The Journal of Mind and Behavior
Summer and Autumn 2024, Volume 45, Numbers 3 and 4
Pages 301–312
ISSN 0271–0137

Understanding Consciousness by Analogy

Dan Bruiger

Hornby Island, British Columbia

Where a causal approach fails, the nature of consciousness can be understood by means of an analogy, by putting oneself imaginatively in the place of the brain. It is further proposed that neural processes evoke sensation and meaning analogously to how words evoke mental images, namely by fiat. Conscious experience is also likened to a virtual reality generated by the brain, while guided by interaction with the external world. An intentional rather than causal explanation is proposed.

Keywords: explanatory gap, intentionality, fiat

The "hard problem of consciousness" (Chalmers, 1995) is the "explanatory gap" (Levine, 1983) between phenomenal experience and causal processes — in other words, between first-person phenomenology and third-person description. In physical science, explanation typically involves a third-person account, especially in terms of efficient causation.[1] But, that is a closed domain insofar as physical events can only be said to cause other physical events — not sensations, images, feelings, or thoughts about physical events (McGinn, 1989). Knowing *about* an experience is not the same as having it (Godfrey–Smith, 2021). Neural processes, understood as such causal processes, can potentially

Correspondence concerning this article should be addressed to Dan Bruiger, 3105 Cowie Road, Hornby Island, British Columbia, Canada V0R1Z0. Email: dbruiger@gmail.com

[1] Among Aristotle's four types of cause, "efficient" cause is the category which was embraced as primary for dynamics by the early scientists and ever since. "Material" cause also figures in modern science, as concerning the properties of materials. A functionalist view of mind might hold that the biological properties of neurons (material cause) are not essential, whereas their dynamical organization is, which could perhaps be realized in a different medium. A modern version of the "formal" cause of an organism might be its DNA; or the formal cause of a machine, its blueprint. "Final" cause is viewed with suspicion in physics because it implies teleology (as in Intelligent Design). Yet, all these types of cause can pertain to organisms.

account for observable behavior (including cognitive behavior), but not for phenomenality,[2] the subjective "what it is like to be" the system itself (Nagel, 1974).

The mind's natural external orientation becomes problematic when the strategy for understanding consciousness is restricted to the terms of the external world. The problem is "hard" because it cannot be solved in the terms in which questions about the external world are normally posed and answered. A physicalist solution to the problem posed by phenomenality remains elusive — for example, a solution which explains phenomenality in terms of chemical, neuro-physiological, or quantum processes. A computational approach is more promising, because it invokes the intentionality of the programmer. Yet, it still retains a third-person perspective that fails to embrace the point of view of the computational system itself (such as a brain, for example).

On the other hand, metaphorical or analogical thought about mind has a long and rich history: from Plato's cave, to Descartes' demon, to the "brain-in-a-vat," to the *Matrix* films, and the "simulation hypothesis" (Bostrom, 2003). Nature itself was once understood to be or be like an organism and later to be or be like a machine. Following changing technologies, mental processing, and even the universe itself, are now considered to be or be like digital computation or information processing. For example, Max Tegmark (2017) claims that "consciousness is a physical phenomenon that feels non-physical because it's like waves and computations: it has properties independent of its specific physical substrate." The integrated information theory of consciousness (Tononi, 2004) is a laudable application of the computational metaphor, which claims to bridge the explanatory gap. It fails in that regard, however, precisely by abandoning analogical thinking to assert actual identity between mental and physical. A recent refinement of the theory (Albantakis et al., 2023) attempts to formalize a rigorous identity theory of correlates of consciousness. However, simply providing a mathematical framework does not overcome the explanatory gap. The paper carefully defines "experience" and "physical existence," but in the end simply *postulates* an "explanatory identity" between them. Defining each domain rigorously, in terms of a common theoretical domain (information), does not make them equivalent in a way that bridges the explanatory gap.[3] Here let us pursue an alternative strategy.

