Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Connectome-derived structural support identifies a compact scaffold for robust embodied control in BIOSEED

Pantaleone Carlucci

pantaleone.carlucci@cnr.it

National Research Council

Giorgia Ghergo

Scuola Universitaria Superiore IUSS di Pavia

Lucia Paciucci

National Research Council

Article

Keywords: connectome, NeuroAI, embodied control, C. elegans, structural priors, recurrent controllers

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11384131/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Connectome-derived structural support identifies a compact scaffold for robust embodied control in BIOSEED

Pantaleone Carlucci[1], Giorgia Ghergo[2], Lucia Paciucci[1]

1. National Research Council Department of Earth System Sciences and Environmental Technologies, Piazzale Aldo Moro, 7 - 00185 Roma, Italy
2. Scuola Universitaria Superiore IUSS di Pavia - Palazzo del Broletto, Piazza della Vittoria, 15, 27100 Pavia, Italy

Corresponding author: Pantaleone Carlucci, pantaleone.carlucci@cnr.it

Abstract

Biological nervous systems combine sparse connectivity with structured recurrent organization, but it remains unclear whether such structure can serve as a functional prior for artificial embodied control. Here we use BIOSEED, a bounded food/toxin foraging framework, to test whether connectome-derived structural support has behaviorally consequential value under controlled perturbations and learned-controller controls. A staged audit demotes historical shuffled/random controls to fixed-support perturbations, shows that corrected density-matched binary rewiring collapses performance, verifies transfer across implemented environmental shifts, and shows that causal lesions degrade the TRUE structure. A compact top-in 10% scaffold is statistically non-inferior to the full topology in locked validation and transfer, so 85% of the connectome-derived edges can be removed at no behavioral cost. Matched non-biological reductions of equal sparsity reach at most about two thirds of the full-topology score, and connectome-derived selections based on out-degree or synaptic weight also fail, so the sufficiency is specific to incoming-connectivity structure rather than to sparsity or connectome provenance alone. Learned MLP/RNN/GRU controls show that the task is learnable and that intake can be increased, including by a compact GRU approaching the conservative band; however, they do not reduce the scaffold result to a trivial optimization artifact. Live and rover-style demonstrations provide secondary closed-loop diagnostics. BIOSEED therefore identifies a connectome-derived structural prior for conservative embodied control within the tested framework.

Keywords: connectome, NeuroAI, embodied control, C. elegans, structural priors, recurrent controllers

Connectome-derived structural support identifies a compact scaffold for robust embodied control in BIOSEED

Introduction

Biological nervous systems solve sensorimotor problems with connectivity that is sparse, recurrent and strongly structured. [1,2] A central question for computational neuroscience and NeuroAI is whether this structure carries functional information detectable in an embodied task [3], or whether the same behavior could equally be explained by generic recurrence, matched density, arbitrary sparsity, or a learned controller of similar size. This question sits within a longer tradition in embodied artificial intelligence [4,5,6] and dynamical-systems approaches to behavior [7,8], both of which treat adaptive action as an emergent property of an agent's morphology, nervous system and environment rather than of explicit internal representation. Answering it is difficult. A simulated connectome-inspired controller can easily be overinterpreted: a performance gap after perturbation may reflect damage to a tuned controller rather than a meaningful contribution of biological topology; a shuffled control may preserve the relevant support; a shaped score may

hide trade-offs between intake and safety; and a single trajectory can be mistaken for a general behavioral regime.

The C. elegans connectome is an unusually compact, experimentally mapped nervous-system graph. [1,9] Its original wiring diagram has since been reassessed and reconstructed at higher resolution. [10,11] Complementary genomic and circuit-level work has mapped decision-relevant subcircuits onto that structure. [12,13] The connectome also has a well-characterized low-dimensional organization, and minimal circuits within it are sufficient for basic gradient-following behavior, which together make it tractable as a compact computational substrate. [14,15] Yet most connectome-derived models still treat this wiring diagram as a descriptive object to analyse, rather than as a causal, testable prior embedded in a closed embodied control loop. Using it this way creates a methodological obligation: the analysis must not confuse biological inspiration with biological simulation. The work must separate structural support, edge density, weight/sign assignment, controller co-adaptation and learned-controller alternatives, and it must distinguish behavioral claims made inside a simplified task from claims about the living animal. We use BIOSEED under this bounded interpretation.

BIOSEED is a two-dimensional embodied foraging framework in which an agent must approach food, avoid toxin, maintain survival and recover behavior over stochastic trials. This task instantiates a minimal version of the classical foraging problem studied in behavioral ecology, and echoes evidence that C. elegans itself performs quantifiably informative foraging decisions under uncertainty. [16,17] The controller uses a sparse connectome-derived support as a structural prior, but the work does not claim to simulate membrane dynamics, muscles, detailed proprioception, or the complete behavior of C. elegans. Related connectome-derived simulations have shown that this reference structure can drive interactive simulation and klinotaxis-like locomotion. [18,19,20] Robotic implementations of the same wiring diagram have separately produced emergent obstacle-avoidance dynamics, providing prior evidence that it can support qualitatively different embodied behaviors depending on task and embodiment. [21] BIOSEED serves here as a reproducible NeuroAI test case: if a connectome-derived support is placed into a locked embodied controller, which components of that support are behaviorally consequential, and which simpler explanations can be ruled out?

An internal audit revealed that some historical structural controls had preserved the intact binary connectivity, confounding any claim about topology. We therefore rebuilt the evidence around a corrected sequence: density-matched rewiring, replication under real environmental shifts, causal lesions of the intact structure, and tests of whether a compact, structurally selected scaffold is sufficient and transfers across held-out conditions. We benchmark this scaffold against matched non-biological controls and against learned MLP, RNN and GRU policies, to situate BIOSEED's structural prior within the broader space of intake-maximizing and safety-preserving embodied control. [22,23] This question also has a direct machine-learning counterpart: whether a sparse mask, selected without training, can carry task-relevant information beyond its sparsity budget [24,25].

Here we show that the connectome contains an identifiable, compressible control core. A 10% subgraph selected on incoming connectivity is non-inferior to the full connectome-derived support and transfers to held-out environments, whereas equally sparse non-biological reductions do not approach it. The biological signal lies in which sparse subgraph the connectome makes selectable, not in a performance margin over the intact wiring. Density, edge count, generic sparsity, fixed-support perturbation and the learned controllers tested here do not account for the result, and neither do connectome-derived selections built on other structural criteria.

