SUPERFLY: LEARNING LOCOMOTION WITH A FLY CONNECTOME

Anonymous authors
Paper under double-blind review

ABSTRACT

We present SUPERFLY, a sensorimotor network built from the adult *Drosophila* brain and ventral nerve cord connectome for embodied locomotion. The policy maps inputs to body controls through about 23,000 neurons and 2.2 million directed edges, fixing each connection's endpoints, synapse count, and excitatory or inhibitory sign. Training learns nonnegative scales that convert fixed synapse counts into synaptic currents. Mappings based on known sensorimotor pathways route body feedback, movement commands, and previous controls to corresponding neuron populations. Each neuron maintains voltage across body steps, and decoder readouts map its activity to controls for a MuJoCo fly model. Behavioral cloning learns about 37,000 parameters from demonstrations by a flybody walking expert. Compared with a parameter-matched graph neural network (GNN), the policy achieves lower validation action error (MSE 0.069 versus 0.096) and travels farther in five-second walking trials (8.5 versus 0.4 cm toward a 10 cm target). Although training provides only expert body controls and no neural activity targets, turn commands selectively modulate neuron types with established roles in navigation and steering. These responses persist after rewiring, showing that turn-command selectivity does not require the exact synaptic partners. The rewired policy fits actions better offline but travels less in closed loop. This executable model links learned locomotion to identified biological neurons and predicts their activity during behavior. We make our code, checkpoints, and an interactive website for exploring neuron activity available online.

1 INTRODUCTION

Evolution supplies a fly with sparse neural wiring shaped to support its sensorimotor behaviors (Zador, 2019). Electron-microscopy connectomes now expose this structure at synaptic resolution. FlyWire reconstructs 139,255 neurons and about 54.5 million synapses in an adult *Drosophila* brain (Dorkenwald et al., 2024; Schlegel et al., 2024). The Brain and Nerve Cord connectome (BANC) extends reconstruction through the ventral nerve cord, connecting central computation to leg control (Bates et al., 2026). Cell annotations, synapse counts, and predicted transmitter signs provide a biological scaffold for a policy (Eckstein et al., 2024).

The navigation circuit provides concrete targets for interpretation. EPG neurons form an internal compass whose activity peak tracks heading, including in darkness (Seelig & Jayaraman, 2015; Turner-Evans et al., 2020). PFN neurons combine heading and self-motion to represent travel direction (Hulse et al., 2021; Lyu et al., 2022). PFL3 compares heading with a goal and relays steering output toward paired DNa02 descending neurons (Westeinde et al., 2024; Mussells Pires et al., 2024). Recordings and stimulation link DNa02 activity to turning toward the more active side (Rayshubskiy et al., 2025; Yang et al., 2024). Superfly retains these identities within a brain-and-cord circuit rather than prescribing their responses as training targets.

Anatomy alone does not specify how observations drive neurons or how activity controls a simulated body. FlyGM establishes that a whole-brain connectome can support walking, turning, and flight (Jin et al., 2026). We pursue the complementary goal of optimizing a fixed brain-and-cord connectome for locomotion while preserving the biological identity of every modeled neuron. Each state variable should remain interpretable as the activity of a particular cell, enabling direct comparison with neuroscience experiments and targeted perturbations.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

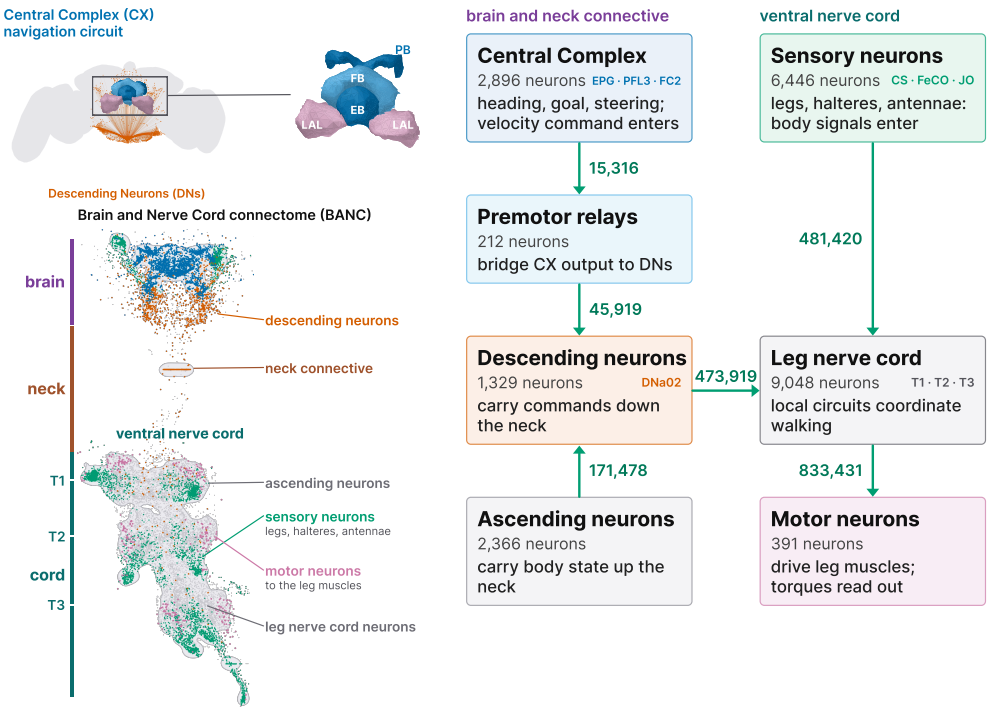

Figure 1: **From navigation circuits to leg control.** The FlyWire brain view locates the Central Complex (CX) and descending neurons (DNs), with navigation-related regions in the inset. BANC cell-body locations show the modeled brain-and-cord populations; T1–T3 identify leg neuromeres. Population cards summarize selected pathways and their synapse counts. Computation uses all 2.2 million retained edges, including pathways omitted from the schematic. Twenty-eight sensory routes deliver body signals, commands target 87 FC2 cells, and descending/motor readouts generate 59 controls. Every modeled activity trace remains attached to an identified BANC neuron.

We present **Superfly**, a recurrent policy with 22,688 BANC neurons and 2,222,379 fixed, signed connections. Anatomical interfaces deliver 75 body readings, 59 previous controls, and three velocity commands to matched populations. Passive point-neuron dynamics retain scalar voltage across body steps; seven linear readouts produce 59 controls. Behavioral cloning fits 37,296 parameters while retaining every edge's endpoints, count, and sign. Figure 2 summarizes the policy.

Our experiments distinguish action imitation, closed-loop control, and neural interpretability. Under matched training, Superfly attains 0.0686 validation MSE and travels 8.476 cm toward a target 10 cm ahead, compared with 0.0957 and 0.402 cm for the GNN. The rewired Null fits actions better (0.0595) but travels less (4.505 cm); the dense MLP is strongest overall (0.0361, 9.768 cm). Thus, low offline error does not determine reliable gait. Superfly's scalar dynamics also reduce counted operations by 103× relative to generic feature message passing on the same graph.

The model additionally exposes neuron-level responses without neural supervision. With body inputs and initial state fixed, opposite yaw commands recruit PFL2, PFL3, DNa03, and DNa02, cell types implicated in fly steering. At the final update, their mean current responses are 11.7, 14.5, 131, and 9.05 times their respective population averages. The same families remain turn responsive after degree-preserving rewiring. These observations support interpretable, cell-indexed predictions while limiting any claim that selectivity uniquely identifies the native topology.

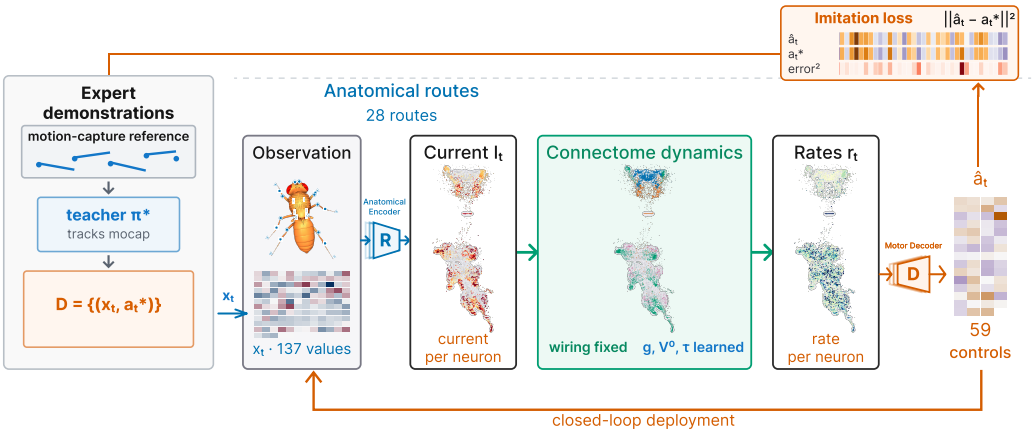

Figure 2: **Superfly's closed-loop policy.** Anatomical interfaces convert 137 inputs into current at selected BANC neurons. Four 0.5 ms updates propagate scalar voltage over fixed signed connections; the final voltage persists into the next 2 ms body step. Nonnegative learned scales calibrate fixed synapse counts. Seven linear readouts convert rectified activity into 59 controls. Behavioral cloning fits interfaces, cellular dynamics, synaptic scales, and readouts using motor targets alone. Each internal state retains its biological neuron identity.

2 Related Work

Connectomes as executable models. Lappalainen et al. freeze measured visual-system connectivity and optimize cellular and synaptic quantities for optic-flow estimation, recovering identified cell responses without neural targets (Lappalainen et al., 2024). Shiu et al. simulate the adult fly brain with a fixed leaky integrate-and-fire model (Shiu et al., 2024). Superfly adopts task-optimized mechanistic modeling for locomotion, with scalar state and anatomically restricted interfaces spanning sensory input through motor output. FlyGM instead uses learned graph-neural-network updates for broad embodied control (Jin et al., 2026). Our matched GNN tests that update family in the present task and training budget; it is not a reproduction of FlyGM's full training procedure.

Structure in embodied control. Morphology-based graph policies encode relationships between a robot's joints and limbs, while imitation supplies control targets. Superfly changes the source of structure from the designed body to the reconstructed nervous system. NeuroMechFly and flybody provide biomechanical fly environments, including the walking expert cloned here (Lobato-Rios et al., 2022; Vaxenburg et al., 2025). Anatomical interpretation follows established navigation pathways: EPG represents heading, PFN combines heading and self-motion, and PFL populations translate heading-goal relationships into descending steering commands (Seelig & Jayaraman, 2015; Hulse et al., 2021; Lyu et al., 2022; Westeinde et al., 2024). DNa02 activity predicts turning and shapes leg kinematics (Rayshubskiy et al., 2025; Yang et al., 2024).

Testing the wiring prior. Inherited structure may reduce what must be learned (Zador, 2019), but topology requires controlled comparisons. Degree-preserving nulls distinguish exact partners from coarse graph statistics (Dhiman, 2026). We compare a parameter-matched MLP against the entire structured policy and a matched rewire against native synaptic partners. The rewire preserves degrees, population-pair counts, source-assigned signed synapse counts, interfaces, dynamics, and training budget. We evaluate one rewire and one GNN configuration, so the findings do not estimate variation across rewires or a hyperparameter search. Appendix A gives extended context.

