

Geometric Phase Transition Enables Extreme Hippocampal Memory Capacity

Prashant C. Raju

rajuprashant@gmail.com

Significance Statement

Biological memory systems face a fundamental scaling paradox: how to expand information capacity without a proportional increase in neurons. Resolving this paradox in food-caching birds, we discover that extreme spatial memory arises from a phase transition in how neural activity is collectively organized: below the critical threshold, networks suffer catastrophic interference; above it, crystalline geometry enables storage beyond 1,000 locations. High-capacity circuits maintain this rigid geometry through synergistic excitatory-inhibitory dynamics, confirmed across 10,000 network configurations. To sustain the transition, the system pays a “geometric tax”: representational redundancy that stabilizes the manifold against noise. These findings establish geometric stability as an organizing principle linking circuit architecture to cognitive limits, with implications for any system that must store many memories without interference.

Abstract

Memory systems can store vastly different amounts of information despite similar hardware constraints. Here, we show that superior spatial memory emerges from a discrete stiffening of hippocampal population geometry—a transition from disorganized to crystalline collective coding. Comparing food-caching chickadees to non-caching zebra finches, we found that the caching hippocampus maintains a topologically rigid, “crystalline” geometry with significantly higher geometric stability (Shesha: 0.245 vs. 0.166) and nearly two-fold greater temporal coherence (Shesha: 0.393 vs. 0.209), while the non-caching hippocampus resembles a disorganized “mist.” This stability is actively constructed by synergistic circuit dynamics: excitatory neurons form the spatial scaffold while inhibitory populations contribute orthogonal decorrelation, a circuit motif in which excitatory and inhibitory populations occupy largely non-overlapping representational subspaces. A double dissociation with Valiant’s Stable Memory Allocator, a model predicting that dedicated neuron ensembles underlie each memory, confirms this advantage reflects continuous topological organization rather than discrete neuron allocation: caching networks exhibit near-zero split-half allocation reliability despite their geometric superiority. Computational modeling across 10,000 configurations reveals topological rigidity as the mathematical prerequisite for scale: crystalline codes sustain high-fidelity readout beyond $M = 1,000$ locations while mist codes fail below $M = 10$, a >100 -fold capacity advantage. This capacity requires a 169-fold representational redundancy: a “geometric tax” stabilizing the manifold against biological noise. These results establish geometric stability as a candidate organizing principle of biological memory: evolution achieves high-capacity memory not by proliferating neurons, but by engineering the geometry of the neural code itself.

The capacity of biological memory systems presents a fundamental puzzle: how do organisms store vastly different amounts of information using apparently similar neural hardware (Bailey, Bartsch, and Kandel 1996; Roxin and Fusi 2013; Squire 1992)? Theoretical frameworks have long sought to quantify these limits (Fusi 2024; Marr 1971; Treves and Rolls 1994), yet the network-level principles enabling high-capacity storage while preventing catastrophic interference (Bakker et al. 2008; French 1999; McCloskey and Cohen 1989) as memory load grows remain poorly understood. The hippocampus, conserved across vertebrates (Jacobs 2003; Thome et al. 2017), supports spatial memory demands ranging from modest to extraordinary scale, making it an ideal system in which to investigate how biological circuits navigate this capacity-interference tradeoff.

We address this question by exploiting a remarkable natural experiment: the ethological divide between food-caching birds and non-caching species (H. L. Payne, G. F. Lynch, and Aronov 2021). Black-capped chickadees cache and retrieve thousands of spatial locations across a foraging season—demands that drive pronounced hippocampal enlargement (Garamszegi and Lucas 2005; Sherry, Vaccarino, et al. 1989)—while zebra finches do not cache and show no comparable spatial memory requirements (Brodbeck and Shettleworth 1995; Clayton and John R Krebs 1995; J R Krebs et al. 1989; Pravosudov et al. 2015; Sherry and Vaccarino 1989; Smulders 2017). Both species possess hippocampal neurons with spatial tuning (Aronov, Nevers, and Tank 2017; Marcus K. Benna and Fusi 2021; Gulli et al. 2019; H. L. Payne, G. F. Lynch, and Aronov 2021), yet their behavioral capacities differ by orders of magnitude (Balda and Kamil 1992). This divergence provides a controlled system for isolating the neural basis of extreme memory capacity.

We propose that superior memory in caching birds does not arise from simply scaling neural resources, such as allocating more neurons, expanding ensembles, or increasing firing rates (Marcus K Benna and Fusi 2016; Dayan and L. F. Abbott 2005; Fusi and L F Abbott 2007; Valiant 2005). Instead, we demonstrate that extreme capacity emerges from a topological phase transition in the hippocampal population code: a discrete reorganization in how neural activity is collectively structured across the population. Chickadee hippocampal activity organizes onto a geometrically rigid, continuously organized manifold (Bernardi et al. 2020; Boyle et al. 2024; Burak and Fiete 2009; Chaudhuri et al. 2019; Courellis et al. 2024; Cunningham and Yu 2014)—a crystalline code—exhibiting high geometric fidelity and temporal coherence across sessions. The finch population code, by contrast, resembles a disorganized mist: low-dimensional structure is absent and spatial relationships between locations are poorly preserved in the population response. This geometric rigidity is not incidental; it is actively constructed by specific excitatory-inhibitory circuit dynamics that expand the representational dimensionality of the population beyond what either cell class achieves alone.

Computational modeling across ten thousand network configurations confirms that this topological transition is a prerequisite for high-capacity storage rather than a correlate of it (Burak and Fiete 2012). Below a critical stability threshold, memory systems fail catastrophically at moderate loads; above it, reliable recall extends to scales orders of magnitude larger (Fusi, Miller, and Rigotti 2016; Rigotti et al. 2013). To sustain the crystalline geometry, the network pays a 169-fold representational redundancy—the “geometric tax”—relative to the minimum needed to encode a single spatial point (Figure 3). Yet the core finding is clear: evolution achieves extreme memory capacity not by proliferating neurons, but by engineering the geometry of the population code itself (Jazayeri and Ostojic 2021; Saxe, Nelli, and Summerfield 2020). We argue that geometric stability (P. C. Raju 2026a) is a candidate organizing principle of biological memory that may generalize well beyond the avian system studied here.

Results

A Topological Phase Shift in the Population Code

Crucially, this stability is explicitly topographic. The chickadee RDM, sorted by physical arena location, displays a highly structured pattern: population vectors at nearby spatial locations are highly similar (forming a dark, low-distance diagonal), while distant locations are sharply distinct (forming warm-toned, high-distance off-diagonal blocks). The finch RDM under identical sorting is comparatively featureless (Figure 1B, C; see Figure 4A, B for hierarchical cluster-based sorting). Using a non-parametric Mantel test evaluated over 1,000 permutations, we found that chickadee population codes maintain a significantly stricter correspondence between neural geometry and the physical layout of the environment ($r_{\text{mean}} = 0.420$, 39/39 sessions significant) compared to the less structured, “mist-like” geometry of the finch ($r_{\text{mean}} = 0.230$, 8/8 sessions significant; $p < 0.001$; Figure 1D, E). This structural anchoring translates directly into temporal robustness: chickadee spatial maps exhibited higher within-session temporal cross-correlation compared to finches (0.399 vs. 0.241, $p < 0.001$), as if the population is repeatedly reading from the same printed map rather than reconstructing one from scattered landmarks.

Dissociation from Tuning Heterogeneity and Sample Size Comparing neural geometries across species requires strict controls for physiological and experimental confounds. First, to ensure the observed topological rigidity was not a statistical artifact of differences in recorded neuron counts, we randomly downsampled chickadee ensembles to match the median finch neuron count ($N = 6$; 30 of 39 chickadee sessions eligible). The geometric stability gap persisted with a medium effect size (Cohen’s $d = 0.560$, $p = 0.025$; Figure 6A), confirming that the difference reflects genuine representational structure rather than recording depth. Furthermore, *in silico* map-shuffling completely abolished geometric stability in both species ($\text{Shesha}_{\text{FS}} \approx 0.000$). Circular shifting, which preserves individual neuron statistics but disrupts the popu-

lation code, substantially reduced stability in the highest-yield sessions (chickadee: 0.241 \rightarrow 0.154; finch: 0.171 \rightarrow 0.072; **Figure 6B**), verifying that maximal rigidity requires the authentic spatial continuity of the full population code rather than any incidental statistical property of individual neurons.

Second, we asked whether the manifold advantage could arise spuriously from simple spatial tuning heterogeneity. We evaluated traditional “Stable Memory Allocation” (SMA) metrics (Valiant 2012), which hypothesize that networks prevent interference by allocating fixed, consistent ensembles of neurons to each memory (**Methods A.6**). Rather than supporting this discrete model, the data revealed a striking double dissociation: caching networks exhibited highly heterogeneous place field sizes (Coefficient of Variation = 1.03 vs. 0.23, $p = 0.003$; **Figure 5A**), and chickadee split-half allocation reliability was near zero and significantly *lower* than that of finches (mean $r = -0.064$ vs. 0.264, $p = 0.022$ one-tailed; **Figure 5B**). This indicates that caching neurons are actively inconsistent in their ensemble membership across neuron subsets. Caching hippocampus does not operate by reserving dedicated cells for each location; rather, geometric stability emerges from the population-level manifold structure while individual neuron assignments remain fluid.

Third, we examined whether the species difference is detectable by established linear metrics. Split-half population vector (PV) correlation, which measures the reproducibility of mean firing rates per location across neuron subsets, showed no species difference and numerically favored the finch (chickadee: 0.058; finch: 0.167; $p = 0.894$). This dissociation is theoretically significant: it confirms that the chickadee advantage is not carried by stronger individual place fields, but resides in the higher-order relational structure of the population code. Geometry-sensitive measures—our geometric stability metric ($\text{Shesha}_{\text{FS}}$) and an independent CCA stability measure (chickadee 0.554 vs. finch 0.481, $p = 0.087$; **Figure 6D**)—begin to capture this rigidity that the linear rate-map metric cannot (**Figure 6C**), validating the hypothesis that food-caching memory capacity is encoded in manifold topology: that is, in the geometry of the collective code rather than in the precision of any individual neuron’s tuning.

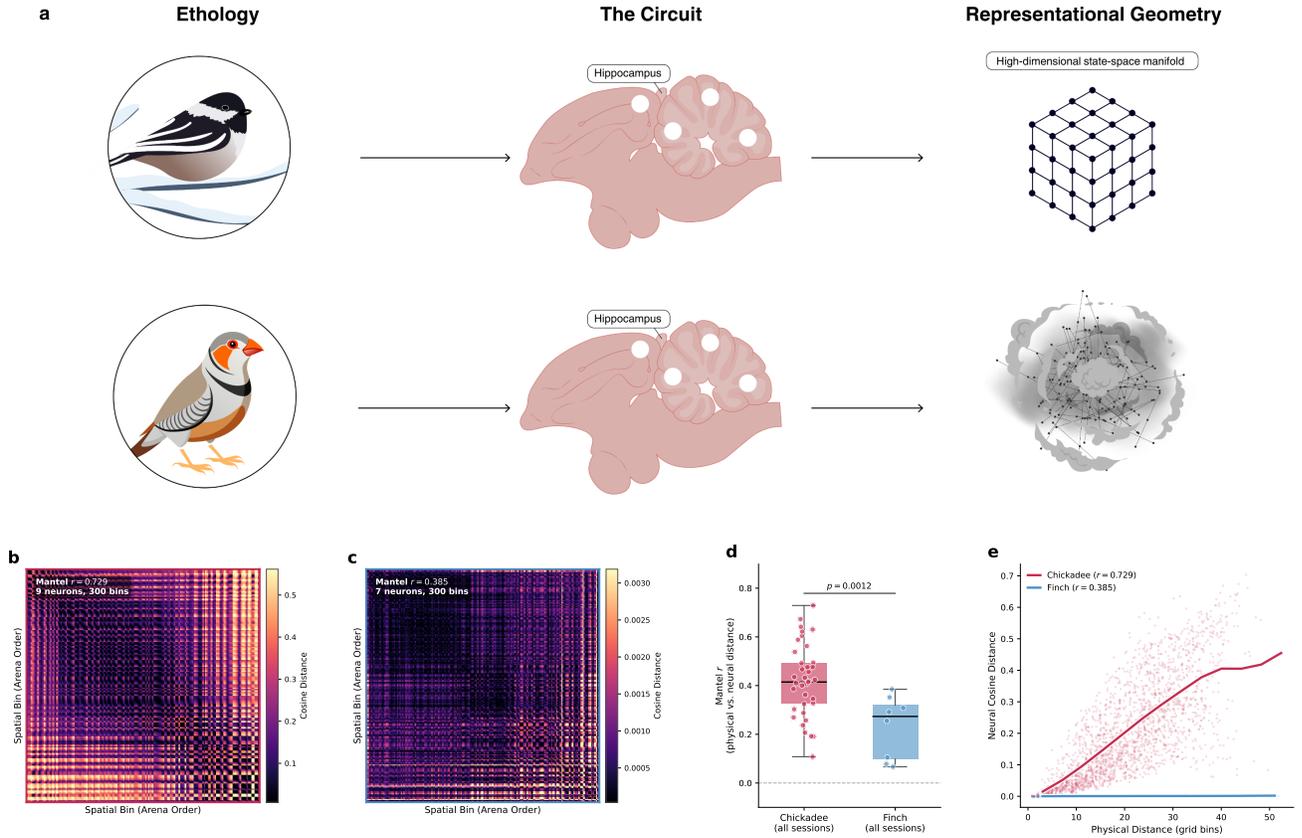


Figure 1: Geometric rigidity distinguishes hippocampal population codes in food-caching and non-caching birds. (a) Schematic overview. The food-caching black-capped chickadee (*Poecile atricapillus*) and non-caching zebra finch (*Taeniopygia guttata*) possess hippocampi of comparable neuron count and gross connectivity, yet the caching species must store and retrieve thousands of cache locations across a winter season. Conceptually, the chickadee hippocampus operates as a “Geometric Crystal” — a low-dimensional, topologically rigid manifold — whereas the finch hippocampus resembles a “Mist” of disorganized representations susceptible to catastrophic interference. (b, c) Representational dissimilarity matrices (RDMs) for the best-session chickadee (b; 9 neurons, 300 spatial bins, $r_{\text{Mantel}} = 0.729$) and finch (c; 7 neurons, 300 bins, $r_{\text{Mantel}} = 0.385$), sorted by row-major arena order. Matrix entries reflect cosine distance between population vectors at pairs of spatial locations. Note that color scales differ between species: the chickadee RDM spans [0, 0.6] while the finch RDM spans [0, 0.003], reflecting the greater dynamic range of the chickadee spatial code. A warm block-diagonal structure is visible in the chickadee RDM, indicating that physically adjacent locations have dissimilar population vectors from distant locations; the finch RDM is comparatively featureless. (d) Mantel r (Spearman correlation between pairwise physical distance and pairwise neural cosine distance) across all sessions. Chickadees show significantly higher physical–neural correspondence than finches (Mann–Whitney U , $p = 0.0012$; chickadee median $r = 0.42$, finch median $r = 0.27$). Individual session values overlaid as points. (e) Mantel scatter for the best session per species. Each point represents one pair of spatial bins; trend lines show binned means. The chickadee exhibits a monotonically increasing physical–neural distance relationship ($r = 0.729$); the finch trend is flat and near zero ($r = 0.385$), consistent with a disorganized population code.

Orthogonal Inhibition Stiffens the Geometric Attractor

Excitatory-Inhibitory Dissociation

Excitatory cells were the primary spatial information carriers, encoding nearly 20-fold more spatial information per spike than inhibitory cells (E: 0.169 bits/spike; I: 0.009 bits/spike; $p < 10^{-74}$; **Figure 7A**), and E subpopulations maintained significantly stronger topographic structure at the session level (Mantel $r_E = 0.418$ vs. $r_I = 0.343$, $p = 0.023$; **Figure 7C**). Yet despite carrying almost no spatial information individually, inhibitory cells exhibited within-session temporal stability statistically indistinguishable from excitatory cells ($\bar{x}_{\text{corr},E} = 0.399$ vs. $\bar{x}_{\text{corr},I} = 0.405$, $p = 0.659$; **Figure 7B**) and were spatially selective at a slightly higher rate (68.0% vs. 62.5%). Furthermore, the geometric stability of I cells is negatively coupled with E cells across sessions ($r = -0.333$, $p = 0.381$; **Figure 2B**), inconsistent with a redundancy model where both cell types track the identical spatial signal. The combination of near-zero bits per spike, negative geometric coupling, and high temporal stability is consistent with broad, low-amplitude place fields that tile the arena without sharp spatial

tuning—a profile more consistent with a global gain or contextual signal than with precise location coding.

Emergent Stability Through Selective Orthogonality

To characterize the geometric relationship between E and I population codes, we computed the principal angles between their respective top-3 principal subspaces ($n = 21$ chickadee sessions). The angle structure revealed a “one shared, two orthogonal” architecture (**Figure 2C**). The first principal angle averaged 14° , indicating one strongly shared representational dimension consistent with global gain co-modulation between E and I circuits. The second and third angles averaged 64° and 82° respectively—the latter approaching the $\sim 83^\circ$ expected for random independent subspaces in the ambient neural state space—indicating that the spatial map dimensions are statistically independent between cell types. This pattern was robust across sessions: 16 of 21 sessions exhibited an overall mean principal angle exceeding 45° , with per-session means ranging from 39.2° to 66.8° .

Mechanistic Interpretation

This selective orthogonality has a direct mechanistic interpretation. If inhibitory cells acted purely as subtractive normalization, they would occupy the same subspace as E cells and all principal angles would approach 0° . Instead, the observed pattern—one shared axis of co-modulation embedded in an otherwise orthogonal relationship—is consistent with divisive normalization: inhibitory cells globally rescale excitatory activity along one dimension (preventing runaway excitation and attractor overlap) while leaving the spatial geometry encoded in the remaining dimensions intact. The inhibitory network does not erase the map; it adds orthogonal dynamical constraints that prevent distinct attractor states from colliding in representational space. Consequently, this inhibitory decorrelation significantly expands the intrinsic dimensionality of the population code, as measured by the participation ratio, from $D = 7.5$ to $D = 9.8$ (**Figure 2D**). Manifold rigidity thus emerges as a synergistic circuit property: the excitatory population provides the spatial geometry, the inhibitory population stabilizes its boundaries, and the combination achieves a geometric rigidity consistent with the crystalline code observed at the population level—one that neither population can sustain alone.

Topology Determines Capacity

Theoretical Capacity Advantage of Crystalline Codes To situate the empirical geometric stability findings within a normative framework, we simulated three idealized population codes – crystal (topology = 1.0), mist (topology = 0.5), and noise (topology = 0.0) – in a $N = 500$ neuron network with biological sparsity ($s = 0.15$) and measured memory retrieval error (normalized circular decoding error, expressed relative to a chance baseline of 1.0) across varying numbers of stored memories.

Geometric organization produced a continuous, monotonic improvement in retrieval fidelity. Across eleven topology values from 0 to 1, retrieval error decreased from 0.994 (near chance) to 0.264 at $n = 100$ stored locations, while the correspondence between neural geometry and physical space (Mantel r) increased from 0.001 to 0.700 (**Figure 3A**). The dose-response relationship was smooth throughout, demonstrating that capacity benefits accrue continuously with geometric organization rather than emerging at a discrete threshold.

