Biological Competency

What Must Any Explanation Posit?

Project: Return to Consciousness
Author: Bruno Tonetto
Authorship Note: Co-authored with AI as a disciplined thinking instrument—not a replacement for judgment. Prioritizes epistemic integrity and truth-seeking as a moral responsibility.
Finalized: February 2026

Abstract

This essay examines what any adequate explanation of biological development and regeneration must minimally posit. The target is not physicalism per se, but **pure bottom-up sufficiency**—the claim that local molecular interactions, without control-level primitives, fully explain how organisms reliably achieve and maintain complex anatomical form. The phenomena examined—embryogenesis, regeneration, bioelectric patterning, and developmental plasticity—are mainstream science with strong empirical grounding. The argument is structural, not metaphysical: certain biological systems exhibit **competency**—the reliable achievement of global outcomes under perturbation—and competency requires control architecture (goal states, error signals, corrective dynamics). These primitives cannot be eliminated into purely local microcausation without being reintroduced under different names. The essay concludes by noting how different ontologies interpret this constraint, but the central contribution is the constraint itself, not the ontological comparison.

Keywords: biological competency · embryogenesis · regeneration · bioelectric patterning · bottom-up causation · control architecture · teleology · philosophy of biology

I. Methodological Note: A Constraint Analysis

This essay differs from typical philosophy-of-biology discussions. It does not ask which ontology "feels more natural" given biological phenomena. It asks:

> **What must any adequate explanation of biological development minimally posit?**

This is a **constraint analysis**. The goal is to identify structural features that any successful explanation must accommodate, regardless of metaphysical commitments.

What This Essay Does

- Defines the explanandum precisely: not mechanism, but **organization toward ends**
- Introduces the distinction between mechanism and competency
- Examines mainstream biological phenomena that exhibit competency
- Identifies what competency minimally requires (control-level primitives)
- Shows that these primitives resist elimination into pure bottom-up causation
- Notes how different ontologies interpret this constraint

What This Essay Does Not Claim

- **Proof of any ontology:** The argument identifies constraints and their differential implications, not proof of any ontology
- **Rejection of mechanism:** Mechanisms are real and important; the question is whether they are sufficient
- **Vitalism or new physics:** No non-physical forces are introduced
- **That biology is mysterious:** The phenomena are well-documented; the question is explanatory structure

The Target: Pure Bottom-Up Sufficiency

The claim this essay examines is:

> Local molecular interactions, governed by physics and chemistry, are sufficient to explain how organisms reliably develop, regenerate, and maintain complex anatomical form. No additional explanatory primitives (goal states, targets, error correction) are needed except as useful shorthand.

If this claim is correct, then "goal-directed" language in biology is merely pragmatic—a convenient way to describe what are fundamentally targetless dynamics.

If this claim is incorrect, then adequate biological explanation requires control-level primitives that cannot be eliminated into microcausation. This would not refute physicalism outright, but it would force physicalism to abandon microphysical sufficiency—the core explanatory aspiration that distinguishes it from frameworks that treat organization as primitive.

II. Competency vs. Mechanism

Three Levels of Description

To avoid confusion, this essay distinguishes three levels of biological description:

Level	What it describes	Status in this essay
Mechanistic implementation	Molecules, channels, pathways—how physical interactions occur	Real and important; not disputed
Control architecture	Feedback loops, attractors, setpoints—how systems achieve outcomes	The focus of this essay's constraint argument

Level	What it describes	Status in this essay
Semantic gloss	"Goals," "meanings," "interpretations"—cognitive vocabulary	Used cautiously; not essential to the argument

The essay's central claim concerns **Level 2**: that control architecture is a real, irreducible feature of biological organization, not merely convenient shorthand for Level 1. The essay does *not* require Level 3 claims about biological systems having literal goals or meanings in a mental sense.

