PROCESS, CONSTRAINT, AND CONSCIOUSNESS IN A SELF-ORGANISING COSMOS:

TOWARD A NATURAL-PHILOSOPHICAL RESEARCH PROGRAMME

Nobuyuki Otsubo

ABSTRACT: Modern science often treats matter, life, and consciousness as separate domains. Taking Whitehead's process philosophy as its axis and drawing together non-equilibrium thermodynamics, complexity science, and consciousness studies, this paper develops CAT (Cosmic Architecture Theory) as a natural-philosophical research programme. Its central schema articulates the relation between multi-scale constraint architecture ($\mathcal{A}$) and embedded local modes (μ), together with a cross-scale embedding relation through which local modes at one scale participate in the constraint environment of the next. CAT proposes that matter, life, and consciousness may be understood as phases of a single processual continuum, and formulates the bridge from complex organisation to experience as an explicit, speculative principle rather than leaving it tacit. It re-reads Gödelian undecidability and computational irreducibility, within an inference to the best explanation rather than a deduction, as considerations bearing on ontological openness. CAT is positioned as a Lakatosian research programme with specified empirical contact surfaces and explicit degeneracy conditions.

KEYWORDS: Process philosophy; Whitehead; Self-organisation; Constraint-closure; Consciousness; research programme; Living universe

1. INTRODUCTION: THE PERSISTENCE OF THE MIND–COSMOS PROBLEM

One of the most deep-seated difficulties confronting contemporary thought concerns a question that antedates modern science yet remains unresolved within it: how does mind belong to the world? On the standard picture inherited from seventeenth-century mechanism, matter is inert, extended, and governed by deterministic laws. Mind is interior and qualitative, appended to matter in an explanatorily opaque fashion. This picture has proved spectacularly productive in physics and engineering, yet it generates what David Chalmers termed the "hard problem"—the explanatory gap between a system's physical-functional description and the felt quality of experience.[1] The hard problem is hard, not because neuroscience is incomplete, but because the mechanistic picture is framed so as to exclude subjectivity from the outset. No accumulation of third-personal data concerning firing rates or informational-dynamic organisation, taken by itself, explains why there is *something it is like* to be a conscious creature.

This paper does not attempt to dissolve the hard problem by philosophical sleight of hand. It argues, rather, that the problem appears especially intractable so long as one starts from the mechanistic picture. Following Alfred North Whitehead, it explores an alternative *process* standpoint in which events, relations, and actual occasions—rather than inert substances—may be taken as ontologically basic. From within such a standpoint, consciousness can at least be reconsidered not as a late and wholly anomalous intrusion into an otherwise mindless world, but as a mode of organisation arising within a universe characterised by activity, self-organisation, and relational process. Whether such a standpoint ultimately warrants a stronger claim about interiority is treated here not as settled in advance, but as a matter for explicit metaphysical argument.[2]

The CAT (Cosmic Architecture Theory) framework formulates this alternative starting point through a minimal $\mathcal{A}$–μ schema. $\mathcal{A}$ represents global constraint architecture at scales from molecular to cosmological; μ represents a local mode of integrated informational-dynamic organisation embedded within it. The central working hypothesis is that life and consciousness are best understood within a self-organising, non-equilibrium universe rather than as anomalies added to inert reality. In its stronger, explicitly speculative form, CAT asks whether they may be interpreted as structural tendencies of such a cosmos. That claim is a metaphysical proposal, assessed by coherence, explanatory scope,

[1] David J. Chalmers, 'Facing Up to the Problem of Consciousness', *Journal of Consciousness Studies*, vol. 2, no. 3, 1995, pp. 200–219.
[2] Alfred North Whitehead, *Process and Reality: An Essay in Cosmology*, corrected edn, ed. David Ray Griffin and Donald W. Sherburne, New York, Free Press, 1978.

and heuristic fruitfulness.

The claims are organised in three layers. As a metaphysical orientation, CAT takes process rather than substance as basic. As a descriptive schema, it treats multi-scale organisation through $\mathcal{A}$–μ relations. As a research programme, it generates questions that contact empirical inquiry through operational proxies. Strong ontological theses are presented in Section 6, but, as Section 8 explains, they function as a Lakatosian hard core rather than as individually falsifiable predictions.

The argument proceeds as follows. Sections 2–3 develop the process and constraint-closure background; Sections 4–5 examine consciousness and formal limits; Section 6 presents the CAT architecture; Sections 7–8 address philosophical positioning, objections, and methodological status; and Section 9 concludes.

2. THE UNIVERSE AS PROCESS: NON-EQUILIBRIUM AND SELF-ORGANISATION

2.1. From Substance to Process

The dominant tradition in Western metaphysics has characterised reality in terms of substances that persist through time and bear properties. Whitehead identified this as the "fallacy of misplaced concreteness"—the treatment of the abstractions of Newtonian mechanics as though they were the concrete constituents of reality.[3] Against this, he proposed that the fundamental units of reality are *actual occasions*—events in which the cosmos "takes account of" its own past and advances creatively towards novelty. Each actual occasion involves a kind of *prehension*—an active grasping of its environment—and thus a proto-experiential dimension.

This is neither animism nor vitalism. On Whitehead's account, the categories appropriate to the description of conscious experience—responsiveness, self-relatedness, environment-dependent selectivity—may be read as generalisations

[3] Alfred North Whitehead, *Science and the Modern World*, New York, Free Press, 1967 (original work published 1925), pp. 50–55.

of features that appear, in attenuated form, across self-organising processes.[4]

Seen in this light, contemporary cosmology discloses a universe in which process is constitutive rather than incidental. The cosmos did not begin in static equilibrium and develop structure by accident. It began in a state of extremely low gravitational entropy and has continued to generate structure through the gradual release of its thermodynamic gradient.[5] Stars, galaxies, planets, biospheres, and technological civilisations are manifestations of a single cascade of structure-formation, driven by the thermodynamic gradient between the hot initial state and the cold, expanding void. As Eric Chaisson's account of "cosmic evolution" traces in detail, each stage of this cascade is characterised by increasing energy rate density.[6] It should be emphasised that this cascade is guided by no teleological intention; it unfolds under thermodynamic constraint.

2.2. Dissipative Structures and the Architecture of Order

The thermodynamics of this cascade was elucidated by Ilya Prigogine and his collaborators.[7] Prigogine showed that systems driven sufficiently far from thermodynamic equilibrium by energy flows can spontaneously generate and maintain ordered, coherent structures—convection cells, chemical oscillators, biological membranes. Such *dissipative structures* are not anomalies that violate the second law; they are manifestations of that law under far-from-equilibrium conditions. The cosmos is not merely running down toward heat death. Even as its overall entropy continues to increase, it is also generating, in local dissipative pockets, ever more complex structures.

