

How Much Route Fits in 3.2 kB?

Capacity Limits of Quantized Kenyon-Cell Familiarity

Roy Vaid

Abstract—Visual route familiarity can replace explicit maps with task-specific memory, but how much route can such a system actually retain when its state is limited to only a few kilobytes? We characterize this capacity for an insect-inspired panoramic familiarity circuit constrained to less than 3.2 kB of fixed-width algorithmic state. A 360° panorama is reduced to a 32-sector azimuthal representation, projected into a sparse population of Kenyon-cell-like units using procedural connectivity, and stored through quantized depression onto a single familiarity output. Circular shifts of the current view then recover the locally most familiar heading without storing individual image templates. On the Seville ant-habitat dataset, the best configuration achieves an 11.25° median heading error with 72.0% of queries recovered within one 11.25° bin using 3,130 bytes of algorithmic state. A memory-matched template baseline achieves substantially higher heading accuracy, revealing the cross-talk cost of superimposing multiple route views into a single weight vector. Despite a non-monotonic familiarity landscape, a local scanning controller converges in 70.0% of perturbed simulated closed-loop trials. These results characterize the capacity, robustness, and architectural limits of single-vector familiarity memory under severe resource constraints.

I. Introduction

Reliable autonomy becomes qualitatively different when sensing, memory, and compute are severely constrained. For tiny mobile robots, conventional mapping and localization pipelines can become prohibitive under severe on-board sensing and compute limits [1]. Conventional visual navigation systems typically retain feature descriptors, keyframes, learned representations, or metric map state whose resource requirements scale with environmental complexity [2]. For insect-scale agents, an alternative is to retain only information that is directly actionable for steering.

Visual route memories in insects provide one such model. Panoramic visual familiarity can support homing and route guidance without explicit Cartesian reconstruction [3], [4], and computational models of the insect mushroom body (MB) show that sparse Kenyon-cell (KC) representations combined with depression at KC-to-output synapses can distinguish familiar from novel views [5]. Simple comparison against stored panoramic views has also been shown computationally to provide a visual compass for route guidance in natural ant habitats [6]. Behavioral experiments further demonstrate that coarse panoramic skyline information provides sufficient directional information for ant navigation [7]. Sub-

sequent embodied implementations have demonstrated familiarity-based visual route navigation on small robotic platforms [8]–[10]. The insect central complex additionally provides well-supported mechanisms for representing heading and converting desired directions into steering commands [11].

Our goal is not to introduce familiarity-based mushroom-body navigation itself. Instead, we examine an engineering question motivated by insect-scale autonomy: what route-following capability remains when familiarity-module algorithmic state is restricted to under 3.2 kilobytes?

We introduce a minimal azimuthal Kenyon familiarity circuit with three resource-oriented properties:

- 1) An equirectangular panorama is reduced to a circular 32-sector signal such that discrete body rotations correspond directly to index shifts.
- 2) Sparse KC fan-in is generated procedurally from a deterministic integer hash, eliminating connectivity lookup tables (0 B lookup ROM) while executing directly from program flash.
- 3) Route memory resides in a single vector of quantized KC-to-output synapses, making learned-state size independent of the number of explicitly stored image templates.

We evaluate the circuit across route length, KC activity sparsity, and synaptic learning rate. Our primary contribution is an empirical characterization of the memory-capacity trade-off rather than a claim that mapless familiarity universally outperforms metric localization.

II. Minimal Panoramic Familiarity Circuit


A. Circular Azimuthal Visual Representation

Prior computational analyses have shown that field of view and angular resolution directly affect the navigational information available to panoramic view-matching strategies [12]. We assume an equirectangular panoramic observation with horizontal field of view 360°. For physical camera inputs, after grayscale conversion and a fixed vertical crop $[0.2H, 0.8H]$, the panorama is divided into $N = 32$ equal azimuthal sectors, where sector j is represented by its spatial mean:


$$p_j = \frac{1}{|\Omega_j|} \sum_{(u,v) \in \Omega_j} I(u,v). \quad (1)$$

For the Seville experiments, the lightweight renderer directly produces this 32-sector azimuthal profile without full ray casting: at each simulated viewpoint, polygon

(a) Minimal Panoramic Familiarity Pipeline



(b) Algorithmic State (k=128: 3,130 B Total)



(c) Heading Accuracy at 2,560-B Budget (30 Training Views in Seville)




