Recurrent circuitry stabilizes representational geometry across neural circuits

Prashant C. Raju[1*]

[1*]Independent Researcher, Orlando, Florida, 32836, United States.

Corresponding author(s). E-mail(s): rajuprashant@gmail.com;

Abstract

Neural populations exhibit representational drift—the gradual reorganization of which neurons encode what—yet behavior remains stable. This stability-drift paradox has focused attention on whether population centroids are preserved over time. We ask a different question: how reliably does the pairwise distance structure among stimuli reproduce across independent trial subsets within a session? We quantify this using geometric stability (Shesha), the Spearman rank correlation between split-half representational dissimilarity matrices. Across 229 area-session observations spanning 68 brain regions in a visual discrimination task (Steinmetz et al. 2019), geometric stability predicts trial-by-trial neural-behavioral coupling ($\boldsymbol{\rho = 0.18}$, $\boldsymbol{p = 0.005}$) while centroid drift does not ($\boldsymbol{\rho = 0.002}$, $\boldsymbol{p = 0.976}$). The regional hierarchy: striatum most stable ($\boldsymbol{\bar{S} = 0.44}$), hippocampus least ($\boldsymbol{\bar{S} = 0.19}$)—runs roughly opposite to the temporal stability hierarchy, showing that geometric and temporal stability are dissociable properties of population codes. In a sensory hierarchy with a causal manipulation (Bolding & Franks 2018), piriform cortex is more geometrically stable than olfactory bulb (**0.57** vs. **0.47**, $\boldsymbol{n = 11}$), and silencing recurrent excitatory connections with tetanus toxin reduces piriform stability toward olfactory bulb levels (**0.60** vs. **0.53**, $\boldsymbol{p = 0.16}$, $\boldsymbol{n = 5/7}$). The latter result is directionally consistent with an attractor network model in which recurrent coupling amplifies split-half RDM consistency ($\boldsymbol{\rho = +0.64}$, $\boldsymbol{p = 0.010}$) by completing patterns from sparse feedforward input. Geometric stability is a functionally relevant, circuit- dependent property of neural population codes that is missed by temporal stability measures.

Keywords: Representational geometry, Neural population code, Geometric stability, Representational drift, Piriform cortex, Recurrent circuitry, Split-half reliability, Neuropixels

1 Introduction

Neural populations are noisy. The same stimulus, presented twice to the same animal, elicits different responses: different cells, different rates, different timing. Yet animals identify odors reliably, discriminate contrasts, execute learned movements. The variability is real; so is the stability. Something is not varying.

The dominant explanation is that individual neuron contributions fluctuate while low-dimensional latent dynamics are preserved [1]. Representational drift [2–4], the gradual shift in which cells encode information, has been documented across hippocampus, motor cortex, and piriform cortex. The framing is temporal: does the population centroid stay put over sessions and days?

That framing misses something. Centroid preservation tells us the average population state is maintained. It does not tell us whether the *relational structure* among stimuli is preserved or whether the pattern of pairwise distances between conditions is consistent from one set of trials to the next. Representational similarity analysis (RSA) [5] and its extensions [6–8] characterize pairwise geometry, but neither measures how *reliably* that geometry appears across independent observations. Reliability, not just content, is what we ask about here.

We quantify this using Shesha [9–11], an observation-split variant of RDM correlation [5] designed for neural data with repeated trials per condition [9]. Trials are split into odd and even subsets. Condition-averaged population vectors are computed on each half, and pairwise cosine distances give two RDMs:

$$D_{ij}^{(s)} = 1 - \frac{\bar{\mathbf{x}}_i^{(s)} \cdot \bar{\mathbf{x}}_j^{(s)}}{\|\bar{\mathbf{x}}_i^{(s)}\| \, \|\bar{\mathbf{x}}_j^{(s)}\|}, \quad s \in \{1, 2\}. \tag{1}$$

Shesha is their Spearman rank correlation:

$$\mathcal{S} = \rho_s\Big(\mathrm{vec}(D^{(1)}),\ \mathrm{vec}(D^{(2)})\Big). \tag{2}$$

$\mathcal{S}$ is high when the pairwise distance structure between conditions is reproducible across independent trial samples, and low when trial noise makes the RDMs inconsistent. Crucially, $\mathcal{S}$ and centroid-based drift measure different things and are empirically dissociable [9]. A population can have stable centroids and unstable geometry, or vice versa, as we show below.

