Optogenetic Sweet Neuron Activation Supports Spatial Place Conditioning in Drosophila melanogaster, with Preliminary Evidence for Goal-Directed Outcome Sensitivity

NEU 350 Independent Project — Final Lab Report

Ajay Donthula | May 5, 2026

Abstract

Whether *Drosophila melanogaster* can form an appetitive spatial memory using a non-caloric, optogenetically delivered reward, and whether such a memory remains sensitive to the action–outcome contingency, has not been directly addressed in freely walking flies. Here, sweet-sensing gustatory receptor neurons were activated optogenetically using Gr64f-Gal4 > UAS-CsChrimson, and 617 nm illumination was delivered in closed loop when flies entered a preassigned reward-arm ROI in a V-shaped maze to test whether localized sweet-neuron activation supports place conditioning. Ten flies were recorded — six retinal-fed (ATR+) experimental flies and four non-retinal-fed (ATR–) controls — and three of the six ATR+ flies were carried into a Phase 2 manipulation in which the reward was either omitted (extinction, n = 1) or reassigned to the opposite arm (reversal, n = 2). Bonsai-tracked centroid trajectories were used to compute two preference indices (entry-count PI; dwell-time PI), time-resolved learning curves, and arm-specific movement speed. In the Phase-1-only ATR+ subgroup (n = 3) versus ATR– controls (n = 4), ATR+ flies developed a positive dwell-time preference for the stimulated arm (PI = +0.34 ± 0.26) while ATR– controls did not (PI = –0.22 ± 0.32; Mann–Whitney U = 11.0, p = 0.057, r = 0.83), and individual learning curves were well fit by sigmoidal or Gaussian functions (R^2 = 0.64–0.98). When the reward contingency was altered in Phase 2, dwell-time PI fell in every fly tested (ΔPI = –0.02 to –0.18), indicating that the learned preference weakened after omission or reassignment of the reinforcer. These results indicate that optogenetic Gr64f activation is sufficient to drive measurable spatial conditioning within a single 10–15 minute session, and that the resulting preference is sensitive to altered contingency — consistent with outcome sensitivity, a prerequisite for goal-directed control, though the small Phase 2 cohort prevents a definitive distinction between goal-directed and habit-like responding. The small cohort precludes strong claims of statistical significance, but the convergent direction of effects across PI, learning-curve, and locomotor measures supports the central hypothesis.

Introduction

Drosophila melanogaster has long served as a tractable model for the cellular and circuit logic of associative learning. Its compact nervous system, defined neuromodulatory projections, and genetic accessibility make it possible to causally link single populations of neurons to specific behavioral outputs

(Tully and Quinn, 1985; Kahsai and Zars, 2011). Appetitive learning in flies depends on coincidence detection in the mushroom bodies, where Kenyon cell synapses are modulated by reward signals carried by dopaminergic neurons of the protocerebral anterior medial (PAM) cluster (Aso et al., 2014; Kim et al., 2007). Sweet taste drives PAM activity through a sensory pathway that begins at peripheral gustatory receptor neurons (GRNs) labeled by the Gr64f promoter, which respond broadly to sugars and project centrally via the gnathal ganglion. Sweet taste is detected by gustatory receptor neurons expressing members of the Gr5a/Gr64 receptor family, with Gr64f functioning broadly with other sugar receptors in sugar detection (Dahanukar et al., 2007; Jiao et al., 2008; Scott, 2018; Kirkhart and Scott, 2015).

Optogenetics has made it possible to substitute defined neural activation for natural reward delivery, thereby decoupling the reinforcing signal from caloric content. The red-light–gated cation channel CsChrimson, expressed in sweet GRNs via Gr64f-Gal4, is sufficient to drive immediate appetitive responses and to reinforce associative memory (Klapoetke et al., 2014; Yapici et al., 2016). This approach has been extended to operant choice paradigms in which Gr64f-driven activation directs flies toward specific arms of a maze, with reward expectation shaping learning rates and choice statistics (Rajagopalan et al., 2023). A clear advantage of this approach is that the reward signal can be tied precisely to spatial location rather than to ingestion, which permits clean experimental control over what is reinforced. While *Drosophila* can form spatial preferences using aversive heat-based reinforcement (Wustmann et al., 1996; Foucaud et al., 2010) and freely walking flies can learn appetitively reinforced operant tasks using sucrose reward (Wiggin et al., 2021), it remains less clear whether a purely non-caloric, optogenetically delivered sweet-neuron signal can support spatial place conditioning when reward is tied directly to location.

A separate but related question concerns the nature of any preference that does emerge: is it goal-directed — i.e., dependent on the current value or current presence of the reward — or is it a habit, expressed independently of the outcome (Dickinson, 1985; Balleine and Dickinson, 1998; Bouton, 2024)? The classical assay for this distinction is the contingency omission test, in which the reward is removed and behavior is tracked across extinction trials. Goal-directed responding declines when the contingency is broken; habitual responding persists. In flies, omission has been shown to engage dopaminergic prediction-error circuitry in the mushroom body, generating an opposing extinction memory (Felsenberg et al., 2018), and Brembs (2009) has further shown that the mushroom body itself regulates the transition from goal-directed to habitual behavior with extended training. Whether a freely walking, optogenetically conditioned spatial preference is sensitive to omission of the action–outcome contingency, however, remains untested.

