

ARTICLE

Methods and preliminary results

MaleCNS Pong Lab: an inspectable connectome-constrained platform for closed-loop sensorimotor experiments

Mirośław Bober

Manuscript revision: 20 September 2026. Independent research draft; not peer reviewed.

The entire paper was automatically generated by GPT-6 Astra.

Connectomic reconstructions constrain neuronal connectivity but do not uniquely determine dynamics, sensory transduction or behaviour. Here we describe MaleCNS Pong Lab, a browser-executed platform that couples a retained adult male *Drosophila* central nervous system graph to an inspectable sensorimotor task. The runtime contains 166,700 annotated neurons, 25,582,938 directed neuron-pair edges and 124,177,617 aggregated synaptic contacts. Sixty-four image channels stimulate annotation-matched photoreceptors; a sparse phenomenological spiking model drives tibia motor pools through an engineered antagonist-muscle and joint plant. A bounded subset of existing sensory-to-motor edges supports event-gated eligibility-trace plasticity. We specify data provenance, numerical update order, sensory coverage, transmitter uncertainty, embodiment, telemetry and paired evaluation. In an archived three-seed pilot, ten seconds of training altered 418–477 distinct weights, but trained and original weights produced identical held-out hit/miss counts within each seed. These observations establish operational integration and synaptic modification, without demonstrating improved control. The platform provides a tractable experimental substrate for testing how anatomical constraints and explicit modelling assumptions jointly shape closed-loop computation.

Keywords: connectomics; computational neuroethology; sparse spiking networks; sensorimotor integration; three-factor plasticity; reproducible simulation.

From reconstructed anatomy to a dynamical experiment

A synaptic connectome specifies a structural substrate for neural computation. It does not specify a unique generative model of behaviour: the same anatomical graph can support different responses under different membrane constants, transmitter-to-receptor mappings, sensory transfer functions and mechanical loads. Consequently, the scientific interpretation of an embodied connectome simulation depends on an explicit account of the transformations between observed anatomy and modelled function.

The adult male *Drosophila* central nervous system resource provides a substrate spanning brain and ventral nerve cord, with cell identities, anatomical annotations, transmitter predictions and reconstructed connectivity.^{1,2} Connectome-constrained modelling has already been used to investigate sensorimotor processing in the fly brain.³ Such work motivates the use of anatomical topology as a prior, while leaving the validity of any new dynamics, task or learning rule to be established independently.

MaleCNS Pong Lab instantiates this problem as a closed loop between a rendered court, a coarse retinal adapter, a persistent spiking network and an artificial actuator. The task is deliberately simple at the environmental level: vertical paddle motion changes the likelihood of contacting a moving ball. Complexity instead resides in the path through the retained neural graph and in the explicit interfaces needed to connect that graph to observations and actions. Both paddles are outputs of one shared nervous-system model. They are not independently trained adversaries.

The system is designed to expose intermediate variables rather than present a successful animation as an explanation.

A user can inspect anatomical position, annotation identity, regional activity, selected-cell membrane state, motor-pool activity, joint state, synaptic update telemetry and exported batch logs. The conceptual contribution is an integrated, inspectable experimental apparatus. Priority over all previous embodied connectome simulations, reconstruction of living-fly physiology and successful reinforcement learning are not established by the present evidence.

We address three questions. First, can released annotations define an auditable sensory-to-actuator interface? Second, can a commodity browser execute a retained graph while maintaining a clear boundary between neural state, environmental state and visualization? Third, does the implemented event-gated plasticity improve performance in a paired held-out comparison? The first two are questions of implementation and experimental access. The third is an empirical efficacy question whose answer cannot be inferred from the first two.

Anatomical substrate and operational interfaces

The retained conversion contains 166,700 neuronal records and approximately 25.6 million directed neuron-pair edges (Table 1). The 124.2 million contact count is a sum over anatomical contacts represented by these edges; it is not the number of independently modelled synaptic compartments. Multiple contacts between a presynaptic and postsynaptic neuron contribute to one edge magnitude. Individual contact locations, dendritic cable properties and receptor distributions are not dynamically simulated.

The conversion preserves annotation fields and records

Table 1. Retained substrate and interface audit. Values describe the locally converted MaleCNS v1.0 runtime used by the archived pilot.

Quantity	Count
Retained neurons	166,700
Directed neuron-pair edges	25,582,938
Summed anatomical contacts	124,177,617
Known soma positions	139,662
Positive / negative sign cells	103,720 / 28,167
Unresolved-sign cells	34,813
Zero-current outgoing edges	5,837,334
Mapped / candidate receptors	5,868 / 6,091
Covered / available image channels	49 / 64
Left / right proprioceptors	735 / 718
Selected tibia motor neurons	33
Mapped dopamine candidates	394
Maximum plastic edge slots	100,000

source hashes. Missing soma coordinates exclude a cell from coordinate-based visualization, not from the computational graph. Known positions are available for 139,662 cells; 27,038 retained cells have no displayed soma. This separation avoids introducing synthetic anatomical coordinates into a dataset that already supports neural identities independently of position.

Functional sign remains a separate modelling decision. The declared convention assigns positive current to acetylcholine, negative current to GABA and negative current to annotated optic sensory histamine sources. The histamine convention is motivated by established photoreceptor transmission through histamine-gated chloride channels.⁴ Neuron-level transmitter identity alone is insufficient to resolve all postsynaptic effects. The 34,813 unresolved-sign cells therefore remain represented, but their 5,837,334 outgoing edges carry no current. The effective transmission graph differs substantially from the stored anatomical graph.

Visual input is distributed to 5,868 of 6,091 annotation-matched receptor candidates, covering 49 of 64 camera channels. Proprioception is mapped to 735 left and 718 right cells. Four motor pools contain 15 left flexors, two left extensors, 14 right flexors and two right extensors. The teaching gateway contains 394 dopamine candidates. These are operational annotation assignments. They do not demonstrate physiological completeness, calibrated visual fields or functional sufficiency of the selected pathways.