Crystalline and mist codes diverged sharply in their capacity scaling. At $n = 100$ stored memories, crystal error (0.285) was approximately half that of mist (0.529), while noise codes remained near chance (0.996) at all memory loads tested (**Figure 3C**). Crystal codes maintained sub-threshold retrieval error (< 0.640) across the full range tested ($n = 10$ to $n = 1000$ memories), while mist and noise codes exceeded this threshold at the smallest tested load ($n = 10$). The threshold of 0.640 was defined as the midpoint between crystal and noise performance at $n = 100$ memories, providing a principled boundary between reliable and unreliable retrieval. This separation demonstrates that the *presence* of topological geometric structure – not simply network size – is the primary determinant of memory capacity in this regime.

To confirm that this capacity advantage is a fundamental property of the topological regime rather than an artifact of specific network parameters, we evaluated the topology advantage (ΔError) across a 10,000-configuration parameter sweep varying population size (N), trial counts (T), and sparsity (ρ). This revealed a striking invariance: the representational advantage of the crystalline code is overwhelmingly governed by the network’s sparsity level, while scaling the neural hardware (N) or trial count (T) yields substantially smaller changes in relative performance (**Figure 3B**). The topology advantage follows a pronounced, non-linear trajectory as a function of sparsity, rising steeply between $\rho = 0.02$ and $\rho = 0.10$ before saturating at a plateau of median $\Delta\text{Error} \approx 0.183$ for $\rho \geq 0.11$. Critically, the empirical sparsity of

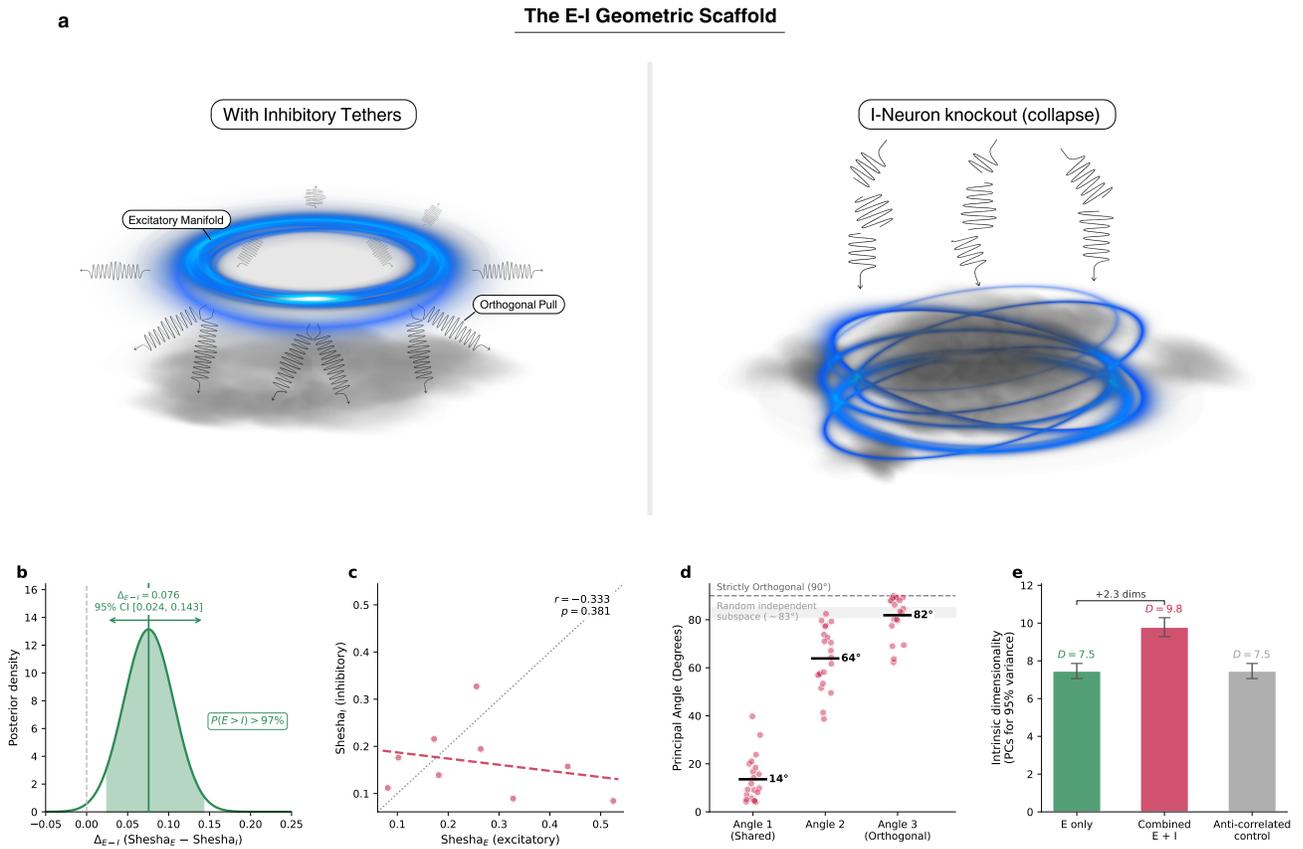


Figure 2: Synergistic excitatory-inhibitory circuit dynamics underlie geometric stability. (a) Conceptual illustration of the E-I geometric scaffold. (Left) The continuous excitatory spatial manifold (blue ring) is stabilized by inhibitory populations acting as structural tethers (springs). By exerting an orthogonal pull, inhibition stiffens the network into a rigid, high-dimensional configuration. (Right) Simulated knockout of these inhibitory tethers causes the representational geometry to collapse, demonstrating that manifold rigidity is an actively constructed, emergent property of the combined E-I architecture rather than an autonomous excitatory phenomenon. (b) Bayesian posterior for the E-I geometric fidelity difference (Δ_{E-I}) across $N = 13$ sessions with simultaneous E and I recordings. The credible interval [0.024, 0.143] excludes zero; posterior probability of $E > I$ exceeds 97.5%. (c) Scatter of excitatory vs. inhibitory geometric stability (Shesha) across sessions, showing negative coordination ($r = -0.333$). Each point is one session; the negative slope is inconsistent with a redundancy model in which both cell types track the identical spatial signal. (d) Principal angles between the top-3 excitatory and inhibitory subspaces ($n = 21$ chickadee sessions). The architecture reveals a single shared axis of co-modulation (Angle 1, mean 14°), while the remaining spatial map dimensions are structurally independent (Angle 3, mean 82° , approaching the $\sim 83^\circ$ expected for random independent subspaces). (e) Intrinsic dimensionality of the excitatory-only ($D = 7.5$), combined E-I ($D = 9.8$), and anti-correlated inhibitory control state spaces. The expansion from $D = 7.5$ to $D = 9.8$ requires genuine orthogonal inhibitory structure; the anti-correlated control produces no expansion, ruling out simple subtractive suppression as the mechanism.

chickadee hippocampal neurons ($\rho \approx 0.15$) falls squarely within this saturated high-advantage regime, confirming that caching networks operate at a sparsity level that maximizes the structural benefits of the crystalline code—and that this optimization is robust rather than finely tuned to a precise parameter value.

Empirical Redundancy To connect these theoretical predictions to the neural data, we estimated population redundancy – the ratio of summed single-cell information to population information – as a proxy for the degree to which spatial information is distributed across the ensemble. Sessions with near-zero population information ($I_{\text{pop}} < 0.05$ bits) were excluded from the primary analysis to avoid undefined ratios at floor-level denominators (14 chickadee and 4 finch sessions excluded; see below).

In the remaining sessions, chickadee populations showed substantially higher redundancy than finch populations (chickadee: mean = 19.3, median = 12.0, $n = 21$ sessions; finch: mean = 5.7, median = 2.2, $n = 4$ sessions; Mann-Whitney $p = 0.057$; **Figure 3C**). The unfiltered analysis (chickadee: median = 14.5, $n = 25$; finch: median = 2.2, $n = 6$; $p = 0.041$) showed the same direction, though unfiltered means are inflated by a small number of sessions with very low I_{pop} denominators (maximum chickadee redundancy = 3311 in one session with $I_{\text{pop}} = 0.001$ bits) and should be interpreted with caution. Medians are reported as the primary summary statistic throughout.

The high proportion of near-zero I_{pop} sessions (14/39 chickadee; 2/8 finch) reflects a known limitation of the redundancy estimator in small neural ensembles: when neuron counts are low, the population mutual information estimator collapses toward zero before single-cell estimates do, producing undefined or extreme ratios. The mean absolute population information was nonetheless higher in chickadee sessions (mean = 0.113 bits, $n = 39$) than in finch sessions (mean = 0.046 bits, $n = 8$), consistent with richer ensemble-level spatial encoding in the caching species.

Higher redundancy in caching hippocampus is consistent with the crystalline coding framework: topologically organized codes distribute each memory trace across the ensemble, creating structured redundancy that protects representations against single-neuron failure. The filtered result ($p = 0.057$) represents directional evidence; Given the limited filtered finch sample ($n = 4$ filtered sessions), the result should be interpreted as directional; the 5.5-fold median difference is consistent with the crystalline coding framework but replication with larger finch samples is needed to establish significance.

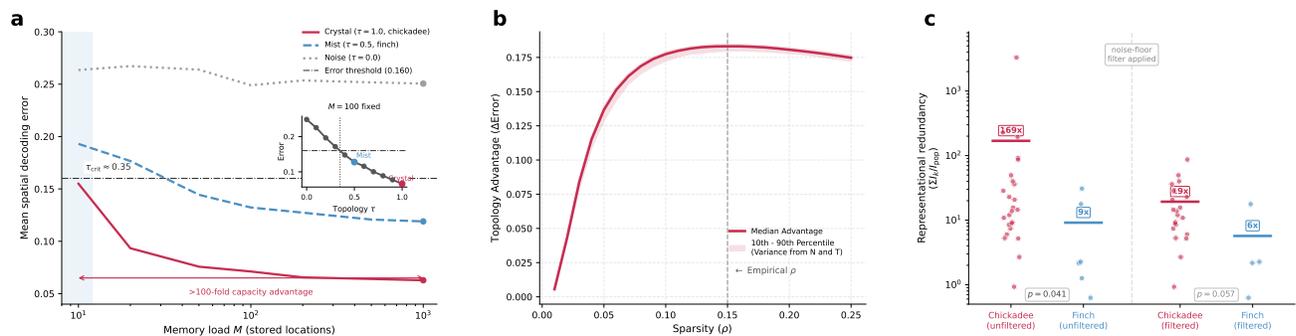


Figure 3: Topological structure determines memory capacity and exacts a geometric tax. (a) Decoding error as a function of memory load M for crystal, mist, and noise regimes, with the critical error threshold and estimated phase transition τ_{crit} marked. Inset shows decoding error as a function of topology strength τ at fixed $M = 100$, illustrating the non-linear inflection. (b) Topology advantage ($\Delta\text{Error} = \text{Error}_{\text{Random}} - \text{Error}_{\text{crystal}}$) across a 10,000-configuration parameter sweep. The solid line represents the median advantage across all population sizes (N) and trial counts (T) for a given sparsity (ρ). The narrow shaded band (10th–90th percentiles) demonstrates that the advantage is highly invariant to N and T , being overwhelmingly driven by sparsity. The advantage peaks precisely at the empirical chickadee sparsity of $\rho \approx 0.15$. (c) Empirical representational redundancy for chickadee and finch sessions, shown as individual points with species means. Filtered results (noise-floor $I_{\text{pop}} \geq 0.05$ bits) are shown alongside unfiltered values. To sustain the optimal crystalline topology shown in (b), the caching network operates with a 169-fold redundancy.

Discussion

Evolutionary Sculpting and the Geometric Tax The extreme selection pressure for cache retrieval, a behavior critical for winter survival, forced the caching avian brain to bypass the standard biological limits of memory capacity (Amit, Gutfreund, and Sompolinsky 1987; Canning and Gardner 1988; Hopfield 1982; Marr 1971; Treves and Rolls 1994)—the interference and crosstalk that normally cause one memory to degrade another as load increases. Our findings suggest

that the chickadee hippocampus achieves this not by simply expanding its physical hardware, but by evolving a highly specific, inhibition-stabilized circuit architecture (Marcus K Benna and Fusi 2016; Fusi and L F Abbott 2007; Rolls and Treves 1998; Vogels et al. 2011) that actively constructs a geometrically rigid manifold (Bernardi et al. 2020; Boyle et al. 2024; Chaudhuri et al. 2019). The difference lies not in the hardware itself, but in the *phase* in which that hardware operates. As demonstrated by our parameter sweep, simply proliferating neurons (increasing N) does not meaningfully improve the network’s resistance to catastrophic interference if the underlying topology remains unstructured. Instead, the network must transition discontinuously into a rigid, topologically stiff lattice.

To appreciate what this means, consider an analogy from condensed matter physics. Both the caching chickadee and the non-caching zebra finch work with the same fundamental biological hardware: billions of neurons, the same synaptic machinery, the same electrochemical signals. In the finch, hippocampal population activity resembles a disorganized gas: the neurons are highly active and drift dynamically across sessions, but they lack any stable, low-dimensional structure. Like water vapor, the “mist” code can sustain a handful of fleeting, localized representations, but as memory load increases toward $M = 10$ and beyond, these representations inevitably collide, blur together, and dissipate (Bakker et al. 2008; Colgin, E. I. Moser, and M.-B. Moser 2008; French 1999; McCloskey and Cohen 1989). Local decoders suffer catastrophic interference precisely because there is no stable geometric reference frame to hold distinct memories apart. In the chickadee, the precise “temperature” of the network, set by the balance of excitatory drive and inhibitory stabilization, causes the population code to lock into a rigid lattice, analogous to an ice crystal. This phase transition is not merely metaphorical (Burak and Fiete 2012; Compte 2000; Renart et al. 2010). Just as water molecules transition discontinuously from gas to solid at a critical temperature, our computational modeling reveals a sharp, non-linear inflection in decoding error at a critical topology threshold $\tau_{\text{crit}} \approx 0.35$: below this threshold, capacity is bounded at $M \leq 10$; above it, reliable readout extends beyond $M = 1,000$ locations.

The circuit architecture of this crystalline code maps naturally onto the physical analogy. The excitatory neurons function as the heavy structural nodes of the lattice, the oxygen atoms that define the foundational geometry of the spatial map. The inhibitory populations act as the hydrogen bonds: they enforce strict geometric angles between these nodes, pushing representations apart orthogonally and preventing collapse (Denève and Machens 2016; Meissner-Bernard, Zenke, and Friedrich 2024). It is this inhibitory decorrelation that expands the intrinsic dimensionality of the population code from $D = 7.5$ (excitatory alone) to $D = 9.8$ (combined E-I), creating a vast, high-dimensional geometry capable of holding thousands of non-interfering representations simultaneously—a geometry with more independent axes along which distinct memories can be separated.

Critically, the capacity of this lattice does not arise from etching each memory onto a specific, isolated neuron—the discrete allocation model proposed by classical frameworks such as Valiant’s Stable Memory Allocator. A single water molecule within the lattice, however precisely positioned, is both fragile and informationally impoverished. Instead, the memory is stored in the continuous shape and angular relationships of the lattice itself: a topological code that is inherently distributed and inherently robust. Our double dissociation confirms this directly. Caching networks exhibit near-zero split-half allocation reliability (the discrete metric) despite maximal geometric stability (the topological metric), demonstrating that the two strategies are not merely different points on a continuum, but fundamentally distinct computational regimes (Bernardi et al. 2020; Fusi, Miller, and Rigotti 2016).

However, this architectural robustness is not computationally free. Consider placing one thousand heavy objects onto an ice structure subject to constant random vibrations, the biological noise of synaptic stochasticity and spontaneous firing. A single, fragile filament of ice molecules cannot bear this load; the structure will shatter, and the memories will be lost. To ensure stability, the caching brain does not simply trace the minimal geometric path. Instead, it freezes a hyper-redundant, extraordinarily thick block of ice, using 169 times more representational capacity than would be strictly necessary to encode a single spatial point. This is the “geometric tax”: the caching system sacrifices coding efficiency to physically stiffen the neural manifold, ensuring that noise-induced perturbations shift the readout toward adjacent spatial neighbors rather than triggering complete representational collapse. The tax is steep, but the return is remarkable: a greater than 100-fold expansion in reliable memory capacity relative to an unstructured code operating under identical hardware constraints.

The Architecture of Variation What this geometric tax ultimately buys is a dramatic suppression of catastrophic interference, sustaining high-fidelity readout at memory loads that would cause complete decoding collapse in unstructured networks (Fusi 2024; Langdon, Genkin, and Engel 2023; Sorscher et al. 2023). By heavily over-parameterizing the spatial map, the network resolves the capacity challenge through a structural mechanism reminiscent of the passacaglia

in Brahms’s Fourth Symphony (Brahms 1886; Buzsáki 2006; Knapp 1989; Pascall 1989): a repeating, immovable harmonic foundation over which elaborate variations are freely constructed. Once the spatial scaffold is locked in place, the remaining flexibility in population activity is available to encode episodic details (such as the specific food item cached, its location, or the temporal context of the event). The stiffness of the underlying geometric scaffold ensures that these continuous episodic updates remain strictly orthogonal to the baseline spatial map. Consequently, the continuous encoding of thousands of new, overlapping cache memories cannot distort or compromise the structural integrity of the baseline coordinate system.

Limitations and Future Directions Several limitations warrant consideration. The E-I circuit analysis is based on $N = 13$ sessions, which is underpowered for frequentist inference; while the Bayesian estimation provides robust credible intervals, replication in larger datasets with simultaneous E-I recordings will be necessary to establish the inhibition-stabilized architecture as a general principle. The computational model operates on a 1D circular track for tractability; direct extension to 2D proved intractable due to concentration of measure, a geometric phenomenon in which random points in high-dimensional spaces become uniformly distant from one another, making topology recovery unreliable. This limitation does not affect the core empirical claims: the Mantel test is geometry-agnostic, and the species difference in geometric stability was measured directly in the 2D open-field arena used for recording, providing empirical confirmation that the crystalline geometry identified by the model is present in the actual recording geometry. Finally, the empirical comparison rests on a single publicly available dataset contrasting two avian species. Convergent evidence from additional caching species, mammalian hippocampal recordings, or perturbation experiments that selectively disrupt inhibitory stabilization would substantially strengthen the causal claims. An alternative theoretical framework proposes that hippocampal spatial representations are fundamentally sequential rather than Euclidean (R. V. Raju et al. 2024)—that is, the brain may represent space as a series of transitions between locations—raising the possibility that place field geometry partially reflects the sequential structure of the animal’s experience rather than a static coordinate system. Our findings are agnostic on this point: the species difference in geometric stability would be expected under either framework, since more rigid sequential context learning and more rigid Euclidean geometry are not easily distinguished in a 1D track recording paradigm. Future work could examine whether the geometric stability differences observed here reflect upstream differences in entorhinal grid cell organization, given that grid scale directly influences hippocampal place field stability and long-term memory (Hardcastle, Ganguli, and Giocomo 2015; Mallory et al. 2018).