3 A CONNECTOME-CONSTRAINED LOCOMOTION POLICY

3.1 DEMONSTRATIONS AND ANATOMICAL GRAPH

The velocity-tracking SAC teacher released with flybody (Vaxenburg et al., 2025) executes three commands from each of 25 motion-capture postures. Straight walking requests 2.0 cm/s with zero yaw; turns request 1.75 cm/s with yaw ± 0.35 rad/s. Each rollout spans 2,500 steps of 2 ms. The resulting 187,500 observation-control pairs are partitioned by complete posture: 150,000 frames from 20 postures for training and 37,500 from five disjoint postures for validation. All commands from a posture stay in one split.

At step t, the input is $x_t = (s_t, a_{t-1}, c_t) \in \mathbb{R}^{137}$: 75 body readings, 59 preceding controls, and a three-dimensional velocity command. Targets $a_t^\star \in [-1, 1]^{59}$ comprise 48 leg-joint, six adhesion, and five head/abdomen controls. Unlike the teacher, Superfly does not observe reference or future target poses. Inputs and targets use training-set means and standard deviations, with a 10^{-3} floor on each standard deviation.

BANC supplies 22,688 neurons in 3,430 cell types across sensory, Central Complex (CX), ascending, descending, leg-cord, motor, and premotor-relay populations. Its 2,222,379 directed edges aggregate about 8.5 million chemical synapses. Following the canonical table, we exclude autapses and retain all other edges among these populations. For edge $u \to v$, the source's transmitter fixes $s_u \in \{-1, +1\}$ and the reconstructed count fixes n_{uv}. The effective coupling is

$$W_{uv} = s_u n_{uv} g_{\kappa(u),\rho(v)}, \qquad g_{\kappa,\rho} \geq 0, \tag{1}$$

where $\kappa(u)$ is source cell type and $\rho(v)$ is target population. The 14,235 learned scales calibrate current without changing endpoints, signs, or relative counts within a group. This sharing adapts the cell-type parameterization of Lappalainen et al. (2024).

3.2 INPUT INTERFACES AND NEURAL DYNAMICS

Twenty-eight anatomical sensory routes, split by recipient type into 462 channels, reach 6,858 neurons (6,446 sensory and 412 ascending). Each channel q selects fixed observation entries $\mathcal{S}_q$ and recipients $\mathcal{V}_q$. Its learned affine map produces a shared current:

$$\iota_t^q = p_q^\top \widetilde{x}_t[\mathcal{S}_q] + b_q, \qquad I_t^{\text{sens}}[v] = \sum_{q: v \in \mathcal{V}_q} \iota_t^q. \tag{2}$$

Organ, side, and leg annotations determine the route. For example, a front-left joint-speed channel reads only that leg's speed and drives its annotated femoral chordotonal neurons. Optimization changes its magnitude but not its selected inputs or recipients.

A bilateral command map sends forward drive equally to left and right FC2 neurons and reverses lateral/yaw drive between sides. It reaches 87 FC2A/B/C cells upstream of steering pathways. Previous controls return downstream: five body controls reach 1,329 descending neurons, and six leg-specific groups of nine controls reach matching motor neurons (391 total). Total external current is $I_t = I_t^{\text{sens}} + I_t^{\text{cmd}} + I_t^{\text{eff}}$. There is no observation-to-readout shortcut.

Each neuron has scalar voltage $V_{t,k}[v]$, with body step t and within-step update k. Threshold-linear release and recurrent input are

$$r_{t,k}[v] = \text{ReLU}(V_{t,k}[v]), \qquad m_{t,k}[v] = \sum_{u \to v} W_{uv} r_{t,k}[u]. \tag{3}$$

A passive leak integrates recurrent and external current:

$$V_{t,k+1}[v] = V_{t,k}[v] + \frac{\Delta t}{\tau_{\kappa(v)}} \left[-V_{t,k}[v] + V^0_{\kappa(v)} + m_{t,k}[v] + I_t[v] \right]. \tag{4}$$

Each cell type shares a fitted resting voltage V^0 and time constant τ. Four Euler updates with $\Delta t = 0.5$ ms span one body step; I_t stays fixed while release and recurrent current are recomputed. The final state initializes the next step, $V_{t+1,0} = V_{t,4}$, and each rollout begins at fitted rest. After optimization steps, scales are clamped nonnegative and $\tau \geq \Delta t$. Initial time constants are 5 ms, resting voltages follow $\mathcal{N}(0, 0.05^2)$, and each synaptic group starts with mean absolute edge weight 0.01. This is a circuit-level model, not a simulation of spikes or dendritic compartments.

Table 1: **Offline action error and five-second gait.** MSE uses standardized controls. Travel, absolute lateral drift, and absolute heading error average five straight-command rollouts. The best numerical value in each column is bold; the GNN's small drift reflects near immobility. All policies have 37,296 parameters.

Policy	Train MSE	Val. MSE	Travel (cm)	Drift (cm)	Heading (rad)
Superfly	0.1042	0.0686	8.476	2.008	0.655
GNN	0.1268	0.0957	0.402	**0.133**	0.741
MLP	**0.0480**	**0.0361**	**9.768**	1.273	**0.198**
Null	0.0860	0.0595	4.505	5.940	1.886

3.3 Motor decoding and learning

Seven linear readouts decode $r_{t,4}$. One reads all descending neurons to predict five body controls; each of six others reads only one leg's motor neurons to predict eight joint controls and one adhesion control. The concatenated standardized prediction $\widehat{\widetilde{a}}_t$ is trained with

$$\mathcal{L}_{\mathrm{BC}} = \frac{1}{59|\mathcal{B}|} \sum_{t \in \mathcal{B}} \left\| \widehat{\widetilde{a}}_t - \widetilde{a}_t^{\star} \right\|_2^2 . \tag{5}$$

The 37,296 trainable scalars comprise 5,978 input-interface parameters, 21,095 circuit parameters (14,235 scales and 3,430 each of resting voltages and time constants), and 10,223 readout parameters. Nonlinear feature construction occurs in the recurrent circuit.

Training processes 16 streams of up to 16 consecutive frames per step. Voltage persists across chunks from a rollout but is detached at chunk boundaries, giving 32 ms truncated backpropagation. Adam uses weight decay 10^{-5} and learning rates 10^{-4} for circuit dynamics and 5×10^{-3} for interfaces/readouts. A cosine schedule reduces both rates to 1% over 50 epochs (7.5 million frame presentations). We select the lowest validation-MSE checkpoint before closed-loop evaluation. During deployment, predictions are unstandardized, clipped to $[-1, 1]$, and applied to MuJoCo for 2 ms; executed controls and persistent voltage feed the next step. Appendix B provides the full routing, normalization, and update procedure.

4 Experiments and results

4.1 Matched baselines and evaluation

All four policies have 37,296 parameters and share observations, control targets, posture splits, training presentations, and validation-based checkpoint selection. The **MLP** maps standardized observations directly to controls. The **GNN** retains BANC nodes, edges, interfaces, and readouts but replaces scalar dynamics with learned message passing over 49–54 features per neuron. The **Null** swaps two postsynaptic targets within a fixed source/target population pair, preserving each neuron's in- and out-degree, population-pair edge counts, and each source's signed synapse counts. Cell types, parameter sharing, interfaces, dynamics, and readouts remain unchanged.

The five-second evaluation averages five held-out-posture rollouts commanded straight at 2.0 cm/s toward a 10 cm target (Table 1). A separate ten-second evaluation uses the same cohort of 25 held-out postures for every policy, with a 20 cm target (Figure 3). These protocols have different durations and cohort sizes; their displacement statistics should not be interchanged.

4.2 Action imitation, balance, and gait differ

Superfly achieves the strongest travel among the evaluated graph policies, while the MLP is best overall. Relative to the GNN, the validation-MSE gap (0.0686 versus 0.0957) becomes a 21.1× five-second travel gap. The GNN completes 96% of its travel in the first second, then holds a support posture. It generates alternating leg commands on teacher states and with constant sensory inputs; failure arises when its own actions change posture and contact feedback. Thus, an offline stepping rhythm does not ensure corrective behavior along the policy's own trajectories.

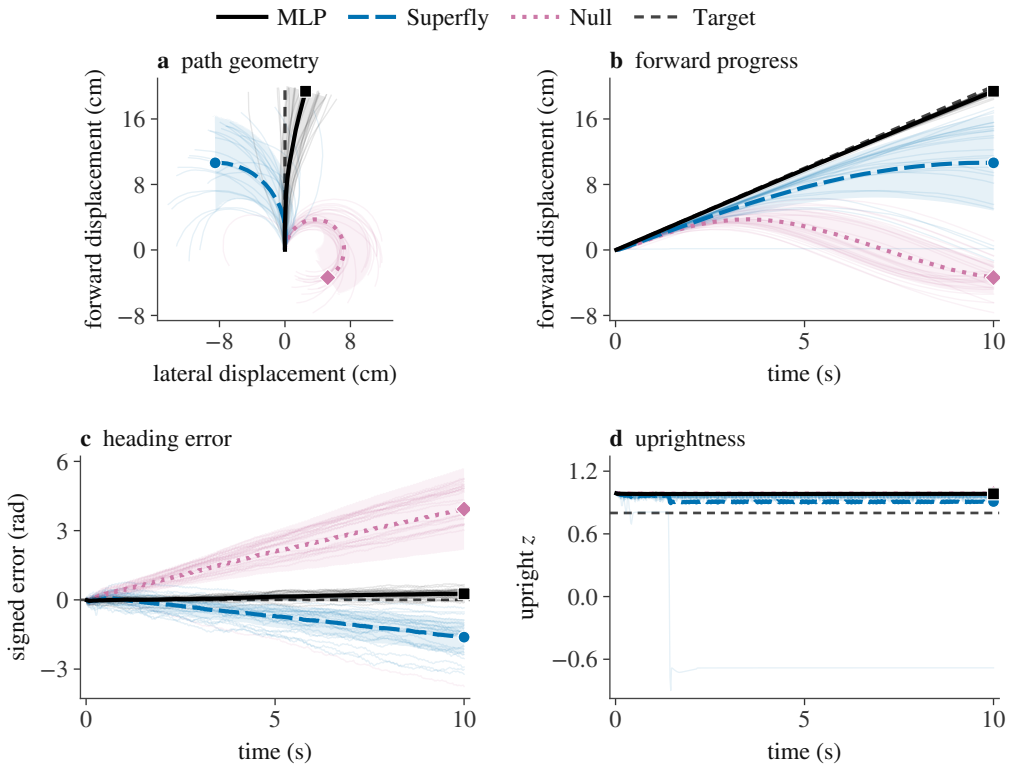

Figure 3: **Ten-second closed-loop control from 25 held-out postures.** All policies receive 2.0 cm/s forward commands toward a 20 cm target. Panels show aligned paths, forward displacement, signed heading error, and uprightness. Thin lines are individual rollouts, thick lines are means, and bands show one sample standard deviation. The MLP advances 19.4 cm, Superfly 10.6 cm, and the Null ends 3.4 cm behind the start. The omitted GNN advances 0.269 ± 0.126 cm and stays upright in $21/25$ rollouts. Gray references mark target motion, zero heading error, and the 0.8 uprightness threshold.