---

[2] Because "phenomenal experience" is a mildly redundant phrase, hereafter *phenomenality* will be used instead to denote the totality of experience inclusively: all possible contents of consciousness — all experience that can occur to a subject, such as sensations, feelings, thoughts, emotions, imagination, dreams, hallucinations, mental images, qualia, etc.

[3] From the sufficiency of the definition of experience, it does not follow that "Thus, no additional physical property is a necessary requirement for being a substrate of consciousness." Moreover, a biological or relational property may also be required: namely, *embodiment* as an evolutionary relationship with an environment. The authors claim that "The identity is not between two different substances or realms — the phenomenal and the physical — but between intrinsic (subjective) existence and extrinsic (objective) existence." This does no more than rename phenomenal and physical as "intrinsic" and "extrinsic" existence.

There exists a parallel explanatory gap between words and the mental images and feelings they evoke, which does not seem to have given rise explicitly to a "hard problem of language."[4] In fact, these gaps are examples of the same conundrum, which may ultimately be unresolvable simply because we cannot stand outside the dilemma. From a first-person perspective, both language processing and perceptual processing seem transparent. (That is, we are hardly aware of the processes themselves, only of their result.) While transparency itself merits explanation, the hard problem of consciousness amounts to a symbol grounding problem[5] for human beings.

Lacking a scientific resolution of either conundrum, perhaps we can at least gain some comfort from the familiarity of language, by asserting that consciousness arises from neural processes *analogously to how meaning arises from words*, which ultimately involves grounding in sensory experience. Thus, the explanatory burden can be shifted from causal explanation to intentional explanation.[6] The burden falls on the cognitive system as an embodied agent that creates meaning (i.e., significance to itself) for its own purposes.

Cause Versus Intention

The term "neurological" suggests two aspects. There are *neural* events, such as the chemical discharges of nerve cells, propagated along axons. These can be viewed as events in the physical world that happen through causal processes in space and time. However, these are also *logical* events or intentional acts: something the organism *does* for its own reasons, as part of its survival strategy.[7] In that sense, we may think of neurological events as simultaneously intentional

---

[4] Unless for the "symbol grounding problem" for artificial intelligence (Harnad, 1990; Searle, 1980). The problem considered here is not about the relationship between language and *concepts*.

[5] "The symbol grounding problem is a concept in the fields of artificial intelligence, cognitive science, philosophy of mind, and semantics. It addresses the challenge of connecting symbols, such as words or abstract representations, to the real-world objects or concepts they refer to. In essence, it is about how symbols acquire meaning in a way that is tied to the physical world." [Wikipedia: symbol grounding problem].

[6] Note that "intention" here does not refer specially or only to consciously experienced human intentions, but more broadly to connections an agent makes within itself for its own reasons. Like Dennett's (1990) "intentional stance," Vollmer (1986) describes intentional explanation as a way to explain why some behavior may occur at a certain time — how it is "rational, understandable, and to be expected," so that reasons "constitute premises from which agents (and observers) can deduce that certain actions are desirable and appropriate." Note that the reasons of the cognitive system observed may not correspond to the reasons of the observer. Note also that teleology and design are implied, in contrast to efficient causation. Here, however, we are not concerned with accounting for *behavior* but for subjective *experience*. Intentional explanation may be necessary to account for behavior too.

[7] Logical, in the broad sense of a (potentially formalizable) system with elements and rules, not in the narrow sense of a given human logic such as Boolean algebra. Intentional, in the broad sense of being effected by an agent for a reason, not in the narrow sense of conscious human intention.

and causal.[8] Indeed, as noted by Lakoff and Johnson (1980), and earlier by Piaget (1967), causality merges with intentionality in the early development of the human mind. For the infant, the sense of one external event *causing* another derives from, extends, and externalizes the sense of direct manipulation of objects and of willing one's own limbs to move.