A coupled neural, sensory and mechanical system

The application separates three state owners: the neural worker, the body and game controller, and the display. The worker owns the retained graph, membrane and synaptic states, random keys, plastic weights, timing traces, pending teaching events and motor threshold biases. The main thread owns image acquisition, ball and paddle physics, muscle activation, joint integration and event generation. The visualization layer consumes committed states without supplying interpolated frames to the sensory adapter (Fig. 1).

A default batch integrates four neural steps of 4 ms each. Its 16 ms simulated duration also advances the body and task. The resulting image and proprioceptive measurements be-

come the next observation. Ball position and velocity remain environmental variables rather than direct neural inputs. Nevertheless, reward events carry side and action labels, so the complete learning interface includes engineered task information in addition to image and proprioceptive channels.

The network implements a sparse leaky membrane model with a signed synaptic accumulator, stochastic release, refractory state and bounded conductance-like input. The accumulator is not a receptor-resolved biophysical conductance: its positive and negative increments enter the membrane equation through a fixed gain. Sparse activation thresholds and dormant-state handling reduce the number of membranes integrated per tick. These choices create a computational approximation whose accuracy must be tested against an appropriately matched dense reference.

Motor-pool counts are normalized by population size and simulated duration, then passed through rate filters, saturating activation and antagonist mechanics. The plant includes passive restoring forces, joint limits, bilateral gain adaptation and endpoint-bias memory. This produces a dynamical actuator with its own history and feedback. It is not a reconstruction of the fly's full musculoskeletal system. The rendered three-dimensional fly is a visualization of this limited embodiment, with no aerodynamic, locomotor or tissue-mechanical solver.

This decomposition is experimentally consequential. Paddle motion can arise from intrinsic neural drive, homeostatic threshold adaptation, persistent muscle activation, inertia or passive restoring forces. Contact with the ball cannot by itself identify which mechanism generated the movement. Conversely, an apparent lack of movement does not imply neural silence; opposing antagonist activity can cancel at the actuator interface.


Plasticity and preliminary behavioural evidence

Plasticity is confined to a deterministic sample of at most 100,000 existing edges on bounded sensory-to-motor paths. The mask conserves anatomical adjacency: learning changes edge magnitude rather than adding connections. Pre- and postsynaptic timing traces update an eligibility variable; a task event snapshots preceding eligibility before teaching stimulation. A subsequent spike in the annotated teaching population gates each queued event once within a 32 ms window. This is a computational three-factor mechanism, motivated by eligibility-trace approaches to distal reinforcement.⁵


Two distinctions are essential. First, the eligibility decay time is approximately 200 ms, whereas 32 ms is the latency allowed for a teaching spike after an event. These are different clocks. Second, the structural action mask records reachability, not measured causality. A shared upstream edge can belong to multiple motor masks, and increasing an inhibitory magnitude need not promote the motor action with which the edge is associated.

The archived pilot used training seeds 11, 29 and 47, with 10 simulated seconds of training followed by paired 10-second evaluations using learned or original weights. Each evaluation arm reset neural, game and body state and used

a Committed neural-body-task loop



b Event-local teaching credit



c Observation and inspection

Anatomical section | Sampled raster | Selected-cell trace | Fly pose | JSON batch log

Displays consume committed states; interpolated animation does not enter the sensory image.

Figure 1. The implemented closed loop and its measurement boundaries. **a**, Offline conversion supplies annotation metadata and compressed sparse connectivity to the neural worker. Committed court images and joint feedback enter the worker through explicit sensory adapters. Motor-pool rates advance an engineered body and task. **b**, The event-credit path stores pre-teaching eligibility and uses subsequent activity in the annotated teaching population to gate a signed, bounded update once. **c**, Visualization receives state and telemetry without feeding interpolated display frames into the experiment. Arrows specify implemented information flow, not anatomically reconstructed pathways.

Table 2. Complete paired pilot. H/M denotes contacts/escapes on either side. Every segment lasts 10 simulated seconds. $N_{\Delta w}$ counts distinct final weight changes after training.

Seed	Train H/M	$N_{\Delta w}$	Learned H/M	Original H/M	Δh
11	0/5	477	1/4	1/4	0.00
29	3/1	436	1/3	1/3	0.00
47	2/2	418	0/5	0/5	0.00

the same held-out seed within its pair. Plasticity was disabled during evaluation; motor homeostasis and actuator adaptation remained active. The task observations were produced by a software raster using the application’s area-pooling adapter, without browser antialiasing.

Training changed 477, 436 and 418 distinct final edge weights, respectively. The learned and original-weight arms nevertheless produced identical contact and escape counts for every seed (Table 2; Fig. 2). Both arms accumulated two contacts and twelve escapes. The opportunity-weighted aggregate hit fraction is $2/14 \approx 0.143$, whereas the unweighted mean of the three seed-level fractions is 0.15. These summaries differ because event counts differ among seeds.

The cumulative number of update operations was 1,264, 1,017 and 996, respectively. These values exceed the cor-

responding numbers of distinct altered weights because a synapse can be updated repeatedly. No teaching events expired in the archived training segments. This establishes that the implemented credit gateway was exercised in these runs; it does not establish that the credited synapses caused the preceding behaviour.

Only four or five opportunities occurred in each individual segment. The observations therefore support neither an efficacy claim nor a statistical equivalence claim. Identical event counts do not imply identical neural or mechanical trajectories. They also do not establish that extended training, alternative sensory interfaces or different hyperparameters would fail. No new training runs were introduced for this manuscript revision; all behavioural plots and tabulated outcomes derive from the archived pilot report.

Interpretation and experimental reach


The platform turns otherwise implicit model assumptions into accessible experimental variables. Retinal mapping can be interrogated separately from membrane dynamics; motor-pool selection can be distinguished from mechanical gain; anatomical reachability can be compared with causal perturbation. This decomposition makes the system useful as a methods platform even before a behavioural advantage has been demonstrated.

