

How the olfactory bulb maintains stable odor manifolds amid adaptation and representational drift

CiCi Xingyu Zheng^{1,2}, Bin Yu³, Saket Navlakha¹, and Alexei Koulakov^{1}*

¹ Cold Spring Harbor Laboratory, Cold Spring Harbor, NY

² Allen Institute, Seattle, WA

³ California Institute of Technology, Pasadena, CA

Abstract

Adaptive coding in sensory circuits enables stable perception while accommodating experience-dependent changes. In the olfactory bulb (OB), repeated odor exposure reshapes population activity even without explicit behavioral feedback, but the underlying circuit mechanisms remain unclear. By analyzing longitudinal two-photon calcium imaging datasets from the mouse OB, we identified three concurrent forms of representational change: gain adaptation, similarity-dependent pattern separation or convergence, and a rotation of encoding subspace resulting in the representational drift. Using a computational model of the mitral cell-granule cell circuit, we showed that Hebbian plasticity and structural connectivity constraints are sufficient to reproduce these transformations. Despite global representational drift, the relative geometry of odor response vectors remained stable, preserving a low-dimensional odor manifold. Together, our results reveal how local plasticity and network structure jointly enable both stability and flexibility in early sensory coding.

* Correspondence: koulakov@cshl.edu

Introduction

Adaptive coding allows sensory systems to extract behaviorally relevant information from continuously changing environments. In olfaction, adaptation occurs across multiple timescales from rapid habituation within a few breaths (Chaudhury et al., 2010; Kim et al., 2020; Shen et al., 2020) to long-term adjustments spanning days or even seasons (Hamdani et al., 2008; Teşileanu et al., 2019). Adaptation can decorrelate sensory representations, facilitating flexible behaviors such as distinguishing food sources from predators or compensating for persistent background odors (Dalton, 2000). Adaptive coding is not a single process: it can take distinct forms across multiple levels, from receptor and synaptic plasticity to circuit-level reorganizations. Understanding how odor representations evolve over days, and what computational consequences these changes entail, is therefore essential for linking sensory experience to long-term learning.

The mammalian olfactory system transforms chemical stimuli from odorants into structured neural representations through a series of processing stages. Odorants are detected by olfactory sensory neurons (OSNs) in the nasal epithelium, each expressing one odorant receptor (OR) gene out of hundreds of possibilities (Buck & Axel, 1991; Chess et al., 1994; Malnic et al., 1999). OSNs expressing the same receptor send their axons to a small number of glomeruli located in the glomerular layer of the olfactory bulb (OB) (Mombaerts et al., 1996). The glomeruli form a dense two-dimensional sheet-like structure across the OB surface, creating a stereotypical spatial map of inputs from the olfactory epithelium (Bozza et al., 2009; Giessel & Datta, 2014; Wang et al., 1998). Within each glomerulus, OSNs expressing a specific OR gene type form excitatory synapses onto the principal OB projection neurons, mitral and tufted cells (MCs and TCs), as well as onto local interneurons (Mombaerts, 2006).

Before transmitting information to cortical targets such as the piriform cortex, MCs and TCs interact with inhibitory interneurons in the deeper granule cell (GC) layer of the OB (Egger & Urban, 2006). GCs form reciprocal dendrodendritic synapses with MC/TCs, where excitation and inhibition occur at the same contact site (Aghvami et al., 2022; Aroniadou-Anderjaska et al., 1999; Rall et al., 1966). These reciprocal connections mediate recurrent inhibition that may dynamically regulate overall network gain (Shepherd et al., 2007; Yokoi et al., 1995). Because GCs integrate inputs from MC/TCs associated with multiple glomeruli, they implement a powerful form of lateral inhibition that can reshape the principal neurons' activity based on combinatorial odor inputs (Arevian et al., 2007; Cleland & Linstner, 2012; Koulakov & Rinberg, 2011). MCs in particular, exhibit more concentration-invariant responses than TCs and provide the dominant projection to piriform cortex, making them a good substrate for studying long-term changes in population coding (Chae et al., 2022). MCs and TCs subsequently convey odor information to the higher brain regions, including piriform cortex (PCx) and anterior olfactory nucleus (AON), for further processing (Shepherd, 2003).

Circuits within these cortical regions are known to support odor learning, associative processing, and integration with other sensory and contextual information (Boyd et al., 2012; Calu et al., 2007; Large et al., 2016; Otazu et al., 2015) via long-term changes in olfactory representations (Li et al., 2006; Wilson, 2000; Wilson et al., 2006). Numerous studies have shown that plasticity in the PCx recurrent connectivity supports perceptual and discriminative learning, as well as reward-associated memory formation (Bao et al., 2016; Cohen et al., 2008; Franks et al., 2011; Gadziola et al., 2020). At the same time, population activity in PCx has been reported to drift

over days, even during infrequent or passive odor exposure (Schoonover et al., 2021). It remains unclear, however, how much of this representational change arises within the cortex itself versus being inherited from earlier processing stages. One step upstream, the OB already exhibits time and exposure-dependent adaptations in its output patterns (Chu et al., 2016; Kato et al., 2012). Since recordings from OSNs have shown stable synaptic release properties during passive odor experience (Chu et al., 2017), these longer-timescale changes in neural responses could emerge within the OB networks. Such bulb-level plasticity could thus serve as a source of the representational drift observed in the cortex.

Despite its role as an early sensory relay, OB processes olfactory information in a dynamic and experience-dependent manner (Chu et al., 2016; Wu et al., 2020; Yu et al., 2025). During wakefulness, repeated passive exposure to the same odorants leads to a gradual, odor-specific reduction in MC responses (Chu et al., 2016; Kato et al., 2012; Yamada et al., 2017). Moreover, olfactory perceptual learning, where animals learn to discriminate initially indistinguishable odorants, has been shown to involve changes as early as the bulb, including the decorrelation of MC and TC population activity for similar odors, particularly pronounced in the MCs (Chu et al., 2016; Yamada et al., 2017). These findings indicate that the OB is not merely a stable encoder of sensory inputs but a flexible circuit capable of long-term adaptation.

Such experience-dependent decorrelation can be mediated by inhibitory plasticity that refines local circuits to selectively suppress overlapping activity patterns among MCs. In agreement with this idea, GCs enhance pattern separation during olfactory perceptual learning: optogenetic activation of GCs increases separation of MC population responses and improves odor discrimination performance, whereas GC inactivation reduces separation and impairs discrimination (Alonso et al., 2012; Li et al., 2018). In rodents, OB is also one of the few brain regions where new inhibitory neurons are continually incorporated throughout life, suggesting an additional source of long-term circuit plasticity (Lledo & Valley, 2016). However, adult-born GCs integrate gradually and may contribute primarily to later phases of adaptation (Whitman & Greer, 2007); their ablation does not abolish all forms of learning-related plasticity (Li et al., 2018; Nissant et al., 2009). Together, these findings point to the MC-GC recurrent network as a candidate for shaping odor representations over extended timescales.

Here, we investigate how odor representations by the MCs, one of the OB's principal output neurons, evolve across days of repeated passive exposure. By analyzing longitudinal two-photon calcium imaging data from mice repeatedly exposed to either similar or dissimilar odor pairs and to an eight-odor mixture panel, we quantified how population codes transform over time. We then developed a recurrent inhibitory feedback model of the MC-GC circuit to test whether local plasticity and circuit structure can account for the observed transformations. Together, our work unifies several distinct forms of representational changes – pattern separation, pattern convergence, gain adaptation, and representational drift – under a single circuit mechanism.

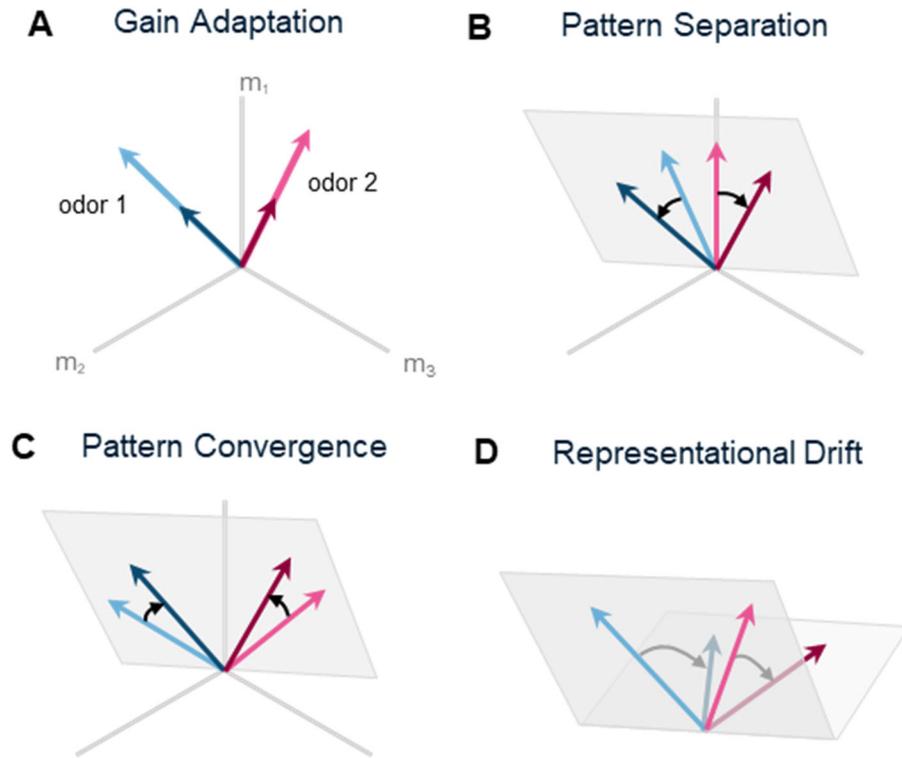


Figure 1. Schematic representation of four types of adaptation of odor representations resulting from odor exposure. (A) Gain adaptation. Repeated exposure to an odor leads to a reduced response. Light blue and pink arrows represent population responses (m_1 , m_2 , m_3 , ... are responses of the individual mitral cells) to odor 1 and odor 2 during early odor exposure; darker arrows show representations of the same odors after extended exposure. **(B) Pattern separation.** Initially similar odor representations (light arrows, small angular separation) become more distinct over time, increasing the angle between representation vectors (darker arrows). The vectors are limited to move in the same subspace (shaded). **(C) Pattern convergence.** Initially distinct odor representations become more similar over time, decreasing the angle between representation vectors, while remaining in the same subspace (shaded). **(D) Representational drift.** Odor representation subspace may rotate over time while maintaining relative relationships between population response vectors. Curved arrows indicate the drift direction.

Results

Our goal is to investigate how odor representations of MCs in the OB change under repeated stimulus exposure over several days. Before proceeding with our analyses, we define four types of long-term plasticity in odor representations: gain adaptation, pattern separation, pattern convergence, and representational drift (Fig. 1). To characterize these changes, we consider the population response vector, whose components correspond to the responses of individual MCs to an odorant. *Gain adaptation* is defined as a reduction in the length of the population response vector (Fig. 1A). *Pattern separation* can occur when pairs of odors are presented intermittently across multiple experimental sessions over several days. It corresponds to a rotation of the population response vectors away from each other while remaining within the same plane (Fig. 1B). *Pattern convergence* is the opposite effect and consists of a rotation of population responses to odor pairs toward each other (Fig. 1C). Finally, *representational drift* is defined as a rotation of the plane containing population responses to the presented odors as a result of odor exposure (Fig. 1D). While the first three phenomena have been reported in the OB (Chu et al., 2016; Kato et al., 2012), representational drift remains unexplored. Below, we combine a computational model of OB circuitry with longitudinal calcium imaging data from mice repeatedly exposed to odors (similar or dissimilar pairs, as well as an eight-odor mixture panel) over 7–9 days (Chu et al., 2016; Li et al., 2018). This integrated approach enables us to identify the circuit features underlying each type of plasticity and to test theoretical predictions against experimental observations.

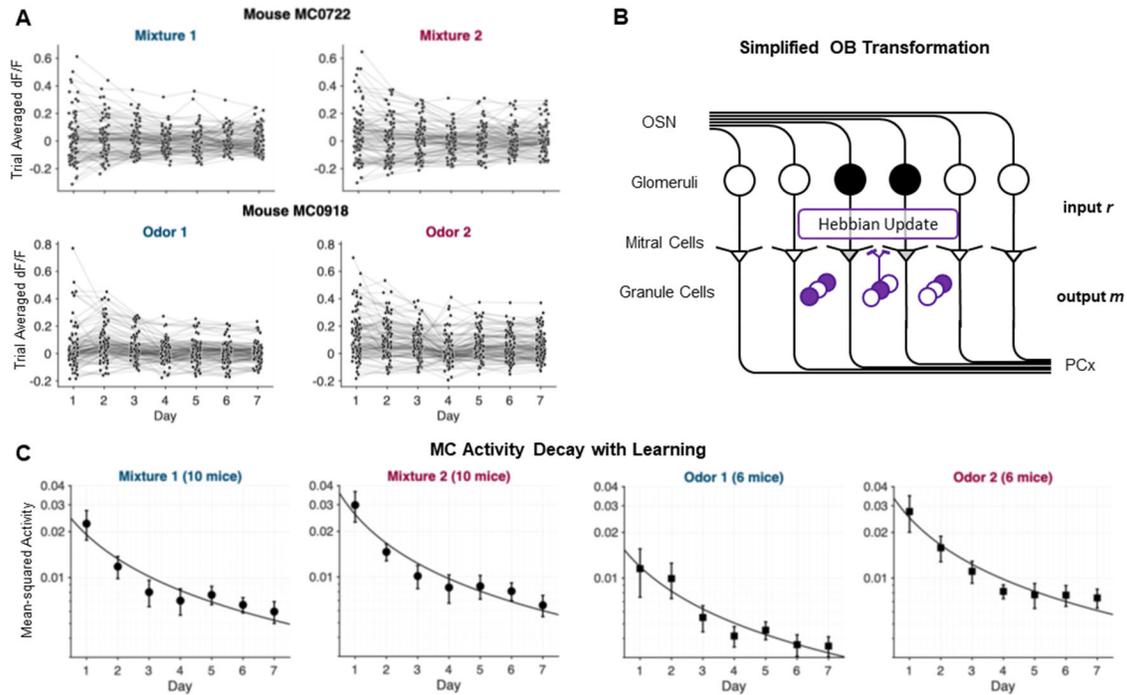


Figure 2. Gain adaptation in the MC population activity.

