When Connectomic Wiring Hinders Imitation: Interface Constraints and the Latency of Sensory Access

Anonymous authors
Paper under double-blind review

Abstract

Fly connectomes are increasingly proposed as architectural priors for neural controllers, with reports that a policy structured by the wiring diagram learns better than random graphs. We test this on four connectomes (male and female ventral nerve cords, two brains) with 14 null models, two input/output (IO) designs and three motor tasks. Under gradient-based imitation the real wiring never outperforms a random graph. With a dense trainable IO we find no consistent effect of topology; with a biologically wired IO (per-leg sensing and actuation, commands through descending neurons) the real wiring is worse, error rises monotonically with structural fidelity, and the effect replicates in the other sex and on a hexapod robot. Hubs, modularity, reciprocity and spectral properties do not explain the penalty; the scarcity of direct afferent→efferent edges does. Degree- and strength-preserving swaps that add such edges bring the connectome to the level of a degree-matched random graph, sham swaps do not, and only sensory→motor shortcuts help. The penalty is unchanged under a teacher with no causal dependence on other legs, is reduced when the teacher tolerates two to three ticks of sensory delay, and lies predominantly below 50 Hz. The wiring diagram withholds one-tick paths from sensory populations that predict the target; a fitted accessibility–demand statistic ranks 27 graph variants ($\rho = 0.92$). Under our implementation the published positive result does not survive an Erdős–Rényi baseline.

1 Introduction

Connectome-constrained modelling has two strands. In the first, connectivity is fixed and unknown cellular parameters are fitted to a task; the resulting model predicts neural activity at single-neuron resolution (Lappalainen et al., 2024; Shiu et al., 2024). In the second, the wiring diagram is used as an architectural prior for an artificial controller: the graph is instantiated as a message-passing network, trained end to end, and compared with random graphs (Jin et al., 2026). The second strand makes a strong claim—that evolved wiring carries an inductive bias useful for control—but the evidence rests on two null models (a degree-preserving rewiring and an Erdős–Rényi graph) and a single, untested IO design.

We ask what inductive bias, if any, a fly connectome provides to a gradient-trained controller, and we answer with interventions rather than correlations. The short answer is that the bias is real but is one of information access and its latency: the wiring diagram routes most sensory input to motor neurons through interneurons, so one-tick access from any sensory population to a motor pool is scarce, and other legs, the largest such population, are reached only through long, weak routes. Every leg's afferents carry signals that predict the target, and a student without one-tick paths to them is at a disadvantage in our experiments whether or not the teacher uses them; that this disadvantage is one of latency is supported but not settled (Sec. 4.7). Our contributions are:

1. A null-model family that separates degree sequence, hub structure, modularity, reciprocity, block structure at three anatomical granularities, sign, weight, and flow architecture (Sec. 3.2), applied to four connectomes.

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.

2. Under dense IO no connectome beats a random graph (Sec. 4.1); under biologically wired IO structural fidelity is monotonically harmful, on both sexes and on a second body (Sec. 4.2).

3. Interventions in both directions (re-targeting, and weight-matched exact-degree swaps with a sham control) showing that direct sensory→motor connectivity accounts for most of the penalty, with an audit of what each operation preserves (Sec. 4.4); teacher-side tests showing that the penalty does not depend on what the teacher causally uses and is reduced when the teacher tolerates sensory delay, and a frequency decomposition placing it predominantly below 50 Hz (Sec. 4.7); and an accessibility–demand model with representability bounds (Appendix E) that ranks all 27 graph variants from the graph and the teacher alone.

4. A re-examination of the published positive result showing that, under our implementation, its ordering depends on the choice of null model (Sec. 4.5).

5. Characterisation of what the frozen connectome does carry (command-driven rhythm), of the one task where the penalty vanishes (Sec. 4.6), and of an extracted walking module that gives no advantage under evolutionary interface training (Appendix H).

2 Related work

Connectome-constrained models of neural activity. Fixing connectivity from electron microscopy and fitting cellular parameters to a task predicts single-neuron responses in the fly visual system (Lappalainen et al., 2024); leaky integrate-and-fire models of the whole FlyWire brain recover sensorimotor circuits (Shiu et al., 2024); rate-model simulations of the nerve cord isolate a walking central pattern generator (Pugliese et al., 2025). These works ask whether structure predicts the animal's computation; ours asks whether it transfers to a controller trained for a different objective.

Connectomes as architectures for artificial networks. Small connectomes (C. elegans) have been used as fixed architectures or reservoirs and reported to outperform random networks (Bardozzo et al., 2024; Schmidgall et al., 2022); the Drosophila brain as a reservoir for time-series prediction (Costi et al., 2025); human structural connectomes as reservoirs whose performance tracks network statistics (Suárez et al., 2021). "Model connectomes" transfer sparse wiring across generations of language models (Kotar & Tuckute, 2025). Closest to us, FlyGM instantiates the FlyWire brain as a graph controller for the flybody model and reports better sample efficiency than degree-preserving and ER graphs (Jin et al., 2026). We reproduce its setting on the same data and show that the reported ordering depends on the null model (Sec. 4.5).

Null models for network structure. Degree-preserving rewiring (Maslov & Sneppen, 2002) and motif statistics (Milo et al., 2002) are standard; we add block-preserving rewiring at three granularities, synthetic modularity and reciprocity controls, and direct interventions on afferent→efferent shortcuts.

Embodied fly simulation and hierarchical robot control. NeuroMechFly (Lobato-Rios et al., 2022; Wang-Chen et al., 2024) and flybody (Vaxenburg et al., 2025) provide biomechanical fly bodies; the latter ships reinforcement-learned walking and flight policies that we use as teachers. In robotics, fast/slow architectures pair a slow vision-language planner with a fast low-level controller (Bjorck et al., 2025; Chen et al., 2025; Zhang et al., 2024; Kim et al., 2024). Our results speak to the design of the fast layer: in our experiments the reflex-like, shortcut-free structure that biology uses there was not obtained by imitation of a global policy.

3 Methods

3.1 Connectomes

MaleCNS v1.0 (Berg et al., 2026) is the complete central nervous system of one male (166,700 neurons). We extract the thoracic ventral nerve cord (VNC): 21,507 traced neurons and 858 k connections with ≥ 5 synapses (1,314 descending neurons, DN; 8,191 afferents; 12,786 intrinsic; 530 motor neurons, MN). 86.5% of MN input comes from VNC-intrinsic neurons and 10.3% directly from DNs, consistent with published analyses of the male nerve cord (Cheong et al., 2026; Takemura et al., 2024). Edge sign is the consensus neurotransmitter of the presynaptic neuron (Eckstein et al., 2024) (ACh +; GABA, Glu −; 96% coverage). From BANC (Bates et al., 2026) (female brain-and-nerve-cord) we extract the homologous thoracic VNC (21,554 proofread neurons, 391 k edges; 88.2% of MN input intrinsic, 8.8% from DNs), mapping leg nerves, neuromeres and effector labels onto the same IO groups. We also extract the MaleCNS central brain (50,570 neurons, 2.43 M edges) and, to reproduce the setting of Jin et al. (2026), the FlyWire v783 female brain (Dorkenwald et al., 2024; Schlegel et al., 2024) (54,778 neurons without optic lobes; 138,625 with), using the FlyWire flow labels (afferent/intrinsic/efferent).

3.2 Null models

Degree-preserving rewiring (DP; post-synaptic endpoints permuted, multi-edges merged), exact-degree Maslov–Sneppen swaps ($\mathrm{DP}_{\mathrm{swap}}$) (Maslov & Sneppen, 2002), Erdős–Rényi (ER), partial DP rewiring (10/30/60% of edges), block-preserving DP within cell-type, hemilineage, or flow (afferent/intrinsic/efferent) blocks, sign shuffle, weight shuffle, synthetic modular graphs (random 24-module partition; within-module fraction $q \in \{0.3, 0.6, 0.9\}$), synthetic reciprocal graphs ($r \in \{0.05, 0.12\}$), and two interventions: +shortcuts (a fraction of afferent→intrinsic edges re-targeted to efferents) and −shortcuts (all afferent→efferent edges re-targeted to intrinsic neurons). All nulls keep the node table—roles and IO identities—fixed.

3.3 Controllers

Connectome GNN. Node state $H \in \mathbb{R}^{N \times d}$ ($d = 16$). One control step computes $M = WH$, with W the row-normalised signed synapse-count matrix, followed by a shared gated update $H' = \mathrm{GRUCell}([M, U, \delta_i], H)$ with trainable per-node descriptors δ_i (as in Jin et al. (2026); the gated update prevents divergence over long rollouts). Frozen core. A rate model $r \leftarrow r + \frac{dt}{\tau}\big(-r + \tanh \mathrm{relu}(g(Wr + u) - \theta)\big)$ with W, $g = 8$ ($g = 6$ for the extracted walking module of Appendix H), $\tau = 20$ ms and θ fixed; only the IO is trained. A DMN-style variant (Lappalainen et al., 2024) additionally trains per-neuron gain and threshold. IO designs. Dense: a trainable encoder from the full observation onto every afferent node and a decoder from every efferent node. Bio-wired: each leg's 29 proprioceptive channels enter only that leg's sensory afferents (entry nerve × side); the command channel enters only DNs through a rank-32 bottleneck; vestibular channels enter haltere/wing afferents; each leg's 9 actuators are decoded only from that leg's MNs (effector subclass × side). Baselines. MLP (4×512) and GRU.

3.4 Tasks

T1 Walking imitation in the flybody physics model (Vaxenburg et al., 2025) at 500 Hz (59 actuators), tracking synthetic constant-speed/yaw references; the teacher is the published DMPO policy; behaviour cloning with truncated BPTT through whole episodes (3,000 updates, batch 16, window 32). T2 Tethered walking on a ball, cloned from the same policy driven by synthetic forward references. T3 Real fly gait continuation: 16,252 motion-capture snippets of walking females; input is root velocity plus a 32-step proprioceptive warm-up; output is 66 leg joint angles; we report explained variance on the generation phase. Closed-loop evaluation (heading error, falls, perturbations) is reported for controllers that walk.

Table 1: Dense IO, walking imitation (T1), validation MSE (lower is better; mean ± s.d. over 3 seeds where shown).

Graph	connectome	degree-preserving	flow-block DP	ER
MaleCNS VNC (21.5 k)	0.171	0.168	0.171	0.164
MaleCNS brain (50.6 k)	0.179 ± 0.012	0.175 ± 0.007	0.189	0.172 ± 0.008
FlyWire brain w/o optic lobes (54.8 k)	0.170 ± 0.004	0.182 ± 0.008	0.175 ± 0.007	0.168 ± 0.003
FlyWire full brain (138.6 k)	0.193	0.205	–	0.194

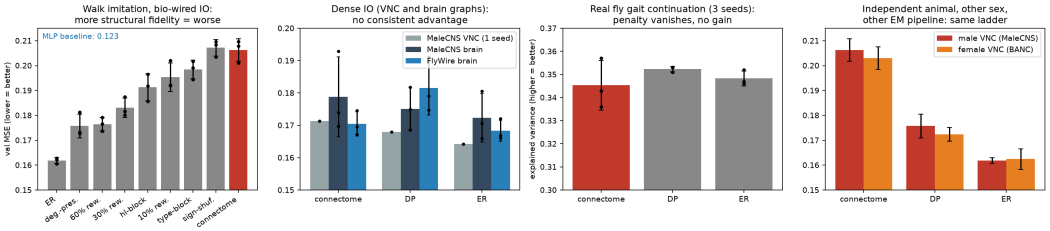


Figure 1: (a) Bio-wired IO on the MaleCNS VNC: error increases with structural fidelity (3 seeds for connectome, DP, ER; red = real wiring). (b) Dense IO: no consistent ordering on the VNC or either brain graph. (c) Real fly gait continuation (T3), trainable bio-wired GNN: the penalty vanishes ($p = 0.32$) but no gain appears. (d) The female VNC (BANC) reproduces the male ladder.