The Distinction

Mechanisms explain *how* interactions occur: - Gene regulatory networks describe which transcription factors bind which promoters - Signaling pathways describe how molecules propagate information - Ion channels describe how voltage gradients are maintained

Competency explains *what* a system can reliably achieve: - An embryo reliably produces a complex organism despite perturbations - A planarian fragment reliably regenerates a complete worm - A salamander limb reliably regrows to the correct size and shape

The distinction is not mysterious. A thermostat has mechanisms (bimetallic strip, electrical contacts) and competency (maintaining room temperature). A car has mechanisms (engine, transmission) and competency (reaching destinations). In both cases, understanding the mechanisms does not automatically explain the competency—you must also understand the control architecture. But these are engineered systems whose targets are externally specified by a designer. Biological competency is harder to explain, not easier: the target—"become a complete, correctly proportioned worm"—is not set by any external agent. It is somehow intrinsic to the system itself, and the explanatory demand is to account for how.

What Competency Requires

Any system that reliably achieves a specific outcome under perturbation must implement, explicitly or implicitly:

1. **A target state** — what the system is organized to achieve
2. **An error signal** — information about deviation from the target
3. **Corrective dynamics** — processes that reduce the error

This is not a metaphysical assumption but a control-theoretic necessity. A system without these features cannot exhibit reliable goal-achievement under perturbation; it can only exhibit whatever outcome its initial conditions and dynamics happen to produce.

This is more specific than the general observation that higher-level sciences use ineliminable vocabulary. Economics needs "price," meteorology needs "pressure systems"—but these describe aggregate dynamics, not active correction toward a target. A market does not detect deviation from a target state and self-correct; a weather system does not restore a disrupted pattern. Biological competency is distinctive because it involves deviation detection and corrective return to a specific outcome.

The Eliminability Question

The central question is whether these control-level primitives can be **eliminated** into purely local microcausation:

The eliminativist position: "Target," "error," and "correction" are convenient descriptions of what are fundamentally just molecular interactions. There is no target represented anywhere; there is only chemistry. What looks like "error correction" is just the system doing what chemistry does.

The non-eliminativist position: Control-level primitives cannot be eliminated without being reintroduced under different names. Any account of how development "converges" to a specific outcome, how regeneration "restores" proper form, or how the system "knows when to stop" implicitly invokes the very primitives it claims to eliminate.

This essay argues for the non-eliminativist position—not on metaphysical grounds, but by examining what biological explanation actually requires.

III. The Explanandum: What Needs Explaining

Before examining specific phenomena, we must be precise about what needs explaining. The question is not:

- How do molecules interact? (Biochemistry answers this)
- What genes are involved? (Genomics answers this)
- What signaling pathways operate? (Developmental biology answers this)

The question is:

> **How do systems reliably achieve and maintain specific global outcomes—complex anatomical form, proper proportions, functional organization—under noise, damage, and perturbation?**

This is a question about **organizational achievement**, not molecular interaction.

Key Features of the Explanandum

Reliability: Development produces the same basic anatomy across individuals despite genetic variation, environmental differences, and stochastic noise. This is not merely "the system does what it does"—it is convergence to a specific outcome across variable conditions.

Robustness: Perturbations (damage, transplantation, experimental manipulation) are often corrected. The system does not merely proceed from its perturbed state; it returns toward the original trajectory. This requires information about what that trajectory is.

Stopping conditions: Development and regeneration stop at appropriate points. A regenerating limb does not grow indefinitely; it stops when it reaches the correct size. A tumor is precisely the failure of this stopping—growth without appropriate termination. What specifies "correct size"?

Global coordination: Local cells respond to local signals, yet the outcome is a globally coordinated form. No single cell "knows" the whole plan, yet the whole plan is realized. How does information about the whole inform local decisions?

These features define the explanandum. Any adequate explanation must account for them.

IV. Empirical Anchors: Mainstream Biology

The phenomena examined here are not contested or fringe. They are textbook biology with strong experimental support. The question is not whether they occur, but what their occurrence requires explanatorily.