Fig. 1. Minimal panoramic familiarity circuit and memory accounting under a 3.2 kB budget. (a) End-to-end panoramic familiarity loop: 360° panorama is downsampled into a 32-sector azimuthal int8 vector (32 B), projected via procedural hashing with odd strides (0 B ROM), filtered by streaming Top- k WTA (128–512 B scratch RAM), depressed onto a quantized MBON (2,560 B weights), and steered via a 5-heading circular scan. (b) Fixed-width algorithmic state breakdown: $k = 128$ requires 3,130 B total; $k = 32$ requires 2,746 B total, both strictly below the 3.2 kB ceiling. (c) Heading accuracy at 2,560-B learned-state budget across 30 training views in Seville LiDAR.

reflectance values are binned by relative azimuth and aggregated using inverse-distance weighting, with distance-dependent blending toward the sky value. To reduce sensitivity to global illumination offsets, the sector vector is centered and normalized via median absolute deviation (MAD). The normalized values are mapped to int8 via linear scaling and clipping:

$$x_j = \text{clip} \left(\left\lfloor 40 \times \frac{p_j - \text{med}(\mathbf{p})}{\text{MAD}(\mathbf{p}) + \epsilon} \right\rfloor, -128, 127 \right), \quad (2)$$

where $\epsilon = 10^{-6}$. The factor of 40 maps ± 3 MADs to the full 8-bit dynamic range. The vector $\mathbf{x} \in \text{int8}[32]$ occupies 32 bytes and preserves circular azimuth ordering: $\mathbf{x}(\theta + 11.25^\circ) \approx \text{roll}(\mathbf{x}(\theta), 1)$. This representation acts as an engineered, AOTU-inspired azimuthal pathway rather than a detailed anatomical model [13].

B. Sparse Kenyon Expansion

The 32-dimensional input projects to $M = 2560$ Kenyon-cell-like units, each receiving $K = 6$ azimuthal inputs: $z_i = \sum_{j \in \mathcal{C}_i} x_j$. The connection set $\mathcal{C}_i$ is generated on-the-fly from KC index i and a global seed. A deterministic integer hash (MurmurHash3 32-bit finalizer) generates an azimuthal offset $b_i \in [0, 31]$ and an odd stride $a_i \in \{1, 3, \dots, 31\}$:

$$\mathcal{C}_i = \{(b_i + a_i s) \bmod 32 \mid s = 0, \dots, 5\}. \quad (3)$$

Because $\text{gcd}(a_i, 32) = 1$, all 6 indices are guaranteed to be distinct without rejection sampling, eliminating lookup tables (0 B ROM). With 32 offsets and 16 odd strides, this construction admits 512 ordered pairs and at most 256 distinct unordered 6-sector receptive fields across the 2,560 KCs ($\sim 10\times$ population redundancy). Per view, Top- k spans 3–4 ($k = 32$) to 11–16 ($k = 128$) unique receptive fields (~ 12 effective visual features). Across 30 route views, 448 (17.5%, $k = 32$) to 1,244 (48.6%, $k = 128$) distinct KCs are activated, uniformly distributed across $[0, 2559]$ (mean 1268, median 1255). A global inhibitory operation selects the k largest activations via a streaming min-heap, breaking ties deterministically by lower KC index i , avoiding full activity vector buffering.

C. One-Shot Quantized Familiarity Learning

Each KC projects to one familiarity output through weight $w_i \in [0, 255]$ (initialized to 255). During route learning, active KCs undergo multiplicative depression: $w_i \leftarrow Q_b((1 - \eta)w_i)$ for $r_i = 1$, where Q_b denotes 8-bit rounding to uint8. For a query panorama, unfamiliarity (MBON output) is:

$$V(\mathbf{x}) = \frac{1}{k} \sum_{i \in \text{TopK}(\mathbf{z}, k)} w_i. \quad (4)$$

D. Circular Heading Recovery and Steering

At a query location, candidate orientations are evaluated by circularly shifting the azimuthal input: $s^* = \arg \min_{s \in \mathcal{S}} V(\text{roll}(\mathbf{x}, s))$, where $\mathcal{S} = \{0, \dots, 31\}$ during global initialization and $\mathcal{S} = \{-2, -1, 0, 1, 2\}$ during normal forward motion. The visual steering correction $\Delta\theta = 11.25^\circ s^*$ drives kinematic point-mass motion ($\theta_{t+1} = \theta_t + 0.6\Delta\theta$, forward step 0.05 m). Normalized novelty contrast margin is $M = (V_{\text{novel}} - V_{\text{familiar}})/V_{\text{novel}}$, where V_{novel} is unfamiliarity on off-route habitat views and V_{familiar} on the learned route view.