(A) Single-neuron activity trajectories across days of odor exposure. Each dot represents the trial-averaged activity of an individual MC. Top panels show responses from a mouse exposed to a pair of similar odor mixtures (Mixture 1: 52% Heptanal/48% Ethyl Tiglate; Mixture 2: 48% Heptanal/52% Ethyl Tiglate). Bottom panels show responses from a mouse exposed to a pair of dissimilar odors (Odor 1: 100% Heptanal; Odor 2: 100% Ethyl Tiglate).

(B) Simplified OB circuit model. MCs (triangles) receive excitatory input r from glomeruli (circles; filled = active, empty = inactive). Mitral cell population output m is sent to downstream areas including piriform cortex (PCx). GCs (purple circles, filled/empty -- active/inactive) form reciprocal dendrodendritic synapses with the MCs, providing feedback inhibition onto the MCs.

(C) Comparison of model results with experimental data. In the mean-field model (solid curve), population activity decays as described by Equation 1. In the data (dark circles), mean-squared population activity (\pm SEM) across all mice for each odor condition follow similar dynamics.

Gain adaptation reflects long term population-level activity decay.

We define long-term gain adaptation as the gradual reduction in the MC population activity over days of repeated odor exposure. This phenomenon manifests itself as decreased amplitude of MC population responses (Fig. 2A; Fig. S2A-D). We hypothesized that this long-term adaptation arises from activity-dependent plasticity in the dendrodendritic synapses between MCs and GCs (Fig. 2B). Odor-evoked MC activation drives GC responses, which in turn provide inhibitory feedback to MCs. Hebbian strengthening of these reciprocal connections progressively enhances inhibitory feedback, leading to reduced MC responses over time.

To test this hypothesis quantitatively, we developed a computational mean-field model that predicts the temporal evolution of population activity (Supplementary Information Section 2). The model considers OSN activation represented by a vector \mathbf{r} driving the population of MCs. The MCs are connected to a single GC that provides inhibitory feedback. This simplification is motivated by the sparseness of GC responses, which implies that very few GCs are responsive to any given odor. It also allows us to obtain a closed-form solution for the MC activity. The network implements the following mechanisms: (1) GC activity is proportional to weighted mitral cell responses, (2) MC outputs \mathbf{m} are represented by the difference between OSN inputs and inhibitory feedback they receive from the GC, and (3) dendrodendritic synapses are modified via Hebbian plasticity. Since the observed population average of the MC activity is close to zero (Fig. S2A-D), in the model, we computed the expected variance of the MC population activity as a function of time measured in the number of odor presentations:

$$\langle m^2 \rangle(t) \approx \frac{\tau}{4t + C\tau} \quad (1)$$

This equation predicts a power-law decay in population activity variance, where τ represents the characteristic timescale of synaptic plasticity and C is the fitting parameter that accounts for initial MC responses. We fit this model to the mean-squared MC activity of responsive neurons (on average 37 neurons per animal; Methods) across all recorded animals (Fig. 2C). We observed an agreement of the observed decay in the MC response variance with Equation (1). This agreement was observed across mice and odor conditions, with similar time constants τ recovered across experiments (Table S1; Fig. S2E-H for individual animal fits).

Overall, using experimental data of MC population responses to repeated odor presentations across several days, we observed a robust decrease in response amplitudes. This decrease was quantitatively described by the theoretical result [Equation (1)] derived using plasticity in dendrodendritic MC-GC synapses. Good agreement between experimental and theoretical results suggest that MC-GC plasticity may play a role in gain adaptation.

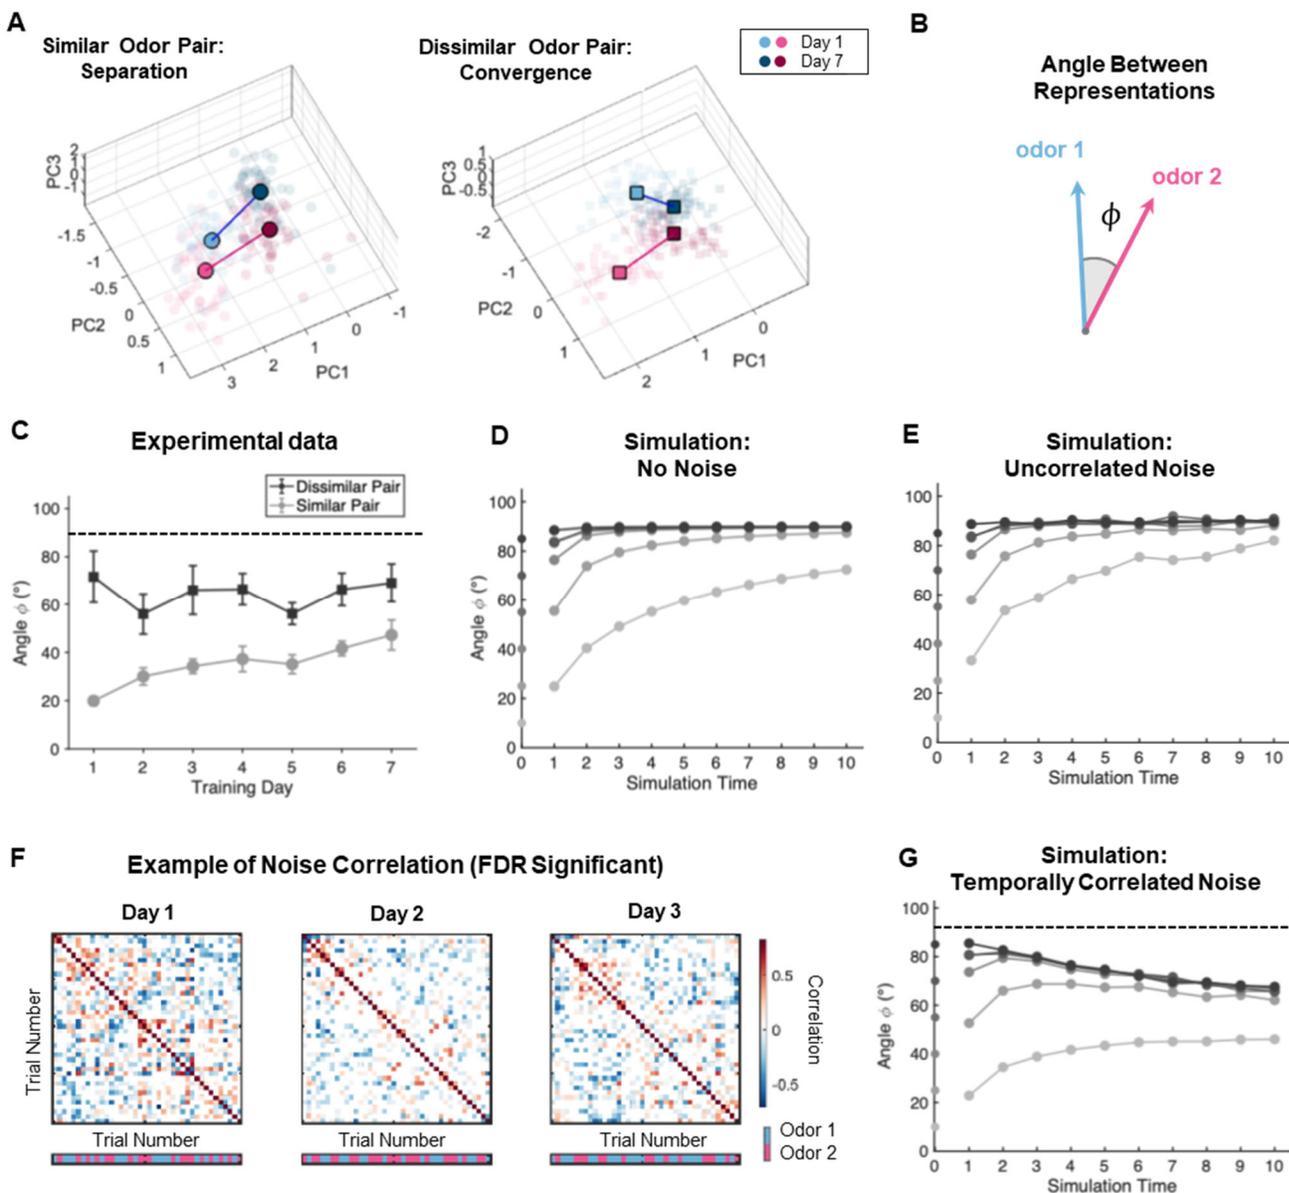


Figure 3. Pattern separation and convergence in OB odor responses.

(A) Similarity-dependent MC activity changes in population responses. 3D visualization of PC1-PC3 of mitral cell population vectors across all trials from representative animals. Left: mouse exposed to similar odor mixtures (Mixture 1 vs. Mixture 2); Right: mouse exposed to dissimilar odors (Odor 1 vs. Odor 2). Transparent points show individual trials; solid points show daily trial averages. Data from Day 1 (light) and Day 7 (dark) illustrate opposing dynamics: similar odor representations separate over time while dissimilar odor representations converge.

(B) Angle between odor representation vectors of two odors ϕ .

(C) Mean angular separation (\pm SEM) between odor population vectors across training days. Similar odor pairs (gray) progressively separate, approaching a larger angle, while dissimilar pairs (black) remain at a non-orthogonal angle.

(D) Model simulation with no noise. Evolution of the angle between odor pairs for varying initial similarities (grayscale levels represent different starting angles). Without noise, all odor pairs separate toward orthogonality regardless of initial similarity.

(E) Model simulation with uncorrelated Gaussian noise. Same format as panel E, but with independent Gaussian noise (10% of input variance). Uncorrelated noise does not prevent separation toward orthogonality and fails to reproduce experimental convergence of dissimilar pairs.

(F) Trial-to-trial noise correlations in experimental data. Correlation matrices of noise residuals (trial responses minus average response) for the first 40 trials per day from one representative animal. First three days shown. For each day, significant trial-by-trial correlations are displayed (two-tailed t-test, FDR corrected using Benjamini-Hochberg procedure); non-significant correlation entries are filled with white. Block structure along the diagonal reveals temporally correlated noise in adjacent trials.

(G) Model simulation with temporally correlated noise. Same format as panels E and F, but with temporally correlated Gaussian noise (10% of input variance, adjacent-trial correlations). Correlated noise enables convergence of initially dissimilar representations to a finite asymptotic angle (dashed line), matching experimental observations.

Pattern separation and convergence in MC responses.

The adaptation of MC population responses under long-term odor exposure differs between similar and dissimilar odor pairs. For highly similar odor mixtures presented over a week, the angle between MC population representation pairs progressively increased (Fig. 3C, light trace). At the same time, as shown in the previous section, the lengths of the population vectors decayed. As a result, the Euclidean distance between the representations of similar odor pairs by MCs remained approximately constant over time (Fig. 3A, left; Fig. S3A, light trace).

For dissimilar odor pairs, the angle between population vectors remained approximately constant, though not equal to 90 degrees (Fig. 3C, dark trace). Here too, the lengths of the population vectors decreased, leading to an overall convergence of dissimilar odor representations (Fig. 3A, right; Fig. S3A, dark trace).

The functional consequences of such an adaptation depend on the readout mechanism implemented in MC target regions. For similar odor pairs, because the angle between representations diverges, normalization can make the downstream signals more distinct. For example, if downstream processing normalizes population vectors to equal lengths, adaptation causes representations of similar odors to move further apart. For dissimilar odor pairs, under this normalization scenario, the distance between representations remains unchanged. If, on the other hand, downstream regions rely on non-normalized population responses, we expect no divergence for similar pairs and a convergence for dissimilar pairs. Thus, the functional interpretation of these transformations depends on the nature of the readout mechanism.

We next interpreted these results in light of our computational model. In the model, a pair of input vectors \mathbf{r} was delivered intermittently to a population of MCs, which interacted with a population of GCs via dendrodendritic synapses (Fig. 2B). The connection weights between MCs and GCs were modified according to Hebbian plasticity rule (Supplementary Information Section 1). The model then computed the output population vectors \mathbf{m} resulting from the interdigitated exposure to odor pairs.

In our simulations, we used 10 MCs and 750 GCs and verified that the conclusions are robust to the number of neurons, provided the population size is sufficiently large (Supplemental Information Section 4). MCs inputs from the odorant receptors (represented by the vector \mathbf{r}) for each smell remained unchanged over days, while the responses on MCs combined in the vector \mathbf{m} were modified by GC inhibition. The network was presented with pairs of input vectors \mathbf{r} , ranging from highly similar to highly dissimilar. Similarity within a pair was defined by the angle between the input vectors \mathbf{r} , which varied from 10° (similar pairs) to 85° (dissimilar pairs, Fig. 3D). We then monitored how the angle between the output vectors \mathbf{m} , representing MC responses, evolved as a function of exposure duration.

We observed that, for *all* odor pairs, both similar and dissimilar, the angle between MC representations systematically increased with odor exposure time (Fig. 3D). In our computational model, this angle asymptotically approached 90° , implying that odor representations became orthogonalized. This behavior was evident both when MC inputs from receptor neurons contained no noise (Fig. 3D) and when MCs, in addition to odorant-dependent inputs, received uncorrelated white noise (Fig. 3E; Methods; Table S2). Thus, our model explained the divergence of similar odor representations over time (Fig. 3C). This occurred because GC feedback learned the overlaps between odor pairs and subtracted them from MC

responses, leading to orthogonalization (Supplemental Information Section 4.1). By contrast, in the experimental data, the angle between representations of dissimilar pairs never reached 90° and instead remained between 50° and 80° (Fig. 3C). This result was *inconsistent* with the predictions of our simple computational model (Fig. 3D, E).

To refine our model, we examined the structure of trial-to-trial noise in MC responses. In the experimental data, we observed substantial noise correlations across trials (Fig. 3F). Incorporating temporally correlated noise into our model successfully reproduced the convergence of initially dissimilar representations to a stable intermediate angle between 0° and 90° (Fig. 3G; Supplementary Information Section 4.2; Fig. S3C-D). These results demonstrate that realistic noise correlations, in addition to the learning rule, are essential for explaining the representational dynamics observed in the OB.

Overall, we showed that the angle between representations of similar odors by the MCs diverge over time. This behavior is replicated by the computational model which included Hebbian learning in the dendrodendritic MC-GC synapses. By incorporating a realistic temporal model for noise, the same model accounts for stable angle between representation of dissimilar odors. Together, these results point to a single circuit mechanism capable of generating two very distinct effects, depending on the relationship between the odor pairs.

