The Shape a Worm Remembers

Bioelectric target morphology as a defended setpoint, and a test that could break the claim

Jaxon Costa, Researcher in Brain and Frontiers

International Academy for Consciousness Studies

Researched and prepared with the Academy. Working draft 1.0

Abstract

Where is the shape of a body stored? The default answer in developmental biology is the genome: sequence specifies proteins, proteins build tissue, tissue settles into form. Regeneration in the planarian flatworm complicates this answer. A worm cut in half rebuilds exactly one head and one tail, in the right places, at the right sizes, and then stops. That "stops" is the interesting part. It implies a target, a shape the tissue is regenerating toward and defending against deviation. A now well-replicated body of work from the Levin laboratory shows that transient perturbation of cell membrane voltage (Vmem) patterns and gap junctional connectivity can rewrite this target: worms regenerate two heads, or heads shaped like those of related species, without any change to genomic sequence, and in the two-headed case the altered outcome persists through subsequent rounds of cutting in plain water. This paper treats that stored, defended target as a candidate biological closure: a setpoint the tissue holds, with a measurable remainder between current form and target form. I argue that closure is not idle relabeling only if it makes a prediction that local biochemical determinism does not, and I state one, together with the condition that would kill it. I keep the strong "cognition" and "memory" readings tagged as disputed, and I firewall a defended setpoint (a structural and functional claim) from any claim about experience.

1. The problem: a target that is defended, not just followed

FACT. Cut a planarian into pieces and each fragment regenerates a complete, correctly proportioned worm: one head anterior, one tail posterior, organs scaled to the new body size. The genome in every fragment is identical to the genome in every other fragment and identical to the genome of the intact animal. Sequence alone therefore cannot be what tells a middle fragment to build a head at its anterior wound and a tail at its posterior wound, because the same sequence is present in a piece that must do the opposite at each end. Something spatial and physiological, laid over the genome, carries the polarity and the stopping rule.

The phrase worth pausing on is "stopping rule." A gradient that instructs cells to divide and differentiate is a familiar object. A system that divides and differentiates *until a particular whole-body configuration is reached and then halts, and that resists being pushed away from that configuration*, is a different kind of

object. It behaves as though it holds a representation of the endpoint and works to reduce the difference between the current state and that endpoint. In control-theoretic language this is a setpoint with error correction. The empirical question this paper is built around is whether planarian regeneration is genuinely setpoint-like, defending a stored whole-shape target, or whether the appearance of a defended target dissolves, on inspection, into a chain of purely local causes with no stored endpoint anywhere in the system.

2. What is established: Vmem and connectivity edit the target

FACT. Several independent lines of evidence show that bioelectric state, not only transcription, instructs large-scale planarian morphology.

Beane, Morokuma, Adams, and Levin (2011) used a chemical-genetics screen to show that membrane voltage set by the H+/K+-ATPase is required for head regeneration, and that manipulating it changes head and organ size. Depolarization and hyperpolarization of the regeneration blastema shifted anatomical outcomes in a voltage-dependent way, establishing Vmem as an instructive variable upstream of gross morphology, not merely a correlate of it.

Emmons-Bell and colleagues (2015) blocked gap junctional communication in *Girardia dorotocephala* with octanol during regeneration. Genetically wild-type worms regenerated head shapes and brain morphologies resembling those of *other* planarian species (including forms resembling *Schmidtea*, *Dugesia*, and *Polycelis*), stochastically, without any edit to the genome. The distribution of physiologically accessible head shapes recapitulated a set of related species-typical anatomies. The effect was reversible: worms remodeled back toward their native form.

Durant, Morokuma, Fields, Williams, Adams, and Levin (2017) provide the sharpest case for stored-target language. A brief perturbation of endogenous bioelectric signaling produced a stable fraction of two-headed regenerates. The crucial result is what happened next: when two-headed worms were amputated *again*, in plain water with no further treatment, they regenerated as two-headed. The one-headed-looking worms from the same treated cohort also carried the altered pattern latently and could yield two-headed progeny on recutting. The altered anatomy bred true through cutting. The authors could also reset the two-headed animals back to one head by targeting the H+/K+-ATPase circuit, restoring the native outcome.

