciousness [24].

The claim that consciousness is a feature of all life may be quite difficult to accept without an awareness of the importance of the problem of subjectivity to our fundamental theory. Many biologists currently doubt the validity of an inclusive explanation of subjectivity. But an explanation is necessary, even if it is to insist that subjectivity is entirely an epiphenomenon. Currently genes-first theories do not include subjectivity; there is no account of decision-making that isn’t controlled by inheritance, i.e. by evolutionary history. The genes-first explanation focuses on traits moving through Darwinian populations, and skips the individual level entirely [25]. As a result, there has been little insight into the fundamental rules of biological development which involves independent, subjective, contingent decision-making by cells. Despite enormous progress in our understanding of the molecular basis of inheritance and genetics, we still cannot explain exactly how it is that a plastic, living phenotype emerges from a largely invariant molecule, DNA.

This problem is greatly aggravated by any observable influences the individual might appear to have over the process of evolution. These are the problems the EES was designed to address. The EES, though it lacks coherence as an explanation, at least includes an acknowledgment that there are valid questions about how portions of the DNA sequence are selectively activated and silenced during ontogeny, why the ontogeny of many species involves their intimate interaction with the conditions of selection itself, and why many organisms and cancers are prone to suspicious meddling with the sequences of their own genetic materials [23]. These phenomena take the subjectivity of the individual out of a merely descriptive role in our explanations and place them in a directly causal role. Because they affect the course of evolution, these observations give physical, causal importance to mental events, and they therefore encourage the view that all cells can give rise to minds.

What are minds?

Minds are strange biological constructs. Entities of any sort are only scientifically real if they play an important role in our best causal explanations of reality; that is why the idea of minds has a current status that is, at best, hazy. In some versions of biological thought, they are accepted though not explained, while in more rigorous theories they lack any reality whatsoever. Such a universal “mental” property in all cells could be ruled out if we could locate a firm boundary between what

type of cells can give rise to mental functions and which cannot. If so, where would this boundary lie? Are neurons alone sufficient? Neurons and glial cells? Neurons, glial cells, peripheral nerves, and gut microbiota? Even humans depend, to whatever slight degree, on an acquired community of simple, mostly prokaryotic microbiota to accomplish our superior cognition [26]. In a learning-first theory of life, minds are a vital part of the explanation for how biology works; therefore if this view is accepted they will acquire the status of real, causally important entities that can act on the physical world. In a learning-first theory minds are not something outside of physics, but rather a name for the way physical matter-energy is causally disposed [27].

Dealing with minds and subjectivity is an essential part of a complete physical explanation of life either from genes-first or learning-first perspectives. In a learning-first theory, explanations for subjective attributes like memories, beliefs, intentions, and desires are necessarily a part of a complete theory because they are part of explaining individual behavior. Subjective experiences like love and fear, experiences which have meaning for a subject, are forces that drive organisms toward goals. These goals cannot be gene or trait-level goals that play out in populations over generations but must be subjective goals pertinent at the individual level, in the cognitive domain of an individual. Survival and reproduction are not events that an individual can learn from in making choices. It requires a highly complex awareness of the world outside one's own life and a great deal of storytelling to even form such concepts. This advanced storytelling is unlikely to occur in simpler life forms. If the learning-first alternative is correct, then individual choices are not epiphenomenal and inconsequential, rather they are the actual forces that shape evolutionary destiny.

On the other hand, if the genes-first alternative is true, then subjectivity is an illusion, and beliefs, intentions, and desires are unreal. They are unreal not in the sense that we cannot experience them but rather in the sense that they have no physical consequences at all. They do not affect the events we can measure. Minds, in this version, are not causally active and therefore even their existence is in grave doubt. Leaving aside the obvious contradictions this raises, let us accept provisionally that the genes-first view is a coherent theory but also be sure to articulate its consequences. An insistence that organisms are completely understandable as sets of genetic programs filtered by natural selection is an explicit denial of the power of subjectivity to shape the physical world. It is an explicit claim that organisms have no agency. As existentially

unpleasant as this may be, this acknowledgment is vitally important in an honest accounting of the full consequences of genes-first theory. It is unproductive to both agree and disagree with a causal claim in a physical theory because accepting causal contradictions makes empirical testing impossible. Either subjective experiences are informing the results we observe in the physical world or they are not; we can't have it both ways. Either way a complete physical theory of life must grapple with this dimension of living phenomena.