The Null reverses its favorable offline ranking in closed loop. Its yaw bias averages 0.8 mrad per 2 ms step with the same sign in 24 of 25 ten-second rollouts. Across 5,000 steps this accumulates to about 3.9 rad. It stays upright in all 25 rollouts but ends 3.37 cm behind the start, whereas Superfly advances 10.64 cm. Superfly also veers, in the opposite direction, with a $2.4\times$ smaller bias (about 0.3 mrad per step). These results isolate a directional failure within the tested matched pair, not a universal advantage over all rewires.

Balance, propulsion, turn selection, and path regulation are distinct. Superfly sustains gait but ends about 8.5 cm left of the straight-command path. It follows right-turn commands in $5/5$ rollouts and left-turn commands in $3/5$, with endpoint errors of 4.6 and 7.6 cm, respectively. Behavioral cloning establishes locomotion but leaves systematic steering errors. Closed-loop corrective training is a natural next step. Training curves and additional interpretation appear in Appendix C.

4.3 Scalar dynamics reduce graph computation

With E edges, N neurons, and feature width $d \approx 50$, generic GNN updates cost $O(Ed + Nd^2)$, whereas scalar Superfly updates cost $O(E+N)$. Counted work falls from 1.839 to 0.0178 GFLOPs per action, a $103\times$ reduction. On the same GPU, median latency falls from 3.558 to 1.868 ms and reported board energy from 1,260 to 493 mJ. Superfly's median is below the 2 ms control period;

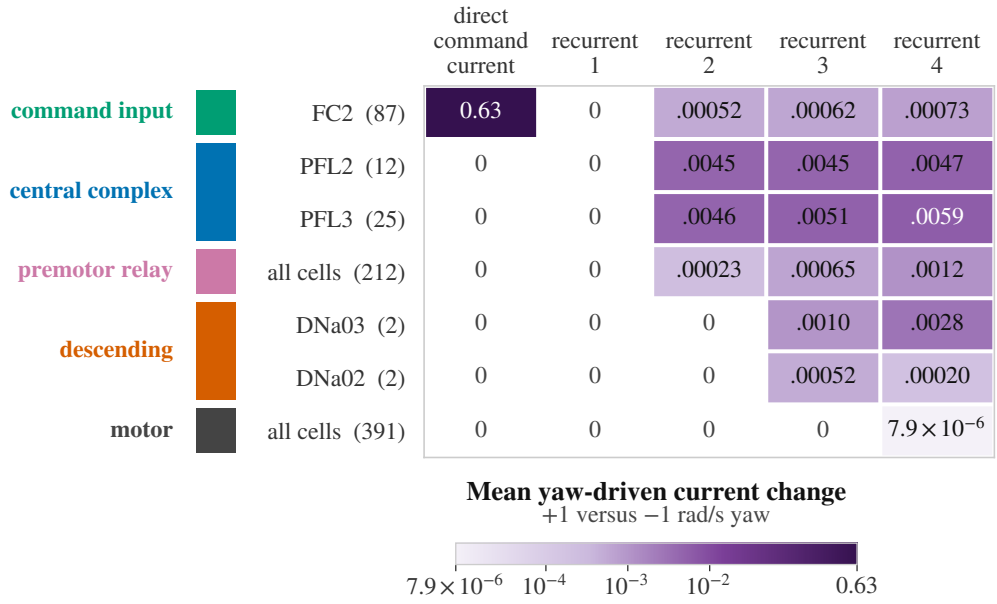

Figure 4: **A fixed-state yaw intervention produces staged current responses.** Each cell is the mean absolute contrast $|I_{+1} - I_{-1}|$ across five validation states and neurons in the named row. Columns show direct command current and four recurrent-current calculations. Yaw first affects PFL2/PFL3 in recurrent calculation 2, descending cells in calculation 3, and motor cells in calculation 4. Color shows magnitude on a logarithmic scale; white denotes zero. These are model responses without neural activity supervision.

tail latency and deadline-miss rates are not reported. The MLP remains fastest, with cost scaling with its 37,296 parameters instead of millions of graph edges.

The advantage is therefore specific to the compared graph representations. We do not establish that Superfly is cheaper than the MLP or that it transfers across tasks. Its distinct capability is linking computation and possible interventions to identified neurons. Appendix E reports the inference and energy comparison.

4.4 Yaw responses are traceable to known steering neurons

Prior neurobiology identifies PFL3 with left-right steering, PFL2 with steering gain, and DNa03 upstream of DNa02 (Westeinde et al., 2024; Mussells Pires et al., 2024). Bilateral DNa02 activity predicts turning velocity (Rayshubskiy et al., 2025; Yang et al., 2024). We use these findings to select the populations and sequence to inspect. Training supplies motor targets alone: analyzed voltages, currents, and release are model predictions, not neural recordings.

For each of five fixed validation states, we branch from identical body inputs and neural voltages with yaw commands $+1$ and -1 rad/s. The command directly enters only 87 FC2 cells (0.38% of the full graph). Figure 4 reports mean absolute current contrasts across states and cells. Recurrent current is initially identical; contrast appears in PFL2/PFL3 and premotor cells in recurrent calculation 2, DNa03/DNa02 in calculation 3, and motor cells in calculation 4. Fixed interfaces and connectivity constrain the stage order, while learned dynamics determine response magnitude. The update sequence is not a measurement of biological timing.

At the final update, PFL2 and PFL3 have $11.7\times$ and $14.5\times$ the average CX response; DNa03 and DNa02 have $131\times$ and $9.05\times$ the average descending response. Thus, the control signal can be traced through named populations associated with steering in living flies. However, the same four families remain turn responsive after degree-preserving rewiring. Selectivity alone therefore does not establish a causal requirement for exact native partners. Neural recordings and cell-specific

perturbations are needed to test the predicted computation. Appendix D maps the modeled cell types to biological roles and interfaces.

5 Discussion and conclusion

Superfly makes the BANC sensorimotor connectome executable as a recurrent controller while retaining an explicit identity for every modeled neuron. The experiments reveal a tradeoff: the dense MLP travels farthest and runs fastest; Superfly combines sustained gait with neuron-resolved analysis and lower cost than the tested GNN; the Null predicts actions well but accumulates greater directional error. One-step MSE, balance, and accurate locomotion should consequently be assessed separately.

The evidence is limited to imitation of walking, one parameter-matched GNN configuration, one rewired graph, and the reported evaluation cohorts. The model omits spikes and dendritic structure, and its cell responses have not been validated against neural recordings. Anatomical command placement contributes to the observed activity sequence. Neither broad transfer nor a biological energy advantage is demonstrated. A reusable embodied backbone remains a future direction: fixed circuitry could support new task-specific interfaces while preserving experimentally testable internal states. Closed-loop learning and targeted neural interventions are the next tests of that possibility.

The longer-term direction is a reusable biological architecture for embodied intelligence. A connectome fixes which neurons communicate; learning calibrates how observations enter that circuit and how its activity drives a body. New interfaces could adapt the same scaffold to other sensory conditions or tasks, while persistent neuron identities keep its predictions open to experimental tests. FlyGM provides precedent for broader behavioral scope, but the present results establish only the walking setting studied here. Extending Superfly across tasks, bodies, and environments remains a proposal rather than an observed result.

AI use statement

AI tools assisted with initial code and experiment-pipeline setup, manuscript drafting and editing, literature checks, and feedback on methodology and interpretation.

Reproducibility statement

Section 3 specifies the policy, objective, and training setup. Appendix B provides observation normalization, anatomical input routes, synaptic parameterization, recurrent updates, and deployment details. Section 4 distinguishes evaluation cohorts and matched baselines; Appendices C and E provide training curves and the inference-cost comparison.

REFERENCES

Florian Aymanns et al. Descending neuron population dynamics during odor-evoked and spontaneous limb-dependent behaviors. *eLife*, 11:e81527, 2022. doi: 10.7554/eLife.81527.

Anthony Azevedo, Ellen Lesser, Jasper S Phelps, et al. Connectomic reconstruction of a female Drosophila ventral nerve cord. *Nature*, 631:360–368, 2024. doi: 10.1038/s41586-024-07389-x.

Aishwarya Balwani, Alex Q Wang, Farzaneh Najafi, and Hannah Choi. Constructing biologically constrained RNNs via Dale's backpropagation and topologically informed pruning. *Science Advances*, 11(50):eadw4970, 2025. doi: 10.1126/sciadv.adw4970.

Alexander S. Bates, Jasper S. Phelps, Minsu Kim, Helen H. Yang, et al. Distributed control circuits across a brain-and-cord connectome. *Nature*, 2026. doi: 10.1038/s41586-026-10735-w.

Salil S Bidaye, Till Bockemühl, and Ansgar Büschges. Six-legged walking in insects: how CPGs, peripheral feedback, and descending signals generate coordinated and adaptive motor rhythms. *Journal of Neurophysiology*, 119(2):459–475, 2018. doi: 10.1152/jn.00658.2017.

Kwabena Boahen. Dendrocentric learning for synthetic intelligence. *Nature*, 612:43–50, 2022. doi: 10.1038/s41586-022-05340-6.

Jonas Braun, Femke Hurtak, Sibo Wang-Chen, and Pavan Ramdya. Descending networks transform command signals into population motor control. *Nature*, 630:686–694, 2024. doi: 10.1038/s41586-024-07523-9.

Han S J Cheong, Katharina Eichler, Tomke Stürner, et al. Organization of circuits linking descending input to motor output in the Drosophila male adult nerve cord connectome. *eLife*, 13:RP96084, 2026. doi: 10.7554/eLife.96084.

Jonathan Cornford, Damjan Kalajdzievski, Marco Leite, Amélie Lamarquette, Dimitri M Kullmann, and Blake A Richards. Learning to live with Dale's principle: ANNs with separate excitatory and inhibitory units. In *ICLR*, 2021. URL https://openreview.net/forum?id=eU776ZYxEpz.

Mike Davies et al. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1):82–99, 2018. doi: 10.1109/MM.2018.112130359.

Nalin Dhiman. Topological sensitivity in connectome-constrained neural networks, 2026. arXiv:2604.04033.

Sven Dorkenwald, Arie Matsliah, Amy R Sterling, et al. Neuronal wiring diagram of an adult brain. *Nature*, 634:124–138, 2024. doi: 10.1038/s41586-024-07558-y.

Nils Eckstein, Alexander S Bates, Andrew Champion, et al. Neurotransmitter classification from electron microscopy images at synaptic sites in Drosophila melanogaster. *Cell*, 187(10):2574–2594.e23, 2024. doi: 10.1016/j.cell.2024.03.016.

Jonathan Green, Atsuko Adachi, Kunal K Shah, Jonathan D Hirokawa, Pablo S Magani, and Gaby Maimon. A neural circuit architecture for angular integration in Drosophila. *Nature*, 546:101–106, 2017. doi: 10.1038/nature22343.

Agrim Gupta, Linxi Fan, Surya Ganguli, and Li Fei-Fei. MetaMorph: Learning universal controllers with transformers. In *ICLR*, 2022.

Scott Harris, Aarav Sinha, Viktor Toth, Alexis Pomares, and Philip Shiu. How the Eon team produced a virtual embodied fly. Eon Systems, https://eon.systems/updates/embodied-brain-emulation, 2026. Blog post, March 10, 2026; accessed August 15, 2026. No peer-reviewed report at the time of writing.

Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research*, 21(248):1–43, 2020.