The pattern across these studies is consistent and, for the purposes of this paper, the load-bearing empirical claim: the target a planarian regenerates toward can be edited by changing a physiological (bioelectric) variable, the genome held constant, and at least in the two-headed case the edited target is heritable across cuts and separately reversible. Reviews by Mathews and Levin (2018) and Levin and Martyniuk (2018) collect the wider vertebrate and invertebrate evidence for Vmem as an instructive layer in patterning.

3. The closure reading: a stored setpoint with a remainder

C'S READING (framework lens, tagged; the biology carries the argument). The Academy's closure framework describes openness settling into a "this" that then runs on rules to stay itself. Its terms: C for presence, c for content, M for openness, m for horizon, Cl for the closure itself, R for the remainder that

closure does not absorb. A closure is a bounded self-maintaining "this" that defends what it is against perturbation, and that has a remainder: the difference between what it holds itself to and what it currently is, the gap that its own activity works on.

A defended morphological target is a clean candidate for a biological closure, a low rung on a ladder that has nothing yet to do with minds: a living tissue holding a form. Read this way, the planarian's regenerative target is Cl, the stored "this" the animal keeps returning to. The current anatomy after a cut is the content c. The remainder R is exactly the mismatch the blastema is working to close, the difference between "one head, one tail, organs to scale" and "raw wound." The openness M is the space of anatomies the tissue *could* be pushed into, and the Emmons-Bell result is striking precisely because it shows that space is structured: perturbation does not yield noise, it yields a discrete menu of coherent, species-like heads. The horizon m is the boundary of what this physiological network can measure and correct, which is why a worm defends worm-scale shape and not, say, the shape of its enclosure.

HYPOTHESIS. The value of the closure reading is not the vocabulary. It is that "defended setpoint with a remainder" is a claim with teeth: it says the system tracks a whole-shape target and acts to reduce a global error, and that this global behavior is not just the summed shadow of local rules. If that is true, editing the setpoint should produce reorganizations that are *coherent at the scale of the whole shape* rather than a patchwork of independent local changes. Section 5 turns this into a test.

4. What would make closure trivial, and what would make it real

The honest worry is that "stored setpoint" is redescription. Any physical system relaxes toward some configuration; a ball in a bowl "defends" the bottom. Calling the bottom of the bowl a stored target teaches us nothing. So the closure claim must be more than "the system returns to a stable state."

The distinguishing content is this. In a purely local, bottom-up account, morphology is the fixed point of local rules: each cell reads its immediate chemical neighborhood, expresses accordingly, and global form is whatever those local interactions happen to sum to. There is no representation of the whole anywhere; the "target" is an epiphenomenon of local dynamics, and editing outcomes means editing local rules one neighborhood at a time. In a stored-setpoint (closure) account, there is a distributed physiological variable, a Vmem pattern, that encodes the *whole-shape* endpoint semi-independently of the local transcriptional machinery, and that the tissue reads and defends. On this account you can change the whole outcome by editing the pattern without editing the local genetic rules, and the change will be globally coherent because you edited a global variable.

Pezzulo and Levin (2016) make the general version of this argument: some living systems are more tractably explained and controlled at a level above the molecular, with variables and goal states defined over the whole, the way control theory and least-action principles work in other sciences. That is a methodological claim about the right level of description. The closure reading adds a commitment that is checkable: the higher-level variable is not merely a convenient summary but is *doing work the local level does not do*, and it leaves a remainder (global error) that the system measurably reduces.

The two accounts are not always cleanly separable, and much bioelectric signaling surely does reduce to reaction-diffusion and gene-regulatory dynamics. The claim is not that all of morphogenesis is closure. It

is narrower: that at least the target-storage step in planarian regeneration behaves as a stored global setpoint, and that this is decidable by experiment.

5. A test that risks the claim

WAGER. I will bet that the setpoint is real: that Vmem patterns store a whole-shape target that cannot be fully reduced to local gene-expression rules. Here is the prediction that follows and the observation that would sink it.