4 Results

4.1 Dense IO: no consistent connectome advantage

Table 1 reports T1 validation MSE (normalised action units; MLP baseline 0.123). Connectome, DP and ER are statistically indistinguishable on the MaleCNS VNC, the MaleCNS brain and the FlyWire brain; validation curves coincide throughout training (Appendix D). With a dense trainable encoder and decoder we detect no consistent effect of topology; the confidence intervals (Appendix C) do not establish that the effect is negligible. Throughout, "seed" varies both the sampled null graph and the training initialisation; non-significant differences are not claims of equivalence (confidence intervals in Appendix C).

4.2 Bio-wired IO: structural fidelity is monotonically harmful

With the biologically wired IO on the MaleCNS VNC (Fig. 1a, Table 2), error increases with structural fidelity (monotonically within the partial-rewiring series; the ordering across null families is descriptive): ER < DP ≈ DP_{swap} < 60% < 30% < hemilineage-block < 10% rewired < flow-block ≈ type-block < sign-shuffled ≈ connectome. Sign is irrelevant. The same ordering holds with frozen dynamics (T2 MSE: connectome 0.78 vs. ER 0.65; T3 frozen core, explained variance, higher is better: 0.115 vs. 0.161) and with DMN-style neuron-parameter training (T3 explained variance, 3 seeds: connectome 0.143 ± 0.010 vs. DP 0.189 ± 0.009, $p = 0.004$). The female VNC (BANC) reproduces the ladder almost exactly—connectome 0.203 ± 0.005 > DP 0.172 ± 0.003 > ER 0.162 ± 0.004 ($p = 6 \times 10^{-4}$, connectome vs. DP)—with the same flow signature (Fig. 1d). The effect is therefore not specific to one individual, sex, or reconstruction pipeline. It also transfers to a second body: on an 18-DoF hexapod cloned from an engineered tripod gait (Appendix J) the same ordering appears (connectome 0.0045 ± 0.0003 vs. ER 0.0032 ± 0.0004, $p = 0.014$) with a much smaller gap, as the predictive reading of Sec. 4.7 expects for a teacher without causal sensory dependence; all variants control the robot, and under a 20 ms observation delay the connectome-wired controller has the smallest speed error (0.034 vs. 0.046, $p = 0.025$), a single-task hint.

Table 2: Bio-wired IO on the MaleCNS VNC (T1). Direct mass: fraction of motor-neuron synaptic input arriving directly from afferents. Mean±s.d. over 3 seeds (multi-value rows: means only; s.d. ≤ 0.006).

Variant	What is preserved	Direct mass	val MSE
connectome	everything	0.132	0.206 ± 0.005
sign-shuffled	topology, weights	0.132	0.207 ± 0.003
type-block DP	cell-type block counts	–	0.198 ± 0.004
flow-block DP	afferent/intrinsic/efferent block counts	0.132	0.197 ± 0.003
hemilineage-block DP	hemilineage block counts	–	0.191 ± 0.005
10 / 30 / 60% rewired	degrees, partial topology	–	0.195 / 0.183 / 0.176
degree-preserving (DP)	degrees (approx.)	0.266	0.176 ± 0.005
$\mathrm{DP}_{\mathrm{swap}}$	exact degrees, hubs	–	0.176 ± 0.003
synthetic modular $q = 0.3/0.6/0.9$	degrees + modularity	–	0.173 / 0.172 / 0.172
synthetic reciprocal $r = 0.05/0.12$	degrees + reciprocity	–	0.172 / 0.171
ER	–	0.376	0.162 ± 0.001
connectome + 10% shortcuts (re-targeting)	source endpoints, out-degrees, weights	0.260	$\mathbf{0.169 \pm 0.003}$
connectome + 30% shortcuts		0.429	0.172 ± 0.001
DP − shortcuts (re-targeting)	DP out-degrees, no shortcuts	0.000	$\mathbf{0.194 \pm 0.006}$
swap shortcuts (weight-matched)	exact degrees, in/out-strengths	0.260	$\mathbf{0.171 \pm 0.001}$
sham swaps (same number, no new shortcuts)	as above, plus flow blocks and MN rows	0.132	0.200 ± 0.001
+ descending→MN only (matched swaps)	exact degrees, in/out-strengths	0.260	0.203 ± 0.002
+ same-leg sensory→MN only (matched swaps)		0.260	0.187 ± 0.002
+ cross-leg sensory→MN only (matched swaps)		0.260	$\mathbf{0.175 \pm 0.002}$
connectome, unit unsigned weights	topology	–	0.202 ± 0.004
DP, unit unsigned weights	degrees	–	0.169 ± 0.001
ER, unit unsigned weights	–	–	0.161 ± 0.003

4.3 What is not sufficient to explain it

Path length (DN→MN 2.85 vs. 2.73 hops), spectral radius (matching both graphs at $\rho = 0.5$ leaves the gap unchanged), propagation depth (three sub-updates per control step, i.e. two edge traversals of the current observation: both improve, gap unchanged), decoder capacity, hub structure ($\mathrm{DP}_{\mathrm{swap}}$ keeps the exact in-degree distribution and behaves like DP), modularity and reciprocity (synthetic modular graphs with 5.1 M 3-cycles and reciprocity 0.10—both above the connectome—score 0.172–0.173; synthetic reciprocal graphs 0.171–0.172). Correlations between error and these statistics across real-derived nulls (Spearman +0.88 to +0.93) are therefore confounded with anatomical organisation; matching or exceeding a statistic shows it is not a sufficient explanation, not that it contributes nothing in interaction with others. The real connectome's operator does carry ≥ 6 marginally stable modes ($|\lambda| = 1$; closed excitatory loops absent from every null), but damping them does not close the gap.

4.4 What is the cause: flow architecture

Preserving only the edge counts between afferent, intrinsic and efferent blocks (random inside) recovers most of the penalty (0.197 ± 0.003). The real connectome routes 13% of motor-neuron input directly from afferents; DP and ER create 27% and 38% (Table 2). We then intervene on this connectivity. Two operations are used. Re-targeting: 10% of afferent→intrinsic edges are re-targeted to random efferents (direct mass 0.26; out-degrees preserved, efferent in-degrees and hence the row normalisation of their inputs change), giving 0.169 ± 0.003, at or below DP; the reverse operation on the DP graph (all afferent→efferent edges re-targeted to intrinsic neurons, direct mass 0) gives 0.194 ± 0.006, most but not all of the way to the connectome's 0.206. Weight-matched swaps: pairs (afferent→intrinsic, intrinsic→efferent) with the same synapse count are swapped into (afferent→efferent, intrinsic→intrinsic) until direct mass reaches 0.26, so every node keeps its exact in- and out-degree and its exact in- and out-strength, and the normalisation of every untouched edge is unchanged (Appendix C); a sham performs the same number of matched swaps among afferent→intrinsic edges only (flow blocks unchanged, motor-neuron rows untouched). Swap-shortcuts score 0.171 ± 0.001 vs. sham 0.200 ± 0.001 (3 seeds each; $p = 10^{-5}$; sham vs. connectome $p = 0.15$). Restricting the afferent side of the swap to one source, each matched to direct mass 0.26, gives descending→MN 0.203 ± 0.002 (no improvement; $p = 0.31$ vs. connectome), same-leg sensory→MN 0.187 ± 0.002, cross-leg sensory→MN 0.175 ± 0.002 (the DP level; $p = 0.91$). Unmatched swaps and re-targeting

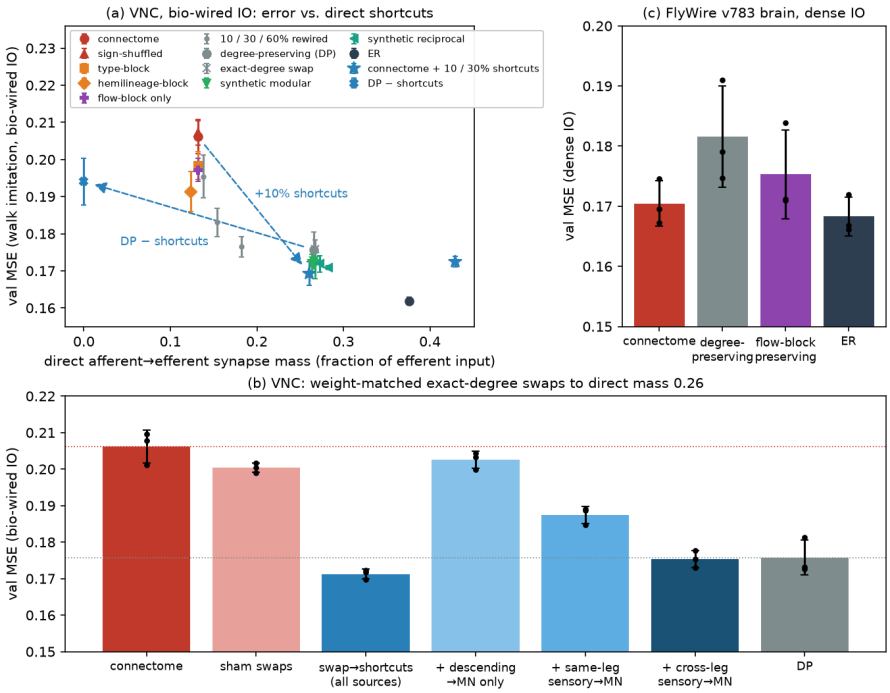


Figure 2: (a) Error against direct afferent→efferent synapse mass for all VNC variants (bio-wired IO); dashed arrows are the two re-targeting interventions. Modularity, reciprocity and hub controls (green, teal, grey) sit with DP. (b) Weight-matched exact-degree swaps (every node's degrees and strengths preserved); the four shortcut-creating interventions are brought to direct mass 0.26, the sham keeps 0.132: sham swaps leave the penalty, swaps that create shortcuts remove it, and by source only cross-leg sensory→MN access does so (dotted lines: connectome and DP means; dots: seeds). (c) FlyWire brain, dense IO: the degree-preserving null underperforms, but ER and a flow-preserving null match the connectome.

versions give the same ordering (Appendix C). Direct sensory→motor connectivity therefore accounts for most of the penalty. The real VNC gives each leg's motor pool essentially no route to other legs' sensory afferents within three synapses ($M_3 = 0.002$ vs. 0.03–0.08 in random graphs), and a linearised accessibility model (Appendix E) predicts that adding that route removes the penalty, as observed; measured access curves (Fig. 4) show how own-leg and cross-leg access grow with depth in each graph. Why cross-leg access matters when the teacher could be leg-local is the subject of Sec. 4.7.