Results

Historical controls required a stricter topology audit

We first revisited the evidence base that had originally motivated this audit: a large N=100 BIOSEED package containing TRUE, shuffled, random-same-sparsity, no-output and internal-ablation variants. These data remain useful background evidence because they show that perturbing the controller can strongly alter foraging, survival and safety metrics. However, they could not by themselves support a final binary-topology claim. The support audit showed that the historical shuffled and random-same-sparsity matrices had support

Jaccard similarity of 1.0 relative to TRUE. [26] Thus, those controls preserved the same non-zero edge coordinates and were fixed-support perturbations of weight order, sign or magnitude rather than true binary rewiring controls.

This audit separates two mechanistically different questions: do weights, signs or magnitudes matter on the same support, or does the binary coordinate support itself matter? The historical controls can address only the first; the final topology claim is assigned to corrected experiments in which edge coordinates are actually changed, following the staged sequence in Fig. 1.

Density-matched rewiring and environmental shifts support a structural effect

The density-matched rewiring control tested whether edge count and density alone were sufficient (an earlier, non-normalized version of this run is retained only as audit provenance, not as a claim). The corrected rewiring contrast randomized non-zero edge coordinates at matched density [27] and applied the appropriate column-wise normalization. Under this corrected design, the TRUE reference remained high while the corrected density-matched rewired control collapsed: TRUE mean score was 28.794, the corrected rewired control was -0.093, and the TRUE-minus-rewired contrast was 28.887.

This result is the first claim-bearing binary-support contrast: preserving only the number and density of non-zero edges does not recover the behavior, so the controller is not simply exploiting a generic sparse recurrent matrix. The coordinate organization of the connectome-derived support itself contributes to performance within the tested simulation and controller design.

The environmental-shift replication then addressed a second risk: that the effect was tied to a nominal environment label rather than real task variation. Across audited implemented environmental shifts summarized in Supplementary Table S1a, connectome-derived support held up while the corrected rewired control remained near collapse. Motor and sensory noise conditions were evaluated separately in the specific held-out transfer block summarized in Supplementary Table S1b. This does not prove that the support will generalize to every possible body or environment, but it removes the narrower explanation that the topology effect was tied only to a single nominal scene.

Causal lesions show that the TRUE structure is behaviorally used

A topology can be statistically associated with performance without being causally used by the controller. The causal lesion test therefore applied targeted lesions to the TRUE structure. [28] Lesions targeting structural components degraded adaptive behavior, supporting the conclusion that the intact connectome-derived organization contributes to the embodied task rather than merely coexisting with a separately sufficient controller.

This result also clarifies how earlier ablation data should be interpreted. Internal-state and output ablations are useful diagnostics, but score changes in shaped metrics can be difficult to interpret when the score includes several engineered terms. The causal lesion test avoids that ambiguity: it perturbs structure directly and measures behavioral degradation, rather than inferring structural relevance indirectly from a composite score.

A compact top-in scaffold is sufficient in validation and transfer

The strongest positive result comes from the scaffold validation and transfer tests. Instead of asking only whether the full connectome-derived support matters, these experiments ask whether a compact substructure can capture a sufficient functional component. Alongside the intact TRUE support we tested two distinct families. The first is connectome-derived: top-in, top-out, top-weight and sensorimotor-boundary selections, each applying a different structural criterion to the same biological wiring. The second is non-biological: reductions matched on density, edge count, degree bins and weight magnitude, which preserve a statistical property of the connectome without preserving which connections it specifies.

In locked validation the 10% top-in scaffold reached 107.8% of the full topology (Fig. 2) (mean 144.497 versus TRUE 133.993; paired difference +10.504, 95% CI +3.195 to +17.813, dz=0.232). [29] Against a pre-specified non-inferiority margin of -5% of the TRUE mean, the scaffold is non-inferior to the intact support (p=3.6e-06), and non-inferiority also holds at margins of -2.5% and -10%. The transfer test reproduces this at 108.9% of TRUE (142.319 versus 130.649; +11.669, 95% CI +6.688 to +16.650, dz=0.266; non-inferiority at -5%:

p=2.6e-12). The validation and transfer environments are disjoint: the scaffold is selected in one set of food/toxin scenes and holds across six held-out environments, including sensory- and motor-noise conditions absent from selection. Removing 85% of the connectome-derived edges therefore costs nothing measurable. The tests were run on 150 pairs in validation (50 seeds across 3 environments) and 300 pairs in transfer (50 seeds across 6 environments); using environments as analysis units, the paired difference is positive in all environments of both blocks (Table 2; validation p=0.0045; transfer p=0.0051).

This does not imply that the remaining edges are irrelevant; it means the relevant behavioral prior is concentrated in a sparse top-in organization. The result is therefore more mechanistic than a full-versus-random comparison: it identifies a sufficient candidate structural backbone.

The working mechanistic hypothesis is that the top-in scaffold captures a high-convergence sensorimotor backbone. Nodes receiving dense incoming connectivity are plausible integration or gating sites where food, toxin and self-state cues can be compressed into motor-relevant control. [30,31] Under this interpretation, the top-in subset is not merely a smaller graph; it preserves the part of the connectome-derived support most likely to transform distributed sensory evidence into coherent action. This hypothesis remains to be tested directly with node-level physiological or circuit analyses, but it provides a concrete mechanistic rationale for why incoming-connectivity concentration can outperform both the full support and matched non-biological reductions in this task.

To probe this hypothesis we mapped the top-in scaffold back onto the source connectome, which is a developmental expansion of the c302 C. elegans model in which neurons are duplicated into functionally labelled modules. Of the 105 top-in nodes, 35 are named biological neurons and the remainder are expansion module copies. Among the biological neurons, most fall within the command and integration layer and the motor output it drives (31 of 35), while primary sensory neurons are rare (4 of 35); named interneurons consistent with this pattern include AVA, AVB, AVD, AVE, PVC, RIA, RIB, AIB and PVP. This class assignment follows the developmental expansion's own module labels and has not yet been cross-checked against the canonical WormAtlas classification, nor tested for statistical enrichment against the full expanded-connectome node background; we therefore report it as a preliminary, qualitative observation consistent with the high-convergence interpretation above, not as a confirmed anatomical enrichment result.

Sparsity does not substitute for connectome structure, and connectome provenance does not substitute for the right criterion

The non-biological control analysis tested whether the top-in result could be explained by simpler non-biological properties. The non-biological controls preserve density, edge count, degree bins or weight magnitude [32] while discarding the identity of the connections the connectome specifies. None approaches the top-in scaffold: the strongest of them, weight-bin matching at 10%, reaches 64.7% of the full topology in validation and 66.4% in transfer, leaving a paired gap of +57.82 (dz=0.907) and +55.52 (dz=0.823). Degree-bin matching reaches 8.9%, random in-node supports 14.0%, and random edge-count supports fall below the no-connectome-output floor.