> **Prediction.** Targeted edits to the endogenous Vmem pattern (not to sequence, and not delivered as a spatial mosaic) will produce whole-shape reorganizations that are globally coherent: a rewritten target such as two-headedness that is (a) proportioned and axially organized as a complete alternative body plan, not a local outgrowth, (b) heritable through subsequent amputations performed in plain water with no further treatment, as in Durant et al. 2017, and (c) independently resettable to the native plan by restoring the pattern. Across a graded series of pattern edits, outcomes will fall into a discrete, low-dimensional set of coherent anatomies (as in Emmons-Bell et al. 2015), and single-cell transcriptomics of the edited blastema, before overt morphological divergence, will show that the *same* local gene-expression programs are deployed in a globally reorganized spatial arrangement, rather than a novel local program at each altered site.
>
> **Kill.** If every bioelectric manipulation that changes morphology is found, on adequate molecular resolution, to act through a local biochemical cause sufficient to specify the changed structure at that site (a locally altered morphogen source or gene-regulatory state that fully accounts for the outcome), with no whole-shape target that persists independently of those local causes, no heritability of the edited plan across untreated cuts, and no discrete coherent menu of outcomes (perturbation instead yielding graded or fragmentary local defects), then there is no stored setpoint and the closure claim fails. Bioelectricity would then be one more local signaling input, correctly described bottom-up, and "defended target" would be redescription with no surplus content.

The prediction is not free. Point (a) forbids the two-headed form from being merely a duplicated local head-inducing focus with the rest of the body indifferent; the second axis must be integrated into a coherent whole. Point (b) is already supported by Durant et al. for two-headedness and is the strongest existing evidence for storage, but the prediction extends it: *other* pattern edits (not only two-headedness) should also breed true if they are genuine setpoint changes. The transcriptomic clause is the sharpest risk. If edited and native blastemas run *different local programs* rather than the same programs rearranged, that is evidence the change lives at the local level after all, and the surplus global variable is doing less than claimed.

6. The cognition dispute, and where the firewall goes

DISPUTED. A prominent framing of these results describes cells and tissues as engaging in "memory," "decision-making," "goals," and "basal cognition," with morphogenesis cast as a collective intelligence navigating an anatomical space, and individuality defined by a "cognitive light cone," the boundary of what a system can measure, model, and affect (Levin 2019). This framing is generative and has motivated real experiments. It is also contested. Critics argue that "cognition," "memory," and "goal" are being used

metaphorically and then cashed as if literal, that homeostatic error correction does not by itself warrant cognitive vocabulary, and that the same data are fully describable in the language of dynamical systems and control without positing anything mind-like. This paper does not adjudicate that dispute and does not depend on it.

I want to be exact about what the closure reading commits to and what it does not. FACT plus HYPOTHESIS: the tissue holds and defends a whole-shape target and reduces a global error. That is a structural and functional claim, testable as in Section 5. It says nothing about whether there is anything it is like to be a regenerating planarian, and nothing about the tissue "wanting" its shape. A thermostat defends a temperature setpoint and has a remainder (the gap it works to close) without experiencing warmth. "Defended setpoint" earns the word "defended" from behavior (it corrects perturbations toward a stored value) and earns nothing more.

FIREWALL. The closure ladder is explicitly a ladder of *forms of self-maintaining organization*, not a ladder of *degrees of experience*. A biological closure, a living thing holding its form, is a rung on the organizational ladder. Placing morphological setpoints on that ladder is a claim about structure and function, that the system stores and defends a "this" with a remainder, and it is firewalled from any claim about consciousness, sentience, or inner life. Where the disputed cognition language imports experiential connotation, I mark it DISPUTED and set it aside. The biology of a defended target stands or falls on Section 5 regardless of how the cognition debate resolves. Conflating the two would let a genuine finding (a stored setpoint) smuggle in an ungrounded one (an experiencing subject), and it is precisely that conflation the firewall exists to prevent.

7. Limits and wager

Several limits bound the claim honestly.

First, scope. The two-headed heritability result (Durant et al. 2017) is the single strongest pillar, and it is one anatomical outcome in one genus under specific perturbations. The Emmons-Bell (2015) species-menu result is reversible and stochastic. Generalizing "stored setpoint" beyond these cases is a hypothesis, not a fact, and the prediction in Section 5 is exactly an attempt to make that generalization risk something.