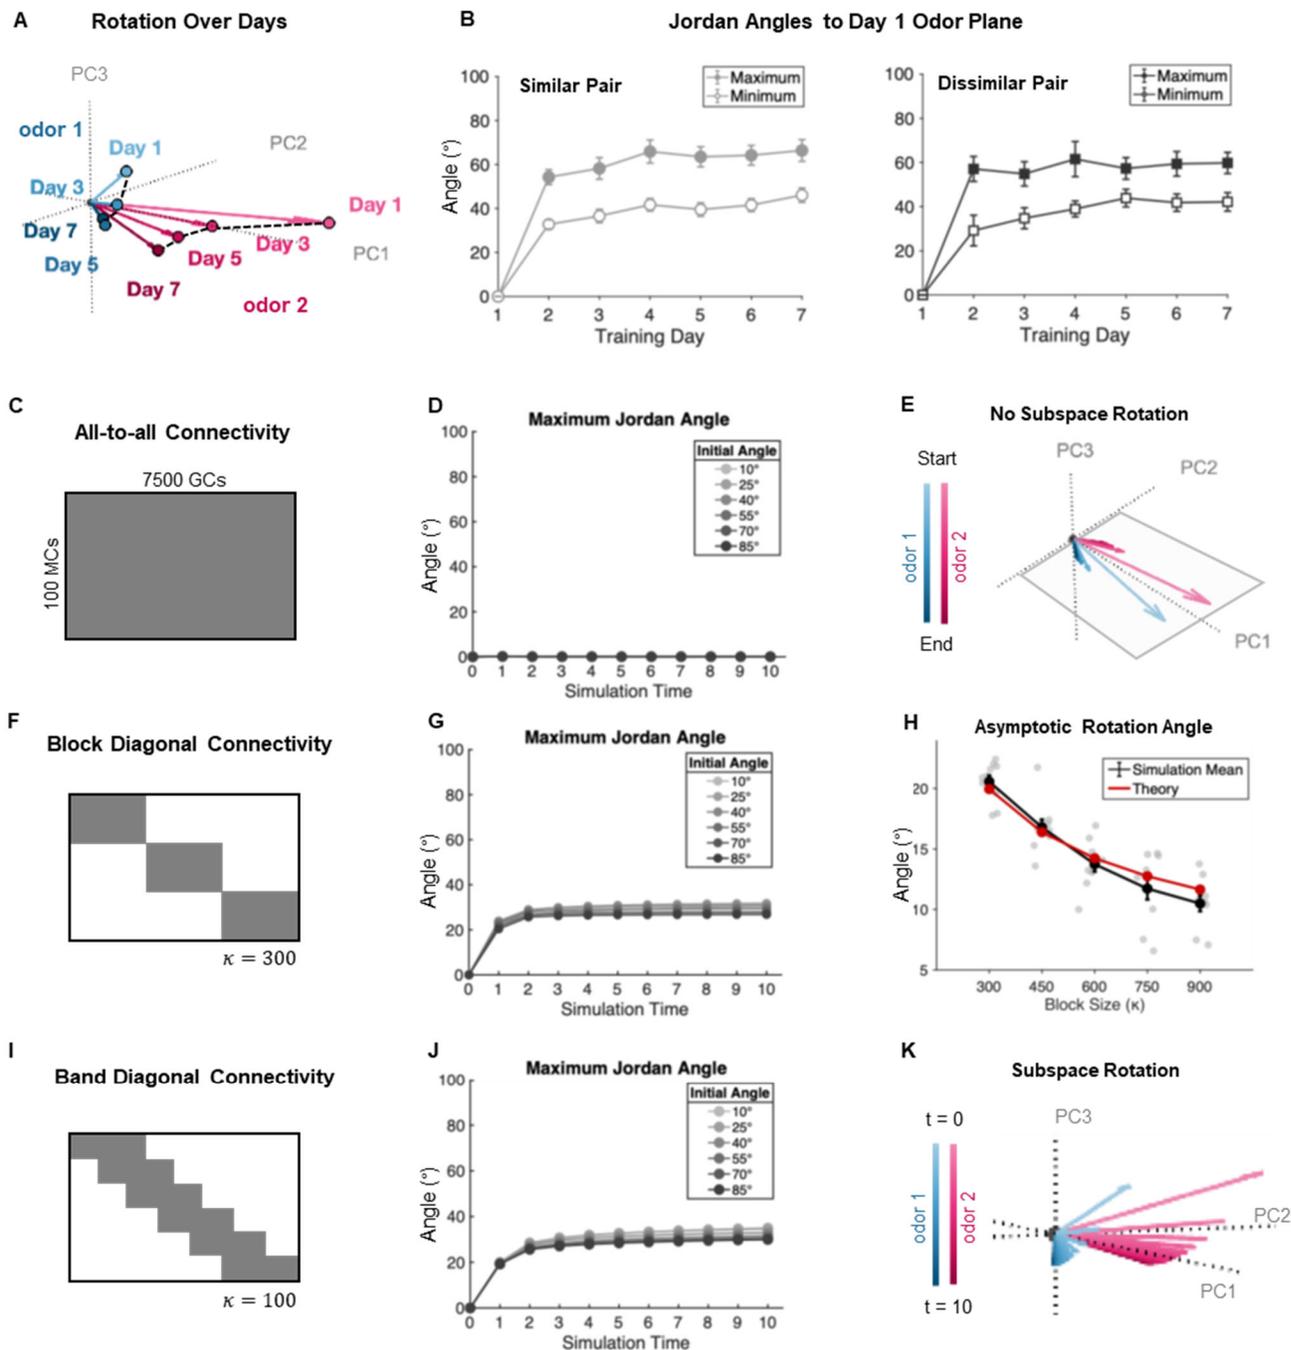


Figure 4. Representational drift results from structured connectivity in the MC-GC circuit.

(A) Experimental evidence of representational drift. Principal component visualization of daily odor trial average from a representative mouse exposed to dissimilar odors. PCA performed on trial average of all days; data for every other day shown for clarity. Odor representations progressively occupy different subspaces over time, indicating systematic rotation distinct from decorrelation within the initial odor plane.

(B) Quantification of subspace rotation using Jordan principal angles. Maximum and minimum Jordan angles between the Day 1 odor subspace and subsequent days measure how the two-dimensional odor representation plane rotates in neural space. Both angles increase progressively for similar (left) and dissimilar (right) odor pairs, indicating that the rotation magnitude is independent of odor similarity. Data show mean \pm SEM across all animals.

(C, D, E) Network simulations for fully connected MC-GC network. **(C)** MC-GC connectivity matrix (gray represents non-zero connections). **(D)** Jordan angles between planes defined by pairs of odor representations. For fully connected network, the representation plane does not rotate. **(E)** The dynamics of odor representations are visualized in PC space. No rotation is evident.

(F, G, H) Network simulations for a sparse MC-GC connectivity. **(F)** In the block-diagonal connectivity matrix, most weights are fixed at zero (white, 90% of weights). The blocks of weights along the diagonal (gray) are allowed to change. **(G)** The maximum Jordan angle between planes defined by pairs of representations and the initial plane as a function of time. Sparse connectivity leads to the representation plane rotation. After certain time of exposure, representation plane stops rotating and the maximum Jordan angle reaches a steady value. **(H)** The asymptotic value of rotation angle of a single input as a function of block size. Rotation is weaker for a bigger block size. Simulation results (black, average over 10 random trials) agree well with theory (red, derived in Supplementary Information Section 3.4).

(I, J, K) Network simulations for band-diagonal sparse MC-GC connectivity. **(I)** The GC-MC connectivity matrix. Gray regions indicate allowed connections (~3% of all weights). **(J)** As in the case of block-diagonal connectivity, the representation planes rotate after repeated odor exposure. **(K)** The dynamics of odor representations visualized in the PC space. The rotation of the representation plane is similar to data in **(A)**.

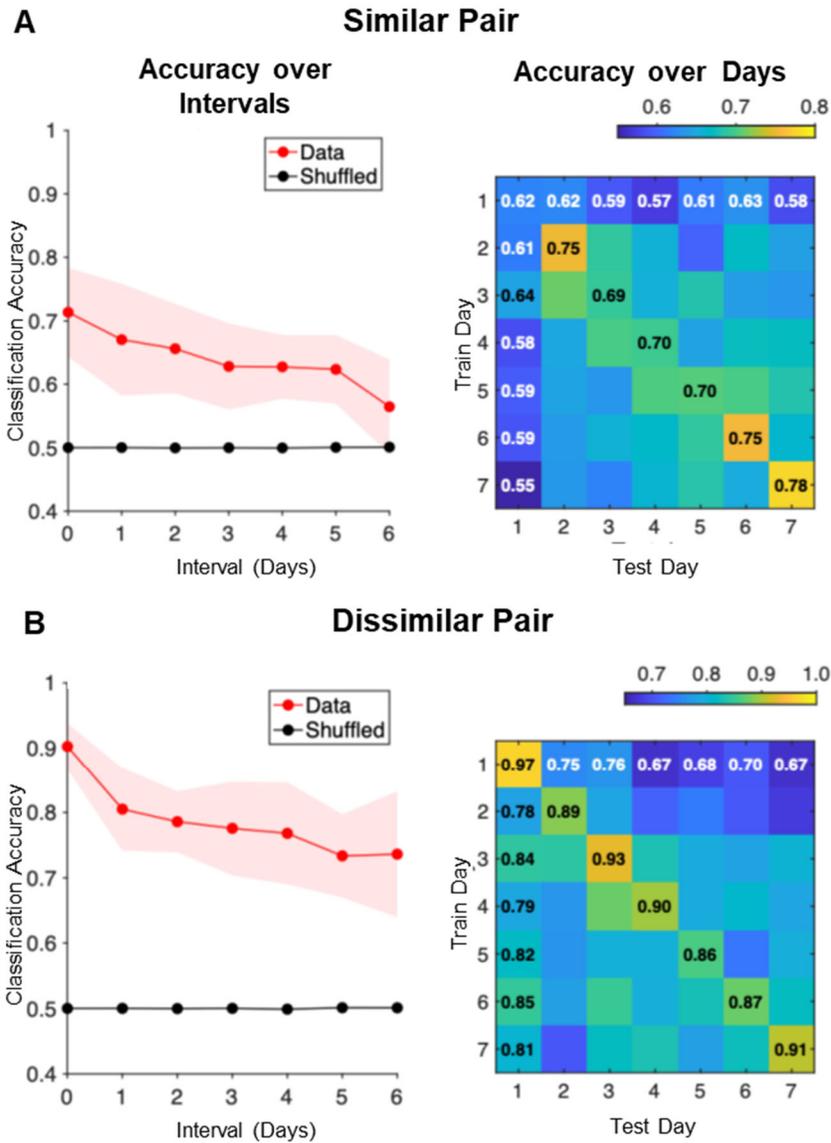


Figure 5. Classification accuracy (linear SVM) within and across days. (A) Classification accuracy for similar odor pairs. Within-day accuracy was obtained using leave-one-out cross-validation; across-day accuracy used models trained on all data from one day and tested on another day. Left: Classification accuracy as a function of time interval between train and test days. For interval equal to zero, within-day accuracy is shown. Mean \pm SD across mice ($n = 10$). Right: Average classification accuracy matrix for all pairs of train and test days. Diagonal elements show within-day accuracy; off-diagonal elements show across-day accuracy. (B) Same as (A), but for dissimilar odor pairs (number of mice, $n = 6$).

Structured connectivity may drive representational drift beyond pattern adaptation, separation and convergence.

In the previous section, we described the relative transformations of representations of pairs of odors under repeated exposure. Here, we examine the global rearrangement of population vectors across several days. To quantify these global changes, we investigated the subspace spanned by the population response vectors to several odors and its evolution over time (Fig. 1B-D, shaded region). In the simplest form of odor adaptation, the two population vectors representing an odor pair move relative to one another within the same subspace (Fig. 1B, C). This type of adaptation simplifies the readout, since the same linear classifier can discriminate between the odors across days. Alternatively, the subspace itself may change, rotating as a result of odor exposure (Fig. 1D). Such a global transformation could complicate discrimination across days. We therefore analyzed how odor exposure affects the population subspace and its stability.

To characterize global transformations, we employed Jordan principal angles (JPAs) (Chae et al., 2019; Jordan, 1875), which quantify the geometric relationship between high-dimensional subspaces. Our goal was to compare the subspaces spanned by population vectors representing the same pair of odors across days. For two 2D planes in 3D space, a single JPA captures the geometric angle between these planes. In higher dimensions, the number of JPAs exceeds one; for instance, the relative position of two 2D planes in 4D space is described by two angles. Intuitively, the minimum (maximum) JPA corresponds to the smallest (largest) angle between any two vectors drawn from the two subspaces. Computing JPAs between odor representations across days provides a convenient way to measure population subspace rotations induced by adaptation.

We computed the minimum and maximum JPAs between population vectors representing responses to odor pairs on day 1 and later days for $n = 16$ mice. For both similar (Fig. 4B, $n = 10$) and dissimilar (Fig. 4C, $n = 6$) odor pairs, JPAs were consistently greater than zero, indicating subspace rotation. Notably, JPAs remain unchanged if population vectors rotate within the same plane. Thus, our findings cannot be explained solely by the relative transformations of odor representations described in the previous section. Instead, the presence of non-zero JPAs across days points to a global transformation of odor representations.

This subspace change resembles representational drift reported in sensory cortices, where stimulus encodings change systematically over timescales of days to weeks (Bauer et al., 2024; Deitch et al., 2021; Marks & Goard, 2021). Piriform cortex data indicates that representational drift depends on odor presentation frequency: less frequently presented odors drift more. To evaluate the frequency dependence of representational drift in the OB, we examined the rotation of odor representations obtained with no or minimal odor exposure, estimated using a different previously published dataset (Shani-Narkiss et al., 2023; Shani-Narkiss et al., 2020). We found that, in the absence of odor exposure, population vectors drift to the extent that is not different with the case of daily odor exposure presented above (Fig. S4A). This observation suggests that OB drift does not depend on odor exposure frequency. This finding further suggests that the exposure frequency dependence observed in cortex may have cortical origins: because cortical and OB drift appear to have different dependences on odor presentation frequencies, the mechanism of the OB subspace rotation described above is likely not to be inherited from the cortex, despite substantial feedback projections from the cortex to the OB.