Several quantities limit the interpretation of the present

a Retained-data and interface coverage



b Distinct weights changed



c Held-out hit fraction




Figure 2. Structural coverage and archived pilot outcomes. **a**, Percentages calculated from the structural audit show retained soma availability, receptor mapping, channel coverage and edges with a resolved source sign. Denominators are printed beside each bar. **b**, Distinct final weights changed after training for each seed. **c**, Paired held-out hit fractions for learned and original weights; labels give contact/opportunity counts. Points are individual seeds, not independent collision-level replicates. No uncertainty intervals or fitted trends are implied. The figure contains no simulated or interpolated behavioural measurements beyond the archived records.

implementation. The mapping covers 76.6% of available image bins, although 96.3% of candidate receptors are assigned. A high fraction of mapped cells therefore does not imply uniformly sampled visual space. Source-sign uncertainty prevents current transmission along 22.8% of stored edges. The default maximum plastic mask represents approximately 0.391% of all retained edges. Learning operates within a small, structurally selected subspace of a much larger network, rather than adapting the whole connectome.

Temporal resolution also shapes the experiment. Spike times are quantized at 4 ms, and same-tick timing updates depend on presynaptic-before-postsynaptic ordering. The retinal adaptation coefficient is applied once per batch using a single-tick duration, yielding an effective time constant of 180 ms under default batching. Motor homeostasis likewise updates once per batch without an explicit duration factor. Batch size consequently changes model dynamics as well as execution cost.

The mechanical controller contributes additional memory through activation decay, bilateral gain normalization and endpoint adaptation. These states can maintain movement after the responsible spike burst. The action label used by plasticity reflects antagonist activation rather than complete torque, so passive forces and inertia can dissociate that label from instantaneous paddle displacement. A future causal analysis must observe or intervene on these mechanical states, not treat the plant as a memoryless decoder.

The short pilot cannot determine why weight modification

failed to improve its chosen endpoint. Plausible hypotheses include insufficient experience, eligibility decay relative to event latency, incomplete sensory coverage, nonexclusive action masks, weak influence of selected edges and dominance of background or mechanical adaptations. These are mechanistic hypotheses derived from code and model structure, not experimentally isolated explanations.

An expanded efficacy study should prespecify a training budget, independent evaluation seeds, event-based endpoints and uncertainty estimation. Comparisons should include stationary paddles, matched random motor drive, learning disabled throughout training, and visual inputs that are static or temporally shuffled. Intrinsic drive, threshold homeostasis and mechanical adaptation should be varied separately. Collision events nested within an episode must not be treated as independent biological replicates; repeated seeds probe a computational ensemble rather than between-animal variation.

Validation against biology would require an additional evidential layer. Relevant constraints include cell-type-specific response properties, receptor-dependent sign, synaptic and conduction delays, sensory field geometry, proprioceptor tuning and measured limb mechanics. An anatomical reconstruction alone cannot supply these parameters. The current apparatus should therefore be understood as a connectome-constrained hypothesis generator, with a preliminary computational result, rather than a demonstrated model of fly cognition or learned natural behaviour.

Methods

Model specification and notation

Let $\mathcal{N} = \{1, \dots, N\}$ index retained neurons and let $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ denote directed edges. Throughout, $i \rightarrow j$ denotes a presynaptic neuron i and postsynaptic neuron j . The integer $c_{ij} > 0$ is the aggregated anatomical contact count; $w_{ij} \geq 0$ is its model magnitude; and $s_i \in \{-1, 0, +1\}$ is the assigned source sign. Sign and magnitude are distinct. Superscript 0 denotes the immutable initial magnitude. Neural steps have duration $\Delta t = 4$ ms; a default batch comprises $B = 4$ steps and duration $\Delta T = B\Delta t = 16$ ms. When body equations use seconds, $\delta = \Delta T/1000$ is the numerical duration in seconds. The function $\text{clip}(x, a, b) = \min[b, \max(a, x)]$ bounds a value.

The system is a hybrid discrete-time dynamical model: continuous-valued membrane and mechanical states evolve on a fixed grid, while threshold crossings, collisions, resets and gated updates generate discrete events. The equations below specify the implemented phenomenology. They should not be interpreted as a fitted conductance-based neuron, detailed muscle model or maximum-likelihood estimate of biological parameters. Table 3 collects principal constants and their update cadence.

Acquisition, retention and graph conversion

The offline pipeline consumes four versioned resources: body annotations, consensus neurotransmitter predictions, neuron-pair contact counts and an optic-column workbook. Source addresses and SHA-256 digests are pinned in the acquisition manifest. The converter retains annotation records with non-null superclass, removes duplicate body identifiers, filters connections to retained endpoints and aggregates repeated directed contacts. Original annotation fields remain available in compressed metadata, preserving the distinction between computational indices and anatomical body identifiers.

Target-oriented compressed sparse row storage uses offsets o_j , source indices and contact counts. Incoming edges to neuron j occupy the half-open interval $[o_j, o_{j+1})$. The loader requires $o_0 = 0$, monotone offsets, a final offset equal to the edge count, valid source indices and positive counts. Metadata cardinality must equal the declared neuronal population. It verifies manifest-declared compressed hashes before gzip decompression and rejects malformed arrays rather than silently substituting a synthetic network.

The active engine constructs source-oriented outgoing adjacency at initialization. Target-oriented storage supports incoming inspection and reverse reachability; outgoing adjacency supports event-driven propagation. Construction requires an additional pass over the retained graph and additional edge-sized arrays. Compression reduces transport size, not resident simulation memory. Two 32-bit arrays of length 25,582,938 alone occupy approximately 195.2 MiB; the outgoing representation and edge-to-plastic-slot index require further storage, in addition to metadata and neuronal state. No peak-memory benchmark is claimed here.

The archived audit records 84,450,521 compressed payload bytes, approximately 80.54 MiB, excluding the JSON manifest. The browser payload consists of 27 compressed numeric chunks and one compressed metadata file. These

are derived runtime artifacts, not the original Feather tables or the conversion environment. The loader assembles the full retained graph before starting neural computation. Runtime chunking is a transport and initialization mechanism; it does not implement graph paging during individual neural ticks.

Annotation mapping and anatomical coordinates

The biological adapter identifies candidates through annotation fields describing type, class, modality, subtype and laterality. It does not assign photoreceptor or muscle roles solely from numeric body identifiers. R7/R8 positions are drawn from the optic-column supplement; R1–R6 cells inherit the column of their strongest assigned L1 partner. Where explicit mapping is unavailable, known receptor soma coordinates can supply a separately reported positional proxy. Coordinates are normalized within an eye and projected to the image grid. This preserves a traceable mapping basis but does not reconstruct optical axes, spectral sensitivity or ommatidial receptive fields.