The connectome-derived family answers a sharper question: whether any selection drawn from the biological wiring would do. It would not. Top-out node selection reaches 32.2% of the full topology and top-weight selection 41.9%, both far below even the non-biological weight-bin control, despite being read off the same connectome. Sensorimotor-boundary selection at 5% behaves differently again, reaching 101.6% and itself non-inferior to the intact support. The connectome therefore contains more than one sufficient sparse core, but it does not make every sparse core sufficient.

Together these analyses bound the claim from both sides. Sparsity matched on density, degree or weight does not reproduce the effect, so the result is not about size. Selections drawn from the same connectome using out-degree or weight do not reproduce it either, so the result is not about provenance. What carries the behavior is the identity of the high-convergence, incoming-connectivity nodes.

Learned controllers expose a trade-off rather than replacing the structural result

The learned-control panel provides a strong competitive control and prevents overclaiming. The comparison is best understood along two axes: absolute task-optimized performance and structural prior. Learned policies

search a task-specific parameter space, whereas the BIOSEED scaffold is fixed and biologically derived; its relevance is a claim about structural organization achieved without task-specific training, not about dominating unlimited optimization budgets. MLP, RNN and GRU policies were optimized in the locked scaffold transfer task under increasing budgets, and they show clearly that BIOSEED is not the best absolute intake or reward maximizer: several learned policies achieved higher mean food intake than the biological baseline, including GRU h=16 (4.292 versus BIOSEED's 2.708) and RNN h=8 configurations above 5.0.

The critical point is not intake superiority. The recurrent sweep reveals a trade-off between foraging gain and safety/coherence (Fig. 3). [33,34,35] Some learned policies increased intake while introducing non-zero toxin hits, wrong-side behavior or collision-like failure modes, but BIOSEED is not uniquely alone in the conservative band either. GRU h=8 GEN4 reached zero toxin hits and full survival, with a wrong-side ratio of 0.0099, while eating more than BIOSEED (3.092 versus 2.708), which rules out “exact clean” or “conservative” as categorical descriptions of BIOSEED alone.

The safe interpretation is that learned controllers can approach or exceed parts of the BIOSEED profile, including raw intake, but no single architecture, budget or matched control reproduces the full structural result. This is not a pass/fail contest; it is a competitive control showing the task admits multiple solutions. The scaffold transfer test defines the boundary of the claim rather than weakening it: the top-in scaffold remains sufficient and transfer-generalizable, unexplained by non-biological controls or learned-controller substitution (Table 3).

Live and rover demonstrations are operational, not primary quantitative evidence

The agent live demo, multi-agent live arena and rover-style embodied package extend the work from offline tables to closed-loop execution. They show that the same local sensorimotor principle can run in real time, accept local food/toxin perturbations and drive an embodied controller. In the multi-agent live arena, BIOSEED and recurrent comparators start from the same common scene and receive identical manual food/toxin injections at matched coordinates. Runtime event-level auditing confirms that apparent mid-run food differences arise from independent per-agent consumption or failure dynamics, not BIOSEED-only spawning (Fig. 4). This setup follows recent work showing that embodied neural agents sharing a common scene can be used to study collective and comparative decision dynamics under controlled perturbation. [36]

The rover demonstration is positioned more cautiously. It is an embodied transposition, showing that the same local control principle can be embedded in a task-oriented body with samples, hazards and return-to-base behavior, rather than a certified autonomous rover-navigation system. Any external hazard-clearance reflex used in a rover variant is an engineering aid, separate from BIOSEED's own behavior, not part of the connectome-derived controller being tested. The quantitative scaffold and control sequence remain the primary evidence; the live and rover tests demonstrate operational closure and embodied plausibility.

Integrated evidence map

The final evidence sequence is therefore cumulative (Fig. 1): the support audit rules out the historical controls as topology evidence, the density-matched rewiring control and its environmental-shift replication supply a corrected binary-support contrast, the causal lesion test shows the intact structure is used rather than merely associated with performance, and the scaffold validation and transfer tests identify a compact sufficient scaffold that non-biological controls and learned policies do not explain away. The live and rover packages close the loop operationally but do not replace the key evidence (Table 1).

Discussion

A compact, connectome-derived support is sufficient in BIOSEED to produce a transferable, safety-preserving sensorimotor regime. This sufficiency survives every simpler explanation tested, matched density, generic sparsity, fixed-support perturbation, non-biological scaffold controls and the learned-controller substitutions used here, none of which reproduce it. The learned-controller results should be read alongside this claim rather than against it: the recurrent sweep does not show that learning fails, nor that BIOSEED is universally superior, but that learned policies can optimize intake and sometimes approach the conservative band without explaining away the biological scaffold result.

The integrated analysis supports a specific conclusion: adaptive behavior in BIOSEED depends on structured support and its co-adaptation with weights, signs and controller dynamics, and a compact top-in scaffold captures a sufficient functional component of that support. The strongest evidence is not the original TRUE-versus-shuffled comparison, since the support audit shows those controls preserved the TRUE binary support, but the corrected evidence sequence summarized above.

This framing makes the result stronger because it is more limited. The work does not claim that BIOSEED is the best possible artificial controller, that C. elegans has been biophysically simulated, or that a universal brain principle has been discovered. This caution is consistent with broader arguments that structural or coding metaphors drawn from small nervous systems should not be generalized to claims about neural computation at large without independent task-level validation. [37] It shows that, in a controlled embodied simulation, a biologically derived structural prior is behaviorally consequential, and that a compact top-in substructure is sufficient and transferable within the tested conditions. That is a meaningful NeuroAI result because it identifies a structural prior that survives a sequence of methodological challenges.

The learned-control results matter here because they show the task is learnable and that biology does not simply win on every metric, since at least one GRU reaches a near-conservative profile while eating more than BIOSEED. This rules out arguments built on exact zeros, without weakening the scaffold analyses' conclusion that learned substitution does not explain away the top-in effect.

The embodied demonstrations follow the same disciplined narrative as the rest of the evidence: the live arena shows the comparator running under symmetric initialization with auditable raw traces, and the rover shows a route toward a more task-oriented body, but neither carries the main causal claim. Engineering reflexes added for hazard clearance in rover variants are external aids, not evidence that BIOSEED alone solves obstacle avoidance.

Connectome-inspired AI requires audit layers, not only larger simulations. [26] The support audit, the density-matched rewiring control, the non-biological control analysis and the scaffold transfer test each remove a different alternative explanation, from fixed-support perturbation to generic sparsity, non-biological scaffolds and weak learned baselines.