A Theoretical Shift in Biological Scaling These findings necessitate a fundamental theoretical shift in how we model scaling and capacity in complex neural systems. Historically, theoretical neuroscience has relied heavily on discrete, allocation-based metaphors, such as quantifying the exact number of neurons recruited to encode a specific memory engram, or assuming that representational capacity scales smoothly with physical hardware allocation. Where earlier work established the molecular and cellular substrates of synaptic plasticity (Kandel 2001, 2007; McNaughton and Morris 1987), the present findings suggest that the organizing principle of high-capacity memory operates at a higher level of abstraction – the geometry of population codes (Kandel 2016). Our observation of a double dissociation, wherein high-capacity caching networks exhibit massive spatial tuning heterogeneity yet precise topological stability, demonstrates the limits of such arithmetic frameworks. To understand memory at biological extremes, we must move toward dynamical, geometric models that measure the topological rigidity of attractor basins rather than the discrete participation of isolated cells (Langdon, Genkin, and Engel 2023; Sorscher et al. 2023). Ultimately, the biological solution to catastrophic interference is not the arithmetic proliferation of neural hardware, but the structural stiffening of representational geometry, a principle that reframes the fundamental unit of memory capacity from the neuron to the geometry of the code, with implications that may extend well beyond the biological systems studied here (Fusi, Miller, and Rigotti 2016; Morales, Muñoz, and Tu 2025; Rigotti et al. 2013).

Materials and Methods

Methods Summary. The detailed description of all analyses, simulations, and statistical procedures is reported in **Appendix A**. Briefly, all data were drawn from a previously published *in vivo* extracellular electrophysiology dataset comprising single units recorded from the hippocampal formation of food-caching black-capped chickadees (*Poecile atricapillus*; 39 sessions, 9 birds) and non-caching zebra finches (*Taeniopygia guttata*; 8 sessions, 9 birds) during open-field foraging (H. Payne, G. Lynch, and Aronov 2021; H. L. Payne, G. F. Lynch, and Aronov 2021). Population geometry was quantified using the Shesha metric, a split-half representational dissimilarity matrix correlation across random neuron partitions

(**Appendix A.4**). Excitatory-inhibitory circuit contributions were assessed via cell-type-specific Shesha decomposition, subspace angle analysis, and a paired full-population vs. excitatory-only synergy test (**Appendix A.12**). Memory capacity and representational redundancy were characterized using a synthetic population code model swept over 10,000 parameter combinations and validated against an open-field 2D arena simulation (**Appendix A.13**). All species comparisons used Mann–Whitney U tests with $n_{\text{perm}} = 10,000$ permutation validation; effect sizes are reported as Cohen’s d with 95% bootstrap confidence intervals ($n_{\text{boot}} = 10,000$).

Code Availability. All custom code including the computational parameter sweeps is available on GitHub (<https://github.com/prashantcraju/hippocampal-stability>).

Acknowledgments

We thank Padma K. and Annapoorna Raju for generously supporting the computational resources used in this work. We thank the authors of Payne et al, 2021 for making their data open-sourced, which made our work possible. The authors acknowledge the use of large language models (specifically the Claude and Gemini families) to assist with processing data, code debugging, improving the visual aesthetic of several plots, and text polishing. All hypotheses, experimental designs, analyses, and interpretations were independently formulated and verified by the authors, and the authors assume full responsibility for all content and claims in this work.

References

- [1] Daniel J Amit, Hanoch Gutfreund, and H Sompolinsky. “Statistical mechanics of neural networks near saturation”. In: *Annals of Physics* 173.1 (Jan. 1987), pp. 30–67. ISSN: 0003-4916. DOI: [10.1016/0003-4916\(87\)90092-3](https://doi.org/10.1016/0003-4916(87)90092-3). URL: [http://dx.doi.org/10.1016/0003-4916\(87\)90092-3](http://dx.doi.org/10.1016/0003-4916(87)90092-3).
- [2] Dmitriy Aronov, Rhino Nevers, and David W. Tank. “Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit”. In: *Nature* 543.7647 (Mar. 2017), pp. 719–722. ISSN: 1476-4687. DOI: [10.1038/nature21692](https://doi.org/10.1038/nature21692). URL: <http://dx.doi.org/10.1038/nature21692>.
- [3] Craig H. Bailey, Dusan Bartsch, and Eric R. Kandel. “Toward a molecular definition of long-term memory storage”. In: *Proceedings of the National Academy of Sciences* 93.24 (Nov. 1996), pp. 13445–13452. ISSN: 1091-6490. DOI: [10.1073/pnas.93.24.13445](https://doi.org/10.1073/pnas.93.24.13445). URL: <http://dx.doi.org/10.1073/pnas.93.24.13445>.
- [4] Arnold Bakker, C. Brock Kirwan, Michael Miller, and Craig E. L. Stark. “Pattern Separation in the Human Hippocampal CA3 and Dentate Gyrus”. In: *Science* 319.5870 (Mar. 2008), pp. 1640–1642. ISSN: 1095-9203. DOI: [10.1126/science.1152882](https://doi.org/10.1126/science.1152882). URL: <http://dx.doi.org/10.1126/science.1152882>.
- [5] Russell P. Balda and Alan C. Kamil. “Long-term spatial memory in clark’s nutcracker, *Nucifraga columbiana*”. In: *Animal Behaviour* 44.4 (Oct. 1992), pp. 761–769. ISSN: 0003-3472. DOI: [10.1016/S0003-3472\(05\)80302-1](https://doi.org/10.1016/S0003-3472(05)80302-1). URL: [http://dx.doi.org/10.1016/S0003-3472\(05\)80302-1](http://dx.doi.org/10.1016/S0003-3472(05)80302-1).
- [6] Marcus K Benna and Stefano Fusi. “Computational principles of synaptic memory consolidation”. In: *Nature Neuroscience* 19.12 (Oct. 2016), pp. 1697–1706. ISSN: 1546-1726. DOI: [10.1038/nn.4401](https://doi.org/10.1038/nn.4401). URL: <http://dx.doi.org/10.1038/nn.4401>.
- [7] Marcus K. Benna and Stefano Fusi. “Place cells may simply be memory cells: Memory compression leads to spatial tuning and history dependence”. In: *Proceedings of the National Academy of Sciences* 118.51 (Dec. 2021). ISSN: 1091-6490. DOI: [10.1073/pnas.2018422118](https://doi.org/10.1073/pnas.2018422118). URL: <http://dx.doi.org/10.1073/pnas.2018422118>.
- [8] Silvia Bernardi, Marcus K. Benna, Mattia Rigotti, Jérôme Munuera, Stefano Fusi, and C. Daniel Salzman. “The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex”. In: *Cell* 183.4 (Nov. 2020), 954–967.e21. ISSN: 0092-8674. DOI: [10.1016/j.cell.2020.09.031](https://doi.org/10.1016/j.cell.2020.09.031). URL: <http://dx.doi.org/10.1016/j.cell.2020.09.031>.
- [9] Lara M. Boyle, Lorenzo Posani, Sarah Irfan, Steven A. Siegelbaum, and Stefano Fusi. “Tuned geometries of hippocampal representations meet the computational demands of social memory”. In: *Neuron* 112.8 (Apr. 2024), 1358–1371.e9. ISSN: 0896-6273. DOI: [10.1016/j.neuron.2024.01.021](https://doi.org/10.1016/j.neuron.2024.01.021). URL: <http://dx.doi.org/10.1016/j.neuron.2024.01.021>.
- [10] Johannes Brahms. *Symphonie Nr. 4 in e-Moll, Op. 98*. First edition. Plate 8686. Engraved by C.G. Röder, Leipzig. Berlin: N. Simrock, 1886.

- [11] David R. Brodbeck and Sara J. Shettleworth. “Matching location and color of a compound stimulus: Comparison of a food-storing and a nonstoring bird species.” In: *Journal of Experimental Psychology: Animal Behavior Processes* 21.1 (Jan. 1995), pp. 64–77. ISSN: 0097-7403. DOI: [10.1037/0097-7403.21.1.64](https://doi.org/10.1037/0097-7403.21.1.64). URL: <http://dx.doi.org/10.1037/0097-7403.21.1.64>.
- [12] Yoram Burak and Ila R. Fiete. “Accurate Path Integration in Continuous Attractor Network Models of Grid Cells”. In: *PLoS Computational Biology* 5.2 (Feb. 2009). Ed. by Olaf Sporns, e1000291. ISSN: 1553-7358. DOI: [10.1371/journal.pcbi.1000291](https://doi.org/10.1371/journal.pcbi.1000291). URL: <http://dx.doi.org/10.1371/journal.pcbi.1000291>.
- [13] Yoram Burak and Ila R. Fiete. “Fundamental limits on persistent activity in networks of noisy neurons”. In: *Proceedings of the National Academy of Sciences* 109.43 (Oct. 2012), pp. 17645–17650. ISSN: 1091-6490. DOI: [10.1073/pnas.1117386109](https://doi.org/10.1073/pnas.1117386109). URL: <http://dx.doi.org/10.1073/pnas.1117386109>.
- [14] György Buzsáki. *Rhythms of the Brain*. Oxford: Oxford University Press, 2006. ISBN: 9780199828234. DOI: [10.1093/acprof:oso/9780195301069.001.0001](https://doi.org/10.1093/acprof:oso/9780195301069.001.0001).
- [15] A Canning and E Gardner. “Partially connected models of neural networks”. In: *Journal of Physics A: Mathematical and General* 21.15 (Aug. 1988), pp. 3275–3284. ISSN: 1361-6447. DOI: [10.1088/0305-4470/21/15/016](https://doi.org/10.1088/0305-4470/21/15/016). URL: <http://dx.doi.org/10.1088/0305-4470/21/15/016>.
- [16] Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. “The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep”. In: *Nature Neuroscience* 22.9 (Aug. 2019), pp. 1512–1520. ISSN: 1546-1726. DOI: [10.1038/s41593-019-0460-x](https://doi.org/10.1038/s41593-019-0460-x). URL: <http://dx.doi.org/10.1038/s41593-019-0460-x>.
- [17] Nicky S Clayton and John R Krebs. “Memory in food-storing birds: from behaviour to brain”. In: *Current Opinion in Neurobiology* 5.2 (Apr. 1995), pp. 149–154. ISSN: 0959-4388. DOI: [10.1016/0959-4388\(95\)80020-4](https://doi.org/10.1016/0959-4388(95)80020-4). URL: [http://dx.doi.org/10.1016/0959-4388\(95\)80020-4](http://dx.doi.org/10.1016/0959-4388(95)80020-4).
- [18] Laura Lee Colgin, Edvard I. Moser, and May-Britt Moser. “Understanding memory through hippocampal remapping”. In: *Trends in Neurosciences* 31.9 (Sept. 2008), pp. 469–477. ISSN: 0166-2236. DOI: [10.1016/j.tins.2008.06.008](https://doi.org/10.1016/j.tins.2008.06.008). URL: <http://dx.doi.org/10.1016/j.tins.2008.06.008>.
- [19] A. Compte. “Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model”. In: *Cerebral Cortex* 10.9 (Sept. 2000), pp. 910–923. ISSN: 1460-2199. DOI: [10.1093/cercor/10.9.910](https://doi.org/10.1093/cercor/10.9.910). URL: <http://dx.doi.org/10.1093/cercor/10.9.910>.
- [20] Hristos S. Courellis, Juri Minxha, Araceli R. Cardenas, Daniel L. Kimmel, Chrystal M. Reed, Taufik A. Valiante, C. Daniel Salzman, Adam N. Mamelak, Stefano Fusi, and Ueli Rutishauser. “Abstract representations emerge in human hippocampal neurons during inference”. In: *Nature* 632.8026 (Aug. 2024), pp. 841–849. ISSN: 1476-4687. DOI: [10.1038/s41586-024-07799-x](https://doi.org/10.1038/s41586-024-07799-x). URL: <http://dx.doi.org/10.1038/s41586-024-07799-x>.
- [21] John P Cunningham and Byron M Yu. “Dimensionality reduction for large-scale neural recordings”. In: *Nature Neuroscience* 17.11 (Aug. 2014), pp. 1500–1509. ISSN: 1546-1726. DOI: [10.1038/nn.3776](https://doi.org/10.1038/nn.3776). URL: <http://dx.doi.org/10.1038/nn.3776>.
- [22] Peter Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Computational Neuroscience Series. Cambridge, MA: MIT Press, 2005. ISBN: 9780262541855.
- [23] Sophie Denève and Christian K Machens. “Efficient codes and balanced networks”. In: *Nature Neuroscience* 19.3 (Feb. 2016), pp. 375–382. ISSN: 1546-1726. DOI: [10.1038/nn.4243](https://doi.org/10.1038/nn.4243). URL: <http://dx.doi.org/10.1038/nn.4243>.
- [24] R French. “Catastrophic forgetting in connectionist networks”. In: *Trends in Cognitive Sciences* 3.4 (Apr. 1999), pp. 128–135. ISSN: 1364-6613. DOI: [10.1016/s1364-6613\(99\)01294-2](https://doi.org/10.1016/s1364-6613(99)01294-2). URL: [http://dx.doi.org/10.1016/S1364-6613\(99\)01294-2](http://dx.doi.org/10.1016/S1364-6613(99)01294-2).
- [25] Stefano Fusi. “Memory Capacity of Neural Network Models”. In: *The Oxford Handbook of Human Memory, Two Volume Pack*. Oxford University Press, July 2024, pp. 740–764. ISBN: 9780190918019. DOI: [10.1093/oxfordhb/9780190917982.013.26](https://doi.org/10.1093/oxfordhb/9780190917982.013.26). URL: <http://dx.doi.org/10.1093/oxfordhb/9780190917982.013.26>.
- [26] Stefano Fusi and L F Abbott. “Limits on the memory storage capacity of bounded synapses”. In: *Nature Neuroscience* 10.4 (Mar. 2007), pp. 485–493. ISSN: 1546-1726. DOI: [10.1038/nn1859](https://doi.org/10.1038/nn1859). URL: <http://dx.doi.org/10.1038/nn1859>.
- [27] Stefano Fusi, Earl K Miller, and Mattia Rigotti. “Why neurons mix: high dimensionality for higher cognition”. In: *Current Opinion in Neurobiology* 37 (Apr. 2016), pp. 66–74. ISSN: 0959-4388. DOI: [10.1016/j.conb.2016.01.010](https://doi.org/10.1016/j.conb.2016.01.010). URL: <http://dx.doi.org/10.1016/j.conb.2016.01.010>.
- [28] László Zsolt Garamszegi and Jeffrey R Lucas. “Continental variation in relative hippocampal volume in birds: the phylogenetic extent of the effect and the potential role of winter temperatures”. In: *Biology Letters* 1.3 (June 2005), pp. 330–333. ISSN: 1744-957X. DOI: [10.1098/rsbl.2005.0328](https://doi.org/10.1098/rsbl.2005.0328). URL: <http://dx.doi.org/10.1098/rsbl.2005.0328>.

- [29] Roberto A. Gulli, Lyndon R. Duong, Benjamin W. Corrigan, Guillaume Doucet, Sylvain Williams, Stefano Fusi, and Julio C. Martinez-Trujillo. “Context-dependent representations of objects and space in the primate hippocampus during virtual navigation”. In: *Nature Neuroscience* 23.1 (Dec. 2019), pp. 103–112. ISSN: 1546-1726. DOI: [10.1038/s41593-019-0548-3](https://doi.org/10.1038/s41593-019-0548-3). URL: <http://dx.doi.org/10.1038/s41593-019-0548-3>.
- [30] Kiah Hardcastle, Surya Ganguli, and Lisa M. Giocomo. “Environmental Boundaries as an Error Correction Mechanism for Grid Cells”. In: *Neuron* 86.3 (May 2015), pp. 827–839. ISSN: 0896-6273. DOI: [10.1016/j.neuron.2015.03.039](https://doi.org/10.1016/j.neuron.2015.03.039). URL: <http://dx.doi.org/10.1016/j.neuron.2015.03.039>.
- [31] J J Hopfield. “Neural networks and physical systems with emergent collective computational abilities.” In: *Proceedings of the National Academy of Sciences* 79.8 (Apr. 1982), pp. 2554–2558. ISSN: 1091-6490. DOI: [10.1073/pnas.79.8.2554](https://doi.org/10.1073/pnas.79.8.2554). URL: <http://dx.doi.org/10.1073/pnas.79.8.2554>.
- [32] Lucia F. Jacobs. “The Evolution of the Cognitive Map”. In: *Brain, Behavior and Evolution* 62.2 (2003), pp. 128–139. ISSN: 1421-9743. DOI: [10.1159/000072443](https://doi.org/10.1159/000072443). URL: <http://dx.doi.org/10.1159/000072443>.
- [33] Mehrdad Jazayeri and Srdjan Ostojic. “Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity”. In: *Current Opinion in Neurobiology* 70 (Oct. 2021), pp. 113–120. ISSN: 0959-4388. DOI: [10.1016/j.conb.2021.08.002](https://doi.org/10.1016/j.conb.2021.08.002). URL: <http://dx.doi.org/10.1016/j.conb.2021.08.002>.
- [34] Eric R. Kandel. “The molecular biology of memory storage: a dialogue between genes and synapses”. In: *Science* 294.5544 (Nov. 2001), pp. 1030–1038.
- [35] Eric R. Kandel. *In search of memory*. New York, NY: WW Norton, Mar. 2007. ISBN: 978-0393329377.
- [36] Eric R. Kandel. *Reductionism in art and brain science*. New York, NY: Columbia University Press, Aug. 2016. ISBN: 978-0231179621.
- [37] Raymond Knapp. “The Finale of Brahms’s Fourth Symphony: The Tale of the Subject”. In: *19th-Century Music* 13.1 (1989), pp. 3–17. ISSN: 0148-2076. DOI: [10.2307/746207](https://doi.org/10.2307/746207). URL: <http://dx.doi.org/10.2307/746207>.
- [38] J R Krebs, D F Sherry, S D Healy, V H Perry, and A L Vaccarino. “Hippocampal specialization of food-storing birds.” In: *Proceedings of the National Academy of Sciences* 86.4 (Feb. 1989), pp. 1388–1392. ISSN: 1091-6490. DOI: [10.1073/pnas.86.4.1388](https://doi.org/10.1073/pnas.86.4.1388). URL: <http://dx.doi.org/10.1073/pnas.86.4.1388>.
- [39] Nikolaus Kriegeskorte, Marieke Mur, and Peter Bandettini. “Representational similarity analysis – connecting the branches of systems neuroscience”. In: *Frontiers in Systems Neuroscience* (2008). ISSN: 1662-5137. DOI: [10.3389/neuro.06.004.2008](https://doi.org/10.3389/neuro.06.004.2008).
- [40] Christopher Langdon, Mikhail Genkin, and Tatiana A. Engel. “A unifying perspective on neural manifolds and circuits for cognition”. In: *Nature Reviews Neuroscience* 24.6 (Apr. 2023), pp. 363–377. ISSN: 1471-0048. DOI: [10.1038/s41583-023-00693-x](https://doi.org/10.1038/s41583-023-00693-x). URL: <http://dx.doi.org/10.1038/s41583-023-00693-x>.
- [41] Caitlin S. Mallory, Kiah Hardcastle, Jason S. Bant, and Lisa M. Giocomo. “Grid scale drives the scale and long-term stability of place maps”. In: *Nature Neuroscience* 21.2 (Jan. 2018), pp. 270–282. ISSN: 1546-1726. DOI: [10.1038/s41593-017-0055-3](https://doi.org/10.1038/s41593-017-0055-3). URL: <http://dx.doi.org/10.1038/s41593-017-0055-3>.
- [42] Nathan Mantel. “The Detection of Disease Clustering and a Generalized Regression Approach”. In: *Cancer Research* 27.2_Part_1 (Feb. 1967), pp. 209–220. ISSN: 0008-5472. eprint: https://aacrjournals.org/cancerres/article-pdf/27/2_Part_1/209/2382183/cr0272p10209.pdf.
- [43] Nathan Mantel and Ranchhodbhai S. Valand. “A Technique of Nonparametric Multivariate Analysis”. In: *Biometrics* 26.3 (Sept. 1970), p. 547. ISSN: 0006-341X. DOI: [10.2307/2529108](https://doi.org/10.2307/2529108). URL: <http://dx.doi.org/10.2307/2529108>.
- [44] D. Marr. “Simple memory: a theory for archicortex”. In: *Philosophical Transactions of the Royal Society of London. B, Biological Sciences* 262.841 (July 1971), pp. 23–81. ISSN: 2054-0280. DOI: [10.1098/rstb.1971.0078](https://doi.org/10.1098/rstb.1971.0078). URL: <http://dx.doi.org/10.1098/rstb.1971.0078>.
- [45] Michael McCloskey and Neal J. Cohen. “Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem”. In: *Psychology of Learning and Motivation*. Elsevier, 1989, pp. 109–165. ISBN: 9780125433242. DOI: [10.1016/s0079-7421\(08\)60536-8](https://doi.org/10.1016/s0079-7421(08)60536-8). URL: [http://dx.doi.org/10.1016/s0079-7421\(08\)60536-8](http://dx.doi.org/10.1016/s0079-7421(08)60536-8).
- [46] B.L. McNaughton and R.G.M. Morris. “Hippocampal synaptic enhancement and information storage within a distributed memory system”. In: *Trends in Neurosciences* 10.10 (Jan. 1987), pp. 408–415. ISSN: 0166-2236. DOI: [10.1016/0166-2236\(87\)90011-7](https://doi.org/10.1016/0166-2236(87)90011-7). URL: [http://dx.doi.org/10.1016/0166-2236\(87\)90011-7](http://dx.doi.org/10.1016/0166-2236(87)90011-7).
- [47] Claire Meissner-Bernard, Friedemann Zenke, and Rainer W. Friedrich. “Geometry and dynamics of representations in a precisely balanced memory network related to olfactory cortex”. In: *eLife* (May 2024). DOI: [10.7554/eLife.96303.1](https://doi.org/10.7554/eLife.96303.1). URL: <http://dx.doi.org/10.7554/eLife.96303.1>.
- [48] Guillermo B. Morales, Miguel A. Muñoz, and Yuhai Tu. “Representational drift and learning-induced stabilization in the piriform cortex”. In: *Proceedings of the National Academy of Sciences* 122.29 (July 2025). ISSN: 1091-6490. DOI: [10.1073/pnas.2501811122](https://doi.org/10.1073/pnas.2501811122). URL: <http://dx.doi.org/10.1073/pnas.2501811122>.