A new telos for life

A learning-first theory of life and a role for the individual in the explanation of evolution presents us with a big dilemma. Having discarded the telos of survival and reproduction as the aim of life, we are left with a large puzzle in explaining goal-directed behavior, agency, and individual choices. If not survival and reproduction, then what? What is it that organizes the aims that all living beings pursue? Are they merely idiosyncratic, erratic choices that emerge spontaneously with no rhyme or reason? If so, why would they result in coordinated functions? Why would cells be capable of cooperating in large swarms, multicellular organisms, holobionts, and ecosystems? Why would any of the separate parts of cells, like mitochondria or organelles, even cooperate with the other parts? There must be some fundamental force, some prime mover, as the evolutionary biologist Richard Watson likes to term it, that shapes the behavior of cells such that they are capable of complex development and cooperation.

Furthermore, this new telos must be compatible with subjectivity and the psychological understanding we have of living motivations. It cannot sit idly by, as survival and reproduction once did, and casually ignore the intimate interactions that organisms have with their worlds. It must derive from and inform those interactions in every aspect of our living existence. This telos must be able to cash out its consequences in terms of individual meanings, individual loves and fears, not just some love and fear for some life forms but all of it for all life forms. It must be interpretable as the root cause of the significance of all memories, the foundation of all beliefs, and the instigation of all desires, for all living cells in all living beings in all times and places in which they have ever lived.

The idea of an alternative telos for life is not new. The entire field of origin of life studies looks at molecular biochemistry and historical geophysical conditions to try to guess the conditions under which life (understood as replicated instructions) might have emerged. But if we take seriously their

premise that self-organization is a key prerequisite for life, then we have to also realize that this condition, once met, never ceased being met. At what point did organisms stop seeking self-organization? Of course they never could have done so. The study of self-organization is a space where an alternative telos of life is continually proposed. Even though they not been explicitly recognized and celebrated as such, ideas such as Kaufmann's "order for free [28], " Friston's free energy principle and active inference [29], Branscomb's anti-entropic machines [30], and others, are conjectures about what is prior to life, about what process allows for inheritance, and therefore about the telos of life. These theories, masquerading as molecular chemistry and physics, are also guesses about fundamental motivations, the purposive goals that lead life to prefer order and purge entropy. These theories implicitly assume that natural selection does not require organisms to dedicate themselves exclusively to survival and reproduction. The removal of nonviable phenotypes from a population does not enforce any particular motivation on the phenotypes that remain. As a result, the search for the principles of self-organization is a tacit admission that the telos of survival and reproduction can and must be replaced for a physical theory of life to emerge.

Having informed our understanding of biology now through the lens of Popperian epistemology, this new telos is no longer hard to identify. Given Popper's insight, the framing of this problem of the mysterious aim of life is now different. We can now interpret biology not as the winnowing of heritable programs, but as process of spontaneous, local knowledge formation. Since we know that natural selection can only operate after the fact as a form of refutation, we now have an open question about what process makes the conjectures. In other words, life now has an epistemological foundation. If Popper is right, self-organization is definitively a matter of conjectural knowledge. Organisms are motivated to understand the world, to develop theories of the entities and forces that exist around and within themselves by open-ended experimentation. In this view, organisms are not only continually making conjectures, organisms actually *are* conjectures. Their development and behavior, their making and remaking of themselves, is *itself* a process of conjecture and refutation. Organisms are guesses about the world, guesses that are reformulated continually, adjusting not to their survival prospects, but to better and better versions of their own embodied knowledge.

Three forms of change, not two

Every entity in the universe, both living and nonliving, has managed to "get itself into the future" both by changing itself and by not changing itself. Organisms are completely determined by their physical past, but in this respect they are no different from nonliving material objects. This fact offers no explanatory power at all. Adaptation means something more than mere continued existence.

