From Connectome Structure to Signal Propagation: Cross-Level Inference Boundaries in *C. elegans* Whole-Brain Dynamics

Zhiqi Shen[1,2]

[1]Guangdong Technion–Israel Institute of Technology, Shantou, Guangdong 515063, China
[2]Southern University of Science and Technology, Shenzhen 518055, China

August 2026

Abstract

Connectomes constrain possible routes of neural communication but do not specify effective propagation or multi-input computation. We evaluated three levels of inference using the *C. elegans* connectome and held-out single-neuron perturbation responses: pairwise propagation, node-level hub ranking, and a conditional dual-input model. Binary connectome support predicted held-out stimulus-triggered average magnitudes with Pearson correlation $\rho = 0.054$, corresponding to approximately 31% of the training-to-test reproducibility benchmark ($\rho = 0.173$). The evaluated degree-preserving rewired control reached $\rho = 0.011$. Homogeneous multi-hop propagation and the tested low-parameter temporal model did not improve on the binary one-hop model. The connectome-constrained edge-specific linear dynamical system of Creamer et al. (2024) reached $\rho = 0.140$, or approximately 81% of the benchmark. This edge-level advantage did not transfer robustly to hub ranking: a larger raw correlation for the edge-specific propagated map disappeared in a sensitivity analysis that conditioned on observation coverage. In the conditional dual-input model, linear internal dynamics remained additive. An imposed local tanh nonlinearity generated a nonzero state-level interaction but little interaction in the selected motor readout. Convergence-based placement was significant relative to degree- and cell-type-matched placements for the chemical–temperature configuration ($p = 0.001$), but not for chemical–oxygen ($p = 0.46$). The chemical–temperature effect was attenuated under a dual-drive-matched control ($p = 0.29$). Within the tested model families, anatomical support contributed limited predictive information, whereas fitted edge heterogeneity accounted for most of the quantitative gain. These results do not identify a biological integration mechanism. Factorial whole-brain $0/A/B/A{+}B$ measurements are required to test multi-input interaction directly.

1 Introduction

A central problem in systems neuroscience is to determine how anatomical organization constrains neural dynamics and computation. A connectome records synaptic contacts and, in some reconstructions, their multiplicity and subcellular location, but it is not itself a dynamical model (Bargmann and Marder, 2013). Effective interactions also depend on quantities that electron microscopy does not fully specify, including synaptic sign and gain, neuronal and synaptic time constants, recent activity, and internal state. Chemical and electrical synapses are also not the only communication channels. Monoaminergic and peptidergic signaling form partially distinct network layers, and transmitter identity need not determine a target-specific postsynaptic sign (Bentley et al., 2016; Fenyves et al., 2020; Ripoll-Sánchez et al., 2023). Functional interactions can therefore vary with time and context on a fixed anatomical scaffold (Yezerets et al., 2025). The

mapping from structure to function is consequently an inference problem involving anatomical support, effective edge dynamics, network propagation, internal state, and measurement.

This problem is unusually accessible in compact nervous systems with cell-resolved connectomes. The synaptic wiring of *Caenorhabditis elegans* was first reconstructed at whole-animal scale by serial electron microscopy and has since been quantified across adult sexes, individuals, and developmental stages (White et al., 1986; Cook et al., 2019; Witvliet et al., 2021). In *Drosophila*, dense central-brain and whole-brain reconstructions provide complementary systems in which structural hypotheses can be evaluated at larger scale (Scheffer et al., 2020; Dorkenwald et al., 2024). Recent fly models illustrate both the value and the limitation of connectome constraints. Connectivity can support detailed neural predictions when combined with task optimization, transmitter information, or physiological validation, but those predictions depend on the additional assumptions and data used to assign dynamics to the wiring diagram (Lappalainen et al., 2024; Shiu et al., 2024).

C. elegans provides a particularly direct setting for structure–function comparison because its neurons have stereotyped identities across animals and a large fraction of the nervous system can be recorded at cellular resolution. Whole-brain imaging has revealed low-dimensional population dynamics associated with motor programs, enabled recordings in freely moving animals, and shown that neural representations vary across behavioral states and timescales (Kato et al., 2015; Nguyen et al., 2016; Atanas et al., 2023). The NeuroPAL color atlas further permits identified neurons to be matched across individuals and experimental modalities (Yemini et al., 2021). These developments allow anatomical wiring, measured neural dynamics, and dynamical models to be compared using corresponding named neurons rather than anonymous population channels.

Randi et al. extended this comparison from spontaneous correlation to direct perturbation by activating individual neurons with two-photon optogenetics while recording brain-wide calcium activity (Randi et al., 2023). Their atlas contains measurements for 23,433 directed neuron pairs and reports the sign, magnitude, direction, and temporal profile of stimulus-evoked propagation. These measurements represent effective signal propagation. A response in neuron j after stimulation of neuron i may contain contributions from a direct synapse, recurrent multi-hop paths, and extrasynaptic communication. Anatomy was informative in aggregate because directly connected pairs and pairs separated by shorter anatomical paths were more likely to show detectable propagation. However, anatomy-based biophysical predictions agreed poorly with the measured response amplitudes. Experiments in animals impaired in dense-core-vesicle release also identified fast propagation that depended on extrasynaptic signaling and was not represented by the wired connectome (Randi et al., 2023).