- [49] Robert Pascall. “Genre and the Finale of Brahms’s Fourth Symphony”. In: *Music Analysis* 8.3 (Oct. 1989), p. 233. ISSN: 0262-5245. DOI: [10.2307/854289](https://doi.org/10.2307/854289). URL: <http://dx.doi.org/10.2307/854289>.
- [50] Hannah Payne, Galen Lynch, and Dmitriy Aronov. *Neural representations of space in the hippocampus of a food-caching bird*. en. 2021. DOI: [10.5061/DRYAD.PG4F4QRP7](https://doi.org/10.5061/DRYAD.PG4F4QRP7). URL: <https://datadryad.org/dataset/doi:10.5061/dryad.pg4f4qrp7>.
- [51] Hannah L. Payne, Galen F. Lynch, and Dmitriy Aronov. “Neural representations of space in the hippocampus of a food-caching bird”. In: *Science* 373.6552 (July 2021), pp. 343–348. ISSN: 1095-9203. DOI: [10.1126/science.abg2009](https://doi.org/10.1126/science.abg2009). URL: <http://dx.doi.org/10.1126/science.abg2009>.
- [52] Vladimir V. Pravosudov, Timothy C. Roth, Lara D. LaDage, and Cody A. Freas. “Environmental Influences on Spatial Memory and the Hippocampus in Food-Caching Chickadees”. In: *Comparative Cognition & Behavior Reviews* 10 (2015), pp. 25–43. ISSN: 1911-4745. DOI: [10.3819/ccbr.2015.100002](https://doi.org/10.3819/ccbr.2015.100002). URL: <http://dx.doi.org/10.3819/ccbr.2015.100002>.
- [53] Prashant C. Raju. “Geometric Stability: The Missing Axis of Representations”. In: *arXiv preprint arXiv:2601.09173* (2026).
- [54] Prashant C. Raju. *Shesha: Self-Consistency Metrics for Representational Stability*. 2026. DOI: [10.5281/zenodo.18227453](https://doi.org/10.5281/zenodo.18227453). URL: <https://doi.org/10.5281/zenodo.18227453>.
- [55] Rajkumar Vasudeva Raju, J. Swaroop Guntupalli, Guangyao Zhou, Carter Wendelken, Miguel Lázaro-Gredilla, and Dileep George. “Space is a latent sequence: A theory of the hippocampus”. In: *Science Advances* 10.31 (Aug. 2024). ISSN: 2375-2548. DOI: [10.1126/sciadv.adm8470](https://doi.org/10.1126/sciadv.adm8470). URL: <http://dx.doi.org/10.1126/sciadv.adm8470>.
- [56] A. D. Redish, F. P. Battaglia, M. K. Chawla, A. D. Ekstrom, J. L. Gerrard, P. Lipa, E. S. Rosenzweig, P. F. Worley, J. F. Guzowski, B. L. McNaughton, and C. A. Barnes. “Independence of Firing Correlates of Anatomically Proximate Hippocampal Pyramidal Cells”. In: *The Journal of Neuroscience* 21.5 (Mar. 2001), RC134–RC134. ISSN: 1529-2401. DOI: [10.1523/jneurosci.21-05-j0004.2001](https://doi.org/10.1523/jneurosci.21-05-j0004.2001). URL: <http://dx.doi.org/10.1523/JNEUROSCI.21-05-j0004.2001>.
- [57] Alfonso Renart, Jaime de la Rocha, Peter Bartho, Liad Hollender, Néstor Parga, Alex Reyes, and Kenneth D. Harris. “The Asynchronous State in Cortical Circuits”. In: *Science* 327.5965 (Jan. 2010), pp. 587–590. ISSN: 1095-9203. DOI: [10.1126/science.1179850](https://doi.org/10.1126/science.1179850). URL: <http://dx.doi.org/10.1126/science.1179850>.
- [58] Mattia Rigotti, Omri Barak, Melissa R. Warden, Xiao-Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi. “The importance of mixed selectivity in complex cognitive tasks”. In: *Nature* 497.7451 (May 2013), pp. 585–590. ISSN: 1476-4687. DOI: [10.1038/nature12160](https://doi.org/10.1038/nature12160). URL: <http://dx.doi.org/10.1038/nature12160>.
- [59] Edmund T. Rolls and Alessandro Treves. *Neural networks and brain function*. Oxford: Oxford University Press, 1998. DOI: [10.1093/acprof:oso/9780198524328.001.0001](https://api.semanticscholar.org/CorpusID:14588175). URL: <https://api.semanticscholar.org/CorpusID:14588175>.
- [60] Alex Roxin and Stefano Fusi. “Efficient Partitioning of Memory Systems and Its Importance for Memory Consolidation”. In: *PLoS Computational Biology* 9.7 (July 2013). Ed. by Jeff Beck, e1003146. ISSN: 1553-7358. DOI: [10.1371/journal.pcbi.1003146](https://doi.org/10.1371/journal.pcbi.1003146). URL: <http://dx.doi.org/10.1371/journal.pcbi.1003146>.
- [61] Andrew Saxe, Stephanie Nelli, and Christopher Summerfield. “If deep learning is the answer, what is the question?”. In: *Nature Reviews Neuroscience* 22.1 (Nov. 2020), pp. 55–67. ISSN: 1471-0048. DOI: [10.1038/s41583-020-00395-8](https://doi.org/10.1038/s41583-020-00395-8). URL: <http://dx.doi.org/10.1038/s41583-020-00395-8>.
- [62] David F. Sherry and Anthony L. Vaccarino. “Hippocampus and memory for food caches in black-capped chickadees.” In: *Behavioral Neuroscience* 103.2 (Apr. 1989), pp. 308–318. ISSN: 0735-7044. DOI: [10.1037/0735-7044.103.2.308](https://doi.org/10.1037/0735-7044.103.2.308). URL: <http://dx.doi.org/10.1037/0735-7044.103.2.308>.
- [63] David F. Sherry, Anthony L. Vaccarino, Karen Buckenham, and Rachel S. Herz. “The Hippocampal Complex of Food-Storing Birds”. In: *Brain, Behavior and Evolution* 34.5 (1989), pp. 308–317. ISSN: 1421-9743. DOI: [10.1159/000116516](https://doi.org/10.1159/000116516). URL: <http://dx.doi.org/10.1159/000116516>.
- [64] Tom V. Smulders. “The Avian Hippocampal Formation and the Stress Response”. In: *Brain, Behavior and Evolution* 90.1 (2017), pp. 81–91. ISSN: 1421-9743. DOI: [10.1159/000477654](https://doi.org/10.1159/000477654). URL: <http://dx.doi.org/10.1159/000477654>.
- [65] Ben Sorscher, Gabriel C. Mel, Samuel A. Ocko, Lisa M. Giocomo, and Surya Ganguli. “A unified theory for the computational and mechanistic origins of grid cells”. In: *Neuron* 111.1 (Jan. 2023), 121–137.e13. ISSN: 0896-6273. DOI: [10.1016/j.neuron.2022.10.003](https://doi.org/10.1016/j.neuron.2022.10.003). URL: <http://dx.doi.org/10.1016/j.neuron.2022.10.003>.
- [66] Larry R. Squire. “Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans.” In: *Psychological Review* 99.2 (1992), pp. 195–231. ISSN: 0033-295X. DOI: [10.1037/0033-295x.99.2.195](https://doi.org/10.1037/0033-295x.99.2.195). URL: <http://dx.doi.org/10.1037/0033-295x.99.2.195>.
- [67] Alexander Thome, Diano F. Marrone, Timothy M. Ellmore, Monica K. Chawla, Peter Lipa, Victor Ramirez-Amaya, Sarah H. Lisanby, Bruce L. McNaughton, and Carol A. Barnes. “Evidence for an Evolutionarily Conserved Memory

- Coding Scheme in the Mammalian Hippocampus”. In: *The Journal of Neuroscience* 37.10 (Feb. 2017), pp. 2795–2801. ISSN: 1529-2401. DOI: [10.1523/jneurosci.3057-16.2017](https://doi.org/10.1523/jneurosci.3057-16.2017). URL: <http://dx.doi.org/10.1523/JNEUROSCI.3057-16.2017>.
- [68] L.T. Thompson and P.J. Best. “Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats”. In: *Brain Research* 509.2 (Feb. 1990), pp. 299–308. ISSN: 0006-8993. DOI: [10.1016/0006-8993\(90\)90555-p](https://doi.org/10.1016/0006-8993(90)90555-p). URL: [http://dx.doi.org/10.1016/0006-8993\(90\)90555-P](http://dx.doi.org/10.1016/0006-8993(90)90555-P).
- [69] Alessandro Treves and Edmund T. Rolls. “Computational analysis of the role of the hippocampus in memory”. In: *Hippocampus* 4.3 (June 1994), pp. 374–391. ISSN: 1098-1063. DOI: [10.1002/hipo.450040319](https://doi.org/10.1002/hipo.450040319). URL: <http://dx.doi.org/10.1002/hipo.450040319>.
- [70] Leslie G. Valiant. “Memorization and Association on a Realistic Neural Model”. In: *Neural Computation* 17.3 (Mar. 2005), pp. 527–555. ISSN: 1530-888X. DOI: [10.1162/0899766053019890](https://doi.org/10.1162/0899766053019890). URL: <http://dx.doi.org/10.1162/0899766053019890>.
- [71] Leslie G. Valiant. “The Hippocampus as a Stable Memory Allocator for Cortex”. In: *Neural Computation* 24.11 (Nov. 2012), pp. 2873–2899. ISSN: 1530-888X. DOI: [10.1162/neco_a_00357](https://doi.org/10.1162/neco_a_00357). URL: http://dx.doi.org/10.1162/neco_a_00357.
- [72] T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner. “Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks”. In: *Science* 334.6062 (Dec. 2011), pp. 1569–1573. ISSN: 1095-9203. DOI: [10.1126/science.1211095](https://doi.org/10.1126/science.1211095). URL: <http://dx.doi.org/10.1126/science.1211095>.
- [73] Stephen Waydo, Alexander Kraskov, Rodrigo Quiñan Quiroga, Itzhak Fried, and Christof Koch. “Sparse Representation in the Human Medial Temporal Lobe”. In: *The Journal of Neuroscience* 26.40 (Oct. 2006), pp. 10232–10234. ISSN: 1529-2401. DOI: [10.1523/jneurosci.2101-06.2006](https://doi.org/10.1523/jneurosci.2101-06.2006). URL: <http://dx.doi.org/10.1523/JNEUROSCI.2101-06.2006>.

A Methods

A.1 Subjects and electrophysiological recordings

All analyses were performed on publicly available *in vivo* extracellular electrophysiological recordings from the avian hippocampal formation (dorsomedial pallium) (H. Payne, G. Lynch, and Aronov 2021; H. L. Payne, G. F. Lynch, and Aronov 2021). The dataset comprised 755 single units recorded from food-caching black-capped chickadees (*Poecile atricapillus*; 39 sessions across 9 birds) and 238 single units from non-caching zebra finches (*Taeniopygia guttata*; 8 sessions across 9 birds), for a total of 993 units. Data were loaded from MATLAB v5 binary files (RESULTS_T.mat, RESULTS_Z.mat) and parsed using a custom recursive miMATRIX reader (for MCOS table objects embedded in the MAT v5 subsystem) or, where possible, using `scipy.io.loadmat`. Behavioral telemetry (position, head angle, movement speed) was extracted from the accompanying B (behavior) variable stored as a MATLAB struct.

A.2 Cell-type classification

Putative excitatory (E) and inhibitory (I) neurons were classified using Ward hierarchical agglomerative clustering on standardized waveform features. For each unit, two features were extracted: (i) spike width (peak-to-trough duration, converted to milliseconds) and (ii) peak-to-peak ratio (`pp_ratio`). Features were z-scored prior to clustering:

$$\mathbf{f}_{\text{std}} = \frac{\mathbf{f} - \bar{\mathbf{f}}}{\sigma_{\mathbf{f}}}$$

Ward linkage was applied to the standardized feature matrix and the dendrogram was cut at $k = 2$ clusters. The cluster with the lower mean spike width was assigned as narrow-spiking (putative inhibitory) and the remaining cluster as broad-spiking (putative excitatory). This procedure yielded a classification concordance of $\sim 99\%$ with the published labels (536/219 versus reported 538/217 for chickadees). Spatial selectivity was assessed by computing the spatial information (`info`) of each unit and comparing it against a null distribution of 200 bin-shuffled surrogates; units with information exceeding the 99th percentile of the shuffle distribution ($p < 0.01$) were classified as spatially selective.

A.3 Construction of population rate maps

For each recording session, two-dimensional arena occupancy was discretized into a 40×40 grid ($S = 1,600$ spatial bins). Firing rate maps were computed for each neuron as the spike count in each bin divided by the occupancy time, yielding a one-dimensional vector $\mathbf{r}_k \in \mathbb{R}^S$ per neuron k . Bins with no occupancy were retained as NaN rather than zero-filled; all pairwise distance computations used pairwise deletion (see below). For place field detection, which requires spatially continuous maps, NaN bins were locally replaced with zero prior to Gaussian smoothing. The population response matrix $\mathbf{M} \in \mathbb{R}^{N \times S}$ was constructed by stacking the rate maps of all N neurons recorded in a given session, where each row represents a neuron and each column represents a spatial bin. Sessions were included in all subsequent analyses only if they contained at least $N_{\text{min}} \geq 5$ excitatory neurons; for the E/I analyses (Tier 2), both cell types required a minimum of $N_{\text{min}} = 3$ neurons per type. Rate maps were z-scored neuron-wise prior to all geometric analyses:

$$\tilde{r}_{k,s} = \frac{r_{k,s} - \bar{r}_k}{\sigma_{r_k}}, \quad \sigma_{r_k} > 0$$

Neurons with zero variance ($\sigma_{r_k} = 0$) were assigned unit standard deviation to avoid division by zero. The per-neuron mean \bar{r}_k and standard deviation σ_{r_k} were computed using NaN-aware estimators, excluding unvisited bins from the normalization. An active-bin filter was applied, retaining only spatial bins for which at least $\max(2, \lfloor N/3 \rfloor)$ neurons had non-zero firing rates; sessions with fewer than 30 active bins were excluded.

A.4 The Shesha metric: geometric stability of the neural manifold

We quantified the geometric stability of the hippocampal spatial code using the Shesha metric (P. C. Raju 2026a), a split-half representational dissimilarity matrix (RDM) correlation framework (Kriegeskorte, Mur, and Bandettini 2008), with pairwise deletion of unvisited spatial bins.

Feature-split stability (Shesha_{FS}). The population of N neurons was randomly partitioned into two equal halves ($N/2$ neurons each). For each half, an RDM was computed over all active spatial bins using *masked cosine distance* with pairwise deletion. Let $V_{ij}^{(h)}$ denote the set of neurons in half h with valid (finite) z -scored rates at both bins i and j . The masked cosine distance is:

$$D_{ij}^{(h)} = 1 - \frac{\sum_{k \in V_{ij}^{(h)}} x_{ik}^{(h)} x_{jk}^{(h)}}{\sqrt{\sum_{k \in V_{ij}^{(h)}} (x_{ik}^{(h)})^2} \sqrt{\sum_{k \in V_{ij}^{(h)}} (x_{jk}^{(h)})^2}}$$

where $x_{ik}^{(h)}$ is the z -scored response of neuron k (from half h) at spatial bin i . Pairs with $|V_{ij}^{(h)}| < 2$ were assigned NaN and excluded from downstream correlations. This pairwise deletion prevents unvisited bins (retained as NaN) from inflating similarity through shared zeros. The Shesha score is the Spearman correlation between the valid upper-triangular entries of the two RDMs:

$$\text{Shesha} = \rho_{\text{Spearman}}(\text{vec}(\mathbf{D}_{\text{valid}}^{(1)}), \text{vec}(\mathbf{D}_{\text{valid}}^{(2)}))$$

computed only over index pairs (i, j) where both $D_{ij}^{(1)}$ and $D_{ij}^{(2)}$ are finite. This was repeated for $n_{\text{splits}} = 100$ random neuron partitions with a fixed random seed (seed = 320), and the mean correlation was taken as the final Shesha score. All active bins were used (no subsampling). The Tier 2 E/I pipeline passes these parameters explicitly to `feature_split`; the Tier 1 pipeline calls `feature_split` with the `shesha-geometry` package defaults, which are `n_splits = 100`, `metric = 'cosine'`, and `seed = 320`, matching the values documented above.