Anna Honkanen, Andrea Adden, Josiane da Silva Freitas, and Stanley Heinze. The insect central complex and the neural basis of navigational strategies. *Journal of Experimental Biology*, 222 (Suppl 1):jeb188854, 2019. doi: 10.1242/jeb.188854.

Wenlong Huang, Igor Mordatch, and Deepak Pathak. One policy to control them all: Shared modular policies for agent-agnostic control. In *ICML*, volume 119 of *PMLR*, pp. 4455–4464, 2020. URL https://proceedings.mlr.press/v119/huang20d.html.

Brad K Hulse, Hannah Haberkern, Romain Franconville, et al. A connectome of the Drosophila central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *eLife*, 10:e66039, 2021. doi: 10.7554/eLife.66039.

Zehao Jin and Yanan Sui. Whole-brain connectomic graph neural networks enable whole-body locomotion control in Drosophila. In *Data on the Brain and Mind, a workshop at NeurIPS 2025*, 2025. URL https://neurips.cc/virtual/2025/131402.

Zehao Jin, Yaoye Zhu, Chen Zhang, and Yanan Sui. Whole-brain connectomic graph model enables whole-body locomotion control in fruit fly, 2026. arXiv:2602.17997.

Reece Keller, Alyn Kirsch, Felix Pei, Xaq Pitkow, Leo Kozachkov, and Aran Nayebi. Intrinsic goals for autonomous agents: Model-based exploration in virtual zebrafish predicts ethological behavior and whole-brain dynamics. In *Advances in Neural Information Processing Systems 38 (NeurIPS)*, pp. 40321–40352, 2025. doi: 10.52202/085713-1205.

Anmo J. Kim, Jamie K. Fitzgerald, and Gaby Maimon. Cellular evidence for efference copy in Drosophila visuomotor processing. *Nature Neuroscience*, 18(9):1247–1255, 2015. doi: 10.1038/nn.4083.

Sung Soo Kim, Hervé Rouault, Shaul Druckmann, and Vivek Jayaraman. Ring attractor dynamics in the Drosophila central brain. *Science*, 356(6340):849–853, 2017. doi: 10.1126/science.aal4835.

Dhireesha Kudithipudi, Catherine Schuman, Craig M. Vineyard, et al. Neuromorphic computing at scale. *Nature*, 637:801–812, 2025. doi: 10.1038/s41586-024-08253-8.

Janne K Lappalainen, Fabian D Tschopp, Sridhama Prakhya, et al. Connectome-constrained networks predict neural activity across the fly visual system. *Nature*, 634:1132–1140, 2024. doi: 10.1038/s41586-024-07939-3.

Ellen Lesser, Anthony W Azevedo, Jasper S Phelps, et al. Synaptic architecture of leg and wing premotor control networks in Drosophila. *Nature*, 631:369–377, 2024. doi: 10.1038/s41586-024-07600-z.

Victor Lobato-Rios et al. NeuroMechFly, a neuromechanical model of adult Drosophila melanogaster. *Nature Methods*, 19(5):620–627, 2022. doi: 10.1038/s41592-022-01466-7.

Alexandra Sasha Luccioni, Yacine Jernite, and Emma Strubell. Power hungry processing: Watts driving the cost of AI deployment? In *ACM Conference on Fairness, Accountability, and Transparency*, pp. 85–99, 2024. doi: 10.1145/3630106.3658542.

Cheng Lyu, L F Abbott, and Gaby Maimon. Building an allocentric travelling direction signal via vector computation. *Nature*, 601:92–97, 2022. doi: 10.1038/s41586-021-04067-0.

Sergei Maslov and Kim Sneppen. Specificity and stability in topology of protein networks. *Science*, 296:910–913, 2002. doi: 10.1126/science.1065103.

Josh Merel, Leonard Hasenclever, Alexandre Galashov, et al. Neural probabilistic motor primitives for humanoid control. In *ICLR*, 2019.

Peter Mussells Pires et al. Converting an allocentric goal into an egocentric steering signal. *Nature*, 626:808–818, 2024. doi: 10.1038/s41586-023-07006-3.

Shigehiro Namiki, Michael H Dickinson, Allan M Wong, et al. The functional organization of descending sensory-motor pathways in Drosophila. *eLife*, 7:e34272, 2018. doi: 10.7554/eLife.34272.

Jeremy E Niven and Simon B Laughlin. Energy limitation as a selective pressure on the evolution of sensory systems. *Journal of Experimental Biology*, 211:1792–1804, 2008. doi: 10.1242/jeb.017574.

Kaushik Panda, Rohith Mittapally, Qian Chen, Abhinandan M. Bhaskaran, Pramod Reddy, Edgar Meyhofer, and Swathi Yadlapalli. Direct quantification of the metabolic heat output of individual Drosophila brains. *Cell Reports Methods*, 2026. doi: 10.1016/j.crmeth.2026.101501.

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 37(4), 2018. doi: 10.1145/3197517.3201311.

Dean A Pospisil, Max J Aragon, Sven Dorkenwald, et al. The fly connectome reveals a path to the effectome. *Nature*, 634:201–209, 2024. doi: 10.1038/s41586-024-07982-0.

Sarah M Pugliese, Grant M Chou, Elliott T T Abe, et al. Connectome simulations identify a central pattern generator circuit for fly walking. *bioRxiv*, 2025. doi: 10.1101/2025.09.12.675944.

Naveen Rao, MeeLan Lee, Michael Carbin, and Sara Achour. Introducing unconventional AI. Unconventional AI, https://unconv.ai/blog/introducing-unconventional-ai/, 2025. December 8, 2025.

Aleksandr Rayshubskiy, Stephen L Holtz, Alexander S Bates, Quinn X Vanderbeck, Laia Serratosa Capdevila, Victoria Rockwell, and Rachel I Wilson. Neural circuit mechanisms for steering control in walking Drosophila. *eLife*, 13:RP102230, 2025. doi: 10.7554/eLife.102230.

Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575:607–617, 2019. doi: 10.1038/s41586-019-1677-2.

Philipp Schlegel, Yijie Yin, Alexander S. Bates, Sven Dorkenwald, Katharina Eichler, Paul Brooks, et al. Whole-brain annotation and multi-connectome cell typing of Drosophila. *Nature*, 634:139–152, 2024. doi: 10.1038/s41586-024-07686-5.

Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. Green AI. *Communications of the ACM*, 63(12):54–63, 2020. doi: 10.1145/3381831.

Johannes D Seelig and Vivek Jayaraman. Neural dynamics for landmark orientation and angular path integration. *Nature*, 521:186–191, 2015. doi: 10.1038/nature14446.

Philip K Shiu, Gabriella R Sterne, Nico Spiller, et al. A Drosophila computational brain model reveals sensorimotor processing. *Nature*, 634:210–219, 2024. doi: 10.1038/s41586-024-07763-9.

Sergey Shuvaev et al. Encoding innate ability through a genomic bottleneck. *Proceedings of the National Academy of Sciences*, 121(38):e2409160121, 2024. doi: 10.1073/pnas.2409160121.

H Francis Song, Guangyu R Yang, and Xiao-Jing Wang. Training excitatory-inhibitory recurrent neural networks for cognitive tasks: A simple and flexible framework. *PLoS Computational Biology*, 12(2):e1004792, 2016. doi: 10.1371/journal.pcbi.1004792.

Stavros Therianos. A frozen rate operator from the complete larval connectome: degree and weight govern the gross response, exact wiring governs input routing and mushroom-body modes, 2026. arXiv:2606.17745.

Daniel Turner-Evans, Stephanie Wegener, Hervé Rouault, et al. Angular velocity integration in a fly heading circuit. *eLife*, 6:e23496, 2017. doi: 10.7554/eLife.23496.

Daniel B Turner-Evans and Vivek Jayaraman. The insect central complex. *Current Biology*, 26(11):R453–R457, 2016. doi: 10.1016/j.cub.2016.04.006.

Daniel B Turner-Evans, Kristopher T Jensen, Saba Ali, Tyler Paterson, et al. The neuroanatomical ultrastructure and function of a biological ring attractor. *Neuron*, 108(1):145–163.e10, 2020. doi: 10.1016/j.neuron.2020.08.006.

František Váša and Bratislav Mišić. Null models in network neuroscience. *Nature Reviews Neuroscience*, 23:493–504, 2022. doi: 10.1038/s41583-022-00601-9.

Roman Vaxenburg, Igor Siwanowicz, Josh Merel, Alice A Robie, Carmen Morrow, et al. Whole-body physics simulation of fruit fly locomotion. *Nature*, 643:1312–1320, 2025. doi: 10.1038/s41586-025-09029-4.

Tingwu Wang, Renjie Liao, Jimmy Ba, and Sanja Fidler. NerveNet: Learning structured policy with graph neural networks. In *ICLR*, 2018.

Sibo Wang-Chen and Pavan Ramdya. The embodied brain: Bridging the brain, body, and behavior with biorealistic neuromechanical models. *Current Opinion in Neurobiology*, 100:103244, 2026. doi: 10.1016/j.conb.2026.103244.

Sibo Wang-Chen, Victor Alfred Stimpfling, Thomas Ka Chung Lam, Pembe Gizem Özdil, Louise Genoud, Femke Hurtak, and Pavan Ramdya. NeuroMechFly v2: simulating embodied sensorimotor control in adult Drosophila. *Nature Methods*, 21(12):2353–2362, 2024. doi: 10.1038/s41592-024-02497-y.

Elena A Westeinde et al. Transforming a head direction signal into a goal-oriented steering command. *Nature*, 626:819–826, 2024. doi: 10.1038/s41586-024-07039-2.

Helen H Yang, Bella E Brezovec, Laia Serratosa Capdevila, Quinn X Vanderbeck, et al. Fine-grained descending control of steering in walking Drosophila. *Cell*, 187(22):6290–6308.e27, 2024. doi: 10.1016/j.cell.2024.08.033.

Anthony M Zador. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10:3770, 2019. doi: 10.1038/s41467-019-11786-6.

Mengdi Zhao, Ning Wang, Xinrui Jiang, et al. An integrative data-driven model simulating C. elegans brain, body and environment interactions. *Nature Computational Science*, 4:978–990, 2024. doi: 10.1038/s43588-024-00738-w.

Zhihao Zheng, J Scott Lauritzen, Eric Perlman, Camenzind G Robinson, et al. A complete electron microscopy volume of the brain of adult Drosophila melanogaster. *Cell*, 174(3):730–743.e22, 2018. doi: 10.1016/j.cell.2018.06.019.

A Extended background and related work

Connectomes expose sensorimotor wiring. A connectome represents each neuron as a node and each directed connection as an edge labeled by synapse count. FlyWire reconstructs 139,255 neurons and about 54.5 million synapses in the adult fly brain (Zheng et al., 2018; Dorkenwald et al., 2024). Its cell-type annotations retain the biological identity of each node (Schlegel et al., 2024). BANC provides a unified reconstruction of the adult brain and ventral nerve cord, including the pathways that connect central computation to local leg control (Bates et al., 2026). Cell annotations and transmitter predictions supply the identities, synapse counts, and signs used by Superfly. BANC therefore provides the anatomical boundary required for an embodied locomotion policy.