To understand the subspace rotation described above, we simulated the adaptation of MC responses to pairs of odorants using our model. We found that, in the case of all-to-all MC-GC connectivity, MC representation subspaces do not rotate (Fig. 4C-E). The lack of representation subspace rotation is observed for all dissimilar and similar odor pairs, suggesting that decorrelation preserves the original odor subspace (Fig. 4D). This finding holds true even in the presence of synaptic update stochasticity, which has been shown to drive drift in sensory and hippocampal systems (Fig. S4B; Supplementary Information Section 3.2) (Bauer et al., 2024; Eppler et al., 2025; Morales et al., 2025; Qin et al., 2023). Overall, our computational model suggests that with all-to-all structureless connectivity between MCs and GCs, Hebbian updates preserve the original odor subspace.

Next, we tested if the presence of structure in the MC-GC connectivity might explain the subspace rotation. Indeed, the OB spatial structure restricts GC dendrites to local interactions with nearby MCs, creating connectivity constraints that deviate substantially from all-to-all connectivity (Arnson & Strowbridge, 2017; Egger & Urban, 2006; Kim et al., 2011; Nagayama et al., 2014). We therefore tested whether structured connectivity patterns alone could reproduce the observed subspace rotation.

We first simulated the circuit model with the simplest form of structure, the block-diagonal connectivity (Fig. 4F). With such a structure, the model successfully generated subspace rotation, with the averaged over days Jordan angles comparable to experimental observations (maximum angle in Fig. 4G; minimum angles see Fig. S4C). The block-diagonal assumption also enabled analytical tracking of MC output vector dynamics, revealing that rotation depends on both input statistics and block size: smaller blocks and greater variability across odorant receptor inputs produced larger drift angles (Fig. 4H; Supplementary Information Section 3.4). Our analysis also shows how structure in MC-GC connectivity induces subspace drift. In the block-diagonal model, plasticity operates independently within each block, producing differential suppression of MC responses across blocks. This block-specific scaling modulates the MC population response, causing representations to rotate away from the original input direction. Overall, the block-diagonal connectivity model yields interpretable subspace rotation in MC population responses, consistent with experimental data.

We then tested a more realistic band-diagonal connectivity, which allows each GC to influence every MC through multi-synaptic pathways. This structure captures the spatially dispersed connectivity patterns of the OB (Arnson & Strowbridge, 2017; Egger & Kuner, 2021; Hong & Wilson, 2015; Willhite et al., 2006). This architecture similarly produces subspace rotation (Fig. 4I, J) with trajectories resembling experimental data (Fig. 4K). Together, these results demonstrate that realistic anatomical constraints on plasticity, not just the learning rule itself, are essential for representational dynamics in the OB.

While our analysis of subspace rotation demonstrates global changes in odor representations across days, it does not directly establish whether these changes affect decoding. To address this, we trained linear classifiers (SVMs) on responses from each day and tested their performance both within and across days. For both similar and dissimilar odor pairs, cross-day classification accuracy declined as the interval between training and testing increased, indicating that subspace rotation leads to representational drift (Fig. 5A, B, left) (Masset et al., 2022; Rule et al., 2019; Schoonover et al., 2021). For example, for similar odor pairs, classifier performance decreased from 71% for same-day training and testing to 56% when training and

testing were separated by six days (Fig. 5A). This decline is comparable in magnitude to drift observed in the piriform cortex [(Schoonover et al., 2021); Extended Data Fig. 10f].

At the same time, our classifier analysis revealed the impact of relative transformations between odor pairs described in the previous section. Within-day decoding of similar odor pairs improved at later days compared to Day 1, reflecting enhanced discriminability resulting from odor exposure (Fig. 5A, right, diagonal). This finding suggests that the pattern separation of similar odors increases their linear separability. By contrast, classifier performance for dissimilar odor pairs remained consistently high (Fig. 5B, right). Overall, these results indicate that global transformations of population vectors drive representational drift, necessitating adjustments to odorant classifiers downstream of the OB.

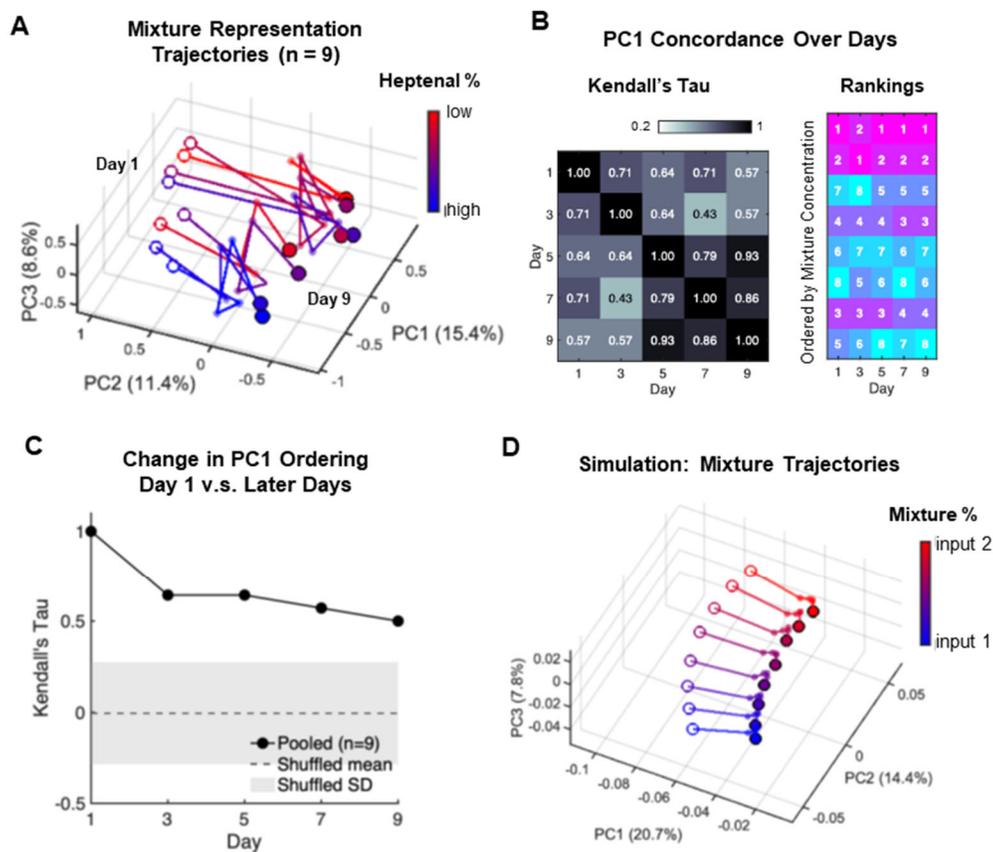


Figure 6. Representational drift preserves geometric structure of the 1D odor manifold.

(A) Mixture representations maintain ordered trajectories during drift. Daily population vector (trial averaged) in response to 8 binary odor mixtures differing in the relative concentrations of Heptanal and Ethyl Tiglate in principal component space across all trials. Data collected on days 1, 3, 5, 7, and 9, with trajectories showing evolution from Day 1 (unfilled circles) to Day 9 (filled circles). Mixtures are color-coded by Heptanal content (red = low, blue = high) (Table S1). Despite substantial drift in PC space, the relative ordering along the concentration gradient remains consistent across days.

(B) Concentration ordering is preserved in PC1 across days. Left: Kendall's tau correlation coefficients measuring preservation of pairwise odor ordering between days. High correlations indicate maintained rank structure. Right: Rank ordering of the 8 mixtures along PC1 for each recording day, with colors representing rank positions (1 = highest PC1 loading, 8 = lowest).

(C) Quantification of ordering preservation. Kendall's tau comparing PC1 ordering on later days versus Day 1 shows significant conservation above chance levels (gray band shows shuffled control mean \pm SD).

(D) Model reproduces ordered drift dynamics. Simulated mixture trajectories using circuit model with block-diagonal connectivity and 10% input noise. Input mixtures created by linearly interpolating between dissimilar odor OSN population response vectors with proportions matching experimental angles on Day 1. Trajectories show trial-averaged evolution every 100 trials.

Concentration gradient structure along representation manifolds is preserved during passive odor exposure.

Our previous results suggest that representational change in the OB under repeated odor exposure can be decomposed into the following components: (i) population-level gain adaptation, (ii) relative pattern separation or convergence, and (iii) global subspace rotation. Although these components were derived from odor pairs, the same framework can be extended to larger stimulus sets. We therefore examined how *stimulus manifolds*, i.e. low-dimensional structures defined by MC population responses to mixtures of the same pure components, are encoded and transformed across days. In one scenario, the representations of individual mixtures are transformed independently across days, and their relative positions are scrambled. In this case, the angles between MC population vectors representing individual mixtures change, and the representation manifold undergoes a geometric deformation. We therefore refer to this form of representational change as *manifold deformation*. In another scenario, the representations of different mixtures rotate as a whole while preserving the relative angles between the MC population vectors. We refer to this hypothetical scenario as *rigid-body manifold rotation* (Pashkhanloo & Koulakov, 2023). Preserving representation manifolds may be important because it ensures that concentration gradients in stimulus space remain systematically reflected in population responses.

Testing the complete geometry of population vectors is not feasible, as it would require measuring odor responses from the entire ensemble of MCs. Indeed, even if population vectors undergo a rigid-body rotation during adaptation, observing only a fraction of the MC population would generally produce an apparent deformation. One possible exception arises when population responses form a one-dimensional (1D) manifold, such as that generated by responses to a binary mixture series. In this case, even partial sampling of the MC population may still reveal a robust rotation of the 1D manifold. We therefore examined 1D trajectories formed in by MC population responses to binary mixture series. We reasoned that if the ordering of mixture responses along the 1D trajectory, as determined by component concentration, is preserved across days, this observation would be consistent with rigid-body manifold rotation. Conversely, if mixture representations become scrambled along the 1D trajectory following adaptation, this would support the manifold deformation hypothesis. Thus, analysis of 1D binary mixture series may provide a first step toward inferring the geometric transformations of MC population representations induced by adaptation.

For a binary mixture of odors A and B, the component concentrations can be written as $C_A = xC_A^p$ and $C_B = (1 - x)C_B^p$, where C_A^p and C_B^p denote the concentrations of the pure components, and x parametrizes the mixture composition. The MC population response vector \mathbf{m} depends on the component concentrations, $\mathbf{m} = \mathbf{m}(C_A, C_B)$, and thus becomes a function of the single parameter x for a given binary mixture series. As the relative concentrations of the components vary with x , the neural population vector \mathbf{m} traces a 1D trajectory, in the population space, which we refer to as a *representation geodesic*. Under repeated exposure to the mixture series over multiple days, this representation geodesic is expected to change. Our goal is to examine whether the ordering of population vectors \mathbf{m} corresponding to different mixture compositions along the geodesic is preserved across days. Previous studies have shown that the gradient between two pure odors projects orderly onto a one-dimensional axis in neural population space, supporting the existence of a representation geodesic (Adefuin et al., 2022; Carlsson et al., 2007; Fletcher, 2011). However, how such geodesics evolve under repeated exposure has

not been investigated. We analyzed MC population responses from mice exposed to a binary mixture series of Heptanal and Ethyl Tiglate over more than a week (8 mixtures; see Table S1 for mixture composition). On Day 1, population responses were ordered according to mixture composition as visualized in the PC space (Fig. 6A). Across days, response amplitudes decreased, consistently with gain adaptation (Fig. S6A). The distribution of pairwise angles between odor population vectors became tighter over time, converging toward $\sim 60^\circ$ (Fig. S6B-C). Although the representation subspace rotated across days, the mixture ordering appears to remain generally preserved as assessed visually (Fig. 6A; Fig. S6D).

To quantitatively evaluate mixture ordering stability along the geodesic, we measured the rank stability of daily odor subspaces. To this end, we computed the first principal component (PC1) of the population response vectors for each day separately, which identifies the direction of the representation geodesic. This direction is expected to rotate across days due to the subspace rotation discussed above. We then projected mixture responses onto the corresponding PC1. If the ordering of representations across days is consistent or random, the ordering of mixtures along PC1 according to the mixture composition parameter x defined above is expected to be conserved or scrambled, respectively. Consequently, Kendall's tau correlation of these orderings across days is expected to be close to one or zero, respectively. We found that Kendall's tau correlations across days were consistently high (Fig. 6B, left), and the overall concordance across all days (Kendall's $W=0.87$, $p<10^{-4}$) indicated robust preservation of mixture ordering. Although occasional local inversions occurred between adjacent mixtures, the global rank order was significantly higher than in shuffled controls (Fig. 6C). Furthermore, the ordering axis aligned with the concentration gradient of the binary mixtures (Fig. 6B, right). Computational modeling incorporating structured connectivity and input noise reproduced these drift trajectories while preserving concentration ordering (Fig. 6D). Together, these results demonstrate that the low-dimensional structure of the stimulus ensemble (here, the binary mixture geodesic) is robustly preserved under representational change. Even as population responses undergo gain adaptation, pattern separation/convergence, and subspace rotation, the encoding of concentration gradients remains stable. This finding is broadly in line with a rigid-body-like rotation of the stimulus manifold, while the remaining local reordering of mixture representations and the tightening of the pairwise angle distribution across days (Fig. S6C) are indicative of additional non-rigid manifold transformations. Together, these results suggest that repeated odor exposure preserves the global organization of the representation manifold.

Discussion

In this study, we examined how odor representations in the OB evolve over days of repeated passive exposure. By analyzing odor responses obtained using longitudinal calcium imaging in the bulb, we identified distinct but concurrent components of representational change: population-level gain adaptation, relative pattern separation or convergence, and global rotation of the encoding subspace. These phenomena could be collectively captured by a mechanistic computational model of the MC-GC circuit. The model demonstrates that dynamics and input-dependent plasticity within this recurrent network are sufficient to reproduce the spectrum of observed representational transformations.

Our first major finding concerns the overall activity of MCs during repeated exposure to similar and dissimilar odor pairs. Across days, we observed a gradual decay in the amplitude of population responses (Fig. 2A; Fig. S2A-D). Using a mean-field approximation of our circuit model, we showed that the variance of MC population activity on later days follows a power-law decay (Equation 1). This model accurately fits the population variance measured in our data across animals and odor conditions (Fig. 2C; Fig. S2E-H), indicating that activity-dependent plasticity in the MC-GC synapses can account for the long-term gain adaptation observed experimentally.