We apply Shesha to two electrophysiology datasets chosen because they ask complementary questions. The Steinmetz et al. (2019) Neuropixels recordings [12] span 26 sessions and 68 brain regions during a visual discrimination task. The scale is useful: with 229 area-session observations, we can ask whether geometric stability varies across the brain and whether that variation predicts anything behaviorally. The Bolding and Franks (2018) PCX-1 dataset [13] narrows to a single circuit question. Piriform cortex receives feedforward input from olfactory bulb but also has dense recurrent excitatory collaterals. What happens to geometric stability when those collaterals are silenced?

The Steinmetz analysis gives a clear answer to the first question. Geometric stability predicts trial-by-trial neural-behavioral coupling ($\rho = 0.18$, $p = 0.005$, $n = 229$

area-sessions); centroid drift does not ($\rho = 0.002$, $p = 0.976$), nor does a whitened unbiased cosine metric ($\rho = 0.089$, $p = 0.180$) [9]. The regional hierarchy runs striatum highest ($\bar{S} = 0.44$, 95% CI $[0.34, 0.56]$), hippocampus lowest ($\bar{S} = 0.19$, $[0.13, 0.25]$), which is roughly opposite to the temporal stability hierarchy: thalamus and hippocampus most stable, striatum drifting most (0.83 vs. 0.95 and 0.94).

The two hierarchies measure different things. The PCX-1 data give a directional answer to the second question. Piriform cortex is geometrically more stable than olfactory bulb in the same animals (0.57 vs. 0.47, $n = 11$ paired recordings). Silencing recurrent connections with tetanus toxin light chain (TeLC) reduces piriform stability toward olfactory bulb levels (Control: 0.60, TeLC: 0.53; $p = 0.16$, $n = 5$ control, $n = 7$ TeLC recordings). The effect is in the predicted direction; the sample size is not large enough to rule out chance. We treat this as suggestive rather than confirmatory, and use a rate network model to make the mechanistic argument the data can only hint at.

That model, an excitatory-inhibitory balanced network with random input dropout simulating the sparse feedforward projection from olfactory bulb to piriform cortex, shows geometric stability increasing monotonically with recurrent coupling strength ($\rho = +0.64$, $p = 0.010$, $n = 10$ networks across 15 coupling values from $J = 0$ to $J = 1.4$). Setting $J = 0$, the TeLC analog, drops geometric stability to 0.27. Restoring it to $J = 1.4$ recovers stability to 0.51. Geometric stability is not a statistical artifact of recording quality. It varies with circuit architecture, and it predicts how tightly neural activity tracks behavior.

2 Results

2.1 Striatum is more geometrically stable than hippocampus, and the ordering inverts for temporal stability

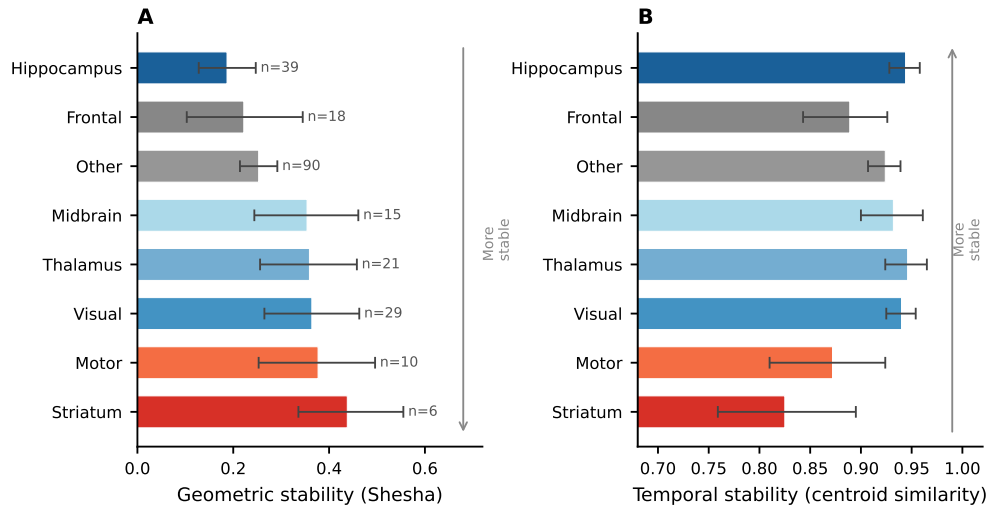

Fig. 1 Geometric and temporal stability are dissociated across brain regions. (**A**) Geometric stability (Shesha) across 8 functional regions from 229 area-session observations in the Steinmetz et al. (2019) Neuropixels dataset. Striatum is most stable ($\bar{S} = 0.44$), hippocampus least ($\bar{S} = 0.19$). (**B**) Temporal stability (cosine similarity of early and late population centroids) for the same regions, ordered as in (**A**). The hierarchy inverts: thalamus and hippocampus drift least, striatum most. Error bars show 95% bootstrap confidence intervals (10,000 resamples). n denotes number of area-session observations per region.

