

Connectome-constrained *Drosophila* neural control and reward-modulated plasticity in a simulated Maxwell-demon particle system

Yash Sharma

yash.sharma@qub.ac.uk

Queen's University Belfast <https://orcid.org/0009-0003-5214-2764>

Method Article

Keywords: Maxwell's demon; information thermodynamics; *Drosophila* connectome; connectome-constrained networks; reward-modulated plasticity; spiking neural networks; closed-loop simulation; particle sorting.

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11441921/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Connectome-constrained *Drosophila* neural control and reward-modulated plasticity in a simulated Maxwell-demon particle system

Yash Sharma

Queen's University Belfast, Belfast, UK

yash.sharma@qub.ac.uk

Abstract

Quantifying how a biologically constrained neural controller acts on a physical environment calls for closed-loop simulations that record neural activity, control decisions and environmental dynamics together. This paper presents a platform where a *Drosophila*-connectome-constrained spiking controller interacts with a simulated two-chamber particle system to sort particles Maxwell-demon style. The platform couples three components: a physics-based ideal-gas simulation, a scripted Maxwell-demon reference controller and a connectome-constrained neural controller whose synapses are updated online by reward-modulated plasticity. Every particle crossing, chamber temperature, gate command and synaptic weight change is logged on a single shared clock, and each run is checked against a 13-point contract; across eight runs of 300–1,200 s the platform logged 41,670 samples and 1,827 crossings. Temperature separation tracks whichever controller holds the gate, the favourable-crossing fraction differs between reward schedules, cumulative synaptic weight change rises in every run, and neuropil-resolved activity is measured throughout. The platform therefore makes it possible to quantify how connectome-constrained, reward-modulated neural control acts on a simulated gas, with the full sensorimotor trace available for analysis.

Keywords

Maxwell's demon; information thermodynamics; *Drosophila* connectome; connectome-constrained networks; reward-modulated plasticity; spiking neural networks; closed-loop simulation; particle sorting.

Abbreviations and symbols

a.u.:	arbitrary units
AI:	artificial intelligence
AL:	antennal lobe
AME:	accessory medulla
BABAM:	Browsable Atlas of Behavior-Anatomy Maps: the adult <i>Drosophila</i> brain compartment atlas of the Branson laboratory, Janelia Research Campus
cg:	the cumulative good crossing counter
cb:	the cumulative bad crossing counter
CSV:	comma-separated values

DEMON:	the gate-controller driver code for the scripted demon
FLY:	the gate-controller driver code for the fly-brain network
dwL1:	the application's cumulative synaptic weight-change estimator (the running $\Sigma \Delta W ; \times 10^4$ a.u. in tables)
dwN:	the estimated number of weights with $ \Delta W $ above 1×10^{-5}
FAFB:	the full adult fly brain electron-microscopy volume underlying FlyWire [16]
HTML:	HyperText Markup Language
JSON:	JavaScript Object Notation
keeper:	the application's real-time crossing evaluator
LA:	lamina
learningSuccess:	the application's one-sided binomial learning KPI
LO:	lobula
LOP:	lobula plate
ME:	medulla
neuroMap:	an interactive graph-visualization tool for the fly neural circuit [35]
PDF:	Portable Document Format
selAll:	the cumulative crossing selectivity (%)
selRecent:	the rolling-20 crossing selectivity (%)
t_sim:	the shared simulation clock on which every row and event is timestamped
ΔW :	synaptic weight change (Equation 1)
W_0:	synaptic weight at the start of a run
W_final:	synaptic weight at the end of a run
$\Sigma \Delta W $:	the cumulative magnitude of synaptic weight change
$\eta(t)$:	the instantaneous learning rate at time t (Equation 2)
lr(t):	the learning-rate schedule evaluated at t
J:	the separation metric: mean absolute left-right temperature difference maintained over the observation window (Equation 3)
T_L(t):	the left chamber temperature at time t
T_R(t):	the right chamber temperature at time t
ΔT :	the left-right temperature difference, $T_L - T_R$
T:	the observation window
t_start:	the start of a run
t_end:	the end of a run
$\eta(\text{jet})$:	the jet duty: percentage of phase rows with the air jet active (Equation 4)
q:	the real-time valence (reward) assigned to each crossing (Equation 6)
dir:	the crossing direction
mean(v^2)_left:	the mean squared particle speed in the left chamber
v^2_{right} :	the squared particle speed on the right at the crossing event
nGood:	the favourable crossing count (Equation 5)
nBad:	the unfavourable crossing count
nCrossings:	the total crossing count

1. Introduction

Maxwell's demon, an agent that sorts gas particles by opening and closing a gate between two chambers, is a paradigm of measurement-based feedback control [1,2]. Information thermodynamics resolved the apparent paradox: measuring, storing and erasing information

carries a minimum heat cost, formalised by Landauer's principle [3,4], and generalised nonequilibrium relations quantify the thermodynamics of feedback and its experimental verification [5–10]. Contemporary work therefore treats the demon as an agent that uses observations to act on a physical system and extract a useful thermodynamic effect [8].