Second, we found that the relative changes between odor representations depend on the similarity of the odor inputs. For highly similar odor pairs, the angle between MC population vectors steadily increased across days, whereas for dissimilar pairs the angle remained stable around approximately a 70-degree angle (Fig. 3A, C). Using our model, we showed that Hebbian plasticity drives orthogonalization of input pairs regardless of their initial similarity (Fig. 3D, E). Incorporating temporally correlated noise, the presence of which we confirmed in the experimental data (Fig. 3F), stabilized the asymptotic angle between representations, producing convergence to a value below 90°, as observed experimentally (Fig. 3G).

Third, we observed a gradual rotation of the population encoding subspace across days, quantified using Jordan principal angles (JPAs). JPAs measured between the initial and subsequent days of odor exposure describes the degree of rotation of odor representation subspaces. Both similar and dissimilar odor pairs exhibited progressive increases in JPAs over time, indicating a global transformation of the representation plane (Fig. 4A, B). This subspace rotation also resulted in a decay of cross-day linear decoding accuracy (Fig. 5A, B). This implies that the linear decoder trained on day 1 of odor exposure could discriminate odors on subsequent days with lower accuracy. This finding suggests that subspace rotation in the OB may need to be compensated by a continual recalibration of downstream decoders, supporting its interpretation as representational drift. The subspace rotation could be accounted for by our MC-GC computational model through introducing biologically motivated structured connectivity constraints into the MC-GC circuit (Fig. 4F-H, J). Our model also predicted that repeated odor exposure leads to a plateau in rotation over time (Figs. 4G, J) as observed experimentally (Fig. 4B). By contrast, variability in the inputs alone, although capable of inducing some output rotation, did not generate the saturating rotation angle we observed experimentally (Fig. S5).

Finally, we extended the framework beyond the analysis of odor pairs to explain how odor representation manifolds defined by mixtures evolve under repeated exposure. One implication of our model is that plasticity and network dynamics in the MC-GC circuit should preserve the relative arrangement of the population vectors as absolute representations drift. This type of drift

was termed here as rigid-body rotation (Pashakhanloo & Koulakov, 2023). In agreement with this hypothesis, the analysis of MC responses to an eight-odor binary mixture panel revealed that the ordering of representation vectors, when projected onto the first principal component, remained approximately stable across days (Fig. 6A-C). Simulations incorporating input noise and structured connectivity reproduced drift trajectories while maintaining the concentration ordering (Fig. 6D), demonstrating that the circuit can undergo substantial representational change without disrupting the low dimensional geometry of representation manifolds. Our findings are generally consistent with the global rigid-body rotation of the representation manifold, although some degree of local manifold deformation is also observed (Fig. S6).

While our model focuses on recurrent interactions between MCs and GCs, other structures in the OB could also contribute to the observed changes. Adaptation at the single-neuron level can occur through receptor desensitization (Antunes et al., 2014) or short-term synaptic plasticity during odor presentation or sniff bouts (Dietz & Murthy, 2005; Zhou et al., 2020). However, these mechanisms operate on short timescales (seconds to minutes) and are unlikely to account for the multi-day gain adaptation reported here. Additional gain control mechanisms are implemented in the glomerular layer by periglomerular cells and short-axon cells, which connect presynaptically to MCs or laterally among glomeruli. These interneurons have been reported to exert normalizing, population-wide gain control and their plasticity could, in principle, contribute to multi-day decreases in activity (Banerjee et al., 2015; Burton et al., 2017; Vaaga et al., 2017). However, our data show heterogeneous changes across individual MCs: some neurons decrease activity over days, while others increase (Fig. S1), a pattern more consistent with our model in which inhibitory strengthening is input and MC-specific. Nevertheless, we cannot rule out the possibility that the observed population-level changes arise from a combination of mechanisms across multiple interneuron types (Zavitz et al., 2020).

Numerous previous studies have modeled OB circuits (Cavarretta et al., 2018; Gilra & Bhalla, 2015; Kersen et al., 2022; Migliore et al., 2015; Migliore & Shepherd, 2008) to reveal how this network shapes inputs into the cortex. We propose that pattern separation in the OB over days arises from recurrent inhibition and plasticity at the MC-GC dendrodendritic synapses. In our model, local Hebbian plasticity progressively increases the angle between population vectors, orthogonalizing odor representations. Previous studies have shown that nonlinearity in neuronal responses can decorrelate input patterns via thresholding facilitated by sparse recurrent connectivity, providing a biologically plausible explanation for pattern decorrelation in the OB (Cayco-Gajic & Silver, 2019; Dasgupta et al., 2017; Wiechert et al., 2010). Several studies have proposed that learning and adaptation in the OB can further enhance pattern separation, especially for fine odor discrimination. For example, (Wick et al., 2010) compared feedforward and feedback (recurrent) inhibition and showed that the recurrent circuit produces stronger decorrelation. Linster and Cleland implemented a spike timing-dependent plasticity rule in a spiking model of vertebrate OB and insect antennal lobe networks, demonstrating that progressive sharpening of odor representations emerges from learning between principal neurons and their postsynaptic inhibitory partners, though they did not specifically examine the MC-GC circuit (Linster & Cleland, 2010). Several studies (Cecchi et al., 2001; Chow et al., 2012; Kersen et al., 2022) have built detailed anatomically and physiologically realistic models to study fine discrimination in the context of OB neurogenesis alone. These models show that OB neurogenesis helps decorrelate MC responses. Most recently, Meng and Riecke (Meng & Riecke, 2022) modeled plasticity between GCs and MCs during repeated exposure and assumed a non-monotonic relationship between synaptic spine formation rate and neuronal

activity. In the context of input decorrelation, our work is consistent with these findings but differs in approach and further extends their results. By using a tractable linear system, we provide an analytical account of how population vectors evolve over the course of several days. We also propose that specific correlations in the MC input noise limit the levels of decorrelation in a manner consistent with experimental data.

The OB contains the first neural network that processes olfactory inputs and is generally thought to maintain relatively stable stimulus representations (Hirata, 2024), making our observation of representational drift at this early stage particularly noteworthy. Representational drift, on the other hand, has been previously reported in the piriform cortex (Morales et al., 2025; Schoonover et al., 2021). We found that, in the bulb, the degradation of linear classifier performance across days was numerically comparable to that reported in the cortex (Fig. S4A). Although we attribute the observed drift of the representation manifold to mechanisms within the OB, an alternative possibility is that the drift originates in the piriform cortex (Morales et al., 2025; Schoonover et al., 2021) and is conveyed to the bulb via corticobulbar feedback. Several arguments suggest, however, that the drift we observe may arise from intrinsic OB mechanisms. First, prior work indicates that during passive odor exposure, i.e. under the experimental conditions used here, cortical feedback remains relatively weak (Wu et al., 2020), making a cortical origin of OB drift unlikely. Second, both our experimental data and theoretical model show that the rotation of the MC representation plane quantifying the drift in the OB saturates after seven days of exposure, whereas cortical drift is reported to progress continuously without saturation. It is possible that longer-term structural connectivity changes in the OB driven by ongoing neurogenesis could in turn drive longer-term representational drift beyond the seven-day period, a possibility that warrants further investigation. Third, we find that the rate of drift in the OB is similar in overall magnitude to that seen in the cortex, but, unlike cortical drift, is not affected by odor exposure frequency. Overall, we propose that OB drift may be partially influenced by cortical feedback, but such feedback alone is unlikely to fully account for the phenomenon that we report.

At first glance, this aspect of drift appears to challenge classic neural coding theory, which emphasizes the stability of sensory representations in early processing stages (Morales et al., 2025; Patterson et al., 2013; Shani-Narkiss et al., 2023). To examine representation stability, we analyzed responses to a series of binary mixtures with gradually varying concentrations interpolating along a one-dimensional geodesic. We found that this geodesic rotated as a whole while preserving the relative positions of the population vectors representing individual mixtures. Preserving such a manifold may be important because it ensures that concentration gradients in stimulus space remain systematically mapped to population responses, even as plasticity alters their absolute configuration. Our computational model reproduced this coexistence of drift and manifold preservation, providing a mechanistic explanation for how stability and plasticity can co-occur in the OB.

Since representational drift was first described, most studies have emphasized stochasticity in synaptic weight updates as its primary source (Bauer et al., 2024; Eppler et al., 2025; Pashakhanloo & Koulakov, 2023; Qin et al., 2023). Here, we highlight a complementary mechanism, drift emerging from fixed structural constraints. In the OB, connectivity between MCs and GCs is both sparse and topologically constrained, and our model shows that such structure alone can drive systematic subspace rotations in the absence of noise (Fig. 4F-K). This rotation occurs from imperfect suppression of shared components between two odor

representations due to the constraints on network connectivity. In a fully connected unconstrained MC-GC network, inhibitory feedback can selectively suppress overlapping components of activity patterns, and no subspace rotation occurs. By contrast, when each GC samples only a restricted subset of MCs, inhibitory feedback cannot selectively cancel shared components of two odor representations without also affecting their unique components. As a result, learning reduces overlap only approximately, leaving residual responses in the perpendicular to the original representation subspace directions. This residual activity manifests as a rotation of the encoding subspace, or representational drift (Fig. 4C-E). Similar ideas have been proposed elsewhere. For example, Kong et al. (Kong et al., 2024; Zabehe et al., 2025) suggested that differences in structural connectivity sparsity could account for the contrasting levels of drift observed in hippocampal CA1 versus CA3. The existence of structural constraints does not rule out stochastic fluctuations, but our findings suggest that, in the OB, fixed architecture may be a major contributor to drift. More broadly, such mechanisms may operate in other circuits where connectivity is locally organized. For instance, this may occur in primary visual cortex, where “like-to-like” connectivity is prominent. In this way, our results extend studies of representational drift by demonstrating that anatomical structure alone can generate long-term changes in population codes.

To further validate our proposed local learning mechanism between MCs and GCs, more studies of long-term synaptic plasticity at the dendrodendritic synapse are needed (Abraham et al., 2010; Egger & Kuner, 2021; Gao et al., 2009; Satou et al., 2005). Such plasticity would provide a compelling local mechanism for unsupervised perceptual learning. In sensory systems, decorrelation of neural representations for similar inputs can facilitate downstream classification by simple readout mechanisms, thereby accelerating subsequent task learning, as shown in visual cortex (Fink et al., 2025; Zhong et al., 2025). We hypothesize that analogous unsupervised changes in the OB may form the basis for reward association, odor identity reconstruction and other behaviorally relevant learning in the olfactory system (Hiratani & Latham, 2020; Krishnamurthy et al., 2022; Zavatone-Veth et al., 2023). These processes likely involve stronger and more specific cortical feedback than those studied here and may require considering the piriform cortex and OB together as an integrated circuit. Further work will be needed to integrate these mechanisms across bulb and cortex to understand how long-term learning shapes olfactory processing (Kumar et al., 2021; Wu et al., 2020; Yu et al., 2025).

Taken together, our findings provide a framework for understanding how the OB adapts to repeated odor exposure. By integrating quantitative analysis of longitudinal calcium imaging with a mechanistic MC-GC circuit model, we show that population-level gain adaptation, pattern separation, and subspace rotation can emerge from local plasticity and structural constraints. This work extends the concept of representational drift to an early sensory stage while demonstrating that low-dimensional manifold structure can be preserved despite ongoing change. More broadly, our results illustrate how stability and plasticity can coexist in recurrent sensory circuits and suggest that similar principles may govern other brain regions with sparse connectivity.

Acknowledgements

This work was supported by the National Institutes of Health BRAIN Initiative grant number U19NS112953 and, in part, by National Science Foundation grant PHY-1748958 to the Kavli Institute for Theoretical Physics. C.X.Z. was supported by William R. Miller Fellowship at the School of Biological Sciences at Cold Spring Harbor Laboratory. S.N. thanks funding from the Simons Center for Quantitative Biology at Cold Spring Harbor Laboratory.

Methods

Experimental Dataset Details

Chu et al., 2016: Two-photon calcium imaging for mitral cell-specific calcium indicator was performed using longitudinal recordings of the same mitral cell populations over 7 consecutive days. During passive exposure sessions, mice experienced odorant presentations (4 s duration) of either a dissimilar pair (Odor 1, Heptanal; Odor 2, Ethyl Tiglate) or a similar pair ((mixture 1, 52% Heptanal/48% Ethyl Tiglate; mixture 2, 48% Heptanal/52% Ethyl Tiglate)) or followed by 15 s inter-trial intervals, with 75 trials per odorant delivered pseudo-randomly per session. No behavioral task or water reward was provided during passive exposure.

Li et al., 2018: Two-photon calcium imaging was performed every other day using longitudinal recordings of the same mitral cell populations across multiple sessions over 9 days. During passive exposure sessions, mice experienced odorant presentations (4 s duration) of eight different binary mixtures (54% Heptanal/46% Ethyl Tiglate through 46% Heptanal/54% Ethyl Tiglate in 2-3% increments), followed by 15 s inter-trial intervals, with 150 trials delivered pseudo-randomly per session. No behavioral task or water reward was provided during passive exposure.

Shani-Narkiss et al., 2020 & Shani-Narkiss et al., 2023: Two-photon calcium imaging was performed on awake head-fixed mice using longitudinal recordings of the same mitral cell populations across two sessions separated by 4 weeks. During odor presentation sessions, mice experienced odor panel presentations (2 s duration), in which Ethyl Tiglate was included, presentation is followed by 26 s inter-stimulus intervals, and each session contains 5 trials per odor. No behavioral task or water reward was provided during passive exposure. This dataset served as a control condition examining long-term stability of olfactory responses without chronic odor exposure.

Data Analysis

Calcium Activity and Population Vector Formation

Following Chu et al. (2016) and Li et al. (2018), for each trial and each mitral cell, the calcium signal time series was normalized to the average fluorescence value during the baseline period (5 s before odorant onset) to calculate dF/F . Mitral cells that are classified as responsive in at least one session are considered (See criteria in (Chu et al., 2016)). The MC population response in each trial was expressed as an activity vector, which was a concatenation of the time-averaged dF/F from the first and second 2 s windows during the 4 s odorant period of each mitral cell. We used the trial average within each day's session as the representation of odor response on that day. Using the full 4 s period to form population vectors did not change the main conclusions.

For Shani-Narkiss et al. (2020, 2023) data, we followed the original paper's method to calculate relative fluorescence change (dF/F), defining baseline fluorescence for each cell in each trial as its mean fluorescence measured 5-2 s before odor onset. All traces were smoothed with a moving average filter (span = 5). To form the population vector, each entry represents the average dF/F during the 2-second odor presentation.