A. Embryogenesis: Reliable Convergence to Form

The phenomenon: A single fertilized cell produces a complex organism with specific anatomy. Despite starting from a single cell, despite noise in gene expression, despite environmental variation, development reliably converges to the species-typical form.

Key observations:

- **Convergence under perturbation:** Early embryos can be split, and each half develops into a complete organism (regulation). Cells removed from early embryos are replaced, and development proceeds normally.
- **Positional plasticity:** Cells transplanted to different locations often adopt fates appropriate to their new position, not their origin.
- **Error correction:** Developmental errors introduced experimentally are often corrected later, suggesting the system tracks deviation from a target.

The mechanistic account: Gene regulatory networks, morphogen gradients, cell-cell signaling, and mechanical forces coordinate development. These mechanisms are real and well-characterized.

What remains unexplained by mechanism alone: Mechanisms explain the *interactions*. They do not obviously explain the *target*—why development converges to this specific outcome rather than any of the countless configurations that molecular interactions could produce. The phrase "development converges to the species-typical form" already presupposes what needs explaining.

B. Regeneration: Restoration Toward a Target

The phenomenon: Some organisms can regenerate complex structures after damage. This is not merely wound healing; it is the reliable production of anatomically correct replacement structures.

Key observations:

- **Planarian regeneration:** Flatworms can regenerate complete organisms from fragments as small as 1/279th of the body. Each fragment produces a correctly proportioned complete worm, not a fragment-sized worm.
- **Salamander limb regeneration:** Amputated limbs regrow with correct structure, stopping at the appropriate size.
- **Deer antler regrowth:** Annual regeneration of complex branched structures following species-typical patterns.

The mechanistic account: Stem cells, dedifferentiation, positional information gradients, and gene expression programs are involved. These mechanisms are real.

What remains unexplained by mechanism alone: The target. A planarian fragment "knows" to become a complete worm, not a fragment-sized piece. The regenerating system behaves as if it has access to information about what the complete form should be. Where is this information? What constitutes it? Saying "the genes encode it" pushes the question back: how do genes encode a three-dimensional anatomical target, and how does the system read this encoding to guide regeneration?

C. Bioelectric Patterning: Large-Scale Information Integration

The phenomenon: Research by Michael Levin and colleagues has demonstrated that bioelectric signaling—voltage gradients across cells and tissues—plays a central role in specifying anatomical form.

Key observations:

- **Voltage patterns encode anatomy:** The spatial pattern of membrane potentials across tissue specifies what structures will form.
- **Reprogrammability:** Altering bioelectric patterns can induce ectopic structures (eyes, limbs, heads) without genetic modification. The same genome produces different anatomies depending on bioelectric state.
- **Memory and stability:** Bioelectric patterns exhibit persistence. They can store information about target morphology and guide correction when disrupted.
- **Independence from genetics:** These effects occur without changing DNA sequence.

Levin's framework: Levin describes these systems using cognitive vocabulary—"anatomical decision-making," "morphogenetic intelligence," "goal states"—while emphasizing that this language describes functional organization, not necessarily consciousness.

The mechanistic account: Ion channels, gap junctions, and voltage-sensitive signaling molecules implement bioelectric signaling. These are well-characterized molecular components.

What remains unexplained by mechanism alone: The *informational* character of bioelectric patterning. Voltage gradients do not merely cause local effects; they encode large-scale anatomical information. The gap between "ion channels change voltage" and "voltage patterns specify what anatomy to build" is substantial. The latter description already invokes goal-states (what to build) and representation (patterns that specify outcomes).

D. Developmental Plasticity: Context-Sensitive Reinterpretation

The phenomenon: Biological development exhibits remarkable context sensitivity. The same genetic and molecular signals produce different outcomes depending on context.

Key observations:

- **Polyphenism:** Single genotypes produce radically different morphologies depending on environmental cues (castes in social insects, seasonal forms in butterflies).
- **Cellular reinterpretation:** Cells respond to the same molecular signal differently depending on their developmental history and position.
- **Late-stage correction:** Errors introduced early in development can be corrected later, as if the system maintains information about what should be built.