The fly's navigation-to-locomotion circuit. The Central Complex (CX) combines heading, travel direction, and goals into steering commands (Honkanen et al., 2019; Turner-Evans & Jayaraman, 2016). EPG cells form an internal compass: the position of one activity bump represents the fly's heading (Seelig & Jayaraman, 2015; Kim et al., 2017). PEN cells move that bump as the fly turns (Turner-Evans et al., 2017; Green et al., 2017). FC2 cells represent goal direction upstream of steering, while fan-shaped-body populations represent the fly's goal and travel direction (Mussells Pires et al., 2024; Hulse et al., 2021). PFL3 selects left or right from the heading-goal comparison, and PFL2 adjusts turn strength (Westeinde et al., 2024). Their output reaches DNa01 and DNa02, which carry steering signals toward the ventral nerve cord (Rayshubskiy et al., 2025). Around 1,300 descending neurons carry brain output as coordinated populations rather than isolated channels (Namiki et al., 2018; Braun et al., 2024; Aymanns et al., 2022). Premotor circuits, reflexes, and rhythm generators in the nerve cord transform that output into coordinated walking (Azevedo et al., 2024; Lesser et al., 2024; Cheong et al., 2026; Pugliese et al., 2025; Bidaye et al., 2018). The named roles motivate Superfly's FC2 command interface and its PFL2, PFL3, and DNa02 response analysis.

Graph-structured policies in robot learning. NerveNet represents a robot's morphology as a graph, assigning one node to each joint and one edge to each physical link (Wang et al., 2018). Shared modular policies (Huang et al., 2020) and morphology-conditioned transformers (Gupta et al., 2022) extend this principle across different bodies. Motion-capture imitation provides training targets for such controllers (Peng et al., 2018; Merel et al., 2019). In these approaches, the robot's designed morphology determines the policy graph. Superfly retains the imitation setting but changes the source of structure: the animal's nervous system supplies the graph. The question therefore shifts from how to encode a robot's body to whether inherited neural wiring provides a useful control prior.

Biologically grounded models of brains and bodies. NeuroMechFly (Lobato-Rios et al., 2022) and NeuroMechFly v2 (Wang-Chen et al., 2024) reconstruct the adult fly body in MuJoCo, including vision, olfaction, and leg adhesion. The flybody simulator adds a motion-capture walking task and releases the trained expert that Superfly clones (Vaxenburg et al., 2025). Most simulated fly bodies use artificial controllers whose connections do not follow the fly's reconstructed nervous system (Wang-Chen & Ramdya, 2026). A concurrent Eon Systems demonstration similarly connects a numerical model of the full central brain to NeuroMechFly v2 (Harris et al., 2026). Its leaky integrate-and-fire brain remains fixed, while descending activity steers pretrained imitation controllers. Superfly follows the same embodied neural-modeling direction but uses a task-optimized BANC model spanning the brain and nerve cord as the locomotion policy itself. Beyond the fly, BAAIWorm couples a connectome-constrained *C. elegans* nervous system to a body and reproduces its locomotion (Zhao et al., 2024). Virtual zebrafish agents trained with an intrinsic objective predict whole-brain dynamics without neural training data (Keller et al., 2025). These projects primarily ask whether a neural model reproduces animal behavior or brain dynamics.

Connectome-constrained neural models. Lappalainen et al. freeze measured visual-system connectivity and optimize unknown cellular and synaptic quantities for optic-flow estimation (Lappalainen et al., 2024). Their models receive no target neural activity during task training, yet reproduce known responses of identified visual cell types. Shiu et al. simulate the complete adult fly brain with a leaky integrate-and-fire model and recover sensorimotor mappings without fitted parameters (Shiu et al., 2024). Pospisil et al. treat the signed synapse-count matrix as a linear dynamical operator (Pospisil et al., 2024). Concurrently, Therianos studies a frozen rate operator on

the larval connectome (Therianos, 2026). Its global response follows degree and weight statistics, while exact edges determine where inputs travel. Sign-constrained networks instead learn excitatory and inhibitory roles from data (Song et al., 2016; Cornford et al., 2021; Balwani et al., 2025).

Superfly adopts the task-optimized mechanistic strategy for embodied locomotion. It keeps every BANC edge endpoint, synapse count, and transmitter sign fixed. Anatomical interfaces route body signals into identified neurons, while training fits group-shared synaptic scales, cell-type physiology, and linear motor readouts. Persistent scalar voltage keeps every modeled response attached to a named cell throughout a rollout. The design turns behavioral training into cell-indexed predictions for PFL2, PFL3, DNa02, and wider motor populations.

From inherited wiring to a testable prior. Zador (2019) argues that evolution stores behavioral knowledge in neural wiring, reducing what each animal must learn during its lifetime. Shuvaev et al. (2024) show that an analogous genomic bottleneck regularizes artificial networks. Energy provides a second motivation: sparse biological circuits communicate through synaptic events, and neuromorphic hardware executes such events without dense matrix multiplication (Niven & Laughlin, 2008; Roy et al., 2019; Davies et al., 2018). Related efforts pursue physical computing substrates intended to approach biological energy efficiency (Rao et al., 2025). These arguments motivate wiring as a prior, but they do not establish that a particular connectome topology matters.

Testing topology requires matched nulls. Degree-preserving edge swaps change synaptic partners while retaining each neuron's incoming and outgoing edge counts (Maslov & Sneppen, 2002). Váša & Mišić (2022) formalize which network properties a null must preserve before researchers attribute an effect to topology. Dhiman (2026) apply this discipline to a connectome-constrained visual model and find that its topology advantage largely disappears under shared initialization and a degree-preserving rewire.

Superfly controls trainable capacity and synaptic partners in two matched comparisons. A parameter-matched MLP tests the complete structured policy against a dense controller. The topology Null changes BANC partners within the same neural model. Each accepted swap exchanges two targets inside one population-to-population block. The Null preserves every neuron's in-degree and out-degree, every block count, each source neuron's signed synapse-count values, the neural dynamics, input interfaces, motor readouts, parameter count, training frames, and optimization budget. Only postsynaptic partners change within the tested pair. Because we evaluate one rewired realization, the comparison controls these confounds for that instance but does not estimate variation across possible rewires.

Connectomes as policies. FlyGM provides the direct policy precedent for Superfly. Jin et al. (2026) instantiate the complete adult fly brain as a directed message-passing controller and train it through imitation followed by reinforcement learning on flybody. FlyGM achieves walking, turning, and flight, extending an earlier workshop system (Jin & Sui, 2025). It compares the connectome policy with rewired, random, MLP, and conventional GNN baselines. FlyGM therefore establishes that a whole-brain connectome supports broad embodied control.

FlyGM and Superfly ask different questions about how a fly connectome controls a body. FlyGM tests one whole-brain model across walking, turning, and flight. Superfly instead focuses on walking and follows signals from sensory neurons through the brain and nerve cord to motor neurons. Each Superfly neuron stores one voltage that leaks toward rest while integrating incoming signals; FlyGM uses a learned graph-neural-network update. Superfly routes each body signal only to neurons associated with the matching pathway, whereas FlyGM sends a shared encoded input to its afferent neurons. A dense MLP with exactly the same fitted parameter count tests whether structure helps beyond model size. The strict rewire preserves each neuron's incoming and outgoing connection counts, plus the number, sign, and synapse count of connections between major neural groups. Only the connected neuron pairs change. Together, these controls test whether the reconstructed wiring affects walking, while every simulated response remains associated with a specific biological neuron.

B DETAILED METHOD AND IMPLEMENTATION

SUPERFLY turns a mapped sensorimotor circuit spanning the adult *Drosophila* brain and nerve cord into a recurrent locomotion policy. Measured anatomy fixes which neurons communicate, each connection's synapse count, and whether it is excitatory or inhibitory. Training fits how body signals enter the circuit, how strongly neurons interact, how quickly they respond, and how neural activity produces motor control. A passive point-neuron equation governs how each neuron integrates incoming synaptic current over time. Walking imitation learns synaptic scales, cell-type physiology, and the maps between body and circuit. The method has four parts: walking demonstrations, an embodied connectome, anatomical body interfaces, and recurrent neural dynamics with linear motor readouts.

B.1 WALKING DEMONSTRATIONS

We use the velocity-tracking SAC policy released with flybody (Vaxenburg et al., 2025) as the demonstration teacher. We sample 25 motion-capture postures, yielding 75 rollouts after replaying three commands from each posture. Straight walking requests 2.0 cm/s forward velocity and zero yaw. Left and right turns request 1.75 cm/s forward velocity with yaw rates of $+0.35$ and -0.35 rad/s, respectively. Each rollout lasts 2,500 control steps, or 5 s at 2 ms per step. The 75 rollouts contain 187,500 observation-control pairs. Every posture contributes all three commands from the same initial body state, which separates the requested motion from variation among postures.

At control step t, SUPERFLY predicts the teacher's 59-dimensional control from a 137-dimensional observation:

$$x_t \equiv (s_t, a_{t-1}, c_t) \in \mathbb{R}^{137}, \qquad a_t^\star \in [-1, 1]^{59}. \tag{6}$$

Here, $s_t \in \mathbb{R}^{75}$ contains proprioceptive and inertial body readings. The vector $a_{t-1} \in \mathbb{R}^{59}$ is the control applied to the body at the preceding step. Feeding it back identifies the motor command that immediately preceded the current sensor readings. An efference copy is an internal record of a recently issued motor command. Drosophila visual neurons receive such motor-related signals during flight (Kim et al., 2015). The command $c_t \in \mathbb{R}^3$ specifies the requested forward velocity, lateral velocity, and yaw rate. The target $a_t^\star$ comprises six adhesion commands, five head and abdomen controls, and 48 leg-joint controls. SUPERFLY receives only current body feedback, the preceding control, and the three requested motion components. The teacher also observes a motion-capture reference pose and future target pose; SUPERFLY does not. Equation (6) groups the 137 inputs by their role in control. For anatomical routing, we further organize the 75 body readings by sensory organ, side, and leg. Equation (9) sends each group only to its corresponding sensory neurons. Figure 5 maps the 75 body readings to their sensor locations and shows all 137 inputs in storage order.

We partition complete postures rather than individual frames. Training uses 60 rollouts from 20 postures, totaling 150,000 frames. Validation uses 15 rollouts from five disjoint postures, totaling 37,500 frames. All three commands from one posture remain in one partition. The split tests unseen body configurations without sharing adjacent states across training and validation. For each observation and control dimension j, we standardize inputs and targets as

$$\widetilde{x}_{t,j} = \frac{x_{t,j} - \mu_j^x}{\bar{\sigma}_j^x}, \qquad \widetilde{a}_{t,j}^\star = \frac{a_{t,j}^\star - \mu_j^a}{\bar{\sigma}_j^a}, \qquad \bar{\sigma}_j^z = \max(\sigma_j^z, 10^{-3}), \quad z \in \{x, a\}. \tag{7}$$

We compute every mean and standard deviation in Equation (7) from training frames only. The 10^{-3} floor prevents nearly constant channels from amplifying numerical noise. We then apply those fixed statistics unchanged to both training and validation frames.