4.5 Under our implementation the published ordering depends on the null model

On FlyWire, under our implementation, the DP null underperforms the connectome (0.182 ± 0.008 vs. 0.170 ± 0.004, $p = 0.11$), in the direction reported by Jin et al. (2026), but ER matches the connectome (0.168 ± 0.003, $p = 0.49$) and a flow-preserving DP null lies in between (0.175 ± 0.007). FlyWire contains 29–43% afferent neurons with no output at the synapse threshold; DP rewiring keeps them output-less and scrambles the afferent→efferent flow (direct mass 0.126 → 0.05–0.08), whereas ER does not. FlyGM's non-connectome baselines carry unweighted, unsigned edges while its connectome carries signed synapse counts; unit-weight baselines give the same ordering in our implementation (Appendix F), so weighting does not explain the discrepancy.

FlyGM's code is not public, so our replication follows its paper; the shared and differing settings are listed in Appendix F, and we make no claim about FlyGM's closed-loop results. Under our implementation the apparent benefit over a degree-preserving null does not extend to ER, is reduced by a flow-preserving null, and is unchanged by giving the baselines unit weights.

4.6 Frozen dynamics and task alignment

Clamping the walking command neuron DNg100 in the frozen rate model (Appendix I, Fig. 5), only the real connectome sustains rhythmic activity (5–10 Hz, the fly stepping range) in a balanced sparse regime; DP and ER nulls are silent or saturated. This is consistent with connectome CPG simulations (Pugliese et al., 2025), yet confers no advantage under gradient training. On real fly gait continuation (T3) with the trainable bio-wired GNN the penalty is absent (connectome 0.345 ± 0.011, DP 0.352 ± 0.001, ER 0.348 ± 0.003; $p = 0.32$) but no gain appears, and the connectome shows ten-fold higher seed variance. T3 differs from T1 in target, protocol and data, so the absence of a penalty is compatible with task alignment but does not isolate it. The accessibility model offers a sharper reading: during generation the only external input is the root velocity, which enters through descending neurons every motor pool reaches within one step ($M_1 = 0.073$), so no penalty is predicted for the trainable student, as observed.

Evolutionary optimisation of an interface to an extracted walking module does not rescue it either: the module scores no better than its rewired version and below an engineered clock (Appendix H).

4.7 Testing the demand side: access, latency and prediction

Every intervention above changes the graph with the teacher fixed; we now fix the graphs and change the teacher (Appendix G; Table 3). We distil the expert into two MLP teachers of identical architecture: a local teacher, in which each leg's nine actuators are a function of that leg's 29 channels and the 467-channel command block only, and a global teacher that sees every channel a bio-wired student can access. The local teacher reproduces the expert to a validation error of 0.136 and the global one to 0.079; permuting other legs' channels changes the local teacher's output by exactly zero and the global teacher's by 0.97 per leg. Each teacher's actions replace the expert's and the bio-wired students are trained as in T1 with the expert's normalisation.

The penalty does not depend on what the teacher causally uses. Against the local teacher the students score connectome 0.136 ± 0.004, DP 0.106 ± 0.002, ER 0.105 ± 0.001; against the global teacher 0.132 ± 0.002, 0.104 ± 0.001, 0.102 ± 0.000; against their mean 0.120, 0.093, 0.089 (3 seeds each). The connectome-to-DP gap is 0.030, 0.028 and 0.028 ($p \leq 0.003$), against 0.030 under the expert; only the DP-to-ER difference (0.014 under the expert) collapses under the local teacher (0.001). A teacher that never uses other legs is not easier for the connectome, and the benefit of cross-leg shortcuts is equally teacher-independent: under the local teacher, weight-matched cross-leg sensory→MN swaps bring the connectome from 0.136 to 0.110 ± 0.003 ($p = 0.002$; DP 0.106) and descending→MN swaps do not (0.131 ± 0.006, $p = 0.30$), as under the expert. A combined intervention on the student, restricting the loss to the left front leg and holding other legs' inputs at their mean, gives a gap to DP of 0.018 ($p = 0.09$; to ER 0.027, $p = 0.04$; 3 seeds). The same one-leg loss with all sensory inputs retained gives a gap of 0.029 ($p = 0.009$), identical to the full objective, so the reduction to 0.018 is attributable to removing the other legs' inputs rather than to the changed objective (reduction 0.011, about one standard error). Own-leg access does not explain the residual: one-step own-leg access is similar in the connectome and DP (0.012 vs. 0.011) and cumulative own-leg access is higher (Fig. 4), so we do not attribute it.

Tolerated sensory delay reduces the penalty by about half. Delaying the sensory channels the local teacher sees by L ticks (2 ms each; command channels undelayed; all four conditions trained at batch 8) gives connectome−DP gaps of 0.028 ± 0.007, 0.023, 0.015 and 0.014 for $L = 0, 1, 2, 3$ (3 seeds each; reduction from $L = 0$ to $L = 3$: 0.014 ± 0.007 s.e., $p = 0.10$, about half of the zero-delay gap). The observation at tick t enters the afferent nodes' state at t and reaches a motor pool's state through a direct edge at $t + 1$ and through h interneurons at $t + 1 + h$ (Appendix G), so a zero-delay teacher is reachable by no graph and a one-tick teacher only through direct edges, of which the connectome has few (own-leg routes: 31% of sensory–MN pairs within two hops, mean 2.8; cross-leg: 2.4% within two, mean 3.4; Fig. 4). This is consistent with an important role for temporal access, but not the whole story: three sub-updates within a tick, which let the current observation traverse two

Table 3: Demand-side test, bio-wired IO on the MaleCNS VNC, validation MSE in expert-normalised units (3 seeds; teachers defined in the text and Appendix G; the last four rows use batch 8, with the batch-8 local row as comparator; per-graph delayed values in Appendix G). Cross-leg dependence: causal, mean squared output change when another leg's channels are permuted; ridge, predictive demand D per other leg.

Teacher	cross-leg dep. (causal / ridge)	connectome	DP	ER	conn.−DP
expert (DMPO)	– / 0.165	0.206 ± 0.005	0.176 ± 0.005	0.162 ± 0.001	0.030
global MLP	0.97 / 0.159	0.132 ± 0.002	0.104 ± 0.001	0.102 ± 0.000	0.028
mixed	0.24 / 0.144	0.120 ± 0.002	0.093 ± 0.004	0.089 ± 0.001	0.028
local MLP	0 / 0.134	0.136 ± 0.004	0.106 ± 0.002	0.105 ± 0.001	0.030
local, batch 8	0 / 0.134	0.143 ± 0.006	0.114 ± 0.004	0.111 ± 0.003	0.028
local, one leg only, inputs kept	0 / 0.134	0.110 ± 0.008	0.080 ± 0.005	0.082 ± 0.003	0.029
local, one leg only, others blind	0 / 0	0.118 ± 0.011	0.100 ± 0.009	0.091 ± 0.005	0.018
local, delayed 1 / 2 / 3 ticks	0 / –	–	–	–	0.023 / 0.015 / 0.014
smoothed expert	– / –	0.060 ± 0.006	0.042 ± 0.001	0.036 ± 0.002	0.018

edges without delay (Appendix G), left the gap unchanged against the expert (Sec. 4.3). A matched comparison (local teacher, batch 8, 3 seeds) confirms this: three sub-updates at zero delay, which give the current observation the same two-edge reach as a two-tick delay, leave the gap at 0.030 ± 0.010 (change $+0.002 \pm 0.007$ s.e.), whereas the two-tick delay gives 0.015. Within-tick reach is therefore not what delay supplies; the delayed target depends on observations the recurrent state has already integrated, a hop-count account of latency is insufficient, and the mechanism remains open.

The excess error is predominantly low-frequency. Target smoothing reduces the absolute gap but does not eliminate it: against a causally smoothed expert (exponential average, 10 ms) the connectome−DP gap is 0.018 (connectome 0.060 ± 0.006, DP 0.042 ± 0.001, ER 0.036 ± 0.002; $p = 0.03$), smaller than 0.030 in absolute terms and larger relative to the DP error (44% vs. 17% for the unsmoothed expert). An orthogonal frequency decomposition of the error (Fourier bands, additive by Parseval; Table 4) places the connectome's excess predominantly below 50 Hz: the connectome's excess over DP is 0.024 in the low band and 0.007 in the high band (totals 0.206 vs. 0.176; 3 seeds), and the weight-matched shortcut swap removes 0.023 of low-band and 0.006 of high-band error relative to the sham. This identifies where the error lies, not why; a fast-input, slow-output mechanism is a hypothesis these analyses do not test. What they do establish is that low-frequency prediction suffers despite access to correlated sensory signals, and that demand in the model must be read predictively (ridge coefficients, Appendix E): ridge cross-leg demand barely differs between the local and global teachers (0.134 vs. 0.159) and predicts similar gaps; a causal demand of zero would have predicted none.

5 Discussion

What the wiring diagram is a prior for. A connectome is a strong prior for predicting the activity and behaviour of the animal that owns it (Lappalainen et al., 2024), but a weak or negative prior for controlling something else with gradient descent. The feature that hurts—the scarcity of direct afferent→efferent shortcuts—is a design principle of nervous systems: most sensory input reaches motor neurons through interneuron circuits, under a low-dimensional descending command. It conflicts with imitation learning, which rewards shallow paths from every input that predicts the target.

An accessibility–demand principle. The experiments are summarised by one statement (Appendix E): for a structured student with masked IO and K propagation steps, the maps it can realise from input group g to output group e are confined to the walks $A_{\leq K}[\mathcal{E}_e, \mathcal{A}_g]$; a teacher whose output is predictable from input group g ($B_{e,g}$ large) where the accessible walk mass $M_K(e, g)$ is small incurs, for a gain-bounded linear student with uncorrelated inputs, an imitation error that no optimiser can remove (Propositions 1–2). Neither assumption holds for the trained controllers, for which the score is an empirical diagnostic; K is the number of ticks of delay the teacher tolerates (Sec. 4.7). The resulting score ranks 27 graph variants ($\rho = 0.92$–0.93) and the source decomposition of Sec. 4.4, and Sec. 4.7 shows that B must be read predictively (correlated inputs are interchangeable routes to the same

signal) and supports the latency reading of K only in part. Two design consequences follow: distilling a teacher into a modular student needs a teacher that tolerates the student's path latency or explicit direct access (masking the teacher's inputs is not enough), and a fast/slow hierarchy's command bottleneck must have rank at least that of the teacher's locally inaccessible dependence (Corollary 1). The principle applies equally to per-limb controllers, routed mixtures of experts and depth-limited message passing.

Implications for evaluation. Claims that connectome-derived architectures improve learning should be tested against ER graphs and flow-block-preserving nulls with equally weighted baselines, not only against degree-preserving rewiring, which scrambles flow blocks and can yield apparent gains on graphs with many output-less afferents. Lower imitation loss is not better control: on the hexapod the MLP has the lowest loss and the worst closed-loop speed error (Appendix J).