In the CAT framework, this is expressed by saying that the universe constitutes a *constraint architecture* $\mathcal{A}$—a multi-scale pattern of boundary conditions, energy gradients, and causal organisation that determines which configurations of matter and information are stable. As Kleidon analysed in detail

[4] Alfred North Whitehead, *Adventures of Ideas*, New York, Free Press, 1967 (original work published 1933), pp. 176–179.
[5] Roger Penrose, *The Road to Reality: A Complete Guide to the Laws of the Universe*, London, Jonathan Cape, 2004, pp. 726–732.
[6] Eric J. Chaisson, *Cosmic Evolution: The Rise of Complexity in Nature*, Cambridge MA, Harvard University Press, 2001.
[7] Grégoire Nicolis and Ilya Prigogine, *Self-Organization in Non-Equilibrium Systems: From Dissipative Structures to Order through Fluctuations*, New York, Wiley-Interscience, 1977.

for the Earth system, each level of this architecture tends to self-organise in ways that often increase energy throughput.[8] $\mathcal{A}$ is neither a material field nor a hidden variable. It is an abstract description of the relational structure of a far-from-equilibrium cosmos. Stars, atmospheric circulation cells, and biospheric trophic networks are manifestations of different levels of $\mathcal{A}$, layered as nested constraint architectures.

2.3. Suggestions from Network Topology

Complementing thermodynamics, complex-network science has shown that certain topological properties—modularity, hub-and-spoke structure, small-world connectivity—recur across physically diverse systems: the cosmic web, the neural connectome, metabolic networks, and the internet.[9] Particularly noteworthy is the work of Vazza and Feletti, who quantitatively compared the topological properties of the cosmic web and neural networks, reporting striking similarities in spectral density, inter-node connectivity, and clustering coefficients.[10] This convergence is suggestive, but it demands cautious interpretation. To claim that the cosmic web and neural networks are governed by the same laws would be a clear overreach. Yet the recurrence of similar topological solutions under far-from-equilibrium conditions at different scales poses an empirical question: is this structural similarity mere coincidence, or does it reflect some constraint common to non-equilibrium self-organisation? What the paper draws from this evidence is, at most, *heuristic support*—not direct evidence of ontological continuity. The CAT framework provides a descriptive language for exploring this possibility further.

[8] Axel Kleidon, 'Life, Hierarchy, and the Thermodynamic Machinery of Planet Earth', *Physics of Life Reviews*, vol. 7, no. 4, 2010, pp. 424–460.
[9] Albert-László Barabási, *Network Science*, Cambridge, Cambridge University Press, 2016.
[10] Franco Vazza and Alberto Feletti, 'The Quantitative Comparison Between the Neuronal Network and the Cosmic Web', *Frontiers in Physics*, vol. 8, 2020, Article 525731.

3. LIFE AS CONSTRAINT-CLOSURE: TELEODYNAMICS AND AUTOPOIESIS

3.1. The Problem of Teleology

The transition from physics to biology has traditionally posed a challenge for the mechanistic worldview. Organisms *behave* as though they had purposes, yet mechanism allows no place for genuine purpose in nature. The standard Darwinian response is to show how apparent purposiveness arises from blind variation and selection. This is partially successful, yet it does not explain what distinguishes a living system from a merely dissipative one, nor why living systems maintain their own boundary conditions in ways that flames and hurricanes cannot.

Whitehead's process philosophy offers a different approach. The CAT framework draws on Whitehead for its event ontology, its notion of creativity, and its insight that series of events form sustained patterns—'societies.' For Whitehead, every actual occasion, however minimal, involves what he calls a *subjective aim*.[11] CAT reads 'subjective aim' as a minimal orientation towards each occasion's own immediate outcome. From this standpoint, a living organism may be construed as a *society* of actual occasions with a particularly rich internal organisation, sustained across successive events by a dominant nexus.[12] What distinguishes it from a rock is not the possession of an additional vital entity, but the depth of coordinated integration achieved beyond its constituent events.

3.2. Constraint-Closure and Teleodynamics

Contemporary theoretical biology has developed naturalistic accounts that partially converge with Whitehead's intuitions. Terrence Deacon's concept of *teleodynamics* offers a naturalistic account of what is distinctive about living systems: they can be understood as achieving a form of *constraint-closure*—a network of mutually enabling constraints that help regenerate the conditions of their own persistence.[13] A flame sustains itself so long as fuel and oxygen are externally

---

[11] Whitehead, *Process and Reality: An Essay in Cosmology*.

[12] Whitehead, *Process and Reality: An Essay in Cosmology*.

[13] Terrence W. Deacon, *Incomplete Nature: How Mind Emerged from Matter*, New York, W. W. Norton & Company, 2011, pp. 264–308.

supplied; an organism actively maintains the boundary conditions that make its own metabolism possible. This self-referential loop constitutes a minimally naturalised form of purposiveness—not an externally imposed purpose, but an intrinsic organisational tendency. An ontological tension must, however, be noted. Deacon is also critical of Whiteheadian panexperientialism, taking it to risk reintroducing mental properties too early into nature.[14] The CAT framework does not overlook this conflict: what it borrows is Deacon's *dynamical analysis* of how constraint-closure generates a minimal form of purposiveness, not his emergentist ontology. That analysis is reinterpreted within a Whiteheadian process ontology: the self-maintenance realised through constraint-closure is understood not as a case in which experiential qualities are *added after the fact*, but as a process in which they *gradually become manifest* as integration deepens. Independent closure-of-constraints accounts in theoretical biology provide a way of formulating the organisational pattern without relying solely on Deacon's broader metaphysical commitments.[15]

The theory of autopoiesis, developed by Maturana and Varela, makes a structurally parallel point. A living system is a network of processes that continually produces the components sustaining the network and its boundary.[16] An organism is not a container of molecules but a *unity of process*. Its identity resides in recursive organisation, not in any particular material substrate.

Building on the structural and dynamical definitions of constraint-closure provided by these precedents, the CAT framework reformulates it as the $\mathcal{A}$–μ architecture. At each scale n, a system is described as a pair comprising a global constraint architecture $\mathcal{A}_n$ and a self-sustaining local mode μ_n realised within it. The mode μ_n is not a passive trajectory through the phase space of $\mathcal{A}_n$; it is a self-maintaining dynamical regime that selectively exploits—and in some cases modifies in return—a local region of $\mathcal{A}_n$. The relation $\mu_n \hookrightarrow \mathcal{A}_{n+1}$, in turn, expresses the fact that a local mode at scale n contributes to the formation and maintenance of a more global constraint architecture at scale n+1 while itself being conditioned within it. Here $\hookrightarrow$ does not introduce a strict mathematical

[14] Deacon, *Incomplete Nature: How Mind Emerged from Matter*, pp. 77–79.

[15] Maël Montévil and Matteo Mossio, 'Biological Organisation as Closure of Constraints', *Journal of Theoretical Biology*, vol. 372, 2015, pp. 179–191.

[16] Humberto R. Maturana and Francisco J. Varela, *Autopoiesis and Cognition: The Realization of the Living*, Dordrecht, D. Reidel, 1980, pp. 73–135.

mapping; it is a symbolic expression indicating a non-symmetric, directed relation in which an organisational pattern at a lower scale is nested within the constraint environment at a higher scale while retaining its character. Accordingly, $(\mathcal{A}_n,\mu_n)$ denotes the constraint–local-mode pair at a single scale, and $\mu_n \hookrightarrow \mathcal{A}_{n+1}$ denotes the cross-scale embedding relation. This nested hierarchy is not a logical deduction but an empirically grounded structural claim about the way biological organisation achieves robustness through multi-scale coupling.