Adaptation is inextricably linked to survival and reproduction. Adaptation is defined by being caused by natural selection, therefore it can only be properly considered to be changes that occur *between* organisms, changes that accrue at the population level, not changes *within* organisms. No observations can be made that demonstrate adaptation other than survival and reproduction, and a single organism cannot have more than one lifetime reproductive rate. The goals of adaptation are thus entirely meted out at the population level, therefore nothing on this axis explains goal-oriented change within an organism. Indeed, a key feature of the Modern Synthesis was to expunge goal-oriented processes at the individual level from the genes-first theory of life.

In a genes-first view, there are only two sorts of change with respect to knowledge, adaptation and chaotic disintegration, also known as stochasticity or blind variation or entropy. A genes-first theory bakes in all the knowledge an individual ever possesses at the genetic level; it denies that the phenotype can truly change itself adaptively at all except by blind accident.

In a learning-first theory, adaptation still happens. It describes the changes between individuals that over time lead to survival and reproduction in a lineage. As Darwin put it in a phrase which sounds inelegant to our modern ears, adaptation is the "preservation of favoured races in the struggle for life." In a learning-first theory, however, there is a third possible form of change with respect to knowledge, an intentional change at the level of the individual. Prefiguring a learning-first theory, Popper once imagined a world without natural selection, a world of expanding infinite resources where every organism was immortal [31]. He noticed that even in this world, evolution would still occur as organisms and populations changed. Popper's thought experiment thus neatly separates epistolution from adaptation.

Epistolution is a new term for intentional change in an organism at the level of the individual. Since adaptation cannot possibly apply at this level, we require a new word. Some candidates are obvious, but they all have fatal flaws. "Learning," "cognition," "intelligence,"

“sentience,” and “problem-solving” are all referencing a similar process, but these terms are already in common use, and they all are defined in terms of a human observer. Nothing is considered to be learning, intelligent, etc. unless it is capable of solving a problem that a human can recognize as a problem. Learning is currently defined and tested in terms of problem-solving ability, or identified in incipient forms of problem-solving such as sensitization, habituation, or associative learning [12–14]. This conception of learning is implicitly based on the assumption that learning is an evolved trait which arose late in the history of life to solve complex problems related to survival and reproduction.

These old terms thus live intrinsically in a world of objective problems and solutions. But a learning-first theory does not posit that sort of world. A learning-first theory is a theory of many subjective perspectives. In a learning-first theory there are many more domains of problems that humans cannot access simply because they are not the organisms for which these problems exist. A bacterium constructs its world entirely differently from a human, therefore we would be entirely unable to see which of its physiological changes are accumulating knowledge and which of them are disintegrative. The term epistolution is necessary to refer to all these changes. Learning, intelligence, cognition and so forth are subsets of epistolution in the domain of the human observer, but the concept of epistolution recognizes that there are many additional domains of knowledge.

This is why the term epistolution is necessary, both because in a learning-first theory there are significant changes organized at the level of the individual, and because those changes occur in the subjective domain of the organism itself. Despite the fact that we cannot directly inhabit that alien cognitive domain of another organism, acknowledging that it exists and is crucial to life might help us find new ways to test and understand epistolution, especially in nonhuman minds. Epistolution refers to the effects of the universal natural learning algorithm. In a learning-first theory this learning process represents a fundamental union of *evolution* with *epistemology*, the sources of knowledge.

Unlike problem-solving, which presupposes a goal implanted by natural selection, Popperian understanding means open-ended experimentation to develop an explanation of what entities and causal forces exist in one’s surroundings [16,17]. This would mean that all life, every living cell and also every whole organism, would contain within its phenotype some representation of the world that conflicts, in some respects, with experience. These conflicts present meaningful

problems, significant areas which call out for investigation. The urge to understand an *umwelt* could thus serve as an intrinsic motivation for all living behavior, even morphological development. As strange as it may seem, this would mean for example that a lineage of cells in a developing embryo, as it differentiates into distinct somatic cells with specific functions, would be also developing into unique perspectives. Differentiated cells would be in effect forming stable *opinions* about how they each should live and behave as individuals and collectively, opinions based on experience, memory, and learning.