The present experiments address two hypotheses. First, that *Drosophila* expressing Gr64f-Gal4 > UAS-CsChrimson will develop a measurable arm preference when entry into one arm of a V-shaped maze triggered red-light stimulation, but only in flies fed all-trans-retinal (ATR+); ATR– sibling controls, lacking the chromophore required for CsChrimson function, should show no preference. Second, that any such preference reflects goal-directed responding and will therefore decline when the contingency is altered — either by removing the reward (extinction) or by reassigning it to the opposite arm (reversal). The

original proposal called for an extinction-only Phase 2; in practice, three of the six ATR+ flies were carried into a Phase 2 manipulation, with two assigned to reversal and one to extinction, providing complementary tests of behavioral flexibility within the same dataset. The combined Phase 1 / Phase 2 design therefore tests both whether the conditioning works at all and whether the resulting preference shows the outcome sensitivity expected of goal-directed behavior.

Methods

Fly stocks and rearing. All experiments used Gr64f-Gal4 × UAS-CsChrimson flies (lab vial "2-12 / Sweet Response"). Flies were housed on standard cornmeal–dextrose food; the experimental (ATR+) cohort was raised in the dark on food supplemented with 0.5 mM all-trans-retinal (ATR) for at least two days prior to recording, since CsChrimson requires the all-trans-retinal chromophore to gate light-driven cation flux (Klapoetke et al., 2014; Yapici et al., 2016). The control (ATR–) cohort was the same genotype reared on standard food without ATR, so any place preference in the experimental group could be attributed to functional CsChrimson activity rather than to the red light itself or to baseline phototactic biases. Flies were food-deprived for ~16–20 hours before testing to heighten appetitive motivation (Kirkhart and Scott, 2015). Across recording sessions, 10 flies were recorded in total: n = 4 ATR– control flies (Phase 1 only), n = 3 ATR+ Phase 1–only flies, n = 1 ATR+ Phase 1 → extinction fly, and n = 2 ATR+ Phase 1 → reversal flies. The six ATR+ flies were therefore split into a Phase-1-only subgroup used for the ATR+ vs. ATR– group comparison and a Phase-2 subgroup used for the within-fly omission/reversal analyses described below.

Maze apparatus and optogenetic stimulation. Behavior was recorded in a V-shaped chamber consisting of two ~5 cm arms joined at a central entry stem (Fig. 0A). The chamber was placed under the existing teaching-lab observation platform; an FLIR Firefly camera (FFY-U3-16S2M, ~60 fps) mounted overhead acquired video under ambient illumination. A 617 nm LED was positioned above the maze and controlled through an Arduino interface running within the Bonsai workflow. Exact irradiance at the maze surface was not calibrated in mW/cm²; the experiment can be reproduced by matching the controller setting (level 5) on the same hardware, though future work should report calibrated intensity to enable cross-lab comparison. During closed-loop stimulation, entry into the preassigned reward-arm ROI triggered 617 nm illumination, while entries into the opposite arm or central stem did not trigger stimulation. The stimulation-contingent arm therefore served as the “reward arm,” defined operationally as the arm in which entry produced CsChrimson activation in ATR+ flies. Arm membership was initially monitored in Bonsai using two independently defined rectangular ROIs (ROI1, ROI2) placed over each arm in the camera frame; post hoc Python analysis assigned centroid positions to reward or non-reward arms using fly-specific boundaries aligned to the two maze arms. For each video, arm boundaries were drawn once using the static maze geometry visible in the camera frame, not adjusted based on the fly’s behavior. Centroid positions falling within the left or right arm polygon were scored as arm occupancy; frames in the central stem or outside both arm masks were excluded from dwell-time PI calculations. The reward arm was counterbalanced across flies (left vs. right) to control for chamber-side biases. The genetic logic of the

optogenetic circuit, the resulting ATR+/ATR– contrast, and the four behavioral subgroups are summarized in Fig. 0B–C.

Phase 1 (conditioning). Each fly was loaded into the central stem of the chamber and allowed to explore freely for 10–15 minutes under closed-loop stimulation. During Phase 1, entry into the preassigned reward-arm ROI triggered 617 nm illumination, while entries into the opposite arm did not trigger stimulation (Fig. 0D). No baseline-only probe was used; the conditioning trial was treated as a continuous learning episode whose internal dynamics could be analyzed post hoc through the rolling preference index (see below). Because no separate baseline probe was collected, Phase 1 preference was interpreted relative to ATR– sibling controls and, where possible, to within-session changes in rolling PI rather than as a pure pre/post-acquisition measure. For ATR+ flies, every entry into the reward arm therefore co-occurred with sweet-GRN activation, while for ATR– controls the same physical stimulus was present but neurally silent.

Phase 2 (extinction or reversal). Three of the six ATR+ flies were carried into a Phase 2 manipulation immediately after Phase 1, with no rehandling or re-loading. In the extinction condition (Fly 4, n = 1), the LED was switched off in both arms for the remainder of the session, removing the reinforcer entirely; this is the standard contingency-omission probe of goal-directedness (Dickinson, 1985; Felsenberg et al., 2018). In the reversal condition (Flies 5 and 6, n = 2), the LED was instead switched to illuminate the previously non-rewarded arm, so that the reward contingency was inverted. Reversal provides a complementary test of behavioral flexibility: if the Phase 1 preference is goal-directed, the fly should track the reward to its new location, whereas a habit-like preference should persist toward the original arm. The decision to add reversal to the original extinction-only design was made on a fly-by-fly basis during data collection, on flies that were still tracking cleanly at the Phase 1 endpoint.