3.3. Distributed Intelligence and the Extension of Mind

If life is constraint-closure, cognition is constraint-closure that has become *self-referentially recursive*. A system not only maintains itself—it also forms internal models of itself and its environment, thereby guiding and improving its own maintenance. Karl Friston's Free-Energy Principle (FEP) formulates this process in information-theoretic terms.[17] According to the FEP, self-organising systems maintain a statistical boundary (Markov blanket) that partitions them from their environment and behave so as to minimise prediction error (free energy) of sensory input. Translated into the CAT vocabulary, maintenance of the Markov blanket can be understood as one aspect of the self-sustaining dynamics of μ within a local region of $\mathcal{A}$. The FEP thus provides a *dynamical description* of how constraint-closure forms and updates predictive models of its environment; this role is revisited in Section 8. This extension of self-reference is not confined to neural systems. The slime mould *Physarum polycephalum*, lacking any neurons whatsoever, can nevertheless find minimum-length routes through a maze—an optimisation that, in human engineers, would ordinarily require deliberate reasoning.[18] Such cases suggest that informational integration and adaptive response can be realised in distributed material configurations independently of the presence or absence of a centralised brain. As the extended-mind thesis of Clark and Chalmers maintains, cognitive processes are not confined within the brain but can be realised across body and environment.[19] Evan Thompson develops a related perspective, arguing that mind is inseparable from life—

[17] Karl Friston, 'The Free-Energy Principle: A Unified Brain Theory?', *Nature Reviews Neuroscience*, vol. 11, no. 2, 2010, pp. 127–138.
[18] Toshiyuki Nakagaki, Hiroyasu Yamada and Ágota Tóth, 'Maze-Solving by an Amoeboid Organism', *Nature*, vol. 407, no. 6803, 2000, p. 470.
[19] Andy Clark and David J. Chalmers, 'The Extended Mind', *Analysis*, vol. 58, no. 1, 1998, pp. 7–19.

cognition is always an embodied, environmentally embedded, living process.[20]

In the CAT framework, these phenomena are understood as different points along a continuum of $\mathcal{A}$–μ coupling. Yet it does not follow from FEP, slime-mould behaviour, or the extended-mind thesis that the existence of experience is thereby proved. What such cases at least suggest is that nature contains distributed systems exhibiting boundary-maintenance, recursive self-relation, and temporal integration. The argumentative status of the next step must therefore be made explicit. CAT adds a *speculative bridge principle*: a sufficiently integrated and recursively self-maintaining organisation is a necessary, though not sufficient, structural condition for experience, and may in some cases possess an inner aspect not exhausted by third-personal functional description. This bridge principle is not read directly off the empirical sciences; it is a metaphysical postulate whose merit depends on the coherence and fruitfulness of the framework it enables. The central vulnerability of the framework is accordingly clear: the move from complex organisation to experience is not a deductive inference. CAT presents that move openly as a speculative bridge principle rather than as a result established by the cited sciences themselves. Non-neural distributed systems are treated not as instances of finished consciousness, but as lower-bound cases that help identify the kinds of natural processes that may instantiate some structural preconditions for experience.

4. CONSCIOUSNESS AS AN $\mathcal{A}$–M CONFIGURATION

4.1. The Hard Problem Reconsidered

As Chalmers argues, one can, in principle, explain all the functional, behavioural, and neural correlates of consciousness—why particular stimuli elicit particular responses, why information is integrated in particular ways, why subjects report particular contents—without thereby explaining why any of this is accompanied by *experience*.[21] The gap appears not merely epistemological but conceptual: no third-personal physical description seems, by itself, to entail a first-personal

[20] Evan Thompson, *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*, Cambridge MA, Belknap Press of Harvard University Press, 2007, chap. 'Autopoiesis: The Organization of the Living'.

[21] David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory*, Oxford, Oxford University Press, 1996.

experiential fact.

The standard physicalist responses deny that the gap is real. Either experience reduces to functional organisation (functionalism), or the intuition generating the hard problem is itself mistaken (illusionism, as defended by Daniel Dennett and Keith Frankish).[22] Both responses have powerful advocates, but both face well-known objections.[23] Among the objections commonly raised in this broader dialectic are the familiar thought that functionalism struggles to account for the specific qualitative character of experience (qualia), and the charge that illusionism risks undercutting the very appearance it seeks to explain away.

The CAT framework proposes a different strategy. Rather than attempting to derive consciousness from a prior description of matter, it follows Whitehead in treating experience not as an anomaly requiring derivation but as a *primordial feature* of process. At least four levels must, however, be distinguished here. First, a *proto-experiential dimension* consisting of responsiveness, self-relatedness, and selectivity. Second, *phenomenal consciousness* as unified feeling. Third, *self-referential consciousness* involving the modelling of self and environment. Fourth, *reflective consciousness* involving linguistic and conceptual reflection. The CAT framework posits only the first of these at the base of nature; the second through fourth are realised in graduated stages as constraint-closure, informational integration, and self-referential depth increase. The position of this paper is accordingly not one that attributes human consciousness as such to everything, but a process-philosophical panexperientialism that distinguishes the minimal form of experience from its higher-order developments. This approach also differs from mechanism-specific quantum theories of consciousness that locate conscious episodes in particular quantum processes.[24]

4.2. Integrated Information Theory and the $\mathcal{A}$–μ Schema

Giulio Tononi's Integrated Information Theory (IIT) provides a useful

[22] Daniel C. Dennett, *Consciousness Explained*, Boston, Little, Brown and Company, 1991; Keith Frankish (ed.), *Illusionism as a Theory of Consciousness*, Exeter, Imprint Academic, 2017.

[23] For a systematic critique of physicalist positions, see Philip Goff, *Consciousness and Fundamental Reality*, Oxford, Oxford University Press, 2017.

[24] Stuart Hameroff and Roger Penrose, 'Orchestrated Reduction of Quantum Coherence in Brain Microtubules: A Model for Consciousness', *Mathematics and Computers in Simulation*, vol. 40, no. 3–4, 1996, pp. 453–480.

operational framework for the CAT account of consciousness.[25][26] IIT proposes that consciousness corresponds to *integrated information*—the irreducible causal power of a system as a unified whole, quantified by a measure called Φ. A system with high Φ cannot be decomposed into independent subsystems without loss of information. It is genuinely more than the sum of its parts.

The CAT framework does not endorse IIT as a complete theory; it treats Φ and related measures as *operational proxies* for consciousness-related $\mathcal{A}$–μ configurations. Although IIT tends to treat Φ as a physically real ontological entity, the CAT framework does not accept this commitment; Φ is invoked only as one of the most operationally advanced proxies currently available. IIT 4.0 takes causal irreducibility as its core measure.[27] The claim, then, is that conscious experience correlates with μ modes that achieve high integration, differentiation, and recursive self-referential coupling within a local region of $\mathcal{A}$—not a definition of consciousness but a structural constraint on the physical configurations that may underlie experience.