Understanding as the aim of life

The ability to contain a representation of the world inside one’s phenotype could be a result, as the CBC theory has it, of consciousness in all living cells. Single cells may pass many of the fundamental tests of consciousness, including sensitivity to anesthetics, stable memory formation, and navigational behavior [32,33]. But representing the world in a subjective perspective is an even lower bar than becoming conscious of that representation. Humans are widely considered to have many attitudes, preferences, and habits of which we are unconscious. Consciousness is also largely extinguished during sleep even though sleep is considered critical for learning [18]. Like consciousness, there is currently no definitive empirical test to determine whether organisms or cells contain representations of their *umwelts*, yet we assume that some biological structures (humans) certainly do. It is currently impossible to rule out the possibility that other cells and organisms contain representations, whether these representations are conscious or not.

Within a learning-first theory all organisms also have minds [17,33,34]. If ideas in our human minds are what determine *our* behavior, with all the concomitant gene expression activity that this entails, then on what possible basis would a *non-agential* organism, one without any sort of mind, manage its patterns of gene expression? Humans seek understanding, but other life may seek it just as well. Understanding-seeking systems would be easy to mistake for survival-seeking genetic programs; the only such systems that could have persisted for long must incidentally also be compatible with a successful chain of material inheritance. Certainly DNA and other heritable materials carry influence, but perhaps they do so only in the presence of a natural learning system that gives a mental meaning to them in the context of a subjective perspective. It does not violate the objectivity of scientific inquiry to acknowledge that subjectivity may exist and prove

explanatory. It is quite difficult to imagine an umwelt from the point of view of another life form, and yet this is exactly what may be required to adequately explain a specific organism's actual development and behavior.

Why do umwelts repeat themselves?

There is also no doubt that the environmental conditions in which offspring develop tend to resemble in many major respects the environments in which their parents developed. This might explain, to some extent, how epigenetics could reliably trigger genetic programs and thereby help strengthen a genes-first view. But what is responsible for the maintenance of this orderly recurrence? The details of the mechanism, in either theory, is unknown. Since, according to the genes-first account, the material that is inherited is a matter of contingent historical accident compounded over trillions of generations, there is no reason to believe that any general principle could be extracted that might elucidate how ontogeny occurs. In this view genes contain encoded instructions for traits, but exactly how they are decoded is left blank. In a genes-first view life may be intractably complex and inscrutable because, as Branscomb writes, “the system and its historical antecedents are mechanistically and causally inseparable [30]”. The organismic ability to derive a trait from a gene requires a physical process that we cannot recreate in a lab. It requires nothing less than a living cell. If we were able to create artificial genes and artificial organisms from abiotic materials that followed encoded genetic instructions, in other words Von Neumann machines [35], this would weaken the learning-first theory and sustain the notion of genetic programs.

Since this attempt has never succeeded, the failure leaves a mystery at the heart of the process that creates the repetitive environmental conditions. We observe that the umwelts that organisms develop within resemble the umwelts of their ancestors. Attributing this effect to genes may confuse the effect with the cause. It is no accident that this occurs; it requires the presence of contextual knowledge to recapitulate the sameness of those umwelts again and again. Small differences in living conditions from generation to generation have to be accommodated, not erased from the process, but incorporated in a way that still functions.

Functions for what? In a genes-first view the only relevant function means replicating genes, while in a learning-first view function means accumulating knowledge and incidentally replicating genes as well. In a genes-first theory umwelts are recapitulated only by programmatic inherited behavior, which would

seem to be a much less reliable controller because of the lack of locally produced Bateson information. It is true that the fundamental idea of a genetic program means repetitive behaviors emerge by definition. But if the only aspect of control resides in an inflexible molecule, why is this molecule used in a contingent fashion that results in a repeating behavior?

In a learning-first theory on the other hand umwelts are recapitulated by being accommodated and actively shaped by the interactions of organisms. This also means that in a learning-first theory, organisms learn reliable lessons because there is a stable set of truths underlying physical reality, and similar inherited materials provides a similar sort of access to those truths. The sorts of opinions cells can form about reality are limited by the tools at their disposal, including not only the genetic tools, but the nutrients and other conditions in the local environment as well. Thus genes can appear to lead to heritable traits, while in fact only providing materials critical in assembling those traits.