Principal Component Analysis

For individual mice, principal component analysis was performed for response vectors of all trials across all recording days in the Chu et al. (2016) and Li et al. (2018) datasets. For analysis and visualization across mice, a matrix of trial response vectors was created using all mitral cells that are classified as responsive in at least one session for each mouse using the first n trials (determined by the session with the lowest number of trials for a single odorant). Pooled response vectors were created by concatenating response vectors across animals. Given that calcium responses represent magnitude changes from baseline (dF/F) and mitral cell activation patterns are naturally sparse, PCA was performed without centering to preserve the meaningful zero baseline and activation magnitudes.

Jordan Angle between Subspaces

To quantify changes in odor representation geometry across days, we computed Jordan angles between odor subspaces. For each recording day, we defined the odor subspace as that spanned by the trial-averaged population vectors for each odorant. Specifically, for day i with k odorants, the subspace S_i was spanned by k vectors, where each vector represents the mean population response to one odorant across all trials on that day. Jordan angles between subspaces S_i and S_j (corresponding to different recording days) were computed using the following procedure. Let A and B be matrices whose columns are orthonormal bases for subspaces S_i and S_j , respectively. The Jordan angles $\phi_1, \phi_2, \dots, \phi_k$ are defined as:

$$\cos(\phi_1) = \sigma_1$$

where σ_1 are the singular values of $A^T B$, obtained from the singular value decomposition:

$$A^T B = U \Sigma V^T \text{ with } \Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_k) \text{ where } \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k.$$

For the two-odorant datasets (Chu et al., 2016), this yielded two principal angles per day-pair comparison: a minimum angle ($\phi_{min} = \phi_1$) and maximum angle ($\phi_{max} = \phi_2$). For the eight-odorant dataset (Li et al., 2018), eight principal angles were obtained, representing the full spectrum of geometric relationships between the higher-dimensional odor subspaces.

Linear Classifier Analysis for Odor Pairs

Binary classification was performed using a linear Support Vector Machine with soft margins. Features were standardized (z-scored), and the soft margin parameter C was set to 1. For within-day classification, we used leave-one-out cross-validation, testing classifier prediction on each trial that was left out. Accuracy was reported as the average across all left-out trials. For across-day classification, we trained the SVM on all trials from one day and tested on all trials from another day. Performance on shuffled data was assessed by randomly permuting the odor labels ($n = 1000$ shuffles).

Modeling Details

Model Fitting

We fit the mean-field decay model equation (1) to mitral cell activity in both individual animals and pooled populations across animals. For each responsive neuron, we calculated the squared trial-averaged activity and obtained the mean squared activity across the population for each day. The daily mean squared activity was then fit using least squares regression to estimate model parameters (Table S1).

Simulation

We initialized the network with random small weights and assumed fast dynamics for mitral cell-granule cell interactions coupled with slow Hebbian learning and weight decay (detailed modeling set-up see Supplementary Information Section 1). Inputs r are drawn from zero-centered Gaussian distribution.

The simulation protocol consisted of two phases:

- *Pre-exposure phase*: The network was presented with random inputs drawn from the same distribution as the experimental inputs to ensure proper weight initialization before chronic exposure
- *Chronic exposure phase*: Fixed input patterns were presented in an interleaved manner over many trials, mimicking the experimental paradigm

Network metrics, including principal angles between output representations, were evaluated every 100 trials for two-odor inputs and every 50 trials for eight-odor input panels, chosen to approximate the trial numbers in the experimental data. Input pattern (r), input noise, and weight update noise were drawn from Gaussian distributions. Detailed parameter values are provided in Table S2.

References

- Abraham, N. M., Egger, V., Shimshek, D. R., Renden, R., Fukunaga, I., Sprengel, R., Seeburg, P. H., Klugmann, M., Margrie, T. W., Schaefer, A. T., & Kuner, T. (2010). Synaptic Inhibition in the Olfactory Bulb Accelerates Odor Discrimination in Mice. *Neuron*, 65(3). <https://doi.org/10.1016/j.neuron.2010.01.009>
- Adefuin, A. M., Lindeman, S., Reinert, J. K., & Fukunaga, I. (2022). State-dependent representations of mixtures by the olfactory bulb. *eLife*, 11. <https://doi.org/10.7554/eLife.76882>
- Aghvami, S. S., Kubota, Y., & Egger, V. (2022). Anatomical and Functional Connectivity at the Dendrodendritic Reciprocal Mitral Cell–Granule Cell Synapse: Impact on Recurrent and Lateral Inhibition. *Frontiers in Neural Circuits*, 16. <https://doi.org/10.3389/fncir.2022.933201>
- Alonso, M., Lepousez, G., Wagner, S., Bardy, C., Gabellec, M.-M., Torquet, N., Lledo, P.-M., Alonso, M., Lepousez, G., Wagner, S., Bardy, C., Gabellec, M.-M., Torquet, N., & Lledo, P.-M. (2012). Activation of adult-born neurons facilitates learning and memory. *Nature Neuroscience* 2012 15:6, 15(6). <https://doi.org/10.1038/nn.3108>
- Antunes, G., Sebastião, A. M., & Simoes de Souza, F. M. (2014). Mechanisms of Regulation of Olfactory Transduction and Adaptation in the Olfactory Cilium. *PLOS ONE*, 9(8), e105531. <https://doi.org/10.1371/journal.pone.0105531>
- Arevian, A. C., Kapoor, V., Urban, N. N., Arevian, A. C., Kapoor, V., & Urban, N. N. (2007). Activity-dependent gating of lateral inhibition in the mouse olfactory bulb. *Nature Neuroscience* 2008 11:1, 11(1). <https://doi.org/10.1038/nn2030>
- Arnson, H. A., & Strowbridge, B. W. (2017). Spatial Structure of Synchronized Inhibition in the Olfactory Bulb. *Journal of Neuroscience*, 37(43). <https://doi.org/10.1523/JNEUROSCI.1004-17.2017>
- Aroniadou-Anderjaska, V., Ennis, M., & Shipley, M. T. (1999). Dendrodendritic Recurrent Excitation in Mitral Cells of the Rat Olfactory Bulb. *Journal of Neurophysiology*, 82(1). <https://doi.org/10.1152/jn.1999.82.1.489>
- Banerjee, A., Marbach, F., Anselmi, F., Koh, M. S., Davis, M. B., Garcia da Silva, P., Delevich, K., Oyibo, H. K., Gupta, P., Li, B., & Albeanu, D. F. (2015). An Interglomerular Circuit Gates Glomerular Output and Implements Gain Control in the Mouse Olfactory Bulb. *Neuron*, 87(1), 193-207. <https://doi.org/10.1016/j.neuron.2015.06.019>
- Bao, X., Raguette, L. L., Cole, S. M., Howard, J. D., & Gottfried, J. A. (2016). The role of piriform associative connections in odor categorization. *eLife*, 5. <https://doi.org/10.7554/eLife.13732>
- Bauer, J., Lewin, U., Herbert, E., Gjorgjieva, J., Schoonover, C. E., Fink, A. J. P., Rose, T., Bonhoeffer, T., & Hübener, M. (2024). Sensory experience steers representational drift in mouse visual cortex. *Nature Communications*, 15(1), 9153. <https://doi.org/10.1038/s41467-024-53326-x>
- Boyd, Alison M., Sturgill, James F., Poo, C., & Isaacson, Jeffrey S. (2012). Cortical Feedback Control of Olfactory Bulb Circuits. *Neuron*, 76(6). <https://doi.org/10.1016/j.neuron.2012.10.020>
- Bozza, T., Vassalli, A., Fuss, S., Zhang, J.-J., Weiland, B., Pacifico, R., Feinstein, P., & Mombaerts, P. (2009). Mapping of Class I and Class II Odorant Receptors to Glomerular Domains by Two Distinct Types of Olfactory Sensory Neurons in the Mouse. *Neuron*, 61(2). <https://doi.org/10.1016/j.neuron.2008.11.010>
- Buck, L., & Axel, R. (1991). A novel multigene family may encode odorant receptors: a molecular basis for odor recognition. *Cell*, 65(1), 175-187. [https://doi.org/10.1016/0092-8674\(91\)90418-x](https://doi.org/10.1016/0092-8674(91)90418-x)
- Burton, S. D., LaRocca, G., Liu, A., Cheetham, C. E. J., & Urban, N. N. (2017). Olfactory Bulb Deep Short-Axon Cells Mediate Widespread Inhibition of Tufted Cell Apical Dendrites. *The Journal of Neuroscience*, 37(5), 1117-1138. <https://doi.org/10.1523/JNEUROSCI.2880-16.2016>

- Calu, D. J., Roesch, M. R., Stalnaker, T. A., & Schoenbaum, G. (2007). Associative Encoding in Posterior Piriform Cortex during Odor Discrimination and Reversal Learning. *Cerebral Cortex*, 17(6). <https://doi.org/10.1093/cercor/bhl045>
- Carlsson, M. A., Chong, K. Y., Daniels, W., Hansson, B. S., & Pearce, T. C. (2007). Component Information Is Preserved in Glomerular Responses to Binary Odor Mixtures in the Moth *Spodoptera littoralis*. *Chemical Senses*, 32(5), 433-443. <https://doi.org/10.1093/chemse/bjm009>
- Cavarretta, F., Burton, S. D., Igarashi, K. M., Shepherd, G. M., Hines, M. L., Migliore, M., Cavarretta, F., Burton, S. D., Igarashi, K. M., Shepherd, G. M., Hines, M. L., & Migliore, M. (2018). Parallel odor processing by mitral and middle tufted cells in the olfactory bulb. *Scientific Reports* 2018 8:1, 8(1). <https://doi.org/10.1038/s41598-018-25740-x>
- Cayco-Gajic, N. A., & Silver, R. A. (2019). Re-evaluating Circuit Mechanisms Underlying Pattern Separation. *Neuron*, 101(4), 584-602. <https://doi.org/10.1016/j.neuron.2019.01.044>
- Cecchi, G. A., Petreanu, L. T., Alvarez-Buylla, A., Magnasco, M. O., Cecchi, G. A., Petreanu, L. T., Alvarez-Buylla, A., & Magnasco, M. O. (2001). Unsupervised Learning and Adaptation in a Model of Adult Neurogenesis. *Journal of Computational Neuroscience* 2001 11:2, 11(2). <https://doi.org/10.1023/A:1012849801892>
- Chae, H., Banerjee, A., Dussauze, M., & Albeanu, D. F. (2022). Long-range functional loops in the mouse olfactory system and their roles in computing odor identity. *Neuron*, 110(23). <https://doi.org/10.1016/j.neuron.2022.09.005>
- Chae, H., Kepple, D. R., Bast, W. G., Murthy, V. N., Koulakov, A. A., & Albeanu, D. F. (2019). Mosaic representations of odors in the input and output layers of the mouse olfactory bulb. *Nature Neuroscience*, 22(8), 1306-1317. <https://doi.org/10.1038/s41593-019-0442-z>
- Chaudhury, D., Manella, L., Arellanos, A., Escanilla, O., Cleland, T. A., & Linster, C. (2010). Olfactory bulb habituation to odor stimuli. *Behavioral neuroscience*, 124(4). <https://doi.org/10.1037/a0020293>
- Chess, A., Simon, I., Cedar, H., & Axel, R. (1994). Allelic inactivation regulates olfactory receptor gene expression. *Cell*, 78(5). [https://doi.org/10.1016/S0092-8674\(94\)90562-2](https://doi.org/10.1016/S0092-8674(94)90562-2)
- Chow, S.-F., Wick, S. D., & Riecke, H. (2012). Neurogenesis drives stimulus decorrelation in a model of the olfactory bulb. *PLoS Computational Biology*, 8(3), e1002398. <https://doi.org/10.1371/journal.pcbi.1002398>
- Chu, M. W., Li, W. L., & Komiyama, T. (2016). Balancing the Robustness and Efficiency of Odor Representations during Learning. *Neuron*, 92(1), 174-186. <https://doi.org/10.1016/j.neuron.2016.09.004>
- Chu, M. W., Li, W. L., & Komiyama, T. (2017). Lack of Pattern Separation in Sensory Inputs to the Olfactory Bulb during Perceptual Learning. *eNeuro*, 4(5). <https://doi.org/10.1523/ENEURO.0287-17.2017>
- Cleland, T. A., & Linster, C. (2012). Frontiers | On-Center/Inhibitory-Surround Decorrelation via Intraglomerular Inhibition in the Olfactory Bulb Glomerular Layer. *Frontiers in Integrative Neuroscience*, 6. <https://doi.org/10.3389/fnint.2012.00005>
- Cohen, Y., Reuveni, I., Barkai, E., & Maroun, M. (2008). Olfactory Learning-Induced Long-Lasting Enhancement of Descending and Ascending Synaptic Transmission to the Piriform Cortex. *Journal of Neuroscience*, 28(26). <https://doi.org/10.1523/JNEUROSCI.0178-08.2008>
- Dalton, P. (2000). Psychophysical and Behavioral Characteristics of Olfactory Adaptation. *Chemical Senses*, 25(4). <https://doi.org/10.1093/chemse/25.4.487>
- Dasgupta, S., Stevens, C. F., & Navlakha, S. (2017). A neural algorithm for a fundamental computing problem. *Science*. <https://doi.org/10.1126/science.aam9868>
- Deitch, D., Rubin, A., & Ziv, Y. (2021). Representational drift in the mouse visual cortex. *Current Biology*, 31(19), 4327-4339.e4326. <https://doi.org/10.1016/j.cub.2021.07.062>