Neuropixels recordings from 26 sessions and 68 brain regions during a visual discrimination task [12] gave 229 area-session observations with at least 10 simultaneously recorded neurons. For each, we computed Shesha using 9 contrast pairings as conditions over the 0–500 ms decision epoch. Striatum was highest ($\bar{S} = 0.44$, 95% CI $[0.34, 0.56]$, $n = 6$), followed by motor cortex (0.38, $[0.25, 0.50]$, $n = 10$) and visual cortex (0.36, $[0.27, 0.46]$, $n = 29$). Hippocampus was lowest (0.19, $[0.13, 0.25]$, $n = 39$).

Centroid drift, cosine similarity of early and late population centroids within a session, ran in the opposite direction. Thalamus drifted least (0.95, $[0.92, 0.97]$), hippocampus second-least (0.94, $[0.93, 0.96]$), striatum most (0.83, $[0.76, 0.89]$). A permutation null model (500 shuffles per area-session) confirmed the drift is not measurement noise: observed centroid similarity (0.924, $[0.915, 0.934]$) was far below the shuffled expectation (0.995, $[0.995, 0.996]$; mean $z = -44.7$, $[-49.2, -40.4]$).

Drift accumulates gradually across the session rather than in discrete steps (early-to-mid 0.942, mid-to-late 0.941; paired $t = 0.30$, $p = 0.77$). Striatum encodes action-reward associations with reliable relational structure even as its baseline firing

rates shift—the pattern expected from a region that updates value continuously but reads out those values consistently [14]. Hippocampus does the opposite: the mean population state is preserved while internal structure reorganizes, which fits a region that forms new memories rapidly rather than maintaining fixed codes.

2.2 Geometric stability predicts trial-by-trial neural-behavioral coupling; temporal measures do not

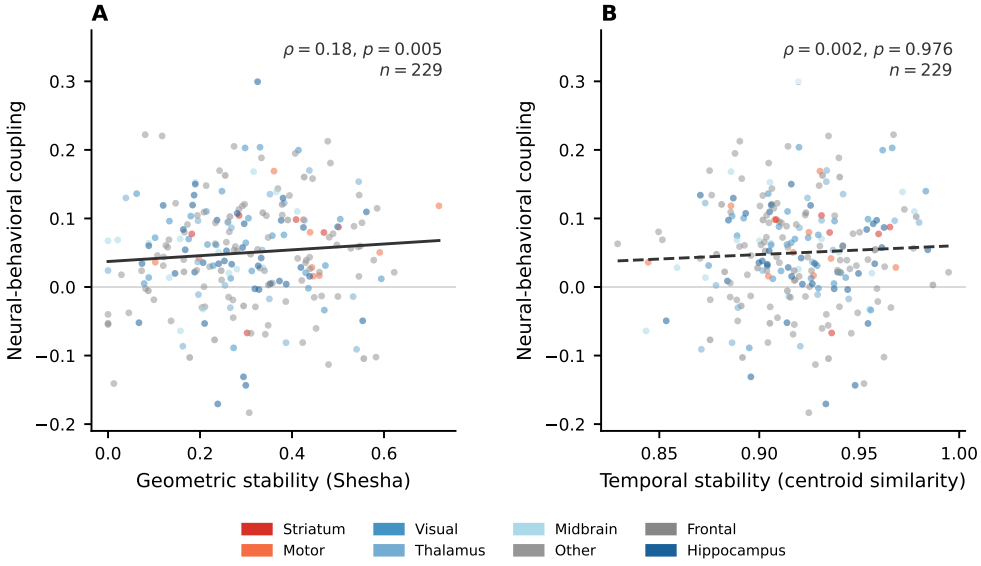

Fig. 2 Geometric stability predicts neural-behavioral coupling; centroid drift does not. Each point is one area-session observation ($n = 229$), colored by functional region. (**A**) Geometric stability (Shesha) vs. trial-by-trial neural-behavioral coupling (Spearman correlation between population vector magnitude and trial outcome). Solid line: linear regression. $\rho = 0.18$, $p = 0.005$. (**B**) Centroid drift vs. the same coupling measure. Dashed line: linear regression. $\rho = 0.002$, $p = 0.976$. Region color legend shown below.