The crucial test of a learning-first theory

While the only way to test the plausibility of a genes-first theory would be to build self-replicating artificial organisms (Von Neumann machines), a learning-first theory of life could be tested by building any artificial system that understood its umwelt without using programs or codes to do so. This subjective, agential machine capable of natural learning would be, in my terms, an *epistevolver*. A test of true understanding, of the embodiment of a conjecture that accurately identifies causal entities and forces in the local world of an epistevolving system, is difficult but not impossible to imagine. Any organism that learns must be a system with an open-ended constitution that makes it capable of informing its own internal organization by reading its external surroundings. Some abstract aspect of these surroundings must be incorporated into the organism functionally and purposefully. If a general physical mechanism was discovered to explain this phenomenon, or even a mathematical model of this behavior was shown to work in an artificial environment, it would obviate the need to posit supernatural origins for the presence of meaning, intention, and purpose in the biosphere. Such a model would also obviate the need to deny that such phenomena exist in order to preserve an attitude of scientific materialism. It would unite physics and biology.

What could be the general mechanism, common to all cell types on earth, that might produce epistolution? Could it be the networked connection of biological oscillators and servomechanisms? The CBC authors

have proposed that the fundamental glue that cements life may be clocks, including circadian clocks [36]. These clocks are considered essential to cognition because cognition is not only about *what*, but crucially also about *when* (presumably these might lead to *how* and *why*.) In cybernetics, life was conceived of as a set of servos and oscillators driven by an inherited program, but perhaps instead it is “use and disuse” in interactions with the *umwelt* itself that drives them. Gene expression (and everything that it leads to, including self-organization, development, and cognition) is controlled by an interaction between the phenotype and the environment.

Before the twentieth century genes-first paradigm was hardened into dogma, nearly all biologists presupposed that some form of “use and disuse” influenced development. By Denis Noble’s account, Darwin mentions this several critical times in the *Origin of Species* [37]. By taking this mysterious process of “use and disuse” seriously as a developmental principle, universal regularities might be discovered and formalized that could model epistolution in a network of oscillators linked by adjustable connections. Perhaps some form of program like this could interpret novel data and find new problems rather than acting simply as a readout of evolved, inherited signals. This type of algorithm, unlike all existing evolutionary algorithms, would have no intrinsic genotype-phenotype map and no pre-programmed survival or reproductive instinct. It would instead develop its form and its behavior strictly according to what it learned about the *umwelt*. If so, this would support the learning-first theory and resolve the question of why and how biological genotypes lead to phenotypes.

Whatever the mechanism turns out to be, the solution is constrained by several factors. Anything that cannot conceivably be instantiated by every living cell cannot be a candidate. It must be universally present in all life. The algorithm embodies a general mechanism for producing and comparing representations without containing a specific fitness function in any relevant domain. The algorithm trains only on its *umwelt*, not from any large set of training data, yet it recognizes patterns and from them it produces memories and pertinent goals to attain more knowledge. It is enactive and time-sensitive, and able to recover from disruption and amputation to a limited degree. It recognizes and defines its own external boundary. These parameters may limit the possible algorithms that we have to choose from. A careful examination of universal cellular phenomena such as circadian oscillation in light of these constraints might lead to candidate models.

A formalized model of epistolution, if it could be discovered, might serve as a general mathematical principle which in combination with Darwinian principles could more adequately explain life. If so, this conceptual breakthrough would lead to more effective medical intervention, regenerative therapies, and eventually to the ability to design and build whole organisms fit for therapeutic, commercial, or creative purposes. It would also lead to the discovery of genuine creative knowledge by machines for the first time. In contrast to the current statistical mimicry derived statistically from a large body of human knowledge yet termed “artificial intelligence” [38], machines programmed according to the (still currently unknown) principles of epistolution might gain causal understanding themselves. These machines could invent and test their own subjective theories of what entities and forces exist in the world. Though not necessarily conscious, these would be the first artificial systems with the ability to develop their own subjective cognitive perspectives, their own agency. This would comprise a powerful new technology amplifying scientific, technological, aesthetic, and moral progress in many intellectual domains.