Implications for robot controllers. End-to-end policies (vision-language-action models, monolithic MLP or diffusion policies) are all-shortcut architectures. Our expert and MLP baselines fail under 10 ms sensory delay or a locked limb (Appendix B), but since the bio-wired connectome controllers do not walk we cannot attribute that fragility to shortcuts. The biological alternative—shortcut-poor low-level circuits under a narrow command bottleneck—is what gradient-based imitation penalises. Our results quantify the price of that structure under imitation; show that, for the interfaces, teachers and training procedures tested, hierarchical fast/slow controllers (Bjorck et al., 2025; Chen et al., 2025) do not obtain reflex-like low-level layers by imitating a policy that reacts to sensory input within a control tick (Sec. 4.7); and abstract a reusable per-limb interface under a narrow command bottleneck that does not depend on any particular connectome. The frozen connectome's command-driven rhythm is a candidate untrained pattern-generator template for neuromorphic hardware (Wang et al., 2025; Costi et al., 2025), though it does not beat an engineered clock (Appendix H).

6 Limitations

Our reproduction of FlyGM is by description, not by code (Sec. 4.5); a released implementation could change the numbers in Table 1, though not the null-model argument. Two individuals (one per sex) for the VNC comparison but a single animal per brain graph; rate-model parameters that are homogeneous except in the DMN variant (per-neuron gain and threshold only); locomotion tasks only (fly body plus one simple hexapod); on the fly body, perturbation robustness could not be compared because bio-wired connectome controllers do not walk, and the closed-loop behaviour of the matched-shortcut variants was not evaluated; dense and bio-wired IO differ in accessible channels and parameter count as well as in masks, so their contrast is not attributable to masks alone; the ES comparison uses three seeds with large variance (Table 5); a gap of 0.014–0.018 remains when the teacher tolerates three ticks of delay or when other legs are blind, and the small hexapod ladder under a sensory-independent teacher points the same way; the linear model does not capture this residual, and our aggregate access statistics do not attribute it.

7 Conclusion

Across four fly connectomes, two IO designs, three tasks and three kinds of dynamics, the wiring diagram does not outperform a random graph as an architectural prior for gradient-trained imitation. Under biologically wired IO it is worse, and sham-controlled interventions show that the penalty is carried by the scarcity of direct sensory→motor connectivity: the connectome denies each motor pool short paths from the sensory populations that predict the gait. The penalty is unchanged under a teacher with no causal dependence on other legs, is reduced when the teacher tolerates two to three ticks of sensory delay, and lies predominantly below 50 Hz. The published counter-example is, under our implementation, sensitive to the null model. The wiring diagram's command-driven rhythm is real but confers no controller advantage under gradient or evolutionary interface training.

AI use statement

In this work we used a generative AI coding assistant (Claude, Anthropic) for tasks with required disclosure: it implemented the methods (graph construction, null-model generators, controllers, training and evaluation code), proposed and refined hypotheses and the design of experiments and controls (including the null-model family, the shortcut interventions and the sham control), helped formulate the accessibility–demand model of Appendix E, including drafting the statements and proofs of its propositions (which the authors checked), cleaned and reformatted the connectome and motion-capture datasets, and assisted in interpreting results. We did not use generative AI to generate synthetic datasets, for translation, or for qualitative data analysis. Additionally, we used it for tasks with recommended disclosure: creating and editing software code, creating scientific figures, drafting and editing parts of this paper, and summarising related literature. We have reviewed all AI-assisted work: experimental designs were specified or approved by the authors before execution, all code paths that produce reported numbers were inspected by the authors, reported values were checked against raw run logs, and the literature summary was verified against the cited sources. We take responsibility for the final content of this work, including text, claims and artifacts produced with the aid of generative AI.

Ethics statement

This work uses publicly released invertebrate connectomes and physics simulation only; it involves no human subjects, personal data, or foreseeable harmful application. Its main practical implication is methodological: caution in claiming benefits of biological wiring priors without adequate null models.

Reproducibility statement

All connectome tables are public and downloaded without credentials (MaleCNS gs://flyem-male-cns; BANC gs://lee-lab_brain-and-nerve-cord-fly-connectome; FlyWire Zenodo record 10676866); the flybody simulator and its DMPO walking policy are public. Section 3 and Appendix C specify graph construction (partitions, synapse threshold, neurotransmitter signs, operator normalisation), every null model and intervention, the controllers and IO groups, the optimiser and training schedule, the seed semantics, and confidence intervals on the central comparisons; Appendix A gives data sources and compute. Anonymised code for data extraction, null-model generation and interventions, controllers, training, evaluation and figures, together with the logged numbers behind every table, is provided as supplementary material; a public repository will accompany the camera-ready version.

References

Francesco Bardozzo, Andrea Terlizzi, Claudio Simoncini, Pietro Liò, and Roberto Tagliaferri. Elegans-AI: How the connectome of a living organism could model artificial neural networks. Neurocomputing, 584:127598, 2024.

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; preprint bioRxiv doi:10.1101/2025.07.31.667571.

Stuart Berg, Isabella R Beckett, Marta Costa, Philipp Schlegel, Michał Januszewski, Elizabeth C Marin, Aljoscha Nern, Stephan Preibisch, Wei Qiu, Shin-ya Takemura, et al. Sexual dimorphism in the complete Drosophila male central nervous system connectome. Cell, 2026. PII S0092-8674(26)00942-6; preprint bioRxiv doi:10.1101/2025.10.09.680999. MaleCNS v1.0, https://male-cns.janelia.org.

Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, et al. GR00T N1: An open foundation model for generalist humanoid robots. arXiv preprint arXiv:2503.14734, 2025.

Hao Chen, Jiaming Liu, Chenyang Gu, Zhuoyang Liu, Renrui Zhang, Xiaoqi Li, Xiao He, Yandong Guo, Chi-Wing Fu, Shanghang Zhang, and Pheng-Ann Heng. Fast-in-slow: A dual-system VLA model unifying fast manipulation within slow reasoning. In Advances in Neural Information Processing Systems, 2025. arXiv:2506.01953.

Han S J Cheong, Katharine N Boone, Marryn M Bennett, Farzaan Salman, Jacob D Ralston, Kaitlyn Hatch, Raven F Allen, Jasper S Phelps, et al. Organization of circuits linking descending input to motor output in the Drosophila male adult nerve cord connectome. eLife, 13:e96084, 2026. doi:10.7554/eLife.96084.

Leone Costi, Alexander Hadjiivanov, Dominik Dold, Zachary F Hale, and Dario Izzo. The Drosophila connectome as a computational reservoir for time-series prediction. Biomimetics, 10(5):341, 2025.

Sven Dorkenwald, Arie Matsliah, Amy R Sterling, Philipp Schlegel, Szi-chieh Yu, Claire E McKellar, Albert Lin, Marta Costa, Katharina Eichler, Yijie Yin, et al. Neuronal wiring diagram of an adult brain. Nature, 634:124–138, 2024.

Nils Eckstein, Alexander Shakeel Bates, Andrew Champion, Michelle Du, Yijie Yin, Philipp Schlegel, et al. Neurotransmitter classification from electron microscopy images at synaptic sites in Drosophila melanogaster. Cell, 187:2574–2594, 2024.

Zehao Jin, Yaoye Zhu, Chen Zhang, and Yanan Sui. Whole-brain connectomic graph model enables whole-body locomotion control in fruit fly. arXiv preprint arXiv:2602.17997, 2026. Presented at NeurIPS 2025.

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, et al. OpenVLA: An open-source vision-language-action model. arXiv preprint arXiv:2406.09246, 2024.

Klemen Kotar and Greta Tuckute. Model connectomes: A generational approach to data-efficient language models. In ICLR 2025 Workshop on Representational Alignment (Re-Align), 2025. arXiv:2504.21047.

Janne K Lappalainen, Fabian D Tschopp, Sridhama Prakhya, Mason McGill, Aljoscha Nern, Kazunori Shinomiya, Shin-ya Takemura, Eyal Gruntman, Jakob H Macke, and Srinivas C Turaga. Connectome-constrained networks predict neural activity across the fly visual system. Nature, 634:1132–1140, 2024.

Victor Lobato-Rios, Pembe Gizem Özdil, Shravan Tata Ramalingasetty, Jonathan Arreguit, Stéphanie Clerc Rosset, Graham Knott, Auke Jan Ijspeert, and Pavan Ramdya. NeuroMechFly, a neuromechanical model of adult Drosophila melanogaster. Nature Methods, 19:620–627, 2022.

Sergei Maslov and Kim Sneppen. Specificity and stability in topology of protein networks. Science, 296(5569):910–913, 2002.

Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. Science, 298(5594): 824–827, 2002.

Sarah M Pugliese, Grant M Chou, Elliott T T Abe, Denis Turcu, Jackson K Lancaster, John C Tuthill, and Bingni W Brunton. Connectome simulations identify a central pattern generator circuit for fly walking. bioRxiv, 2025. doi:10.1101/2025.09.12.675944.

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. AISTATS, 2011.

Philipp Schlegel, Yijie Yin, Alexander S Bates, Sven Dorkenwald, Katharina Eichler, Paul Brooks, Daniel S Han, Marina Yasmin, Chloe Cesar, et al. Whole-brain annotation and multi-connectome cell typing of Drosophila. Nature, 634:139–152, 2024.

Samuel Schmidgall, Catherine Schuman, and Maryam Parsa. Biological connectomes as a representation for the architecture of artificial neural networks. bioRxiv, 2022. doi:10.1101/2022.09.30.510374; arXiv:2209.14406.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017.

Philip K Shiu, Gabriella R Sterne, Nico Spiller, Lucy Blenkinsop, et al. A Drosophila computational brain model reveals sensorimotor processing. Nature, 634:210–219, 2024.

Laura E Suárez, Blake A Richards, Guillaume Lajoie, and Bratislav Misic. Learning function from structure in neuromorphic networks. Nature Machine Intelligence, 3:771–786, 2021.

Shin-ya Takemura, Kenneth J Hayworth, Gary B Huang, Michal Januszewski, Zhiyuan Lu, Elizabeth C Marin, Stephan Preibisch, C Shan Xu, John Bogovic, Andrew S Champion, et al. A connectome of the male Drosophila ventral nerve cord. eLife, 13:e97769, 2024.

Roman Vaxenburg, Igor Siwanowicz, Josh Merel, Alice A Robie, Carmen Morrow, Guido Novati, Zinovia Stefanidi, Gert-Jan Both, Gwyneth M Card, Michael B Reiser, Matthew M Botvinick, Kristin M Branson, Yuval Tassa, and Srinivas C Turaga. Whole-body physics simulation of fruit fly locomotion. Nature, 643:1312–1320, 2025.

Felix Wang, Bradley H Theilman, Fred Rothganger, William Severa, Craig M Vineyard, and James B Aimone. Neuromorphic simulation of Drosophila melanogaster brain connectome on Loihi 2. arXiv preprint arXiv:2508.16792, 2025.

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, 21:2353–2362, 2024.

Jianke Zhang, Yanjiang Guo, Xiaoyu Chen, Yen-Jen Wang, Yucheng Hu, Chengming Shi, and Jianyu Chen. HiRT: Enhancing robotic control with hierarchical robot transformers. arXiv preprint arXiv:2410.05273, 2024.