Creamer et al. subsequently asked how much of the perturbational dataset could be described by dynamics whose recurrent support was restricted to chemical synapses and gap junctions (Creamer et al., 2024). They fitted a latent linear dynamical system in which an interaction weight could be nonzero only for anatomically connected neuron pairs, while the sign and magnitude of each allowed weight were learned from whole-brain perturbation recordings. Recurrent evolution allowed the model to generate responses between neurons without a direct synapse through multi-hop propagation. On held-out animals, predicted stimulus-triggered response magnitudes correlated with measured responses at approximately 0.14. The corresponding training-to-test reproducibility was approximately 0.17, giving a relative correlation of approximately 0.82. This result supports the importance of the specific anatomical support, but it does not show that raw synapse counts or an unparameterized topology provide the same prediction. The successful model combines anatomical sparsity with edge-specific effective weights learned from functional data.

This distinction leaves three inference questions. First, how much held-out single-stimulus propagation is supplied by anatomical edge support, and how much additional performance is

associated with edge-specific effective dynamics? Second, does improved pairwise prediction transfer to node-level summaries such as hub ranking, given that anatomical and empirical signaling networks can differ at that scale (Dvali et al., 2025)? Third, can single-stimulus propagation determine where two inputs interact? Randi et al. and Creamer et al. both identify simultaneous multi-neuron perturbations and nonlinear extensions as necessary next steps (Randi et al., 2023; Creamer et al., 2024). The present project contains no factorial $0/A/B/A{+}B$ whole-brain recordings. We therefore treat the dual-input analysis as an in silico sensitivity study rather than evidence for an endogenous integration mechanism. We use a nested model ladder to separate binary support, homogeneous multi-hop propagation, low-parameter temporal dynamics, and edge-specific fitted dynamics. We then test whether edge-level predictive differences persist after aggregation to hub ranking and examine what structural convergence can establish when a local nonlinearity is imposed explicitly.

2 Results

2.1 Binary connectome support predicts a limited fraction of held-out propagation

We first asked how much of the held-out STAM matrix could be predicted without assigning an independent effective weight to every anatomically allowed edge. The model ladder compared one-hop propagation, a homogeneous multi-hop resolvent, a low-parameter continuous-time model, and the edge-specific Creamer LDS. The first three families were evaluated using both the project-supplied signed operator and the corresponding binary support.

The signed variants performed near zero on the held-out test data: $\rho = 0.010$ for one-hop propagation, -0.003 for the resolvent, and 0.009 for the low-parameter temporal model (figure 1 and table 1). The binary one-hop and binary resolvent models reached $\rho = 0.054$ and 0.053, respectively. Relative to the training-to-test reproducibility benchmark $\rho_{\mathrm{tt}} = 0.173$, both values correspond to approximately 31%. The evaluated degree-preserving rewired control reached $\rho = 0.011$. Binary support therefore exceeded this negative control, although a distribution over independent rewirings would be required for a formal topology test.

The homogeneous multi-hop resolvent did not improve on binary one-hop propagation. The tested low-parameter temporal model also failed to bridge the gap to the edge-specific LDS, which reached $\rho = 0.140$, or approximately 81% of the benchmark. Within the model families tested here, binary support contained limited held-out information, whereas the main quantitative gain was associated with edge-specific fitted dynamics and recurrent propagation. This comparison does not establish that edge-specific fitting is the only possible intermediate description.