E. Fixed-Width Algorithmic State Budget

As summarized in Table I, logical runtime state remains strictly below 3.2 kB ($k = 128$ uses 3,130 B; $k = 32$ uses 2,746 B, excluding runtime interpreter, program code, transient sensor DMA buffers, and ≤ 256 B preprocessing temporaries multiplexed into the Top- k scratch buffer prior to projection). Procedural connectivity eliminates 15.36 kB of lookup table ROM.

III. Experimental Evaluation

We evaluate the circuit on four core questions: (RQ1) heading recovery in LiDAR-reconstructed natural terrain; (RQ2) route-length scaling and discriminability against a memory-matched template baseline; (RQ3) robustness to illumination and occlusions; and (RQ4) closed-loop route convergence.




Fig. 2. Directional heading recovery error and route discriminability on the Seville LiDAR ant habitat across 5 random seeds (mean $\pm$ SEM). (a) Route-length scaling: heading error falls from 46.8° ($T = 5$) to 17.4° – 21.2° ($T \in [30, 35]$; 22.0° at $T = 40$), with a persistent gap relative to the Budgeted Template Memory baseline (≤ 80 -view capacity). Inset: Inter-route AUROC summary table. (b) Inter-route AUROC curves across route lengths $T \in [5, 40]$: sparse KCs outperform templates on short routes ($T = 5$, 0.774 vs. 0.656), while templates dominate on longer routes ($T \geq 20$, 0.700–0.729 vs. 0.578–0.669).

TABLE I
Fixed-Width Algorithmic State Accounting

Component	Data Type	$k = 128$	$k = 32$
Weights ($\mathbf{w}$)	uint8[2560]	2,560 B	2,560 B
Azimuth ($\mathbf{x}$)	int8[32]	32 B	32 B
Top- k Heap	int16[k]	256 B	64 B
Top- k Indices	uint16[k]	256 B	64 B
Scan Buffer	uint16[5]	10 B	10 B
Controller State	struct	16 B	16 B
Total State		3,130 B	2,746 B
Procedural ROM	Hash	0 B	0 B
LUT ROM (Ablation)	uint8[2560x6]	(15,360 B)	(15,360 B)

A. Benchmark Environment & Baselines

Evaluations are conducted on the Seville 2009 Terrestrial LiDAR 3D Ant Habitat (world5000_gray.mat, reconstructed from 56 terrestrial laser scans covering 1,018 m² of natural Cataglyphis velox foraging terrain [14]) and real ant homing trajectories (Ant1_Route1, Ant1_Route2) recorded by Mangan and Webb [15]. We benchmark against Budgeted Template Memory, which stores raw 8-bit, 32-sector profiles under the identical 2,560-byte learned weight budget (≤ 80 views), performing circular exhaustive L_1 distance matching with int16 difference accumulation.

IV. Results

A. Directional Familiarity

On a 3-meter segment of Ant1_Route1 (60 positions, 30 training views, 30 held-out test views), each held-out view was evaluated across 8 rotation offsets ($s \in \{0, 4, \dots, 28\}$ sectors, every 45°), yielding $N = 240$ queries for Budgeted Template Memory and $N = 2,400$ across 10 random seeds for Kenyon Cell Memory (providing the denominator for 1-bin recovery percentages). The Budgeted Template Memory baseline achieved a median circular error of 0.00° (IQR: 11.25°), with 90.0% (216/240) within one 11.25° bin and margin $M = 0.3249$. Proposed Kenyon Cell Memory ($k = 128, \eta = 0.65$) achieved 11.25° median error (IQR: 22.50°), with 72.0% (1,728/2,400) within one bin and margin $M = 0.4733$ (6.33% saturated weights) in 3,130 B state. At

$k = 64, \eta = 0.10$: median error was 11.25° (IQR: 22.50° , 67.0% within one bin); at $k = 32, \eta = 0.10$: median error was 11.25° (IQR: 11.25° , 52.0% within one bin) in 2,746 B state. Superimposing memories into a single KC weight vector causes route-wide visual interference, visible in the proportion within one bin rather than the median.