Sample-split stability (Shesha_{SS}). An analogous procedure was applied using sample splits rather than feature splits: subsets comprising 40% of spatial bins were randomly drawn and RDMs computed from each subsample, yielding a measure of representational robustness across spatial subsamples (sample_split(X , n_splits=30, subsample_fraction=0.4, metric='cosine', seed=320)).

A.5 Mantel test for spatial topology

To assess whether population-level geometric structure faithfully preserves the physical layout of the environment, we computed the Mantel test (Mantel 1967; Mantel and Valand 1970). Physical distances were computed as Euclidean distances between the 2D grid coordinates (r_i, c_i) of each active bin:

$$d_{\text{phys}}(i, j) = \sqrt{(r_i - r_j)^2 + (c_i - c_j)^2}$$

Neural distances were computed as the root-mean-square distance over mutually valid neurons (pairwise deletion):

$$d_{\text{neural}}(i, j) = \sqrt{\frac{1}{|V_{ij}|} \sum_{k \in V_{ij}} (\tilde{m}_{k,i} - \tilde{m}_{k,j})^2}$$

where $V_{ij} = \{k : \tilde{m}_{k,i} \text{ and } \tilde{m}_{k,j} \text{ are both finite}\}$ and $\tilde{m}_{k,i}$ is the z -scored firing rate of neuron k at bin i . Pairs with $|V_{ij}| < 2$ were excluded. When all neurons have valid data at both bins, this reduces to $\|\tilde{\mathbf{m}}_i - \tilde{\mathbf{m}}_j\|_2 / \sqrt{N}$. The Tier 1 Mantel analysis uses the masked formulation above (pairwise deletion over V_{ij}). In the Tier 2 E/I pipeline, rate maps are zero-filled prior to distance computation (`nan_to_num` replaces all NaN entries with zero before calling `pdist`), so $V_{ij} = \{1, \dots, N\}$ for all bin pairs and the expression reduces to standard Euclidean distance divided by \sqrt{N} without pairwise deletion. This zero-filling is appropriate in Tier 2 because the active-bin filter has already removed unvisited bins, and the remaining bins have valid rates for all neurons within each cell-type subpopulation. The observed Mantel statistic was the Spearman rank correlation r_{obs} between all valid $\binom{n}{2}$ pairwise distances. Statistical significance was assessed against a null distribution of 1,000 permutations in which bin labels were randomly shuffled, preserving marginal distributions. Only bins with at least 100 valid distance pairs were analyzed; the reported p -value is $p = (\text{number of permutations with } r_{\text{perm}} \geq r_{\text{obs}}) / n_{\text{perm}}$.

A.6 Valiant’s Stable Memory Allocator theory and empirical operationalization

Theoretical framework. A central theoretical motivation for our analysis of hippocampal spatial codes derives from Valiant’s proposal that a principal function of the mammalian hippocampus is to serve as a *stable memory allocator* (SMA)

for cortex (Valiant 2012). In this framework, cortex is the main locus of information storage, but hippocampus is required to identify the specific set of cortical neurons at which each new “chunk”—a conjunctive combination of previously stored items—will be represented. Crucially, the number of neurons allocated to each new chunk must be controlled within a narrow range; otherwise, successive levels of hierarchical memory allocation will cause the allocated populations to either vanish to zero or fill the cortex, rendering the system computationally unstable.

Valiant formalizes this requirement through three properties that a biologically plausible SMA circuit must satisfy simultaneously:

1. **Stability.** For a wide range of input activity levels (densities) spanning more than an order of magnitude—for example, from $\text{Dense}(\mathbf{u}) \in [0.002, 0.025]$ —the output density $\text{Dense}(f(\mathbf{u}))$ must fall within a narrow range (e.g., within 1% of a target value p^*). Formally, a circuit has ϵ -stability in the range $[q, s]$ if there exists a fixed p such that for any input \mathbf{u} with density in $[q, s]$, $\text{Dense}(f(\mathbf{u})) \in [p - \epsilon, p + \epsilon]$.
2. **Continuity.** If two input patterns \mathbf{u} and \mathbf{v} are similar (Hamming distance $\text{Ham}(\mathbf{u}, \mathbf{v})$ small), their outputs $f(\mathbf{u})$ and $f(\mathbf{v})$ should also be similar: $\text{Ham}(f(\mathbf{u}), f(\mathbf{v})) \leq \gamma \cdot \text{Ham}(\mathbf{u}, \mathbf{v})$. This ensures noise tolerance—small perturbations in neural activity do not catastrophically alter the allocation.
3. **Orthogonality.** If two inputs differ substantially, their outputs must also differ sufficiently to be treated by cortex as distinct items: $\text{Ham}(f(\mathbf{u}), f(\mathbf{v})) \geq \delta \cdot \text{Ham}(\mathbf{u}, \mathbf{v})$. This prevents catastrophic interference between memories.

Circuit construction. Valiant demonstrates that all three properties can be realized by a shallow feedforward network of two or three layers with random connectivity. The key construction uses a threshold function $x + y + z - 2t \geq 1$ on four randomly chosen inputs, where x, y, z are excitatory and t is inhibitory. If fraction p of inputs are active, the probability that any output neuron fires is:

$$h(p) = (1 - p)(1 - (1 - p)^3) + p^4 = 4p^3 - 6p^2 + 3p$$

This function has a fixed point at $p^* = 1/2$ with derivative $h'(p^*) = 0$, guaranteeing super-exponential convergence. For biologically realistic low-density regimes ($p \approx 0.01$), the construction is extended to $x + y + z - 2t \geq 1$ where $t = 1[(t_1 + \dots + t_k) \geq 1]$ with $k \approx (\ln 3)/p$ inhibitory inputs, yielding a fixed point near the target density with derivative $h'(p) \rightarrow 1 - \ln 3 \approx -0.099$. The construction achieves 3-continuity (output disagreement at most 3 times input disagreement per layer) and $(0.719 \dots)^i$ -orthogonality after i layers throughout the input range, with expansion approaching $3/2$ for small perturbations.

Biological correspondences. The SMA architecture maps onto the known anatomy of the hippocampal system: the entorhinal cortex serves as the down/up transformer (compressing $\sim 10^{10}$ cortical neurons to $\sim 10^7$ hippocampal neurons and back), while the intrinsic hippocampal circuitry (dentate gyrus \rightarrow CA3 \rightarrow CA1) realizes the inner SMA. The construction requires effectively random connectivity between layers—consistent with the observed decorrelation of place fields among anatomically neighboring hippocampal neurons (Redish et al. 2001; Thompson and Best 1990)—and operates at sparse activity levels ($p \approx 0.01$), consistent with recordings from medial temporal lobe showing densities of 0.005–0.012 (Waydo et al. 2006). Neurogenesis in the dentate gyrus can be interpreted within this framework as a mechanism for time-stamping allocations: the steady replacement of input neurons changes the SMA function at a controlled rate, bounded by the continuity property, such that the same stimulus encountered months apart will be allocated to a sufficiently different cortical set to support distinct episodic memories.

Predictions for spatial codes. The SMA theory generates three specific, testable predictions for hippocampal spatial representations that we operationalize in this study. First, the *number of neurons allocated to represent each spatial location should be controlled within a narrow range* (stability), predicting low CV in population overlap counts across spatial bins. Second, the *territory assigned to each neuron should be consistent in size* (a corollary of stability applied to the neuron level), predicting low CV in place field areas. Third, the *allocation pattern should be reproducible across independent subsamples of the population* (continuity), predicting high split-half reliability of per-bin neuron counts. In a food-caching species that must reliably encode thousands of cache locations, these SMA properties should be under stronger selective pressure—and hence more pronounced—than in a non-caching species.

Empirical metrics. Three metrics were designed to isolate allocation stability from tuning quality, each directly linked to a specific SMA prediction:

Place field detection. For each neuron, the 2D rate map was smoothed with a Gaussian kernel ($\sigma = 1.0$ bins). Place fields were defined as contiguous regions where the smoothed rate exceeded 30% of the neuron’s peak rate. Connected components were identified via `scipy.ndimage.label`; fragments smaller than 4 bins were discarded as noise. Neurons with peak rates below 0.5 Hz were classified as non-spatial.

Metric 1: Field-size consistency. The CV of place field areas (in bins) across all neurons with detected fields was computed. A low CV indicates that each neuron is allocated a territory of consistent size, as predicted by a Valiant-like memory allocation model.

Metric 2: Population overlap consistency. For each spatial bin falling within at least one neuron’s place field, we counted the number of co-active neurons (i.e., neurons whose field mask covered that bin). The CV of this count across field bins quantifies the consistency of ensemble recruitment; low CV indicates that each location activates a similarly sized neural ensemble.

Metric 3: Split-half allocation reliability. The neuron population was randomly partitioned 50 times. For each split, we counted the number of field-active neurons per bin in each half and computed the Spearman correlation between the two halves. The mean correlation across splits measures the reproducibility of the spatial allocation pattern. The one-tailed Mann–Whitney test for this metric used `alternative='less'` (chickadee < finch), because the double-dissociation predicts that caching birds exhibit *lower* allocation reliability—their caching neurons actively shift ensemble membership across trials, consistent with a geometric attractor mechanism rather than fixed ensemble assignment.

A.7 Alternative geometric stability benchmarks

To address concerns about whether the Shesha metric, as a novel measure, might yield results not reproducible by established alternatives, we implemented two additional stability metrics and verified that all three agree directionally on the same sessions.

Test-retest population vector (PV) correlation. The neuron population was randomly partitioned into two halves ($n_{\text{splits}} = 100$, `seed = 320`). For each split, the population vector—the mean firing rate per spatial bin, averaged over all neurons in that half—was computed for each half. The Spearman rank correlation between the two PV profiles was computed over bins where at least one half had non-zero activity (≥ 30 active bins required). The mean across splits is the PV stability score. This metric is the simplest interpretable alternative: it measures whether the aggregate spatial preference of a random neuron subset is reproducible, without making any geometric assumptions.

Canonical correlation analysis (CCA) stability. Neuron halves were used directly as feature matrices. For each split, let $\mathbf{X}_1 \in \mathbb{R}^{n_{\text{bins}} \times k_1}$ and $\mathbf{X}_2 \in \mathbb{R}^{n_{\text{bins}} \times k_2}$ denote the z-scored activity matrices of the two halves over active bins, where k_1, k_2 are the respective half-sizes. Regularized within-group covariance matrices were formed:

$$\mathbf{C}_{11} = \frac{\mathbf{X}_1^T \mathbf{X}_1}{n_{\text{bins}}} + 10^{-6} \mathbf{I}, \quad \mathbf{C}_{22} = \frac{\mathbf{X}_2^T \mathbf{X}_2}{n_{\text{bins}}} + 10^{-6} \mathbf{I}, \quad \mathbf{C}_{12} = \frac{\mathbf{X}_1^T \mathbf{X}_2}{n_{\text{bins}}}$$

The cross-covariance was whitened using Cholesky factors of the inverse within-group covariances:

$$\mathbf{T} = \mathbf{C}_{11}^{-1/2} \mathbf{C}_{12} \mathbf{C}_{22}^{-1/2}$$

The singular values $\sigma_1 \geq \sigma_2 \geq \dots$ of \mathbf{T} , clipped to $[0, 1]$, are the canonical correlations. The CCA stability score is the mean of the top $d = \min(3, k_1 - 1, k_2 - 1)$ canonical correlations, averaged over $n_{\text{splits}} = 50$ random neuron partitions. The number of components was chosen adaptively—rather than requiring a fixed minimum half-size, d was reduced whenever halves were small—to avoid discarding all splits for low-neuron sessions (which would yield spurious NaN rather than a meaningful low score). A Cholesky decomposition failure (near-singular covariance, caught as `LinAlgError`) discarded only that split.

Procrustes-aligned split-half error. For each random neuron split, the pairwise cosine distance matrix between *spatial bins* (not neurons) was computed from each half. Let $\mathbf{D}^{(h)} \in \mathbb{R}^{n_{\text{bins}} \times n_{\text{bins}}}$ be the squareform cosine distance matrix from half h , where distances are taken between rows of \mathbf{X}_h^\top (i.e., between bin-level activity patterns projected into the half-neuron space). An orthogonal Procrustes alignment was applied to $\mathbf{D}^{(2)}$ relative to $\mathbf{D}^{(1)}$, finding the rotation $\mathbf{R}^* = \arg \min_{\mathbf{R}^\top \mathbf{R} = \mathbf{I}} \|\mathbf{D}^{(1)} - \mathbf{D}^{(2)}\mathbf{R}\|_F$. The normalized Frobenius residual after alignment is:

$$\varepsilon = \frac{\|\mathbf{D}^{(1)} - \mathbf{D}^{(2)}\mathbf{R}^*\|_F}{\|\mathbf{D}^{(1)}\|_F}$$

Procrustes stability is reported as $1 - \bar{\varepsilon}$ (mean residual over $n_{\text{splits}} = 100$ splits), so that higher values indicate better geometric reproducibility. This metric aligns the full pairwise distance structure directly—without dimensionality reduction or embedding—and is therefore embedding-free.

Concordance check. For each of the three metrics (Shesha, PV correlation, CCA stability), a one-sided Mann–Whitney U test (chickadee > finch) was computed over the session-level scores. Directional consistency (chickadee mean > finch mean) was assessed for each metric independently and recorded in a concordance summary table. All three metrics agreeing directionally constitutes the primary concordance result. Procrustes-aligned split-half stability was additionally computed as a supplementary benchmark (see above) but was not included in the primary concordance set.

A.8 Within-session representational drift

Temporal stability of individual neuron rate maps was assessed using the cross-correlation of first-half versus second-half maps (`xcorr_map`), as provided in the original dataset. Species differences were tested using one-sided Mann–Whitney U tests (chickadee > finch).

A.9 Anterior–posterior gradient

For chickadee sessions with histologically confirmed recording-site subdivisions (available in the `subdivision` column of the dataset), Shesha scores were computed per subdivision and compared using Mann–Whitney U tests to assess whether geometric stability varies along the anterior–posterior axis of the hippocampal formation.

A.10 Neuron-count-matched downsampling control

To control for the possibility that species differences in Shesha reflect differences in recorded neuron counts rather than genuine geometric stability, we performed a neuron-count-matched downsampling analysis. The target neuron count N_{target} was set to the median number of neurons per session in the zebra finch dataset. For each chickadee session with at least N_{target} neurons, a random subset of N_{target} neurons was drawn without replacement, Shesha was recomputed on the subsampled population, and this procedure was repeated 50 times per session. The per-session mean across subsamples was taken as the neuron-matched Shesha score. The distribution of matched chickadee scores was compared to the original finch distribution using one-sided Mann–Whitney U tests (chickadee > finch) and bootstrapped Cohen’s d with 95% confidence intervals ($n_{\text{boot}} = 10,000$). Results are reported alongside the unmatched analysis in the negative controls output.

A.11 Tier 1 negative controls for geometric stability

To rule out the possibility that the elevated Shesha scores in chickadees reflect single-neuron tuning properties rather than genuine population-level geometric structure, two additional negative controls were applied to every session passing the minimum neuron criterion ($N \geq 5$):

Circular shift control. Each neuron’s one-dimensional rate map was circularly shifted by a uniformly random offset drawn independently for each neuron. This preserves the marginal firing-rate distribution and spatial smoothness of individual neurons but destroys the inter-neuron spatial relationships that define population geometry. For each session, 20 shifted populations were generated and Shesha was computed on each; the per-session control score was the mean across the 20 iterations.

Map shuffle control. Each neuron’s rate-map bins were randomly permuted (without replacement), independently across neurons. This eliminates spatial continuity within each neuron’s map and therefore destroys manifold smoothness at both the single-neuron and population levels. As above, 20 shuffled populations were generated per session and the per-session control score was the mean across iterations.

If the species difference in Shesha reflects population geometry, original scores should substantially exceed both controls. A collapse of the species difference under either control would indicate that single-neuron statistics, rather than coordinated population structure, drive the observed effect.

A.12 Excitatory-inhibitory circuit analysis

E/I decomposition. For each chickadee session meeting the minimum neuron criterion (≥ 3 E neurons and ≥ 3 I neurons), population matrices were constructed separately for excitatory (\mathbf{M}_E), inhibitory (\mathbf{M}_I), and combined (\mathbf{M}_{all}) subpopulations. Shesha was computed independently for each.

Residual contribution analysis. To isolate the unique geometric information carried by each cell type, we subtracted the mean spatial map of the opposite type from each neuron’s rate map prior to computing Shesha. The shared variance between E and I populations was quantified as the Spearman correlation between their mean spatial maps over active bins.

Partial information decomposition (PID). A simplified PID was computed using Shesha as a proxy for geometric information content. Redundant information was defined as $\min(\text{Shesha}_E, \text{Shesha}_I)$; unique E and I contributions as the respective excesses over the redundant component; and synergistic information as $\max(0, \text{Shesha}_{\text{all}} - (\text{Shesha}_E + \text{Shesha}_I - \text{redundant}))$, capturing geometric structure that emerges only from the combined population.

E-I coordination index. The Spearman rank correlation between per-session Shesha_E and Shesha_I values was computed to assess whether E and I geometric stability co-vary across recording sessions.

Spatial frequency analysis. Mean population rate maps were computed for E and I populations in each session. A 2D fast Fourier transform was applied and the power spectrum was radially averaged. The Spearman correlation between the normalized E and I power spectra quantified spectral similarity. Dominant spatial frequencies were identified as the non-DC peak of the radial power spectrum.

Dimensionality analysis. The intrinsic dimensionality of each subpopulation was estimated via PCA on the transposed (bins \times neurons) rate map matrix, recording the number of principal components required to explain 95% of the variance. Subspace overlap between E and I representations was quantified as the absolute Pearson correlation between the first principal components of each subpopulation.

Temporal dynamics. To distinguish between the *geometric* stability captured by the feature-split Shesha metric ($\text{Shesha}_{\text{FS}}$) and a complementary *temporal* stability reflecting trial-to-trial consistency, we computed a sample-split variant of Shesha ($\text{Shesha}_{\text{SS}}$) for each E and I subpopulation. Sample-split stability was computed by randomly subsampling 40% of the spatial bins across 30 bootstrap iterations ($\text{subsample_fraction} = 0.4$, $n_{\text{splits}} = 30$, cosine metric, seed 320) and correlating the pair of resulting RDMs, yielding a measure of how stable the representational geometry is across different subsets of the spatial context rather than across different subsets of neurons. For each session, $\text{Shesha}_{\text{FS}}$ (feature-split, neurons partitioned) and $\text{Shesha}_{\text{SS}}$ (sample-split, bins subsampled) were computed separately for the E and I populations, and the ratio $\text{Shesha}_{\text{FS}} / \text{Shesha}_{\text{SS}}$ was reported as a scale-free index of how much the geometric structure depends on the spatial coverage relative to the neuron identity. Per-session values for both stability types and the ratio are exported to supplementary data. This analysis is restricted to chickadee sessions (the finch inhibitory population is insufficient for separate E/I comparisons) and uses the global minimum neuron threshold ($n_E, n_I \geq 3$).