B.2 EMBODIED BANC GRAPH

SUPERFLY operates on a sensorimotor graph extracted from the Brain and Nerve Cord connectome (BANC) (Bates et al., 2026). The graph comprises seven populations: peripheral sensory neurons, Central Complex neurons, ascending neurons, descending neurons, leg nerve cord neurons, motor neurons, and premotor relays. The graph contains 22 thousand neurons assigned to 3,430 cell types. A chemical synapse is one reconstructed contact. A directed edge groups every such contact from

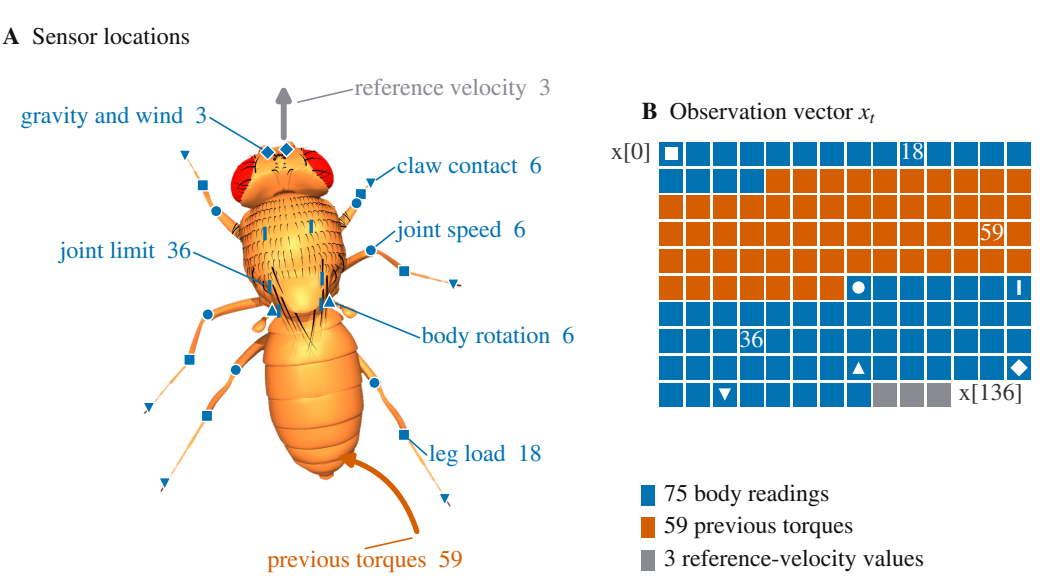

Figure 5: **SUPERFLY's 137 inputs combine body feedback, previous controls, and motion commands.** (A) A top view of flybody locates the six groups that provide 75 body readings. Squares mark leg load, circles joint speed, and ticks joint limits. Upright triangles mark body rotation, diamonds gravity and wind, and inverted triangles claw contact. Blue marks body readings, and orange marks the 59 controls applied at the previous step. Gray marks the three-dimensional command: requested forward velocity, lateral velocity, and yaw rate. (B) The observation in array order, read from left to right and top to bottom. Blue cells contain body readings, orange cells contain previous controls, and gray cells contain the three command values. Body readings occupy $x[0{:}18]$ and $x[77{:}134]$, previous controls occupy $x[18{:}77]$, and command values occupy $x[134{:}137]$. All index ranges are half-open.

one source neuron to one target neuron. The graph's 2.2 million directed edges therefore aggregate 8.5 million individual chemical synapses. Unlike a brain-only graph, this sensorimotor graph extends from peripheral sensory input through central brain circuits and the ventral nerve cord to motor neurons. Figure 6 draws the seven populations and the edge rule.

Each directed edge $u \to v$ aggregates all chemical synapses from source neuron u to target neuron v. The count n_{uv} records its anatomical magnitude. The predicted neurotransmitter of u fixes $s_u \in \{-1, +1\}$, where negative and positive values denote inhibitory and excitatory output. The canonical BANC edge table excludes autapses, which are connections from a neuron back onto itself. We follow that convention and retain every other edge among the seven populations. Training preserves every remaining edge's source, target, count, and sign.

Synapse count provides a relative measure of anatomical strength: a ten-synapse edge contains ten reconstructed contacts between two neurons, whereas a one-synapse edge contains one. The connectome does not specify how much current one contact produces. SUPERFLY converts each fixed signed count into an effective coupling:

$$W_{uv} = s_u n_{uv} g_{\kappa(u),\rho(v)}, \qquad g_{\kappa,\rho} \geq 0. \tag{8}$$

SUPERFLY does not fit a separate strength for every edge. Instead, it groups edges by two fixed biological labels: the cell type of the sending neuron $\kappa(u)$ and the population of the receiving neuron $\rho(v)$. All edges in one group share a learned scale $g_{\kappa,\rho}$, which converts one reconstructed synapse into modeled current. For example, every edge from a given cell type into descending neurons uses the same scale. The sharing rule follows the cell-type parameterization of Lappalainen et al. (2024), whose visual-system model uses one unitary strength for each source-target cell-type pair. SUPERFLY retains exact source types but groups receivers into seven populations to limit the fitted parameter count. Because $g_{\kappa,\rho} \geq 0$, optimization changes a group's magnitude, including to zero,

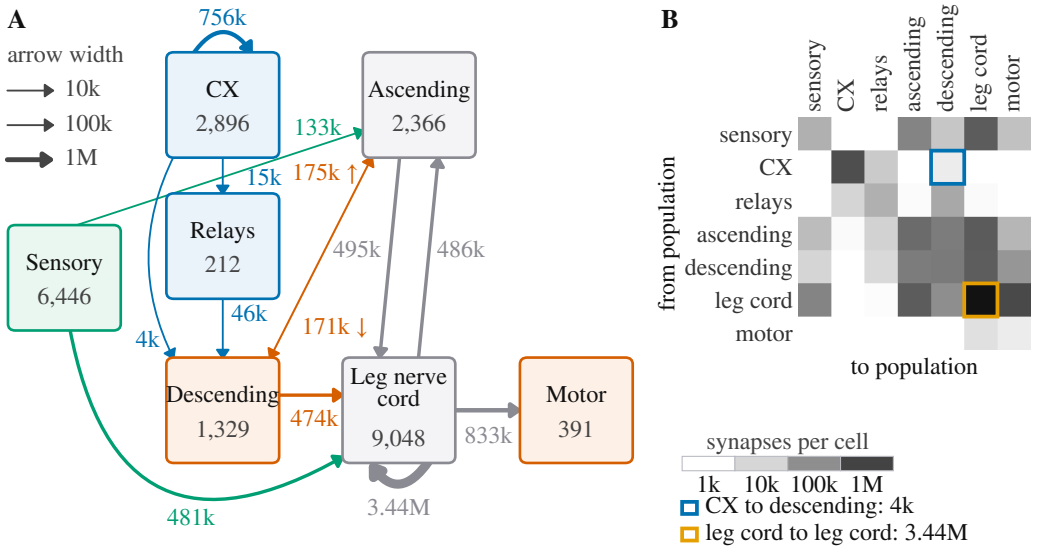

Figure 6: **Embodied BANC graph.** (A) Boxes name the seven populations with their neuron counts. Arrows carry synapse counts, and arrow width grows with that count. The drawing keeps every flow above 100,000 synapses plus the two Central Complex paths into descending neurons. (B) The full matrix of synapses from each row population to each column population on a logarithmic gray scale. Blue outlines the 3,620 CX-to-descending synapses, and gold outlines the 3,442,245 recurrent synapses within the leg nerve cord. Cells below 1,000 synapses stay white. Equation 8 turns each edge's fixed sign and count into a coupling through one learned scale per source-type and target-population group, 14,235 scales for 2,222,379 edges.

without reversing its fixed excitatory or inhibitory sign. The grouping fits about 14 thousand scales rather than 2.2 million separate edge weights. Within a group, an edge containing five synapses retains five times the coupling magnitude of an edge containing one.

B.3 Anatomical body interfaces

The connectome identifies the neurons associated with each sensory organ, while flybody supplies numerical sensor readings. Neither resource specifies how strongly a reading should drive those neurons. SUPERFLY defines the missing interface through 28 anatomical routes, each tied to an organ, body side, and leg. Splitting each route by recipient cell type produces 462 channels that reach 6,858 neurons: 6,446 sensory neurons and 412 ascending neurons. For channel q, the fixed index set $\mathcal{S}_q$ selects the corresponding entries of observation x_t. The fixed set $\mathcal{V}_q$ identifies the recipient neurons. BANC supplies the organ and body-part annotations that define $\mathcal{V}_q$ (Bates et al., 2026); SUPERFLY introduces the map from simulator values to neural current. A learned vector p_q and bias b_q convert $x_t[\mathcal{S}_q]$ into one scalar current ι_t^q. Every neuron in $\mathcal{V}_q$ receives that current. When multiple channels target the same neuron, their currents add, as in passive point-neuron connectome models (Lappalainen et al., 2024):

$$\iota_t^q = p_q^\top x_t[\mathcal{S}_q] + b_q, \qquad I_t^{\text{sens}}[v] = \sum_{q:v\in\mathcal{V}_q} \iota_t^q. \tag{9}$$

Figure 7 makes the routing rule concrete. Consider the front-left joint-speed channel. It selects only the front-left speed value from x_t and converts that value into one scalar current. The channel delivers the current only to annotated front-left femoral chordotonal neurons, which carry leg-joint motion signals. It neither reads another leg nor targets another body segment. Training fits (p_q, b_q), which set the current's magnitude. The selected observation index and recipient neurons remain fixed.

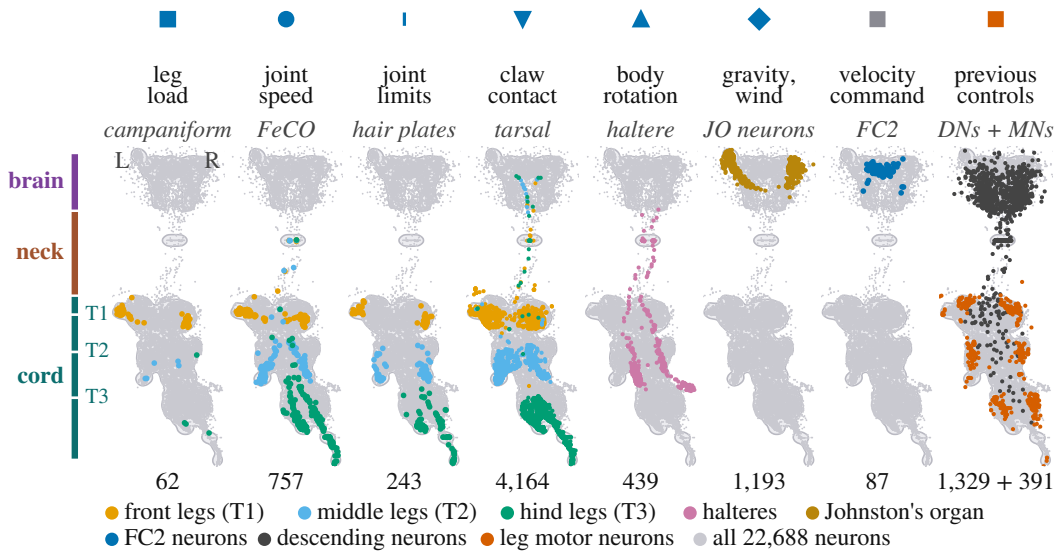

Figure 7: **Anatomical entry points for SUPERFLY's 137-value observation.** Each panel overlays one input group on the same dorsal map of all 22,688 Brain and Nerve Cord connectome neurons, shown in gray. The head points upward, and the fly's left is at left. Purple, brown, and teal bars mark the brain, neck, and leg nerve cord, with T1 through T3 ticks on the three leg segments. Italic labels name the receiving neuron class: campaniform sensilla, femoral chordotonal organ (FeCO), hair plates, tarsal mechanosensors, haltere afferents, Johnston's organ (JO) neurons, FC2 neurons, and descending plus motor neurons (DNs + MNs). Blue symbols reuse Figure 5's sensor key: square, leg load; circle, joint speed; tick, joint limits; inverted triangle, claw contact; upright triangle, body rotation; diamond, gravity and wind. Their observation blocks are $x[0{:}18]$, $x[77{:}83]$, $x[83{:}119]$, $x[128{:}134]$, $x[119{:}125]$, and $x[125{:}128]$, respectively. The gray square marks the velocity command $x[134{:}137]$, while the orange square marks previous controls $x[18{:}77]$. Panel order groups anatomical roles; Figure 5 gives vector storage order. Recipient colors preserve leg identity: orange front, sky blue middle, and green hind. Pink marks halteres, ochre Johnston's organ neurons, blue FC2 neurons, dark gray descending neurons, and vermilion leg motor neurons. Counts below panels report recipients, including 87 command cells and 1,329 descending plus 391 motor recipients for previous controls. Every leg-specific route preserves side and enters its matching T1, T2, or T3 hemisegment.