B. Route-Length Scaling & Inter-Route Discriminability

We compared familiar views (Ant1_Route1) against a novel route in the same habitat (Ant1_Route2) across route lengths $T \in [5, 40]$ (5 random seeds, Fig. 2):

The Budgeted Template Memory baseline achieved AUROC of 0.656 at $T = 5$ rising to 0.729 at $T = 40$, with mean heading error falling from 15.75° ($T = 5$) to 6.38° ($T = 30$; 6.47° at $T = 40$). For single-vector Kenyon Cell Memory ($k = 128, \eta = 0.65$), AUROC was 0.558 at $T = 5$ and 0.425 at $T = 10$ (recovering to 0.605–0.669 for $T \geq 20$; on trained views AUROC is 0.998, but test views along early sharp turns experience visual mismatch while cross-talk depresses novel-view KCs), with mean heading error improving from 46.8° ($T = 5$) to 17.36° ($T = 35$), 21.23° ($T = 30$), and 21.99° ($T = 40$). With sparser activation ($k = 64, \eta = 0.10$), AUROC spanned 0.492–0.774 (0.578 at $T = 40$) with mean error falling from 76.5° ($T = 5$) to 20.25° ($T \in [30, 35]$; 26.44° at $T = 40$). For $k = 32$ ($\eta = 0.10$), AUROC ranged 0.492–0.725 (0.568 at $T = 40$) and error dropped from 123.8° ($T = 5$) to 30.98° ($T = 30$; 37.29° at $T = 40$). The 123.8° error at $T = 5$ stems from hash collisions under severe undersampling (5.6% KCs depressed, 0.4% margin), clustering queries deterministically at discrete offsets (+11, +10, −13, −15, −6 bins) matching spurious mixtures of depressed KCs.

This scaling reveals an informative memory–capacity trade-off: while templates maintain lower heading error across all route lengths (6.4° – 15.8° vs. 17.4° – 46.8°), route discriminability exhibits a crossover (Fig. 2(b)). Sparser KCs at $T = 5$ reach higher AUROC (0.774 for $k = 64$, 0.725 for $k = 32$) than templates (0.656) due to sharp feature selectivity under low synaptic occupancy;




Fig. 3. (a) Directional unfamiliarity tuning curves on Seville LiDAR held-out test views, showing ensemble mean $\pm$ SEM (familiar route views exhibit a distinct minimum at 0° , while novel habitat views remain elevated; dataset-derived rendering variability reflects vegetation density). (b) Seville 3D LiDAR desert arena showing reference ant homing route (Ant1_Route1) and all 30 closed-loop robot simulation trials (21 converging, 9 deflected into local minima; 70.0% success rate).

as route length scales ($T \geq 20$), cumulative superposition cross-talk degrades KC discriminability (0.578–0.669 at $T = 40$), where templates consistently prevail (0.700–0.729).

C. Perturbation Robustness

On panoramic camera image sequences (360×90 px) prior to pooling, brightness scaling ($0.5 \times -1.5 \times$) retained 100.0% 1-bin accuracy via median centering; gamma scaling ($0.6-1.5$) yielded 96.7%–100.0%; Gaussian blur ($\sigma \leq 4$ px) and additive noise ($\sigma \leq 0.10$) retained 100.0% accuracy. In contrast, azimuthal occlusion dropped accuracy to 23.3% at 10% occlusion and 0.0% at $\geq 25\%$, revealing sensitivity to sector-wide blockages.

D. Closed-Loop Navigation & Catchment Basin

In the Seville 3D LiDAR world, across lateral transects perpendicular to the route, the spatial familiarity surface is not globally monotonic: 54.0% of lateral steps violate monotonicity, exhibiting 4.4 false minima per transect ($\rho = -0.151$ between lateral offset and 32-heading minimum unfamiliarity; $+0.379$ route-aligned). In 30 closed-loop trials under perturbations (± 0.5 m lateral offset, $\pm 22.5^\circ$ heading offset), the local-scan controller achieved a 70.0% homing success rate (21/30 converged within 80 steps; median lateral tracking error 0.55 m, p95 1.25 m, Fig. 3(b)), confirming that closed-loop steering succeeds via locally consistent 5-heading directional contrast despite a rugged global landscape.

E. Resource Cost & Microbenchmarks

On Apple Silicon CPU (Python 3.11), desktop vectorized evaluation required 44.6 μ s median; local 5-heading scan required 286.5 μ s; global 32-heading scan took 2.03 ms. A pure-Python streaming min-heap pass ($O(k)$ scratch RAM, 0 B ROM) took 6.51 ms over 2,560 KCs. Fixed algorithmic state (excluding interpreter and DMA buffers) is 3,130 B ($k = 128$) and 2,746 B ($k = 32$), strictly below 3.2 kB, while procedural projection eliminates 15.36 kB of connection ROM.

V. Discussion & Limitations

Our results characterize the empirical memory–capacity trade-off of insect-inspired familiarity circuits under severe constraints:

1) Single-vector cross-talk: Superimposing all route views into a single MBON weight vector incurs cross-talk between distant route segments: templates retain superior heading precision across all lengths, whereas single-vector superposition degrades route discriminability beyond the crossover at $T \approx 20$.