A Data sources and compute

All graphs are built from public releases (MaleCNS gs://flyem-male-cns, BANC gs://lee-lab_brain-and-nerve-cord-fly-connectome, FlyWire Zenodo 10676866). Code, null-model generators, environment wrappers and training scripts are provided as supplementary material. A single 24 GB GPU suffices: a 21.5 k-node connectome GNN with truncated BPTT (batch 16, window 32) uses ≈10.5 GB; the 138 k-node FlyWire graph uses 17 GB at batch 4.

B Closed-loop evaluation

Expert (DMPO) on the clean command grid: heading error 12.2°, 4% falls; MLP after BC+PPO: 18.4°, 17% falls; dense-IO connectome GNN after BC+PPO: 14.9°, 33% falls; bio-wired connectome GNN after BC, DAgger (Ross et al., 2011) and PPO (Schulman et al., 2017): 100% falls. Under 10 ms observation delay or a locked middle leg, expert and MLP fall in 88–100% of episodes.

C Implementation details

Graph operator. Edges with ≥ 5 synapses are kept. $W \in \mathbb{R}^{N\times N}$ has $W_{\text{post,pre}} = s_{\text{pre}}\, w_{\text{pre}\to\text{post}} / \sum_{\text{pre}'} |s_{\text{pre}'}\, w_{\text{pre}'\to\text{post}}|$, where w is the synapse count and $s \in \{+1, -1\}$ the presynaptic neurotransmitter sign (unknown signs, 4% of edges, are treated as $+1$); rows with no input are left zero. Messages are $M = WH$ (row = postsynaptic neuron).

Roles. Following FlyGM's flow partition, afferent = descending neurons (1,314) + sensory neurons (6,877) = 8,191; intrinsic = VNC-intrinsic + ascending neurons; efferent = motor and other VNC efferents (530). "Direct afferent→efferent mass" therefore includes descending→MN synapses: in the connectome it is 0.132 in total, of which 0.103 is descending→MN and 0.021 leg-sensory→MN.

Bio-wired interface. T1 has 59 actuators: 6 legs × (8 joint targets + 1 adhesion) = 54 are decoded from that leg's motor neurons (subclass fl/ml/hl × side, 58–68 MNs each); the 3 head actuators from neck MNs (subclass nm, 24) and the 2 abdomen actuators from abdominal MNs (ad, 6). Inputs: each leg's 29 channels (11 joint angles, 11 velocities, 3 contact forces, touch, claw position) into that leg's sensory afferents (entry nerve × side, 443–754 neurons); reference trajectory, gravity vector, velocimeter, accelerometer and gyro (467 channels) into descending neurons through a rank-32 encoder; gyro and accelerometer into haltere/wing/notum afferents; abdominal and head joints into abdominal and neck afferents. Every encoder writes two of the d channels of its target nodes; intrinsic neurons receive no direct input. Each decoder is a per-node linear map $d \to 1$ followed by a per-group linear map to that group's actuators.

Training. AdamW, learning rate 10^{-3} with cosine decay, weight decay 10^{-4}, gradient-norm clip 1.0; truncated BPTT through whole 300-step episodes in windows of 32 with state carried and detached between windows; batch 16 episodes; 3,000 updates. An update is one optimiser step on a batch of 16 episodes over one 32-step window (512 samples), with state carried and detached between consecutive windows; 3,000 updates process ≈1.5 M samples, about 13 passes over the 384 training episodes. Validation MSE is computed by running the controller sequentially through whole held-out episodes from a zero state; actions are standardised by the training-set mean and s.d. Frozen core: $\theta = 0.05$, $r_{\max} = 1$, $dt = 2\,\text{ms}$.

Null models. ER: the same number of edges as the thresholded connectome, with the source and target of every edge drawn uniformly over all nodes (self-loops removed, duplicate pairs merged by summing counts) and weights a random permutation of the original synapse counts, so density and weight distribution match the connectome and the node table is unchanged. DP: post-synaptic endpoints are permuted across all edges (weights travel with the pre-synaptic slot), self-loops removed and duplicate edges merged by summing counts, so in-degrees are preserved as a multiset before merging and approximately after (max in-degree 897 → 560). DP_{swap}: Maslov–Sneppen double-edge swaps on the thresholded edge set (5 swaps per edge; exact degrees, no duplicates). Block-preserving DP: the same permutation restricted to edges within each ($\text{group}_{\text{pre}}$, $\text{group}_{\text{post}}$) block. Partial DP: a random fraction of edges is permuted. Synthetic modular: a random 24-module node partition; each edge keeps its source and is re-targeted within the source's module with probability q, otherwise anywhere, targets drawn proportional to in-degree. Synthetic reciprocal: DP followed by re-targeting a random fraction r of edges to make them reciprocal. Interventions are described in Sec. 4.4. Seeds: the seed of a run sets both the sampled null graph and the training initialisation and data order, so reported spreads combine graph-sampling and optimisation variability.

What each intervention preserves. Exact-degree swap (the variant of the 09-18 version), pseudocode: repeat until direct mass reaches 0.26: draw an afferent→intrinsic edge $(a \to i, p)$ and an intrinsic→efferent edge $(j \to e, q)$ uniformly from the ≥ 5-synapse edge set; if $a \neq j$, $i \neq e$ and neither $a \to e$ nor $j \to i$ exists, replace them by $(a \to e, p)$ and $(j \to i, q)$. Weights travel with the source, so every in- and out-degree, the multiset of weights and every out-strength are exact, but the in-strength of e changes by $p - q$ and that of i by $q - p$, and with it their row-normalisation denominators. The sham draws both edges from afferent→intrinsic. Audit on the seed-0 graphs (outputs/swap_audit.json): swap shortcuts change in-strength for 93% of motor neurons (median 8%, 90th percentile 20% of the original), while the intrinsic nodes i that lose an edge change by $q - p$ each (the median over all 12,786 intrinsic nodes is zero because most are untouched); the sham leaves motor-neuron rows untouched and changes intrinsic in-strengths by 0.5% (median). Re-targeting interventions (connectome + shortcuts, DP − shortcuts, and the three source-specific variants) preserve out-degrees only: they change the in-degree of 21–36% of nodes

and inflate motor-neuron in-strength (median +20–37%, 90th percentile up to 8×). Weight-matched swaps (the variant reported in Table 2 and Fig. 2b) draw the partner edge from those with the same synapse count, so $p = q$ and every node's in- and out-strength is exact (audit: zero change for all nodes). They reproduce the unmatched results: shortcuts 0.171 ± 0.001 vs. sham 0.200 ± 0.001 (unmatched swaps: 0.177 ± 0.002 vs. 0.200 ± 0.002); by source, descending 0.203 ± 0.002, same-leg 0.187 ± 0.002, cross-leg 0.175 ± 0.002 (re-targeting: 0.210 ± 0.003, 0.192 ± 0.001, 0.173 ± 0.001). The sham matches the number of swaps and degree preservation but not the perturbation of motor-neuron rows; the source-specific swaps, which perturb motor-neuron rows equally and differ only in the afferent kind, are the control for that.

Confidence intervals. Welch 95% intervals on the mean differences of validation MSE (T3: explained variance), 3 seeds each:

Comparison	difference	95% CI
VNC bio: connectome − DP	+0.030	[+0.020, +0.041]
VNC bio: connectome − ER	+0.044	[+0.034, +0.055]
VNC bio: connectome − (+10% shortcuts)	+0.037	[+0.028, +0.046]
VNC bio: (DP − shortcuts) − DP	+0.018	[+0.005, +0.031]
BANC bio: connectome − DP	+0.031	[+0.021, +0.040]
MaleCNS brain dense: connectome − ER	+0.006	[−0.019, +0.032]
FlyWire dense: connectome − ER	+0.002	[−0.006, +0.010]
FlyWire dense: connectome − DP	−0.011	[−0.029, +0.007]
T3 gait EV: connectome − DP	−0.007	[−0.033, +0.018]

Synapse threshold. The bio-wired gap is insensitive to the edge threshold: at ≥ 3 synapses connectome 0.200 vs. DP 0.173; at ≥ 10, 0.211 vs. 0.174 (one seed each).

Gait dataset (T3). The 16,252 flybody walking snippets are cut into 256-step windows; windows from the same source trajectory are assigned together to train or validation (10% of trajectories held out), so no trajectory crosses the split.

D Learning curves

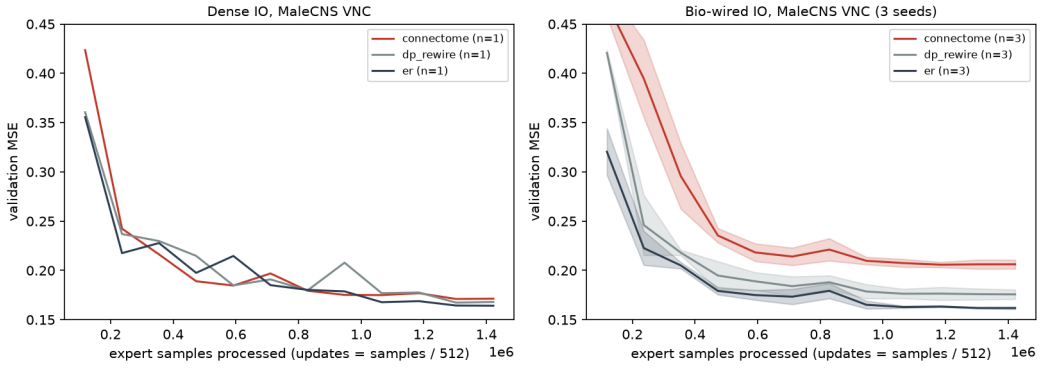


Figure 3: Validation MSE against expert samples processed (512 samples per update). Left: dense IO, the three graphs coincide. Right: bio-wired IO, mean ± s.d. over 3 seeds; the gap between connectome and nulls is present from the first evaluation and persists to the end of the budget, i.e. it is not explained by slower early optimisation within this budget (the connectome curve is still decreasing slowly, so a longer budget could narrow it).

E Accessibility–demand mismatch: statements and proofs

Setting. A structured student has N nodes, a propagation operator $W \in \mathbb{R}^{N\times N}$ with $\|W\|_\infty \le 1$ (row-normalised), input groups g written only onto afferent sets $\mathcal{A}_g$ through encoders $C_g^{(k)}$, output groups e read only from efferent sets $\mathcal{E}_e$ through decoders $D_e^{(k)}$, and K propagation steps per control tick. For an idealised walk-linear student (endpoint encoders and decoders around powers of W; not a linearisation of the gated update), the output of group e is

$$y_e = \sum_g \sum_{k=1}^{K} D_e^{(k)} A_k[\mathcal{E}_e, \mathcal{A}_g]\, C_g^{(k)} u_g, \qquad A_k := W^k, \tag{1}$$

where u_g is the (possibly delayed) input of group g. A teacher acts as $y_e^\star = \sum_g B_{e,g} u_g$ (a linearisation of the expert). Write $\mathcal{R}_K(e,g) = \{\sum_{k\le K} D_k A_k[\mathcal{E}_e, \mathcal{A}_g] C_k : D_k, C_k \text{ arbitrary}\}$ for the realisable maps from u_g to y_e, and $M_K(e,g) = \frac{1}{|\mathcal{E}_e|}\sum_{k\le K} \| \, |A_k|[\mathcal{E}_e, \mathcal{A}_g] \, \|_1$ for the accessible walk mass.