- Dietz, S. B., & Murthy, V. N. (2005). Contrasting short-term plasticity at two sides of the mitral–granule reciprocal synapse in the mammalian olfactory bulb. *The Journal of Physiology*, 569(Pt 2), 475-488. <https://doi.org/10.1113/jphysiol.2005.095844>
- Egger, V., & Kuner, T. (2021). Olfactory bulb granule cells: specialized to link coactive glomerular columns for percept generation and discrimination of odors. *Cell and Tissue Research*, 383(1), 495-506. <https://doi.org/10.1007/s00441-020-03402-7>
- Egger, V., & Urban, N. N. (2006). Dynamic connectivity in the mitral cell-granule cell microcircuit. *Seminars in Cell & Developmental Biology*, 17(4), 424-432. <https://doi.org/10.1016/j.semcdb.2006.04.006>
- Eppler, J.-B., Lai, T., Aschauer, D., Rumpel, S., & Kaschube, M. (2025). Representational drift reflects ongoing balancing of stochastic changes by Hebbian learning. In: bioRxiv.
- Fink, A. J. P., Muscinelli, S. P., Wang, S., Hogan, M. I., English, D. F., Axel, R., Litwin-Kumar, A., & Schoonover, C. E. (2025). Experience-dependent reorganization of inhibitory neuron synaptic connectivity. In: bioRxiv.
- Fletcher, M. L. (2011). Analytical Processing of Binary Mixture Information by Olfactory Bulb Glomeruli. *PLOS ONE*, 6(12), e29360. <https://doi.org/10.1371/journal.pone.0029360>
- Franks, K. M., Russo, M. J., Sosulski, D. L., Mulligan, A. A., Siegelbaum, S. A., & Axel, R. (2011). Recurrent circuitry dynamically shapes the activation of piriform cortex. *Neuron*, 72(1). <https://doi.org/10.1016/j.neuron.2011.08.020>
- Gadziola, M. A., Stetzik, L. A., Wright, K. N., Milton, A. J., Arakawa, K., Cortijo, M. d. M., & Wesson, D. W. (2020). A Neural System that Represents the Association of Odors with Rewarded Outcomes and Promotes Behavioral Engagement. *Cell reports*, 32(3). <https://doi.org/10.1016/j.celrep.2020.107919>
- Gao, Y., Strowbridge, B. W., Gao, Y., & Strowbridge, B. W. (2009). Long-term plasticity of excitatory inputs to granule cells in the rat olfactory bulb. *Nature Neuroscience* 2009 12:6, 12(6). <https://doi.org/10.1038/nn.2319>
- Giessel, A. J., & Datta, S. R. (2014). Olfactory maps, circuits and computations. *Current Opinion in Neurobiology*, 24. <https://doi.org/10.1016/j.conb.2013.09.010>
- Gilra, A., & Bhalla, U. S. (2015). Bulbar Microcircuit Model Predicts Connectivity and Roles of Interneurons in Odor Coding. *PLOS ONE*, 10(5), e0098045. <https://doi.org/10.1371/journal.pone.0098045>
- Hamdani, E. H., Lastein, S., Gregersen, F., & Døving, K. B. (2008). Seasonal Variations in Olfactory Sensory Neurons—Fish Sensitivity to Sex Pheromones Explained? *Chemical Senses*, 33(2). <https://doi.org/10.1093/chemse/bjm072>
- Hirata, T. (2024). Olfactory information processing viewed through mitral and tufted cell-specific channels. *Frontiers in Neural Circuits*, 18. <https://doi.org/10.3389/fncir.2024.1382626>
- Hiratani, N., & Latham, P. E. (2020). Rapid Bayesian learning in the mammalian olfactory system. *Nature Communications*, 11(1), 3845. <https://doi.org/10.1038/s41467-020-17490-0>
- Hong, E. J., & Wilson, R. I. (2015). Simultaneous encoding of odors by channels with diverse sensitivity to inhibition. *Neuron*, 85(3), 573-589. <https://doi.org/10.1016/j.neuron.2014.12.040>
- Jordan, C. (1875). Essai sur la géométrie à n dimensions. *Bulletin de la Société Mathématique de France*, 3, 103-174. <https://doi.org/10.24033/bsmf.90>
- Kato, H. K., Chu, M. W., Isaacson, J. S., & Komiyama, T. (2012). Dynamic Sensory Representations in the Olfactory Bulb: Modulation by Wakefulness and Experience. *Neuron*, 76(5), 962-975. <https://doi.org/10.1016/j.neuron.2012.09.037>
- Kersen, D. E. C., Tavoni, G., & Balasubramanian, V. (2022). Connectivity and dynamics in the olfactory bulb. *PLoS Computational Biology*, 18(2), e1009856. <https://doi.org/10.1371/journal.pcbi.1009856>

- Kim, D. H., Phillips, M. E., Chang, A. Y., Patel, H. K., Nguyen, K. T., & Willhite, D. C. (2011). Lateral Connectivity in the Olfactory Bulb is Sparse and Segregated. *Frontiers in Neural Circuits*, 5. <https://doi.org/10.3389/fncir.2011.00005>
- Kim, K., Bae, J., Jin, Y., Moon, C., Kim, K., Bae, J., Jin, Y., & Moon, C. (2020). Odor habituation can modulate very early olfactory event-related potential. *Scientific Reports* 2020 10:1, 10(1). <https://doi.org/10.1038/s41598-020-75263-7>
- Kong, E., Zabeh, E., Liao, Z., Mihaila, T. S., Wilson, C., Santhirasegaran, C., Peterka, D. S., Losonczy, A., & Geiller, T. (2024). Recurrent Connectivity Shapes Spatial Coding in Hippocampal CA3 Subregions. *bioRxiv*, 2024.2011.2007.622379. <https://doi.org/10.1101/2024.11.07.622379>
- Koulakov, A. A., & Rinberg, D. (2011). Sparse Incomplete Representations: A Potential Role of Olfactory Granule Cells. *Neuron*, 72(1), 124-136. <https://doi.org/10.1016/j.neuron.2011.07.031>
- Krishnamurthy, K., Hermundstad, A. M., Mora, T., Walczak, A. M., & Balasubramanian, V. (2022). Disorder and the Neural Representation of Complex Odors. *Frontiers in Computational Neuroscience*, 16. <https://doi.org/10.3389/fncom.2022.917786>
- Kumar, A., Barkai, E., & Schiller, J. (2021). Plasticity of olfactory bulb inputs mediated by dendritic NMDA-spikes in rodent piriform cortex. *eLife*, 10. <https://doi.org/10.7554/eLife.70383>
- Large, A. M., Vogler, N. W., Mielo, S., Oswald, A.-M. M., Large, A. M., Vogler, N. W., Mielo, S., & Oswald, A.-M. M. (2016). Balanced feedforward inhibition and dominant recurrent inhibition in olfactory cortex. *Proceedings of the National Academy of Sciences*, 113(8). <https://doi.org/10.1073/pnas.1519295113>
- Li, W., Luxenberg, E., Parrish, T., & Gottfried, J. A. (2006). Learning to Smell the Roses: Experience-Dependent Neural Plasticity in Human Piriform and Orbitofrontal Cortices. *Neuron*, 52(6). <https://doi.org/10.1016/j.neuron.2006.10.026>
- Li, W. L., Chu, M. W., Wu, A., Suzuki, Y., Imayoshi, I., & Komiyama, T. (2018). Adult-born neurons facilitate olfactory bulb pattern separation during task engagement. *eLife*, 7, e33006. <https://doi.org/10.7554/eLife.33006>
- Linster, C., & Cleland, T. A. (2010). Decorrelation of Odor Representations via Spike Timing-Dependent Plasticity. *Frontiers in Computational Neuroscience*, 4. <https://doi.org/10.3389/fncom.2010.00157>
- Lledo, P.-M., & Valley, M. (2016). Adult Olfactory Bulb Neurogenesis. *Cold Spring Harbor Perspectives in Biology*, 8(8). <https://doi.org/10.1101/cshperspect.a018945>
- Malnic, B., Hirono, J., Sato, T., & Buck, L. B. (1999). Combinatorial receptor codes for odors. *Cell*, 96(5), 713-723. [https://doi.org/10.1016/s0092-8674\(00\)80581-4](https://doi.org/10.1016/s0092-8674(00)80581-4)
- Marks, T. D., & Goard, M. J. (2021). Stimulus-dependent representational drift in primary visual cortex. *Nature Communications*, 12(1), 5169. <https://doi.org/10.1038/s41467-021-25436-3>
- Masset, P., Qin, S., & Zavatone-Veth, J. A. (2022). Drifting neuronal representations: Bug or feature? *Biological Cybernetics*, 116(3), 253-266. <https://doi.org/10.1007/s00422-021-00916-3>
- Meng, J. H., & Riecke, H. (2022). Structural spine plasticity: Learning and forgetting of odor-specific subnetworks in the olfactory bulb. *PLoS Computational Biology*, 18(10), e1010338. <https://doi.org/10.1371/journal.pcbi.1010338>
- Migliore, M., Cavarretta, F., Marasco, A., Tulumello, E., Hines, M. L., & Shepherd, G. M. (2015). Synaptic clusters function as odor operators in the olfactory bulb. *Proceedings of the National Academy of Sciences*, 112(27), 8499-8504. <https://doi.org/10.1073/pnas.1502513112>
- Migliore, M., & Shepherd, G. M. (2008). Dendritic action potentials connect distributed dendrodendritic microcircuits. *Journal of computational neuroscience*, 24(2), 207-221. <https://doi.org/10.1007/s10827-007-0051-9>

- Mombaerts, P. (2006). Axonal Wiring in the Mouse Olfactory System. *Annual Review of Cell and Developmental Biology*, 22(Volume 22, 2006). <https://doi.org/10.1146/annurev.cellbio.21.012804.093915>
- Mombaerts, P., Wang, F., Dulac, C., Chao, S. K., Nemes, A., Mendelsohn, M., Edmondson, J., & Axel, R. (1996). Visualizing an Olfactory Sensory Map. *Cell*, 87(4). [https://doi.org/10.1016/S0092-8674\(00\)81387-2](https://doi.org/10.1016/S0092-8674(00)81387-2)
- Morales, G. B., Muñoz, M. A., & Tu, Y. (2025). Representational drift and learning-induced stabilization in the piriform cortex. *Proceedings of the National Academy of Sciences*, 122(29), e2501811122. <https://doi.org/10.1073/pnas.2501811122>
- Nagayama, S., Homma, R., & Imamura, F. (2014). Neuronal organization of olfactory bulb circuits. *Frontiers in Neural Circuits*, 8. <https://doi.org/10.3389/fncir.2014.00098>
- Nissant, A., Bardy, C., Katagiri, H., Murray, K., Lledo, P.-M., Nissant, A., Bardy, C., Katagiri, H., Murray, K., & Lledo, P.-M. (2009). Adult neurogenesis promotes synaptic plasticity in the olfactory bulb. *Nature Neuroscience* 2009 12:6, 12(6). <https://doi.org/10.1038/nn.2298>
- Otazu, G. H., Chae, H., Davis, M. B., & Albeanu, D. F. (2015). Cortical feedback decorrelates olfactory bulb output in awake mice. *Neuron*, 86(6), 1461-1477. [https://www.cell.com/neuron/fulltext/S0896-6273\(15\)00430-4](https://www.cell.com/neuron/fulltext/S0896-6273(15)00430-4)
- Pashakhanloo, F., & Koulakov, A. (2023). Stochastic gradient descent-induced drift of representation in a two-layer neural network | Proceedings of the 40th International Conference on Machine Learning. *Proceedings of the 40th International Conference on Machine Learning, PMLR*. <https://doi.org/10.5555/3618408.3619547>
- Patterson, M. A., Lagier, S., & Carleton, A. (2013). Odor representations in the olfactory bulb evolve after the first breath and persist as an odor afterimage. *Proceedings of the National Academy of Sciences*, 110(35), E3340-E3349. <https://doi.org/10.1073/pnas.1303873110>
- Qin, S., Farashahi, S., Lipshutz, D., Sengupta, A. M., Chklovskii, D. B., & Pehlevan, C. (2023). Coordinated drift of receptive fields in Hebbian/anti-Hebbian network models during noisy representation learning. *Nature Neuroscience*, 26(2), 339-349. <https://doi.org/10.1038/s41593-022-01225-z>
- Rall, W., Shepherd, G. M., Reese, T. S., & Brightman, M. W. (1966). Dendrodendritic synaptic pathway for inhibition in the olfactory bulb. *Experimental Neurology*, 14(1). [https://doi.org/10.1016/0014-4886\(66\)90023-9](https://doi.org/10.1016/0014-4886(66)90023-9)
- Rule, M. E., O'Leary, T., & Harvey, C. D. (2019). Causes and consequences of representational drift. *Current Opinion in Neurobiology*, 58, 141-147. <https://doi.org/10.1016/j.conb.2019.08.005>
- Satou, M., Anzai, S., Huruno, M., Satou, M., Anzai, S., & Huruno, M. (2005). Long-term potentiation and olfactory memory formation in the carp (*Cyprinus carpio* L.) olfactory bulb. *Journal of Comparative Physiology A* 2005 191:5, 191(5). <https://doi.org/10.1007/s00359-005-0600-5>
- Schoonover, C. E., Ohashi, S. N., Axel, R., & Fink, A. J. P. (2021). Representational drift in primary olfactory cortex. *Nature*, 594(7864), 541-546. <https://doi.org/10.1038/s41586-021-03628-7>
- Shani-Narkiss, H., Beniaguev, D., Segev, I., & Mizrahi, A. (2023). Frontiers | Stability and flexibility of odor representations in the mouse olfactory bulb. *Frontiers in Neural Circuits*, 17. <https://doi.org/10.3389/fncir.2023.1157259>
- Shani-Narkiss, H., Vinograd, A., Landau, I. D., Tasaka, G., Yayon, N., Terletsy, S., Groysman, M., Maor, I., Sompolinsky, H., Mizrahi, A., Shani-Narkiss, H., Vinograd, A., Landau, I. D., Tasaka, G., Yayon, N., Terletsy, S., Groysman, M., Maor, I., Sompolinsky, H., & Mizrahi, A. (2020). Young adult-born neurons improve odor coding by mitral cells. *Nature Communications* 2020 11:1, 11(1). <https://doi.org/10.1038/s41467-020-19472-8>
- Shen, Y., Dasgupta, S., Navlakha, S., Shen, Y., Dasgupta, S., & Navlakha, S. (2020). Habituation as a neural algorithm for online odor discrimination. *Proceedings of the National Academy of Sciences*, 117(22). <https://doi.org/10.1073/pnas.1915252117>

