

Insect-model steering for a simulated search-and-return drone: two pre-registered studies against equal-budget rivals

Mutaqin Aryawijaya

Independent researcher ORCID 0009-0008-5199-2413

Abstract

We tested whether published insect navigation models, combined, improve a simulated quadrotor that must find a pictured target, stop at it and return without GPS or a map. Study 1 (pre-registered) combined a mushroom-body target memory, a central-complex path integrator and a ring-attractor steering hub, and compared the hub with five equal-budget arbiters. In 2D kill tests, recurrence traded holding a hidden target against releasing a look-alike, but a leaky accumulator did better on both. In 45,000 3D missions the ring scored slightly below the best rival (success-curve area 0.623 vs 0.631), as did two arms without memory; the mushroom body was at chance (ROC area 0.502) with a look-alike on the target's spot; and compass and odometry faults defeated every arm. Study 2 (pre-registered on fresh seeds) added an insect-style heading ring, a nest learning flight with view-based landing, a colour and edge front end and an evidence-gated ring. The integrated brain found the target, returned and landed on the pad in 0.990 of 300 fault-free missions and kept middle-severity success at or above 0.71 under six of seven faults. The pre-registered goal was not met: with a swapped look-alike, success was 0.39. Post hoc, it beat Study 1 only on the three path-integrator faults. Component tests missed their criteria, the arbiter barely mattered, and among non-definitional ablations removing the home view memory cost most. In an exploratory post-study extension, size-aware search raised fault-free success with a small target from 0.000 to 0.980. All results are from simulation.

1 Introduction

Several insect navigation circuits have published computational models. The mushroom body can be modelled as a sparse random expansion whose output reports how familiar a view is after one-shot learning [1]. A central-complex circuit can integrate compass and speed signals into a home vector [2], and a ring-attractor model of the central complex can blend or select navigational cues by their certainty [3, 4]. In the fly, a bump of activity in the compass ring persists in darkness for seconds to tens of seconds [5, 6]; it is a heading memory, not a goal memory. In walking flies, descending neurons combine steering from the central complex with steering from sensory reflexes [7]; using this arrangement in flight is our extrapolation. Bees and wasps fly structured learning flights around a new goal and record views of it [8, 9].

These models invite an engineering test. A small drone is shown a picture of a target, for example a tent at a search-and-rescue supply drop. It must find the target without GPS or a map, stop at it and fly home through trees, wind and sensor faults. We asked whether a controller built from insect models beats simple engineered parts given the same tuning effort.

Parts of this question have been studied. Sun, Yue and Mangan combined path integration, visual homing and route following in a central-complex ring for a simulated walker in two dimensions without obstacles [4]. Goulard, Heinze and Webb summed path integration with a seen-target memory in a central-complex model [10]. Lu and Webb navigated to coordinate goals with a mushroom body and central complex, including obstacle avoidance [11]. Hoinville and Wehner showed that a circular vector sum also leans toward the more reliable cue at large conflicts [12]. That recurrent memory trades persistence against flexibility is consistent with attractor-network theory: a ring network holds a localized activity profile through its recurrent weights [13], and

neuromodulation of recurrent synapses has been proposed to set how well bump-attractor memories resist distractors [14]. Bicho, Mallet and Schöner gave a robot a dynamic-field target memory that holds a lost target and deletes it after a delay [15]. Bee-Nav trained a tiny network from robotic learning flights to home a real quadrotor [16], and Firlefyn and colleagues learned home-vector direction directly from learning-flight images [17]. AntFlie flew a mushroom-body memory under 4 kB on a 33 g drone [18], and van Dijk and colleagues followed routes on a small drone with very compact snapshot memories [19]. Gao and colleagues flew a quadrotor to the exact object instance shown in a query image [20], and ForaNavig addressed insect-inspired online target-oriented navigation for micro aerial vehicles in tree plantations [21]. Voting [22], histogram methods [23], behaviour trees [24] and ring attractors for action selection [25] are established arbitration schemes.

This paper reports two pre-registered studies in one simulator. Its contributions are:

- an equal-budget comparison of a ring-attractor steering hub with five arbiters, in kill tests and in six-degree-of-freedom missions, with a measured hold-versus-release frontier;
- a recognition test of mushroom-body memories under a swap control, with and without learning flights, against pixels and DINOv2 [26];
- a second study, on fresh seeds, of insect-inspired fixes for the failures of the first, with an ablation of every module.

2 Methods

2.1 Task, simulator and vehicle

The drone flew in a 100×100 m forest (obstacle density 0.3) to a target tent 25–40 m from home, with a 300 s

limit. An operator tip gave the target’s rough bearing and distance (bearing error drawn from $N(0, 8^\circ)$, distance scaled by $1 + N(0, 0.1)$). The vehicle is a six-degree-of-freedom rigid body integrated at 500 Hz with the Crazyflie 2.0 mass from the MuJoCo Menagerie model [27, 28], inertia identified by Landry [29], and motor maps, motor lag and drag from Förster’s system identification [30]. Turbulence follows the low-altitude Dryden form of MIL-F-8785C [31] with MathWorks intensities [32], an engineering choice for a 30 g vehicle. A cascade proportional-integral-derivative (PID) loop (velocity, attitude, rate) is shared by every arm. The brain runs at 50 Hz on a ray-cast low-resolution panoramic eye. Obstacle avoidance by optic-flow balancing and looming is computed geometrically from the ray-casts and added after the hub, at the level of descending neurons [7]. Every result comes from a vectorised Python reference. A browser port runs the same equations; a parity check run on the mission code that produced the Study 2 results matched the reference in closed loop for all six Study 1 arbiters and for the Study 2 brain.