- **Adaptive plasticity:** Plastic responses typically produce biologically functional outcomes, not random variations.

The mechanistic account: Epigenetic modifications, conditional gene expression, and environmental signaling pathways implement plasticity. These are real mechanisms.

What remains unexplained by mechanism alone: The *semantic* character of plasticity. Signals are molecular; responses are adaptive. The system appears to respond not just to signals but to their *meaning*—their significance in context. "The cell interprets the signal based on its history" is standard biological language, but "interpretation" is a cognitive term. Either this language is misleading shorthand, or it points to something that mechanism alone does not capture.

V. The Explanatory Fork

The empirical phenomena establish that biological systems exhibit competency: reliable achievement of global outcomes under perturbation. The question is what this requires explanatorily.

Option A: Accept Control-Level Primitives as Real Explanatory Features

On this view, goal states, error signals, and corrective dynamics are **real features** of biological organization, not merely convenient descriptions.

This does not mean these features are non-physical or mysterious. A thermostat's goal state (target temperature) is physically implemented in the setting dial; its error signal is physically implemented in the temperature sensor; its corrective dynamics are physically implemented in the heating element. But explaining the thermostat requires identifying these control features, not just describing the molecular composition of its components.

Similarly, explaining biological competency requires identifying how goal states, error signals, and corrective dynamics are implemented—not eliminating them into component interactions.

Implications:

- Biological explanation must include control-theoretic primitives
- These primitives are not eliminable into purely local microcausation
- "What is the system organized to achieve?" is a legitimate scientific question
- Any ontology must accept irreducible control-level organization—for physicalism, this abandons microphysical sufficiency

Option B: Attempt Elimination

On this view, control-level language is useful shorthand but does not name real features. There are no goal states, only molecular interactions. What looks like "error correction" is just chemistry doing what chemistry does.

The problem with elimination:

Elimination typically reintroduces the eliminated concepts implicitly. Consider how eliminativist accounts describe the phenomena:

- "The system *converges* to the correct form" — Convergence to a specific outcome implies a target.

- "The network *computes* the appropriate response" — Computation implies goals and correctness conditions.
- "Development is *robust* to perturbation" — Robustness implies a standard from which deviation is measured.
- "The system *knows when to stop*" — This explicitly invokes knowledge and goal-states.

These phrases are not sloppy language; they are ineliminable. Describing what happens without them yields descriptions that fail to capture the phenomenon. "The molecules interact according to physical laws" is true but does not explain why the outcome is a complete worm rather than any other configuration compatible with those laws.

The Upshot

Either control-level primitives are real explanatory features, or they are eliminable shorthand. If the former, pure bottom-up sufficiency is false. If the latter, elimination must actually succeed—descriptions must be possible that capture the phenomena without implicitly invoking targets, goals, and error correction.

This essay argues that elimination does not succeed. The phenomena require explanatory resources beyond local microcausation.

The Strongest Objection: Attractors as Targets

The most sophisticated physicalist response does not attempt elimination. Instead, it reframes control-level primitives in dynamical systems terms:

> "Goal-directedness is a property of dynamical systems with attractors. 'Targets' are shorthand for basins of attraction shaped by evolutionary constraint. No extra posits are required beyond physics and selection."

This objection correctly warns against sliding from control architecture (Level 2) to semantic gloss (Level 3) without argument—a distinction introduced in Section II. Biology often *speaks* semantically while *working* control-theoretically.

However, the objection concedes the central point. An attractor is not a molecule or local interaction. It is a **global property of the system's state space**—defined over possible states, temporally extended, counterfactually robust. When the critic says "targets are shorthand for basins of attraction," they have conceded that adequate explanation requires global state-space structure, not only local interaction descriptions. Attractors are control-level primitives, even if physical. This concession is not unique to the present essay: philosophers of biology working within physicalist frameworks—including organizational accounts of biological autonomy (Mossio & Moreno, 2010) and mechanistic explanation (Bechtel, 2007; Craver, 2007)—independently conclude that organizational constraints are irreducible to their molecular constituents. Their disagreement with the present analysis begins not at the structural claim, but at its ontological implications.