The demon also need not be hand-programmed: reinforcement-learning and machine-learning agents discover measurement-and-feedback policies that optimise the trade-off between measurement-powered cooling and measurement efficiency and have been framed as 'artificially intelligent Maxwell's demons' [10–12]. Optimal finite-time demon protocols have been derived for Langevin systems [12], and neural-network feedback protocols have been learned for fluctuating nano systems [10]. These studies establish that a demon can be learned by an agent that receives observations, selects actions and improves its policy through a reward signal [11,12].

In parallel, the complete adult *Drosophila melanogaster* connectome has made it possible to build controllers whose architecture is constrained by biological wiring rather than arbitrary network design [13–22]. FlyWire provides a synapse-level wiring diagram of the adult fly brain [16,18–20], and recent studies turn such structural maps into function: connectome-constrained networks predict activity across the fly visual system [23], connectome-derived reservoirs support time-series prediction [24], connectome-driven inventories complete visual-system maps [25], connectome-based graph neural networks control whole-body locomotion [26], additional connectome reservoirs [27] and connectome-based orientation-map models [28] extend the approach, and fly-brain structure has been implemented on neuromorphic hardware [29].

These two directions have largely remained separate: intelligent Maxwell-demon models use generic reinforcement-learning agents [11,12], whereas connectome-constrained models address locomotion, sensory processing, reservoir computing or circuit dynamics [23–29]. It remains unclear how a fly-connectome-constrained network behaves when placed in a thermodynamic particle-sorting task and updated by reward.

A *Drosophila*-connectome-constrained neural controller is coupled to a simulated two-chamber gas system to perform a Maxwell-demon particle-sorting task. The controller observes particle crossings and chamber states, commands a gate and air jet, and updates its synapses through reward-modulated plasticity [30–34]; its behaviour is compared against an idealised demon policy, with control alternating every 150 s. A single time-stamped log records every crossing, chamber temperature, controller activity and air-jet state, enabling joint analysis of neural activity, sorting decisions and gas dynamics (**Figure 1**).

Across eight runs of 300–1,200 s the platform recorded 41,670 samples and 1,827 crossings, of which 44.9% were favourable (48.1% under a linear and 33.9% under a random schedule). Demon-controlled phases produced left–right temperature differences of 0.011–0.051 simulation units, whereas fly-brain phases remained near balance (0.001–0.009); cumulative synaptic weight change rose in every run, and all eight runs passed a 13-point validation contract. This paper provides a verified experimental platform for studying how connectome-constrained, reward-modulated neural control behaves in a Maxwell-demon task.

The paper develops this in four steps. Section 2 sets out the simulation framework: the two-chamber gas, the scripted demon and the connectome-constrained fly-brain controller, together with the shared-clock recording and the separation metric J (**Figures 1–4**). Section 3 reports the results in four parts: R1 shows that controller and gas are recorded together, R2 compares temperature separation under each gate controller (**Figures 5–7 and Table 1**), R3 relates the favourable-crossing fraction to reward-modulated plasticity (**Figures 8–9 and Table 2**), and R4 examines neuropil-resolved brain activity (**Figures 10–12**). Section 4 discusses these findings and the limits of the present design, Section 5 concludes, and Section 6 collects the methods.

2. Approach (simulation framework)

The platform couples three components in one closed loop: a simulated two-chamber gas, an idealised Maxwell-demon controller and a spiking neural controller constrained by the adult *Drosophila* connectome [15,18]. An ideal-gas particle system (320 particles) occupies two chambers separated by a sliding gate; an air jet on the left wall dominates the energy input, and temperature probes T_L and T_R monitor the two sides. Two commanders take alternating 150-s turns at the gate: the scripted demon logic, which sorts particles to sustain a temperature separation, and the fly-brain controller. Each controller therefore acts on the same physical system, on the same clock, within the same runs. **Figure 1** shows the apparatus and a recorded session, establishing the simulated particle system the controller has to sort.