Video tracking and pose estimation. Video was acquired at ~60 fps using the FLIR Firefly camera (FFY-U3-16S2M, Mono8 format) and processed in real time using Bonsai (Lopes et al., 2015). Within the Bonsai workflow, each camera frame was cropped to the maze region, converted to grayscale, and passed through background subtraction to isolate the fly. A Range Threshold node binarized the resulting image, and BinaryRegionAnalysis extracted the centroid (x, y) of the largest connected region per frame — corresponding to the fly's body. Centroid coordinates were logged continuously to CSV via CsvWriter nodes at full camera frame rate. The LED was triggered via an Arduino microcontroller interfaced through Bonsai, with ROI-entry events detected by the DistinctUntilChanged/RisingEdge nodes in the workflow driving the stimulus logic.

Behavioral metrics. Two preference indices (PI) were computed, both bounded on [–1, +1]. The *entry-count PI* was defined as (entries into reward arm – entries into non-reward arm) / total entries, scored from arm-occupancy state transitions. The *dwell-time PI* (also called the "centroid-time PI") was defined as (frames in reward arm – frames in non-reward arm) / (frames in either arm) and is the more sensitive measure because it uses every tracked frame rather than discrete entries. Two-dimensional trajectory

density maps were generated by binning the centroid coordinates and smoothing with a Gaussian kernel; these visualizations were used to inspect spatial structure that summary indices alone could miss (Fig. 1). Within-session learning dynamics were captured by a *rolling PI* computed in 60-second sliding windows. Curve fits were used descriptively to summarize temporal structure within individual flies rather than as inferential tests of distinct learning mechanisms; for ATR+ Phase 1 flies, rolling PI(t) was fit with four candidate models — sigmoidal, Gaussian, linear, and exponential saturating — and the best-fitting model was selected by R^2 (Fig. 4, top row). Per-frame movement speed was computed as centroid displacement between consecutive frames multiplied by frame rate. To suppress tracking artifacts from occasional centroid jumps, speeds above the 85th percentile for each fly × arm combination were excluded before plotting and computing arm-specific means for the Phase 1 locomotor analysis (Fig. 5).

Statistical analysis. Given the small cohort and non-normal distributions, all inferential tests were non-parametric. Because the prior prediction was directional, that ATR+ flies would show greater reward-arm preference than ATR– controls: ATR+ vs. ATR– group comparisons in Phase 1 were performed with planned one-sided Mann–Whitney U tests on each PI metric, with rank-biserial r reported as effect size. As a complementary test that ATR+ PI exceeded chance (PI = 0), a one-sample Wilcoxon signed-rank test was applied to the ATR+ group. Phase 2 changes were examined within-fly by computing ΔPI = PI(Phase 2) – PI(Stim ON) for each manipulated fly. All analyses were performed in Python (pandas, NumPy, SciPy); figures were generated in matplotlib. † $p < 0.10$, * $p < 0.05$.