2.2 Study 1 brain

Fig. 1 shows the parts. *Target memory*: a mushroom body after Ardin and colleagues [1] (20,000 Kenyon cells, 1% active, one-shot learning), loaded from the target photo and a 10 s learning flight. Its familiarity was mapped to object evidence in [0, 1]. *Home vector*: a path integrator reimplemented from the text of Stone and colleagues [2], fed by a sun compass and odometry. *Ring hub*: a rate ring of 32 neurons after Sun, Yue and Mangan [3, 4]; each drive (object, home, search) enters as a von Mises bump scaled by its gain and strength, recurrent excitation is scaled by a gain g ($g = 0$ removes recurrence), and the heading command is the population vector.

2.3 Arbiters and equal-budget tuning

Six arms replaced the hub in turn: the ring; the ring without recurrence; a static vote (argmax of the same input stage); a vector sum; a leaky evidence accumulator; and a state machine (search, approach, hold-last-bearing with a timeout, home) written before any fault data and frozen by hash. Its author knew the hypotheses and the fault list. All arms saw, per object, a bearing, a strength and a visibility flag; the state machine used the flag, while the ring and the vote saw only strength times visibility. Every arm received the same 64 configurations drawn from one uniform matrix, evaluated in two runs on training seeds (3D: all 64 on seeds 0–4, the best 4 on seeds 0–19), and its winner was frozen before any test seed ran. For the ring the budget was equal in number but not in effect. 53 of the 64 draws had a gain of 0.7 or more, and most of these latched onto their first choice and never turned home. Only 13 of 64 were usable in 2D and 24 of 64 in 3D.

2.4 Study 1 tests

Kill tests. These are cheap pre-registered tests, here in a two-dimensional point-mass world, each with a threshold (the kill line) at which a hypothesis is dropped. KT1 (H2) compared the ring’s heading with a static vote across two-cue conflicts. KT2 (H1 pilot) flew 28 fault conditions on 300 test seeds.

KT3 (H3) measured two times. The hold is the time after a target is hidden for 10 s until the command leaves its bearing by 20° (cap 10 s). The release is the time from the noise-free crossing, where a look-alike’s evidence falls below the target’s, until the command reaches the target (cap 20 s). KT3 swept the ring gain, the accumulator leak and the state-machine timeout, with other parameters fixed at their tuned values. *3D mission (H1, H3)*. Seven faults at up to four severities gave 25 conditions per seed: target occlusion (1–10 s), a swapped look-alike on the reported spot (similarity 0.5–0.95), a half-silenced compass (one level), dead eye columns (10–40%), a false sun (30 – 120°), a speed-scale error and Dryden gusts. Success required a target stop, a return within 2 m of home and no collision. H1 used test seeds 1000–1299. *Recognition (H4)*. MuJoCo renders of a tent and three look-alikes in the target’s exact pose, from 0 – 180° azimuth at three distances, two new suns and camera jitter, on ten test layouts. The score is the area under the receiver operating characteristic curve (AUROC) of target versus look-alike probes per azimuth.

2.5 Study 2 modules

Study 2 was planned after diagnosing Study 1 (Section 3.1). Every sensory and motor module was given to all six arbiters. *Heading ring (H5)*. An insect-inspired model, designed in this study after the fly’s E-PG compass neurons [5, 6]: a ring of 32 neurons whose bump is shifted by the gyro yaw rate and corrected by the sun compass only while their conflict stays inside a tuned gate, with a slow remapping of the cue offset when the conflict persists. Its output is the path integrator’s compass input. Its rivals were a complementary filter and a Kalman filter with a gyro-bias state, each with a gate and the same 64-configuration budget; they were tested on 2D replays of recorded flights (KT5), scored by root-mean-square (RMS) heading error. *Home view memory and landing (H6)*. At take-off the drone flew one outward, rising loop around a striped home pad and stored its views one-shot in a second mushroom body, after insect nest learning flights [8, 9]. On return, once the path integrator placed home within a few metres, the best-matching view locked and a visual vector replaced the home vector. An engineered landing controller then centred the drone over the pad and descended at 0.3 m/s to touchdown. *Visual front end (H8)*. An insect-inspired engineered front end fed the unchanged mushroom body. It cropped the image around the object, applied log receptors, formed one luminance and two colour-opponent channels (approximated from RGB), normalised local contrast, added oriented-edge channels and pooled to 6×6 , after work showing that simple feature detectors support generalisation in a bee-like model [33]. The target learning flight became multi-scale: 30 views at three distances, two heights and three suns (the suns stand for repeated flights at different times of day, an assumption). An inspection rule stopped only after consistent familiarity from two viewpoints at least 20° apart. *Evidence-gated ring (H7)*. The ring’s recurrent gain drops when the input away from the bump exceeds the input at it, and hidden objects keep their last evidence, decaying with a tuned time constant. Every rival received the matching visibility channel. KT7 repeated KT3 with a third trial that penalises premature switching. *Variable altitude*. The drone searched at 3.5 m (chosen on training seeds between 2.5 and 3.5 m), descended to in-




Figure 1. The controller. Shading shows which parts are reimplemented published insect models, which are insect-inspired components designed in this study, and which are engineered; dashed outlines mark Study 2 additions. The ring hub (orange outline) is under test: each arm replaces it with one of five rival arbiters that receive the same inputs. The heading ring feeds the path integrator’s compass input, and its estimate also sets the frame of all bearings (dashed arrow). The avoidance reflex is added to the hub’s heading command.

spect a close candidate, climbed to return and descended to land.