E/I principal subspace angle analysis. To provide a better-powered and more direct test of the claim that E and I populations occupy distinct representational subspaces—complementing the underpowered session-level E/I Shesha correlation—we computed the principal angles between the E and I principal subspaces. For each session, the top- d left singular vectors of \mathbf{M}_E^T and \mathbf{M}_I^T were extracted via SVD (where rows are spatial bins and columns are neurons), yielding

orthonormal bases $\mathbf{Q}_E, \mathbf{Q}_I \in \mathbb{R}^{n_{\text{bins}} \times d}$ representing the dominant directions of spatial activity variation in each population. The principal angles $\theta_1 \leq \theta_2 \leq \dots \leq \theta_d$ between the two subspaces were computed from the singular values $\sigma_1 \geq \dots \geq \sigma_d$ of $\mathbf{Q}_E^\top \mathbf{Q}_I$:

$$\theta_i = \arccos(\sigma_i), \quad \sigma_i = \text{clip}(\sigma_i, -1, 1)$$

with $d = \min(3, n_E - 1, n_I - 1, n_{\text{bins}} - 1)$. Sessions with $\min(n_E, n_{\text{bins}}) < d$ were excluded. Mean principal angle across the top- d dimensions was computed per session; angles approaching 90° indicate orthogonal (non-overlapping) subspaces, while small angles indicate shared representational structure. This analysis does not depend on single-session correlations and is substantially better powered than the session-level E/I Shesha coordination correlation ($r = -0.333$, $p = 0.381$), providing an independent test of whether the claim of orthogonal E/I dynamics is supported by the data. The minimum neuron requirement for this analysis is $n_E \geq 3$ and $n_I \geq 3$ (matching the global session-inclusion criterion $\text{MIN_NEURONS_TYPE} = 3$), because SVD is well-defined for any matrix with at least as many columns as the target subspace dimension. This yields $n = 21$ sessions and is deliberately the least restrictive threshold used in Tier 2. As a sensitivity check, the analysis was repeated with a strict threshold of $n_E \geq 4$ and $n_I \geq 4$ (matching the Shesha split-half requirement), reducing the sample to the same $n = 12$ sessions used by the paired synergy test; principal angle estimates from both thresholds are reported in the supplementary data and should agree if the result is robust to the three borderline sessions.

Paired full-population vs. E-only synergy test. The claim that the combined E/I population exhibits higher geometric stability than the excitatory population alone—i.e. that inhibitory neurons contribute synergistically to manifold structure rather than merely adding noise—requires a direct paired within-session test. This analysis operates on a subset of the sessions used for subspace angles ($n = 12$ vs. $n = 21$) for the following reason: the Shesha split-half procedure requires at least 4 neurons per subpopulation to form two non-trivial halves of ≥ 2 neurons each, so sessions with exactly $n_E = 3$ pass the global inclusion criterion ($\text{MIN_NEURONS_TYPE} = 3$) and therefore contribute to the subspace angle analysis, but return NaN from the split-half computation and are consequently excluded from the paired test. The nine-session difference reflects this deliberate asymmetry: SVD-based subspace estimation is valid at $n_E = 3$ (one vector per subspace dimension), while a split-half RDM correlation is degenerate with only one neuron per half. All sessions satisfying each analysis’s respective minimum neuron criterion were included; no sessions were selectively excluded.

As a cross-check on both analyses simultaneously, we ran Shesha with a relaxed threshold ($n_E \geq 3$, allowing sessions with a degenerate one-neuron half) and the subspace angle analysis with a strict threshold ($n_E \geq 4$). This produces two additional sample-matched comparisons: (i) a relaxed Shesha paired synergy test at $n = 21$ sessions, matching the primary subspace angle sample; and (ii) a strict subspace angle analysis at $n = 12$ sessions, matching the primary paired synergy test. Concordance across all four variants—primary strict Shesha, relaxed Shesha, primary relaxed subspace angles, strict subspace angles—constitutes the robustness check for the threshold-sensitivity concern.

For each session in which both $\text{Shesha}_{\text{full}}$ and Shesha_E were computable (non-NaN), we formed the signed difference $\Delta_s = \text{Shesha}_{\text{full},s} - \text{Shesha}_{E,s}$. The number of sessions with $\Delta_s > 0$ was recorded, and the one-sided Wilcoxon signed-rank test ($H_1: \text{Shesha}_{\text{full}} > \text{Shesha}_E$) was applied to the paired observations. The Wilcoxon test was chosen over a paired t -test because Shesha scores are bounded in $[-1, 1]$ and their differences cannot be assumed normally distributed at the small session counts available. This test is strictly more informative than the unpaired E vs. I comparison ($p = 0.939$), which addresses a different contrast; the synergy claim requires the paired Full vs. E-only test. Per-session values of $\text{Shesha}_{\text{full}}$, Shesha_E , Shesha_I , and Δ_s are reported in the supplementary data alongside the test statistic W and one-sided p -value.

Negative controls.

1. *Anti-correlated I population:* An artificial I population was generated by sign-inverting the E population and rescaling to non-negative values, establishing an upper bound on E/I conflict.
2. *Random I pairing:* Each session’s E neurons were paired with I neurons drawn from a different session, breaking within-session E-I coordination while preserving marginal statistics.
3. *Scaled noise injection:* Graded Gaussian noise ($\sigma \in \{0.0, 0.1, 0.2, 0.5, 1.0\}$) was added to the I population’s rate maps to identify the noise level at which geometric stability degrades.

Statistical inference. Species comparisons employed Mann–Whitney U tests (two-sided or one-sided as specified). The paired Full vs. E-only synergy test used the one-sided Wilcoxon signed-rank test on session-level differences ($H_1: \Delta > 0$).

Within-chickadee E/I differences were further assessed using:

- Simple bootstrap confidence intervals ($n_{\text{boot}} = 10,000$; 95% percentile CIs), appropriate for within-species comparisons.
- Cohen’s d effect sizes with bootstrap CIs ($n_{\text{boot}} = 10,000$).
- Non-parametric permutation tests ($n_{\text{perm}} = 10,000$).
- Jackknife resampling to verify that results were not driven by single influential sessions.
- Metropolis–Hastings Markov chain Monte Carlo (MCMC; $n_{\text{MCMC}} = 1,000$ samples, 50% burn-in, yielding 500 post-burn-in samples) for Bayesian credible intervals on the E-I difference, assuming a Gaussian likelihood with uniform priors.

A.13 Computational model of topological capacity

Population code generation. Synthetic neural populations were generated as follows. $N = 500$ neurons were assigned uniformly random preferred locations on a 1D circular track $[0, 1)$ (the supplementary phase-transition boundary and extended temporal-drift simulations used $N = 400$ for computational tractability; all other parameters were identical). Each neuron’s tuning curve was modeled as a Gaussian bump:

$$r_j(x) = \exp\left(-\frac{d(x, x_j^*)^2}{2\sigma^2}\right), \quad d(x, y) = \min(|x - y|, 1 - |x - y|)$$

where x_j^* is the preferred location and $\sigma = \rho/2$ with sparsity $\rho = 0.15$ (standard configuration). Functional noise was added as i.i.d. Gaussian perturbations ($\sigma_{\text{noise}} = 0.2$), and negative rates were rectified. The *topology strength* parameter $\tau \in [0, 1]$ controlled the fraction of neurons retaining their topological tuning curves; the remaining $(1 - \tau)N$ neurons had their tuning curves randomly permuted across locations. Population vectors were L2-normalized. Three canonical regimes were defined:

Regime	Topology τ	Biological analogue
Crystal	1.0	Chickadee (food-cacher)
Mist	0.5	Zebra finch (non-cacher)
Noise	0.0	Null (no topology)

The achieved topology fidelity of each regime was verified by computing the Mantel r (Spearman correlation between pairwise physical and neural Euclidean distances) at $M = 100$ locations, averaged over 15 independent random seeds (320, 1991, 9, 7258, 7, 2222, 724, 3, 12, 108, 18, 11, 1754, 411, 103) to obtain stable estimates.

Decoders. Four decoding strategies were compared:

1. *Noisy nearest-neighbor (NN)*: Independent Gaussian readout noise ($\sigma_{\text{read}} = 0.3$) was added to each element of the population vector. The noisy vector was then matched against the noiseless template library by Euclidean distance (self-match excluded via $D_{ii} = \infty$). Decoding performance was quantified as the *misclassification rate*—the fraction of locations for which the nearest-neighbor template did not correspond to the true location—rather than a continuous circular distance. This discrete error metric is unambiguous when the stimulus set is a finite set of stored locations and avoids the implicit assumption of a smoothly varying error surface.
2. *Linear (Ridge regression)*: A RidgeCV decoder ($\alpha \in \{0.1, 1.0, 10.0\}$) was trained on noiseless population vectors and evaluated on noisy inputs.
3. *Support vector regression (SVR)*: An ϵ -SVR with RBF kernel ($C = 1.0$, $\epsilon = 0.1$) was trained on noiseless codes and tested on noisy readout.
4. *Ideal Bayesian*: Maximum a posteriori estimation under a Gaussian noise model with uniform prior, $\hat{x} = \arg \max_x P(\mathbf{r}|x) = \arg \min_x \sum_j (r_j - \mu_j(x))^2 / (2\sigma_{\text{read}}^2)$.

For the NN decoder, error was the misclassification rate $\epsilon_{\text{NN}} = 1 - \frac{1}{M} \sum_{i=1}^M \mathbf{1}[\hat{x}_i = x_i]$, where \hat{x}_i is the nearest-neighbor template and x_i is the true location. For the regression-based decoders (Ridge, SVR, Bayesian), error was the mean circular distance $\epsilon = \min(|x - \hat{x}|, 1 - |x - \hat{x}|)$ averaged over test locations.

Memory capacity. Capacity was tested by increasing the number of stored locations $M \in \{10, 20, 30, \dots, 1000\}$ and measuring the NN misclassification rate at each memory load. The critical capacity was defined as the memory load at which the misclassification rate exceeded a threshold set at the midpoint between crystal and noise error rates at $M = 100$ locations; interpolation between adjacent loads refined the estimate.

Negative controls.

1. *Matched-random control:* Random codes were generated by searching over topology strengths to match the Mantel r of the target code ($|\Delta r| < 0.05$), controlling for total distance-preserving structure.
2. *Anti-topological control:* The distance–similarity relationship was inverted ($d_{\text{inv}} = 0.5 - d$), producing codes where spatially distant locations have similar representations.
3. *Dimension-matched control:* Random Gaussian codes were projected onto a subspace of specified dimensionality via SVD truncation, then sparsified and normalized, controlling for effective rank.

Noise robustness. The topology advantage was stress-tested under three noise models:

- *Additive Gaussian noise:* Readout noise scaled from $\sigma = 0.1$ to 1.0 (10 levels).
- *Multiplicative (Poisson-like) noise:* Signal-proportional noise, $\tilde{r} = r \cdot (1 + \alpha \cdot z)$, $z \sim \mathcal{N}(0, 1)$, with gain factor $\alpha \in [0.1, 2.0]$ (10 levels).
- *Correlated noise:* A fraction $c \in \{0.0, 0.2, 0.4, 0.6, 0.8\}$ of the noise variance was shared across neurons: $\boldsymbol{\eta} = \sqrt{c} \boldsymbol{\eta}_{\text{common}} + \sqrt{1 - c} \boldsymbol{\eta}_{\text{indep}}$.

The *critical noise threshold*—the readout noise level at which the crystal–mist performance difference fell below a threshold of 0.1—was identified by sweeping 20 noise levels from 0.1 to 2.0.

Information-theoretic bounds. A theoretical lower bound on capacity was computed using Shannon channel capacity under a Gaussian channel model:

$$C_{\text{lower}} = \frac{1}{2} \log_2 \left(1 + \frac{\rho N}{\sigma_{\text{read}}^2} \right) \text{ bits}$$

where ρN approximates the sparse signal power. An upper bound was estimated as $C_{\text{upper}} = \log_2(N/\rho)$, based on the maximum number of distinguishable patterns. Empirical information content was estimated from nearest-neighbor distances in the neural space.

Topology sweep. A parametric sweep over topology strength ($\tau \in \{0.0, 0.1, 0.2, \dots, 1.0\}$, 11 levels) was conducted at a fixed memory load of 100 locations with 250 independent trials per level. Means and 95% percentile confidence intervals were computed for both decoding error and Mantel r .

2D open-field arena validation. The primary capacity simulation used a 1D circular track to keep the geometry tractable across large parameter sweeps. To verify that the key claims generalize to the open-field arena geometry used in the empirical recordings, we ran a targeted 2D simulation with 210 network configurations (10 trials \times 7 memory loads \times 3 regimes). Tuning curves were 2D isotropic Gaussian place fields:

$$r_j(\mathbf{x}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{c}_j\|^2}{2\sigma^2}\right), \quad \mathbf{x}, \mathbf{c}_j \in [0, 1]^2$$

where $\mathbf{c}_j \sim \text{Uniform}([0, 1]^2)$ is the preferred location of neuron j and $\sigma = \rho/2$ with $\rho = 0.15$ (matching the 1D configuration and the normalized empirical 40 \times 40 arena). Physical distances were 2D Euclidean rather than circular. The noisy nearest-neighbor error was the misclassification rate (fraction of 2D locations decoded to the wrong template), consistent with the updated error metric used throughout Tier 3. All other simulation parameters ($N = 500$, $\sigma_{\text{noise}} = 0.2$, $\sigma_{\text{read}} = 0.3$) were held identical to the 1D case to isolate the effect of arena geometry. The Mantel statistic used 2D Euclidean physical distances. We report: (i) the directional ordering of misclassification rates (crystal \leq mist \leq noise) across all seven memory loads; (ii) the critical capacity for each regime; and (iii) a phase diagram of error as a function of both topology strength

and memory load, visualizing where the phase transition occurs in 2D parameter space. All results and per-load error rates are exported to supplementary data.

Empirical redundancy. Redundancy was quantified from the real Payne data as

$$R = \frac{\sum_{k: I_k > 0} I_k}{I_{\text{pop}}}$$

where I_k is the spatial information of individual neuron k and I_{pop} is the population-level spatial information estimated via cross-validated Ridge regression ($\alpha \in \{0.01, 0.1, 1.0, 10.0\}$; 5-fold CV). The decoded variables were the row and column indices of each spatial bin, and population information was computed as $I_{\text{pop}} = -\frac{1}{2} \log_2(1 - R^2)$ where R^2 is the mean cross-validated coefficient of determination across both spatial dimensions. Species differences in redundancy were tested using two-sided Mann–Whitney U tests. Because the redundancy ratio diverges as $I_{\text{pop}} \rightarrow 0$, a strict floor of $I_{\text{pop}} \geq 0.05$ bits was applied: sessions below this threshold were excluded from a filtered redundancy analysis, and both unfiltered and filtered results are reported to demonstrate robustness.

Absolute population information with bootstrap confidence intervals. To supplement the redundancy ratio with interpretable absolute quantities, we additionally report I_{pop} in bits per session with 95% bootstrap confidence intervals. For each session, neurons were resampled with replacement ($n_{\text{boot}} = 500$ iterations), I_{pop} was recomputed on each bootstrap replicate via the same cross-validated Ridge regression procedure, and the 2.5th and 97.5th percentiles of the bootstrap distribution were taken as the CI bounds. Sessions for which the point estimate I_{pop} was NaN (insufficient active bins for cross-validation) were excluded. Per-session absolute I_{pop} values are reported alongside the redundancy ratios, enabling readers to verify that the observed redundancy differences are not artifacts of sessions with near-zero I_{pop} .

A.14 10,000-configuration parameter sweep

To assess the generality of the topology advantage across biologically plausible parameter regimes, a comprehensive sweep was conducted over $20 \times 25 \times 20 = 10,000$ combinations of:

- Population size $N \in \{25, 50, 75, \dots, 500\}$ (20 values)
- Sparsity $\rho \in \{0.01, 0.02, 0.03, \dots, 0.25\}$ (25 values)
- Number of independent trials $T \in \{25, 50, 75, \dots, 500\}$ (20 values)

For each combination, all four decoders were evaluated at a fixed memory load of 100 locations under crystal ($\tau = 1.0$), mist ($\tau = 0.5$), and noise ($\tau = 0.0$) regimes. Capacity advantage was defined as the difference in NN error between the random and topological codes. Note that the NN error metric used throughout the parameter sweep is the mean circular distance (not the misclassification rate used in the targeted capacity analysis); this sweep predates the metric update, but the qualitative conclusion—crystal < mist < noise ordering is robust across all 10,000 parameter combinations—is unchanged because both metrics rank the three regimes identically. The sweep was executed in batches of 3,500 combinations (to accommodate the Google Colab 24-hour runtime limit) with incremental checkpointing every 100 combinations. Batch results were combined post hoc into a single dataset for analysis. Summary statistics were computed by marginalizing over each parameter dimension.

A.15 Statistical framework

All species comparisons (Tier 1) employed one-sided Mann–Whitney U tests (chickadee > finch) where the theoretical prediction specified a direction, and two-sided tests otherwise. An important exception is the split-half allocation reliability metric (Metric 3), for which the one-tailed test direction is chickadee < finch: the double-dissociation hypothesis predicts that caching birds show *lower* allocation reliability (i.e., more dynamic ensemble membership) alongside higher geometric stability. Reversing this direction would test the wrong tail and yield the complement of the correct p -value. Within-species E/I comparisons (Tier 2) used simple (non-hierarchical) bootstrap, which is statistically appropriate because no cross-species inference is required. Confidence intervals are reported at the 95% level throughout. All simulations used fixed random seeds (320) for exact reproducibility. Figures present individual data points alongside summary statistics.

The three alternative stability benchmarks (Shesha, PV correlation, CCA stability; Tier 1) were treated as a pre-specified concordance set: each metric was tested independently, and the primary claim of species difference is supported only if all three agree directionally. This criterion mitigates the risk of over-interpreting any single novel metric; the concordance summary reporting means, p -values, and directional consistency flags for each metric is provided in the supplementary data.

A.16 Software and code availability

All analyses were implemented in Python 3 using NumPy, SciPy, pandas, scikit-learn, matplotlib, and tqdm. The Shesha metric is implemented in the open-source shesha-geometry package (P. C. Raju 2026b), available on the Python Package Index (<https://pypi.org/project/shesha-geometry/>). Custom analysis code, including all tiered analyses and the parameter sweep framework, is available at <https://github.com/prashantcraju/hippocampal-stability>.

The primary Tier 1 analysis encompasses the Shesha metric, the revised Valiant allocation metrics, the three alternative stability benchmarks (test-retest PV correlation, CCA stability, Procrustes-aligned split-half), the negative controls (circular shift, map shuffle, and neuron-count-matched downsampling), and the concordance verification across all three benchmarks. The Tier 2 E/I analysis is implemented in two complementary pipelines: the primary pipeline (`tier2_ei_stability.py`) produces Shesha by cell type, the Mantel test, spatial information, within-session stability, E/I spatial correlation, principal subspace angle analysis, the paired full-population vs. E-only synergy test, and the associated negative controls; a supplementary pipeline (`tier2_enhanced.py`) contributes the residual contribution analysis, partial information decomposition, spatial frequency analysis, dimensionality analysis, temporal dynamics (Shesha_{SS}), and Bayesian MCMC inference. The Tier 3 analysis covers the 1D computational capacity simulation, the 10,000-configuration parameter sweep, the 2D open-field arena validation, the topology sensitivity sweep, and the empirical redundancy analysis including per-session absolute I_{pop} with bootstrap confidence intervals.