A final word on Darwinism

It is essential to understand that a learning-first theory of life is not a refutation of Darwinism, only of the idea of genetic programs that result in heritable traits (sometimes called Neo-Darwinism [7]). In a learning-first theory traits are constructed with genes as templates, much like in music performances songs are constructed from piano keys, and evolution by natural selection still occurs. Even without a model of the specific mechanism that underpins epistolution, it strengthens the Darwinian theory of evolution by natural selection immensely to posit that this mechanism of epistolution does indeed exist. The dispute between genes-first and learning-first theories is a dispute over the instructions that take us from a heritable sequence of nucleotides embodying Shannon information to a trait which is a descriptive aspect of a living being. Without a learning-first interpretation, the theory of life contains no explanation for why plasticity arises and is able to address problems that are not identical to those experienced by a lineage during the evolutionary past, why despite a single set of heritable materials a germline cell differentiates into orderly multicellular phenotypes, why organisms actively influence their own conditions of selection to make it more likely that their lineage persists, and why genes are selectively turned on and off by epigenetic signaling that functionally adjusts to environmental conditions. All these features of life

contradict the straitjacket of inheritance and seem to require a learning-first explanation.

The organizational logic of organisms is simply presupposed rather than explained by existing mathematical models of the evolution of traits. The reigning interpretation of inheritance experiments presupposes heritable instructions. Usually this is accomplished in a genes-first view by surreptitiously defining a “gene” as the complete cause of a trait rather than as a particular sequence of nucleotides, or by positing not only a genetic code but additional codes in the germline cell such as the genomic code, mitochondrial code, epigenetic code, or karyotype code. But inventing more codes in the cell only multiplies the homunculi present in the theory of traits [39]. Adding an alternative purposive aim or telos of knowledge acquisition to the theory of life, one which can be pursued *de novo* by agents with perspectival subjectivity, at least theoretically resolves this explanatory gap. Above I ventured a tentative guess about oscillators as the universal natural learning mechanism, but many more such guesses should be proposed and tested experimentally.

It is not reasonable to require that a mechanism for epistolution be discovered before a problem with the existing genes-first explanation is acknowledged and a better one accepted. I have tried to do so above, but this is only a rough sketch. Mechanisms do not have to be understood in detail in order to become vital components of a good explanatory theory. A learning-first theory relies on a universal mechanism for natural learning that has not yet been discovered. But the reigning genes-first theory relies on a detailed physical map of instructions for how to get from genes to traits that has also not yet been discovered, in fact as I discussed above this has been falsified by the Human Genome Project. Neither theory has been confirmed as plausible by an abiotic experiment. Given these conditions, a learning-first theory is superior because it explains more, and explains it better.

The theoretical addition of epistolution to Darwinian biology does not require a supernatural or theistic claim of any sort. Although it does make sense of subjectivity and even morality in a new way, it does not involve any nonmaterialist claims. Though it also does not finally work out all the philosophical problems associated with the mind-body problem, even minds, in a learning-first theory, do play a physical role. An epistolution algorithm is not a deity or a divine influence but an explanation on a physical materialistic level for unexplained purposive phenomena that are impossible to deny. It assigns no cosmic plan or agency to evolution itself, only to individuals who assimilate knowledge. Biology-

derived learned knowledge is obviously influential in the physical world. Proof of this can be seen in the rapid development of human societies in the past century, a time period considered insignificant for the action of natural selection. Proof of this can also be seen in the fact that laws of physics, which do not include a theory of knowledge or of life, can only be rigorously tested by explicitly eliminating all the influences of biology on experimental apparatus; only then do the results of physical experiments conform to the current known laws of nature. Incorporating epistolution into our fundamental theory of life, and thus accomplishing a more satisfactory theoretical extension of physics, merely requires the supposition that we have not yet discovered and understood every physical process that exists in biological systems. A learning-first view also defines specific experimental projects for us, experiments by which we might find out much more.