2.6 Study 2 mission (H9)

Success now required a target stop, touchdown within 0.5 m of the pad centre and no collision within 300 s. The six arbiters were retuned on the Study 2 stack with the same budget (training seeds 100–104, then the best 4 on 100–119). The test ran 17 arms on test seeds 2000–2299 that had never been used: the six arbiters, the Study 1 brain, a Kalman-heading variant, one ablation per module, a 0 s learning flight, and two operator-tip sensitivity arms (no tip, with an outward spiral search; a poor tip of 30° and 30%). H9 required (i) fault-free success ≥ 0.90 , (ii) an AUC at least 0.15 above the Study 1 brain, and (iii) success ≥ 0.50 for every fault at its middle severity. The middle severity is the level closest to half the severity range, the more severe one on ties: occlusion 7 s, look-alike similarity 0.8, eye columns 30%, false sun 90°, speed-scale error 0.15 and moderate gusts; the half compass has one level. This rule was fixed before testing.

2.7 Post-study extension: size-aware search (exploratory)

After both studies we added a module, not pre-registered, because the Study 2 brain never stopped at a small target (a backpack). The brain estimates the target’s size from its learning-flight views: for a view at horizontal distance d_k and altitude z_k , the height is $H_k = d_k(\tan e_{\text{top}} - \tan e_{\text{bot}})$ from the elevations of the object’s top and base, and the size is $D_k = \rho_k \tan \alpha_k$ from its angular size α_k at range ρ_k , with angles quantised to the eye’s 2° acuity; $\hat{D}$ and $\hat{H}$ are medians over the views. With the eye’s 4° detection threshold plus a 2° margin, the detection range is $R = \hat{D}/\tan 6^\circ$, the search spiral’s pitch is $p = \min(p_0, 2R \cos 35^\circ)$ (p_0 the Study 2 pitch; 35° is the usable depression of the eye’s lower field) and the search altitude is $h = \hat{H}/2 + (p/2) \tan 35^\circ$, clipped to

1–3.5 m. The attention window scales with $\hat{D}$, the evidence window is re-calibrated in each mission from unstored look-back views at altitude h , and within range R the drone dips toward the learning-flight altitude to inspect a candidate. The test used seeds 3000–3199, never used before (200 per cell), the Study 2 ring brain with its frozen parameters, a tent or a backpack as target, the module off or on, and four conditions (no fault, false sun 90°, moderate gusts, a look-alike of the same kind at similarity 0.8), scored by Study 2 full success with Wilson intervals and paired on-minus-off bootstraps.

2.8 Pre-registration and statistics

For each seed and fault, success was integrated over severity with the trapezoid rule (fault-free run at 0, worst level at 1); this area is the AUC, and the aggregate AUC is the mean over the seven faults. Differences are paired bootstraps over seeds (10,000 resamples), proportions carry Wilson intervals, and pairwise tests were Holm-corrected. The pre-registered H1 statistic, the ring’s AUC minus the largest of the other five, is biased against any arm when arms tie. Amendment 4, made before the 3D test data were reduced and prompted by an independent review, therefore made a within-seed arm-permutation test of that statistic the decision rule; the bootstrap interval is still reported. Permutation p -values are Monte Carlo estimates from 5,000 permutations and are given as approximate.

Chronology. The Study 1 hypotheses were drafted before any data, but the record is an internal log without independent timestamps, and the public commits start after the kill tests and the recognition study. The 3D tuning freeze and Amendment 4 were commits made before the 3D test data were reduced. For Study 2, a SHA-256 hash of the plan was committed (5a6c16e) before any Study 2 data; the plan itself is unpublished. The tuned parameters were committed (8ba817f) before any Study 2 test seed ran. These commits are in the public repository history; commit dates are set by the author, and the history was published only after the studies were complete. The support for this order that can be

checked from outside is the cloud service’s server clock: the first test record is stamped 08:48:37 UTC on 2026-10-12, after the freeze commit time (08:44:29 UTC on 2026-10-12). For the H5 and H8 component tests no separate commit lies between the tuning freeze and the test, so their order rests on the internal log. Study 2 design choices (search altitude, landing constants, the inspection default) were fixed on training seeds and logged before testing. The results of H5, H7 and H8 were known before the mission test; all modules were nevertheless kept as pre-registered. Analyses added after the data are labelled post hoc.

3 Results

3.1 Study 1: the insect parts did not beat the rivals

Ring versus vote (H2). The pre-registered statistic was $D^* = 62.9^\circ$, above the 5° kill line. Post hoc, it depended on the vote-width grid ($48.9\text{--}62.9^\circ$) and fell to $5.4\text{--}6.9^\circ$ without equal-certainty cells, so, excluding those cells, the ring–vote difference is at the 5° threshold.