The motor mapper requires motor identity, laterality and tibia flexor/extensor annotation. It selects foreleg scope when resolvable, with a reported aggregate fallback otherwise. The audited population uses foreleg scope. Only 33 of 708 annotation-matched motor neurons enter the four actuator pools; the remaining 675 are not interpreted as paddle actuators. Proprioceptive candidates include chordotonal, campaniform and related labels, pooled by side. Of 1,454 candidates, 1,453 resolve into the two input pools. The mapping is an operational analysis of available annotations, not an independent anatomical reassignment.

Available soma coordinates in 8 nm units are preserved alongside a centered display transform. The anatomy panel uses a common projection and region colours. A neuron lacking a coordinate still participates in connectivity and dynamics. Optional official SWC skeletons are requested for selected-cell inspection; their geometry does not determine transmission delay, membrane capacitance or cable propagation. An SWC visualization and a single-compartment simulation therefore represent different levels of biological detail.

Committed observation and execution order

Each cycle starts from a committed game state. The main thread draws that state to a sensory canvas and pools its pixels. The worker receives image channels, the preceding joint angles and velocities, reward and event labels, event side and action, and an episode seed. It advances neural state for a batch and returns simulated duration, population and motor rates, event-credit diagnostics and display data. The body then advances by that duration, the game advances from the resulting paddle targets, and any resulting event is delivered to the following neural batch.

The sensory observation is held fixed during a neural batch. Display interpolation may generate additional frames for smooth animation, but those frames are not observations used by the brain. Wall-clock latency therefore affects interactive responsiveness without redefining the nominal neural timestep. The shared clock does not make all adaptive mechanisms invariant to batching: retinal adaptation and homeostasis retain batch-specific semantics described below.

Image measurement and retinal adaptation

Each camera crop spans 86% of the court width. Left and right crops overlap and are independently reduced to 8×4 bins, yielding 64 channels. For a pixel with channels $R, G, B \in [0, 255]$, the measured intensity is

$$I = \frac{0.2126R + 0.7152G + 0.0722B}{255}. \quad (1)$$

An integral-image implementation computes area averages over each bin. This is an RGB-weighted software measurement. It is neither spectrally resolved phototransduction nor calibrated retinal irradiance, and it does not explicitly linearize a biological optical signal.

For bin k in batch b , intensity $L_{k,b}$ updates an adaptive baseline and then a bounded drive:

$$\alpha_a = 1 - \exp(-\Delta t / \tau_a), \quad (2)$$

$$a_{k,b} = a_{k,b-1} + \alpha_a (L_{k,b} - a_{k,b-1}), \quad (3)$$

$$d_{k,b} = \text{clip}(0.01 + 0.8L_{k,b} + 4(L_{k,b} - a_{k,b}), 0, 1). \quad (4)$$

The first observation initializes $a_{k,0} = L_{k,0}$. The nominal $\tau_a = 45$ ms is used with the single-tick duration even though the update occurs once per batch. At constant B , the effective time constant in simulated time is $B\tau_a = 180$ ms. The model combines tonic luminance drive with signed temporal contrast, then clips negative or excessive drive. It does not explicitly implement ON/OFF pathways or local motion detectors at the observation interface.

A receptor assigned to channel k receives a per-tick injection opportunity with probability

$$P_k = 220d_{k,b} \frac{\Delta t}{1000 \text{ ms}}. \quad (5)$$

Numerically, the maximum is 0.88 at the default timestep. Trials are Bernoulli events on the integration grid and are subject to refractory eligibility; the nominal 220 Hz input scale is not a guaranteed output firing rate. Input is injected after the current membrane-integration stage, so its effect appears at a subsequent integration opportunity.

Proprioceptive afference

For side $s \in \{L, R\}$, measured angle θ_s and angular velocity ω_s determine the side-pooled rate

$$r_s = \text{clip}\left(25 + 35 \frac{\theta_s}{\pi/2} + 25 \frac{\omega_s}{2.2}, 0, 180\right). \quad (6)$$

The numerical output is in hertz, angles are in radians and velocity in radians per second; the denominators are the model's normalization scales. Each mapped cell on a side receives this same rate converted to a per-tick probability. The interface supplies joint feedback, but not receptor-specific directional tuning, tendon force, cuticular strain or distinctions among leg segments. Variation among such afferent classes is therefore collapsed into an engineered side-wise signal.

Membrane integration and sparse activation

Let g_i denote a signed synaptic accumulator and V_i the membrane-like state. For an active neuron at tick q , accu-

mulator decay precedes membrane integration:

$$g_i^- = g_i^q \exp(-\Delta t / \tau_g), \quad (7)$$

$$\widehat{V}_i = V_{\text{rest}} + (V_i^q - V_{\text{rest}}) \exp(-\Delta t / \tau_m), \quad (8)$$

$$V_i^* = \max\{-80, \widehat{V}_i + \kappa g_i^- + \xi_{i,q}\}. \quad (9)$$

Here $\tau_g = 5$ ms, $\tau_m = 20$ ms, $V_{\text{rest}} = -52$, $\kappa = 0.8$, and the deterministic seeded noise lies approximately in $[-0.015, 0.015]$. Values have a phenomenological voltage interpretation, without fitted capacitance or receptor reversal potentials. Float32 state writes are preserved by the implementation; arithmetic reordering can therefore alter trajectories even when real-valued algebra is unchanged.

A nonrefractory neuron fires when $V_i^* > -45 - b_i$, where b_i is a motor threshold bias and is zero outside selected motor pools. Firing resets V_i to rest, clears g_i and sets a two-tick refractory deadline. Synaptic state decays during refractory ticks, but membrane integration is skipped. Refractory targets are skipped during edge propagation.