Proposition 1 (representability; linear, uncorrelated inputs). Every map in $\mathcal{R}_K(e,g)$ has rank at most $\rho_K(e,g) := \sum_{k\le K} \operatorname{rank} A_k[\mathcal{E}_e, \mathcal{A}_g]$. In particular, if $A_k[\mathcal{E}_e, \mathcal{A}_g] = 0$ for all $k \le K$ while $B_{e,g} \neq 0$, then $\mathcal{R}_K(e,g) = \{0\}$ and, for any student of the form equation 1, $\mathbb{E}\|y_e - y_e^\star\|^2 \ge \lambda_{\min}\big(\operatorname{Cov}(\tilde u_g)\big)\, \|B_{e,g}\|_F^2$, where $\tilde u_g$ is the residual of u_g after linear regression on the other input groups (and, for a recurrent student, on their history). If u_g is linearly predictable from accessible inputs, $\lambda_{\min}(\operatorname{Cov} \tilde u_g) = 0$ and the bound is vacuous: a student can substitute the proxy.

Proof. $\operatorname{rank}(D_k X C_k) \le \operatorname{rank} X$ for any D_k, C_k, and the rank of a sum is at most the sum of the ranks. If all $A_k[\mathcal{E}_e, \mathcal{A}_g]$ vanish, the student's dependence on u_g is identically zero, so $y_e - y_e^\star = -B_{e,g}u_g + \sum_{g'\neq g} X_{g'} u_{g'}$ for some matrices $X_{g'}$; decompose $u_g = \hat u_g + \tilde u_g$ with $\hat u_g$ the linear prediction from the other groups; the student's terms can absorb $B_{e,g}\hat u_g$ but not $B_{e,g}\tilde u_g$, which is uncorrelated with everything the student sees, so $\mathbb{E}\|y_e - y_e^\star\|^2 \ge \mathbb{E}\|B_{e,g}\tilde u_g\|^2 \ge \lambda_{\min}(\operatorname{Cov} \tilde u_g)\|B_{e,g}\|_F^2$. □

Proposition 2 (bounded-gain mismatch bound). Let $n_e = \dim y_e$ and assume the input groups are uncorrelated. If encoders and decoders are gain-bounded, $\|C_g^{(k)}\|_{1\to 1} \le c$ and $\|D_e^{(k)}\|_{1\to 1} \le d$, then the student's map $X_{e,g}$ from u_g to y_e satisfies $\|X_{e,g}\|_{1\to 1} \le c\,d\,|\mathcal{E}_e|\, M_K(e,g)$, and hence

$$\min_{\text{students}} \ \mathbb{E}\|y_e - y_e^\star\|^2 \ \ge \ \frac{\lambda_{\min}\big(\operatorname{Cov}(u_g)\big)}{n_e} \Big(\|B_{e,g}\|_{1\to 1} - c\,d\,|\mathcal{E}_e|\, M_K(e,g) \Big)_+^2 .$$

Proof. Sub-multiplicativity of the induced ℓ_1 norm gives $\|D_k A_k[\mathcal{E}_e, \mathcal{A}_g] C_k\|_{1\to 1} \le d\,c\,\|A_k[\mathcal{E}_e, \mathcal{A}_g]\|_{1\to 1} \le d\,c\,\| \, |A_k|[\mathcal{E}_e, \mathcal{A}_g]\|_1$ (the maximum column sum is at most the total entry sum); summing over k gives $c\,d\,|\mathcal{E}_e| M_K(e,g)$, the factor $|\mathcal{E}_e|$ undoing the normalisation in M_K. Hence $\|B_{e,g} - X_{e,g}\|_{1\to 1} \ge (\|B_{e,g}\|_{1\to 1} - c\,d\,|\mathcal{E}_e| M_K)_+$ by the triangle inequality. As in Proposition 1, uncorrelated groups give $\mathbb{E}\|y_e - y_e^\star\|^2 \ge \lambda_{\min}(\operatorname{Cov} u_g)\|B_{e,g} - X_{e,g}\|_F^2$, and $\|Y\|_F^2 \ge \|Y\|_{1\to 1}^2 / n_e$ for any Y with n_e rows (Cauchy–Schwarz on the column attaining the maximum column sum). □

Remark (nonlinear students). No bound is claimed for the gated student. Signed cancellation in W^k need not survive node-specific gating: a gate that suppresses one of two cancelling paths transmits a signal the linear power hides, so a nonlinear statement would need absolute path contributions and intermediate Jacobians. The linear statements motivate the statistic; they do not govern the trained controllers.

The bound says that representing a demand $B_{e,g}$ requires gain $\gtrsim \|B_{e,g}\|/M_K(e,g)$. Without a gain bound any non-zero accessibility suffices in principle ($\hat y = d\,\varepsilon\, c\, x$ represents $y = x$ with $d = 1/\varepsilon$), so the statement is about students with an explicit gain bound; weight

decay and gradient clipping do not supply such a bound for the trained models, and small accessibility then forces error or ill-conditioning, which we do not separate. It motivates the mismatch score $S_K = \sum_{e,g} D(e,g) \exp(-M_K(e,g)/m_0)$ used below (a smooth surrogate of $(\|B\| - \kappa M)_+^2$ with $m_0 \sim 1/(cd)$).

Block access across depth (measured). A global spectral statistic does not bound the access of a particular input–output block: a near-unit second eigenvalue can belong to a component unrelated to the block. We therefore measure block access directly. Figure 4 plots the per-step access $m_k(e,g) = \frac{1}{|\mathcal{E}_e|} \| \, |A_k|[\mathcal{E}_e, \mathcal{A}_g] \, \|_1$ of leg motor pools to their own leg's afferents, to other legs' afferents and to descending neurons for $k \le 64$, in the connectome, DP, ER and the cross-leg matched-swap graph, together with shortest-path (hop) statistics from leg sensory afferents to motor neurons. In the connectome, one-step cross-leg access is 10^{-4}, against 0.011 in DP and 0.028 in ER; it rises to 1.4×10^{-3} at $k = 3$ and then decays, as it does in every graph beyond $k \approx 4$ because signed paths cancel (the unsigned operator, dashed, is an upper bound that does not decay). Own-leg one-step access is similar in the connectome and DP (0.012 vs. 0.011; ER 0.027). Shortest sensory→MN routes: own-leg, mean 2.80 hops in the connectome (0.6% direct, 31% within two hops) against 3.0 hops in DP and ER (0.2% direct, 7–11% within two); cross-leg, mean 3.36 hops (none direct, 2.4% within two) against 3.0 hops (0.2%, 7–11%). Matched cross-leg swaps raise one-step cross-leg access to 0.078. Per-step signed access falls below 10^{-3} in every graph beyond $k \approx 6$ and its ordering there reverses, so the deficit is one of immediate access, not of access at every depth. The score uses the cumulative M_K: over 64 steps cross-leg M_{64} is 0.006 in the connectome against 0.032 (DP) and 0.092 (ER), five to sixteen times lower, whereas own-leg M_{64} is not lower (0.085 vs. 0.031 and 0.091). BANC shows the same curves. These are properties of the linear operator, not measurements of transmission through the trained gated controller. The predictor of error is M_K relative to the IO masks, not any spectral quantity: synthetic modular graphs have $|\lambda_2|$ up to 0.91 yet, because their modules are not aligned with the leg masks, cross-leg access is unchanged and no penalty appears.

Corollary 1 (command bandwidth in a fast/slow hierarchy). Suppose the fast layer receives, besides its local inputs, a command $z \in \mathbb{R}^r$ computed by a slow layer from all inputs, and let $\mathcal{G}_e$ be the input groups that output group e cannot access locally ($A_k[\mathcal{E}_e, \mathcal{A}_g] = 0$ for all $k \le K$, $g \in \mathcal{G}_e$). Then the fast layer can realise the teacher only if $r \ge \operatorname{rank} \left[B_{e,g}\right]_{g \in \mathcal{G}_e}$ for every e, the rank of the teacher's dependence on the inputs that are locally inaccessible to e (concatenated over $g \in \mathcal{G}_e$). If the inaccessible sets coincide, $\mathcal{G}_e = \mathcal{G}$ for all e, the condition tightens to $r \ge \operatorname{rank} \left[B_{e,g}\right]_{e,\, g \in \mathcal{G}}$, the rank of the block matrix stacked over all output groups, because the same z serves every e.

Proof. By Proposition 1 the student's map from $(u_g)_{g \in \mathcal{G}_e}$ to y_e receives no contribution through the local walks, so it factors through z and has rank at most r; the teacher's map on the same inputs has rank $\operatorname{rank}[B_{e,g}]_{g \in \mathcal{G}_e}$, and equal maps have equal rank. The statement concerns the instantaneous map; a controller that infers inaccessible inputs from history is outside its scope. When $\mathcal{G}_e = \mathcal{G}$ for all e, the stacked map from $(u_g)_{g \in \mathcal{G}}$ to $(y_e)_e$ also factors through the single vector z, giving the stacked bound. □

Remark (recurrence across ticks). The trained student carries state between control ticks, so over T ticks the walks available to it have length up to KT, each attenuated by the per-tick gain. Gradient truncation does not erase the forward state, so the information available to the controller is that of walks up to length KT over the whole episode (T up to 300), not K; a $K \le 3$ statistic is therefore a proxy for the latency of the shortest usable routes, not a bound on what the recurrent controller can access. Propositions 1–2 apply with K replaced by KT for what the student can access from past observations. What a zero-delay teacher demands is the current observation, which reaches a motor pool no sooner than its shortest usable path allows (time indexing in Appendix G); recurrent depth increases access to older observations without shortening that latency. Figure 4 reports the measured growth of block access with depth, and the delayed-teacher runs of Sec. 4.7 test the latency reading directly.