These results also connect to a standing question in machine learning about how much of a network's competence lies in its connectivity rather than its weights. The lottery ticket hypothesis established that dense networks contain sparse subnetworks matching full-network performance, though those subnetworks are defined by trainability and recovered by magnitude pruning of an already-trained model [24]. Weight-agnostic architecture search pushed further, finding networks that solve control tasks with no weight training at all [25]. The connectome-derived scaffold occupies a third position: the sparse subgraph is neither trained nor searched, but read off a nervous system mapped by experiment. The comparison sharpens against pruning at initialization, where structural or gradient criteria select a mask before training: such artificial masks can be layerwise-shuffled or reinitialized at little cost to accuracy, which indicates they encode a per-layer sparsity budget rather than the identity of specific connections [38,39,40]. The connectome-derived scaffold behaves the opposite way, since matched reductions collapse and the retained nodes carry information that sparsity alone does not. In this sense artificial pruning criteria appear to encode how much to keep, whereas the connectome encodes what to keep.

Several limitations remain. The task is two-dimensional and simplified. The controller is connectome-derived, not a living-animal simulation. [41,42] The top-in scaffold has not yet been tested in richer three-dimensional bodies, alternative sensory encodings, independent simulators or unlimited learned-controller budgets. The learned sweep uses the tested MLP/RNN/GRU architectures and optimization budgets, not every possible reinforcement-learning or recurrent architecture. The rover and live demos are useful closed-loop demonstrations, but they remain subordinate to the controlled quantitative evidence.

Biological structure should therefore be treated as a source of candidate priors, not as a complete controller blueprint. [3,43] The 10% top-in scaffold identifies a high-convergence subgraph that is sufficient in this task family and motivates targeted tests of integration, gating and transfer in richer bodies. This protocol, audit historical controls for confounded support, rewire at matched density, lesion causally, then test a structurally selected scaffold against non-biological alternatives, does not depend on C. elegans specifically and could be

applied to any compact, mapped nervous system used as a computational prior. It also converts a diffuse claim (“the connectome matters”) into a falsifiable one (“this specific subgraph is sufficient”), which is the kind of claim that later circuit-level or physiological work can directly test or refute.

Connectome-derived structure, in short, can be experimentally interrogated as a causal prior that is not reducible to edge density, generic sparsity, fixed-support perturbation or matched non-biological scaffolds. This audit-to-transfer protocol offers a reproducible framework for studying biological structural priors in artificial embodied agents.

Figures

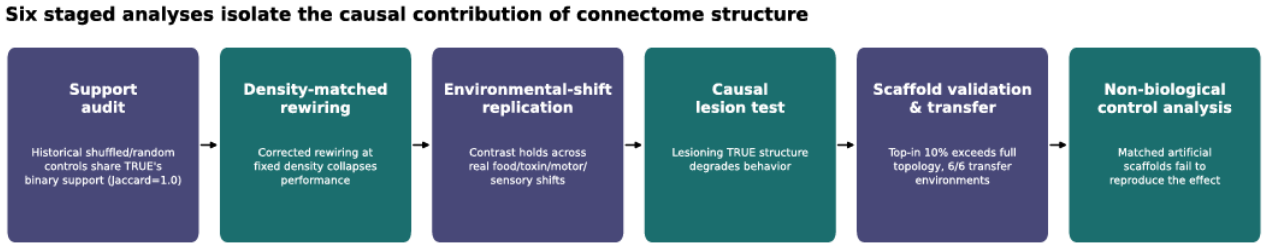


Fig. 1 | Six staged analyses isolate the causal contribution of connectome structure. The sequence moves from a support audit through density-matched rewiring, environmental-shift replication, causal lesion testing, scaffold validation and transfer, to non-biological control analysis. Each step removes one alternative explanation for the behavioral effect.

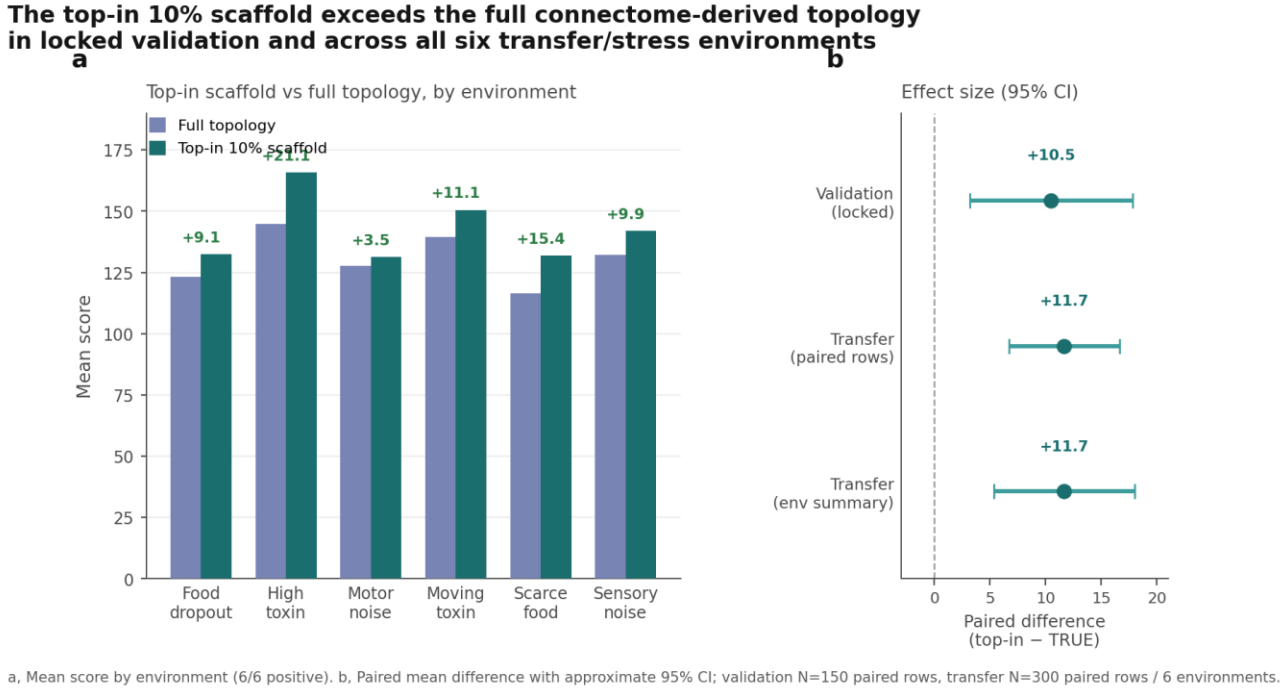


Fig. 2 | The top-in 10% scaffold exceeds the full connectome-derived topology in locked validation and across all six transfer/stress environments. a, Mean score for the full topology and the top-in 10% scaffold across six held-out environments (positive in 6/6). b, Paired mean difference with approximate 95% confidence intervals for validation (n=150 paired trials) and transfer (n=300 paired trials across 6 environments).