Stubbornness (H3, 2D). As the ring gain rose from 0, the hold of a hidden target lengthened from 0.42 s to at most 9.51 s and the release of a fading look-alike from 1.22 s to the 20 s cap from gain 0.7 (Fig. 2a). At every gain where the comparison was defined, the accumulator released sooner than the ring at the same hold time (-1.87 s [$-1.91, -1.83$] at gain 0.1, down to -18.7 s), so the pre-registered claim that the ring’s frontier lies above its rivals was rejected. The state machine’s zero release is a clamp: unclamped (post hoc), it switched -0.72 s before the crossing, at an evidence ratio of 0.99, and no trial penalised early switching. With the ring’s other parameters varied (post hoc, training seeds; Fig. 2b), the best ring configuration held 9.93 s and released in 3.10 s, still -1.19 s [$-1.24, -1.15$] behind the accumulator’s envelope. The margin shrank from tens of seconds to about one second but did not reverse.

3D mission (H1). Over 45,000 missions on 300 test seeds, the tuned ring (gain 0.022, almost no recurrence) reached an aggregate AUC of 0.623 [0.612, 0.632]; the state machine reached 0.631 and the accumulator 0.630. The pre-registered ring-minus-best-other difference was -0.0088 [$-0.0157, -0.0025$], and the Amendment 4 permutation test gave $p \approx 0.003$. H1 was supported: the arms differed, and the ring was lower. The effect was not specific to the ring: post hoc, the vector sum ($p \approx 0.003$) and the ring without recurrence ($p \approx 0.027$) failed the same test. Post hoc, the difference came from the swapped look-alike alone (fault AUC ring 0.491, state machine 0.552); without that fault all six arms lay between 0.643 and 0.645 (post hoc; Fig. 3). The false sun (0.12), the speed-scale error (0.36) and the half compass (0.49) defeated every arm alike, because arbitration cannot repair a wrong home vector.

Stubbornness in the mission (H3, 3D). The pre-registered comparison was inconclusive. The ring’s mean hold peaked at 3.25 s at gain 0.3. From gain 1 the ring stopped at a hidden target in 0.00 of runs and 0.99 of runs hit the time limit (post hoc), so high-gain points describe a latched controller. No arm’s mean hold exceeded 3.82 s although the target was hidden for 10 s, because the occlusion began on first sight within 10 m while the drone was closing (a ceiling of the metric). At

gain 0 the ring’s input stage beat the accumulator’s frontier by 1.06 s [0.44, 1.71]; recurrence did not add to that.

Recognition (H4). At the stored azimuth, pooled over distance, height jitter and two new suns, the mushroom body’s AUROC was 0.502 [0.494, 0.510] and raw pixels 0.506, with or without the learning flight. DINOv2 patch features, given an object-centred memory, reached 0.761 and 0.799 with the learning flight. H4 was rejected for the mushroom body.

Diagnosis. Four causes of failure were identified, and Study 2 addressed each: the path integrator trusted a single sun compass and speed signal; landing used the path integrator alone; the recogniser could not separate look-alikes; and the ring’s single gain coupled hold and release.

3.2 Study 2: the integrated brain completed most missions

Components. Across the false-sun and half-compass faults, the heading ring cut the RMS heading error from 57.6° with the Study 1 compass to 3.64° [3.38, 3.91] and beat the complementary filter (4.79°), but the Kalman filter did better (0.88° ; Kalman minus ring -2.76° [$-3.03, -2.50$]), so H5 was not supported (Fig. 8a). H7 was rejected. Once every arm had a visibility channel, the rivals held and released without a trade-off. In the premature-switch trial, at gains up to 0.7, the gated ring spent 1.5–4.9 s longer off the target than the state machine (whose time off target was 0.015 s). From gain 0.8 to 2 the excess was at most 0.03 s, within the non-inferiority margin, but there the ring lay below the state machine’s dominance frontier (sensitivity analysis; the pre-registered interpolation was undefined at these gains); at gain 3 the excess rose to 0.66 s. In the 2D mission world the gated ring scored 0.824 against 0.865 for the gated vote ($p < 0.001$). The visual front end lifted the mushroom body from chance (0.501 on the new layouts) to 0.686 [0.655, 0.720] at 0° , 0.185 [0.150, 0.219] above raw pixels, but missed the 0.75 criterion, so H8 failed (Fig. 4). A nearest-neighbour readout of the same features reached 0.732 and DINOv2 0.797; outside the flown arc the mushroom body fell below chance (0.32 at 90°), and the multi-scale flight did not beat the single goal view at 0° (-0.019 [$-0.059, 0.024$]).

Integrated mission (H9). Of the 300 fault-free test missions, the Study 2 brain completed the full mission in 0.990 [0.971, 0.997], meeting criterion (i) (Fig. 5). Its aggregate AUC was 0.852 [0.840, 0.863] against 0.009 for the Study 1 brain, which has no landing (difference 0.843 [0.831, 0.855]); scored by its own 2 m definition the Study 1 brain reached 0.627, still 0.225 [0.212, 0.239] lower. Criterion (ii) passed. At middle severity, success was 0.95 under occlusion, 0.86 with the half compass, 0.81 with dead eye columns, 0.98 under a false sun, 0.98 with a speed-scale error, 0.71 in moderate gusts and 0.39 with a swapped look-alike. Criterion (iii) failed on the look-alike alone, so H9 was not met. Look-alike success stayed between 0.34 and 0.40 at every similarity; the drone stopped at the wrong object in 0.097 of all missions (state machine 0.078, accumulator 0.074).