The CAT framework does not adopt IIT's Exclusion Postulate, which attributes consciousness to the single complex with maximal Φ among overlapping physical substrates, excluding sub-systems and super-systems. The exclusion postulate is a constitutive element of IIT's mathematical structure: it serves to determine the boundary of consciousness uniquely and to prevent combinatorial explosion; removing it deprives Tononi's strict Φ of its definitional presupposition. The proxies to which the framework appeals when it admits the possibility of multi-scale coexisting experiential modes are therefore not IIT's strict ΦMax itself but more general measures of informational integration from which the exclusivity condition has been removed—complexity measures such as the causal density of Seth and colleagues.[28] This distinction allows the CAT

---

[25] Giulio Tononi, Melanie Boly, Marcello Massimini and Christof Koch, 'Integrated Information Theory: From Consciousness to Its Physical Substrate', *Nature Reviews Neuroscience*, vol. 17, no. 7, 2016, pp. 450–461.

[26] For a recent overview of competing theories of consciousness, including IIT, global workspace theories, higher-order theories, and predictive-processing approaches, see Anil K. Seth and Tim Bayne, 'Theories of Consciousness', *Nature Reviews Neuroscience*, vol. 23, no. 7, 2022, pp. 439–452.

[27] Larissa Albantakis et al., 'Integrated Information Theory (IIT) 4.0: Formulating the Properties of Phenomenal Existence in Physical Terms', *PLOS Computational Biology*, vol. 19, no. 10, 2023, e1011465.

[28] Anil K. Seth, Adam B. Barrett and Lionel Barnett, 'Causal Density and Integrated Information as Measures of Conscious Level', *Philosophical Transactions of the Royal Society A*, vol. 369, no. 1952, 2011, pp. 3748–3767.

framework to draw on the structural insights of IIT without committing itself to its exclusionary single-scale ontology.

4.3. Consciousness, Cosmopsychism, and the CAT Framework

Recent philosophy has witnessed a revival of *cosmopsychism*—the view that the cosmos as a whole is the primary seat of consciousness.[29] Cosmopsychism is motivated by the combination problem: if the fundamental unit is the cosmos rather than particles with proto-experience, the combination problem dissolves (replaced by a "decomposition problem").

The CAT framework is structurally consonant with cosmopsychist motivations but does not commit to a unified cosmic subject. Any teleological interpretation—such as "the cosmos produced consciousness in order to know itself"—is emphatically rejected. The global constraint architecture $\mathcal{A}$ is not itself a subject, but the physical-informational structure within which subjects are realised as μ modes. Individual minds are locally integrated patterns arising when $\mathcal{A}$–μ coupling achieves sufficient self-referential depth. The framework loosely recalls the position on which the physical and the experiential are two descriptions of a single reality, though this is a structural analogy, not a strict interpretive claim about Spinoza.[30]

5. FORMAL LIMITS AND THE QUESTION OF ONTOLOGICAL OPENNESS

5.1. Gödel, Turing, and the Limits of Formalisation

A persistent temptation in the philosophy of science is to equate a complete theory of nature with a complete *formalisation* of nature—a finite set of axioms and inference rules from which all physical truths could, in principle, be derived. Gödel's incompleteness theorems show that this ideal cannot be realised for any sufficiently rich formal system.[31] For any consistent system rich enough to encode

[29] Itay Shani and Joachim Keppler, 'Beyond Combination: How Cosmic Consciousness Grounds Ordinary Experience', *Journal of the American Philosophical Association*, vol. 4, no. 3, 2018, pp. 390–410.
[30] Baruch Spinoza, *Ethics*, trans. Edwin Curley, London, Penguin Classics, 1996 (original work published 1677).
[31] Kurt Gödel, 'Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I', *Monatshefte für Mathematik und Physik*, vol. 38, 1931, pp. 173–198.

arithmetic, there exist true propositions that the system cannot itself prove. Turing's undecidability result extends the point to computation: no general algorithm can determine, for an arbitrary program and input, whether that program will halt.[32]

Roger Penrose argued that these results may bear on the philosophy of mind: in his view, if human mathematical intuition can recognise certain Gödelian truths that no algorithm can establish within a formal system, then human mathematical thought cannot be fully captured in purely computational terms.[33] This argument remains controversial; critics have particularly questioned the inference from Gödel's theorem to the non-computability of mind. The CAT framework does not rely on the mind argument itself. Instead, it draws a more limited conclusion: that principled caution is warranted against the ideal of reducing nature to a completely closed and finite formal description. This does not immediately deduce the ontological openness of the cosmos, but it does provide a philosophical reason for considering that possibility.

5.2. Undecidability and Computational Irreducibility in Physics

The bearing of these logical limits on physics need not be merely metaphorical. Cubitt, Perez-Garcia, and Wolf showed that the question whether a quantum spin lattice has a spectral gap—a foundational question in condensed-matter physics—is undecidable: no algorithm exists for solving this class of problems.[34] Bennett's concept of *logical depth* provides a complementary perspective: some structures, though briefly describable, require unavoidably long computational histories.[35] Kauffman and Roli, on related grounds, argue that the world is not a theorem: the space of biological possibilities is unprestatable, and the trajectory of evolution cannot, in general, be derived from any fixed set of initial conditions

---

[32] Alan M. Turing, 'On Computable Numbers, with an Application to the Entscheidungsproblem', *Proceedings of the London Mathematical Society*, series 2, vol. 42, no. 1, 1937, pp. 230–265.

[33] Roger Penrose, *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*, Oxford, Oxford University Press, 1989, chap. 4.

[34] Toby S. Cubitt, David Perez-Garcia and Michael M. Wolf, 'Undecidability of the Spectral Gap', *Nature*, vol. 528, no. 7581, 2015, pp. 207–211.

[35] Charles H. Bennett, 'Logical Depth and Physical Complexity', in Rolf Herken (ed.), *The Universal Turing Machine: A Half-Century Survey*, Oxford, Oxford University Press, 1988, pp. 227–257.

and laws.[36] In a broader context, Kauffman and Gare argue that a process-philosophical view of nature is needed if one is to move beyond the Cartesian–Newtonian paradigm and recover life and humanity.[37]

These are not three instances of a single formal theorem. Gödel and Turing concern formal systems and algorithms; Cubitt, Perez-Garcia, and Wolf concern a class of quantum many-body problems; Kauffman and Roli address biological possibility spaces. CAT therefore does not infer a metaphysical thesis directly from any one of them. The point is comparative and negative: each blocks the simple picture that novelty is already captured by a single finite formal closure. The positive claim of ontological openness arises only when these considerations are read with the process ontology and non-equilibrium account developed in Sections 2–3.

These results do not by themselves entail ontological openness or cosmic creativity. At least three interpretations remain available: an epistemic reading, on which limits constrain description rather than being; a deterministic but irreducibilist reading, on which novelty is practically uncompressible within a closed ontology; and the CAT reading, on which these limits, biological unprestatability, and recurrent novelty in non-equilibrium systems jointly support openness as a coherent metaphysical hypothesis.

CAT does not claim that the formal results deduce this third reading. It argues only that they make it a defensible metaphysical option and, in conjunction with the broader processual picture developed here, a potentially fruitful one. Ontological openness is therefore presented not as a theorem but as a comparative philosophical hypothesis.