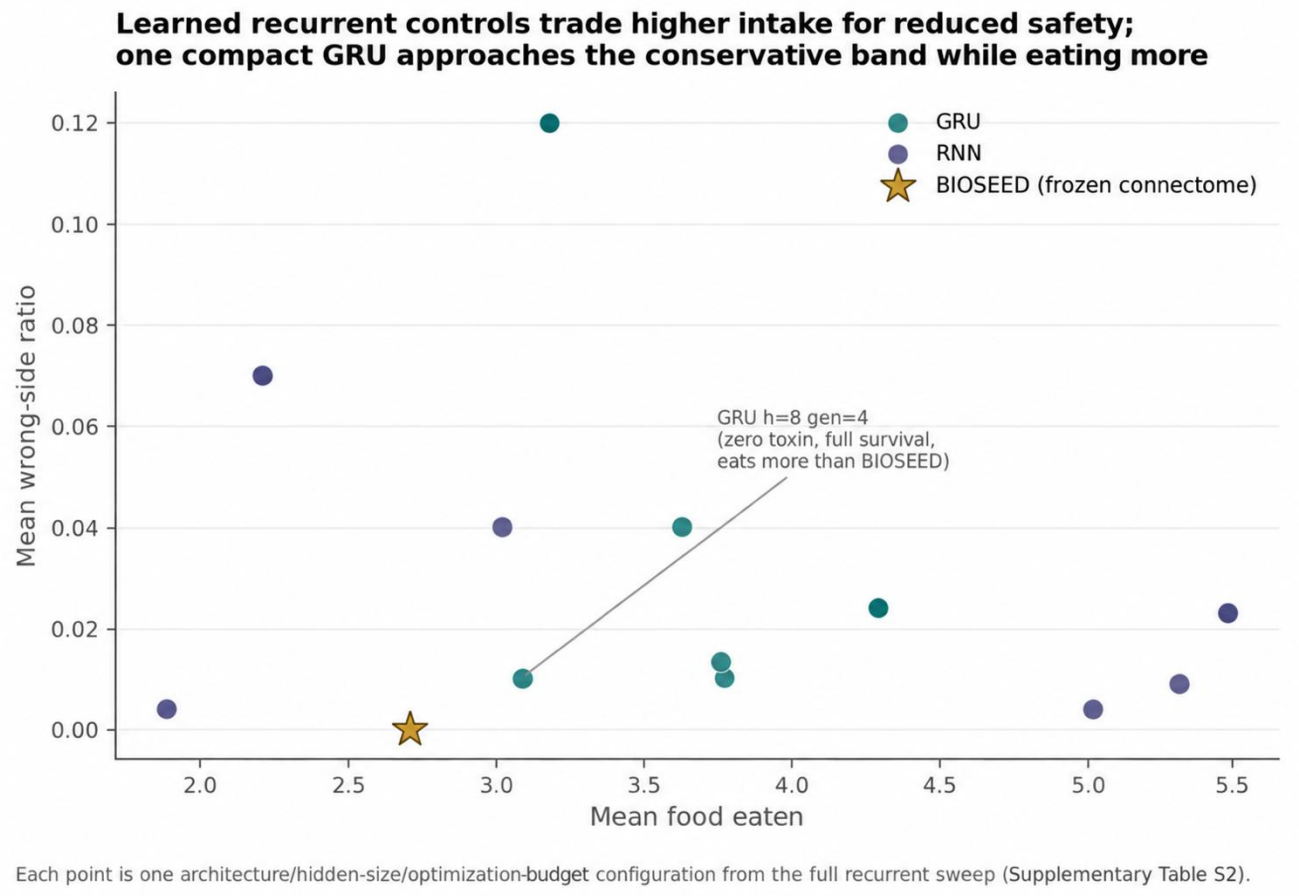


Fig. 3 | Learned recurrent controls trade higher intake for reduced safety, and one compact GRU approaches the conservative band while eating more. Each point is one architecture/hidden-size/optimization-budget configuration from the full recurrent sweep (Supplementary Table S2); the star marks BIOSEED with the frozen connectome-derived controller.

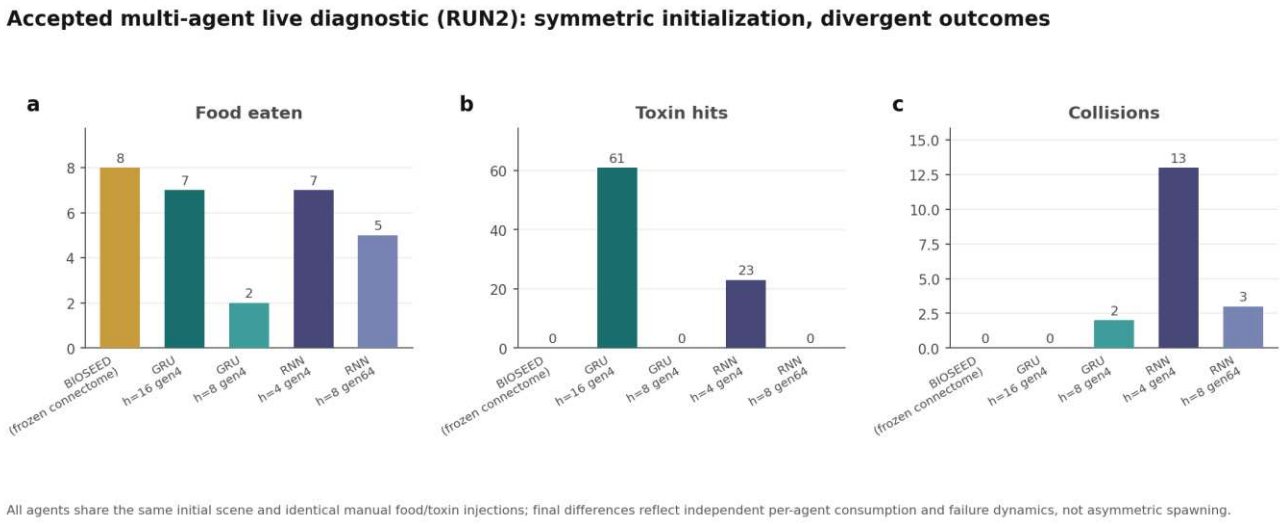


Fig. 4 | Accepted multi-agent live diagnostic (RUN2): symmetric initialization, divergent outcomes. a, Food eaten. b, Toxin hits. c, Collisions. All agents share the same initial scene and identical manual food/toxin injections; differences reflect independent per-agent consumption and failure dynamics, not asymmetric spawning.