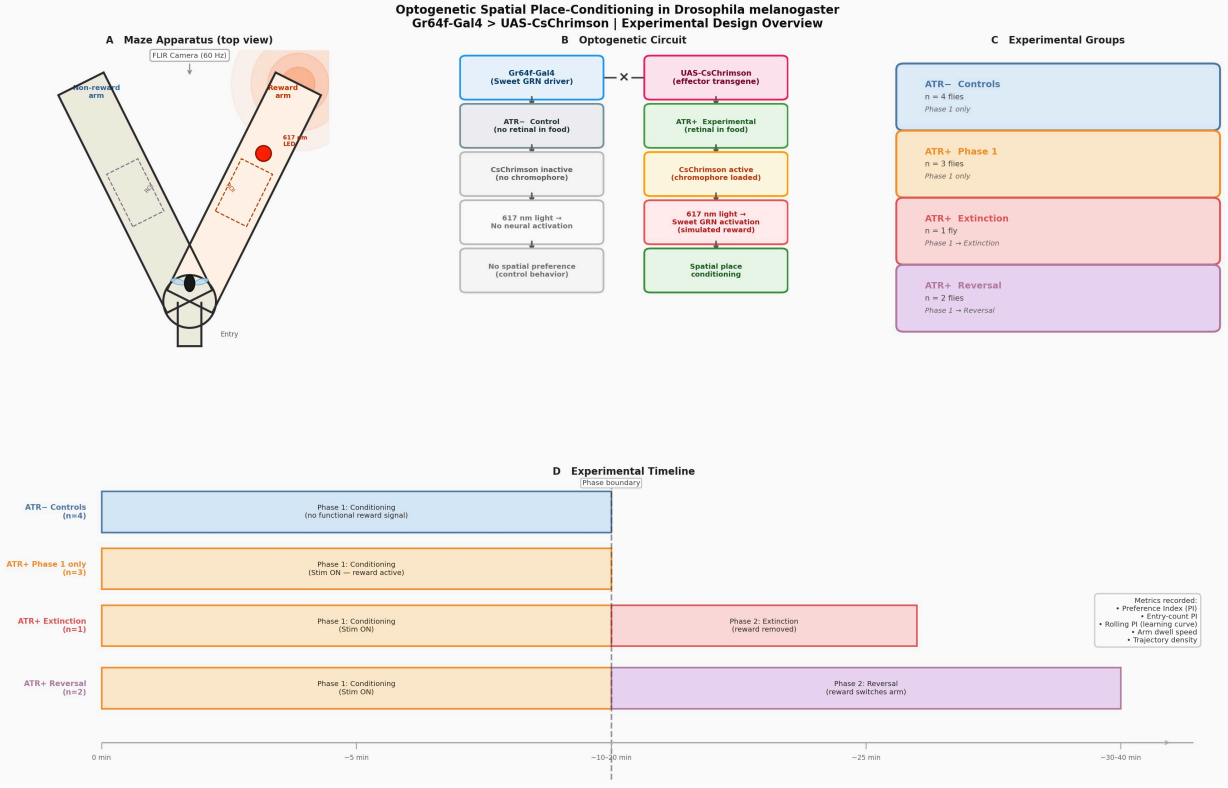

Figure 0. Experimental design overview. (**A**) Top-down schematic of the V-shaped maze. A 617 nm LED delivered closed-loop stimulation contingent on reward-arm ROI occupancy; the FLIR camera records from above at 60 fps. Dashed boxes indicate the arm regions of interest used for occupancy scoring. (**B**) Optogenetic logic. Gr64f-Gal4 drives UAS-CsChrimson expression in sweet-sensing GRNs. (**C**) The four behavioral subgroups: ATR– controls (n = 4, Phase 1 only), ATR+ Phase 1 only (n = 3), ATR+ Extinction (n = 1), and ATR+ Reversal (n = 2). (**D**) Session timelines per group. All flies underwent ~10–15 min of Phase 1 conditioning; the Extinction fly then had the LED switched off in both arms, and the Reversal flies had the LED relocated to the previously non-rewarded arm.

Results

Spatial trajectories reveal arm-level structure consistent with stimulation-associated occupancy.

Centroid trajectory density maps (Fig. 1) show where each fly spent its time. ATR– controls show broadly distributed density along both arms with no consistent bias toward the reward arm. In contrast, ATR+ Phase 1 flies show higher density concentrated in the reward arm, most clearly in Flies 1 and 2. Phase 2 panels show flatter distributions; reversal flies do not show a clear switch toward the new reward arm, an observation taken up in the Discussion.

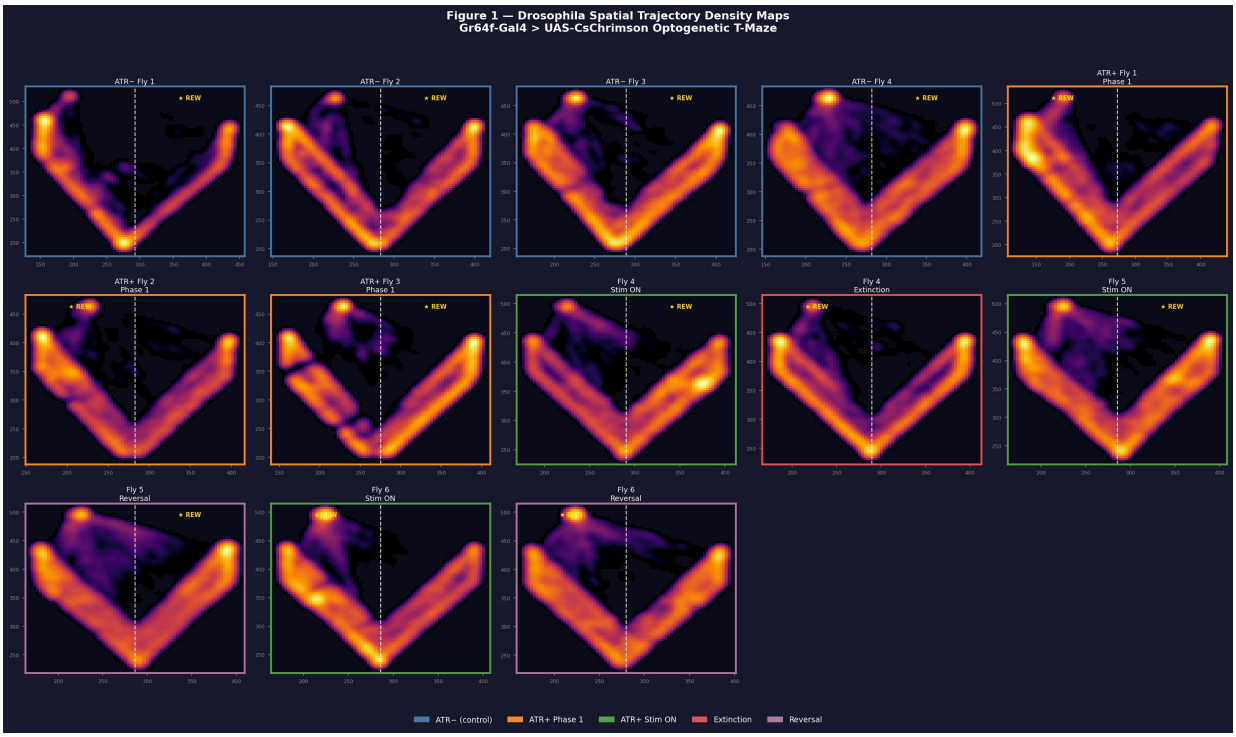

Figure 1. Spatial trajectory density maps for all flies across all phases. Each panel shows a Gaussian-smoothed 2D histogram of the Bonsai-tracked centroid for one fly under one phase. The yellow "REW" marks the rewarded (LED-illuminated) arm for that phase; the dashed white line indicates the central midline used to score arm membership. Top row: ATR– controls (Phase 1 only). Middle row: ATR+ Phase 1 conditioning (Flies 1–3) and the Stim ON segments of Phase 2 flies. Bottom row: Phase 2 outcomes — the single Extinction fly and both Reversal flies, with the new reward arm relabeled. ATR+ flies show denser occupancy of the rewarded arm during Stim ON than ATR– controls, while Phase 2 panels show flatter, less biased distributions.

ATR+ flies develop a positive arm preference during Phase 1.