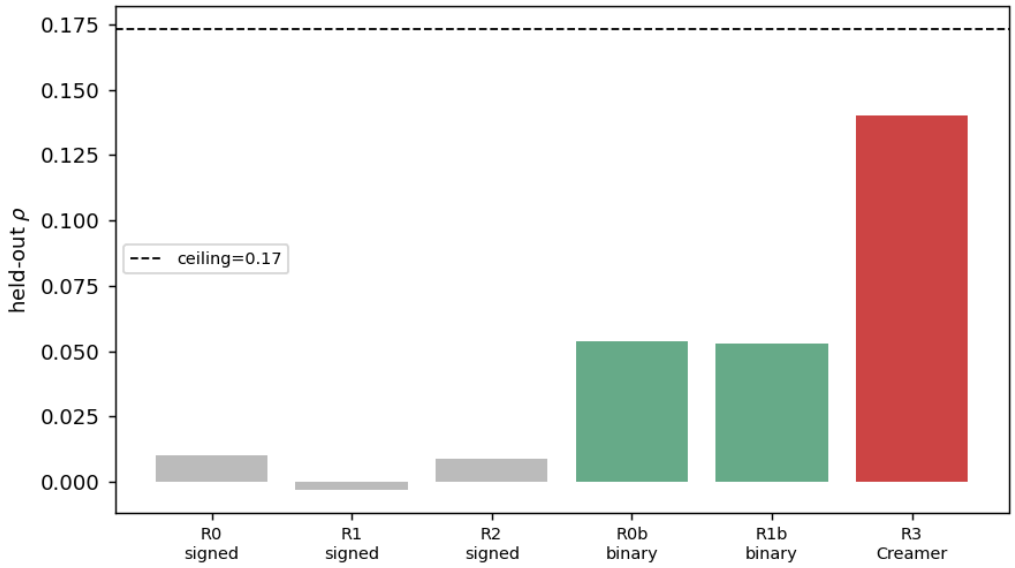

Figure 1: Held-out single-stimulus prediction across the model ladder. Bars show Pearson correlations between predicted and measured test-set STAMs. The dashed line is the training-to-test reproducibility benchmark ($\rho_{\mathrm{tt}} \approx 0.173$), not a strict upper bound. Signed anatomical models performed near zero, binary one-hop and homogeneous multi-hop models reached $\rho \approx 0.054$, and the connectome-constrained edge-specific LDS reached $\rho = 0.140$.

Table 1: Held-out single-stimulus prediction. The final column is normalized by the training-to-test reproducibility benchmark $\rho_{\mathrm{tt}} = 0.173$.

Model	Free parameters	Test ρ	ρ/ρ_{tt}
Degree-preserving rewired support	—	0.011	0.06
Signed one-hop (R0)	1	0.010	0.06
Signed resolvent (R1)	2	-0.003	≈ 0
Signed temporal dynamics (R2)	3	0.009	0.05
Binary one-hop (R0b)	1	0.054	0.31
Binary resolvent (R1b)	2	0.053	0.31
Binary temporal dynamics (R2b)	3	-0.024	< 0
Creamer edge-specific LDS (R3)	edge-specific	**0.140**	**0.81**

2.2 Edge-level prediction does not transfer robustly to hub ranking

The first result concerns prediction of directed neuron pairs. We next asked whether the same ordering persists after aggregation to node-level in-strength. Motivated by the reported divergence between anatomical and empirical signaling architectures (Dvali et al., 2025), we compared direct anatomical support A, homogeneous multi-hop propagation $G_A = (I - \gamma P_A)^{-1} - I$, and the full-path response magnitude G_C generated by the edge-specific Creamer dynamics. Each representation was compared with the mean absolute held-out STAM received by each neuron.

Before conditioning on observation coverage, the edge-specific representation had the larger Spearman correlation with empirical hub ranking: 0.265 for G_C and 0.164 for G_A (figure 2 and table 2). Empirical in-strength was strongly anticorrelated with the number of available incoming observations ($\rho = -0.51$), and G_C shared part of this dependence. A partial-rank sensitivity analysis conditioned on observation coverage gave correlations of 0.202 for G_C and 0.239 for G_A. The raw ordering therefore did not persist.

A shuffled-support negative control gave a coverage-conditioned correlation of 0.005. This

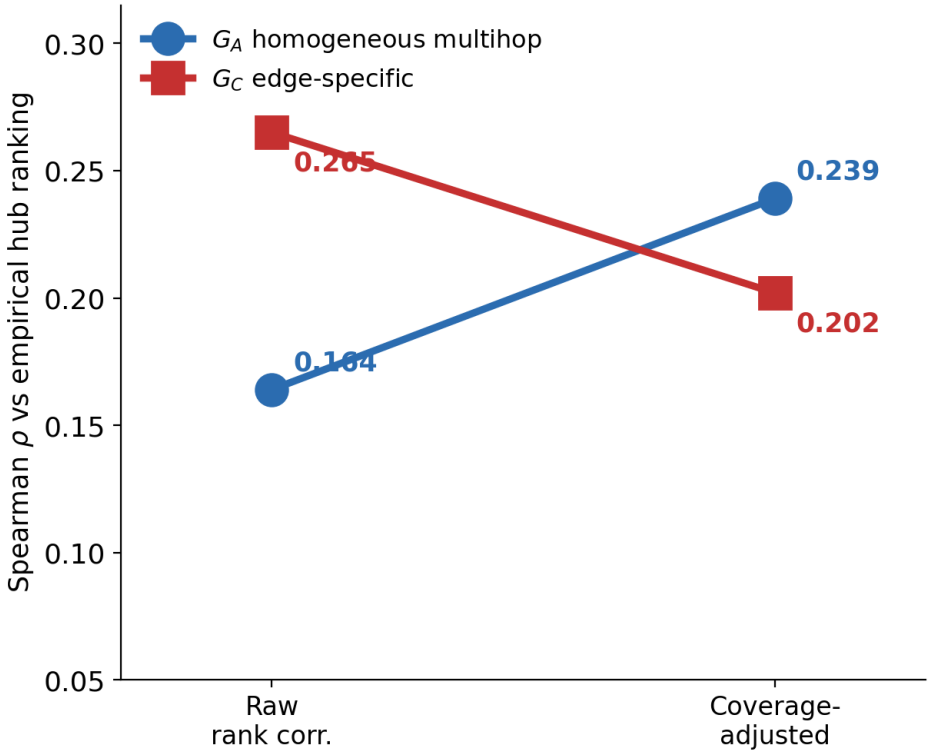

Figure 2: Sensitivity analysis of empirical hub ranking. The edge-specific propagated map G_C had a larger raw Spearman correlation than homogeneous multi-hop propagation G_A. The ordering was numerically reversed after conditioning on per-neuron observation coverage. This analysis is exploratory because coverage is not a randomized covariate and no independent network replicates were available.