The fly-brain controller is a spiking network constrained by the FlyWire FAFB connectome [14,15]: 106,765 neurons and 553,807 chemical edges arranged in five layers. Recorded sensors (T_L , T_R , ΔT , jet and door state) drive the network; its motor layer commands the gate; and its synapses are updated online by reward-modulated plasticity [30–34] rather than offline training. Every crossing is scored in real time and converted into a reward, and a simulated dopamine signal drives a synaptic weight change (Equation 1). The cumulative change is tracked as the $\Sigma|\Delta W|$ estimator (dwL1). The learning-rate schedule is a linear ramp from 0.004 to 0.020 or a random walk of comparable range (Equation 2). **Figure 2** shows the experimental window (the controls, the fly-side gas and the brain-side FAFB connectome with 13 marked neuropils) and shows that the controller is a real connectome-constrained network.

$$\Delta W = W_{\text{final}} - W_0 \quad (1)$$

$$\eta(t) = \text{lr}(t), t \in [t_{\text{start}}, t_{\text{end}}] \quad (2)$$

Because every sensor value, gate command, crossing verdict, brain-activity estimate and weight change is written to one shared t_{sim} clock, the controller’s decisions and the gas’s response can be analysed together at full resolution. All eight runs passed a 13-point validation contract (row cadence, clock span, exact 150-s changeovers, populated tails), so the results rest on a recording of the whole sensorimotor loop that can be checked end to end. **Figure 3** traces that loop, from the sensors through the connectome network and motor layer to the gate, crossing evaluation, reward and plasticity. The separation metric J measures how far the controller holds the two chambers apart. It is defined in Equation 3, illustrated in Figure 4 and used throughout.






Figure 1: The two-chamber apparatus and a recorded session. (a) Annotated schematic of the simulated apparatus: two chambers separated by a sliding gate, an air jet on the left wall as the dominant energy input, temperature probes T_L and T_R , and the fly positioned at the gate. (b) A series of apparatus snapshots across one run showing the fly and the 320 gas particles redistributing between the chambers as the gate opens and closes. This is the simulated particle system the controller must sort, the Maxwell-demon setting the connectome-constrained network acts upon.

a, Experiment window — how one run is set up



b, Fly side — fly (3), demon (2), gate (5), jet (4)



c, Brain side — FAFB v783 connectome, marked neuropils (1)




Figure 2: The experimental window couples a connectome-constrained network to the particle gas. (a) Application controls: recorded sensors, the demon driver, the fly view and the gas and air-jet settings. (b) Fly side: the live two-chamber gas with the fly (3), demon (2), gate (5) and jet (4) labelled. (c) Brain side: the FlyWire FAFB v783 connectome with 13 marked neuropils. The panel shows that a real *Drosophila* connectome, rather than a generic network, is the controller acting on the gas.




Figure 3: The complete closed loop and study overview. (a) Closed-loop controller; (b) the alternating-control protocol (demon and fly-brain alternate every 150 s); (c) apparatus; (d) example recorded sessions; (e) separation metric J and synaptic weight change ΔW per phase; (f) brain maps; (g) neuropil firing rates; (h) J by gate holder; (i) balance versus sorting. The composite traces the full sensorimotor chain (sensors → connectome network → motor layer → gate → crossing evaluation → reward → dopamine → plasticity) on one shared clock.

$$J = \frac{1}{T} \int_0^T |T_L(t) - T_R(t)| dt$$

Figure 4: The separation metric J . J is the mean absolute temperature difference maintained between the two chambers over the observation window T , with $T_L(t)$ and $T_R(t)$ the left and right chamber temperatures, and is used throughout to score how well a controller performs the demon's sorting. The schematic shows the two chambers, the fly at the gate and the temperature probes.

$$J = \frac{(1/T) \int |T_L(t) - T_R(t)| dt}{T} \quad (3)$$

3. Results

3.1 R1: Alternating-control sessions record controller and gas together

Eight independent runs were recorded (300, 600, 900 and 1,200 s), with the gate alternating between the scripted demon and the fly-brain controller every 150 s. Each run logs T_L , T_R , ΔT , the air-jet state, gate commands, every crossing with its real-time valence verdict,

population activity and the synaptic weight-change estimator on one shared clock (schema v2.2; ~6.6 Hz rows; 41,670 samples in total). The full chain runs as a single loop, from the sensors through the FAFB-constrained network (106,765 neurons; 553,807 chemical edges) and its motor layer to the gate, crossing evaluation, reward, dopamine and synaptic plasticity (**Figure 3**); every session is captured in the apparatus view (**Figure 1**). All eight runs passed the full 13-point validation contract, so controller and gas are recorded together in every run.

3.2 R2: Demon-controlled phases show temperature separation; fly-brain phases stay near balance

Across all eight runs, phases in which the demon held the gate sustained a temperature separation (per-run mean $\Delta T = 0.011$ – 0.051 simulation units), whereas fly-brain phases remained near balance (per-run mean $\Delta T = 0.001$ – 0.009 ; Table 1). **Figure 5** shows $\Delta T(t)$ for every run, coloured by the controller holding the gate: the difference climbs under demon control and returns toward balance whenever the fly holds it, at every run length and under both schedules. **Figure 6** summarises chamber balance, and **Figure 7** compares the demon and the network head-to-head at 900 s (top, $\Delta T(t)$; middle, balance; bottom, activity map). Taken together, these panels show that the controller holding the gate determines whether the gas is sorted. The scripted demon sustains the separation, whereas the connectome-constrained network keeps the chambers closer to balance, with a 56–94% smaller mean ΔT in fly-held than in demon-held phases across the eight runs (**Table 1**).