While the organism is a material object, subject to causal processes within and without, it is also an *agent*. As a physical thing in the natural world, it can be acted upon by other things. Yet, because it is an autonomous, self-sustaining and self-defining (autopoietic[9]) system, the organism can also act on the world and on those parts of the world that constitute its own physical being. Its actions, whether internal or external, can be viewed as intentional (if not as consciously intended) as well as caused.

In physicalist terms, causality occurs within the domain of physical description of events occurring in the "external" world. The explanatory gap is the fact that consciousness does not seem to occur in that domain, but in some other domain or category, the "phenomenal" realm. While causal explanation may not be able to bridge these domains, here we propose intentional explanation, since *bridging* itself is an intentional action.

An *intentional connection* is a mapping (in the mathematical sense) from one domain to another, made by an agent for some purpose. Physicalism is at a loss to rationalize the domains of the physical and the phenomenal (and thus provide a causal explanation of consciousness), because it confines itself, in effect, to third-person description (Bruiger, 2016). Yet, our brains bridge the explanatory gap on a daily basis, so that somehow neural processes result in conscious experience in the first person. Activity within the brain is normally projected as experience of a real world outside the skull. The challenge is to understand how this happens.[10] The proposal here is that this feat is of the same nature as the (also mysterious) act by which the brain conjures mental images and other experiences upon hearing or reading words. A convenient name for this act is *fiat*.[11] The term conveys the sense of defining into being the elements and operations involved in mapping.[12]

Let us view the organism as an intentional agent making intentional connections. In one sense, these are internal connections that supervene, for example,

---

[8] Similarly, one can consider the connections on a circuit board either as physical (soldered wires) or as intentional (its intended design and use).

[9] The term autopoiesis (literally "self-creation") was introduced by Maturana and Varela (1980).

[10] As illustrated even by the word understand, language itself is largely metaphorical (Lakoff and Johnson, 1980).

[11] I.e., decree, as when the Red Queen orders "Off with her head!" Or declaring into existence, as when God says "Let there be light!" Or supposition, as when the mathematician proposes, "Let x stand for…."

[12] It is interesting in this context that the word *metaphor* derives from a Greek verb meaning to "carry over" or "transfer," which conveys the sense of intention as a mapping from one domain to another.

on physical synaptic connections. In another sense, they are symbolic elements to map a putative external world, including the physical body as part of that world. The explanation of phenomenality we seek is not in terms of physical processes, or causes originating in the external world, but in terms of intentional processes that originate within the organism, which *makes* intentional connections for its own purposes, often in response to the external world.[13] As such, they constitute a form of internal communication and a basis for action. The meaning (to the organism) of these connections supervenes on physical connectivity in the way that semantic meaning does on the symbols and syntax of language.

The Simile of the Submarine

Being intentional must not presume being conscious, since consciousness is what we hope to explain, and circular reasoning is generally considered empty. Nevertheless, we can allow ourselves the subterfuge of analogy — which does, of course, presume our own consciousness as language users. Here we hope to understand the brain's challenge by putting ourselves in its place. The brain is sealed inside the skull, connected via nerve fibers with the world outside. By presumption, there is no other way for information to enter or leave the brain. The skull is a "black box," whose functional content can, to some extent, be inferred by an outside observer from comparing inputs and outputs.[14] From a point of view within it, the world outside the skull is equally a black box.

The skull is not a room with windows through which an imaginary inner occupant views the outside world. Despite the ancient trope, the eyes are not literally portals, but more like remote sensors connected by wires that supply digital feeds to an underground bunker. Out-of-body experiences notwithstanding, there is no door through which to exit to gain experience outside this room. Our metaphorical task is rather to explain *seeing* and *experiencing*, and to arrive at a concept of the *world*, without already presuming any of these: in other words, to explain how the brain constructs phenomenality and a concept of an external world, purely in terms of processes taking place inside the black box. For, otherwise, we are again caught in circular thinking. If we invoke the metaphor of a room, it is a room without portals and exits. If we invoke an imaginary occupant, it is an agent who has never been outside the room and has no prior knowledge of a world outside (or that there is even such a thing as "outside"). With those provisos, we deliberately invoke a hypothetical conscious agent — a homunculus — who can explore this interior environment and do things within it.