Quantitatively, ATR+ flies showed a positive arm preference on both PI metrics, while ATR– controls did not (Fig. 2). For the entry-count PI, ATR+ flies averaged +0.21 ± 0.20 (mean ± SD), with all three flies at or above zero, while ATR– controls averaged +0.00 ± 0.13 and were distributed symmetrically around zero. The Mann–Whitney comparison did not reach the conventional 0.05 threshold (U = 9.0, p = 0.200, rank-biserial r = +0.50), which is unsurprising given n = 3 vs. n = 4. The dwell-time PI gave a stronger, more separable signal: ATR+ flies averaged +0.34 ± 0.26, with all three flies clearly above zero, while ATR– controls averaged –0.22 ± 0.32 — i.e., the controls showed a slight bias *against* the LED arm, plausibly a residual avoidance of the brighter side of the chamber. The ATR+ vs. ATR– contrast on the dwell-time PI was U = 11.0, p = 0.057, r = +0.83 (Fig. 2B; Fig. 3); a one-sample Wilcoxon test of ATR+ against PI = 0 yielded W = 6, p = 0.125 — direction-consistent but underpowered at n = 3. The dwell-time PI was consistently the more sensitive readout.

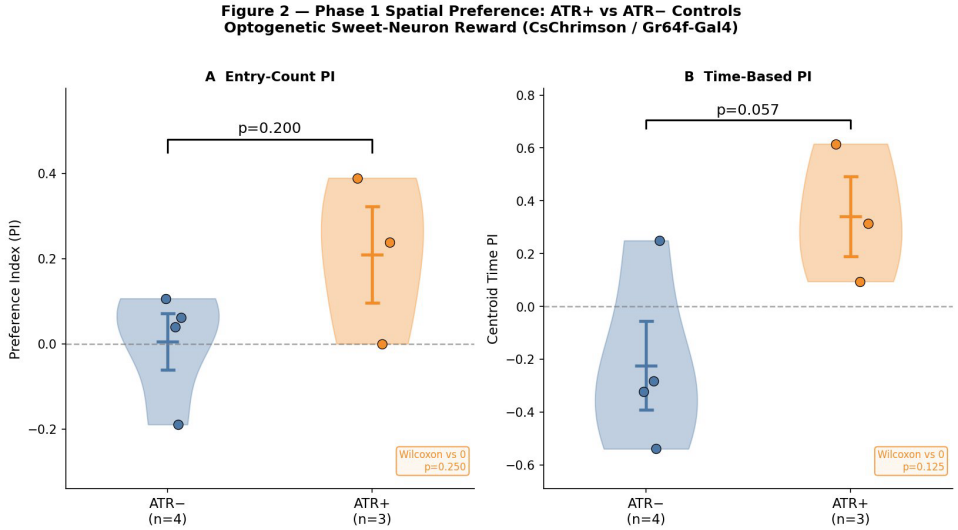

Figure 2. Phase 1 spatial preference: ATR+ vs. ATR– controls. (**A**) Entry-count PI. (**B**) Time-based (dwell-time) PI. Filled dots are individual flies, violins are kernel density estimates, and error bars are mean ± SEM. Dashed line at PI = 0 indicates no preference. p-values are from one-sided Mann–Whitney U tests; the inset boxes report the one-sample Wilcoxon test of the ATR+ group against PI = 0.

Statistical Summary — Spatial Place-Conditioning

Gr64f-Gal4 > UAS-CsChrimson · ATR+ (n = 3) vs ATR− (n = 4)

Metric	ATR+ Mean ± SD	ATR− Mean ± SD	Test Statistic	p-value	Effect size (r)
Phase 1 · Inferential Statistics					
Group comparison · ATR+ vs ATR− (Mann–Whitney U)					
Entry-Count PI	+0.209 ± 0.198	+0.004 ± 0.130	U = 9.0	0.200	+0.500
Dwell-Time PI	+0.341 ± 0.263	−0.223 ± 0.319	U = 11.0	0.057 †	+0.833
Above-chance test · ATR+ vs PI = 0 (Wilcoxon Signed-Rank)					
Entry-Count PI	+0.209 ± 0.198	—	W = 5	0.250	—
Dwell-Time PI	+0.341 ± 0.263	—	W = 6	0.125	—

Phase 2 · Within-Fly Changes (Dwell-Time PI, Stim → Phase 2)				
Fly	Phase 2 Type	Stim PI	Phase 2 PI	Δ PI
Fly 4	Extinction	+0.351	+0.195	**−0.156**
Fly 5	Reversal	+0.295	+0.118	**−0.177**
Fly 6	Reversal	+0.118	+0.094	**−0.024**

Entry-Count PI — arm entries scored by observer

Dwell-Time PI — centroid-tracked time in each arm

† p < 0.10 * p < 0.05 ** p < 0.01 *** p < 0.001 | Effect size = rank-biserial r | PI range: −1 to +1

Figure 3. Statistical summary across all measures. Top block: Phase 1 inferential statistics, including the ATR+ vs. ATR– group comparison (Mann–Whitney U) and the ATR+ vs. PI = 0 above-chance test (Wilcoxon signed-rank). The dwell-time PI shows a marginal group difference (p = 0.057 †) with a large effect size (r = +0.83). Bottom block: Phase 2 within-fly changes for the three ATR+ flies that received a Phase 2 manipulation. Stim PI is the dwell-time PI during the Phase 1 portion of the same session; Phase 2 PI is computed identically over the post-manipulation interval. ΔPI is negative in all three flies, consistent with outcome sensitivity to the broken contingency.

Within-session PI dynamics show structured temporal patterns.