Figure 5: Temperature separation is phase-locked to the controller holding the gate. Left–right temperature difference $\Delta T(t)$ for every run length (300–1,200 s), coloured by which controller held the gate (red = demon, blue = fly). ΔT climbs under demon control and returns toward balance whenever the fly-brain controller holds the gate, at every run length and under both value-increment schedules. This is the sorting pattern the demon role produces in this system.




Figure 6: Chamber temperature balance under each controller. Mean compartment temperatures for demon-held versus fly-held phases at each run length (300–1,200 s), with the separation metric J annotated. Demon-held phases maintain a clear left–right temperature gap, whereas fly-held phases sit near balance, quantifying the temperature sorting that each controller achieves.




Figure 7: Head-to-head comparison of the demon and the connectome-constrained network (900 s). (top) $\Delta T(t)$ through representative 900-s runs for the demon (red) and fly (blue) cases; (middle) the corresponding chamber balance; (bottom) the brain-side activity map. The stacked panels show in one view that the demon sustains the temperature separation while the connectome-constrained network holds the chambers closer to balance, and they link that thermal effect to a neural readout.

Table 1 | Chamber temperatures by gate controller

Duration (s)	Schedule	Demon TL	Demon TR	Demon ΔT	Fly TL	Fly TR	Fly ΔT	ΔT difference (Demon – Fly)	ΔT reduction in fly phase (%)
300	Linear	0.1862	0.1630	0.0232	0.3354	0.3340	0.0014	0.0218	94.0
300	Random	0.1347	0.0939	0.0408	0.0938	0.0900	0.0038	0.0370	90.7
600	Linear	0.2343	0.2140	0.0203	0.3046	0.3026	0.0020	0.0183	90.1
600	Random	0.1413	0.0901	0.0512	0.0977	0.0914	0.0063	0.0449	87.7
900	Linear	0.2425	0.2281	0.0144	0.2943	0.2879	0.0064	0.0080	55.6
900	Random	0.1412	0.0926	0.0486	0.0963	0.0908	0.0055	0.0431	88.7
1200	Linear	0.2520	0.2409	0.0111	0.2840	0.2820	0.0020	0.0091	82.0
1200	Random	0.1395	0.0924	0.0471	0.0979	0.0894	0.0085	0.0386	82.0

T_L and T_R are mean left and right chamber temperatures (simulation units) over the rows in which the named controller held the gate. $\Delta T = T_L - T_R$. Each run alternated Demon and Fly-brain control every 150 s, with demon observation on.

3.3 R3: Favourable-crossing fraction depends on the schedule; synaptic weight change grows in every run

Of 1,827 crossings, 821 (44.9%) were favourable by the application's real-time valence rule (Table 2). **Figure 8 maps brain-side activity as the mean firing rate of the most active neuropils during fly-held versus demon-held phases.** **Figure 9** shows the sorting and plasticity record across all runs: the favourable fraction under each controller and the cumulative synaptic weight change ($\Sigma|\Delta W|$). Because $\Sigma|\Delta W|$ is a running sum, it cannot fall by construction; it rises from 1.17×10^4 to 3.33×10^4 a.u. and tends to be larger in longer runs. The favourable fraction is higher under the linear schedule (48.1% pooled) than under the random schedule (33.9%). This pattern is consistent with reward-modulated learning [30–34], in which reward feedback influences synaptic changes following each sorting decision. These results are consistent with reward feedback shaping the network, but they do not show that it

has learned a reliable sorting strategy. That requires replicated runs, a plasticity-off control and evidence of improvement over time.




Figure 8: Brain-side activity map of the fly network. Mean firing rate of the most active neuropils (AL, AME, LA, LOP, ME, LO) during fly-held versus demon-held phases. The map shows which neuropils are most active while the connectome-constrained controller acts in the demon role, giving a neural readout alongside the sorting results.




Figure 9: Sorting performance and reward-modulated plasticity across all runs. Favourable-crossing fractions under each gate controller and the cumulative synaptic weight change ($\Sigma|\Delta W|$) across all runs (300–1,200 s; linear and random value-increment schedules). $\Sigma|\Delta W|$ rises with run duration ($1.17\text{--}3.33 \times 10^4$ a.u.), while the favourable fraction is higher under the linear (48.1%) than the random (33.9%) schedule.