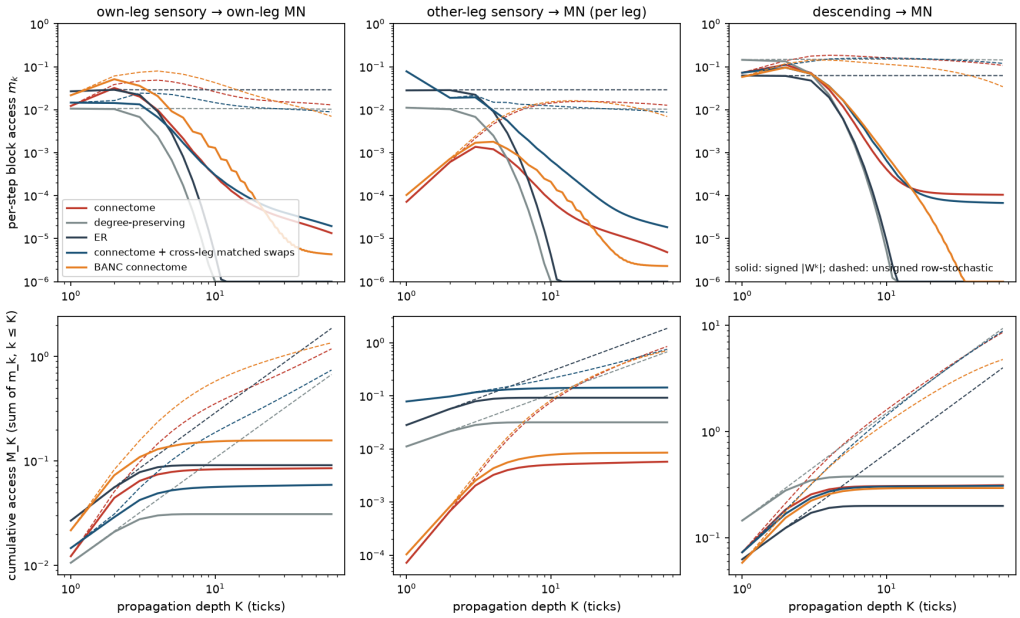


Figure 4: Measured per-step block access m_k of leg motor pools to their own leg's sensory afferents, to other legs' afferents (mean per leg) and to descending neurons, as a function of propagation depth k (one tick per step), for the MaleCNS connectome, its DP and ER nulls, the connectome with matched cross-leg swaps, and the BANC connectome. Top: per-step m_k; bottom: cumulative M_K. Solid: signed operator $|W^k|$; dashed: unsigned row-stochastic operator (an upper bound). Immediate cross-leg access in the real graphs is an order of magnitude below the random graphs and stays five to sixteen times lower cumulatively; per-step signed access collapses in every graph beyond $k \approx 6$.

Scope. These statements are about representability of linearised students with bounded gain; they say nothing about optimisation, and the gated nonlinear extension is loose. Proposition 1 is not the operative regime in our data: in a phase-locked gait every leg's afferents are largely predictable from the others, so the residual $\tilde{u}_g$ is small and a student can use whichever sensory population it reaches as a proxy, which is what Sec. 4.7 finds. Proposition 2 with K read as the delay the teacher tolerates is the version the delayed-teacher runs test, and support only in part (Sec. 4.7). The role of the statements is to motivate a graph statistic (M_K against a predictive D) that is computable before training; the score below is an empirical diagnostic, not a proven lower bound for these experiments.

Statistics. Accessibility $M_K(e, g)$ as defined above, computed on the row-normalised signed W (≥ 5 synapses). Demand $D(e, g) = \|B_{e,g}\|_F^2$ from a ridge regression of the teacher's actions (expert-normalised units) on all 741 observation channels, each standardised to unit variance, with penalty $0.1N$ on $N = 40{,}000$ samples. Coefficient magnitude depends on this scaling, on the penalty and on how redundant information is split across correlated groups, so D is a fitted ranking statistic, not a measure of necessary information. D is predictive: it measures how well group g linearly predicts the teacher's output for e, not whether the teacher causally uses g; Sec. 4.7 shows that this is the right notion, because a student with access to g can use it as a proxy for correlated inputs it cannot reach. Mismatch $S_K = \sum_{e,g} D(e, g) \exp(-M_K(e, g)/m_0)$, with the single scalar m_0 fitted by rank correlation.

What the statistics say. The teacher is global: for the left front leg, demand on the right hind leg's sensors (0.28) exceeds demand on its own (0.18), and demand on the command is 0.34. In the real connectome, motor pools have essentially no access to other legs' afferents (M_1 cross-leg $= 10^{-4}$, $M_3 = 0.002$) but do reach their own leg ($M_1 = 0.012$) and the descending command ($M_1 = 0.073$, $M_3 = 0.26$). Degree-preserving and ER graphs give

cross-leg access of 0.011–0.028 at one step, two orders of magnitude more. Across 27 graph variants (all null models, partial rewirings, both intervention types, sham, and the source-decomposed shortcuts), the mismatch score correlates with the measured validation error at Spearman $\rho = 0.92$–0.93 ($K = 1$–3; m_0 fitted in-sample, one scalar); cross-leg accessibility alone gives $\rho = -0.88$ ($K = 1$) and -0.92 ($K = 3$), so the demand weighting adds little in-sample. Fitting m_0 on the 16 null-model variants alone and scoring the 11 interventions held out gives $\rho = 0.85$ ($K = 1$) and 0.92 ($K = 3$), against -0.74 and -0.86 for cross-leg accessibility alone, so out of sample the weighting does add something, whereas spectral statistics (second eigenvalue of the random-walk operator, number of near-unit modes) correlate more weakly ($\rho \approx 0.6$) and fail on the synthetic modular graphs, which mix slowly yet carry no penalty. The model also predicts the source decomposition of Sec. 4.4: adding cross-leg sensory→MN edges (cross-leg M_1: $10^{-4} \to 0.08$ for the matched swaps) brings the penalty to the DP level, adding descending→MN edges (command already accessible at $K = 3$; M_1 command $0.07 \to 0.50$ in the re-targeting variant) does not, and same-leg shortcuts help partly; and it predicts that three propagation steps per tick should not close the gap (cross-leg M_3 remains 0.002), as observed. Under dense IO the group masks impose no restriction (every input group reaches every efferent through the encoder), so the score is not designed to rank graphs there; the walk masses M_K still differ between graphs, and Table 1 shows they do not translate into a consistent ordering. All of these tests vary M_K under a single teacher; Sec. 4.7 and Appendix G vary the teacher under fixed graphs.

Interpretation. The penalty is an inductive-bias mismatch in the teacher–student sense: the real VNC is organised so that a leg's motor pool sees its own sensors and the descending command within one or two synapses, and other legs only through long, weak routes—a local-reflex architecture—while the slow component of the behaviour is predictable from every leg. Random graphs give every motor pool one-step access to every sensory population (Fig. 4). What remains unexplained by the linear account is the residual of the "DP − shortcuts" graph (0.194 vs. 0.206), the half of the penalty that survives three ticks of tolerated delay or blinding of the other legs (Sec. 4.7), and the contribution of optimisation, which we do not model.

F Relation to the FlyGM setup

FlyGM's code is not yet public, so our replication follows its paper. Unit-weight baselines: VNC bio-wired connectome 0.202 ± 0.004, DP 0.169 ± 0.001, ER 0.161 ± 0.003; FlyWire dense DP 0.181 ± 0.015, ER 0.172 ± 0.010. Two further differences bear on latency: FlyGM injects the current observation into afferent states before synaptic aggregation, so its current observation traverses one more edge per update than ours (Appendix G), and its imitation objective includes distribution matching rather than action regression alone. Shared: FlyWire v783 with the flow partition, signed synapse-count message passing with trainable per-node descriptors and gated node updates, dense IO onto all afferent/efferent nodes, the flybody walker, and imitation from the published expert. Differences: FlyGM uses 32 node channels and four message-passing layers per control step, a 1,253-dimensional observation including binocular vision, learning rate 10^{-4}, unit-weight unsigned baselines, and an RL stage after imitation; we use 16 channels and one propagation per tick (a 32-channel, four-hop configuration on the VNC gives connectome 0.237 vs. DP 0.187, one seed, so the gap is not specific to the smaller configuration), the 741-dimensional proprioceptive observation, batch 8 with the same number of updates, signed-weight baselines by default, and we evaluate the imitation stage only, so we make no claim about FlyGM's closed-loop results; why the original ordering arose can only be settled with the original implementation.

G Demand-side test and frequency decomposition

Teachers. Each teacher is a set of eight per-output-group MLPs (3×512, SiLU; groups: six legs, neck, abdomen) trained by minibatch regression on the expert's 113,721 training samples (6,000 updates of 1,024, AdamW, cosine schedule) in expert-normalised units. The local teacher's leg group e receives that leg's 29 channels and the 467 command channels (496

inputs); the global teacher's every group receives the union of all bio-accessible channels (675 inputs). The mixed teacher averages the two outputs. Relabelled datasets keep the expert's observations and validity masks; students are trained exactly as in T1 (Appendix C) with the expert's normalisation, so all errors are in the units of Table 2. One-leg control: the loss is restricted to the left front leg's nine actuators and the sensory channels of the other five legs are set to their training mean, so no graph can draw on them; batch 8. Causal dependence: for 20,000 samples, group g's channels are permuted across samples and the mean squared change of leg e's output is recorded (Table 3). Delayed and smoothed teachers: the local teacher applied to sensory channels delayed by $L \in \{1, 2, 3\}$ ticks (command channels undelayed), and a causal exponential average of the expert's actions (coefficient 0.2, time constant 5 ticks); batch 8; 3 seeds each. Temporal structure of the expert: its actions have lag-1 and lag-2 autocorrelations of 0.53 and 0.69 (joint velocities 0.36 and 0.72), a period-two alternation at the control rate that carries 24% of action variance; no student reproduces it (high-band error about half the high-band variance for every graph), so the band split below is not a graph property.

Time indexing. At control tick t: (1) the observation o_t is read; (2) the encoder writes $U_t = \mathrm{Enc}(o_t)$ onto afferent nodes; (3) messages $M_t = W H_{t-1}$ are formed from the previous state; (4) $H_t = \mathrm{GRU}([M_t, U_t, \delta], H_{t-1})$; (5) $a_t = \mathrm{Dec}(H_t[\mathcal{E}])$ is emitted and compared with the teacher's target for tick t, $f(o_{t-L})$. Hence o_t enters afferent states at t and reaches a motor-neuron state at $t+1$ through a direct afferent→MN edge and at $t+1+h$ through h interneurons; with n sub-updates per tick (Sec. 4.3) steps (3)–(4) repeat n times with the same U_t on the updated state, so o_t traverses at most $n-1$ edges within the tick (three sub-updates: one- and two-hop paths, not three). In general the target $f(o_{t-L})$ at tick t can draw on paths of at most $L+n-1$ edges from o_{t-L} for $L \geq 1$, and $n-1$ edges for $L = 0$. Batch size: the delayed, one-leg and smoothed students use batch 8; their zero-delay comparator in Table 3 is a batch-8 rerun of the local-teacher students. Delayed local teacher (batch 8, 3 seeds), $L = 0, 1, 2, 3$: connectome 0.143, 0.118, 0.104, 0.106; DP 0.114, 0.095, 0.089, 0.092; ER 0.111, 0.089, 0.084, 0.089. Three sub-updates per tick, $L = 0$ (batch 8, 3 seeds): connectome 0.117 ± 0.012, DP 0.087 ± 0.002, ER 0.082 ± 0.004. One-leg loss with all inputs retained: connectome 0.110 ± 0.008, DP 0.080 ± 0.005, ER 0.082 ± 0.003.

Frequency decomposition. For trained students (expert teacher, 3 seeds) we take the Fourier transform of the prediction error along time within each valid episode and split the power at 50 Hz; by Parseval the two bands add exactly to the total MSE (an earlier moving-average split was not additive).

Table 4: Validation error by Fourier band (bio-wired IO, MaleCNS VNC, expert teacher; mean over 3 seeds). Bands split at 50 Hz within each valid episode and add exactly to the total. Target variance: 0.467 low band, 0.253 high band.