5.3. Creativity and the Cosmos

Whitehead's metaphysics provides a vocabulary for articulating this working hypothesis. For Whitehead, *creativity* is the ultimate metaphysical principle: "the many become one, and are increased by one."[38] The CAT framework rereads

[36] Stuart A. Kauffman and Andrea Roli, 'The World Is Not a Theorem', *Entropy*, vol. 23, no. 11, 2021, Article 1467.
[37] Stuart A. Kauffman and Arran Gare, 'Beyond Descartes and Newton: Recovering Life and Humanity', *Progress in Biophysics and Molecular Biology*, vol. 119, no. 3, 2015, pp. 219–244.
[38] Whitehead, *Process and Reality: An Essay in Cosmology*, p. 21.

this insight in the light of computability theory and complexity research. Non-computability and undecidability do not prove cosmic creativity, but they do constrain any picture that portrays the world as a completely closed machine, and they leave open an interpretation compatible with Whitehead's concept of creativity. These formal limits may accordingly be read not as defects to be overcome, but as a philosophical occasion for understanding nature as an open process.

6. THE CAT ARCHITECTURE: $\mathcal{A}$, M, AND MULTI-SCALE COUPLING

6.1. Core Propositions

The CAT framework is organised around a minimal conceptual schema. Two symbols, $\mathcal{A}$ and μ, serve as conceptual shorthands for recurrent aspects of multi-scale organisation.

• $\mathcal{A}$ represents the *constraint architecture* of a system at a given scale: a multi-scale pattern of boundary conditions, causal organisation, and structured energy-informational flows that determines which configurations are stable.

• μ represents an *embedded local mode*: a dynamical regime in which relatively integrated informational-dynamic organisation is realised under the constraints imposed by $\mathcal{A}$.

To illustrate the $\mathcal{A}$–μ schema in operational terms, Table 1 presents representative interpretations at characteristic scales.

Table 1. Operational interpretation of the $\mathcal{A}$–μ schema: illustrative (not exhaustive) ($\mathcal{A}_n$, μ_n) pairs at each scale. Evidential quality varies: cell and organism scales are empirically robust, whereas molecular and planetary scales are more hypothetical.

Scale	$\mathcal{A}$ (constraint architecture)	μ (local mode)
Molecular	Thermodynamic gradients, chemical-bond possibility space	Autocatalytic cycles, proto-metabolic networks
Cellular	Membrane permeability, metabolic	Local metabolic regimes, homeostatic

	constraints, genetic regulatory networks	maintenance patterns
Organism	Neural architecture, bodily constraints	Perception–action loops, predictive processing modes
Ecosystem	Trophic structure, biogeochemical cycles	Interspecific interaction patterns, adaptive resource allocation
Planetary	Geophysical boundary conditions, atmosphere–ocean–crust coupling	Biospheric environmental feedback, plus additional modulation by the technosphere

These are not new fundamental fields or hidden variables; they are analytical concepts that permit a cross-cutting description of the nested, non-equilibrium organisation exhibited by the cosmos. What Table 1 presents are ($\mathcal{A}_n$,μ_n) pairs at each scale. The central structural claim, by contrast, is expressed by the relation $\mu_n \hookrightarrow \mathcal{A}_{n+1}$: a local mode at scale n, retaining its organisational character, contributes to the formation and maintenance of the constraint architecture at scale n+1 while being conditioned within it. For example, cellular-level metabolic modes are nested within organism-level physiological constraints, and organism-level behavioural modes within ecosystem-level constraints. This nested hierarchy is not a logical deduction but an empirically grounded structural claim about how biological organisation achieves robustness through multi-scale coupling.

The core propositions of the framework are as follows.

(**P1**) **Non-equilibrium process.** The universe is a far-from-equilibrium historical process in which structures are formed, sustained, and dissolved across multiple scales. The global constraint architecture $\mathcal{A}$ evolves as these structures form and modify boundary conditions.

(**P2**) **Constraint-closure and teleodynamics.** Local modes μ exhibit genuine organisational autonomy when they achieve constraint-closure—

recursively maintaining and reproducing the constraints that make their own persistence possible. This is the minimal sense in which a physical system can naturalistically be said to exhibit goal-directed behaviour. Life is the paradigmatic case, but the CAT framework treats this as a graded rather than threshold property.

(**P3**) **Bidirectional coupling.** The relation between $\mathcal{A}$ and μ is not merely top-down. Local modes μ modify the global constraint architecture $\mathcal{A}$ through their actions, products, and structural effects. Organisms altering ecosystems, and technological civilisations affecting a planet's geosphere, are representative examples. Planetary-scale cases, however, are not as well established as constraint-closure at the cellular or organismal level and remain, at this stage, extensions under investigation. Post-Gaian Earth-system research and recent discussions of "intelligence as a planetary-scale process" provide, at minimum, empirically tractable lines of inquiry.[39] Extrapolation to cosmic scales (the possibility that sufficiently integrated civilisations might alter cosmic-scale constraints) remains a speculative extension, positioned for the present as a conceptual possibility only. This bottom-up feedback is key to understanding the cosmos not as a closed mechanical system but as a multi-scale self-modifying process.

(**P4**) **Consciousness as deep $\mathcal{A}$–μ integration.** Conscious experience correlates with μ modes that achieve high internal integration, differentiation, and self-referential depth within a local region of $\mathcal{A}$. Consciousness is not a separate entity injected into the cosmos from outside. On the working hypothesis advanced here, it may be understood as the "inner view" of a particularly rich $\mathcal{A}$–μ configuration—what it is like to be a self-maintaining, self-referential process.

(**P5**) **Ontological openness.** CAT adopts, as a working hypothesis, that no finite description of $\mathcal{A}$ or μ exhausts the possibilities latent in the processual nature of the cosmos. This is not a proven conclusion but a hypothesis connecting formal and empirical considerations. The openness in question is treated not merely as descriptive failure, but as a structural possibility to be reckoned with in

---

[39] James E. Lovelock and Lynn Margulis, 'Atmospheric Homeostasis by and for the Biosphere: The Gaia Hypothesis', *Tellus*, vol. 26, no. 1–2, 1974, pp. 2–10; Adam Frank, David Grinspoon and Sara Walker, 'Intelligence as a Planetary Scale Process', *International Journal of Astrobiology*, vol. 21, no. 2, 2022, pp. 47–61.

understanding a self-organising universe.

6.2. Three Working Theses

The five propositions entail three substantive theses about the nature of cosmos, life, and mind—the core claims of the CAT framework.

(T1) The universe may be interpreted as a self-organising, life-like system. "Living" here is not restricted to biological constraint-closure, but refers to systems organised by non-equilibrium processes, continuous energy throughput, multi-scale feedback, and the generation of novel forms. The hypothesis is not that the cosmos is literally an organism, but that the same conceptual vocabulary illuminates organisation across scales. "Living" is a stipulative designation for this cluster of properties; the substantive claim stands independently of whether the label is retained.

(T2) Consciousness is not an afterthought. CAT hypothesises that the structural preconditions of experience are continuous with non-equilibrium self-organising processes. A proto-experiential dimension of responsiveness, self-relatedness, and selectivity may have been open within such processes from the outset, developing into phenomenal and reflective consciousness as constraint-closure and self-referential integration deepen. This is a working hypothesis, not a demonstrated result. Proto-experience is posited within CAT as a minimal processual dimension, while constraint-closure and recursive integration are necessary rather than sufficient conditions for phenomenal consciousness; the transition to phenomenal unity remains open.