Table 2: Hub-ranking sensitivity analysis. Partial Spearman correlations condition on the number of observed incoming stimulus-response pairs.

Representation	Raw ρ	ρ vs. coverage	Partial ρ
Direct anatomy A	0.152	−0.020	0.163
Homogeneous multi-hop G_A	0.164	0.080	**0.239**
Edge-specific propagation G_C	**0.265**	−0.185	0.202
Shuffled-support negative control	0.112	−0.211	0.005

control is not a calibrated null because relocating learned weights changes node-strength marginals and recurrent feedback. The partial-rank analysis also cannot correct arbitrary nonrandom missingness and no independent network replicates were available for uncertainty estimation. We therefore interpret the hub-ranking result as an exploratory boundary test. It shows that the edge-specific advantage in pairwise prediction does not automatically imply a corresponding advantage in node-level ranking.

2.3 Structural convergence increases interaction only under an imposed nonlinearity

2.3.1 Internal-state and readout interactions are not equivalent

The Randi atlas contains single-neuron perturbations rather than factorial two-neuron perturbations. We therefore did not estimate an empirical two-stimulus interaction. Instead, we fixed the Creamer edge-specific dynamics as a single-stimulus-calibrated backbone and inserted a local

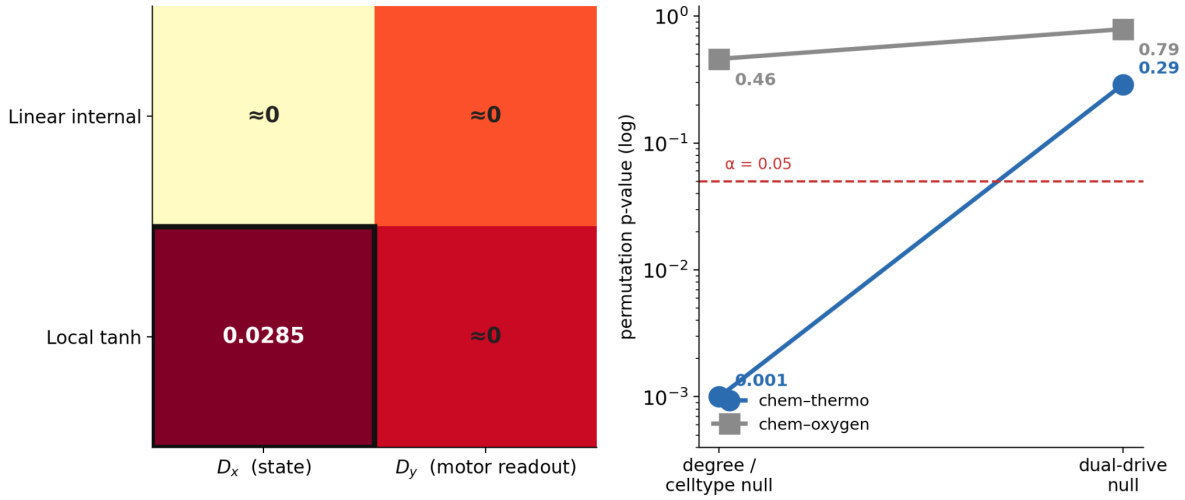

Figure 3: Conditional dual-input analysis. Left: linear internal dynamics obeyed superposition, whereas an imposed local tanh generated state-level nonadditivity. The selected motor readout remained nearly additive. Right: at the configuration level, convergence placement was significant relative to the degree- and cell-type-matched control for chemical–temperature ($p = 0.001$), but not for chemical–oxygen ($p = 0.46$). Under the stricter dual-drive-matched control, the chemical–temperature effect was attenuated ($p = 0.29$), while chemical–oxygen remained nonsignificant ($p = 0.79$). These are permutation results from the in silico model and not measurements of biological two-stimulus integration.

Table 3: Factorial interaction energies for the tested internal dynamics and fixed motor readout.

Internal dynamics	Readout	D_x	D_y
Linear	fixed motor softmax	7.8×10^{-30}	4.9×10^{-13}
Local tanh	fixed motor softmax	**0.0285**	$< 5 \times 10^{-7}$

nonlinearity at selected nodes.