Tables

Table 1 | Claim-to-evidence map for the final BIOSEED evidence hierarchy.

Layer	Purpose	Claim-safe interpretation
Step 1 (audit)	Support audit	Historical shuffled/random supports had Jaccard=1.0 versus TRUE; retained as fixed-support perturbations, not final binary rewiring evidence.
Step 2 (rewiring)	Corrected density-matched binary rewiring	TRUE mean 28.794 versus corrected rewired -0.093; support coordinates matter beyond density.
Step 3 (replication)	Implemented environmental shifts	Topology effect retained across environmental shifts; motor/sensory evaluated in a separate transfer.
Step 4 (lesion)	Causal lesions	Lesioning structural components degrades behavior, supporting causal use of the TRUE structure.
Step 5 (scaffold)	Top-in scaffold sufficiency and transfer	Top-in 10% is non-inferior to full topology in locked validation and transfer; paired difference positive in every held-out environment.
Step 6 (control)	Matched non-biological controls	Non-biological matched controls reach at most ~65% of full topology; connectome-derived out-degree and weight selections also fail. Effect specific to incoming connectivity.
Step 7 (learned sweep)	Learned MLP/RNN/GRU controls	Learned policies can increase intake and one GRU approaches the conservative band; the result is a trade-off, not a winner-only claim.
Step 8 (diagnostics)	Operational closure	Single-agent, multi-agent and rover demos show runtime closed-loop operation; they remain secondary to controlled quantitative evidence.

Table 2 | Core scaffold sufficiency and transfer summary.

Block	TRUE mean	Top-in 10% mean	Top-in minus TRUE	Env count	Top-in > TRUE envs	Best non-biological control	Top-in minus best non-bio
Scaffold validation (locked)	133.993	144.497	10.504	3	3	weight_bin_matched_scaffold	57.820
Scaffold transfer (locked)	130.649	142.319	11.669	6	6	weight_bin_matched_scaffold	55.520

Table 3 | Selected learned recurrent controls and the BIOSEED conservative baseline.

Controller	Eaten mean	Wrong-side ratio	Toxin hits mean	Survival steps	Eaten vs BIOSEED
BIOSEED baseline (frozen connectome)	2.708	0.000	0.000	798.383	0.000
GRU h=4 gen=16	3.767	0.010	2.158	798.425	1.058
GRU h=8 gen=4	3.092	0.010	0.000	800.000	0.383
GRU h=16 gen=4	4.292	0.024	0.242	800.000	1.583
RNN h=8 gen=16	5.308	0.009	2.608	798.650	2.600

RNN h=8 gen=64	5.008	0.004	0.892	800.000	2.300
RNN h=16 gen=16	2.208	0.070	0.000	794.400	-0.500
RNN h=16 gen=64	2.208	0.070	0.000	794.400	-0.500

References and declarations

References

1. White, J. G., Southgate, E., Thomson, J. N. & Brenner, S. The structure of the nervous system of the nematode Caenorhabditis elegans. Phil. Trans. R. Soc. B 314, 1-340 (1986).

2. Varshney, L. R., Chen, B. L., Paniagua, E., Hall, D. H. & Chklovskii, D. B. Structural properties of the Caenorhabditis elegans neuronal network. PLoS Comput. Biol. 7, e1001066 (2011).

3. Lappalainen, J. K. et al. Connectome-constrained networks predict neural activity across the fly visual system. Nature 634, 1132-1140 (2024).

4. Brooks, R. A. Intelligence without representation. Artif. Intell. 47, 139-159 (1991).

5. Pfeifer, R., Lungarella, M. & Iida, F. Self-organization, embodiment, and biologically inspired robotics. Science 318, 1088-1093 (2007).

6. Rossi, F. et al. Embodied intelligence: a synergy of morphology, action, perception, and cognition. ACM Comput. Surv. (2025).

7. Beer, R. D. A dynamical systems perspective on agent-environment interaction. Artif. Intell. 72, 173-215 (1995).

8. Buhrmann, T., Di Paolo, E. A. & Barandiaran, X. A dynamical systems account of sensorimotor contingencies. Front. Psychol. 4, 285 (2013).

9. Cook, S. J. et al. Whole-animal connectomes of both Caenorhabditis elegans sexes. Nature 571, 63-71 (2019).

10. Emmons, S. W. The beginning of connectomics: a commentary on White et al. (1986). Phil. Trans. R. Soc. B 370, 20140309 (2015).

11. Brittin, C. A. et al. A multi-scale brain map derived from whole-brain volumetric reconstructions. Nature 591, 105-110 (2021).

12. Bargmann, C. I. Neurobiology of the Caenorhabditis elegans genome. Science 282, 2028-2033 (1998).

13. Jarrell, T. A. et al. The connectome of a decision-making neural network. Science 337, 437-444 (2012).

14. Kunert, J., Shlizerman, E. & Kutz, J. N. Low-dimensional functionality of the Caenorhabditis elegans connectome. Phys. Rev. E 95, 052417 (2017).

15. Izquierdo, E. J. & Lockery, S. R. Evolution and analysis of minimal neural circuits for klinotaxis in C. elegans. J. Neurosci. 30, 12908-12917 (2010).

16. Stephens, D. W., Brown, J. S. & Ydenberg, R. C. (Eds.) Foraging: Behavior and Ecology. University of Chicago Press (2007).

17. Calhoun, A. J., Chalasani, S. H. & Sharpee, T. O. Maximally informative foraging by Caenorhabditis elegans. eLife 3, e04220 (2014).

18. Kim, J., Leahy, W. & Shlizerman, E. Neural interactome: interactive simulation of a neuronal system. Front. Comput. Neurosci. 13, 8 (2019).

19. Morales, P. & Froese, T. Unsupervised learning facilitates neural coordination across the functional clusters of the C. elegans connectome. Front. Robot. AI 7, 40 (2020).

20. Chen, M. et al. Neural model generating klinotaxis behaviour accompanied by a random walk based on C. elegans connectome. Sci. Rep. 12, 3292 (2022).