Dormant cells do not undergo complete membrane integration each tick. Subthreshold accumulated input is retained in a pending list, decays, and is discarded below magnitude 0.01. Input magnitude at least 0.1 activates a cell. An active cell near rest with negligible residual input can leave the active list. Dormant membranes consequently do not receive continuous noise. Sparse scheduling changes which states are updated and is part of the numerical model, not merely a mathematically exact acceleration of an unrestricted dense simulation.

Synaptic magnitude, release and update order

Initial magnitude scales anatomical counts linearly:

$$w_{ij}^0 = 0.0275c_{ij}. \quad (10)$$

For an internal source spike, transmission is accepted with probability

$$P_{\text{rel}}(i, j) = 1 - \exp(-0.16c_{ij}). \quad (11)$$

An accepted event updates its nonrefractory target as

$$g_j \leftarrow \text{clip}(g_j + \gamma_i s_i w_{ij}, -16, 16), \quad (12)$$

where $\gamma_i = 0.85$ for internal sources. Mapped visual, proprioceptive and teaching sources bypass the release draw and use $\gamma_i = 1$. Unresolved sources have $s_i = 0$. The accumulator bound and contact-dependent release impose nonlinearities, so doubling contact count does not imply doubling the eventual postsynaptic spike response.

Per tick, the engine decays pending accumulators, integrates active cells and identifies spikes, decays plastic traces, processes presynaptic effects and outgoing transmission, then processes postsynaptic traces and spike resets. External injection adds 68.75 to eligible membranes after this stage. The ordering is relevant for simultaneous spikes, refractory checks and the clearing of incoming accumulations on neurons that fire in the same tick. No independent biological axonal-delay distribution is implemented.

Up to 4,096 non-external, sign-resolved neurons are selected for assumed intrinsic stimulation at 1.5 Hz. The reservoir depends on configuration seed and graph ordering. Episode-level random keys make repeated conditions

reproducible within the implementation, but annotation-role invariance under a node permutation does not imply trajectory invariance. Reordering changes reservoir selection and stochastic keys.

For a motor pool with batch-average rate $\bar{r}_m$, each member's threshold bias follows

$$\varepsilon_m = \text{clip}(24 - \bar{r}_m, -96, 24), \quad (13)$$

$$b_i \leftarrow \text{clip}(b_i + 0.012\varepsilon_m, -7.3, 7.3). \quad (14)$$

This adaptation remains enabled when synaptic learning is disabled. It regulates excitability in the model and is not a fitted molecular homeostatic mechanism. Because the update has no explicit batch-duration factor, changing batch size changes its speed in simulated time.

Motor-rate estimation and muscle activation

For pool $\mathcal{P}_m$ with n_m neurons and spike counts z_i during a batch, the reported motor rate is

$$r_m = \frac{\sum_{i \in \mathcal{P}_m} z_i}{n_m \delta}. \quad (15)$$

Normalization prevents the 15-cell flexor pool from receiving greater mean drive solely because it contains more cells than a two-cell extensor pool. It does not equalize sampling variance: a single spike has a larger effect in a smaller pool. Filtering precedes saturation to reduce the impact of that asymmetry.

Filtered rate f_m and muscle activation a_m evolve as

$$f'_m = f_m + (r_m - f_m)(1 - e^{-\delta/0.05}), \quad (16)$$

$$u_m = \min(1, f'_m/45), \quad (17)$$

$$a'_m = a_m + (u_m - a_m)(1 - e^{-\delta/\tau_m^a}), \quad (18)$$

$$\tau_m^a = \begin{cases} 0.026, & u_m > a_m, \\ 0.110, & u_m \leq a_m. \end{cases} \quad (19)$$

Time constants here are in seconds. The rise/fall asymmetry introduces muscle-like persistence but does not model calcium kinetics, cross-bridge cycling or force-length relationships. The four activation states correspond to left flexor/extensor and right flexor/extensor, not a complete set of fly muscles.

Adaptive joint plant and paddle mapping

For a side, antagonist drive is $d_s = a_{s,e} - a_{s,f}$. Its absolute magnitude is filtered with a 0.24 s time constant into D_s . The bilateral reference $\bar{D} = (D_L + D_R)/2$ supplies a gain target

$$\rho_s = \frac{\bar{D}}{\max(0.02, D_s)}, \quad (20)$$

$$G_s^* = \begin{cases} \text{clip}(\rho_s, 0.65, 1.35), & \bar{D} > 0.02, \\ 1, & \text{otherwise.} \end{cases} \quad (21)$$

The actual gain G_s relaxes toward this target with the same 0.24 s time constant. Muscle torque is $T_s^a = G_s d_s$. This normalization mechanically couples the two sides, even apart from their shared connectome, and should be controlled in comparisons attributing bilateral coordination to neural circuitry.

Define a dimensionless endpoint factor

$$h_s = \text{clip}\left(\frac{|\theta_s| - 1.05}{\pi/2 - 1.05}, 0, 1\right). \quad (22)$$

When $h_s > 0$ and outward drive $T_s^a \text{sgn } \theta_s > 0$, endpoint bias β_s approaches that outward drive with coefficient $(1 - e^{-\delta/0.32})h_s$. Otherwise it decays with time constant 0.65 s. Net normalized torque is

$$T_s = T_s^a - \text{sgn}(\theta_s)\beta_s - 0.10\theta_s - 1.05 \text{sgn}(\theta_s)h_s^2. \quad (23)$$

Velocity and angle then advance using a bounded semi-implicit update:

$$\omega'_s = \text{clip}\{\omega_s + (58T_s - 2.2\omega_s)\delta, -3.2, 3.2\}, \quad (24)$$

$$\theta'_s = \text{clip}(\theta_s + \omega'_s \delta, -\pi/2, \pi/2). \quad (25)$$

At a hard endpoint, outward velocity is removed. The coefficients are normalized plant parameters, not measured inertial, damping or stiffness values for an individual fly leg.

For court height H and paddle height P , the top-coordinate target is

$$y_s = \frac{H - P}{2} [1 + \sin(\theta_s)]. \quad (26)$$

With $H = 450$ and $P = 90$, this becomes $180 + 180 \sin \theta_s$. The game separately limits travel to 330 pixels per second. The action label is “up” for $d_s < -0.015$, “down” for $d_s > 0.015$, and otherwise null, with a game-level displacement fallback. It omits passive torque and inertial state; it is therefore a credit-assignment heuristic rather than a complete causal description of movement.