Group-mean comparisons discard within-session dynamics. The rolling PI provides a richer view (Fig. 4, top row); fits are interpreted descriptively rather than as formal evidence of distinct learning mechanisms. Fly 1 was best fit by a Gaussian (R^2 = 0.64), with PI rising from ~0.45 to ~0.65 across the session. Fly 2 was best fit by a Gaussian-shaped trajectory with a declining preference over time (R^2 = 0.98), an early bias that decayed despite ongoing reward availability. Fly 3 showed a sigmoidal acquisition curve (R^2 = 0.93), rising from ≈ –0.10 to a plateau near +0.15 by minute 4. Linear and exponential-saturating fits were uniformly worse in these descriptive fits, suggesting that the observed rolling PI traces were not well summarized by simple monotonic trends alone.

Phase 2: removing or reassigning the reward attenuates the learned preference.

The bottom row of Fig. 4 displays Phase 2 dynamics for the three ATR+ flies that received a Phase 2 manipulation. Fly 4 (extinction) showed the cleanest qualitative effect: during Stim ON the rolling PI fluctuated between approximately +0.15 and +0.65 with a clear positive central tendency, and after the LED was turned off in both arms, the rolling PI shifted downward and became negative during much of

the extinction window, eventually settling near –0.30. However, the aggregate dwell-time PI across the full extinction segment remained positive but reduced, falling from +0.351 during Stim ON to +0.195 during extinction, a within-fly change of ΔPI = –0.156 (Fig. 3, bottom block).

The two reversal flies (Fly 5, Fly 6) both showed a decline in dwell-time PI in Phase 2 (ΔPI = –0.18 and –0.02 respectively), but did not invert the preference toward the new reward arm. Fly 5 hovers near zero with high-amplitude fluctuations (±0.5); Fly 6 drops toward zero but stays slightly positive. Both reversal flies "let go" of the original preference without acquiring a new one within the recording window.

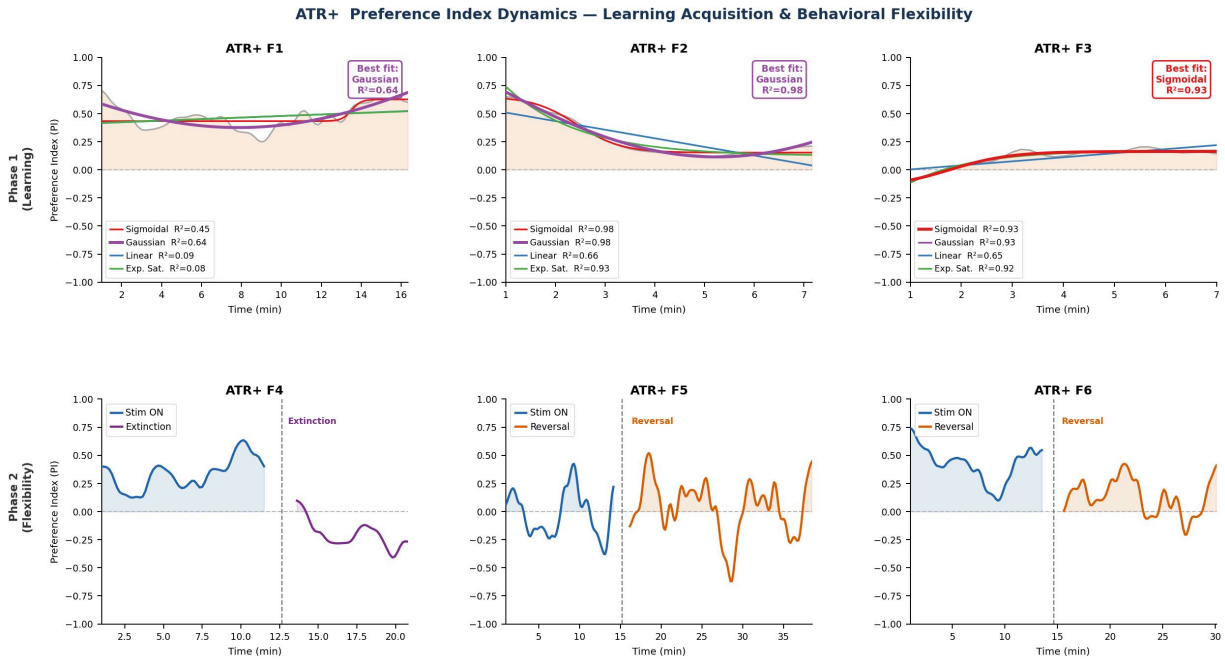

Figure 4. ATR+ rolling preference index dynamics. Top row (Phase 1, learning acquisition): rolling PI in 60-s windows for each ATR+ Phase 1 fly (gray trace), overlaid with sigmoidal (red), Gaussian (purple), linear (blue), and exponential-saturating (green) fits. The best-fit model is highlighted in the legend with the highest R^2. Bottom row (Phase 2, behavioral flexibility): rolling PI for the three ATR+ flies that received a Phase 2 manipulation. The vertical dashed line marks the phase boundary between Stim ON (blue) and Phase 2 (purple = extinction; orange = reversal). All three flies show a reduction in reward-arm preference in Phase 2 relative to Stim ON, consistent with sensitivity to the broken contingency.

Movement-speed analysis argues against a simple locomotor confound.

A common concern with PI-based readouts is that arm bias could in principle be driven by gross locomotor differences (e.g., the fly slowing down whenever it crosses into the bright arm). To address this, per-frame movement speed was filtered to the 85th percentile and split by arm identity (Fig. 5). Across the four ATR– controls, reward and non-reward arm speeds were essentially indistinguishable: ATR– F1 (~70 vs. ~65 px/s), F2 (~52 vs. ~58), F3 (~40 vs. ~50), and F4 (~39 vs. ~46), respectively, with no consistent directional bias. ATR+ flies showed a different pattern: ATR+ F1 and F2 were actually *faster* in the reward arm than in the non-reward arm (F1: 65 vs. 43; F2: 60 vs. 46 px/s), while ATR+ F3 showed similar speeds in both arms (65 vs. 63 px/s). Thus, the Phase 1 speed analysis does not support a simple locomotor-slowing

explanation for the positive dwell-time PI: ATR+ flies did not merely accumulate reward-arm time because they became immobile in the stimulated arm.