For baseline, single-input, and joint-input trajectories, the state-level factorial residual was

$$\Delta \boldsymbol{x}^{AB}(t) = \boldsymbol{x}^{AB}(t) - \boldsymbol{x}^{A}(t) - \boldsymbol{x}^{B}(t) + \boldsymbol{x}^{0}(t), \tag{1}$$

and the same definition was applied to the motor readout $\boldsymbol{y}(t)$. With linear internal dynamics, both interaction energies were zero to numerical precision: $D_x = 7.8 \times 10^{-30}$ and $D_y = 4.9 \times 10^{-13}$ (figure 3 and table 3). Introducing a local tanh nonlinearity produced a nonzero state-level interaction, $D_x = 0.0285$, while the selected motor readout remained near numerical zero, $D_y < 5 \times 10^{-7}$.

This result distinguishes two levels of analysis. A negligible output interaction does not imply a negligible internal-state interaction. Conversely, a nonlinear readout can generate output nonadditivity even when internal dynamics are linear. The present result is limited to the tested stimulus amplitudes, imposed nonlinearity, and fixed motor readout.

2.3.2 Matching local dual drive attenuates the convergence effect

We placed the same gain perturbation on the six highest-ranked structural convergence nodes, defined as neurons receiving direct chemical input from both stimulus modalities, and compared the resulting state interaction with matched placements. At the configuration level, convergence placement exceeded the degree- and cell-type-matched control for chemical–temperature ($p = 0.001$), but not for chemical–oxygen ($p = 0.46$; figure 3). Pair-level standardized effects were

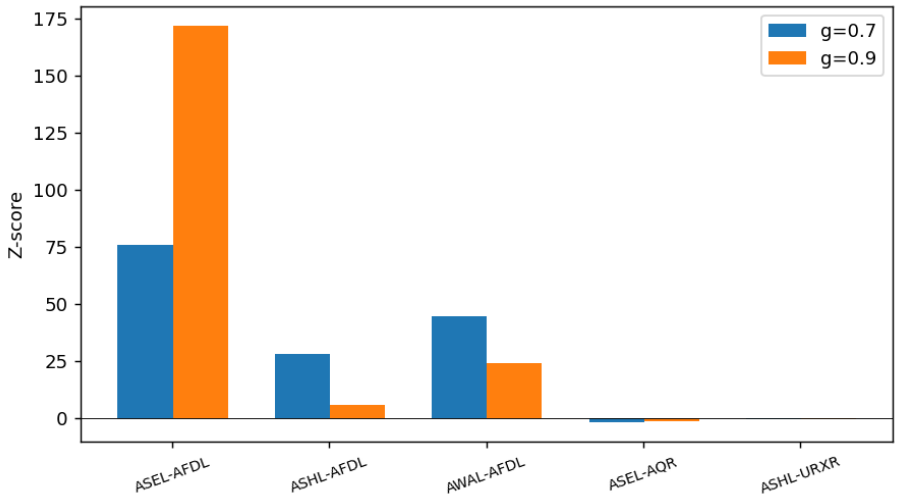

Figure 4: Descriptive pair-level standardized effects of convergence placement relative to degree- and cell-type-matched placements at gains $g = 0.7$ and 0.9. Positive effects were concentrated in the three chemical–temperature pairs, whereas chemical–oxygen pairs were near zero. The five pairs are not independent biological replicates because they share modality groups, candidate nodes, and model parameters.

positive for the three chemical–temperature pairs and near zero for the two chemical–oxygen pairs (figure 4).

The chemical–temperature effect was attenuated when control placements were additionally matched for the two local single-stimulus drives. The configuration-level permutation value increased to $p = 0.29$, while chemical–oxygen remained nonsignificant at $p = 0.79$. Failure to reject under this stricter control does not establish equivalence, and the control matches only the implemented dual-drive statistic. The supported interpretation is therefore limited: structural convergence identifies locations at which two pathways deliver strong local drive, but it does not establish a specialized integrator identity or a general cross-modal mechanism.

3 Discussion

The single-stimulus model ladder separates information in anatomical support from information introduced by fitted edge dynamics. Binary support predicted approximately one third of the training-to-test reproducibility benchmark. Homogeneous multi-hop propagation and the tested low-parameter temporal model did not improve on this value. The larger performance of the Creamer model was associated with edge-specific effective weights fitted from neural data and recurrent propagation on the anatomical support (Creamer et al., 2024). Because the model ladder is not exhaustive, the result does not prove that every anatomical edge requires an independent biological parameter.

Predictive performance must also be evaluated at the scale at which a claim is made. The edge-specific model improved pairwise prediction, but its raw advantage in hub ranking did not persist after conditioning on observation coverage. This sensitivity analysis does not identify the correct node-level architecture because it cannot remove arbitrary missingness and lacks independent network replicates. It does show that pairwise predictive performance cannot be transferred automatically to hub-level interpretation.

The dual-input analysis has a stronger limitation. Shared paths remain additive under linear dynamics, so structural overlap is not itself a nonlinear mechanism. After a local saturating nonlinearity was imposed, structural convergence produced larger interaction than degree- and cell-

type-matched placements in one sensory configuration. Matching local dual drive attenuated this effect. The model therefore supports a conditional architectural statement about concentration of input, not an empirical claim that the selected neurons are specialized integrators.