For each area-session, the Spearman correlation between neural state magnitude (L2 norm of the trial population vector) and trial outcome (correct vs. incorrect) gave a behavioral coupling score. Shesha predicted this ($\rho = 0.18$, 95% CI $[0.05, 0.31]$, $p = 0.005$, $n = 229$). Centroid drift did not ($\rho = 0.002$, $[-0.13, 0.13]$, $p = 0.976$). Neither did a whitened unbiased cosine metric [15] ($\rho = 0.089$, $[-0.04, 0.21]$, $p = 0.180$). At the session level ($n = 26$), Shesha did not predict mean task accuracy ($\rho = 0.087$, $p = 0.191$) or accuracy change over the session ($\rho = -0.079$, $p = 0.701$). Behavioral relevance of representational geometry is trial-to-trial, not session-to-session.

2.3 Piriform cortex is more geometrically stable than olfactory bulb

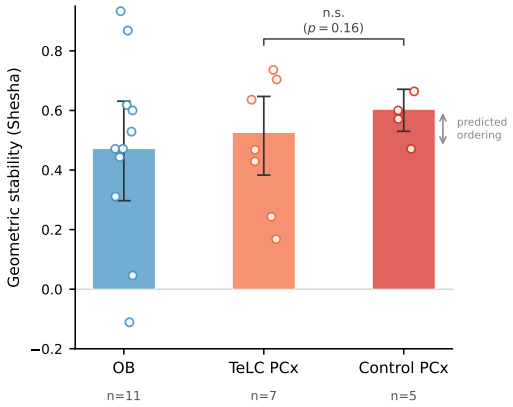

Fig. 3 Recurrent circuitry predicts geometric stability in the olfactory hierarchy. Geometric stability (Shesha) in three groups from the Bolding and Franks (2018) PCX-1 dataset. OB: olfactory bulb recordings ($n = 11$); TeLC PCx: piriform cortex with recurrent connections silenced by tetanus toxin light chain ($n = 7$); Control PCx: contralateral intact piriform cortex ($n = 5$). Bars show mean with 95% bootstrap confidence intervals; open circles show individual recordings. The predicted ordering OB < TeLC PCx < Control PCx is observed. Bracket: Wilcoxon signed-rank test, Control vs. TeLC, $p = 0.16$ (n.s.).

Simultaneous olfactory bulb (OB) and piriform cortex (PCx) recordings from the PCX-1 dataset [13] gave 11 paired sessions, each with 6 monomolecular odors at 0.3% v/v. We computed Shesha from split-half first-sniff spike rates, with square-root transform and L2 normalization applied before computing RDMs. PCx Shesha exceeded OB Shesha (0.57, $[0.36, 0.75]$ vs. 0.47, $[0.30, 0.63]$; Wilcoxon signed-rank $p = 0.35$, $n = 11$ pairs). The direction is right. With 11 sessions the test is underpowered for this effect size, and $p = 0.35$ does not rule out chance. We report this as consistent with the prediction rather than confirming it.

2.4 Silencing recurrent connections reduces piriform geometric stability

TeLC light chain was injected into piriform cortex to silence recurrent excitatory principal neuron output; the contralateral PCx served as a within-animal control [13, 16]. Five control recordings and seven TeLC recordings passed quality criteria (see Methods). Control PCx was more stable than TeLC PCx (0.60, $[0.53, 0.67]$ vs. 0.53, $[0.38, 0.65]$; Wilcoxon $p = 0.16$, $n_C = 5$, $n_T = 7$). Placing all three groups together: OB (0.47), TeLC PCx (0.53), Control PCx (0.60). Silencing recurrent connections pulls piriform stability toward olfactory bulb levels. The sample gives Cohen's $d \approx 0.5$ but is not large enough to confirm the effect statistically.