Post hoc, with both brains scored by the Study 1 2 m definition, the Study 2 brain was far better under the three faults that corrupt the path integrator (AUC difference: half compass 0.43 [0.41, 0.45], false sun 0.86 [0.85, 0.87], speed scale 0.60 [0.57, 0.64]). It was no better under occlusion (-0.01 [$-0.03, 0.00$]) or look-alikes (-0.02 [$-0.08, 0.03$]), and slightly worse




Figure 2. Study 1 stubbornness frontier in the 2D kill test (KT3). (a) Pre-registered sweeps on test seeds ($n = 300$ per point): ring gain g , accumulator leak time and state-machine timeout, other parameters at their tuned values; single points for the arms without a swept memory. Better is lower right (long hold, fast release). Error bars are ± 1.96 standard errors over seeds; release is clamped at 0 and capped at 20 s. (b) Post hoc: all 64 random-search configurations per arm on training seeds 0–99, with each arm’s Pareto envelope (lines); the circle marks the ring configuration closest to the accumulator’s envelope.




Figure 3. Study 1 mission: success-vs-severity AUC per fault for the six arms (test seeds 1000–1299, $n = 300$ seeds per arm and condition). Orange, the ring hub; blue, engineered rivals; grey, controls. Error bars are 95% percentile bootstrap intervals over seeds (10,000 resamples).

under dead eye columns (-0.06 [-0.09 , -0.04]) and gusts (-0.04 [-0.05 , -0.02]); over those four faults the mean difference was -0.034 [-0.050 , -0.018].

Arbitration. With the Study 2 stack, the arbiter barely mattered. The tuned gated ring (gain 0.69) scored 0.852; the five rivals ranged from 0.831 to 0.856 (state machine 0.856), and the permutation test of ring minus best other gave $p \approx 0.90$. Pairwise, only the static vote was lower than the ring after Holm correction (0.021, $p < 0.001$).

Ablations. Removing the home view memory cost 0.645 AUC [0.622, 0.669], the largest drop among the modules whose removal is not definitional (Fig. 6). Removing the landing controller cost 0.849, which is definitional because success requires a touchdown on the pad. Replacing the heading ring with the Study 1 compass cost 0.110 [0.100, 0.118]. The front end (0.030) and the evidence gating (0.020) added little, and the inspection rule nothing (-0.004 [-0.012 , 0.002]). In the closed loop, the heading ring beat the Kalman variant

by 0.047 [0.032, 0.062], the reverse of the 2D replay ranking (Fig. 8b); we did not analyse why.

Sensitivity arms. With a 0 s learning flight the drone made 0 target stops in 300 fault-free missions (AUC 0.000) and collided in 0.34 of all runs. Its evidence mapping had been calibrated with learning-flight memories, so this arm tests the calibrated design without its flight, not a design built for a single photo. Without an operator tip the AUC fell to 0.754 (fault-free 0.89); with a poor tip it was 0.834 (fault-free 0.95). The look-alike AUC was higher without a tip (0.53) or with a poor tip (0.54) than with the default tip (0.46). The look-alike stands on the reported spot, so this fault partly tests trust in the tip.

Landing (H6). Fault-free touchdowns had a median error of 0.009 m and a median touchdown speed of 0.30 m/s, the descent set-point (Fig. 7). In gusts, 0.66 of runs landed on the pad (median error 0.23 m). With the path integrator alone, 0.41 of fault-free runs landed on the pad (median error




Figure 4. Recognition under a swap control: AUROC of target versus look-alike probes on the same spot, per 15° azimuth bin (Study 2 test layouts 1010–1019; $n = 120$ target and 360 look-alike probes per bin). Memories: the Study 1 learning flight for the Study 1 mushroom body, the multi-scale flight for the others. Bands are 95% percentile bootstrap intervals over the 10 layouts (2,000 resamples), which likely under-cover with 10 clusters.

0.54 m), so the home view memory raised the landing rate in the direction H6 predicted (H6 had no numeric gate).

3.3 Post-study extension: size-aware search (exploratory)

This analysis is exploratory and was not pre-registered. Without the module the Study 2 brain never stopped at the backpack (target stops at most 0.000 in every condition; median mission time 300 s, the time limit). With it, full success with the backpack was 0.980 [0.950, 0.992] without faults, 0.985 [0.957, 0.995] under a false sun, 0.720 [0.654, 0.778] in gusts and 0.295 [0.236, 0.362] with a look-alike backpack, where the drone stopped at the look-alike in 0.695 of runs (Fig. 9); pooled over conditions, on minus off was 0.745 [0.715, 0.775]. For the tent the module changed nothing measurable (pooled 0.009 [−0.014, 0.031]; no fault 0.970 off, 0.975 on). The median size estimates were 1.01 m for the backpack and 3.12 m for the tent, against true diagonals of 1.00 m and 3.13 m, and the chosen search altitudes were 2.4 m and 2.7 m.