Table 2 | Crossing classification and plasticity

Duration (s)	Schedule	Crossings (n)	Favourable (n)	Favourable (%)	Learning rate (start → end)	$\Sigma \Delta W $ ($\times 10^4$ a.u.)
300	Linear	194	101	52.1	0.004 → 0.020	1.17
300	Random	65	21	32.3	0.012 → 0.017	1.66
600	Linear	361	159	44.0	0.004 → 0.020	1.57
600	Random	90	35	38.9	0.020 → 0.0135	1.97
900	Linear	277	134	48.4	0.004 → 0.020	1.64
900	Random	106	34	32.1	0.0115 → 0.0165	2.21
1200	Linear	588	289	49.1	0.004 → 0.020	2.62
1200	Random	146	48	32.9	0.018 → 0.0065	3.33

Favourable crossings are those classed as good by the application’s real-time valence rule. $\Sigma|\Delta W|$ is the cumulative sampled plasticity magnitude (the application’s dwL1 estimator). Reward scale was 1.5 and exploration 0.45 in all runs.

3.4 R4: Neuropil-resolved brain activity is recorded in every run

Per-neuropil firing rates were logged throughout all runs. **Figure 10** animates neuropil activity across the runs as BABAM-atlas brain maps, **Figure 11** shows the fly-brain blue-fade animations for the same runs, and **Figure 12** lays out the rolling mean rate at 25/50/75/100% of every session. Activity is dominated by the antennal lobe and optic neuropils (AL, AME, LA, ME, LO, LOP), with a consistent rate ordering across run lengths and schedules. The neuropil visualisation follows the neuroMap approach to fly-circuit graph visualisation [35]. These recordings show the population activity of the connectome-constrained network in every run. They do not by themselves show that reward reshapes this activity; that would need a within-run comparison with plasticity switched off.

NEUROPIIL ACTIVITY — all runs, synchronised by session progress




Figure 10: Neuropil activity across all runs (animation). BABAM-atlas brain maps synchronised by session progress, animating how the connectome-constrained network’s population activity evolves within each run.

Fly-brain blue-fade animations — all runs




Figure 11: Fly-brain blue-fade animations across all runs. Animated brain views for the linear and random value-increment schedules, showing the fly-brain activity that the controller generates while acting in the demon task.

Neuropil activity through all eight runs

BABAM-atlas compartments \u00b7 5-sample rolling mean rate \u00b7 shared 99.3rd-percentile cap (113.4) \u00b7 columns = session progress




Figure 12: Neuropil activity through all eight runs. BABAM-atlas brain maps of the rolling mean firing rate at 25/50/75/100% of session progress for each run length and schedule. Activity is dominated by the antennal lobe and optic neuropils (AL, AME, LA, ME, LO, LOP), with a consistent rate ordering across the eight runs.

4. Discussion

This study contributes a platform, not a demonstration of effective demon behaviour. The platform couples a connectome-constrained *Drosophila* controller, a scripted demon and a simulated two-chamber gas in one closed loop, logged on a shared clock. Three descriptive

observations follow. Temperature separation tracks whichever controller holds the gate: the scripted demon maintains it, while the fly-brain controller lets it fall toward balance. The favourable-crossing fraction differs between the linear and random schedules, and cumulative synaptic weight change rises under both. Neuropil activity follows a consistent ordering across all eight runs.

The fly-brain controller is therefore a weaker sorter than the scripted demon in this setup. This should not be interpreted as evidence against connectome constraints. The experimental design cannot isolate the contribution of the connectome because no size-matched or connectome-shuffled network was tested. Furthermore, the reward schedule and learning-rate schedule were manipulated together, preventing their individual effects from being distinguished. The comparison with reinforcement-learning demons [10–12] is also indirect: those studies optimise an explicit thermodynamic or control objective, whereas reward is the application’s real-time valence heuristic. The favourable fraction is therefore not evidence that the network has learned a sorting policy.

Three limitations bound interpretation. First, the air jet is confounded with the controller. In the 600-s session log, the jet was never active during fly-held rows and was active in about 65% of demon-held rows, so the temperature contrast between controllers is jet-assisted and cannot be attributed to the gating policy alone. A jet-symmetry control, which holds the jet identical across phases, is implemented in the application but has not been run. Second, the study has eight runs, one session per condition, and no plasticity-off control, so no comparison in this paper is tested statistically. Third, temperatures are in simulation units, the favourable-crossing rule is a valence heuristic rather than a measure of thermodynamic work, and $\Sigma|\Delta W|$ is a sampled magnitude estimate. The paper therefore reports no energy, work or entropy balance; its thermodynamic framing motivates the design but is not measured here.