Graph-constrained plastic subspace

Forward breadth-first distances from mapped sensory populations and reverse distances from each motor pool define candidate pathways. An edge $i \rightarrow j$ is eligible for motor group m when

$$d_{\text{sens}}(i) + 1 + d_m(j) \leq 8, \quad (27)$$

with finite reachable distances and a resolved source sign. Traversals exclude unresolved-source transmission. A four-bit mask stores membership for each motor group. Paths can share upstream edges; group membership is neither mutually exclusive nor proof that an edge is functionally influential under the current dynamical state.

A deterministic reservoir samples at most 100,000 candidates to bound plastic-state memory and avoid a simple prefix selection. Sampling is graph-order dependent. The remaining edges retain fixed base magnitudes. The engine stores an edge-to-slot lookup for efficient identification of plastic edges during transmission and source/post mappings for timing updates. Anatomical contact counts remain unchanged throughout training; only model magnitudes in selected slots are altered.

Timing traces and eligibility

Each selected edge has pre-trace p_{ij} , post-trace q_{ij} and eligibility e_{ij} . At each tick these decay multiplicatively:

$$p_{ij} \leftarrow \lambda_p p_{ij}, \quad \lambda_p = 0.995^{40}, \quad (28)$$

$$q_{ij} \leftarrow \lambda_q q_{ij}, \quad \lambda_q = 0.995^{40}, \quad (29)$$

$$e_{ij} \leftarrow \lambda_e e_{ij}, \quad \lambda_e = 0.9995^{40}. \quad (30)$$

On a presynaptic spike, the updates are

$$e_{ij} \leftarrow \text{clip}(e_{ij} - 0.012q_{ij}, -1, 1), \quad (31)$$

$$p_{ij} \leftarrow \min(1, p_{ij} + 1). \quad (32)$$

Postsynaptic spikes are processed subsequently:

$$e_{ij} \leftarrow \text{clip}(e_{ij} + 0.010p_{ij}, -1, 1), \quad (33)$$

$$q_{ij} \leftarrow \min(1, q_{ij} + 1). \quad (34)$$

These saturating traces differ from unbounded linear spike-count traces. Same-tick coincidences depend on the specified ordering because the post update can use the just-incremented pre-trace.

For decay factor λ at timestep Δt , the equivalent exponential time constant is

$$\tau = -\frac{\Delta t}{\ln \lambda}. \quad (35)$$

The defaults imply timing traces of 19.95 ms and eligibility of 199.95 ms. In the absence of replenishment, an isolated eligibility contribution retains approximately 3.35×10^{-4} of its magnitude after 1.6 s. Continuing spikes can replenish eligibility, but the configured decay alone cannot preserve an isolated event across a multi-second rally.

Teaching events and weight modification

The game emits scalar reward $r = +1$ for a contact and $r = -1$ for an escape, with zero otherwise. No continuous alignment reward is used. The event queue captures the current eligibility vector before subsequent teaching stimulation, together with reward, side and action. Each event waits up to 32 ms for any spike in the annotated dopamine population, is applied once if gated, and otherwise expires. This window constrains teaching-response latency; it does not retrospectively extend the eligibility time constant.

Pending positive events request stochastic teaching-cell injection at a nominal 120 Hz. A miss supplies probability-one injection opportunities for the relevant batch through the same population. Actual spiking, refractory state and a nonempty mapped population remain necessary. A teaching spike can gate pending credit even when it was not causally evoked by the reward stimulus; the implementation does not distinguish that causal origin. Multiple queued events may be gated by the same observed teaching activity.

For a side-labelled gameplay event, let $m_{ij} \in \{0, 1\}$ indicate matching structural action membership and let

$$E_{ij} = \max(0, e_{ij}^{\text{snapshot}}). \quad (36)$$

The update is

$$\Delta w_{ij} = 0.08 w_{ij}^0 r E_{ij} m_{ij}, \quad (37)$$

$$w'_{ij} = \text{clip}(w_{ij} + \Delta w_{ij}, 0.25 w_{ij}^0, 2 w_{ij}^0). \quad (38)$$

Unlabelled experimental callers retain signed eligibility instead of this rectification. Bounds scale with the immutable base magnitude; source sign remains fixed. Increasing an inhibitory magnitude increases inhibitory current under the model convention, even if its edge is structurally associated with a selected motor group.

The displayed dopamine pulse is positive telemetry in arbitrary units; it does not determine the reward sign or repeatedly update a stored event. The code's aversive or "pain" fields denote an engineered miss-triggered gateway that reuses dopamine candidates. No anatomically validated nociceptor population, dopamine concentration field, subjective state or biologically calibrated punishment pathway is simulated.

Inspection, visualization and recording

The interface combines a persistent anatomical section, a court with binocular image and receptor-grid views, a procedural three-dimensional fly, a population raster, selected-cell trace and run-status diagnostics. Region-coloured soma positions and sampled connection overlays locate activity within the retained anatomical representation. Display filters affect rendering, not neuronal inclusion in the simulation. The fly's foreleg pose reflects joint/activation state; the WebGL camera transform and mesh do not enter neural or game dynamics.

Quantitative rate and count summaries are computed before display sampling. In the inspected source worker, at most 512 unique batch spike IDs are sent for display; active-edge telemetry is bounded at 384 entries, and the anatomy overlay has a separate 2,048-marker cap. Raster bins collapse IDs into screen columns and anatomical classes and scroll once per batch. Such a raster is an inspection view, not a lossless spike-train dataset. Screen captures from other builds cannot establish the sampling policy of the archived pilot; the source and artifact versions must be checked.

The selected-cell inspector reports anatomical identity, incoming/outgoing degree, recent firing and a finite membrane trace. Optional morphology is fetched separately. Exported JSON records committed batches and event diagnostics, enabling analysis of rates, resets, timing and task outcomes within the resolution actually recorded. It does not implicitly supply all unobserved spikes, release draws or synaptic states. Analyses must distinguish total spikes, unique firing neurons, sampled markers, cumulative update operations and distinct final weight differences.