4 Discussion

What worked. The largest gain came from the home view memory used for landing. Loaded during a take-off learning flight and consulted only at the end of the home vector, it turned path-integrator landings (median error 0.54 m without faults) into touchdowns within centimetres (median 0.009 m), and it absorbed the odometry errors that defeated every Study 1 arm. This is consistent with the role given to nest learning flights in insects [8, 9] and with the engineering results of Bee-Nav and AntFlie [16, 18]. The heading ring was the second contributor: integrating the yaw rate and

gating the sun cue largely removed the false-sun and half-compass failures (Fig. 8b). A standard Kalman filter did better in replays and worse in the closed loop, so neither ranking should be read as a property of the insect design.

What failed. Look-alike discrimination did not improve. With a swapped look-alike at similarity 0.8, success on the 2 m definition was 0.39 for the Study 2 brain and 0.39 for the Study 1 brain, and wrong-object stops were 0.097 against 0.093; post hoc, on the same definition Study 2 beat Study 1 only on the three path-integrator faults and was worse on dead eye columns (−0.06) and gusts (−0.04). The inspection rule did not lower look-alike acceptance in renders. The recognition gain of the front end was real in renders but too small for the stop rule. DINOv2 was the only recogniser above the criterion at the stored azimuth, with an object-centred memory and a far larger pretrained model. The Study 2 brain also missed small objects, because its evidence window and attention window had been calibrated on tents; the exploratory size-aware module derives both, and the search height, from the learning-flight size estimate.

Arbitration mattered little here. In both studies the choice of steering arbiter changed aggregate success little (Study 1: 0.623 to 0.631; Study 2: 0.831 to 0.856, gated ring 0.852). In Study 1, recurrence traded hold against release, consistent with attractor-network theory [13, 14], but rivals with separate knobs did better on both, and the tuner chose almost no recurrence. In Study 2 the rivals reached the same success once they shared the visibility channel and the sensory stack. A ring hub is a workable arbiter for this task; on these data it is not a better one.

Limitations. Everything is simulation, and results describe these models in this simulator, not insects or real drones. The modules are published insect models, insect-inspired designs and engineered parts, not connectomes. Optic flow is computed geometrically from ray-casts, and the mission eye is an analytic ray-cast panorama, while the recognition study used MuJoCo renders; object ranges are noise-free. The kill tests are two-dimensional point-mass worlds. Both missions used a single forest world at one obstacle density. In Study 1 the tuning budget was nominal for the ring, the state machine alone used the visibility flag, and the hold metric in 3D saturated far below the 10 s occlusion. The 0 s learning-flight arm inherits an evidence mapping calibrated with flights. Recognition intervals come from 10 layouts and likely under-cover. Two faults in Study 1 were at zero success at every severity, so their weight in the aggregate carries no arm information. The size-aware search is exploratory and not pre-registered, was tested in the same single forest world, and leaves look-alikes unsolved. Provenance has three further gaps. Early Study 1 result files were stamped with commits that did not yet hold their scripts; they were later re-stamped and, where cheap, re-run. The Study 1 pre-registration is not independently timestamped, and the Study 2 plan is unpublished. Study 2 run records carry no commit because the cloud runs used an archive of the freeze commit; the summary records a SHA-256 hash of every simulation script, and the simulation code is identical at the freeze and results commits.




Figure 5. Study 2 headline: full-mission success (target stop, touchdown within 0.5 m of the pad, no collision) without faults and at each fault’s middle severity, test seeds 2000–2299 ($n = 300$ per bar). Orange, the Study 2 brain; blue and grey, the other arbiters on the same Study 2 stack; hatched, the Study 1 brain scored by its own success definition (return within 2 m of home), shown for reference only. Dashed lines are the pre-registered criteria (i) and (iii). Error bars are Wilson 95% intervals.




Figure 6. Study 2 ablations: aggregate AUC of the full brain minus each arm ($n = 300$ paired seeds). Orange, insect-inspired modules; blue, the Kalman heading swap; grey, engineered parts; light grey, sensitivity arms. Error bars are 95% paired bootstrap intervals (10,000 resamples). The landing ablation is definitional, and the 0 s learning flight used an evidence mapping calibrated with learning-flight memories.

5 Conclusion

In this simulator, the insect-model ring hub did not beat equal-budget engineered rivals, and the mushroom body did not recognise a look-alike on the target’s spot. A second, pre-registered design that added an insect-style heading ring and, above all, a nest learning flight with view-based landing completed 0.990 of fault-free missions and most faulted ones. Its gain over the first design came from the faults that corrupt the path integrator, and it still failed when a look-alike stood where the target was expected.




Figure 7. Landing in the Study 2 mission (test seeds 2000–2299, all fault conditions). (a) Cumulative distribution of horizontal landing error for missions with a target stop; errors below 1 mm are drawn at 1 mm, errors above 30 m (0.2% of these missions) lie beyond the axis, and the dotted line is the pad edge. (b) Touchdown speed of all touchdowns of the Study 2 brain, split by gusts; the descent set-point is 0.3 m/s. Sample sizes are in the legends.




Figure 8. Heading estimation. (a) RMS heading error in 2D replays of 300 recorded test flights per condition (KT5; sun offsets in degrees), log scale; error bars are 95% paired bootstrap intervals (10,000 resamples). (b) Full-mission AUC on the Study 2 stack with the heading ring, the Kalman filter, or the Study 1 sun compass ($n = 300$ seeds); error bars are 95% percentile bootstrap intervals over seeds (10,000 resamples).