The velocity command and previous controls enter the circuit through separate interfaces. The gait policy uses a bilateral command interface. A cell-type-tied map sends forward drive equally to left and right FC2 neurons. The map reverses lateral and yaw drive between the two sides. It reaches 87 FC2A, FC2B, and FC2C neurons. The placement introduces requested motion upstream of known steering pathways (Mussells Pires et al., 2024; Westeinde et al., 2024). The efference-copy interface routes the 59 controls issued at the preceding body step nearer to motor output. One cell-type-tied map sends the five head and abdomen controls to 1,329 descending neurons. Six leg-specific maps send each leg's eight joint controls and one adhesion control to matching motor neurons, 391 in total. The five body controls and six groups of nine leg controls account for all 59 values. Desired motion therefore enters the navigation circuit, while recent motor output re-enters downstream near the body parts it affected.

At neuron v, the routed sensory, command, and efference-copy currents add:

$$I_t[v] = I_t^{\text{sens}}[v] + I_t^{\text{cmd}}[v] + I_t^{\text{eff}}[v]. \tag{10}$$

Equation (10) defines how the observation enters the circuit state. During each neural update, $I_t[v]$ enters Equation (13) as an additive external current at recipient neuron v. An absent route contributes zero current. The interface therefore changes voltage only at the anatomically selected recipients. Every later effect on another neuron must travel through a BANC edge. Gradient descent fits each

route's linear conversion, while its selected input values and recipient neurons remain fixed. No learned connection sends x_t directly to the motor readout.

B.4 NEURAL DYNAMICS ON FIXED SYNAPSES

SUPERFLY adopts the passive, threshold-linear point-neuron dynamics of Lappalainen et al. (2024) and applies them to the BANC graph. Each neuron v has one scalar membrane voltage $V_{t,k}[v]$, representing one electrical compartment. The model omits spatial structure within the neuron's dendrites and axon. The index t identifies the body step, and k identifies a neural update within that step. At update k, $r_{t,k}[v]$ denotes the nonnegative graded-release signal sent by neuron v. The quantity $m_{t,k}[v]$ denotes the total recurrent current received by neuron v through its incoming chemical synapses:

$$r_{t,k}[v] = \mathrm{ReLU}(V_{t,k}[v]), \tag{11}$$

$$m_{t,k}[v] = \sum_{u \to v} W_{uv} r_{t,k}[u]. \tag{12}$$

The rectifier in Equation (11) is the model's only local nonlinearity. It approximates instantaneous, voltage-gated transmitter release: nonpositive voltage yields no release, whereas positive voltage passes unchanged. On one edge $u \to v$, the product $W_{uv} r_{t,k}[u]$ is the current delivered from source u to target v. Equation (12) adds these edge currents over every presynaptic source. As defined in Equation (8), W_{uv} combines the edge's fixed sign and synapse count with its fitted group scale. A positive W_{uv} adds current, a negative W_{uv} subtracts current, and a missing edge contributes nothing. Each source therefore emits one scalar signal, while BANC connectivity determines how each target combines its inputs. The model captures circuit-level signal flow rather than ion channels, spikes, or dendritic integration within individual cells.

The recurrent current $m_{t,k}[v]$ and external interface current $I_t[v]$ jointly determine the next voltage. A passive leak simultaneously pulls that voltage toward its resting value:

$$V_{t,k+1}[v] = V_{t,k}[v] + \frac{\Delta t}{\tau_{\kappa(v)}} \left[-V_{t,k}[v] + V^0_{\kappa(v)} + m_{t,k}[v] + I_t[v] \right]. \tag{13}$$

The label $\kappa(v)$ identifies neuron v's BANC cell type. The leak term $-V_{t,k}[v] + V^0_{\kappa(v)}$ pulls voltage toward that type's resting value. The recurrent term $m_{t,k}[v]$ carries input from other graph neurons, while $I_t[v]$ carries the interface currents defined in Equation (10). The ratio $\Delta t/\tau_{\kappa(v)}$ sets how much these terms change voltage during one update. A shorter time constant therefore produces a faster response to the same input. Each of the $3{,}430$ cell types shares one fitted resting voltage $V^0_{\kappa(v)}$ and one fitted time constant $\tau_{\kappa(v)}$.

After each gradient update, SUPERFLY clamps $g_{\kappa,\rho}$ to nonnegative values, so training never reverses an edge's assigned sign. It also clamps $\tau_{\kappa(v)}$ at Δt, so the Euler fraction never exceeds one. Each synaptic group starts with mean absolute edge weight 0.01. SUPERFLY samples each cell type's resting voltage from $\mathcal{N}(0, 0.05^2)$, and every time constant starts at 5 ms.

One body step spans 2 ms, divided into four Euler updates of $\Delta t = 0.5$ ms each. The observation and its external current I_t remain fixed across those updates. Neural voltage, release $r_{t,k}$, and recurrent current $m_{t,k}$ evolve within the four-update sequence. Before each update, SUPERFLY recomputes release and recurrent current from the current voltage. The policy reads the fourth updated state, which also initializes the next body step:

$$V_{t+1,0} = V_{t,4}. \tag{14}$$

At the start of each rollout, SUPERFLY sets every neuron to its fitted resting voltage. Equation (14) then carries the final state across successive body steps within that rollout. The reset prevents activity from crossing between unrelated episodes. Persistent voltage lets the recurrent BANC circuit integrate body and command history instead of restarting neural activity at every frame.

B.5 MOTOR DECODING

After the fourth neural update, Equation (11) converts the final voltage $V_{t,4}$ into the release vector $r_{t,4} = \mathrm{ReLU}(V_{t,4}) \in \mathbb{R}^{22{,}688}$. Its entries give the nonnegative release of every connectome neuron.

Seven anatomically restricted linear maps decode selected entries into 59 controls:

$$\widehat{\widetilde{a}}_t^{\text{body}} = D_{\text{body}} r_{t,4}[\mathcal{V}_{\text{DN}}] + b_{\text{body}}, \qquad \widehat{\widetilde{a}}_t^{\ell} = D_{\ell} r_{t,4}[\mathcal{V}_{\ell}^{\text{motor}}] + b_{\ell}. \tag{15}$$

One body decoder maps release from all 1,329 DNs to five head and abdomen controls. Six leg decoders each read only the motor neurons associated with one leg. Each leg decoder predicts eight joint controls and one adhesion control. Together, the leg decoders produce 54 controls, which join the five body controls to form $\widehat{\widetilde{a}}_t \in \mathbb{R}^{59}$ in standardized coordinates. Every decoder is linear, so it forms only weighted sums of the selected neural rates. Nonlinear feature construction must therefore occur in the recurrent connectome dynamics rather than in a learned output network. No decoder receives the observation directly, so every input-dependent control must first travel through the anatomically routed circuit.

B.6 End-to-end policy operation

SUPERFLY repeats one recurrent policy step. The step routes the current observation into the connectome, advances neural activity, and decodes one control vector. Training unrolls consecutive steps so control error updates the fitted quantities through time. During deployment, MuJoCo returns the resulting body state, while SUPERFLY carries its final neural voltage into the next step.

B.6.1 Policy update

Algorithm 1 gives one policy update with K neural integration steps. It standardizes the observation, routes external current, repeats the recurrent circuit update, and decodes the final neural release. The reported model uses $K = 4$ and $\Delta t = 0.5$ ms, which span one 2 ms body step.

Algorithm 1 SUPERFLY, one embodied control step. Given an update count K and interval Δt, the policy routes one observation through fixed anatomy, advances neural voltage, and decodes the resulting release.

Require: observation x_t; prior voltage $V_{t,0}$; fixed anatomy $\mathcal{A}$; fitted quantities Θ; update count K; interval Δt

$\widetilde{x}_t \leftarrow \textsc{Standardize}(x_t)$ ▷ normalize input
$I_t \leftarrow \textsc{RouteInputs}(\widetilde{x}_t)$ ▷ route body signals
for $k = 0, \ldots, K-1$ **do** ▷ recurrent updates
  $r_{t,k} \leftarrow \mathrm{ReLU}(V_{t,k})$ ▷ graded release
  $m_{t,k} \leftarrow \textsc{SynapticCurrent}(r_{t,k}, \mathcal{A}, \Theta)$ ▷ sum signed inputs
  $V_{t,k+1} \leftarrow \textsc{LeakyIntegrate}(V_{t,k}, m_{t,k}, I_t; \Theta, \Delta t)$ ▷ leak and integrate
end for
$\widehat{\widetilde{a}}_t \leftarrow \textsc{MotorDecode}(\mathrm{ReLU}(V_{t,K}))$ ▷ decode motor rates
return $\widehat{\widetilde{a}}_t, V_{t,K}$ ▷ carry neural state

B.6.2 Behavioral-cloning training

Training fits the interface, physiology, and decoder quantities while the anatomical graph remains fixed. For a sequence batch $\mathcal{B}$ of valid body steps, the objective is the mean squared error between predicted and teacher controls:

$$\mathcal{L}_{\text{BC}} = \frac{1}{59|\mathcal{B}|} \sum_{t \in \mathcal{B}} \left\| \widehat{\widetilde{a}}_t - \widetilde{a}_t^{\star} \right\|_2^2. \tag{16}$$

The loss averages squared error over all 59 controls and every valid step in $\mathcal{B}$. Equation (7) divides each control by its training-set standard deviation, giving controls with different numerical scales comparable influence. The optimizer fits 37,296 scalars in three groups. The input interfaces use 5,978 to convert observations into neural current. Circuit dynamics use 21,095: 14,235 synaptic scales, 3,430 resting voltages, and 3,430 time constants. The motor decoder uses 10,223 to map neural activity to controls. Control-error gradients pass backward through the decoder, recurrent simulation, and input interfaces. Training changes only these fitted quantities; every BANC edge endpoint, synapse count, and transmitter-derived sign remains fixed.

Learning through time. Each optimization step processes 16 rollout streams in parallel. Every stream contributes up to 16 consecutive 2 ms frames, yielding at most 256 valid targets per step. Training follows every rollout in temporal order and supplies the control executed at step $t-1$ as input at step t. The final voltage of one chunk initializes the next chunk from the same rollout, preserving circuit history in the forward simulation. At each chunk boundary, we detach that voltage from the autograd graph. Gradients therefore span 16 body steps, or 32 ms, while neural state persists across the full rollout. Truncated backpropagation therefore bounds memory without resetting the simulated circuit.