5. Conclusion

This paper verified a Maxwell-demon platform in which a *Drosophila*-connectome-constrained spiking controller, a scripted demon and a simulated two-chamber gas interact in one closed loop, recorded on a shared clock. Across eight runs of 300–1,200 s the platform logged 41,670 samples and 1,827 crossings and passed a 13-point validation contract. Within these runs, temperature separation follows whichever controller holds the gate, the favourable-crossing fraction differs between reward schedules (48.1% linear, 33.9% random), and cumulative synaptic weight change rises in every run. Neuropil activity is consistent across runs.

The contribution is a measurement tool, not a claim that the connectome-constrained controller sorts well or has learned to act as a demon. By coupling a biologically constrained controller to a simulated thermodynamic environment and logging the full sensorimotor loop, the platform makes the controller’s effect on the gas measurable and auditable. The next steps are a jet-symmetry run, replicate sessions for each condition, plasticity-off and connectome-shuffled controls, and jet duty reported for every run. The same platform could then compare alternative connectome-constrained learning rules [30–34] and, eventually, closed-loop tests on neuromorphic hardware [29].

6. Methods

Platform and simulation

An ideal-gas particle system (320 particles) occupies two chambers separated by a sliding gate; an air jet on the left wall is the dominant energy input. Two gate controllers alternate every 150 s: scripted demon logic (Maxwell-demon sorting) and a spiking neural controller constrained by the FlyWire FAFB connectome (106,765 neurons; 553,807 chemical edges) with a motor layer commanding the gate. Demon observation (the Stage-3 registered arm) was on in all runs. Reward scale 1.5; exploration 0.45.

Runs

Eight independent sessions (fresh page loads): durations 300, 600, 900 and 1,200 s, each under a linear value-increment schedule (learning-rate ramp 0.004 $\rightarrow$ 0.020) or a random walk of comparable range. All runs: schema v2.2; 13/13 validation checks passed; 41,670 logged samples at ~ 6.6 Hz; 1,827 crossings.

Recording and schema

Every row of fly_log_<stamp>.csv (94 columns) carries the chamber temperatures T_L/T_R and their difference ΔT , air-jet state, the controller holding the gate (driver), cumulative crossing counters, the keeper's real-time verdicts, selection statistics (selAll, selRecent), plasticity estimators (dwL1, dwN and per-step deltas), the instantaneous learning rate, neuromodulator levels, population spike counts, and the top-ranked neuropil rates, all on the shared t_{sim} clock. events.csv logs gate_open, gate_close, crossing (direction, verdict, q, speed), startle and jet events; keeper_performance.csv holds the authoritative per-run counts (nCrossings, nGood, nBad, selection statistics, z-score, one-sided p, learningSuccess); brain_map_rates.csv holds per-neuropil firing rates. The clock contract (t_{sim} span equals the armed length; exact 150-s changeovers) is verified per run.

Metric definitions

Phase split: rows with driver == DEMON vs driver == FLY. Mean $T_L/T_R/\Delta T$ per phase: arithmetic mean over the phase's rows. Jet duty is given by Equation 4. Crossings and favourable counts: from keeper_performance.csv (the cumulative cg/cb counters can undercount by one at the tail). Favourable % is given by Equation 5. $\Sigma|\Delta W|$ is taken as the last non-null dwL1 value of the run ($\times 10^4$ a.u. in tables). Separation metric: J is the mean absolute temperature difference over the observation window (Equation 3); higher J means the chambers are held further apart, the quantity the demon's sorting sustains. Favourable-crossing rule: the application's real-time valence heuristic (Equation 6), in which each crossing is classed good or bad at the event.

$$\eta(\text{jet}) = \frac{100 \text{ (rows with jet} > 0\text{)}}{\text{phase rows}} \quad (4)$$

$$\text{favourable \%} = \frac{100 \text{ nGood}}{\text{nCrossings}} \quad (5)$$

$$q = \text{dir} \cdot \frac{(\text{mean}(v^2_{\text{left}}) - v^2_{\text{right}})}{\text{mean}(v^2_{\text{left}})} \quad (6)$$

Statistics

Descriptive only: $n = 8$ runs ($n = 1$ session per condition). No population-level inference is performed; pooled percentages are counts over runs, and no hypothesis test is reported. All temperatures are simulation units (uncalibrated), and plasticity is reported as the sampled dwL1 estimator.

Software

Analyses and figures were produced in Python [36] using NumPy [37], Matplotlib [38], pandas [39] and SciPy [40]; brain-map images were composed with Pillow [41], structural-similarity checks of the brain maps used scikit-image [42], PDF export and figure verification used PyMuPDF [43], and connectome edge and synapse tables were processed as Apache Arrow (Feather) files with PyArrow [44]. The interactive 3D supplementary viewer is a self-contained HTML/JavaScript application built with Three.js [45] and D3.js [46]. Neuropil activity maps were drawn on the BABAM adult-brain compartment atlas [47] (Jenett et al. compartment nomenclature [48]) from the Branson laboratory, Janelia Research Campus; the experimental platform additionally streams live telemetry over WebSockets and uses MuJoCo [49] for fly-body embodiment checks. All session logs, event tables and summary statistics were recorded as CSV and JSON [50] documents, with HTML [51] reports for browsing; the manuscript and figure panels were assembled programmatically with lxml [52] and python-docx [53].