21. Valencia Urbina, C. E., Cannas, S. A. & Gleiser, P. M. Emergent dynamics in a robotic model based on the Caenorhabditis elegans connectome. Front. Neurorobot. 16, 1041410 (2023).

22. Cho, K. et al. Learning phrase representations using RNN encoder-decoder for statistical machine translation. Proc. EMNLP 2014, 1724-1734 (2014).

23. Chung, J., Gulcehre, C., Cho, K. & Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv:1412.3555 (2014).

24. Frankle, J. & Carbin, M. The lottery ticket hypothesis: finding sparse, trainable neural networks. Proc. ICLR (2019).

25. Gaier, A. & Ha, D. Weight agnostic neural networks. Adv. Neural Inf. Process. Syst. 32, 5365-5379 (2019).

26. Vasa, F. & Misic, B. Null models in network neuroscience. Nat. Rev. Neurosci. 23, 493-504 (2022).

27. Maslov, S. & Sneppen, K. Specificity and stability in topology of protein networks. Science 296, 910-913 (2002).

28. Fakhar, K. & Hilgetag, C. C. Systematic perturbation of an artificial neural network: a step towards quantifying causal contributions in the brain. PLoS Comput. Biol. 18, e1010250 (2022).

29. Lakens, D. Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. Front. Psychol. 4, 863 (2013).

30. Towlson, E. K., Vertes, P. E., Ahnert, S. E., Schafer, W. R. & Bullmore, E. T. The rich club of the C. elegans neuronal connectome. J. Neurosci. 33, 6380-6387 (2013).

31. Gray, J. M., Hill, J. J. & Bargmann, C. I. A circuit for navigation in Caenorhabditis elegans. Proc. Natl Acad. Sci. USA 102, 3184-3191 (2005).

32. Rubinov, M. & Sporns, O. Weight-conserving characterization of complex functional brain networks. NeuroImage 56, 2068-2079 (2011).

33. Friston, K. J. et al. Synthetic spatial foraging with active inference in a geocaching task. Front. Neurosci. 16, 855811 (2022).

34. Aadithya, K. V. & Ravindran, B. Collective foraging of active particles trained by reinforcement learning. Sci. Rep. 13, 17055 (2023).

35. Achiam, J., Held, D., Tamar, A. & Abbeel, P. Constrained Policy Optimization. Proc. 34th ICML, PMLR 70, 22-31 (2017).

36. Coucke, N. et al. Collective decision making by embodied neural agents. PNAS Nexus 4, pgaf101 (2025).

37. Brette, R. Is coding a relevant metaphor for the brain? Behav. Brain Sci. 42, e215 (2019).

38. Lee, N., Ajanthan, T. & Torr, P. H. S. SNIP: single-shot network pruning based on connection sensitivity. Proc. ICLR (2019).

39. Tanaka, H., Kunin, D., Yamins, D. L. K. & Ganguli, S. Pruning neural networks without any data by iteratively conserving synaptic flow. Adv. Neural Inf. Process. Syst. 33 (2020).

40. Frankle, J., Dziugaite, G. K., Roy, D. M. & Carbin, M. Pruning neural networks at initialization: why are we missing the mark? Proc. ICLR (2021).

41. Gleeson, P. et al. OpenWorm: overview and recent advances in integrative biological simulation of Caenorhabditis elegans. Phil. Trans. R. Soc. B 373, 20170382 (2018).

42. Gleeson, P., Steuber, V. & Silver, R. A. c302: a multiscale framework for modelling the nervous system of Caenorhabditis elegans. Phil. Trans. R. Soc. B 373, 20170379 (2018).

43. Chen, B. L., Hall, D. H. & Chklovskii, D. B. Wiring optimization can relate neuronal structure and function. Proc. Natl Acad. Sci. USA 103, 4723-4728 (2006).

44. Rossi, A., Deslauriers-Gauthier, S. & Natale, E. On null models for temporal small-worldness in brain dynamics. Netw. Neurosci. 8, 377-394 (2024).

Methods

Overview of the experimental design

The study was designed to test whether a compact, connectome-derived scaffold supports closed-loop sensorimotor behavior in a food/toxin navigation task, and whether the observed behavior can be explained by generic sparsity, matched density, fixed-support perturbation, non-biological scaffold controls, or compact learned neural controllers.

The validation framework has three levels. The primary quantitative validation evaluates BIOSEED and its structural variants under controlled food/toxin environments, including the support audit, the density-matched rewiring control, the environmental-shift replication, the causal lesion test, the scaffold validation and transfer tests, the non-biological control analysis, and the learned-controller sweep. The competitive learned-controller controls compare BIOSEED against compact MLP, RNN and GRU policies trained within the same task family, to test whether task-specific learned controllers can match or exceed the connectome-derived controller on intake, reward, safety and survival metrics. Live and embodied demonstrations test whether the same local sensorimotor principle can run in real time and be embodied in a rover-like control loop; these are diagnostic and operational, not the primary quantitative benchmark.

BIOSEED controller and connectome-derived scaffold

BIOSEED uses a fixed, connectome-derived structural support as a sensorimotor controller. Unlike the learned controls, the controller is not trained end-to-end on the target task; its structure is derived from biological connectivity and evaluated as a fixed prior for closed-loop behavior.

To instantiate the BIOSEED control architecture, we used an expanded c302/C. elegans-derived connectome representation as the structural scaffold of the agent. The resulting computational network comprised 1,052 neuronal nodes and 12,707 synaptic connections, including both chemical synapses and gap junctions. Each connection was mapped into a weighted connectivity matrix and used to propagate activity across the agent's internal recurrent network. Sensory variables from the environment, including resource-related, hazard-related and internal-state signals, were injected into defined sensory channels, while downstream network activity was read out through motor-related pools to generate movement decisions. This representation was used as a structural control scaffold rather than as a full organism-level or biophysical simulation of C. elegans.

The study distinguishes the full biological scaffold from compact subsets. The top-in scaffold selects a compact subset of the structural support using a structural criterion, incoming connectivity, rather than task performance, and is evaluated against the full topology, shuffled controls, random controls, matched non-biological scaffolds and learned policies.

The central question is not whether BIOSEED is the best possible intake maximizer, but whether biologically derived structural organization provides a non-trivial, transferable and safety-preserving control prior under the tested conditions.