Acknowledgments

Denis Noble, Frantisek Baluška, Ken Cheng, Wilson Minor, Daniel Phillips, Nate Gaylann, Richard Watson, Henry Heng, Anthony Bucci.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Data availability statement

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

ORCID

Charlie Munford  <http://orcid.org/0000-0002-1538-4268>

References

- [1] Hume D. An enquiry concerning human understanding. Chicago: Open Court Pub. Co; 1921. p. 267.
- [2] Popper KR. The logic of scientific discovery. (NY): Basic Books; 1959. p. 479.
- [3] Deutsch D. The beginning of infinity: explanations that transform the world. Vol. vii. (NY): Viking; 2011. p. 487.
- [4] Bateson G. Steps to an ecology of mind; collected essays in anthropology, psychiatry, evolution, and epistemology. Chandler publications for health sciences. San Francisco: Chandler Pub. Co; 1972. p. 545.
- [5] Slijepcevic P. Principles of cognitive biology and the concept of biocivilisations. *Biosystems*. 2024;235:105109. doi: 10.1016/j.biosystems.2023.105109

- [6] Shannon C. A mathematical theory of communication. *Bell Syst Techn J.* 1948;27(3):379–656. doi: [10.1002/j.1538-7305.1948.tb01338.x](https://doi.org/10.1002/j.1538-7305.1948.tb01338.x)
- [7] Dawkins R. *The selfish gene.* Oxford: Oxford University Press; 1976. p. 224.
- [8] Huxley J. *Evolution, the modern synthesis.* Vol. 645. London: G. Allen & Unwin Ltd; 1942. p. 1.
- [9] Weismann A, Parker WN, Rönnefeldt H. *The germ-plasm: a theory of heredity.* (NY): C. Scribner's sons; 1893. p. 477.
- [10] Noble D. *The illusions of the modern synthesis.* Springer Science+Business Media: Biosemiotics; 2021.
- [11] Sultan SE, Moczek AP, Walsh D. Bridging the explanatory gaps: what can we learn from a biological agency perspective? *BioEssays.* 2022;44(1):e2100185. doi: [10.1002/bies.202100185](https://doi.org/10.1002/bies.202100185)
- [12] Marshall P. Biology transcends the limits of computation. *Prog Biophys Mol Biol.* 2021;165:88–101. doi: [10.1016/j.pbiomolbio.2021.04.006](https://doi.org/10.1016/j.pbiomolbio.2021.04.006)
- [13] Watson R. *Songs of life and mind.* *YouTube.* 2023. <https://www.youtube.com/playlist?list=PLVmjximp0I4OJdT9bsFlebu0HjPAjtlEN>
- [14] Baluska F, Reber AS, Miller WB. Cellular sentience as the primary source of biological order and evolution. *Biosysts.* 2022;218:104694. doi: [10.1016/j.biosystems.2022.104694](https://doi.org/10.1016/j.biosystems.2022.104694)
- [15] Uexküll JV, Uexküll JV. *A foray into the worlds of animals and humans: with a theory of meaning.* 1st University of Minnesota Press ed. Minneapolis: University of Minnesota Press; 2010. p. 272.
- [16] Levin M. Bioelectric signaling: reprogrammable circuits underlying embryogenesis, regeneration, and cancer. *Cell.* 2021;184(8):1971–1989. doi: [10.1016/j.cell.2021.02.034](https://doi.org/10.1016/j.cell.2021.02.034)
- [17] Levin M. Technological approach to mind everywhere: an experimentally-grounded framework for understanding diverse bodies and minds. *Front Syst Neurosci.* 2022;16:768201. doi: [10.3389/fnsys.2022.768201](https://doi.org/10.3389/fnsys.2022.768201)
- [18] Visscher PM, Wray NR, Zhang Q, et al. 10 years of GWAS discovery: biology, function, and translation. *Am J Hum Genet.* 2017;101(1):5–22. doi: [10.1016/j.ajhg.2017.06.005](https://doi.org/10.1016/j.ajhg.2017.06.005)
- [19] Shapiro JA. *Evolution: a view from the 21st century.* Vol. xi. Upper Saddle River, N.J.: FT Press Science; 2011. p. 253.