---

[13] Some connections come about through natural selection, while the individual organism makes connections through real time learning. However they come about, they establish a basis for action and for evaluating stimuli.

[14] It makes little difference for this argument whether the "black box" is literally the skull, containing the brain inside it, or some other boundary within the skin.

The analogy can be made more tangible by likening the sealed chamber of the skull to a submarine without portholes, hatch, or periscope.[15] As outside observers, *we* know that there is an underwater world surrounding the hull. Our homunculus, however, has no such knowledge to start with. Rather, its task is to gain that knowledge in the only way possible: through trial and error within the confines of the vessel. Let's say that the interior of the submarine comes equipped with what we (as outsiders) recognize as "controls" and "instrument panels" — that is, with what our homunculus may eventually come to recognize as inputs and outputs. There are levers and switches to play with, to try to discover any patterned relationships between those actions and the readings on various gauges and dials. In other words: to explore how doing something to the controls might bring about changes in the instrument readings (Oatley, 1978).

In our outsider's view, what causally connects these inputs and outputs is the ontologically real world outside the hull and the fact that the submarine is a part of that real world, can move through it, and can perform actions upon it (for example, with sonar and robotic arms). In other words, the submarine can affect and be affected by its environment. Pulling a lever inside, for example, might activate a propeller causing motion through the water; pushing a button might issue a sonar pulse whose echo is registered by a sensor. In that feedback loop, it is the real underwater world that mediates the patterns between input and output, by completing the loop, so that the patterns identified contain information about that world. It is thus epistemically possible for our homunculus, through such experimenting, to create a model or map of the underwater environment surrounding the submarine, without ever seeing or touching it directly, simply by actions performed within.

Why would this agent bother to do any of this, apart from idle curiosity? The submarine obviously represents a living organism — in this case, a human body. But, of course, the submarine is *not* an organism but simply a machine. While no machine (so far) has a vested interest in its own existence, an organism is *defined* by that vested interest. The organisms that exist have learned how to negotiate their environments and would not exist otherwise. Natural selection is the process of eliminating failures, which drives evolution. So, we must imagine a corresponding principle whereby submarines that lack a "realistic" enough model of the underwater world are potentially eliminated. We must imagine a submariner who knows nothing yet of that principle, but who simply succeeds or fails to preserve the submarine through trial and error. Let us therefore imagine generations of submarines that have adapted (or not) to the underwater world through some equivalent of natural selection combined with learning. Our agent must bother with the modeling process as a condition for its existence qua organism.

[15] The analogy precludes anything like closed circuit tv or video monitor as one of the remote sensors, which would substitute for a porthole by providing a ready-made image of the outside world. The point is to understand how the brain (homunculus) constructs this image for itself.

The challenge for our hypothetical agent is to interpret the inputs from instruments as information about a real external environment, and from that evidence to model that environment. In our analogy, we imagine the homunculus with eyes to see the interior of the submarine and limbs to move about inside it and manipulate controls. But all that is no more than a concession to metaphor, in which we have placed ourselves imaginatively in the situation of the brain. In relation to the reality outside, the submarine is originally *blind*, *ignorant*, and *uncoordinated*. The task of the homunculus is precisely to learn to see and navigate the surrounding world, incidentally learning to treat it as *real* — that is, as holding the power of life and death over the submarine with its occupant. The submariner's ultimate achievement is to see, as it were, straight through the hull, transparently, as though with x-ray vision. In other words, by an act of conjuring, to experience the model as a real world outside the hull.