The main limitations follow directly from the data and design. No factorial whole-brain $0/A/B/A{+}B$ measurements were available. The tanh nonlinearity, gain values, and motor readout were imposed rather than measured. The five stimulus pairs form only two modality configurations and are not independent biological replicates. The node-level analysis is sensitive to incomplete observation coverage. The present models also omit state dependence and wired–wireless interactions known to affect *C. elegans* signaling (Randi et al., 2023; Yezerets et al., 2025).

A decisive test requires a factorial perturbation experiment with baseline, stimulus A, stimulus B, and simultaneous $A{+}B$ conditions; whole-brain single-cell readout; matched stimulus intensities; animal-level replication; and predefined downstream variables. Such data would permit direct estimation of

$$I_{ij\to k}(t) = r_k(A{+}B, t) - r_k(A, t) - r_k(B, t) + r_k(0, t), \tag{2}$$

and would distinguish additive propagation, saturation, competition, state switching, and nonlinear readout. Structural convergence, receptor expression, neuromodulatory layers, and fitted effective dynamics could then be evaluated as held-out predictors of measured multi-input interaction.

4 Methods

4.1 Data sources, neuron alignment, and notation

The analysis used processed data distributed with the `Creamer_LDS_2024` repository (Creamer et al., 2024), derived from the single-neuron perturbation atlas of Randi et al. (2023), together with the chemical-synapse and gap-junction connectome of Cook et al. (2019). Neuron identities were pre-aligned by the source repository. Analyses were restricted to the common set of $N = 154$ head neurons with functional coverage and model parameters. Matrix element M_{ji} maps stimulation or activity of neuron i to the response of neuron j.

The STAM matrix $R \in \mathbb{R}^{N\times N}$ is the area under the stimulus-triggered calcium response during the 30-s post-stimulus window for each directed neuron pair. The source split contained 80 training animals and 30 held-out test animals. Diagonal and nonfinite entries were excluded from matrix-level correlations. The empirical training-to-test correlation was $\rho_{\mathrm{tt}} = 0.173$. Following Creamer et al. (2024), this value was used as a reproducibility benchmark rather than a strict upper bound.

4.2 Single-stimulus model ladder

Let P denote the propagation operator constructed from either the project-supplied signed representation or its binary support. Each model generated a predicted response matrix $\widehat{R}$.

R0: one-hop propagation.

$$\widehat{R}_{\mathrm{R0}} = aP, \tag{3}$$

where the scalar a was fitted using the training STAM matrix.

R1: homogeneous multi-hop propagation.

$$\widehat{R}_{\mathrm{R1}} = a\left[(I - \gamma P)^{-1} - I\right]. \tag{4}$$

The parameter γ was restricted to the stable resolvent regime and selected with a using the training data.

R2: low-parameter continuous-time propagation. The dynamics were

$$\dot{\boldsymbol{x}}(t) = \left(-\tau^{-1} I + \beta P\right) \boldsymbol{x}(t) + \boldsymbol{u}(t). \tag{5}$$

The predicted STAM was obtained by integrating the impulse response over the 30-s post-stimulus window. The global scale a, decay time τ, and coupling β were fitted using the training matrix.

R3: connectome-constrained edge-specific LDS. The R3 prediction was taken from the fitted Creamer LDS. Recurrent weights are nonzero only on chemical-synapse or gap-junction support and are otherwise learned edge by edge from the training recordings (Creamer et al., 2024). The high-dimensional model was not refitted in this project.

For R0–R2, scalar parameters were selected using only the training matrix. Performance was the Pearson correlation between the predicted and held-out test STAM values over common finite off-diagonal entries. Normalized performance was $\rho_{\text{test}}/\rho_{\text{tt}}$.

4.3 Degree-preserving rewired control

The binary support was randomized by directed double-edge swaps that preserve in-degree and out-degree while changing edge identities, following the degree-preserving rewiring principle of Maslov and Sneppen (2002). The one-hop model was then fitted and evaluated using the same procedure. The reported value $\rho = 0.011$ corresponds to the evaluated rewired realization. It is treated as a negative control because a distribution over independent rewirings was not available for formal permutation inference.

4.4 Node-level hub-ranking analysis

For each receiving neuron j, empirical in-strength was defined as the mean absolute held-out response across the set $\mathcal{O}_j$ of observed stimulating neurons,

$$s_j^{\text{emp}} = \frac{1}{|\mathcal{O}_j|} \sum_{i \in \mathcal{O}_j} |R_{ji}^{\text{test}}|, \tag{6}$$

with observation coverage $n_j^{\text{obs}} = |\mathcal{O}_j|$. The anatomical representation A used the binary union of chemical and gap-junction support. Homogeneous multi-hop propagation was

$$G_A(\gamma) = (I - \gamma P_A)^{-1} - I, \tag{7}$$

and the edge-specific propagated representation was the time-integrated absolute response generated by the Creamer dynamics,

$$(G_C)_{ji} = \sum_t \left| h_{ji}^{\text{C}}(t) \right|. \tag{8}$$

Incoming strengths were computed with the same orientation across representations.