Optimization. Adam optimizes all fitted quantities with weight decay 10^{-5}. Neural-dynamics parameters use a learning rate of 10^{-4}, while the input interfaces and motor decoder use 5×10^{-3}. Shared dynamics receive smaller updates because each parameter affects many neurons and recurrent steps. The larger rate learns how simulator measurements become neural current and how neural activity becomes controls. A cosine schedule reduces both rates to 1% of their initial values over 50 epochs. During validation, voltage persists within each complete rollout and resets only at rollout boundaries. We save the checkpoint with the lowest validation MSE. Closed-loop evaluations then begin from every held-out posture. Their outcomes never influence checkpoint selection.

B.6.3 CLOSED-LOOP DEPLOYMENT

Deployment freezes every fitted quantity and no longer uses teacher controls. Equation (17) converts the standardized decoder output back to simulator units and clips each control to the valid range:

$$\widehat{a}_{t,j} = \mathrm{clip}\Big(\bar{\sigma}^a_j \widehat{\widehat{a}}_{t,j} + \mu^a_j, -1, 1\Big). \tag{17}$$

The training-set mean μ^a_j and scale $\bar{\sigma}^a_j$ remain fixed. At body step t, MuJoCo applies the 59-control vector for 2 ms and returns the next body readings. The executed vector $\widehat{a}_t$ becomes the preceding-control block of x_{t+1}, while $V_{t,4}$ initializes the next neural state $V_{t+1,0}$. The following policy step therefore combines updated body readings and command with the executed control and carried neural state. Appendix E audits each 2 ms policy step through analytical operations, synchronized latency, and whole-board energy.

C TRAINING CURVES AND ADDITIONAL LOCOMOTION ANALYSIS

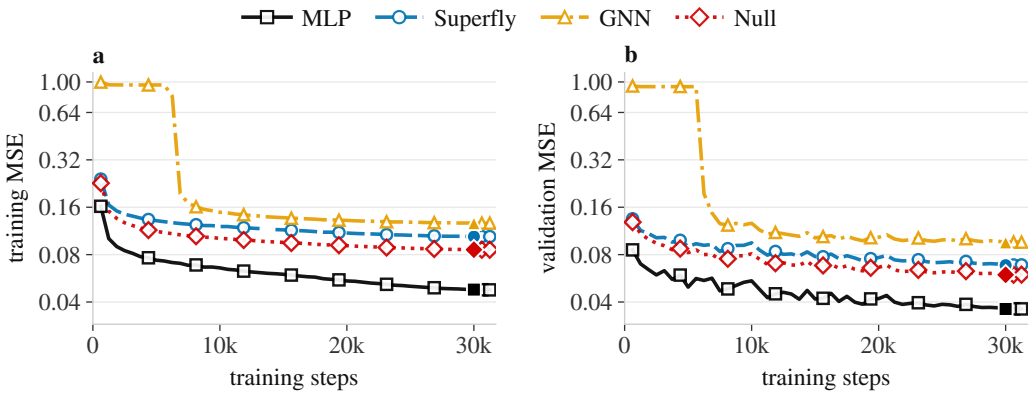

Figure 8: **Matched optimization ranks offline action fit.** (a) Training and (b) validation MSE for a multilayer perceptron (MLP), Superfly, graph neural network (GNN), and degree-preserving Null. Each policy receives 7.5 million frame presentations from the same 150,000-frame corpus and has the same parameter count. Markers sample each history; filled markers identify the validation-selected checkpoint. MLP reaches the lowest validation error, followed by Null, Superfly, and GNN. Closed-loop rollouts separately test whether low offline error sustains gait.

The five-second comparison in Table 1 and the ten-second cohort in Figure 3 test different evaluation horizons. The GNN often supports the body but does not sustain propulsion. On the longer cohort it advances 0.269 ± 0.126 cm and stays upright in 21 of 25 rollouts. The Null stays upright in all 25 rollouts but accumulates 3.93 rad of heading error and ends 3.37 cm behind its start. Superfly advances 10.64 cm but ends about 8.5 cm left of the command line. The MLP advances about 19.4 cm with a rightward veer. These observations distinguish body support from gait and accurate trajectory regulation.

Superfly follows right-turn commands in five of five rollouts and left-turn commands in three of five. The corresponding endpoint errors are 4.6 and 7.6 cm. The behavior supports the use of action imitation to obtain gait, while also showing why closed-loop feedback is needed to correct systematic steering errors. None of these rollout outcomes is used to select checkpoints.

D CELL ATLAS: BIOLOGICAL ROLES AND MODEL INTERFACES

Table 2: **Cell atlas of Superfly's neuron types.** Neurons gives the model count (sensory rows count route recipients). Roles summarize reported biological functions or routed variables; model roles identify command input, body input, action history, readout, or recurrent processing. Swatches match Figures 1, 6, and 7.

Type or class	Neurons	Role or signal	Model role
Central Complex, compass			
EPG	45	internal compass, the activity bump marks heading (Seelig & Jayaraman, 2015; Kim et al., 2017; Turner-Evans et al., 2020)	recurrent
PEN	42	moves the compass bump as the fly turns (Turner-Evans et al., 2017; Green et al., 2017)	recurrent
Central Complex, goal and travel			
FC2A	8	goal direction, upstream of steering (Mussells Pires et al., 2024; Westeinde et al., 2024)	command
FC2B	38	goal direction, upstream of steering (Mussells Pires et al., 2024; Westeinde et al., 2024)	command
FC2C	41	goal direction, upstream of steering (Mussells Pires et al., 2024; Westeinde et al., 2024)	command
PFNd	39	heading with self-motion, travel direction (Hulse et al., 2021; Lyu et al., 2022)	recurrent
PFNv	23	heading with self-motion, travel direction (Hulse et al., 2021; Lyu et al., 2022)	recurrent
Central Complex, steering			
PFL1	14	compares heading and goal direction for corrective flight steering (Bates et al., 2026)	recurrent
PFL2	12	sets turn strength (Westeinde et al., 2024)	recurrent
PFL3	25	picks the turn side, exits through the LAL (Westeinde et al., 2024; Mussells Pires et al., 2024)	recurrent
Descending neurons			
DNa01	2	carries steering toward the nerve cord (Rayshubskiy et al., 2025)	history, readout
DNa02	2	steering, the fly turns toward the more active side (Rayshubskiy et al., 2025; Yang et al., 2024)	history, readout
DNa03	2	at PFL2 and PFL3 convergence, upstream of DNa02 (Westeinde et al., 2024)	history, readout
Premotor relay			
LAL010	2	relays PFL2 and PFL3 steering output to DNa02 (Rayshubskiy et al., 2025)	recurrent
Peripheral sensory classes			
campaniform sensilla (CS)	62	leg load, $x[0{:}18]$	body input
femoral chordotonal (FeCO)	757	joint speed, $x[77{:}83]$	body input
hair plates	243	joint limits, $x[83{:}119]$	body input
tarsal mechanosensors	4,164	claw contact, $x[128{:}134]$	body input
haltere afferents	439	body rotation, $x[119{:}125]$	body input
Johnston's organ (JO)	1,193	gravity and wind, $x[125{:}128]$	body input

E Inference Cost of Neuron-Addressed Control

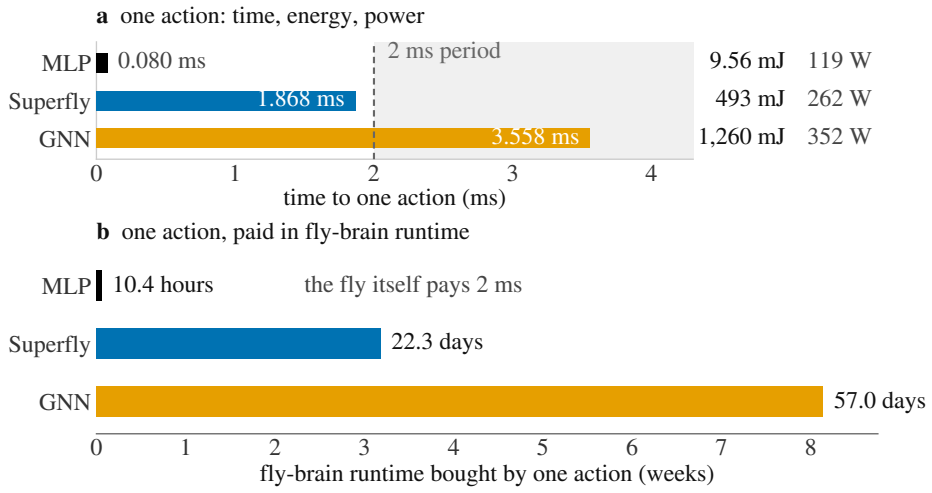

Figure 9: **Matched-parameter inference time and equivalent fly-brain energy.** Every policy contains 37,296 fitted quantities. (a) Black, blue, and orange bars report median latency for the multilayer perceptron (MLP), SUPERFLY, and graph neural network (GNN) across five trials. The dashed line marks the 2 ms control period, and the gray region marks missed deadlines. Right columns report median board energy and power per action. (b) Each action's board energy is expressed as the time a 256 nW fly brain would run on the same energy (Panda et al., 2026). The MLP equals 10.4 hours, SUPERFLY equals 22.3 days, and GNN equals 57.0 days, while the fly brain spends 0.5 nJ in one 2 ms period. Among the graph policies, only SUPERFLY's median latency is below the control period, yet its action still costs about 10^9 fly-brain periods.

Parameter matching does not match computation. SUPERFLY carries one scalar voltage per named neuron, while GNN carries a learned feature vector over the same graph. We therefore measure one deployed action from SUPERFLY, GNN, and MLP under the same training and evaluation workload. Latency and whole-board energy use one shared device and held-out input sequence (Schwartz et al., 2020; Henderson et al., 2020; Luccioni et al., 2024). The archived profile retains every raw power sample, checkpoint hash, and source hash. The audit asks whether neuron-addressed dynamics avoid the cost of generic graph embeddings while preserving action accuracy.

Scalar neuron state avoids generic embedding overhead. Both graph policies traverse the same BANC edges with the same parameter count. SUPERFLY sends one scalar along each edge, while GNN sends a feature vector and remixes it at every neuron. GNN therefore pays once to move features and again to transform them. The representation choice creates a $103\times$ counted-operation gap. Hardware overhead narrows the difference, but SUPERFLY still uses $1.9\times$ less time and $2.6\times$ less board energy. SUPERFLY also fits held-out actions more accurately. MLP remains faster and cheaper than both graph policies, so the result does not favor graph computation in general. For this task, biologically structured scalar dynamics are more efficient than generic node embeddings.

The wiring survives translation, but the energy scale does not. Direct calorimetry places the adult *Drosophila* brain at 256 nW, or 0.5 nJ per 2 ms period (Panda et al., 2026). One SUPERFLY action uses 493 mJ, the energy that brain uses over 22.3 days at the measured rate. The conversion compares energy budgets, not computational equivalence. Even the cheaper graph policy remains about 10^9 fly-brain periods per action. The result locates the remaining gap in the substrate: sparse circuit structure reduces software work, but conventional GPU execution still dominates energy. Event-driven hardware that charges per synaptic event offers the direct route toward closing that substrate gap (Boahen, 2022; Kudithipudi et al., 2025).