The Heuristic Virtue of Virtual Reality

Since the activities of our homunculus, and the data it has gathered, are potentially formalizable in a computer program, we consider a further analogy: virtual reality. The ideal of simulation is to be so like reality that one cannot tell them apart. Normally, a simulation is a computer program that convincingly imitates a real thing or experience. It is itself an analogy. Our submariner's model is a simulation of the undersea world, achieved through a long learning process. Yet, the model cannot be said to *copy* or to literally *resemble* the real thing or situation, to which there is no direct access for comparison. Let us therefore imagine a program that is an original creation, not a copy of something else. Let us further suppose that this original creation is nonetheless guided in an ongoing way by an external reality — for example, through some form of predictive processing (Hohwy, 2013) — just as the development of the submariner's model is guided by the interaction between controls and instruments: through a feedback loop that includes an allegedly real environment. With this modification, we assert that phenomenality is a virtual reality created by the brain, yet continually updated through interaction with a real environment outside the skull (Metzinger, 2009).

A conventional virtual reality is often created as entertainment; but the virtual reality created by the brain is a matter of life and death. A simulation seems real to the degree it is convincing, but the realism of an entertainment does not have the same significance as the brain's natural realism. A conventional virtual-reality headset can be put on or taken off at will by users, who normally will not forget their identity as human beings who can embrace or disengage from the experience. This was not the case for our submariner, who could not leave the submarine and had never had a life outside it to remember. Nor is it the case for the brain. So, in this new analogy we must imagine someone who grew up in the simulation, had never lived outside it, and cannot turn it off. Imagine, therefore,

a simulation like in *The Matrix*, designed to be so comprehensive and convincing that it effectively deceives its captive users. Of course, the film is a fiction and we are not deceived in the same literal way. The point is that the transparency of the *Matrix* simulation resembles that of ordinary perception. In the film, a "glitch" in the computer code belies the situation.

There are glitches in the brain's virtual reality too. The science of cognitive psychology is founded on them. The realization that there *is* processing going on, and that the brain somehow *produces* our conscious experience, began with the recognition of perceptual anomalies. These are glitches in normal perception, such as optical illusions, illusions of shape and figure/ground, motion effects, experimental investigations of sensory adaptation, hallucinations, and cognitive illusions such as the rubber hand effect (Metzinger, 2009). If normal perception were seamlessly transparent, we would all be naïve realists who simply believe that the world exists exactly as we see it and that the brain has nothing to do with the world's appearance. In fact, in daily life it serves us well to believe the brain's virtual reality.

However, such reflections led early thinkers like Descartes to the dread conclusion that it is possible to falsify experience by hacking into the nervous system. That suspicion led to the brain-in-a-vat scenario, the *Matrix* films, and the counterintuitive claim that you probably are, without knowing it, living in a simulation (Bostrom, 2003). Descartes' solution to the dilemma was to trust that God would not allow such systematic deception. In modern thought, we might instead trust that *nature* would not allow it — if by deception we mean a set of ideas that would lead to our elimination through natural selection.[16]

The Hard Problem Metaphorically Mitigated

There remains the hard kernel of the problem of consciousness: how to understand the process through which neurological activity in the brain becomes (or is) phenomenality. We've conceded that a strictly causal explanation is ineffective, because causes do not account for the organism's purposive activity as an agent. Computation may provide a better analogy for mental processing than physics or chemistry, because it invokes the agency of the programmer or user. After all, while you can explain the functioning of a computer on a certain level — in physical terms of wiring, electrical charges and flows — it is the *logical* organization of the device that makes it seem to perform mental operations. It can mimic human thought processes because it was designed by human agents to do so, reflecting their intentionality.

[16] In that context, the human proclivity to tamper with and defy such natural restraints is remarkable and alarming. We seem fascinated by the border between reality and illusion, and committed to replicate nature artificially.