Variant	total	low band ($\leq$ 50 Hz)	high band
connectome	0.206	0.096	0.111
sham swaps (weight-matched)	0.200	0.091	0.110
+ descending→MN only (matched swaps)	0.203	0.096	0.107
+ same-leg sensory→MN only (matched swaps)	0.187	0.079	0.109
+ cross-leg sensory→MN only (matched swaps)	0.175	0.069	0.107
swap shortcuts (weight-matched, all sources)	0.171	0.068	0.104
degree-preserving (DP)	0.176	0.072	0.104
ER	0.162	0.058	0.104

H Evolutionary interface training of an extracted walking module

Because gradient-based imitation of a global teacher may simply be the wrong objective for a shortcut-free circuit, we also optimised the interface with an evolution strategy (OpenAI-ES, 150 generations × 16, antithetic pairs) on the tethered-ball task, keeping the core frozen. We first extracted a walking module from the VNC: all neurons on DNg100→leg-MN paths

Table 5: ES-trained interfaces on the tethered ball (best individual re-evaluated on 8 seeds; mean±s.d. over 3 ES seeds). Standing still scores 0.167.

Rhythm / feature source	fitness	forward (rad/s)	off-axis
pure reflex, no oscillator	0.319 ± 0.172	2.43 ± 2.23	2.67 ± 0.82
engineered tripod clock	0.378 ± 0.072	4.32 ± 0.55	4.62 ± 1.45
connectome walking module	0.036 ± 0.123	3.31 ± 2.43	12.06 ± 3.23
degree-preserving rewired module	0.067 ± 0.035	2.60 ± 1.12	8.96 ± 1.31

of ≤ 3 hops plus the recurrent closure of their interneurons and the leg afferents (13,618 neurons). Under DNg100 drive this module sustains a 4.3 Hz rhythm in all six legs (the feed-forward path set alone does not); a hemilineage lesion screen identifies 08B, 12A, 03B and ascending neurons as necessary for the rhythm and 13A as setting its frequency, yet no subset of "critical" hemilineages is sufficient—the rhythm is a distributed property. The six leg rhythms are, however, nearly synchronous (relative phases 0–0.2 cycles), consistent with connectome CPG simulations lacking sensory feedback (Pugliese et al., 2025); proprioceptive feedback re-organises them (T1 left/right antiphase) but not into a tripod gait.

The ES-trained interface (per-leg readout from causal phase harmonics of the module's motor-neuron rhythm, plus a feedback gain) was compared with the identical readout driven by (i) an engineered tripod clock, (ii) the degree-preserving rewired module, and (iii) no oscillator at all (a per-leg linear reflex from proprioception). Fitness = task reward + capped forward speed − off-axis spin, re-evaluated on fresh seeds (Table 5; 3 ES seeds). Only the engineered clock walks reliably (fitness 0.38 ± 0.07, all seeds ≥ 3.7 rad/s). The pure reflex controller walks in one seed of three. The connectome module (0.04 ± 0.12) is statistically indistinguishable from its degree-preserving rewired version (0.07 ± 0.04) and in two of three seeds scores below standing still (0.167): its synchronous legs push the ball fast but crookedly (off-axis 12±3). The real rhythm therefore confers no advantage over a rewired graph under this evolutionary interface optimisation. This experiment changes optimiser, objective, core and interface representation together relative to T1, so it characterises this frozen-module controller rather than testing whether gradient optimisation per se caused the T1 penalty.

I Latency, memory and activity of candidate low-level controllers

The original motivation for connectome controllers in embodied AI is a "cerebellum-like" layer running at kHz rates under a slow vision-language "brain". Table 6 measures inference cost on one RTX 4090 (PyTorch, sparse CSR operators, batch 1 and 64, 741-dimensional observation). Every candidate, including the 21.5 k-neuron frozen VNC, sustains >1 kHz at batch 1; the frozen connectome cores are the cheapest in memory. Their activity is, however, 34–38% of neurons per step under driven input (20–28% under command-only drive), not the <5% sparsity assumed in neuromorphic energy arguments, and none of the connectome-based controllers produces robust locomotion (Sec. 4.2, Appendix H). kHz operation is therefore feasible but, at present, has nothing robust to run; a hierarchical evaluation (impact recovery, plan–execution fidelity, energy per task) must wait for a low-level connectome controller that beats an engineered clock. Figure 6 shows the spatial pattern of activity under command drive: sparse and structured in the connectome, saturated in the degree-preserving null.

Table 6: Inference cost per 2 ms control step (RTX 4090, fp32). Active fraction: neurons with rate > 0.01 under random driven input.

Controller	params (M)	ms/step (B=1)	max Hz (B=1)	ms/step (B=64)	peak MB (B=64)	active
frozen walking module (13.6 k)	0.16	0.48	2094	0.52	39	0.34
frozen full VNC (21.5 k)	0.29	0.63	1577	0.69	56	0.37
connectome GNN, dense IO (21.5 k, d=16)	1.43	0.12	8184	3.94	1498	–
connectome GNN, bio IO (21.5 k, d=16)	0.72	0.64	1552	4.19	1492	–
MLP 4×512	1.20	0.05	20068	0.06	14	–
GRU 256	0.78	0.03	40294	0.03	12	–

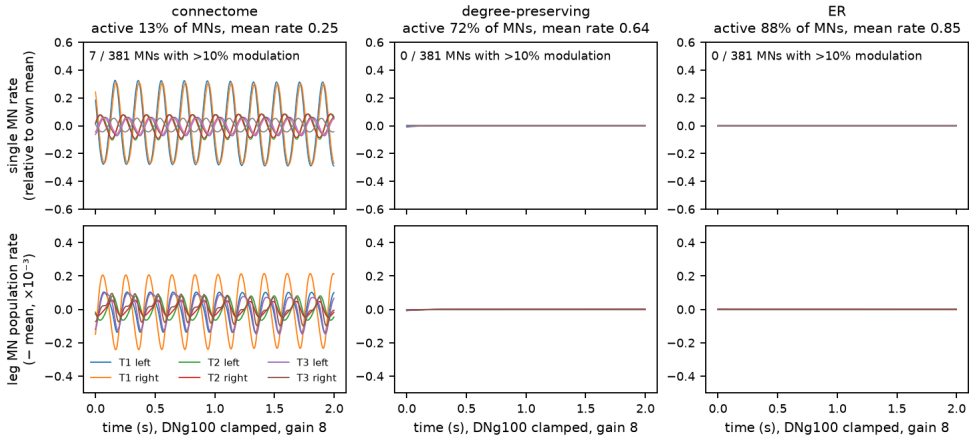


Figure 5: Frozen VNC rate model with DNg100 clamped (gain 8). Top: the eight most rhythmic leg motor neurons of each graph (rate relative to the neuron's own mean); bottom: per-leg MN population rate with its mean removed (same scale in every column). Only the connectome oscillates (≈5 Hz, ±30% in individual MNs); because few MNs are rhythmic the population-level modulation is small. Both nulls sit at fixed points with most MNs saturated (activity maps in Appendix I, Fig. 6).

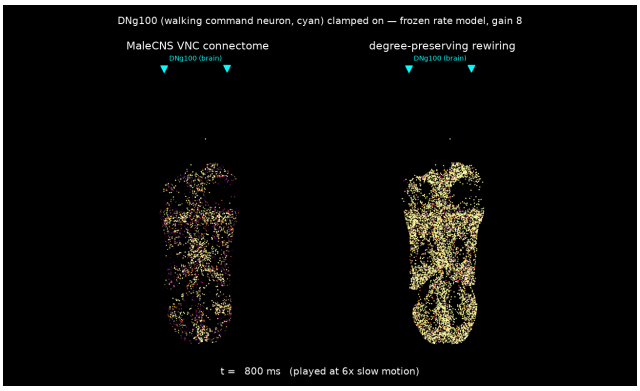


Figure 6: Frozen VNC rate model with DNg100 clamped (gain 8), activity on soma positions: connectome (left) and degree-preserving null (right).

J Second body: an 18-DoF hexapod robot

To test whether the ladder is specific to the fly body, we built a generic hexapod in MuJoCo (torso + six legs × coxa yaw, femur pitch, tibia pitch; position actuators; 200 Hz control; 53-dimensional observation: joint angles and velocities, foot contacts, IMU, speed/yaw command). The teacher is an engineered tripod gait (2.5 Hz, stride scaled by speed command, side-biased by yaw command; 0.09–0.18 m/s, no falls). The bio-wired interface maps each robot leg's 7 sensory channels onto the corresponding fly leg's sensory afferents and each leg's 3 actuators onto that leg's motor-neuron pool; command and IMU enter through descending neurons. Behaviour cloning as in Sec. 3.4 (3 seeds), then closed-loop evaluation on 6 commands × 8 seeds × 600 steps with perturbations.

The imitation ladder reproduces on this body (connectome > DP > ER in error; connectome vs. ER $p = 0.014$), with a much smaller absolute gap because the teacher is a simple periodic function of time and command. Closed-loop performance is saturated except under observation delay, where the connectome-wired controller is best ($p = 0.025$ vs. ER; $p = 0.16$ vs. DP). We report this as a single-task observation.

Table 7: Hexapod: validation MSE and closed-loop speed error ($|\bar{v}_x - v_{\text{expert}}|$, m/s). No variant falls. Mean±s.d. over 3 seeds.

Graph	val MSE	speed err (clean)	push	mass ×1.5	obs. delay 20 ms
connectome	0.0045 ± 0.0003	0.004	0.004	0.007	$\mathbf{0.034 \pm 0.004}$
flow-block DP	0.0038 ± 0.0004	0.004	0.004	0.007	0.041 ± 0.005
degree-preserving	0.0037 ± 0.0004	0.004	0.004	0.007	0.042 ± 0.006
ER	0.0032 ± 0.0004	0.004	0.004	0.008	0.046 ± 0.004
MLP	0.0026	0.010	0.010	–	–

K Frequency decoupling is not free

A fast reflex layer under a slow planner presupposes a low-level controller that tolerates slowly refreshed commands. We tested this with the MLP walking controller (BC+PPO, 500 Hz), refreshing either its command channel (the reference trajectory) or its actions at 500/20/5/2 Hz, on the command grid with perturbations (Table 8). Holding actions (a monolithic slow policy) fails already at 20 Hz. Holding the command while the local loop runs at 500 Hz fails too, though more gracefully: the controller was trained with a dense, fly-frame reference refreshed every 2 ms and does not generalise to stale targets. Decoupling therefore requires the low-level layer to be designed and trained for abstract, slowly varying commands—a narrow descending bottleneck such as the DN interface of Sec. 3.3—rather than obtained by inserting a fast controller under an existing policy. No connectome-based controller enters this experiment, because none walks with a bio-wired interface.

Table 8: Fall rate (300-step episodes, 6 commands × 8 seeds) when the command channel or the action is refreshed at a reduced rate.

Perturbation	refreshed	500 Hz	20 Hz	5 Hz	2 Hz
none	command	0.17	0.67	1.00	0.83
none	action	0.17	1.00	1.00	1.00
lateral push (0.5 BW)	command	0.83	1.00	1.00	0.83
lateral push (0.5 BW)	action	0.83	1.00	1.00	1.00
mass ×1.5	command	0.50	1.00	1.00	1.00
mass ×1.5	action	0.50	1.00	1.00	1.00