Raw rank agreement was evaluated with Spearman correlation. The coverage-conditioned sensitivity analysis applied the standard partial-correlation expression to ranked variables,

$$\rho_{xy \cdot z} = \frac{\rho_{xy} - \rho_{xz}\rho_{yz}}{\sqrt{(1 - \rho_{xz}^2)(1 - \rho_{yz}^2)}}, \tag{9}$$

where $z = n^{\mathrm{obs}}$. The shuffled-support negative control placed learned edge-weight values on a shuffled support without refitting. This operation changes node-strength marginals and recurrent feedback and was not interpreted as a calibrated anatomical null. No independent network replicates were available for uncertainty estimation. The analysis was therefore treated as exploratory. The binary signaling network of Dvali et al. (2025) was not available in this project and was used only to motivate the cross-scale question.

4.5 Conditional dual-input dynamics

The edge-specific matrix W^{C} was rescaled by its spectral radius and used as a fixed recurrent backbone. Internal dynamics were

$$\boldsymbol{x}_{t+1} = (1-\alpha)\boldsymbol{x}_t + \alpha\,\phi\Big[\mathrm{diag}(\boldsymbol{g})\left(W^{\mathrm{C}}\boldsymbol{x}_t + a_u\boldsymbol{u}_t\right)\Big], \tag{10}$$

with $\alpha = 0.3$ and input amplitude $a_u = 2.0$. The linear baseline used $\phi(z) = z$, and the conditional nonlinear model used $\phi(z) = \tanh z$. A fixed motor-neuron projection C followed by a softmax defined the output,

$$\boldsymbol{y}_t = \mathrm{softmax}(C\boldsymbol{x}_t). \tag{11}$$

The projection and simulated time window were held fixed across baseline, single-input, and joint-input conditions.

For each stimulus pair, four trajectories were simulated: zero input, stimulus A, stimulus B, and simultaneous $A + B$. Factorial residuals included the zero-input baseline,

$$\Delta\boldsymbol{x}^{AB}(t) = \boldsymbol{x}^{AB}(t) - \boldsymbol{x}^{A}(t) - \boldsymbol{x}^{B}(t) + \boldsymbol{x}^{0}(t), \tag{12}$$

$$\Delta\boldsymbol{y}^{AB}(t) = \boldsymbol{y}^{AB}(t) - \boldsymbol{y}^{A}(t) - \boldsymbol{y}^{B}(t) + \boldsymbol{y}^{0}(t). \tag{13}$$

The interaction energies were

$$D_x = \sum_t \|\Delta\boldsymbol{x}^{AB}(t)\|_2^2, \qquad D_y = \sum_t \|\Delta\boldsymbol{y}^{AB}(t)\|_2^2. \tag{14}$$

The zero-input terms are required when the baseline state or readout is nonzero.

4.6 Convergence placement and matched controls

Structural convergence candidates were neurons receiving direct chemical-synapse input from both modality-specific input groups. The six highest-ranked convergence nodes were selected with the project convergence score. All nodes had a common baseline gain $g \in \{0.7, 0.9\}$, and candidate nodes received an additional $\delta g = 0.3$. The five stimulus pairs were grouped into two configurations: chemical–temperature (ASEL–AFDL, ASHL–AFDL, and AWAL–AFDL) and chemical–oxygen (ASEL–AQR and ASHL–URXR). Pairs within a configuration were not treated as independent biological replicates.

Two permutation controls were evaluated with $B = 999$ placements. The degree- and cell-type-matched control sampled node sets of the same size with approximately matched structural degree and neuron-class composition. The dual-drive control additionally matched the implemented local drive statistic computed from two independently simulated single-stimulus trajectories. Pair-level placement effects were standardized relative to the matched-control distribution. The configuration-level statistic was the median standardized effect across constituent pairs, and permutation values were obtained by recomputing this statistic for each matched placement.

For an observed configuration statistic T_{obs} and permutation statistics T_b, one-sided permutation

values were

$$p = \frac{1 + \sum_{b=1}^{B} \mathbb{I}(T_b \geq T_{\text{obs}})}{B + 1}. \tag{15}$$

The exact algebraic definition of the dual-drive matching score is specified in the active analysis code rather than in the source draft. Full reproducibility therefore requires the archived code version used to generate the reported controls.

4.7 Software and reproducibility

Analyses were performed in Python 3.11 using NumPy, SciPy, NetworkX, and scikit-learn. The active scripts documented in the project report are `track_a_ladder.py`, `track_b_fixed.py`, `track_c_arch_hierarchy.py`, and `_audit_hub_coverage.py`. A complete archival release should include the exact construction of the signed operator, matrix normalization, parameter grids, fitting objective, random seeds, and dual-drive matching score. These details are not fully recoverable from the manuscript text alone.