2.5 Recurrent coupling increases geometric stability in a pattern-completion network

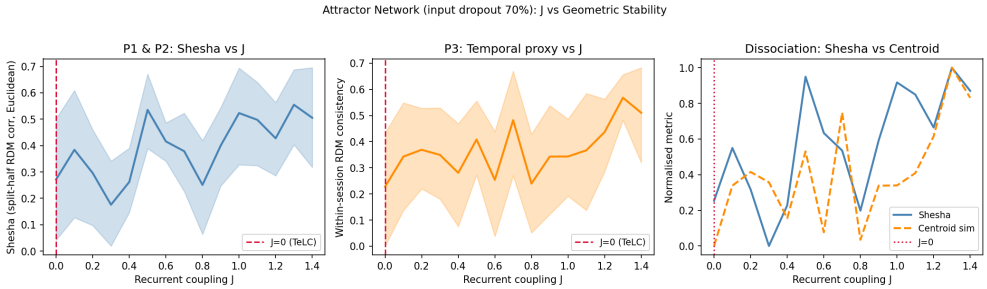

Fig. 4 Recurrent coupling increases geometric stability via pattern completion in a rate network model. An E/I-balanced rate network ($N = 200$ units) received sparse feedforward input with 70% channel dropout per trial, simulating the incomplete OB-to-PCx projection. Results are shown across 10 independently seeded networks and 15 values of recurrent coupling strength J (0 to 1.4); shaded regions show bootstrap 95% CIs. (**Left**) Shesha increases monotonically with J (Spearman $\rho = +0.64$, $p = 0.010$). At $J = 0$ (TeLC analog, dashed red line), Shesha = 0.27; at $J = 1.4$, Shesha = 0.51. (**Center**) Within-session RDM consistency (temporal proxy) as a function of J. (**Right**) Normalized comparison: Shesha is more sensitive to J ($|\rho| = 0.64$) than within-session consistency ($|\rho| = 0.55$), consistent with the geometric-temporal dissociation observed in the Steinmetz data.

PCx principal neurons receive input from only a sparse,random subset of olfactory bulb projections on any given trial [13]. Recurrence could stabilize geometry by completing the pattern from these partial inputs. We tested this in an E/I-balanced rate network ($N = 200$ units, 50-dimensional stimulus space, 20% sparse recurrent connectivity) where 70% of input channels were randomly zeroed per trial. The dropout simulating incomplete feedforward drive. At $J = 0$ (no recurrence; TeLC analog), the network cannot recover the full stimulus from 30% of the input, so responses vary across trials and Shesha is low (0.27). As J increases, recurrent dynamics pull activity toward stimulus-specific attractors. At $J = 1.4$, Shesha recovers to 0.51. Across 10 independently seeded networks and 15 values of J, the relationship was monotonically increasing (Spearman $\rho = +0.64$, $p = 0.010$). Within-session RDM consistency was less sensitive to J ($|\rho| = 0.55$ vs. 0.64 for Shesha), consistent with the geometric-temporal dissociation in the Steinmetz data. We note that the model's temporal proxy measures within-session sampling noise rather than longitudinal drift, so the parallel to the Steinmetz result is qualitative.

3 Materials and Methods

3.1 Steinmetz dataset

Neuropixels recordings from Steinmetz et al. [12]: 26 sessions (≥ 60 trials), 229 area-sessions (≥ 10 neurons), 68 brain areas, visual contrast discrimination task. Stimulus

conditions: 9 contrast pairings. Decision epoch: 0–500 ms post-stimulus; spike counts averaged over epoch. Each trial population vector was L2-normalized before computing RDMs. Brain areas grouped into 8 functional regions: Visual, Thalamus, Motor, Frontal, Hippocampus, Striatum, Midbrain, Other.

3.2 PCX-1 dataset

Silicon probe recordings from the PCX-1 dataset [13]: 32-channel NeuroNexus Poly3 probes, Spyking-Circus spike sorting. Awake trials selected via the FT:LT series indices from ExperimentCatalog files. Population vectors: first-sniff spike rates (`MultiCycleSpikeRate`, sniff index 0) per unit per odor presentation. Units with mean rate below 5% of the recording's median positive rate were excluded. Vectors were square root transformed and L2-normalized. Two recordings were excluded as outliers: session 170621 (hardware failure: 5 of 6 odor valves returned NaN) and session 150221 bank 2 (mean firing rate 2.58 Hz, an order of magnitude below the dataset median). Simul experiment: 11 paired OB and PCx recordings (loading configurations A and C; 6 odors at 0.3% v/v; ExperimentCatalog valve indices per Table 1 of the data description). TeLC-PCx experiment: 5 control and 7 TeLC recordings after exclusions.