AI assistance: the analysis and data reduction and manuscript structuring were drafted with the assistance of AI. All reported quantities were computed from the recorded session logs by rerunnable code, and each generated figure can be reproduced from the released data.

Data and code availability

The complete simulation code (3D fly-particle environment, connectome brain worker and experiment pipeline), the analysis and figure scripts, and the complete session data logs analysed in this study are available in the project repository (version Ver4, Newest-DataLogs):

[https://github.com/Yash9808/Maxwell-s-Demon-Inside-the-Fly-Brain/tree/Ver4\(Newest-DataLogs\)](https://github.com/Yash9808/Maxwell-s-Demon-Inside-the-Fly-Brain/tree/Ver4(Newest-DataLogs))

References

1. Thomson, W. Kinetic theory of the dissipation of energy. *Nature* 9, 441–444 (1874).
2. Szilard, L. Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen. *Z. Phys.* 53, 840–856 (1929).
3. Landauer, R. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* 5, 183–191 (1961).
4. Bennett, C. H. The thermodynamics of computation—a review. *Int. J. Theor. Phys.* 21, 905–940 (1982).
5. Sagawa, T. & Ueda, M. Generalized Jarzynski equality under nonequilibrium feedback control. *Phys. Rev. Lett.* 104, 090602 (2010).
6. Toyabe, S., Sagawa, T., Ueda, M., Muneyuki, E. & Sano, M. Experimental demonstration of information-to-energy conversion and validation of the generalized Jarzynski equality. *Nat. Phys.* 6, 988–992 (2010).

7. Koski, J. V., Maisi, V. F., Sagawa, T. & Pekola, J. P. Experimental observation of the role of mutual information in the nonequilibrium dynamics of a Maxwell demon. *Phys. Rev. Lett.* 113, 030601 (2014).
8. Parrondo, J. M. R., Horowitz, J. M. & Sagawa, T. Thermodynamics of information. *Nat. Phys.* 11, 131–139 (2015).
9. Manzano, G., Subero, D., Maillet, O., Fazio, R., Pekola, J. P. & Roldán, É. Thermodynamics of gambling demons. *Phys. Rev. Lett.* 126, 080603 (2021).
10. Whitelam, S. Demon in the machine: learning to extract work and absorb entropy from fluctuating nanosystems. *Phys. Rev. X* 13, 021005 (2023).
11. Erdman, P. A., Czupryniak, R., Bhandari, B., Jordan, A. N., Noé, F., Eisert, J. & Guarnieri, G. Artificially intelligent Maxwell’s demon for optimal control of open quantum systems. *Quantum Sci. Technol.* 10, 025047 (2025).
12. Kamijima, T., Takatsu, A., Funo, K. & Sagawa, T. Optimal finite-time Maxwell’s demons in Langevin systems. *Phys. Rev. Research* 7, 023159 (2025).
13. Robie, A. A., Hirokawa, J., Edwards, A. W. et al. Mapping the neural substrates of behavior. *Cell* 170, 393–406.e28 (2017).
14. Zheng, Z., Lauritzen, J. S., Perlman, E. et al. A complete electron microscopy volume of the brain of adult *Drosophila melanogaster*. *Cell* 174, 730–743.e22 (2018).
15. Scheffer, L. K., Xu, C. S., Januszewski, M. et al. A connectome and analysis of the adult *Drosophila* central brain. *eLife* 9, e57443 (2020).
16. Dorkenwald, S. et al. FlyWire: online community for whole-brain connectomics. *Nat. Methods* 19, 119–128 (2022).
17. Winding, M., Pedigo, B. D., Barnes, C. L. et al. The connectome of an insect brain. *Science* 379, eadd9330 (2023).
18. Dorkenwald, S., Matsliah, A., Sterling, A. R. et al. Neuronal wiring diagram of an adult brain. *Nature* 634, 124–138 (2024).
19. Schlegel, P., Yin, Y., Bates, A. S. et al. Whole-brain annotation and multi-connectome cell typing of *Drosophila*. *Nature* 634, 139–152 (2024).
20. Lin, A., Yang, R., Dorkenwald, S. et al. Network statistics of the whole-brain connectome of *Drosophila*. *Nature* 634, 153–165 (2024).
21. Pospisil, D. A. et al. The fly connectome reveals a path to the effectome. *Nature* 634, 201–209 (2024).
22. Shiu, P. K. et al. A *Drosophila* computational brain model reveals sensorimotor processing. *Nature* 634, 210–219 (2024).
23. Lappalainen, J. K., Tschopp, F. D., Prakhya, S. et al. Connectome-constrained networks predict neural activity across the fly visual system. *Nature* 634, 1132–1140 (2024).
24. Suárez, L. E. et al. Connectome-based reservoir computing with the conn2res toolbox. *Nat. Commun.* 15, 656 (2024).
25. Nern, A. et al. Connectome-driven neural inventory of a complete visual system. *Nature* 641, 1225–1237 (2025).
26. Jin, Z. & Sui, Y. Whole-brain connectomic graph neural networks enable whole-body locomotion control in *Drosophila*. *NeurIPS 2025 Workshop: Data on the Brain and Mind* (2025).