Figure 9. Exploratory, not pre-registered: full-mission success with the size-aware search module off or on, for a backpack (solid) and a tent (hatched) as target, per condition (seeds 3000–3199, $n = 200$ per bar). Grey, module off; orange, module on; “0” marks bars at zero. Error bars are Wilson 95% intervals.

Data and code availability

Code, result files and run records are in the repository <https://github.com/aryawidjaja/mutant-brain-drone>.

Key commits: Study 1 3D tuning freeze ec18c87; Amendment 4 fa24ee6; Study 2 plan hash 5a6c16e; Study 2 tuning freeze 8ba817f; Study 2 test results 71a9ed1. Every result number in this paper is written by `scripts/export_paper.py` from `results/`, with its source in `numbers_sources.json`; the script's self-check re-reads every source and recomputes the main intervals from the run records. Commit ids recorded inside older result files map to these ids through `results/commit_map.tsv` (the history was rewritten once to remove superseded drafts).

Acknowledgements and licences

MuJoCo (Apache-2.0) and the MuJoCo Menagerie Crazyflie model (MIT) [27, 28]; vehicle constants from Förster [30] and Landry [29], cross-checked against the Crazyflie firmware (GPL-3.0; constants read, no code copied) [34]. The ring follows Sun, Yue and Mangan; their toolkit is MIT-licensed and was not copied [35]. The path integrator was reimplemented from the text of Stone and colleagues, whose code has no licence. DINOv2 ViT-S/14 weights (Apache-2.0) [26]; three.js (MIT) [36]. No funding was received. The author declares no competing interests.

AI assistance

AI coding and writing assistants were used for implementation and drafting; the author reviewed all code, results and text.

References

- [1] P. Ardin, F. Peng, M. Mangan, K. Lagogiannis, and B. Webb. Using an insect mushroom body circuit to encode route memory in complex natural environments. *PLoS Comput Biol*, 12(2):e1004683, 2016. doi: 10.1371/journal.pcbi.1004683.
- [2] T. Stone, B. Webb, A. Adden, N. B. Weddig, A. Honkanen, R. Templin, W. Wcislo, L. Scimeca, E. Warrant, and S. Heinze. An anatomically constrained model for path integration in the bee brain. *Curr Biol*, 27(20):3069–3085.e11, 2017. doi: 10.1016/j.cub.2017.08.052.
- [3] X. Sun, M. Mangan, and S. Yue. An analysis of a ring attractor model for cue integration, 2018. URL https://doi.org/10.1007/978-3-319-95972-6_49. In: *Biomimetic and Biohybrid Systems (Living Machines 2018)*, LNCS 10928, Springer, pp. 459–470.
- [4] X. Sun, S. Yue, and M. Mangan. A decentralised neural model explaining optimal integration of navigational strategies in insects. *eLife*, 9:e54026, 2020. doi: 10.7554/eLife.54026.
- [5] J. D. Seelig and V. Jayaraman. Neural dynamics for landmark orientation and angular path integration. *Nature*, 521(7551):186–191, 2015. doi: 10.1038/nature14446.
- [6] S. S. Kim, H. Rouault, S. Druckmann, and V. Jayaraman. Ring attractor dynamics in the *Drosophila* central brain. *Science*, 356(6340):849–853, 2017. doi: 10.1126/science.aal4835.
- [7] A. Rayshubskiy, S. L. Holtz, A. S. Bates, Q. X. Vanderbeck, L. Serratos Capdevila, V. Rockwell, and R. I. Wilson. Neural circuit mechanisms for steering control in walking *Drosophila*. *eLife*, 13:RP102230, 2025. doi: 10.7554/eLife.102230.
- [8] A. Philippides, N. H. de Ibarra, O. Riabinina, and T. S. Collett. Bumblebee calligraphy: the design and control of flight motifs in the learning and return flights of *Bombus terrestris*. *J Exp Biol*, 216(6):1093–1104, 2013. doi: 10.1242/jeb.081455.
- [9] T. S. Collett and J. Zeil. Insect learning flights and walks. *Curr Biol*, 28(17):R984–R988, 2018. doi: 10.1016/j.cub.2018.04.050.
- [10] R. Goulard, S. Heinze, and B. Webb. Emergent spatial goals in an integrative model of the insect central complex. *PLoS Comput Biol*, 19(12):e1011480, 2023. doi: 10.1371/journal.pcbi.1011480.
- [11] Y. Lu and B. Webb. Insect-inspired visual point-goal navigation, 2026. URL <https://arxiv.org/abs/2601.16806>. arXiv:2601.16806v4.
- [12] T. Hoinville and R. Wehner. Optimal multiguide integration in insect navigation. *Proc Natl Acad Sci USA*, 115(11):2824–2829, 2018. doi: 10.1073/pnas.1721668115.
- [13] K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J Neurosci*, 16(6):2112–2126, 1996. doi: 10.1523/JNEUROSCI.16-06-02112.1996.
- [14] A. Compte, N. Brunel, P. S. Goldman-Rakic, and X. J. Wang. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb Cortex*, 10(9):910–923, 2000. doi: 10.1093/cercor/10.9.910.
- [15] E. Bicho, P. Mallet, and G. Schöner. Target representation on an autonomous vehicle with low-level sensors. *Int J Robot Res*, 19(5):424–447, 2000. doi: 10.1177/02783640022066950.
- [16] D. Ou, J. J. Hagenaars, M. R. Jankowski, M. V. M. Firlfyn, C. De Wagter, F. T. Muijres, J. Degen, and G. C. H. E. de Croon. Efficient robot navigation inspired by honeybee learning flights. *Nature*, 653(8116):1039–1046, 2026. doi: 10.1038/s41586-026-10461-3.
- [17] M. Firlfyn, J. Hagenaars, and G. de Croon. Direct learning of home vector direction for insect-inspired robot navigation, 2024. URL <https://arxiv.org/abs/2405.03827>. arXiv:2405.03827.
- [18] G. Gattaux, H. Diallo, J. R. Serres, A. Wystrach, and F. Ruffier. AntFlie: frugal visual teach and repeat on narrow FoV micro-drones. *IEEE Robot Autom Lett*, 11(4):4617–4624, 2026. doi: 10.1109/LRA.2026.3667486.
- [19] T. van Dijk, C. De Wagter, and G. C. H. E. de Croon. Visual route following for tiny autonomous robots. *Sci Robot*, 9(92):eadk0310, 2024. doi: 10.1126/scirobotics.adk0310.
- [20] J. Gao, Y. Chen, Y. Pan, Y. Deng, J. Xiao, and M. Feroskhan. Flying to image-specified objects: 3D quadrotor navigation via cross-graph memory and viewpoint planning, 2026. URL <https://arxiv.org/abs/2606.29917>. arXiv:2606.29917.
- [21] W. Kuang, H. W. Ho, Y. Zhou, and S. A. Suandi. ForaNavi: insect-inspired online target-oriented navigation for MAVs in tree plantations, 2025. URL <https://arxiv.org/abs/2503.02275>. arXiv:2503.02275.
- [22] J. K. Rosenblatt. DAMN: a distributed architecture for mobile navigation, 1997. URL <https://doi.org/10.1080/095281397147167>. *J Exp Theor Artif Intell* 9(2–3):339–360.