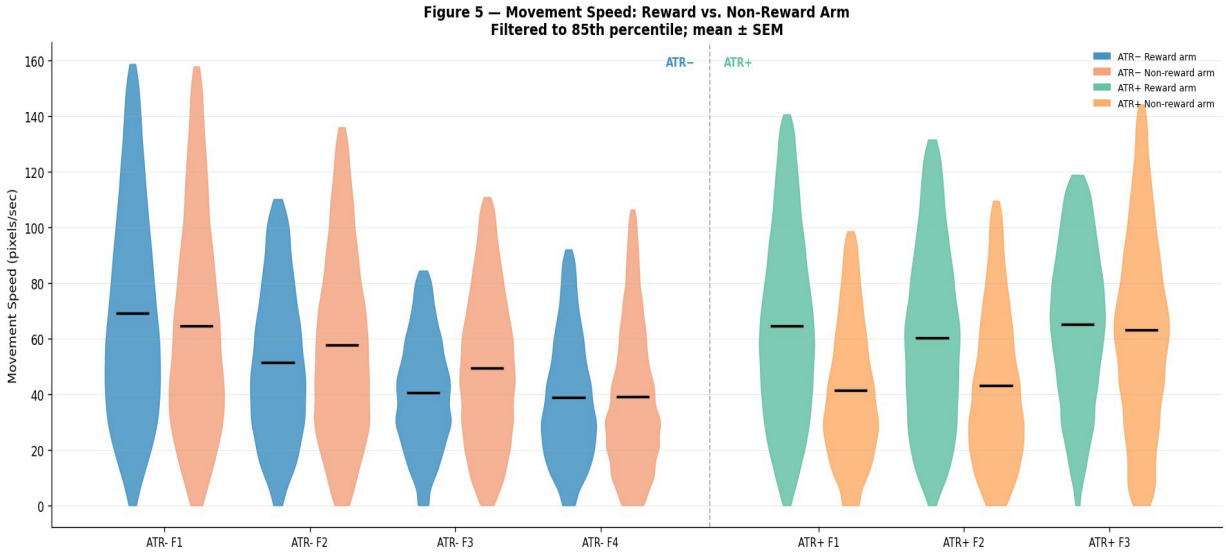

Figure 5. Movement speed during Phase 1, split by arm. Per-frame centroid speed (pixels/sec, filtered to the 85th percentile) for the four ATR– controls (left of dashed line) and three ATR+ flies (right of dashed line). Each fly contributes two violins: reward arm (solid) and non-reward arm (lighter). Black bars are mean ± SEM. ATR– flies show small, inconsistent reward-vs-non-reward speed differences. Two of three ATR+ flies (F1, F2) show faster movement in the reward arm than in the non-reward arm, arguing against the possibility that the dwell-time preference was driven simply by locomotor slowing in the stimulated arm.

Discussion

Three convergent lines of evidence support the central conclusion that optogenetic Gr64f-driven sweet-neuron activation is sufficient to drive a measurable, location-specific spatial preference in freely walking *Drosophila* within a single ~10–15 minute session. First, ATR+ flies showed a positive dwell-time PI relative to ATR– controls with a large effect size (r = +0.83) and a marginal group p-value (p = 0.057) given the small cohort. Second, all three ATR+ Phase 1 flies were well-fit by structured within-session dynamics models — sigmoidal acquisition for Fly 3, Gaussian for Fly 1, and a slow decay from an early bias for Fly 2 — which is a stronger statement than the group means alone; organized PI(t) trajectories of this kind are less consistent with purely unbiased exploration, although these model fits remain descriptive given the small sample. Third, the movement-speed analysis argues against a simple locomotor confound: two of three ATR+ flies were faster, not slower, in the reward arm during Stim ON, making it unlikely that positive dwell-time PI was driven merely by immobility or slowing in the stimulated arm.

The Phase 2 manipulations provide an additional, complementary line of evidence that the resulting preference is sensitive to the action–outcome contingency, consistent with the kind of outcome sensitivity that underlies goal-directed responding (Dickinson, 1985; Balleine and Dickinson, 1998; Bouton, 2024), though the small Phase 2 cohort (n = 1 extinction, n = 2 reversal) prevents a definitive distinction between goal-directed and habit-like control. The Phase 2 data are best interpreted as evidence for outcome

sensitivity, which is necessary for goal-directed control, rather than as a definitive demonstration of goal-directed behavior. In all three flies given a Phase 2 manipulation, dwell-time PI declined relative to Stim ON. The extinction fly (Fly 4) showed the most striking change: a positive, persistent rolling PI during Stim ON shifted downward after the LED was turned off and became negative during much of the extinction window. This is consistent with the signature predicted by the omission-test framework: when the contingency is broken, the behavior tracks the change. The reversal flies showed a partial version of the same effect — they released the original preference but did not acquire a new one within the recording window — which is itself biologically informative. Reversal of an appetitive memory in *Drosophila* is known to engage opposing dopaminergic populations and to involve formation of an extinction-like memory that competes with the original (Felsenberg et al., 2018); under that model, "letting go without re-binding" is exactly what one would expect in the early phase of reversal, before the new contingency has been learned. A within-session reversal of only ~15 minutes is at the lower bound of what such a process would require.