Controlled food/toxin environments

The primary task family consists of closed-loop food/toxin navigation in a two-dimensional environment. Agents receive local sensory information about food and toxin cues, generate motor outputs, and accumulate behavioral outcomes over time. Core behavioral quantities are food eaten, toxin hits, survival steps, wrong-side ratio, energy dynamics, collisions and final alive/dead state. Scaffold comparisons use one composite

score (agent9_v2_score), combining intake, approach and productivity terms with safety penalties and internal-state terms, applied identically across every scaffold family, keep fraction and control, with weights fixed in the versioned analysis script. Because the composite includes internal-state terms whose values depend on the architecture under test, we confirm the scaffold result on foraging_score, a metric with no internal-state reward: the top-in scaffold reaches 106.5% of the intact support in validation and 107.2% in transfer, positive in every environment of both blocks. Absolute score magnitudes are not comparable across the rewiring, lesion and scaffold blocks, which use different score scales and environment sets, so cross-block statements are given as a fraction of the TRUE reference within the same block. The evaluations are designed to distinguish raw intake maximization from conservative, safety-preserving behavior: a controller can be high-intake but less conservative, or lower-intake but safer and more coherent.

Structural evidence sequence

The primary validation follows a corrected sequence of structural controls (see Results): a support audit that verifies the binary connectivity of historical controls against established null-model methodology [44], a density-matched rewiring control, a replication of the topology contrast under real environmental shifts, a causal lesion test of the intact structure, and validation and transfer tests of a structurally selected top-in scaffold against matched non-biological alternatives. Each analysis is based on the corresponding final tables, figures and raw output files.

Non-biological scaffold and perturbation controls

To test whether BIOSEED behavior is explained by generic sparsity or density alone, the study includes non-biological scaffold and perturbation controls. These ask whether matched edge count, matched density, or non-biological graph structure (degree-bin, weight-bin, random in-node and random edge-count supports) can reproduce the same behavioral regime, separating biological structural support from generic sparse-network effects. Connectome-derived selections that are not based on incoming connectivity (top-out, top-weight and sensorimotor-boundary) are analysed separately, as biological alternatives rather than as non-biological controls. The intended claim is not that no artificial controller can ever match BIOSEED, but that the tested non-biological structural controls do not explain away the observed top-in scaffold effect within this framework.

Shifted-environment and transfer tests

Shifted-environment tests evaluate whether the scaffold effect transfers beyond a single fixed scene, distinguishing scene-specific tuning from more general local sensorimotor organization. Shifted-environment variants are described as transfer or robustness checks; where a shifted condition required correction after audit, only the corrected version is used as evidence.

Learned-controller sweep

The learned-controller sweep provides the principal competitive control against compact recurrent policies. It evaluates learned MLP, RNN and GRU policies across architecture and budget variants, compared with BIOSEED under the same task family, to test whether compact task-trained controllers can exceed BIOSEED on intake or reward, and whether they reproduce the same conservative behavioral profile. The sweep uses the tested architectures and budgets and is not exhaustive of all possible optimization strategies; unlimited reinforcement-learning budgets, richer recurrent architectures and scaffold-constrained learned initializations remain future work.

Single-agent live demonstration

A single-agent live demonstration runs the frozen connectome-derived controller in a real-time two-dimensional food/toxin environment, supporting manual inspection of local cue processing, food approach, toxin avoidance, energy dynamics and safety/backout behavior. It provides diagnostic visualization and a sanity check that the final controller operates in closed loop, complementing rather than replacing the controlled quantitative validation.

Multi-agent live comparator

A multi-agent live comparator places BIOSEED and learned recurrent comparators, GRU and RNN policies across architecture variants, in a shared diagnostic scene. All agents are initialized from a common food/toxin layout, with the initial food and toxin counts recorded; manual food/toxin injection events are replicated to

every panel at the same coordinates, after which each agent evolves in its own local copy of the environment, so that food consumed by one agent disappears only in that agent's panel. Mid-run differences in food counts are therefore expected and reflect independent consumption or failure dynamics, not asymmetric spawning. Runtime outputs include per-agent timestep-level state and metrics, the manual injection trace with coordinates and timing, and provenance metadata. This comparator is a diagnostic and fairness visualization, useful for verifying that no controller-specific injection artifact is present, but it is not the primary quantitative benchmark.

Rover embodied demonstration

A rover demonstration evaluates whether the same local sensorimotor principle can be embedded in a more complex embodied system, with task-oriented components such as samples, hazards, home/base return, loading/unloading and battery-like constraints. It is not a certified rover-autonomy benchmark; it is an embodied closed-loop demonstration that the BIOSEED control principle can be placed in a rover-like body. The reference version is reflex-free; any variant that includes an additional hazard-clearance reflex or external engineering assist is treated as a separate engineering aid, not attributed to BIOSEED itself. The rover supports operational and embodied plausibility but does not replace the quantitative food/toxin benchmark.

Statistical reporting

Primary comparisons report paired mean differences, confidence or bootstrap intervals, a paired-test p-value and an effect-size estimate, with sample definitions stated explicitly. For small environment-level comparisons, such as the six transfer/stress families, the p-value reflects the paired environment-level comparison and is interpreted together with the directionality across all environments, not as a standalone universal-stability claim. Within each pre-specified family of contrasts we report Benjamini-Hochberg FDR and Bonferroni corrected p-values. All contrasts of the top-in 10% scaffold against non-biological controls survive both corrections by wide margins. The superiority contrast against the intact support survives FDR in validation ($p=0.0045$) but not Bonferroni ($p=0.063$), which is why that claim is stated as non-inferiority rather than superiority; in transfer it survives both (Bonferroni $p=5.6e-05$). Paired scaffold statistics are reported at both the trial level (150 validation and 300 transfer pairs, from 50 seeds across 3 and 6 environments) and the environment level, since seeds recur across environments.

Data and code availability

Each controlled or live run preserves output sufficient for audit and reproduction: timestep-level results, aggregate summary tables by controller and condition, provenance metadata (seed, input root, environment, variant, engine hash and agent list), live-injection traces for replay, and file-integrity hash manifests.

Acknowledgements

The authors thank their host institutions for their support during this work.

Author contributions

- Conceptualization: P.C., L.P.
- Software: P.C.
- Investigation: G.G., P.C., L.P
- Data curation: G.G., P.C., L.P.
- Visualization: G.G.
- Writing – original draft: P.C., G.G.
- Writing – review & editing: P.C., L.P., G.G.
- Supervision: L.P.

Competing interests

The authors declare no competing interests.

Data availability

The data generated in this study have been deposited in Zenodo under accession code https://doi.org/10.5281/zenodo.21263275. Access is currently under embargo and will be made publicly available upon publication of this article.

Code availability

The code used to generate and analyse the results in this study is available at the same Zenodo repository (https://doi.org/10.5281/zenodo.21263275), currently under embargo and scheduled for public release upon publication.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- BIOSEEDSupplementaryTables.pdf