Evolutionary constraint does not remove the explanatory demand. The move "evolution shaped the basins" explains *why* such systems exist historically—not *how* they function now. Evolution does not specify current target states, real-time error correction, or stopping conditions. One can accept evolutionary shaping and still ask: Why does the resulting system behave as a goal-directed controller rather than a locally equilibrating process?

It is sometimes objected that evolutionary theory relies on implausibly slow random search. Modern evolutionary biology rejects this strawman: variation is heavily constrained by developmental organization, physical law, and inherited regulatory structure. But this concession reinforces the essay's point. Evolutionary explanation already depends on the very control architectures examined here—developmental systems that channel variation toward viable forms. Evolution can explain *why* such architectures were selected, but not *why matter supports* stable, self-maintaining, goal-directed control regimes in the first place. That is a structural precondition of evolutionary dynamics, not a product of them.

The clarification: This essay does not claim biological "goals" are semantic or mental. It claims that goal-like *structure* exists, is causally operative, and is not eliminable into local dynamics. Once this is conceded—as the attractor objection concedes it—the debate shifts from "Are goals real?" to "What kind of thing is a control architecture?" Whether one calls this structure "physical" or not, the grounding question remains: why does reality support stable, self-maintaining, goal-directed organization at all? Dismissing ineliminable structure as "merely useful description" already presupposes that microphysics is the default level of reality—the very assumption under examination. That is where ontological interpretation enters legitimately.

VI. What This Establishes

The Structural Claim

Biological competency—reliable achievement of global outcomes under perturbation—requires control-level primitives: goal states, error signals, corrective dynamics. These primitives:

1. Are implicitly invoked by any account that captures what happens
2. Cannot be eliminated into purely local microcausation without being reintroduced under different names
3. Represent genuine explanatory requirements, not merely pragmatic conveniences

This is a **structural claim** about what explanation requires, not a metaphysical claim about what exists.

What This Does Not Establish

- **Which ontology is correct:** Multiple ontologies can accept the constraint, though not at equal cost. Section VII notes where acceptance is structurally comfortable and where it requires treating irreducible features as brute.
- **That biological systems are conscious:** This analysis argues from control architecture, not from consciousness. The ontological question of what grounds control architecture is addressed in Section VII.
- **That mechanisms are insufficient:** Mechanisms are necessary. The claim is that they are not sufficient without control-theoretic organization.
- **New physics or vitalism:** No non-physical forces are proposed.

VII. Ontological Implications

The constraint established above—that biological explanation requires irreducible control-level primitives—has different implications for different ontologies. The preceding structural claim operates at Level 2—control architecture. Asking what *grounds* that architecture necessarily enters ontological territory. This is not a conflation of levels but a separate inquiry, explicitly marked as such.

Earlier sections referred to "physicalism" as the target. But naive reductive physicalism—which denies any explanatory role for system-level organization—is not a serious contender. The sophisticated physicalist position is **emergentism**, and that is the framework worth engaging.

The Emergentist Response

Emergentism holds that complex organization develops from simpler components through non-linear dynamics, feedback, and systemic interaction. Unlike reductive physicalism, it readily accepts control-level primitives—attractors, feedback loops, goal-directed organization. But accepting the constraint means accepting that these structures are **causally operative**, not merely convenient summaries of micro-events, and that "what is the system organized to achieve?" is a **legitimate scientific question**, not eliminable into "what do the molecules do?"

It is sometimes claimed that far-from-equilibrium thermodynamics already explains why matter supports organized structures: dissipative systems arise naturally from energy flow. But thermodynamic self-organization produces pattern, not competency. A hurricane forms spontaneously; it does not detect deviation from a target state and self-correct. A Bénard cell arises from thermal gradients; it does not regenerate when disrupted. The gap between spontaneous pattern formation and biological error-correcting competency is precisely the explanatory gap this essay identifies.