- [20] Laland KN, Uller T, Feldman MW, et al. The extended evolutionary synthesis: its structure, assumptions and predictions. *Proc R Soc B.* 2015;282(1813):20151019. doi: [10.1098/rspb.2015.1019](https://doi.org/10.1098/rspb.2015.1019)
- [21] Minorsky PV. The “plant neurobiology” revolution. *Plant Signal Behav.* 2024;19(1):2345413. doi: [10.1080/15592324.2024.2345413](https://doi.org/10.1080/15592324.2024.2345413)
- [22] Parise AG, Gagliano M, Souza GM. Extended cognition in plants: is it possible? *Plant Signal Behav.* 2020;15(2):1710661. doi: [10.1080/15592324.2019.1710661](https://doi.org/10.1080/15592324.2019.1710661)
- [23] Ginsburg S, Jablonka E. *The evolution of the sensitive soul: learning and the origins of consciousness.* Cambridge (MA): The MIT Press; 2019. p. 646.
- [24] Reber S, Baluska AFE, Miller WB. *The sentient cell: the cellular foundations of consciousness.* (NY): Oxford University Press; 2023. pages cm.
- [25] Godfrey-Smith P. *Darwinian populations and natural selection.* Vol. viii. Oxford; (NY): Oxford University Press; 2009. p. 207.
- [26] Meyer K, Lulla A, Debroy K, et al. Association of the gut microbiota with cognitive function in midlife. *JAMA Netw Open.* 2022;5(2):e2143941. doi: [10.1001/jamanetworkopen.2021.43941](https://doi.org/10.1001/jamanetworkopen.2021.43941)
- [27] Deacon TW. *Incomplete nature: how mind emerged from matter.* 1st ed. (NY): W.W. Norton & Co. xv; 2012. p. 602.
- [28] Kauffman SA. *The origins of order: self-organization and selection in evolution.* (NY): Oxford University Press. xviii; 1993. p. 709.
- [29] Pezzulo G, Parr T, Friston K. Active inference as a theory of sentient behavior. *Biol Psychol.* 2024;186:108741. doi: [10.1016/j.biopsycho.2023.108741](https://doi.org/10.1016/j.biopsycho.2023.108741)
- [30] Branscomb E. Boltzmann's casino and the unbridgeable chasm in emergence of life research. *ArXiv.* 2023;2312:47.
- [31] Niemann HJ. *Karl Popper and the two new secrets of life: including Karl Popper's Medawar lecture 1986 and three related texts.* Vol. vii. Tübingen: Mohr Siebeck; 2014. p. 157.
- [32] Lyon P. The cognitive cell: bacterial behavior reconsidered. *Front Microbiol.* 2015;6:264. doi: [10.3389/fmicb.2015.00264](https://doi.org/10.3389/fmicb.2015.00264)
- [33] Levin M. Bioelectric networks: the cognitive glue enabling evolutionary scaling from physiology to mind. *Anim Cogn.* 2023;26(6):1865–1891. doi: [10.1007/s10071-023-01780-3](https://doi.org/10.1007/s10071-023-01780-3)
- [34] Levin M. The computational boundary of a “self”: developmental bioelectricity drives multicellularity and scale-free cognition. *Front Psychol.* 2019;10:2688. doi: [10.3389/fpsyg.2019.02688](https://doi.org/10.3389/fpsyg.2019.02688)
- [35] Neumann JV. *Theory of self-reproducing automata.* 1966. https://archive.org/details/theoryofselfrepro0vonn_0/page/n5/mode/2up
- [36] Baluska F, Reber AS. CBC-Clock theory of life - integration of cellular circadian clocks and cellular sentience is essential for cognitive basis of life. *BioEssays.* 2021;43(10):e2100121. doi: [10.1002/bies.202100121](https://doi.org/10.1002/bies.202100121)
- [37] Noble D. *Dance to the tune of life: biological relativity.* Vol. 283. Cambridge; (NY): Cambridge University Press; 2017. p. 12. unnumbered pages of plates.
- [38] Jaeger J. Artificial intelligence is algorithmic mimicry: why artificial “agents” are not (and won't be) proper agents. *Neurons, Behav, Data Anal, and Theory.* 2024;2307:07515. doi: [10.51628/001c.94404](https://doi.org/10.51628/001c.94404)
- [39] Oyama S. *The ontogeny of information: developmental systems and evolution.* Cambridge; (NY): Cambridge University Press; 1985. p. 206.