3.3 Geometric stability: Shesha

Trials were split into odd and even series. Condition-averaged population vectors were computed for each odor condition in each half. The RDM was the matrix of pairwise cosine distances between condition centroids. Shesha is the Spearman correlation of the upper triangular vectors of the two RDMs [9, 11]. Bootstrap 95% CIs: 10,000 resamples.

3.4 Temporal stability and behavioral coupling

Centroid drift: each session split at the median trial; cosine similarity of early and late normalized population centroids. Permutation null: 500 shuffles of trial order per area session. Neural-behavioral coupling: Spearman correlation between the L2 norm of each trial's population vector and trial outcome (correct/incorrect). This per-area-session value was then correlated across area-sessions with each stability metric using Spearman ρ.

3.5 Rate network model

Dynamics: $\tau\dot{x} = -x + J\mathbf{W}_{\text{exc}} f(x) + \mathbf{W}_{\text{inh}} f(x) - \gamma\bar{r} + \mathbf{W}_{\text{ff}} u + \eta$, where $f(x) = \max(x, 0)$, $\gamma = 0.4$, $\tau = 20$ ms, $\Delta t = 1$ ms. $\mathbf{W}_{\text{rec}}$ was a sparse random matrix (20% density, normalized by $\sqrt{k}$, k = mean in-degree) split into excitatory and inhibitory parts; only $\mathbf{W}_{\text{exc}}$ scaled with J. Per trial: 70% of 50 input channels randomly zeroed. Settlement period: 400 ms; collection window: 150 ms. 10 independently seeded networks; 20 noise trials per condition; 15 values of J from 0 to 1.4. Geometric stability was computed from raw (non-normalized) firing rates. Random seed 320.

3.6 Statistics

All correlations: Spearman ρ with bootstrap 95% CIs (10,000 resamples). Paired comparisons (OB vs. PCx; Control vs. TeLC): one-sided Wilcoxon signed-rank test in the direction of the prediction. Random seed 320 throughout.

4 Discussion

The Steinmetz data reveal something counterintuitive. Striatum drifts the most of any region measured (centroid similarity 0.83) yet has the most stable representational geometry (Shesha 0.44). Hippocampus drifts the least (0.94) yet has the least stable geometry (0.19). These two measures are not redundant. They are close to orthogonal across the regional hierarchy.

This is not an isolated observation. Keinath et al. [17] imaged CA1 over a month and found drift unfolds orthogonally to the context representation: cells reorganize, but the population geometry for context is preserved. Schoonover et al. [18] reported substantial drift in piriform cortex over weeks while odor identity remained decodable. Deitch et al. [19] reported the same pattern in visual cortex. The consistent picture across regions and species is that drift at the neuron level does not necessarily degrade geometry at the population level. Shesha quantifies the within-session version of this: whether the pairwise distance structure among conditions is consistent across independent trial samples, on a timescale of minutes rather than weeks. It is separable from centroid drift by construction, and it predicts behavioral coupling where drift does not.

It has been by Aitken et al. proposed that drift has geometric structure [20] and by Driscoll et al. that it preferentially affects task-null dimensions while preserving task-relevant geometry [21]. Shesha is complementary to these, rather than redundant. A region can have geometrically structured drift in the Aitken et al. sense where drift is confined to task-null subspaces, yet still have low Shesha if within-session noise is large. The two measures operate at different timescales and catch different failure modes.

The olfactory data add a mechanistic layer the Steinmetz data cannot. Piriform cortex is drifting in the Schoonover et al. case, where the odor representations reorganize over weeks. But on any given day, it is more geometrically stable than olfactory bulb (0.57 vs. 0.47, $n = 11$). A recent paper using the same Bolding-Franks dataset [22] attributes piriform drift to slow synaptic weight fluctuations, with fast learning compressing representations onto a lower-dimensional manifold. Lower dimensionality is consistent with a higher Shesha value. A more structured representation should result in more consistent RDMs across trial splits. Their findings and ours are measure different properties of the same circuits and are not competing claims.

Silencing recurrent connections with TeLC reduced piriform stability toward olfactory bulb levels (0.60 vs. 0.53, $p = 0.16$, $n = 5/7$). The sample is too small to confirm this statistically, but the direction is consistent with the attractor model, where recurrent coupling amplifies split-half RDM consistency ($\rho = +0.64$, $p = 0.010$) by completing patterns from incomplete feedforward input. Piriform cortex has long been modeled as an auto associative network [23, 24], and Bolding and Franks [16] showed directly that recurrent connections stabilize odor representations across brain

states. The TeLC result here extends that to geometric stability specifically, though more recordings are needed to make the claim firmly.