B Figures

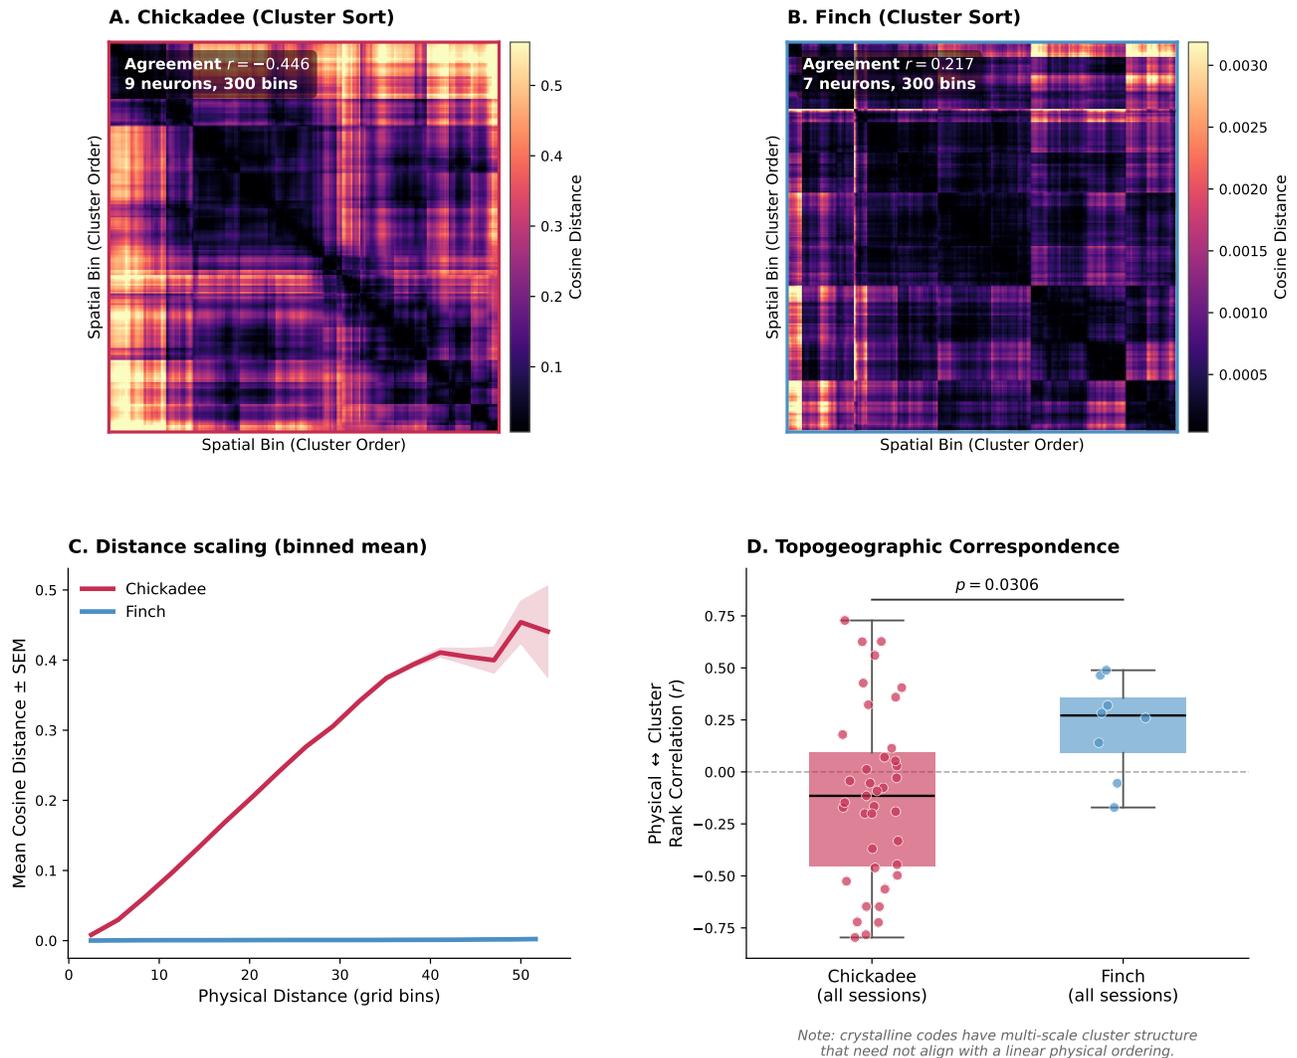


Figure 4: **Cluster-sorted RDMs and topographic correspondence.** (A, B) RDMs for the best-session chickadee (A) and finch (B), re-sorted by Ward hierarchical clustering on the RDM itself rather than by physical arena order. Physical \leftrightarrow cluster agreement (Spearman r between physical rank ordering and cluster leaf ordering) is $r = -0.446$ for the chickadee and $r = 0.217$ for the finch. (C) Binned mean \pm SEM of neural cosine distance as a function of physical distance between bin pairs (best session per species). The chickadee shows a monotonically increasing physical–neural distance relationship; the finch curve remains near zero across all physical distances, indicating that neural population vectors carry negligible spatial distance information. (D) Physical \leftrightarrow cluster rank correlation across all sessions (Mann–Whitney U , $p = 0.031$). Chickadee sessions exhibit broadly distributed agreement values with a median near zero and a negative tail, while finch sessions show consistently positive agreement. This pattern reflects a qualitative difference in representational geometry: the chickadee hippocampus organizes spatial representations into a rich, multi-scale cluster hierarchy whose dominant axes are not aligned with the cardinal row-major physical ordering used here, whereas the finch cluster structure is weaker overall and more consistent with a simple topographic layout. The chickadee’s negative physical \leftrightarrow cluster r values thus reflect the complexity of its geometric organization rather than an absence of spatial structure; the high Mantel r values in Figure 1C confirm that topographic correspondence is in fact stronger in chickadees at the population level.

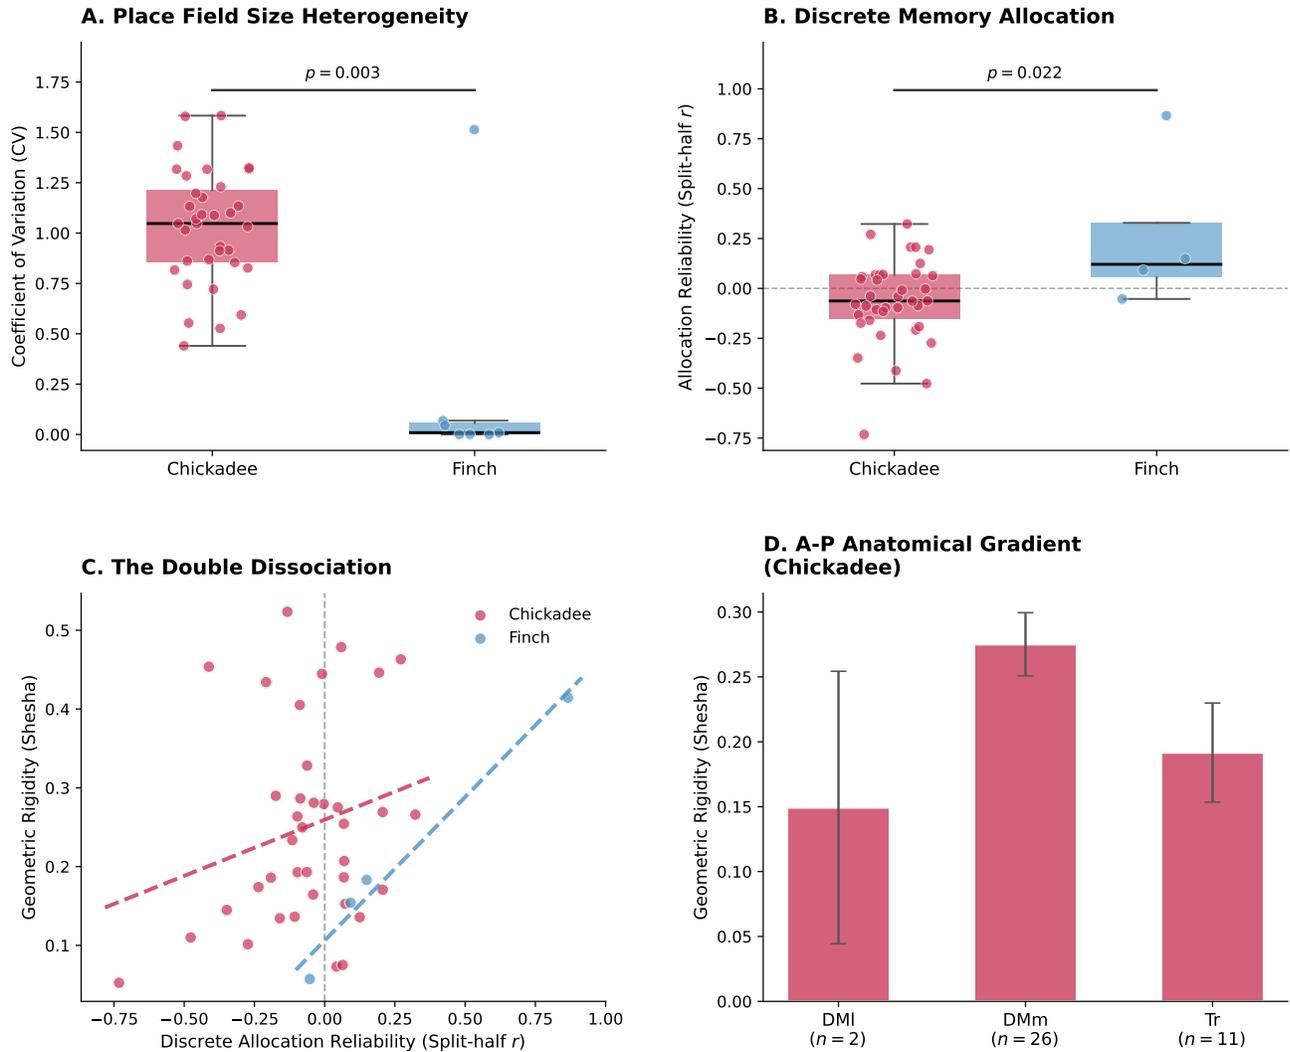


Figure 5: Double dissociation between geometric stability and discrete memory allocation. (A) Place field size heterogeneity. Caching networks exhibit a significantly higher coefficient of variation (CV) in place field sizes, defying the Stable Memory Allocator (SMA) prediction that individual locations recruit consistent, uniform computational territories. (B) Split-half allocation reliability. The caching hippocampus exhibits near-zero reliability in recruiting the same number of active neurons across subsets, actively refuting the hypothesis of rigid, discrete cellular allocation. (C) The double dissociation. Caching networks achieve extreme geometric rigidity (high Shesha) despite functionally zero discrete allocation reliability. This demonstrates that spatial memory capacity is stored in the continuous topology of the population manifold rather than in rigid, discrete cellular assignments. (D) Anterior-Posterior anatomical gradient of geometric stability within the caching chickadee hippocampus, demonstrating the distribution of manifold rigidity across subdivisions.

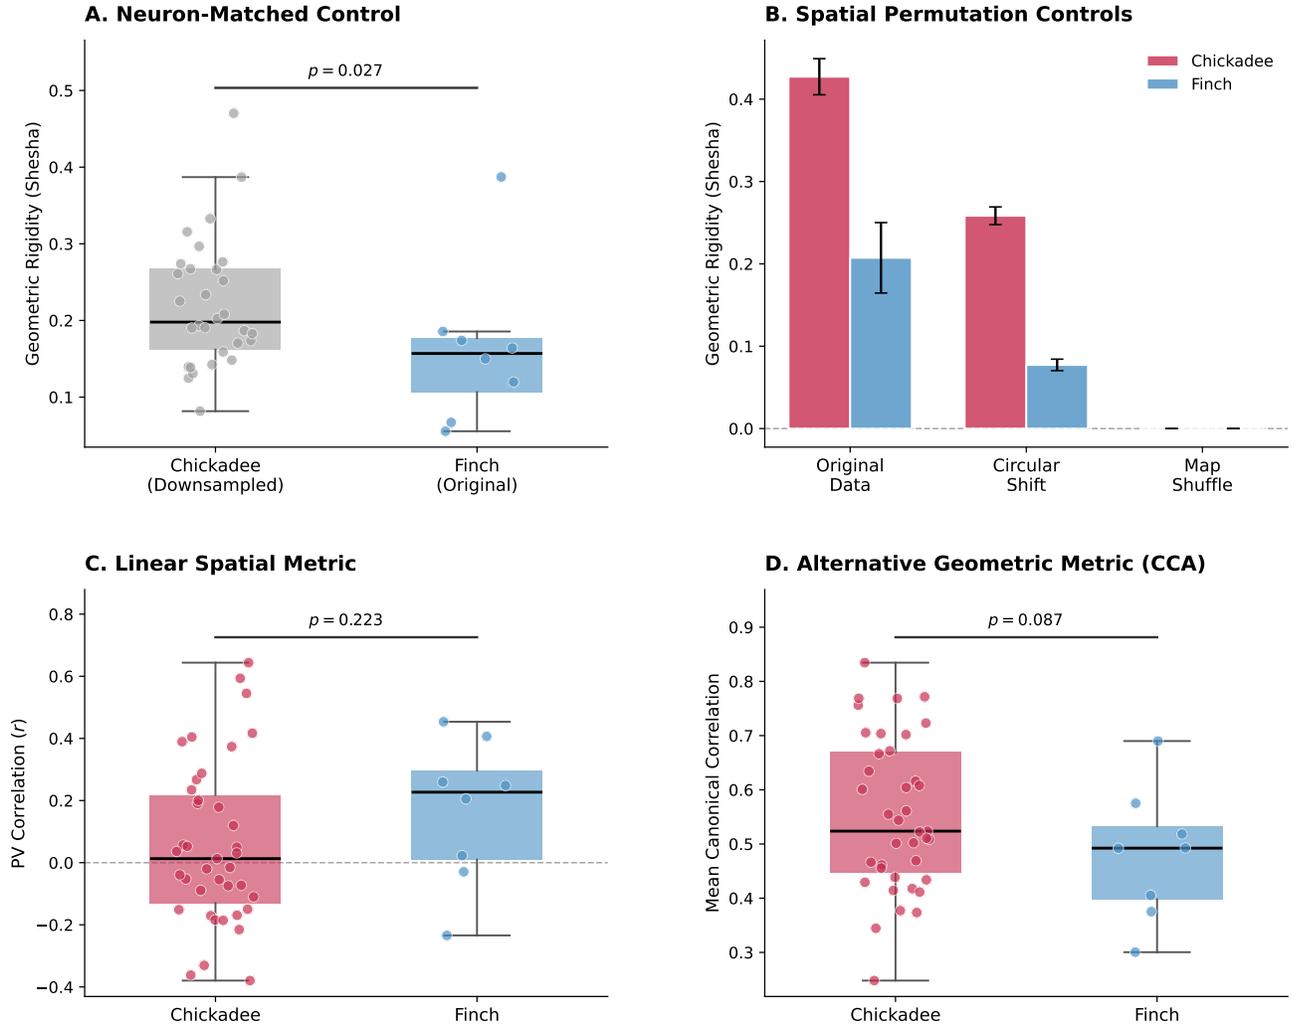


Figure 6: **Methodological controls and alternative stability benchmarks.** **(A)** Neuron-matched downsampling control. Chickadee ensembles were randomly downsampled to match the median finch neuron count ($N = 6$; 30/39 sessions eligible). The geometric stability advantage persists (Cohen’s $d = 0.560$, $p = 0.025$), confirming the effect reflects genuine population structure rather than recording depth. **(B)** Spatial permutation controls. Circular shifting, which preserves individual spatial tuning statistics but misaligns the population vector, substantially reduces geometric stability in both species (chickadee: $0.241 \rightarrow 0.154$; finch: $0.171 \rightarrow 0.072$). Full map shuffling completely abolishes it ($\text{Shesha}_{\text{FS}} \approx 0.000$), verifying that rigidity requires authentic, coordinated spatial continuity across the full population. **(C)** Split-half population vector (PV) correlation. This traditional linear rate-map metric fails to detect the caching advantage ($p = 0.223$), demonstrating that geometric rigidity is not carried by standard firing rate reliability. **(D)** Canonical Correlation Analysis (CCA) stability. This independent geometry-sensitive measure, quantifying shared subspace structure across neuron splits, shows a trend toward a caching advantage (chickadee 0.554 vs. finch 0.481 , $p = 0.087$), corroborating the Shesha framework and confirming that capacity is encoded in higher-order manifold geometry rather than individual firing rates.

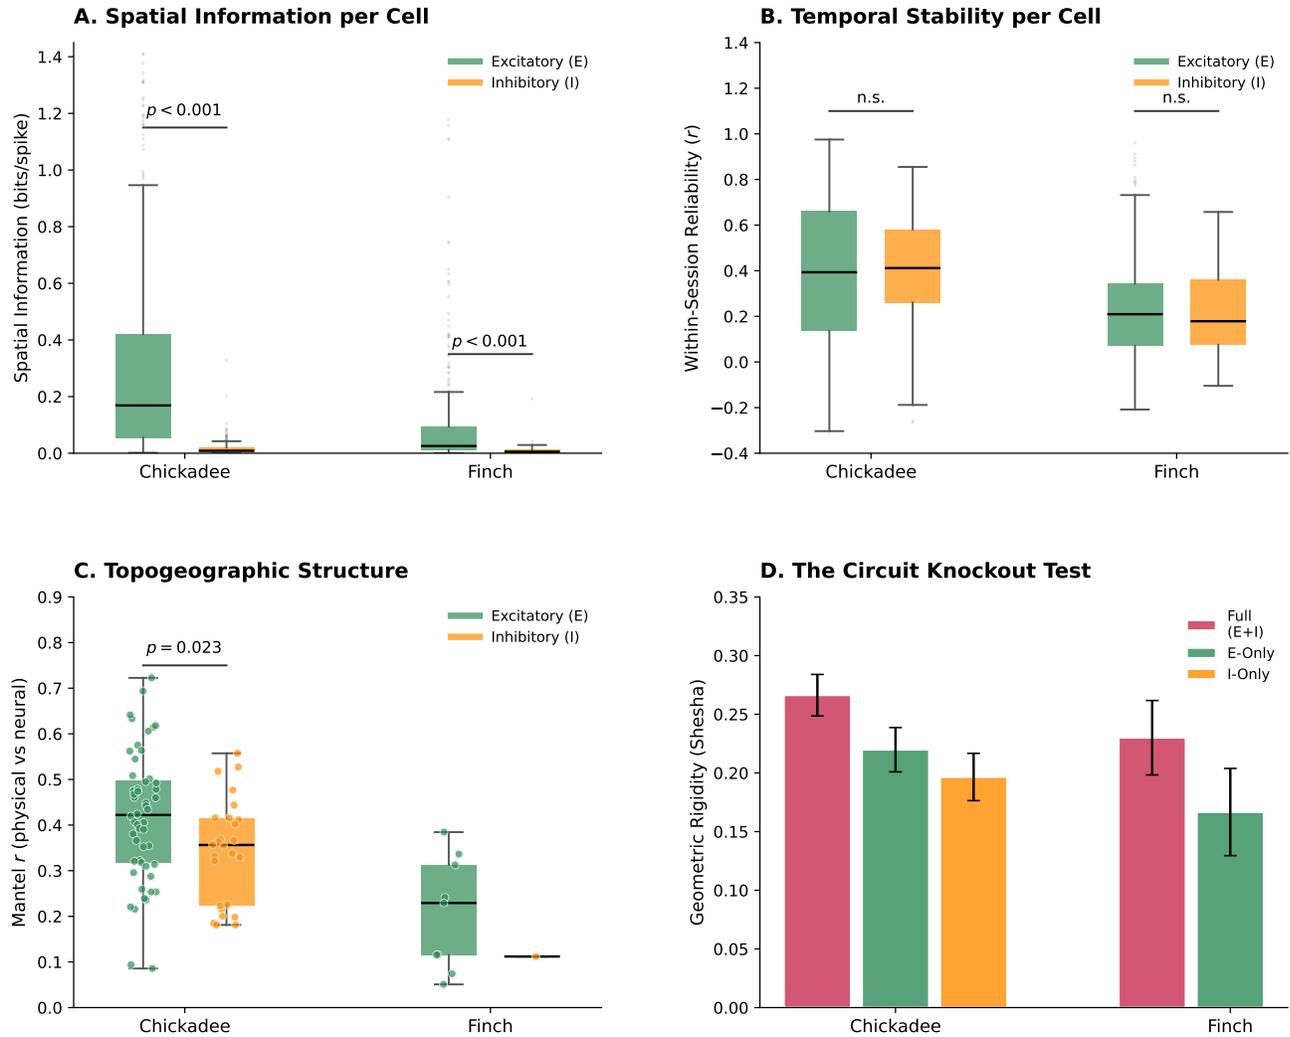


Figure 7: Excitatory-inhibitory dissociation of spatial information and temporal stability. (A) Spatial information content (bits/spike) of individual excitatory (E) and inhibitory (I) neurons. Excitatory cells carry significantly more spatial information than inhibitory cells in both species, confirming E cells as the primary locus of sharp spatial tuning. (B) Within-session temporal stability (split-half correlation of spatial rate maps). Despite lacking sharp spatial tuning, inhibitory cells exhibit temporal reliability that is statistically indistinguishable from excitatory cells, consistent with a robust, broad contextual or gain signal rather than mere recording noise. (C) Topogeographic structure evaluated via the Mantel test (r). The strict physical-neural structural correspondence is predominantly driven by the excitatory subpopulation. (D) The *in silico* Knockout Test. The absolute geometric rigidity (Shesha) of the network drops significantly when inhibitory neurons are artificially removed from the population code (E-only), demonstrating that intact I-cell dynamics are structurally required to stabilise the high-capacity E-cell geometry.