2) Vertical cue loss: Compressing panoramas to 32 azimuthal sectors discards elevation cues, leaving the system vulnerable to sector-wide azimuthal occlusion ($\geq 10\%$).

3) Catchment discontinuity: Sparse thresholding introduces local minima (54.0% monotonicity violations), yet local 5-heading scanning achieves 70.0% convergence from realistic release perturbations.

4) Implementation sensitivity: Sign consistency between sensory shifts and steering kinematics is critical; a flipped sign silently converts a convergent controller into a divergent one. Automated property tests should verify that sensory offsets produce counter-acting corrections.

5) Renderer realism: Azimuthal binning without full ray casting approximates panoramic rendering, though inverse-distance weighting avoids artificial discontinuities of nearest-polygon aggregation.

VI. Conclusion

We presented a quantized Kenyon-cell familiarity memory for panoramic route navigation operating under a 3.2 kB algorithmic state budget. Pairing circular azimuthal pooling with procedural hash projections and quantized LTD eliminates connection matrices while retaining directional familiarity signals. In Seville LiDAR terrain, template storage provides superior heading precision and long-route discriminability under the same budget, while the Kenyon-cell circuit achieves 70.0% closed-loop convergence by exploiting local directional contrast rather than global monotonicity.

References

- [1] K. N. McGuire, C. De Wagter, K. Tuyls, H. J. Kappen, and G. C. H. E. de Croon, “Minimal navigation solution for a swarm of tiny flying robots to explore an unknown environment,” *Sci. Robot.*, vol. 4, no. 35, p. eaaw9710, 2019.
- [2] C. Cadena et al., “Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age,” *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [3] J. Zeil, M. I. Hofmann, and J. S. Chahl, “Catchment areas of panoramic snapshots in outdoor scenes,” *J. Opt. Soc. Am. A*, vol. 20, no. 3, pp. 450–469, 2003.
- [4] B. Baddeley, P. Graham, P. Husbands, and A. Philippides, “A model of ant route navigation driven by scene familiarity,” *PLOS Comput. Biol.*, vol. 8, no. 1, p. e1002336, 2012.
- [5] 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.*, vol. 12, no. 2, p. e1004683, 2016.
- [6] A. Philippides, B. Baddeley, K. Cheng, and P. Graham, “How might ants use panoramic views for route navigation?” *J. Exp. Biol.*, vol. 214, no. 3, pp. 445–451, 2011.
- [7] P. Graham and K. Cheng, “Ants use the panoramic skyline as a visual cue during navigation,” *Curr. Biol.*, vol. 19, no. 20, pp. R935–R937, 2009.
- [8] J. C. Knight et al., “Insect-inspired visual navigation on-board an autonomous robot: Real-world routes encoded in a single layer network,” in *Proc. Conf. Artif. Life (ALIFE)*, 2019, pp. 60–67.
- [9] G. G. Gattaux, A. Wystrach, J. R. Serres, and F. Ruffier, “Route-centric ant-inspired memories enable panoramic route-following in a car-like robot,” *Nat. Commun.*, vol. 16, no. 1, p. 8328, 2025.
- [10] Y. Lu, J. Cen, R. Alkhoury Maroun, and B. Webb, “Insect-inspired embodied visual route following,” *J. Bionic Eng.*, vol. 22, no. 3, pp. 1167–1193, 2025.
- [11] T. Stone et al., “An anatomically constrained model for path integration in the bee brain,” *Curr. Biol.*, vol. 27, no. 20, pp. 3069–3085.e11, 2017.
- [12] A. Wystrach, A. Dewar, A. Philippides, and P. Graham, “How do field of view and resolution affect the information content of panoramic scenes for visual navigation? A computational investigation,” *J. Comp. Physiol. A*, vol. 202, no. 2, pp. 87–95, 2016.
- [13] A. Honkanen, A. Adden, J. da Silva Freitas, and S. Heinze, “The insect central complex and the neural basis of navigational strategies,” *J. Exp. Biol.*, vol. 222, no. Suppl_1, p. jeb188854, 2019.
- [14] H. Risse, M. Mangan, L. Stürzl, and B. Webb, “Software to convert terrestrial lidar scans of natural environments into photorealistic meshes,” *Environ. Model. Softw.*, vol. 99, pp. 88–100, 2018.
- [15] M. Mangan and B. Webb, “Spontaneous formation of multiple routes in individual desert ants (*Cataglyphis velox*),” *Behav. Ecol.*, vol. 23, no. 5, pp. 944–954, 2012.