(T3) Matter, life, and consciousness may be understood within a continuous processual framework. CAT understands the transitions from physics to chemistry, biology, and consciousness as emergent phases in a processual continuum of increasing constraint-closure, informational integration, and self-referential depth. A bacterium differs genuinely from a quartz crystal, and a human mind from a bacterium; the claim is that these differences are degrees along continuous dimensions rather than ontological chasms. This does not eliminate the explanatory gap, but reframes it as a question about degree of integration rather than categorical emergence.

6.3. Operational Proxies and Research Directions

The CAT framework is not primarily an empirical theory; it is a *natural-philosophical map* that clarifies assumptions and generates tractable research questions. It is not, however, empirically idle. The following operational anchoring points indicate how its claims make contact with empirical inquiry.

For T1, relevant proxies include entropy-production rates in cosmic structure-formation, network topology, and multi-scale analyses of energy throughput; at the planetary scale, the identification of constraint-maintaining feedback loops (extensions of Gaia theory)[40] and Frank et al.'s framework of "intelligence as a planetary-scale process"[41]; and comparative studies of self-organisation across astrobiological contexts.

For T2, relevant proxies include integration-sensitive measures, including Φ-derived and causal-density indicators, applied across biological scales; adversarial experimental programmes comparing theories of consciousness[42]; and explorations of proto-experiential organisation in non-neural systems.

For T3, relevant proxies include theoretical and empirical research on origin-of-life transitions understood as the emergence of constraint-closure; comparative studies of teleodynamic organisation across complexity hierarchies; and the development of integrated information-theoretic frameworks spanning physical, biological, and cognitive domains.

Honesty within a research programme requires specifying what would weaken the framework. P1/T1 would be weakened if multi-scale constraint indicators showed no robust association with the persistence of structures after relevant controls. P2/T3 would be weakened if systems without mutually maintaining constraints displayed the same degree of long-term self-maintenance, adaptive error-correction, and boundary preservation as closure systems. P4/T2 would be weakened if pre-registered comparisons showed that integration-sensitive and recursive-coupling measures consistently failed to outperform simpler feedforward indicators or behavioural baselines in predicting

---

40 Lovelock and Margulis, 'Atmospheric Homeostasis by and for the Biosphere: The Gaia Hypothesis', pp. 2–10.
41 Frank, Grinspoon and Walker, 'Intelligence as a Planetary Scale Process', pp. 47–61.
42 Lucia Melloni et al., 'An Adversarial Collaboration Protocol for Testing Contrasting Predictions of Global Neuronal Workspace and Integrated Information Theory', *PLOS ONE*, vol. 18, no. 2, 2023, e0268577.

conscious level or content. Such results would not directly falsify the hard core, but would indicate a degenerative protective belt.

7. PHILOSOPHICAL POSITIONING AND OBJECTIONS

7.1. Whitehead and the Process Tradition

The CAT framework is best understood as a research programme that brings Whitehead's process philosophy into dialogue with contemporary science. It is situated within the tradition of process philosophy initiated by Whitehead and, more broadly, within the lineage of natural philosophy that understands nature as a generative, self-organising process rather than a machine.[43] It draws on Whitehead's processive conception of nature and, more specifically, on his critique of the "bifurcation of nature."[44] Cosmic creativity is accordingly understood as a structural possibility within non-equilibrium dynamics, presented as a working hypothesis rather than a proven conclusion.

Rainer Zimmermann's "systems metaphysics of emergence," which treats emergent properties as real features of systems qua systems, provides a further theoretical resource: the CAT framework may be read as an extension of this systems metaphysics to vital, cognitive, planetary, and cosmological architectures.[45]

7.2. Panpsychism, Panexperientialism, and Cosmopsychism

As discussed in Section 4.3, the CAT framework has affinities with panpsychist and cosmopsychist motivations without coinciding with either.[46] It treats experience as a relational phenomenon at the level of $\mathcal{A}$–μ configuration, not as

[43] For a contemporary reappraisal, see Kauffman and Gare, 'Beyond Descartes and Newton', pp. 219–244. For a systematic defence of speculative naturalism, see Arran Gare, 'Speculative Naturalism: A Manifesto', *Cosmos and History: The Journal of Natural and Social Philosophy*, vol. 10, no. 2, 2014, pp. 300–323. For an attempt to unsnarl the knot of dualism and reductionistic materialism through Kant and Whitehead, see Daniel A. Dombrowski, 'Unsnarling the Metaphysical Knot: Kant and Whitehead on the Real', *Cosmos and History: The Journal of Natural and Social Philosophy*, vol. 20, no. 2, 2024, pp. 263–271.

[44] Alfred North Whitehead, *The Concept of Nature*, Cambridge, Cambridge University Press, 1920, pp. 26–48.

[45] Rainer E. Zimmermann, *Metaphysics of Emergence. Part 1: On the Foundations of Systems*, Berlin, Xenomoi, 2015.

[46] Philip Goff, *Galileo's Error: Foundations for a New Science of Consciousness*, New York, Pantheon Books, 2019.

a property of individual particles as such, and it posits no unified cosmic subject from which individual minds are derived. Relative to emergent physicalism, it does not treat experience as a late categorical addition to an otherwise wholly non-experiential world. Its intended position is process-relational: proto-experience belongs to process in a minimal sense, while phenomenal consciousness requires specific forms of integration, constraint-closure, and self-referential organisation. The combination problem is addressed in Section 8.3.

7.3. Objections

"This is mere metaphor." The phrases "living universe" and "cosmic process" do carry a certain metaphorical charge. The operative content of the paper, however, lies in the structural claims set out in P1–P5 concerning non-equilibrium organisation, constraint-closure, informational integration, and multi-scale feedback. The heuristic vocabulary functions only as an aid to visualisation; if the phrase "living universe" is judged misleading, the substantive argument can stand without it.

"There is no empirical content." The framework is not a finished theory delivering single decisive predictions in the manner of general relativity. It is nevertheless not empirically vacuous. Its empirical contact lies in the operational proxies presented in Section 6.3 and in the comparative research and measurement programmes they organise. More exactly, empirical inquiry does not directly verify the metaphysical core of CAT; it tests whether the framework's protective belt remains progressive, coherent, and explanatorily fruitful.

"Parsimony." Standard physicalism may appear, at first glance, more parsimonious. CAT does not deny this in any straightforward ontological sense. Its claim is conditional: if one judges that the explanatory gap is genuine and that strict mechanistic closure leaves central features of mind insufficiently situated, then a less parsimonious but more integrative framework may be theoretically preferable.

"Panpsychism faces the combination problem." Conceded. The constraint-closure response is a partial answer, not a full solution.

Of these objections, the most serious for the present argument concerns the connection between undecidability and ontological openness developed in Section 5.

"This draws too much ontology from mathematical limits." The criticism has a

legitimate core and is addressed in detail in Section 5. At least three interpretations remain open: an epistemic reading on which formal and computational limits reveal constraints on description rather than on being; a deterministic but irreducibilist reading on which novelty is practically uncompressible while reality remains ontologically closed; and the CAT reading, on which such limits, together with biological unprestatability and recurrent novelty in non-equilibrium systems, support ontological openness as the most coherent metaphysical hypothesis. The framework presents this move as an IBE rather than a deduction, distinguishing a weak version (no finite description exhausts physical reality) from a strong version (genuine novelty), the latter remaining a working hypothesis.