An autopoietic system has its own intentionality. To experience some form of phenomenality, there must be "something it is like" (Nagel, 1974) to be that system.[17] Such a system has priorities. Events *matter* to it, and this mattering is the foundation and prerequisite for the system's ability to make distinctions meaningful to it, upon which it can act. Phenomenality is how the organism represents these distinctions *to itself* and what they signify for it.[18] Input comes from the world (which includes the body), but response comes from the organism as an agent. Outside observers are used to viewing sensory input (stimuli) in third-person causal terms — that is, from their own point of view. But the organism's response must be viewed from *its* point of view — that is, in the intentional terms of an agent acting in its own interests. An outside observer must understand the organism's observable behaviour *and* its phenomenality in terms of actions or connections made within itself for its own reasons.

From a third-person perspective, phenomenality can be viewed as a sort of internal communication, an internal memo about (largely) external input. From a first-person perspective, the world appears to us in consciousness *like* the way that mental images appear to us when evoked by words — that is, by fiat. If that analogy seems more like magic than science, it is the same magic we use every day in language. In any case, the organism constructs its own first-person point of view and phenomenality through acts that an observer can translate as propositional assertions. The brain creates phenomenality in a parallel way to how it creates meaning in language, through the use of symbols to which it *assigns* meaning in the context of an embodied evolutionary history.

While the words of a natural language have relatively transient definitions, the *sensation* of greenness, for example — unlike the word that represents it — is not merely a linguistic convention subject to social change, but a convention of neurological organization, with the force of long genetic precedent. Indeed, the human cognitive system adapts to distorting colored lenses or filters in such a way that subjective experience of verdant foliage, for example, is eventually restored to its normal greenness (Neitz et al., 2002). The sensation of greenness is just what it is, and different from the sensation of redness, precisely because of the real-world things it refers to in our evolutionary history, from which it cannot be arbitrarily dissociated.

---

[17] Note that this expression, which has become popular to characterize the ineffability of phenomenality, is literally a simile — something like something else! We can relate to the experience of another mind only by analogy with our own experience.

[18] A model or representation may be viewed in the third person as a theoretical object, like a program or brain circuit. As an anonymous reviewer has kindly pointed out, a rat navigating a maze may be said to have a "representation" of the maze. (The same may be said of a virtual creature in a virtual maze). In the case of a human being, that representation would be experienced, in the first person, as a mental image of the maze, not a *perceptual* image. The point here is to elucidate how the brain produces the perceptual image: sensory experience in particular, and phenomenality more generally.

One might still wonder what it is about the qualitative "feel" of greenness that commends it to represent foliage, and what about redness commends it to represent things that must stand out against that background — rather than, for example, vice versa. The question may be likened to asking why a specific meaning is denoted in a particular language by a given word, written and pronounced its particular way, rather than by some other symbol. For the native language user, the association seems natural and self-evident, though of course it is actually a social convention and a product of historical accident, subject to change. The internal communication of the organism may be no less arbitrary in its choice of symbols, but is stabilized by the external world to which it refers. *Some* symbol must be chosen, and will inevitably come to seem imbued with the meaning it is made to convey through its connection to the real world. Thus, it is backwards to ask why grass appears green. Rather, greenness is what it is by virtue of the totality of associations related primarily to chlorophyll.[19] Sensory phenomenality is not something gratuitously added to the information it represents, nor caused by it, any more than words are caused by the things they represent. Rather, it is a *version* of that information presented synoptically in consciousness.

The nature of that re-presentation must be understood intentionally. Phenomenality in general involves the sort of act of fiat demonstrated in the visual blind spot and other perceptual completion effects.[20] In the case of the blind spot, the experience of continuity of the visual field is the brain's way to represent to itself its (true) belief that (despite the physiological blind spot) the external world is visibly continuous (Dennett, 1991). The brain affirms that conviction by an act of fiat, which ignores the sensory discontinuity. In other words, it "fills in" phenomenality in the visual field between the enervated retinal areas on either side of the un-enervated area. However, all enervation is ultimately discrete — with gaps between receptors, for example. These gaps *in turn* must be phenomenally filled in, but on a finer scale — temporally as well as spatially (so that there is continuity of motion). In all cases the brain asserts continuity across discrete structures or events when their discreteness is irrelevant to the organism, just as it asserts continuity between frames of a motion picture. Thus, the world has an analog look despite sensory digitation.