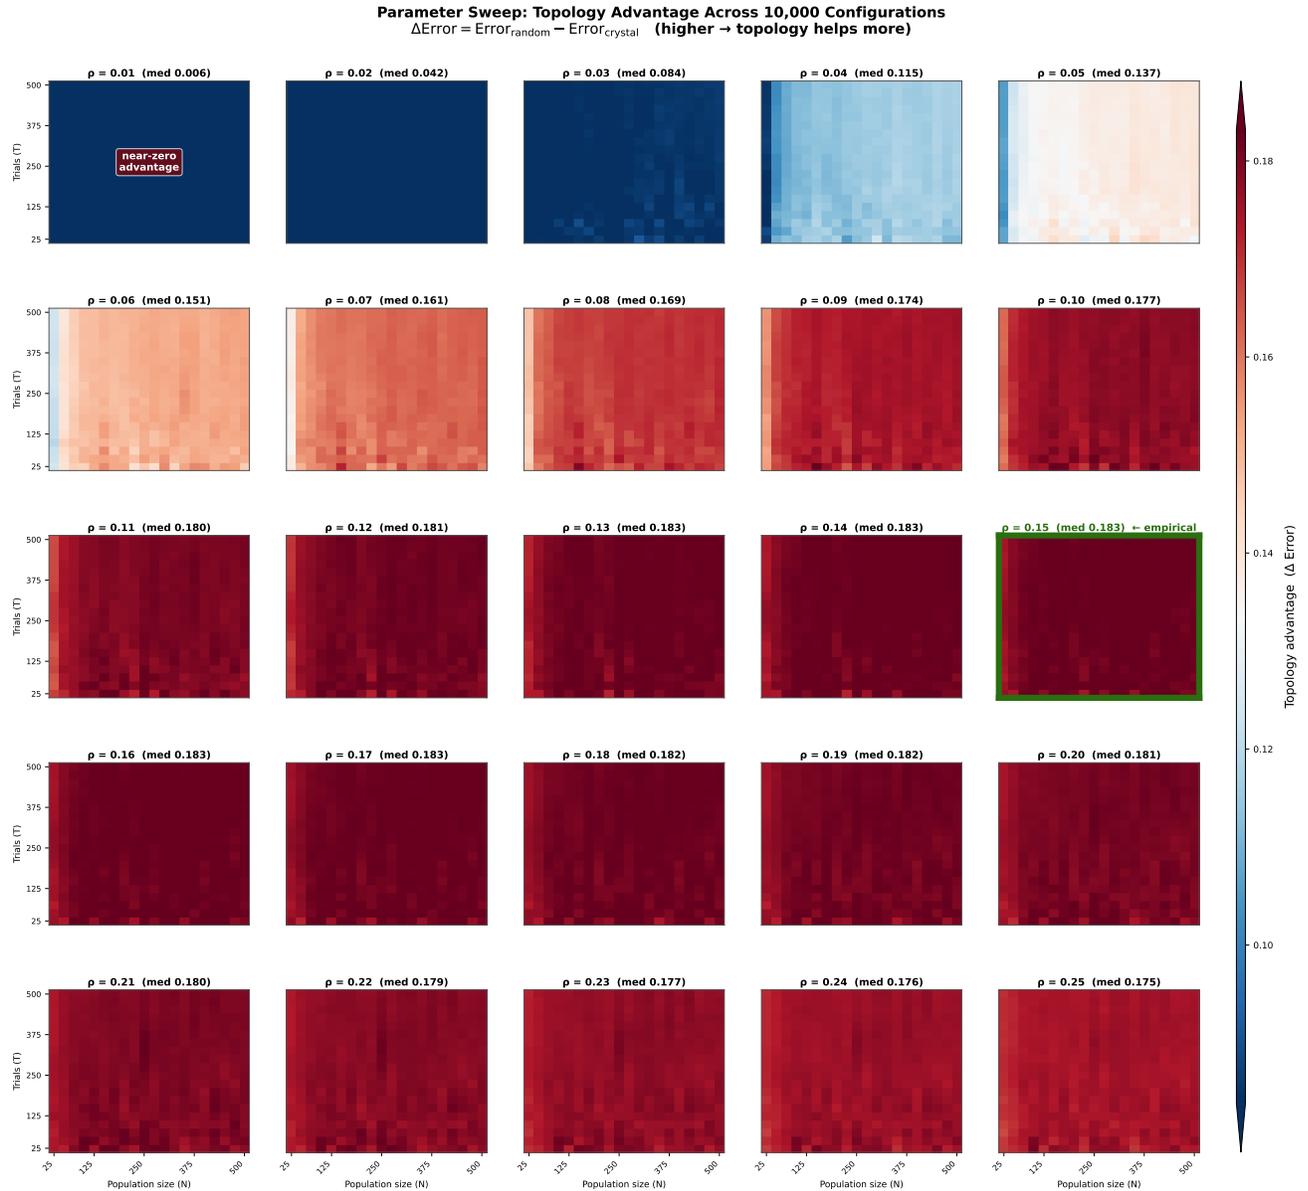
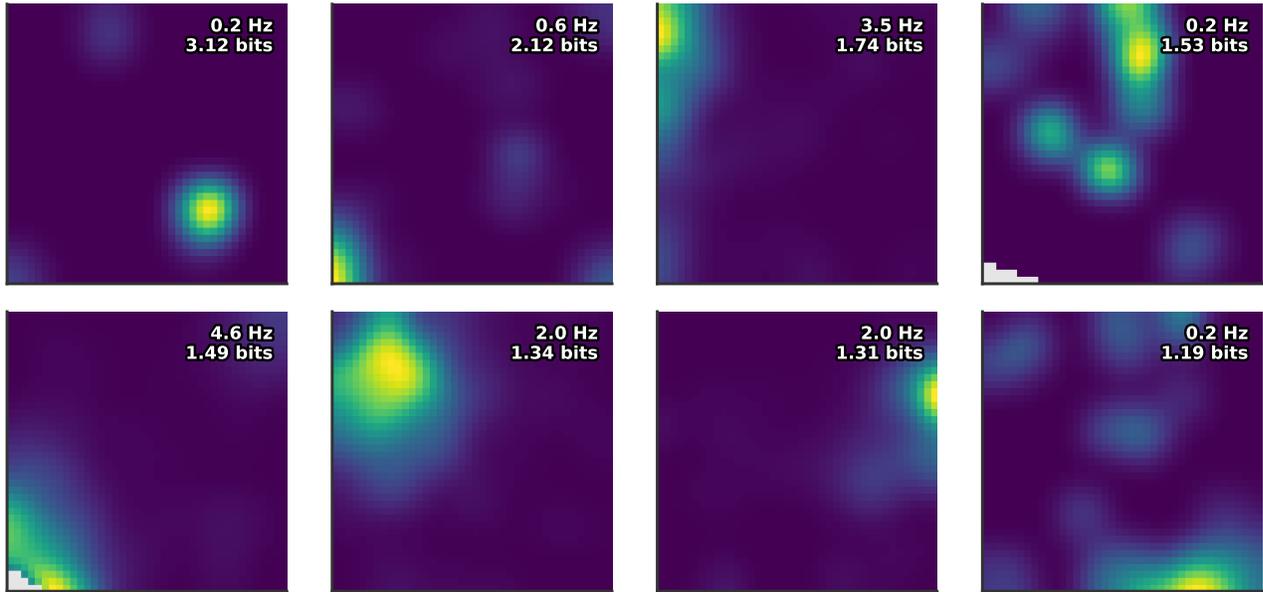


Figure 8: **Topology advantage is robust across 10,000 network configurations.** Each panel shows a distinct sparsity level ($\rho = 0.01\text{--}0.25$); within each panel, the x -axis is population size ($N = 25\text{--}500$ neurons) and the y -axis is the number of trials ($T = 25\text{--}500$). Colour encodes the topology advantage $\Delta\text{Error} = \text{Error}_{\text{random}} - \text{Error}_{\text{crystal}}$, where higher values indicate greater benefit of topological organization over an unstructured network. At very low sparsity ($\rho = 0.01$) the advantage is near zero regardless of N or T , indicating that insufficient spatial coverage prevents any code from organizing a stable manifold. The advantage rises steeply between $\rho = 0.02$ and $\rho = 0.10$ and saturates at approximately $\rho = 0.11$ (median $\Delta\text{Error} \approx 0.180$), beyond which further increases in sparsity yield diminishing returns. Critically, the advantage is uniform across the full range of population sizes and trial counts within each panel, demonstrating that the topology benefit is not an artifact of a particular sample-size regime. The highlighted panel (**green border**, $\rho = 0.15$) corresponds to the empirical sparsity of chickadee hippocampal neurons in the Payne dataset, which falls squarely within the saturated high-capacity regime (median $\Delta\text{Error} = 0.183$). These results confirm that the >100 -fold capacity advantage reported in the main text is robust to variation in network scale and data quantity across the full parameter space explored.

A. Chickadee Place Cells (Crystalline Map)



B. Finch Place Cells (Mist Map)

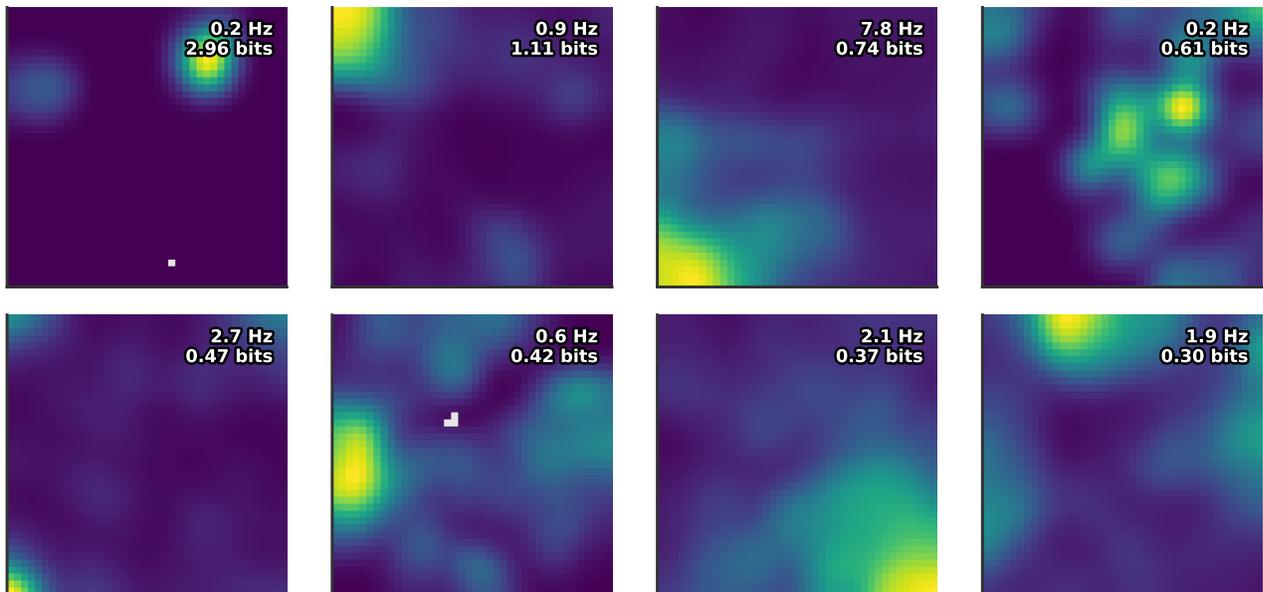


Figure 9: **Representative raw spatial firing rate maps.** (A) Eight highly informative excitatory neurons from the food-caching chickadee. The rate maps demonstrate the highly heterogeneous, multi-field “crystalline” tuning structure characteristic of the caching hippocampus. Unvisited spatial bins are masked in light grey. Peak firing rate (Hz) and spatial information (bits/spike) are inset for each cell. (B) Eight highly informative excitatory neurons from the non-caching finch. Even among the most spatially selective cells in the finch population, spatial tuning is characterised by diffuse, single-field or un-tuned “mist-like” firing that lacks the sharp topographic boundaries required for a high-capacity rigid manifold.

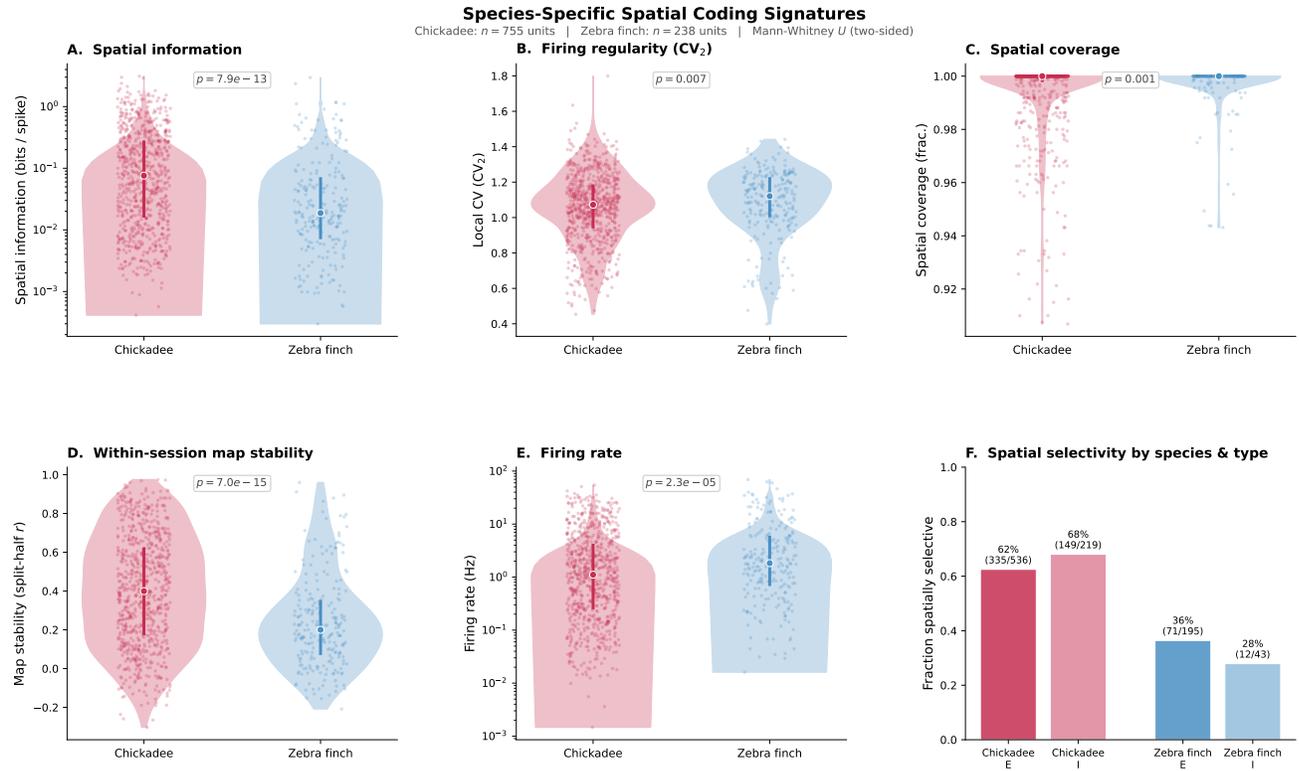


Figure 10: Single-unit correlates of population-level geometric stability in food-caching versus non-caching birds. *In vivo* hippocampal recordings from the food-caching black-capped chickadee (*titmouse*, red; $n = 755$ units) and the non-caching zebra finch (blue; $n = 238$ units) reveal single-unit signatures consistent with the population-level crystalline code identified by Shesha analysis (Figure 1). **(A)** Spatial information (bits/spike) is substantially elevated in chickadees, consistent with higher per-neuron coding fidelity in the high-capacity regime. **(B)** Firing regularity (CV_2) is modestly but significantly elevated in chickadees ($p = 0.007$), reflecting increased temporal precision of place-selective spiking. **(C)** Spatial coverage is near-ceiling for both species, reflecting the open-field recording design rather than a species-specific biological difference. **(D)** Within-session map stability, quantified via split-half spatial map correlation, is markedly higher in chickadees ($p = 7.0 \times 10^{-15}$) — a single-unit analogue of the Shesha effect, confirming that geometric stability is manifest at the level of individual place fields. **(E)** Mean firing rate distributions; chickadees exhibit higher rates on average, consistent with excitatory-dominant drive identified in the circuit analysis (Figure 2). **(F)** Fraction of spatially selective units by species and putative cell type. The excitatory–inhibitory selectivity asymmetry is larger in chickadees (E: 62%, I: 68%) than in zebra finches (E: 36%, I: 28%), consistent with excitatory neurons serving as the primary structural nodes of the crystalline code. For violin plots **(A–E)**, internal dots represent the median, thick vertical lines span the interquartile range, and shaded regions denote kernel density estimates. All comparisons are two-sided Mann-Whitney U tests.