These are irreducible organizational primitives—precisely what emergentism's reductive heritage was designed to eliminate. If anything that resists reduction can be said to "emerge," emergence risks naming the residue of failed reduction rather than doing independent explanatory work. Emergentism explains *how* organization develops given laws that permit attractors, feedback, and self-organization, but not *why* reality has such organization-friendly laws in the first place. That is where it places its brute fact—and in doing so, it moves toward territory idealism already occupies: irreducible organization as a feature of reality's basic nature rather than a product of bottom-up causation.

The Position of Idealism

Under analytic idealism, the constraint is less surprising because goal-directedness is characteristic of mental processes. If biological organization is the extrinsic appearance of underlying mental processes:

- Goal states are intrinsic to mental organization, not puzzling emergents from targetless dynamics
- Error correction reflects the natural coherence of mental processes
- The "target" question dissolves: the system is organized toward specific outcomes because that is what mental organization *is*

Idealism does not explain mechanisms—it does not predict which molecules are involved. But it offers a framework where the *kind* of organization biology exhibits is intelligible rather than anomalous.

The Comparison

The two frameworks accept the same science but differ on grounding:

Dimension	Emergentism	Idealism
Accepts control architecture	Yes	Yes
Accepts attractors, feedback, organization	Yes	Yes
Brute fact	Organization-enabling laws exist	Reality is mental in nature
Organization is…	Real but requiring explanation of its ground	Intrinsic to mind, not anomalous

Both frameworks accept irreducible control architecture. But they differ on what ultimately grounds it—and as the preceding sections show, that difference is not symmetrical.

Where Explanation Stops

Science explains *how* organization unfolds—through non-linear dynamics, feedback loops, attractors, control architectures, and evolutionary shaping. This mechanistic account is **ontologically neutral**. Both emergentism and idealism accept it fully.

The disagreement is not about mechanism but about **grounding**:

- Emergentism treats the existence of organization-enabling laws and structures as **brute facts**—explanation stops there.
- Idealism treats organization as intrinsic to mind-like processes, so that organization-friendliness is not an unexplained feature of reality but an expected one.

This is a metaphysical disagreement about where explanation is allowed to terminate, not a scientific one about how systems work. Both frameworks accept the control-theoretic account developed in this essay. They differ on what that account ultimately rests upon.

Neither ontology *explains* the mechanisms—science does. The ontologies differ on what grounds the reality that science describes—and as the preceding analysis shows, the constraint sits more naturally within one framework than the other.

For a fuller treatment of this theme, see Where Explanation Stops (wes), which develops the "brute fact" analysis across multiple ontological frameworks.

VIII. Conclusions and Limits

What This Essay Establishes

1. Biological systems exhibit **competency**: reliable achievement of global outcomes (complex anatomy, proper proportions, regenerated form) under perturbation.

2. Competency requires **control architecture**: goal states, error signals, and corrective dynamics. These are not merely useful descriptions but structural features that any adequate explanation must posit.
3. These control-level primitives **resist elimination** into purely local microcausation. Attempts to eliminate them reintroduce them implicitly. Pure bottom-up sufficiency is therefore inadequate: explanation requires control-level features that operate through, but are not reducible to, molecular interactions.
4. Emergentism accepts the constraint, but at the cost of treating irreducible organization as brute rather than derived. If anything that resists reduction can be said to "emerge," **emergence names the residue of failed reduction** rather than providing an independent explanatory framework.

What This Essay Does Not Establish

- That emergentism is false
- That consciousness is involved in biological control
- That mechanisms are unimportant or wrong
- Predictions about specific biological systems

The Limits of This Analysis

This analysis identifies what explanation requires; it does not provide that explanation. Knowing that control-level primitives are needed does not tell us how they are implemented in any specific system. That remains the work of empirical biology.