8. METHODOLOGICAL STATUS AND LIMITS OF THE FRAMEWORK

This section clarifies both the methodological constraints the CAT framework imposes upon itself and the scope of its claims.

8.1. The CAT Framework Is Not a Popperian Scientific Theory

By Popper's criterion of falsifiability[47], the CAT framework is not itself a scientific theory: "the cosmos possesses a constraint architecture $\mathcal{A}$" is not directly falsified by any particular observation. The intellectual position it occupies is closer to what Imre Lakatos called a "scientific research programme"—a "hard core" of basic assumptions surrounded by a "protective belt" of auxiliary hypotheses assessed in terms of progressive or degenerative problem-shifts.[48]

The hard core of CAT comprises three points: (i) the cosmos is best described as process; (ii) matter, life, and consciousness are phases of the same processual continuum; (iii) the deepening of constraint-closure supplies the structural condition for consciousness. These are not individually falsifiable, but the protective belt includes auxiliary hypotheses with empirical contact surfaces. IIT, FEP, and Gaia theory function as operational tools, and Section 6.3 specified the

[47] Karl R. Popper, *The Logic of Scientific Discovery*, London, Routledge, 2002; first English edn 1959.
[48] Imre Lakatos, 'Falsification and the Methodology of Scientific Research Programmes', in Imre Lakatos and Alan Musgrave (eds.), *Criticism and the Growth of Knowledge*, Cambridge, Cambridge University Press, 1970, pp. 91–196.

conditions under which those contacts would indicate degeneracy.

8.2. The Relationship to IIT, FEP, and Autopoiesis: Hierarchical Complementation, Not Synthesis

A predictable objection is that the simultaneous reference to IIT, FEP, and autopoiesis amounts to an arbitrary patchwork. The three theories, however, operate at different levels of description—measurement, dynamics, and structure respectively—and the $\mathcal{A}$–μ schema makes their descriptive-level relations explicit rather than synthesising them. As noted in Section 4, CAT brackets the ontological differences among these theories and retains a research strategy employing multiple integration-sensitive proxies.

8.3. The Combination Problem: Clarification of Conditions, Not Solution

Any framework supporting a panexperientialist position inevitably faces the "combination problem": how do micro-level proto-experiential states integrate into the phenomenally unified consciousness that humans experience? This question is unavoidable for the CAT framework.[49]

Two versions of the combination problem must be distinguished.

Weak version (functional integration problem): How do many local processes come to form a single integrated functional unity? Constraint-closure, IIT's Φ, and FEP's hierarchical Markov blankets each provide partial answers at different descriptive levels. The CAT framework, through the relation $\mu_n \hookrightarrow \mathcal{A}_{n+1}$, makes explicit that this integration arises under the structural condition of embedding in higher-scale constraint architecture—not by mere aggregation of parts. This is not a complete answer to the weak version, but it contributes to a more precise formulation of the problem's structure.

Strong version (phenomenal integration problem): Why, when a physical network achieves a particular kind of closure, does *a single subjective boundary of consciousness* arise? This question is, in effect, another version of Chalmers's hard problem, and the CAT framework does not solve it. Nor do IIT, FEP, or any currently existing theory of consciousness.

[49] For an internal sorting of this problem within process-relational thought, tentatively arguing toward a layered ontology regarding the ontological status of multi-scale individuals, see Gregory J. Moses, 'Big Things from Little Things? The Problem of the Compound Individual', *Cosmos and History: The Journal of Natural and Social Philosophy*, vol. 20, no. 2, 2024, pp. 248–262.

What the CAT framework provides, for the strong version, is a *clarification of conditions*, not a solution. Constraint-closure, recursive integration, and embedding in higher-scale $\mathcal{A}$ supply structural preconditions under which unified experience may arise, but why a single subjective boundary appears remains unresolved. The framework's contribution is to shift the question from "why does consciousness arise from matter?" to "why does the proto-experiential dimension develop into what humans call consciousness only in specific forms of constraint-closure?"—a reformulation that indicates the kind of inquiry needed next.

8.4. The Distinctive Role of the CAT Framework: Questions That Become Askable

CAT's distinctive role is not to collect existing theories, but to connect dispersed questions under one vocabulary: to reread limit-results as bearing on openness, and to use the $\mathcal{A}$–μ schema for asking how lower-scale modes modify higher-scale constraint architectures across Earth-system science, theories of life, and consciousness research.

For example, CAT asks whether the origin-of-life transition is merely a new material organisation or the point at which a molecular μ mode begins to constitute a higher-scale constraint architecture that selects its own maintaining conditions. The framework does not answer this in advance; it makes the cross-scale relation visible as an object of inquiry.

9. CONCLUSION

This paper has articulated and partially defended CAT as a speculative natural-philosophical proposal for situating cosmos, life, and consciousness within one processual picture. Drawing together Whitehead's process philosophy, Prigogine's thermodynamics, Deacon's teleodynamics, Chalmers's analysis of consciousness, and formal limit-results, it has advanced three working theses: that the universe may be fruitfully interpreted as a life-like self-organising system (T1); that the structural preconditions of consciousness are continuous with non-equilibrium self-organising processes (T2); and that matter, life, and mind may be described within a continuous processual framework (T3).

The three-layered structure announced in the introduction—metaphysical orientation, descriptive schema, and research programme—has been developed through Sections 2–8, with the framework's operational proxies and degeneracy conditions made explicit in Section 6.3.

The framework is best understood as a philosophical research programme with empirical contact surfaces, not as a narrow Popperian theory. Its strongest claims remain programmatic rather than demonstrative. It does not solve the hard problem or the strong combination problem, nor deduce ontological openness from formal results. Its contribution lies in making the bridge principle explicit, organising disparate findings through the $\mathcal{A}$–μ schema, and specifying conditions under which the framework would lose explanatory force. From a processual and non-equilibrium standpoint, life and mind appear not as miraculous exceptions, but as phenomena to be understood within a broader generative process.

CAT is offered as a natural-philosophical map: it clarifies assumptions, situates empirical findings in a wider context, and motivates questions in cosmology, astrobiology, consciousness science, and philosophy of mind. It may be read as a contemporary practice of Whiteheadian speculative philosophy, re-examining the limits of the mechanistic view of nature. Its adequacy will be measured by the fruitfulness of the inquiry it inspires.

Affiliation: Independent Researcher
Address: Tokyo, Japan
Telephone: +81 90-4582-9192
Email: nobuyuki.otsubo.res@gmail.com
ORCID: 0009-0007-4004-6870

REFERENCES

Albantakis, Larissa et al., 'Integrated Information Theory (IIT) 4.0: Formulating the Properties of Phenomenal Existence in Physical Terms', *PLOS Computational Biology*, vol. 19, no. 10, 2023, e1011465.

Barabási, Albert-László, with Márton Pósfai, *Network Science*, Cambridge, Cambridge University Press, 2016.

Bennett, Charles H., 'Logical Depth and Physical Complexity', in Rolf Herken (ed.), *The Universal Turing Machine: A Half-Century Survey*, Oxford, Oxford University Press, 1988, pp. 227–257.