Pilot protocol and endpoint definition

The archived paired evaluator uses the full installed graph, seeds 11, 29 and 47, and the same CPU neural implementation as the browser. Training lasts 10 simulated seconds with recurrence and plasticity enabled. Learned magnitudes are copied after training. For each training seed s , the held-out seed is $(s \text{ XOR } 0x9e3779b9)$, interpreted as an unsigned 32-bit integer. Each arm resets neural, task and mechanical states; the learned arm then restores only the trained plastic-weight vector. Plasticity is disabled, whereas homeostasis and adaptive mechanics remain active in both arms.

The software camera rasterizes an 800-by-450 court without browser antialiasing and passes its pixels through the same binocular area-pooling routine. Neural batches and

body/game updates share the returned simulated duration. Every reported segment contains 625 default batches. At the fixed endpoint, the last event produced by the game is not flushed through an additional neural batch, so an endpoint reward can remain unapplied.

For seed s and evaluation arm a , contact count $H_{s,a}$ and escape count $M_{s,a}$ give

$$h_{s,a} = \frac{H_{s,a}}{H_{s,a} + M_{s,a}}, \quad (39)$$

$$\Delta h_s = h_{s,\text{learned}} - h_{s,\text{original}}, \quad (40)$$

$$\overline{\Delta h} = \frac{1}{3} \sum_s \Delta h_s. \quad (41)$$

The fraction is undefined if there are no opportunities. The evaluator counts returned events from both sides. The game's separate `state.hits` field counts left contacts only in the inspected source, so it cannot replace this endpoint. A newly constructed game starts with a vertical ball velocity of 120 pixels per second, whereas an explicit reset uses a narrower seeded distribution. The paired evaluator constructs a fresh game in each arm; browser reset protocols are therefore not automatically equivalent.

Seeds and durations were chosen for a short implementation pilot rather than a preregistered confirmatory experiment. No confidence interval, significance test or power calculation is reported. A runner flag requiring at least three seeds and twenty opportunities per segment is only a descriptive sufficiency condition, not a statistical guarantee. The pilot fails that event-count criterion. Raw counts for every seed are retained, allowing alternative descriptive summaries without concealing denominators.

Verification and reproducibility scope

The archived verification record reports 21 Node integration tests and 27 Vitest tests, alongside core assertions and a production build. Tests cover annotation mapping, missing coordinates, corruption detection, recovery, deterministic numerical fixtures, event credit, observation pooling, timing and actuator constraints. These results refer to the archived software record, not a newly executed validation campaign for this typesetting revision. Passing regression tests establishes conformity to specified software behaviour, not physiological realism or convergence under timestep refinement.

The active backend is the CPU sparse engine. A WebGPU capability probe and an exploratory WGSL kernel are present but do not constitute a validated full-graph GPU implementation. Browser responsiveness, wall-clock speed and simulated time are separate quantities. The pilot contains host-specific wall-time measurements, but no controlled cross-device benchmark or real-time guarantee is inferred from them.

The event log should be segmented at resets. A game reset preserves learned weights and much neural state while resetting the task and plant; a brain reset additionally clears learned weights, timing traces, pending credit, homeostatic bias and neural history. Episode-seed changes can reinitialize membrane offsets while preserving other state. These reset semantics must be reported whenever training and evaluation conditions are compared.

Prospective controls and mechanistic tests

The present implementation permits specific follow-up hypotheses, but their outcomes are not supplied by the pilot. A sensory-dependence test should compare intact images with constant-luminance and temporally permuted observations, matching episode seeds and actuator parameters. A plasticity test should separate learning during training from evaluation adaptation and compare learned, original and appropriately shuffled learned magnitudes. Mechanistic attribution requires interventions on selected pathways, beyond a structural reachability mask.

Numerical sensitivity should vary integration timestep, active/pending thresholds, refractory duration and batch size while preserving physical time constants where intended. A dense reference comparison should control floating-point precision and stochastic streams. The existing non-recurrence switch suppresses transmission from non-external sources; it is not a selective removal of graph cycles and should not be described as such.

Mechanical controls should disable bilateral gain adaptation, endpoint bias and threshold homeostasis individually, and compare against fixed or random motor-rate generators. Separate measurements of antagonist activation, net torque and actual displacement would expose credit-label disagreements. Training curves require independent held-out evaluation checkpoints and adequate event counts; the current aggregate pilot cannot support retrospective fabrication of those curves. These experiments would turn plausible interpretations into testable causal claims.

Extended technical specification

Software components and experiment boundaries

The runtime loader performs acquisition of generated browser artifacts, hash verification, decompression, CSR validation and annotation normalization. The biological-interface mapper converts preserved annotations into operational visual, motor, proprioceptive and teaching populations. The active brain engine owns neuronal state, outgoing propagation, the plastic reservoir and motor homeostasis. A separate reward-credit module defines event snapshots, gating and expiry. The synthetic brain implementation used by small tests is distinct from the full MaleCNS engine.

The retinal module owns image area pooling. The body module owns rate filters, muscle activations, mechanical gains, endpoint memory, joint angles and velocities. The game module owns collision detection, serving, bounded paddle motion and task-event labels. The worker protocol separates loading, stepping, selection and resetting. The main thread coordinates these modules using committed states; display modules render anatomy, raster activity, selected morphology and the procedural fly.

This modular decomposition supports local counterfactuals: change a sensory mapping without changing the contact table, change the actuator transfer function without changing synaptic magnitudes, or replace a learning rule while preserving the event stream. Such interventions still require matched initialization and a declared observation protocol. A code-level switch is not itself proof that an experimental contrast isolates one biological mechanism.

Table 3. Principal model constants and cadence. Neural time constants use milliseconds; actuator time constants use seconds. These are implementation parameters rather than estimates from the reconstructed specimen.