27. Costi, L., Hadjiivanov, A., Dold, D., Hale, Z. F. & Izzo, D. The *Drosophila* connectome as a computational reservoir for time-series prediction. *Biomimetics* 10, 341 (2025).
28. Liew, J. N., Lin, S., Chen, B., Zhu, X., Zhang, W. & Hu, X. Connectome-based modelling reveals orientation maps in the *Drosophila* optic lobe. *Adv. Neural Inf. Process. Syst.* 38, 182832–182859 (2025).
29. Wang, F., Theilman, B. H., Rothganger, F., Severa, W., Vineyard, C. M. & Aimone, J. B. Neuromorphic simulation of *Drosophila melanogaster* brain connectome on Loihi 2. *arXiv:2508.16792* (2025).
30. Florian, R. V. Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. *Neural Comput.* 19, 1468–1502 (2007).
31. Frémaux, N. & Gerstner, W. Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Front. Neural Circuits* 9, 85 (2016).
32. Schmidgall, S. & Hays, J. Meta-SpikePropamine: learning to learn with synaptic plasticity in spiking neural networks. *Front. Neurosci.* 17, 1183321 (2023).
33. Fisher, Y. E., Marquis, M., D'Alessandro, I. & Wilson, R. I. Dopamine promotes head direction plasticity during orienting movements. *Nature* 612, 316–322 (2022).
34. Davidson, A. M. et al. Dopamine-dependent plasticity is heterogeneously expressed by presynaptic calcium activity across individual boutons of the *Drosophila* mushroom body. *eNeuro* 10, ENEURO.0275-23.2023 (2023).
35. Sorger, J. neuroMap — interactive graph-visualization of the fruit fly's neural circuit. Diploma thesis, Vienna Univ. Technol. (2013).
36. Python Software Foundation. Python (2025). <https://www.python.org/>.
37. Harris, C. R., Millman, K. J., van der Walt, S. J. et al. Array programming with NumPy. *Nature* 585, 357–362 (2020).
38. Hunter, J. D. Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* 9, 90–95 (2007).
39. The pandas development team. pandas (Zenodo, 2025). <https://doi.org/10.5281/zenodo.3509134>.
40. Virtanen, P., Gommers, R., Oliphant, T. E. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* 17, 261–272 (2020).
41. Pillow development team. Pillow: the friendly Python Imaging Library fork (2025). <https://python-pillow.org/>.
42. van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J. et al. scikit-image: image processing in Python. *PeerJ* 2, e453 (2014).
43. PyMuPDF contributors. PyMuPDF (2025). <https://pymupdf.readthedocs.io/>.
44. Apache Software Foundation. Apache Arrow (2025). <https://arrow.apache.org/>.
45. Three.js contributors. Three.js (2025). <https://threejs.org/>.
46. Bostock, M., Ogievetsky, V. & Heer, J. D³ data-driven documents. *IEEE Trans. Vis. Comput. Graph.* 17, 2301–2309 (2011).
47. Branson lab, Janelia Research Campus. BABAM!: Browsable Atlas of Behavior-Anatomy Maps (2025). <https://kristinbranson.github.io/BABAM/>.
48. Jenett, A. et al. A GAL4-driver line resource for *Drosophila* neurobiology. *Cell Rep.* 2, 991–1001 (2012).
49. Todorov, E., Erez, T. & Tassa, Y. MuJoCo: a physics engine for model-based control. in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems* 5026–5033 (2012).

50. ECMA International. ECMA-404: The JSON Data Interchange Syntax, 2nd ed. (2017). <https://ecma-international.org/publications-and-standards/standards/ecma-404/>.
51. WHATWG. HTML Living Standard (2025). <https://html.spec.whatwg.org/>.
52. lxml project. lxml (2025). <https://lxml.de/>.
53. python-docx developers. python-docx (2025). <https://python-docx.readthedocs.io/>.