Chaisson, Eric J., *Cosmic Evolution: The Rise of Complexity in Nature*, Cambridge MA,

Harvard University Press, 2001.
Chalmers, David J., 'Facing Up to the Problem of Consciousness', *Journal of Consciousness Studies*, vol. 2, no. 3, 1995, pp. 200–219.
Chalmers, David J., *The Conscious Mind: In Search of a Fundamental Theory*, Oxford, Oxford University Press, 1996.
Clark, Andy and Chalmers, David J., 'The Extended Mind', *Analysis*, vol. 58, no. 1, 1998, pp. 7–19.
Cubitt, Toby S., Perez-Garcia, David and Wolf, Michael M., 'Undecidability of the Spectral Gap', *Nature*, vol. 528, no. 7581, 2015, pp. 207–211.
Deacon, Terrence W., *Incomplete Nature: How Mind Emerged from Matter*, New York, W. W. Norton & Company, 2011.
Dennett, Daniel C., *Consciousness Explained*, Boston, Little, Brown and Company, 1991.
Dombrowski, Daniel A., 'Unsnarling the Metaphysical Knot: Kant and Whitehead on the Real', *Cosmos and History: The Journal of Natural and Social Philosophy*, vol. 20, no. 2, 2024, pp. 263–271.
Frank, Adam, Grinspoon, David and Walker, Sara, 'Intelligence as a Planetary Scale Process', *International Journal of Astrobiology*, vol. 21, no. 2, 2022, pp. 47–61.
Frankish, Keith (ed.), *Illusionism as a Theory of Consciousness*, Exeter, Imprint Academic, 2017.
Friston, Karl, 'The Free-Energy Principle: A Unified Brain Theory?', *Nature Reviews Neuroscience*, vol. 11, no. 2, 2010, pp. 127–138.
Gare, Arran, 'Speculative Naturalism: A Manifesto', *Cosmos and History: The Journal of Natural and Social Philosophy*, vol. 10, no. 2, 2014, pp. 300–323.
Gödel, Kurt, 'Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I', *Monatshefte für Mathematik und Physik*, vol. 38, 1931, pp. 173–198.
Goff, Philip, *Consciousness and Fundamental Reality*, Oxford, Oxford University Press, 2017.
Goff, Philip, Galileo's Error: Foundations for a New Science of Consciousness, New York, Pantheon Books, 2019.
Hameroff, Stuart and Penrose, Roger, 'Orchestrated Reduction of Quantum Coherence in Brain Microtubules: A Model for Consciousness', *Mathematics and Computers in Simulation*, vol. 40, no. 3–4, 1996, pp. 453–480.
Kauffman, Stuart A. and Gare, Arran, 'Beyond Descartes and Newton: Recovering Life and Humanity', *Progress in Biophysics and Molecular Biology*, vol. 119, no. 3, 2015, pp. 219–244.
Kauffman, Stuart A. and Roli, Andrea, 'The World Is Not a Theorem', *Entropy*, vol. 23, no. 11, 2021, Article 1467.

Kleidon, Axel, 'Life, Hierarchy, and the Thermodynamic Machinery of Planet Earth', *Physics of Life Reviews*, vol. 7, no. 4, 2010, pp. 424–460.

Lakatos, Imre, 'Falsification and the Methodology of Scientific Research Programmes', in Imre Lakatos and Alan Musgrave (eds.), *Criticism and the Growth of Knowledge*, Cambridge, Cambridge University Press, 1970, pp. 91–196.

Lovelock, James E. and Margulis, Lynn, 'Atmospheric Homeostasis by and for the Biosphere: The Gaia Hypothesis', *Tellus*, vol. 26, no. 1–2, 1974, pp. 2–10.

Maturana, Humberto R. and Varela, Francisco J., *Autopoiesis and Cognition: The Realization of the Living*, Dordrecht, D. Reidel, 1980.

Melloni, Lucia et al., 'An Adversarial Collaboration Protocol for Testing Contrasting Predictions of Global Neuronal Workspace and Integrated Information Theory', *PLOS ONE*, vol. 18, no. 2, 2023, e0268577.

Montévil, Maël and Mossio, Matteo, 'Biological Organisation as Closure of Constraints', *Journal of Theoretical Biology*, vol. 372, 2015, pp. 179–191.

Moses, Gregory J., 'Big Things from Little Things? The Problem of the Compound Individual', *Cosmos and History: The Journal of Natural and Social Philosophy*, vol. 20, no. 2, 2024, pp. 248–262.

Nakagaki, Toshiyuki, Yamada, Hiroyasu and Tóth, Ágota, 'Maze-Solving by an Amoeboid Organism', *Nature*, vol. 407, no. 6803, 2000, p. 470.

Nicolis, Grégoire and Prigogine, Ilya, Self-Organization in Non-Equilibrium Systems: From Dissipative Structures to Order through Fluctuations, New York, Wiley-Interscience, 1977.

Penrose, Roger, The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics, Oxford, Oxford University Press, 1989.

Penrose, Roger, The Road to Reality: A Complete Guide to the Laws of the Universe, London, Jonathan Cape, 2004.

Popper, Karl R., *The Logic of Scientific Discovery*, London, Routledge, 2002; first English edn 1959.

Seth, Anil K., Barrett, Adam B. and Barnett, Lionel, 'Causal Density and Integrated Information as Measures of Conscious Level', *Philosophical Transactions of the Royal Society A*, vol. 369, no. 1952, 2011, pp. 3748–3767.

Seth, Anil K. and Bayne, Tim, 'Theories of Consciousness', *Nature Reviews Neuroscience*, vol. 23, no. 7, 2022, pp. 439–452.

Shani, Itay and Keppler, Joachim, 'Beyond Combination: How Cosmic Consciousness Grounds Ordinary Experience', *Journal of the American Philosophical Association*, vol. 4, no. 3, 2018, pp. 390–410.

Spinoza, Baruch, *Ethics*, trans. Edwin Curley, London, Penguin Classics, 1996 (original work published 1677).

Thompson, Evan, *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*, Cambridge MA, Belknap Press of Harvard University Press, 2007.

Tononi, Giulio, Boly, Melanie, Massimini, Marcello and Koch, Christof, 'Integrated Information Theory: From Consciousness to Its Physical Substrate', *Nature Reviews Neuroscience*, vol. 17, no. 7, 2016, pp. 450–461.

Turing, Alan M., 'On Computable Numbers, with an Application to the Entscheidungsproblem', *Proceedings of the London Mathematical Society*, series 2, vol. 42, no. 1, 1937, pp. 230–265.

Vazza, Franco and Feletti, Alberto, 'The Quantitative Comparison Between the Neuronal Network and the Cosmic Web', *Frontiers in Physics*, vol. 8, 2020, Article 525731.

Whitehead, Alfred North, *The Concept of Nature*, Cambridge, Cambridge University Press, 1920.

Whitehead, Alfred North, *Science and the Modern World*, New York, Free Press, 1967 (original work published 1925).

Whitehead, Alfred North, *Adventures of Ideas*, New York, Free Press, 1967 (original work published 1933).

Whitehead, Alfred North, *Process and Reality: An Essay in Cosmology*, corrected edn, ed. David Ray Griffin and Donald W. Sherburne, New York, Free Press, 1978 (original work published 1929).

Zimmermann, Rainer E., Metaphysics of Emergence. Part 1: On the Foundations of Systems, Berlin, Xenomoi, 2015.