- [23] J. Borenstein and Y. Koren. The vector field histogram: fast obstacle avoidance for mobile robots. *IEEE Trans Robot Autom*, 7(3):278–288, 1991. doi: 10.1109/70.88137.
- [24] M. Colledanchise and P. Ögren. Behavior Trees in Robotics and AI: An Introduction, 2018. URL <https://doi.org/10.1201/9780429489105>. CRC Press.
- [25] M. Negre Saura, R. Allmendinger, W. Pan, and T. Papamarkou. Spatial-aware decision-making with ring attractors in reinforcement learning systems, 2025. URL <https://neurips.cc/virtual/2025/poster/119004>. NeurIPS 2025 (poster); arXiv:2410.03119.
- [26] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, et al. DINOv2: learning robust visual features without supervision, 2023. URL <https://arxiv.org/abs/2304.07193>. arXiv:2304.07193.
- [27] K. Zakka, Y. Tassa, and MuJoCo Menagerie Contributors. MuJoCo Menagerie: a collection of high-quality simulation models for MuJoCo, 2022. URL http://github.com/google-deepmind/mujoco_menagerie. Software.
- [28] E. Todorov, T. Erez, and Y. Tassa. MuJoCo: a physics engine for model-based control, 2012. URL <https://doi.org/10.1109/IR0S.2012.6386109>. In: *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5026–5033.
- [29] B. Landry. Planning and control for quadrotor flight through cluttered environments, 2015. URL https://groups.csail.mit.edu/robotics-center/public_papers/Landry15.pdf. MEng thesis, Massachusetts Institute of Technology.
- [30] J. Förster. System identification of the Crazyflie 2.0 nano quadcopter, 2015. URL <https://doi.org/10.3929/ethz-b-000214143>. Bachelor thesis, Institute for Dynamic Systems and Control, ETH Zurich.
- [31] US Department of Defense. MIL-F-8785C, Military Specification: Flying Qualities of Piloted Airplanes, 1980. URL https://everyspec.com/MIL-SPECS/MIL-SPECS-MIL-F/MIL-F-8785C_5295/. Dated 5 November 1980.
- [32] MathWorks. Dryden Wind Turbulence Model (Continuous), 2026. URL <https://www.mathworks.com/help/aeroblks/drydenwindturbulencemodelcontinuous.html>. Aerospace Blockset documentation.
- [33] M. Roper, C. Fernando, and L. Chittka. Insect bio-inspired neural network provides new evidence on how simple feature detectors can enable complex visual generalization and stimulus location invariance in the miniature brain of honeybees. *PLoS Comput Biol*, 13(2):e1005333, 2017. doi: 10.1371/journal.pcbi.1005333.
- [34] Bitcraze. crazyflie-firmware, 2026. URL <https://github.com/bitcraze/crazyflie-firmware>. GPL-3.0; commit f6e0f0a; constants and mixer layout read, no code copied.
- [35] X. Sun. InsectNavigationToolkitModelling, 2024. URL <https://github.com/XuelongSun/InsectNavigationToolkitModelling>. Software (MIT licence).
- [36] three.js authors. three.js JavaScript 3D library, 2026. URL <https://github.com/mrdoob/three.js>. Software (MIT licence).