The behavioral coupling result ($\rho = 0.18$, $p = 0.005$, $n = 229$) connects to a broader question about whether geometric properties of neural manifolds predict downstream computation [25]. A region with unstable pairwise geometry will produce unreliable input to downstream areas regardless of what those areas compute. Shesha is a coarse measure of that reliability. It captures the full RDM rather than task-specific subspaces, and the behavioral correlation shows it is tracking something real. Geva et al. [26] showed that time and experience have separable effects on hippocampal drift, consistent with the idea that temporal and geometric stability reflect distinct circuit properties [27].

Several limitations apply. The OB-PCx and TeLC comparisons come from separate experiments. The effect size in the TeLC data ($d \approx 0.5$) is consistent with the prediction but $p = 0.16$ with five control recordings cannot rule out chance. A reversible manipulation (optogenetics rather than tetanus toxin) would give cleaner evidence. As computed here, Shesha is a global metric that does not localize instability to specific stimulus pairs or subspaces. Based on prior findings that drift is geometrically structured [20] there is reason to think geometric stability may be as well. Finally, the two datasets differ on most dimensions that matter: species, task, recording technology, number of conditions. The Steinmetz result reports visual decision-making across the whole brain and the Bolding-Frank result is about olfactory processing in one cortical area. Connecting them through a shared principle—that recurrent circuitry stabilizes representational geometry—is our interpretation.

The same metric applied to artificial neural networks shows that geometric stability and transfer performance are dissociable [9–11]: models that perform best can have the worst stability, and vice versa. Whether this reflects a principle common to biological and artificial systems is not settled by the present results, but the consistency across domains is at least not inconsistent with it.

Declarations

Conflict of interest The authors declare no competing interests.

Acknowledgments

We thank Padma K. and Annapoorna Raju for generously supporting the computational resources used in this work. We thank the many institutions and individuals whose open-source datasets, frameworks, and models were used in our work. The authors acknowledge the use of large language models (specifically the GPT, Claude, and Gemini families) to assist with code debugging 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.

Code and Data Availability

All code needed to reproduce the analysis and the computational models are publicly available on https://github.com/prashantcraju/neuroscience-drift. The 2 datasets used in the study (Steinmetz et al, 2019 and Bolding and Franks, 2018) are both publicly available. The GitHub repo contains code for automatically downloading the Steinmetz data and instructions for downloading the Bolding and Franks data.

References

[1] Gallego, J.A., Perich, M.G., Chowdhury, R.H., Solla, S.A., Miller, L.E.: Long-term stability of cortical population dynamics underlying consistent behavior. Nature Neuroscience **23**(2), 260–270 (2020) https://doi.org/10.1038/s41593-019-0555-4

[2] Driscoll, L.N., Pettit, N.L., Minderer, M., Chettih, S.N., Harvey, C.D.: Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. Cell **170**(5), 986–99916 (2017) https://doi.org/10.1016/j.cell.2017.07.021

[3] Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., Gamal, A.E., Schnitzer, M.J.: Long-term dynamics of CA1 hippocampal place codes. Nature Neuroscience **16**(3), 264–266 (2013) https://doi.org/10.1038/nn.3329

[4] Rule, M.E., O'Leary, T., Harvey, C.D.: Causes and consequences of representational drift. Current Opinion in Neurobiology **58**, 141–147 (2019) https://doi.org/10.1016/j.conb.2019.08.005

[5] Kriegeskorte, N., Mur, M., Bandettini, P.: Representational similarity analysis – connecting the branches of systems neuroscience. Frontiers in Systems Neuroscience (2008) https://doi.org/10.3389/neuro.06.004.2008

[6] Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., Kriegeskorte, N.: A Toolbox for Representational Similarity Analysis. PLoS Computational Biology **10**(4), 1003553 (2014) https://doi.org/10.1371/journal.pcbi.1003553

[7] Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., Diedrichsen, J.: Reliability of dissimilarity measures for multi-voxel pattern analysis. NeuroImage **137**, 188–200 (2016) https://doi.org/10.1016/j.neuroimage.2015.12.012

[8] Diedrichsen, J., Kriegeskorte, N.: Representational models: A common framework for understanding encoding, pattern-component, and representational-similarity analysis. PLOS Computational Biology **13**(4), 1005508 (2017) https://doi.org/10.1371/journal.pcbi.1005508