Second, the local/global boundary is not binary. It is likely that bioelectric patterns and local gene-regulatory dynamics are coupled in both directions, Vmem shaping transcription and transcription shaping Vmem. The kill condition is written to catch the case where the global variable does no *independent* work; a finding of tight bidirectional coupling with genuine top-down constraint would still be consistent with a (weaker) setpoint reading, and the paper should not overclaim a clean separation that biology may not honor.

Third, mechanism of storage remains partly open. That a Vmem pattern *encodes* a target is supported; the physical substrate that makes the encoding stable across cuts (the interplay of ion channel expression, gap junctions, and feedback that re-establishes the pattern) is still being worked out (Levin and Martyniuk 2018; Mathews and Levin 2018). "The pattern stores the shape" is a claim at the level of information, and its molecular implementation is a live research question, not a settled one.

WAGER, stated plainly. I am betting that planarian target morphology is a real stored setpoint, a genuine biological closure with a defended whole-shape value and a measurable remainder, and that this will survive single-cell resolution: edited blastemas will show rearranged, not rewritten, local programs. The bet is falsifiable and I have said exactly what loses it. If it loses, the correct description is bottom-up all the way down and bioelectricity is a local input among others; that would be a clean, informative result and not a disaster. The framework earns its keep only if the setpoint is real, and the setpoint's reality is not a matter of framework. It is a matter of what the worm does.

8. Conclusion

A planarian rebuilds one head, one tail, and correctly scaled organs, and then stops, and resists being pushed off that form. Membrane voltage patterns can edit what "that form" is, without touching the genome, and in the two-headed case the edit breeds true through later cuts and can be independently reset. The most economical description of a system that stores a whole-shape endpoint and works to reduce the difference between its current state and that endpoint is a defended setpoint with a remainder, which is what the Academy's framework calls a biological closure. The description is worth adopting only if it predicts something local determinism does not, so I have committed to one: coherent, heritable, resettable, discretely-menued whole-shape reorganizations produced by editing a global variable, with the same local programs rearranged rather than replaced. I have stated the observation that would falsify it. And I have firewalled the structural claim (a stored, defended target) from the disputed and separable claim that the tissue is in any experiential sense a cognizing self. The worm remembers its shape in a sense we can test. Whether it *knows* it is a different question, and not one this paper answers.

References

Beane, W. S., Morokuma, J., Adams, D. S., and Levin, M. (2011). A chemical genetics approach reveals H,K-ATPase-mediated membrane voltage is required for planarian head regeneration. Chemistry and Biology, 18(1), 77-89.

Emmons-Bell, M., Durant, F., Hammelman, J., Bessonov, N., Volpert, V., Morokuma, J., Pinet, K., Adams, D. S., Pietak, A., Lobo, D., and Levin, M. (2015). Gap Junctional Blockade Stochastically Induces Different Species-Specific Head Anatomies in Genetically Wild-Type Girardia dorotocephala Flatworms. International Journal of Molecular Sciences, 16(11), 27865-27896.

Pezzulo, G., and Levin, M. (2016). Top-down models in biology: explanation and control of complex living systems above the molecular level. Journal of the Royal Society Interface, 13(124), 20160555.

Durant, F., Morokuma, J., Fields, C., Williams, K., Adams, D. S., and Levin, M. (2017). Long-Term, Stochastic Editing of Regenerative Anatomy via Targeting Endogenous Bioelectric Gradients. Biophysical Journal, 112(10), 2231-2243.

Levin, M., and Martyniuk, C. J. (2018). The bioelectric code: An ancient computational medium for dynamic control of growth and form. BioSystems, 164, 76-93.

Mathews, J., and Levin, M. (2018). The body electric 2.0: recent advances in developmental bioelectricity for regenerative and synthetic bioengineering. Current Opinion in Biotechnology, 52, 134-144.

Levin, M. (2019). The Computational Boundary of a "Self": Developmental Bioelectricity Drives Multicellularity and Scale-Free Cognition. Frontiers in Psychology, 10, 2688.