The analysis is also silent on mechanism. It does not predict which genes, proteins, or pathways are involved in any particular developmental process. It operates at the level of explanatory structure, not empirical detail.

Finally, the ontological interpretations offered are brief and suggestive, not conclusive. Full defense of emergentism or idealism requires arguments beyond what biology alone provides.

The Bottom Line

Any adequate explanation of biological development and regeneration must posit control-level primitives—goal states, error signals, corrective dynamics—that cannot be eliminated into purely local microcausation. This is a structural constraint on explanation. Multiple ontologies can accept it, though not at equal cost—Section VII shows where acceptance is structurally comfortable and where it requires treating irreducible features as brute.

For readers committed to pure bottom-up sufficiency, this essay poses a challenge: provide an account of competency phenomena that genuinely eliminates control-level language rather than implicitly reintroducing it. If this cannot be done, the structural constraint stands.

References

Bechtel, W. (2007). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. Routledge.

Bongard, J., & Levin, M. (2021). Living things are not (20th century) machines: Updating mechanism metaphors in light of the modern science of machine behavior. *Frontiers in Ecology and Evolution*, 9, 650726.

Brockes, J. P., & Kumar, A. (2008). Comparative aspects of animal regeneration. *Annual Review of Cell and Developmental Biology*, 24, 525-549.

Craver, C. F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press.

Fields, C., & Levin, M. (2022). Competency in navigating arbitrary spaces as an invariant for analyzing cognition in diverse embodiments. *Entropy*, 24(6), 819.

Gilbert, S. F. (2014). *Developmental Biology* (10th ed.). Sinauer Associates.

Gilbert, S. F., & Epel, D. (2015). *Ecological Developmental Biology* (2nd ed.). Sinauer Associates.

Kastrup, B. (2019). *The Idea of the World: A Multi-Disciplinary Argument for the Mental Nature of Reality*. Iff Books.

Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.

Levin, M. (2014). Molecular bioelectricity: How endogenous voltage potentials control cell behavior and instruct pattern regulation in vivo. *Molecular Biology of the Cell*, 25(24), 3835-3850.

Levin, M. (2019). The computational boundary of a 'self': Developmental bioelectricity drives multicellularity and scale-free cognition. *Frontiers in Psychology*, 10, 2688.

Levin, M. (2021). Bioelectric signaling: Reprogrammable circuits underlying embryogenesis, regeneration, and cancer. *Cell*, 184(8), 1971-1989.

Levin, M. (2022). Technological approach to mind everywhere: An experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience*, 16, 768201.

Mathews, J., & Levin, M. (2017). Gap junctional signaling in pattern regulation: Physiological network connectivity instructs growth and form. *Developmental Neurobiology*, 77(5), 643-673.

Mossio, M., & Moreno, A. (2010). Organisational closure in biological organisms. *History and Philosophy of the Life Sciences*, 32(2-3), 269-288.

Nicholson, D. J., & Dupré, J. (Eds.). (2018). *Everything Flows: Towards a Processual Philosophy of Biology*. Oxford University Press.

Reddien, P. W., & Sánchez Alvarado, A. (2004). Fundamentals of planarian regeneration. *Annual Review of Cell and Developmental Biology*, 20, 725-757.

Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.

Walsh, D. M. (2015). *Organisms, Agency, and Evolution*. Cambridge University Press.

West-Eberhard, M. J. (2003). *Developmental Plasticity and Evolution*. Oxford University Press.

Wolpert, L., Tickle, C., & Martinez Arias, A. (2015). *Principles of Development* (5th ed.). Oxford University Press.

Related Essays in This Project

Available at: https://returntoconsciousness.org/

Return to Consciousness (rtc) — The core framework this essay extends

Where Explanation Stops (wes) — Fuller treatment of the "brute fact" analysis introduced in Section VII

Anomalous Phenomena and Consciousness (apc) — Parallel constraint analysis using contested phenomena

Consciousness Structure (cst) — Clinical applications of competency frameworks

License