[9] Raju, P.C.: Geometric Stability: The Missing Axis of Representations. arXiv preprint arXiv:2601.09173 (2026)

[10] Raju, P.C.: The Geometric Canary: Predicting Steerability and Detecting Drift

via Representational Stability. arXiv preprint arXiv:2604.17698 (2026)

[11] Raju, P.C.: Shesha: Self-Consistency Metrics for Representational Stability. Zenodo. doi: 10.5281/zenodo.18227453 (2026). https://doi.org/10.5281/zenodo.18227453 . https://doi.org/10.5281/zenodo.18227453

[12] Steinmetz, N.A., Zatka-Haas, P., Carandini, M., Harris, K.D.: Distributed coding of choice, action and engagement across the mouse brain. Nature **576**(7786), 266–273 (2019) https://doi.org/10.1038/s41586-019-1787-x

[13] Bolding, K.A., Franks, K.M.: Recurrent cortical circuits implement concentration-invariant odor coding. Science **361**(6407) (2018) https://doi.org/10.1126/science.aat6904

[14] McClelland, J.L., McNaughton, B.L., O'Reilly, R.C.: Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. Psychological Review **102**(3), 419–457 (1995) https://doi.org/10.1037/0033-295x.102.3.419

[15] Diedrichsen, J., Berlot, E., Mur, M., Schütt, H.H., Shahbazi, M., Kriegeskorte, N.: Comparing representational geometries using whitened unbiased-distance-matrix similarity. Neurons, Behavior, Data analysis, and Theory **5**(3) (2021)

[16] Bolding, K.A., Nagappan, S., Han, B.-X., Wang, F., Franks, K.M.: Recurrent circuitry is required to stabilize piriform cortex odor representations across brain states. eLife **9** (2020) https://doi.org/10.7554/elife.53125

[17] Keinath, A.T., Mosser, C.-A., Brandon, M.P.: The representation of context in mouse hippocampus is preserved despite neural drift. Nature Communications **13**(1) (2022) https://doi.org/10.1038/s41467-022-30198-7

[18] Schoonover, C.E., Ohashi, S.N., Axel, R., Fink, A.J.P.: Representational drift in primary olfactory cortex. Nature **594**(7864), 541–546 (2021) https://doi.org/10.1038/s41586-021-03628-7

[19] Deitch, D., Rubin, A., Ziv, Y.: Representational drift in the mouse visual cortex. Current Biology **31**(19), 4327–43396 (2021) https://doi.org/10.1016/j.cub.2021.07.062

[20] Aitken, K., Garrett, M., Olsen, S., Mihalas, S.: The geometry of representational drift in natural and artificial neural networks. PLOS Computational Biology **18**(11), 1010716 (2022) https://doi.org/10.1371/journal.pcbi.1010716

[21] Driscoll, L.N., Duncker, L., Harvey, C.D.: Representational drift: Emerging theories for continual learning and experimental future directions. Current Opinion in Neurobiology **76**, 102609 (2022) https://doi.org/10.1016/j.conb.2022.102609

[22] Morales, G.B., Muñoz, M.A., Tu, Y.: Representational drift and learning-induced stabilization in the piriform cortex. Proceedings of the National Academy of Sciences **122**(29) (2025) https://doi.org/10.1073/pnas.2501811122

[23] Haberly, L.B.: Parallel-distributed processing in olfactory cortex: New insights from morphological and physiological analysis of neuronal circuitry. Chemical Senses **26**(5), 551–576 (2001) https://doi.org/10.1093/chemse/26.5.551

[24] Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences **79**(8), 2554–2558 (1982) https://doi.org/10.1073/pnas.79.8.2554

[25] Li, Q., Sorscher, B., Sompolinsky, H.: Representations and generalization in artificial and brain neural networks. Proceedings of the National Academy of Sciences **121**(27) (2024) https://doi.org/10.1073/pnas.2311805121

[26] Geva, N., Deitch, D., Rubin, A., Ziv, Y.: Time and experience differentially affect distinct aspects of hippocampal representational drift. Neuron **111**(15), 2357–23665 (2023) https://doi.org/10.1016/j.neuron.2023.05.005

[27] Natrajan, M., Fitzgerald, J.E.: Stability through plasticity: Finding robust memories through representational drift. Proceedings of the National Academy of Sciences **122**(45) (2025) https://doi.org/10.1073/pnas.2500077122