Acknowledgements

This work was developed during the 2026 PEBBLE BioFusion Workshop at Westlake University in Hangzhou, China, held from July 24 to August 4, 2026.

Data and code availability

The functional perturbation data are available through the Randi atlas and the processed Creamer repository (Randi et al., 2023; Creamer et al., 2024). Connectome data are from Cook et al. (2019). The workshop analysis scripts have not yet been deposited in a public archival release.

References

Atanas, A. A. et al. (2023). "Brain-wide representations of behavior spanning multiple timescales and states in C. elegans". In: *Cell* 186.19, 4134–4151.e31. DOI: 10.1016/j.cell.2023.07.035.

Bargmann, C. I. and Marder, E. (2013). "From the connectome to brain function". In: *Nature Methods* 10.6, pp. 483–490. DOI: 10.1038/nmeth.2451.

Bentley, B. et al. (2016). "The multilayer connectome of Caenorhabditis elegans". In: *PLoS Computational Biology* 12.12, e1005283. DOI: 10.1371/journal.pcbi.1005283.

Cook, S. J. et al. (2019). "Whole-animal connectomes of both Caenorhabditis elegans sexes". In: *Nature* 571.7763, pp. 63–71. DOI: 10.1038/s41586-019-1352-7.

Creamer, M. S., Leifer, A. M., and Pillow, J. W. (2024). "Bridging the gap between the connectome and whole-brain activity in C. elegans". In: *bioRxiv*. Version 4, revised 18 May 2026. DOI: 10.1101/2024.09.22.614271.

Dorkenwald, S. et al. (2024). "Neuronal wiring diagram of an adult brain". In: *Nature* 634, pp. 124–138. DOI: 10.1038/s41586-024-07558-y.

Dvali, S. et al. (2025). "Diverging network architecture of the C. elegans connectome and signaling network". In: *PRX Life* 3.3, p. 033021. DOI: 10.1103/6wgv-b9m6.

Fenyves, B. G. et al. (2020). "Synaptic polarity and sign-balance prediction using gene expression data in the Caenorhabditis elegans chemical synapse neuronal connectome network". In: *PLoS Computational Biology* 16.12, e1007974. DOI: 10.1371/journal.pcbi.1007974.

Kato, S. et al. (2015). "Global brain dynamics embed the motor command sequence of Caenorhabditis elegans". In: *Cell* 163.3, pp. 656–669. DOI: 10.1016/j.cell.2015.09.034.

Lappalainen, J. K. et al. (2024). "Connectome-constrained networks predict neural activity across the fly visual system". In: *Nature* 634, pp. 1132–1140. DOI: 10.1038/s41586-024-07939-3.

Maslov, S. and Sneppen, K. (2002). "Specificity and stability in topology of protein networks". In: *Science* 296.5569, pp. 910–913. DOI: 10.1126/science.1065103.

Nguyen, J. P. et al. (2016). "Whole-brain calcium imaging with cellular resolution in freely behaving Caenorhabditis elegans". In: *Proceedings of the National Academy of Sciences of the United States of America* 113.8, E1074–E1081. DOI: 10.1073/pnas.1507110112.

Randi, F. et al. (2023). "Neural signal propagation atlas of Caenorhabditis elegans". In: *Nature* 623.7986, pp. 406–414. DOI: 10.1038/s41586-023-06683-4.

Ripoll-Sánchez, L. et al. (2023). "The neuropeptidergic connectome of C. elegans". In: *Neuron* 111.22, 3570–3589.e5. DOI: 10.1016/j.neuron.2023.09.043.

Scheffer, L. K. et al. (2020). "A connectome and analysis of the adult Drosophila central brain". In: *eLife* 9, e57443. DOI: 10.7554/eLife.57443.

Shiu, P. K. et al. (2024). "A Drosophila computational brain model reveals sensorimotor processing". In: *Nature* 634, pp. 210–219. DOI: 10.1038/s41586-024-07763-9.

White, J. G. et al. (1986). "The structure of the nervous system of the nematode Caenorhabditis elegans". In: *Philosophical Transactions of the Royal Society of London. B, Biological Sciences* 314.1165, pp. 1–340. DOI: 10.1098/rstb.1986.0056.

Witvliet, D. et al. (2021). "Connectomes across development reveal principles of brain maturation". In: *Nature* 596.7871, pp. 257–261. DOI: 10.1038/s41586-021-03778-8.

Yemini, E. et al. (2021). "NeuroPAL: A multicolor atlas for whole-brain neuronal identification in C. elegans". In: *Cell* 184.1, 272–288.e11. DOI: 10.1016/j.cell.2020.12.012.

Yezerets, E., Mudrik, N., and Charles, A. S. (2025). "Decomposed Linear Dynamical Systems (dLDS) models reveal instantaneous, context-dependent dynamic connectivity in C. elegans". In: *Communications Biology* 8, p. 1218. DOI: 10.1038/s42003-025-08599-3.