Mechanism	Default	Interpretation and update cadence
Neural timestep / batch	4 ms / four ticks	Membrane integration on neural ticks; body/game advance by returned batch duration.
Membrane / accumulator decay	20 ms / 5 ms	Active-state exponential decay; dormant-state rules remain a separate approximation.
Rest / threshold / floor	$-52 / -45 - b_i / -80$	Phenomenological voltage-like units; two-tick refractory deadline.
Contact scaling / recurrent gain	0.0275 / 0.85	Base magnitude per contact and gain for internal-source transmission.
Release coefficient / input bound	0.16 / $[-16, 16]$	Contact-dependent release; signed synaptic accumulator clipping.
Wake / pending discard	0.1 / 0.01	Magnitude thresholds for activating a membrane or discarding residual input.
Visual drive	220 Hz maximum scale	Per-tick injection opportunities; actual spikes remain refractory constrained.
Retinal baseline adaptation	45 ms nominal	Single-tick coefficient applied once per batch; 180 ms effective at $B = 4$.
Proprioception	25 Hz baseline; 180 Hz cap	Uniform side-pooled angle and velocity encoding.
Intrinsic drive	1.5 Hz; up to 4,096 cells	Seed- and order-dependent reservoir outside mapped external populations.
Motor homeostasis	24 Hz target; gain 0.012	Once per batch; threshold bias bounded to ± 7.3 .
Timing / eligibility memory	19.95 ms / 199.95 ms	Multiplicative traces on every neural tick.
Plastic capacity / path bound	100,000 / eight edges	Sample of anatomically existing sensory-to-motor candidate edges.
Teaching response window	32 ms	Maximum waiting time after event snapshot for a teaching-population spike.
Learning gain / magnitude bounds	0.08 / $[0.25, 2]w^0$	Event-local updates relative to immutable base weights.
Rate filtering / saturation	0.050 s / 45 Hz	Pool-normalized rates are smoothed before saturation.
Activation rise / fall	0.026 s / 0.110 s	Asymmetric activation-state relaxation.
Bilateral gain adaptation	0.240 s; $[0.65, 1.35]$	Recent absolute antagonist drive couples mechanical gains across sides.
Endpoint adaptation / decay	0.320 s / 0.650 s	Persistent opposing bias near angular limits.
Angular speed / range	$3.2 \text{ rad s}^{-1} / \pm \pi / 2$	Semi-implicit plant with hard endpoint velocity projection.
Paddle speed cap	330 pixels s^{-1}	Game-level displacement bound after angle-to-target mapping.

Integrity, storage and portability

The runtime’s compressed artifacts are transport representations of metadata and integer arrays. After decoding, offsets, source IDs and contact counts are assembled into typed arrays and validated before engine construction. Serving identical chunks from a local machine or static web host preserves their data content when hashes agree; it does not establish equivalent startup latency or peak memory across browsers. The offline converter and raw source tables are unnecessary for ordinary execution once the derived runtime is prepared.

The manuscript revision was checked against the working-tree implementation and archived evidence. Because historical source hashes can differ after comments or presentation changes, evidence provenance and current source identity are reported separately. Behavioural results remain attached to their recorded hashes. The interface screenshots documented elsewhere in the project are illustrations of user interaction and are not substituted for measurement records in this paper.

Data availability

The source connectome and associated documentation are available from the versioned MaleCNS download service.² The runtime manifest records source-file hashes, conversion output and dataset identity. The pilot’s manifest SHA-256 is the concatenation of the following two lines:

```
0c69aedef3e244c16d983b7a9ea1be01277
af911b8414f3b5c6fa06361a7212ce
```

The local manuscript package retains the original pilot report, structural audit and software-verification record in its evidence directory. It also contains the scripts and vector graphics used for the present descriptive figures. The research-paper directory, including these evidence copies and the PDF, is excluded from GitHub. A public archive of the paper and its evidence has not been deposited. This availability statement distinguishes locally retained evidence from publicly released data.

Code availability

The project source is available at <https://github.com/almera-vs/malecns-pong-lab>. The public source release and the exact source hashes in the archived pilot serve different purposes: the former identifies the project distribution, while the latter identifies the implementation used for the reported measurements. The methods should be reproduced from compatible source, the pinned runtime and the recorded configuration rather than a mutable branch name alone. No separate archival DOI is asserted.

Declarations and provenance

Acknowledgements and declarations

The manuscript author is Mirosław Bober. The original MaleCNS investigators and data contributors are acknowledged through the dataset and primary publication citations. The present study uses released structural data and introduces no new animal experiments. Source data retain their CC BY 4.0 attribution requirements; the software repository carries its existing MIT license. These licenses apply to different materials and are not interchangeable.

Affiliation, correspondence details, funding and competing-interest declarations have not been supplied and remain for author completion. No “none declared” statement is inferred. This document is an independent methods and preliminary-results manuscript. Its journal-inspired typesetting does not indicate submission to, acceptance by or publication in Nature.

Evidence provenance

The three-seed report records training/evaluation durations, seed pairings, complete configuration, event counts, weight-change diagnostics, runtime manifest hash, five central source hashes and Node version. The structural audit additionally records neuronal and edge totals, interface coverage, sign counts, derived time constants, source hashes and host context. The software-verification record documents historical checks. Figure 2 is generated directly from these retained JSON records, without inventing additional trials, trajectories or error estimates.

References

- [1] Berg, S. *et al.* Sexual dimorphism in the complete *Drosophila* male central nervous system connectome. *Cell* (2026). doi:10.1016/j.cell.2026.08.015.
- [2] FlyEM / HHMI Janelia and collaborators. MaleCNS connectome: version 1.0 data downloads. <https://male-cns.janelia.org/download/> (accessed 20 September 2026).
- [3] Shiu, P. K., Sterne, G. R., Spiller, N. *et al.* A *Drosophila* computational brain model reveals sensorimotor processing. *Nature* **634**, 210–219 (2024). doi:10.1038/s41586-024-07763-9.
- [4] Gengs, C. *et al.* The target of *Drosophila* photoreceptor synaptic transmission is a histamine-gated chloride channel encoded by *ort* (*hclA*). *J. Biol. Chem.* **277**, 42113–42120 (2002). doi:10.1074/jbc.M207133200.
- [5] Izhikevich, E. M. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb. Cortex* **17**, 2443–2452 (2007). doi:10.1093/cercor/bhl152.