Several limitations apply. Most importantly, the Phase 1 group comparison rested on n = 3 ATR+ vs. n = 4 ATR– flies, and Phase 2 on n = 3 ATR+ flies (1 extinction, 2 reversal); statistical power is therefore severely limited and the central comparison did not reach $p < 0.05$. A future design should pre-register whether all ATR+ Stim ON segments are included in the Phase 1 comparison, or should collect a larger independent Phase-1-only ATR+ cohort to provide adequate power to test the central group difference. The Phase 2 design was also unbalanced and assigned fly-by-fly rather than randomized, preventing direct between-manipulation comparison. Density maps for Flies 1 and 2 suggest some reward-zone behavior may be local to the LED illumination region; a reward-zone-only ROI analysis would likely strengthen the dwell-time signal. Finally, the V-maze lacks discrete trial structure, making clean trial-averaging impossible, though this also enables naturalistic within-session trajectory analysis.

The most direct next step is a larger cohort (n ≈ 10 per group) with pre-registered Phase 2 assignment, which would provide adequate power to test the central group difference and allow balanced extinction-vs-reversal comparison. Mechanistically, pairing this behavioral assay with thermogenetic silencing of PAM dopaminergic neurons would test whether conditioning depends on the canonical mushroom-body reward pathway, and an extended-training condition would test whether prolonged Phase 1 shifts the preference toward habitual responding as predicted by Brembs (2009).

Taken together, these results extend the use of CsChrimson-driven Gr64f activation from prior choice and operant paradigms (Yapici et al., 2016; Rajagopalan et al., 2023) into a freely walking spatial place-conditioning context and provide preliminary evidence that the resulting preference shows outcome sensitivity consistent with goal-directed control. In doing so, they show that even a small, single-session experiment can produce convergent signals across PI, learning-curve, and locomotor measures, provided the analysis pipeline is built to read out within-session structure rather than relying on group means alone.

This paper represents my own work in accordance with University regulations

Generative AI (Claude Code, Anthropic) was used to assist with the Python-based (.ipynb) data analysis workflow in VS Code. The written report, interpretation, and final figures represent my own work. Relevant Notebooks can be provided upon request (Consulted w/ Prof. Collins)

References

Aso Y, Sitaraman D, Ichinose T, Kaun KR, Vogt K, Belliart-Guérin G, et al. (2014) Mushroom body output neurons encode valence and guide memory-based action selection in Drosophila. eLife 3:e04580.

Balleine BW, Dickinson A (1998) Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. Neuropharmacology 37:407–419.

Bouton ME (2024) Habit and persistence. J Exp Anal Behav 121:88–96.

Brembs B (2009) Mushroom bodies regulate habit formation in Drosophila. Curr Biol 19:1351–1355.

Dahanukar A, Lei YT, Kwon JY, Carlson JR (2007) Two Gr genes underlie sugar reception in Drosophila. Neuron 56:503–516.

Dickinson A (1985) Actions and habits: the development of behavioural autonomy. Philos Trans R Soc Lond B Biol Sci 308:67–78.

Felsenberg J, Jacob PF, Walker T, Barnstedt O, Edmondson-Stait AJ, Pleijzier MW, et al. (2018) Integration of parallel opposing memories underlies memory extinction. Cell 175:709–722.

Foucaud J, Burns JG, Mery F (2010) Use of spatial information and search strategies in a water maze analog in Drosophila melanogaster. PLoS One 5:e15231.

Jiao Y, Moon SJ, Wang X, Ren Q, Montell C (2008) Gr64f is required in combination with other gustatory receptors for sugar detection in Drosophila. Curr Biol 18:1797–1801.

Kahsai L, Zars T (2011) Learning and memory in Drosophila: behavior, genetics, and neural systems. Int Rev Neurobiol 99:139–167.

Kim YC, Lee HG, Han KA (2007) D1 dopamine receptor dDA1 is required in the mushroom body neurons for aversive and appetitive learning in Drosophila. J Neurosci 27:7640–7647.

Klapoetke NC, Murata Y, Kim SS, Pulver SR, Birdsey-Benson A, Cho YK, et al. (2014) Independent optical excitation of distinct neural populations. Nat Methods 11:338–346.

Lopes G, Bonacchi N, Frazão J, Neto JP, Atallah BV, Soares S, et al. (2015) Bonsai: an event-based framework for processing and controlling data streams. Front Neuroinform 9:7.

Kirkhart C, Scott K (2015) Gustatory learning and processing in the Drosophila mushroom bodies. J Neurosci 35:5950–5958.

Rajagopalan AE, Darshan R, Hibbard KL, Fitzgerald JE, Turner GC (2023) Reward expectations direct learning and drive operant matching in Drosophila. Proc Natl Acad Sci USA 120:e2221415120.

Scott K (2018) Gustatory processing in Drosophila melanogaster. Annu Rev Entomol 63:15–30.

Tully T, Quinn WG (1985) Classical conditioning and retention in normal and mutant Drosophila melanogaster. J Comp Physiol A 157:263–277.

Wiggin TD, Hsiao Y, Liu JB, Huber R, Griffith LC (2021) Rest is required to learn an appetitively-reinforced operant task in Drosophila. Front Behav Neurosci 15:681593.

Wustmann G, Rein K, Wolf R, Heisenberg M (1996) A new paradigm for operant conditioning of Drosophila melanogaster. J Comp Physiol A 179:429–436.

Yapici N, Cohn R, Schusterreiter C, Ruta V, Vosshall LB (2016) A taste circuit that regulates ingestion by integrating food and hunger signals. Cell 165:715–729.