---

[19] Which is why there can be no "inverted spectrum" (the idea that color experiences could be systematically interchanged for different human subjects).

[20] Laboratory experiments demonstrate visual completion effects of various sorts. For a "taxonomy of perceptual completion phenomena" see Pessoa et al. (1998). Other experiments demonstrate various forms of spatial and temporal projection. Still others show the adaptability of the nervous system to restore perception that corresponds to functional behavior within an environment.

Conclusion

At least in some fashion, we can understand consciousness through analogy, and perhaps *only* thus. If analogy seems less than satisfying as an explanation of consciousness, it is because we expect "explanation" to invoke causes among objects. But consciousness is no object. The perceiving subject is never within its own field of view. We can explain things appearing in that field of view in terms of other things found within it. However, it is chasing one's tail to try to explain the experiential field itself in terms of appearances within it. By taking intentionality into account, one can at least begin to grasp the brain's challenges in terms we can humanly relate to.

Despite scientific advancements, we may never have an answer to the mystery of why anything exists at all. Similarly, we may never have an answer to the mystery of how there can be consciousness of it. Indeed, it is only the reflexivity of the conscious mind that gives rise to such questions, and it may be this very reflexivity that renders answers to them perpetually elusive.

References

Albantakis, L., Barbosa, L., Findlay, G., Grassol, M., Haun, A., Marshall, W., Mayner, W., Zaeemzadeh, A., Boly, M., Juel, B., Sasai, S., Fujii, K., David, I., Hendren, J., Lang, J., and Tononi, G. (2023). Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLoS Computional Biology, 19*(10), 1–45, *e1011465.*

Bostrom, N. (2003). Are you living in a computer simulation? *Philosophical Quarterly, 53*(211), 243–255.

Bruiger, D. (2016). *The found and the made: Science, reason, and the reality of nature.* Routledge/Transaction.

Chalmers, D. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies, 2*(3), 200–219.

Dennett, D. C. (1990). *The intentional stance.* MIT Press.

Dennett, D. C. (1991) *Consciousness explained.* Little Brown.

Godfrey–Smith, P. (2021). *Metazoa: Animal life and the birth of the mind.* Picador (Pan Macmillan).

Harnad, S. (1990). The symbol grounding problem. *Physica D, 42,* 335–346 [arXiv:cs/9906002].

Hohwy, J. (2013). *The predictive mind.* Oxford University Press.

Lakoff, G., and Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press.

Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly, 64,* 354–361.

Maturana, H., and Varela, F. (1980). *Autopoiesis and cognition: The realization of the living.* D. Reidel.

McGinn, C. (1989). Can we solve the mind–body problem? *Mind* [*New Series*], *98*(391), 349–366.

Metzinger, T. (2009). *The ego tunnel: The science of the mind and the myth of the self.* Basic Books.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review, 83*(4), 435–450.

Neitz, J., Carroll, J., Yamauchi, Y., Neitz, M., and Williams, D. R. (2002). Color perception is mediated by a plastic neural mechanism that is adjustable in adults. *Neuron,* 35(4), 783–792.

Oatley, K. (1978). *Perceptions and representations: The theoretical bases of brain research*. Methuen.

Pessoa, L., Thomson, E., and Noë, A. (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences, 21*(6), 723–802.

Piaget, J. (1967). *La construction du rèel chez l'enfant*. Delachaux et Niestlé.

Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences 3*(3), 417–457.

Tegmark, M. (2017). *Life 3.0: Being human in the age of artificial intelligence.* Alfred A. Knopf.
Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience* 5, 42.
Vollmer, F. (1986). Intentional explanation and its place in psychology. *Journal for the Theory of Social Behaviour, 16*(3), 285–298.