- Shepherd, G. M. (2003). *The synaptic organization of the brain*. Oxford university press.
- Shepherd, G. M., Chen, W. R., Willhite, D., Migliore, M., & Greer, C. A. (2007). The olfactory granule cell: From classical enigma to central role in olfactory processing. *Brain Research Reviews*, 55(2). <https://doi.org/10.1016/j.brainresrev.2007.03.005>
- Teşileanu, T., Cocco, S., Monasson, R., & Balasubramanian, V. (2019). Adaptation of olfactory receptor abundances for efficient coding. *eLife*, 8. <https://doi.org/10.7554/eLife.39279>
- Vaaga, C. E., Yorgason, J. T., Williams, J. T., & Westbrook, G. L. (2017). Presynaptic gain control by endogenous cotransmission of dopamine and GABA in the olfactory bulb. *Journal of Neurophysiology*, 117(3), 1163-1170. <https://doi.org/10.1152/jn.00694.2016>
- Wang, F., Nemes, A., Mendelsohn, M., & Axel, R. (1998). Odorant Receptors Govern the Formation of a Precise Topographic Map. *Cell*, 93(1). [https://doi.org/10.1016/S0092-8674\(00\)81145-9](https://doi.org/10.1016/S0092-8674(00)81145-9)
- Whitman, M. C., & Greer, C. A. (2007). Synaptic Integration of Adult-Generated Olfactory Bulb Granule Cells: Basal Axodendritic Centrifugal Input Precedes Apical Dendrodendritic Local Circuits. *Journal of Neuroscience*, 27(37). <https://doi.org/10.1523/JNEUROSCI.1633-07.2007>
- Wick, S. D., Wiechert, M. T., Friedrich, R. W., & Riecke, H. (2010). Pattern orthogonalization via channel decorrelation by adaptive networks. *Journal of computational neuroscience*, 28(1), 29-45. <https://doi.org/10.1007/s10827-009-0183-1>
- Wiechert, M. T., Judkewitz, B., Riecke, H., & Friedrich, R. W. (2010). Mechanisms of pattern decorrelation by recurrent neuronal circuits. *Nature Neuroscience*, 13(8), 1003-1010. <https://doi.org/10.1038/nn.2591>
- Willhite, D. C., Nguyen, K. T., Masurkar, A. V., Greer, C. A., Shepherd, G. M., & Chen, W. R. (2006). Viral tracing identifies distributed columnar organization in the olfactory bulb. *Proceedings of the National Academy of Sciences of the United States of America*, 103(33), 12592-12597. <https://doi.org/10.1073/pnas.0602032103>
- Wilson, D. A. (2000). Comparison of Odor Receptive Field Plasticity in the Rat Olfactory Bulb and Anterior Piriform Cortex. *Journal of Neurophysiology*, 84(6). <https://doi.org/10.1152/jn.2000.84.6.3036>
- Wilson, D. A., Kadohisa, M., & Fletcher, M. L. (2006). Cortical contributions to olfaction: Plasticity and perception. *Seminars in Cell & Developmental Biology*, 17(4). <https://doi.org/10.1016/j.semcdb.2006.04.008>
- Wu, A., Yu, B., Chen, Q., Matthews, G. A., Lu, C., Campbell, E., Tye, K. M., & Komiyama, T. (2020). Context-dependent plasticity of adult-born neurons regulated by cortical feedback. *Science Advances*. <https://doi.org/10.1126/sciadv.abc8319>
- Yamada, Y., Bhaukaurally, K., Madarász, T. J., Pouget, A., Rodriguez, I., & Carleton, A. (2017). Context- and Output Layer-Dependent Long-Term Ensemble Plasticity in a Sensory Circuit. *Neuron*, 93(5). <https://doi.org/10.1016/j.neuron.2017.02.006>
- Yokoi, M., Mori, K., & Nakanishi, S. (1995). Refinement of odor molecule tuning by dendrodendritic synaptic inhibition in the olfactory bulb. *Proceedings of the National Academy of Sciences of the United States of America*, 92(8). <https://doi.org/10.1073/pnas.92.8.3371>
- Yu, B., Yue, Y., Ren, C., Yun, R., Lim, B., & Komiyama, T. (2025). Cholinergic feedback for modality- and context-specific modulation of sensory representations. *Science*, 388(6753). <https://doi.org/10.1126/science.ads9152>
- Zabeh, E., Jacobs, J., Losonczy, A., & Kong, E. (2025). The recurrency level is a key determinant of representational drift. Computational and Systems Neuroscience (COSYNE) 2025, Montreal, QC, Canada.
- Zavatone-Veth, J. A., Masset, P., Tong, W. L., Zak, J., Murthy, V. N., & Pehlevan, C. (2023, 2023/11/02/). Neural Circuits for Fast Poisson Compressed Sensing in the Olfactory Bulb. Thirty-seventh Conference on Neural Information Processing Systems,

- Zavitz, D., Youngstrom, I. A., Borisyuk, A., & Wachowiak, M. (2020). Effect of Interglomerular Inhibitory Networks on Olfactory Bulb Odor Representations. *The Journal of Neuroscience*, *40*(31), 5954-5969. <https://doi.org/10.1523/JNEUROSCI.0233-20.2020>
- Zhong, L., Baptista, S., Gattoni, R., Arnold, J., Flickinger, D., Stringer, C., & Pachitariu, M. (2025). Unsupervised pretraining in biological neural networks. *Nature*, 1-8. <https://doi.org/10.1038/s41586-025-09180-y>
- Zhou, F.-W., Shao, Z.-Y., Shipley, M. T., & Puche, A. C. (2020). Short-term plasticity in glomerular inhibitory circuits shapes olfactory bulb output. *Journal of Neurophysiology*, *123*(3), 1120-1132. <https://doi.org/10.1152/jn.00628.2019>

Supplementary Tables

Table S1

	Odor Pair Study			
	Cohort 1		Cohort 2	
Odor Name	Odor 1	Odor 2	Mixture 1	Mixture 2
Content	Heptenal	Ethyl Tiglate	52:48	48:52
τ	0.11	0.19	0.17	0.19
C	46.21	19.44	29.28	17.34

	8-odor Panel Study							
	Cohort 3							
Odor Name	Mixture 1	Mixture 2	Mixture 3	Mixture 4	Mixture 5	Mixture 6	Mixture 7	Mixture 8
Content	54:46	52:48	51.5:48.5	51:49	49:51	48.5:51.5	48:52	46:54
τ	0.34	0.46	0.28	0.40	0.54	0.58	0.92	0.93
C	168.71	163.89	159.73	167.22	177.57	168.42	142.75	137.25

Supplementary Table 1. Mean field decay fit.

The average mean-field fitted parameters τ and C for each cohort exposed to each odor.

Table S2

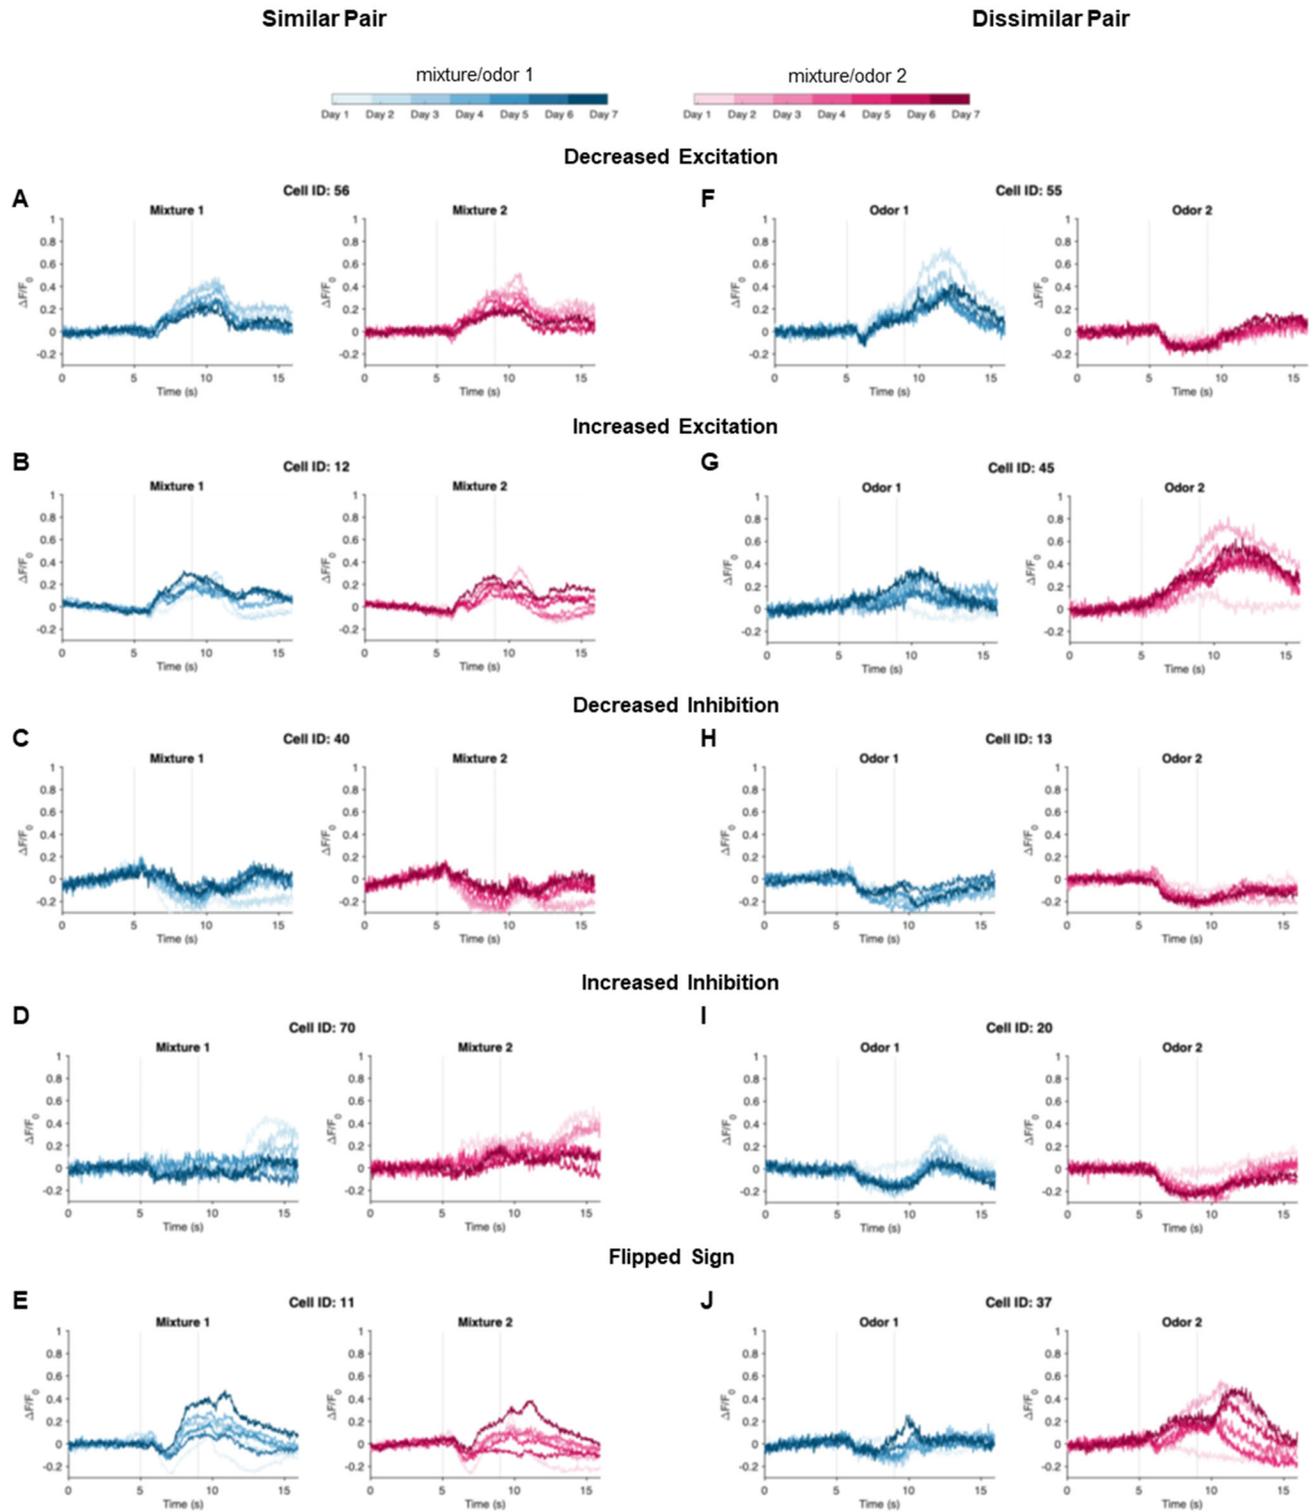
	M	N	θ	η	λ	ϵ_{ind}	ϵ_{sh}	σ	κ
Fig. 3D	10	750	0.5	0.05	0.0001	0	0	0	/
Fig. 3E	10	750	0.5	0.05	0.0001	0.1	0	0	/
Fig. 3G	10	750	0.5	0.05	0.0001	0	0.1	0	/
Fig. 4C	40	3000	0.5	0.05	0.001	0	0	0	/
Fig. 4F	40	3000	0.5	0.05	0.001	0	0	0	300
Fig. 4I	40	3000	0.5	0.05	0.001	0	0	0	100
Fig. 6	20	1500	0.5	0.1	0.012	0	0	0	300
Supp. Fig. 3C	10	750	0.5	0.05	0.0001	0.05	0	0	/
Supp. Fig. 3D	10	750	0.5	0.05	0.0001	0	0.05	0	/
Supp. Fig. 4C	10	750	0.5	0.05	0.0001	0	0	0.01	/
Supp. Fig. 5A	40	3000	0.5	0.05	0.0045	0	0	0	300
Supp. Fig. 5C	40	3000	0.5	0.05	0.0045	0.1	0	0	/
Supp. Fig. 5E	40	3000	0.5	0.05	0.0045	0.1	0	0	300

Supplementary Table 2. Circuit simulation parameters.

M : Number of mitral cells (MCs) in the network. N : Number of granule cells (GCs) in the network. θ : Linear gain parameter relating GC membrane voltage to firing activity. η : Hebbian learning rate controlling the magnitude of synaptic updates. λ : Synaptic decay (turnover) rate enforcing weight regularization. ϵ_{ind} : Independent input noise magnitude; Gaussian white noise applied independently to each neuron, with variance proportional to input strength. ϵ_{sh} : Shared input noise magnitude; Gaussian white noise shared between inputs, with variance proportional to input strength. σ : Synaptic noise amplitude in Hebbian updates; Gaussian white noise with standard deviation calibrated to be on the same order of magnitude as the average absolute deterministic Hebbian update. κ : in block constraint, GC block size (columns per block); MC block size set proportionally. In band constraint, represent half-width (in GC indices) around each row's center, with periodic wrap-around.

Supplementary Figures

Fig. S1



Supplementary Figure 1. Individual mitral cells show diverse response dynamics during learning, despite population-level gain adaption.

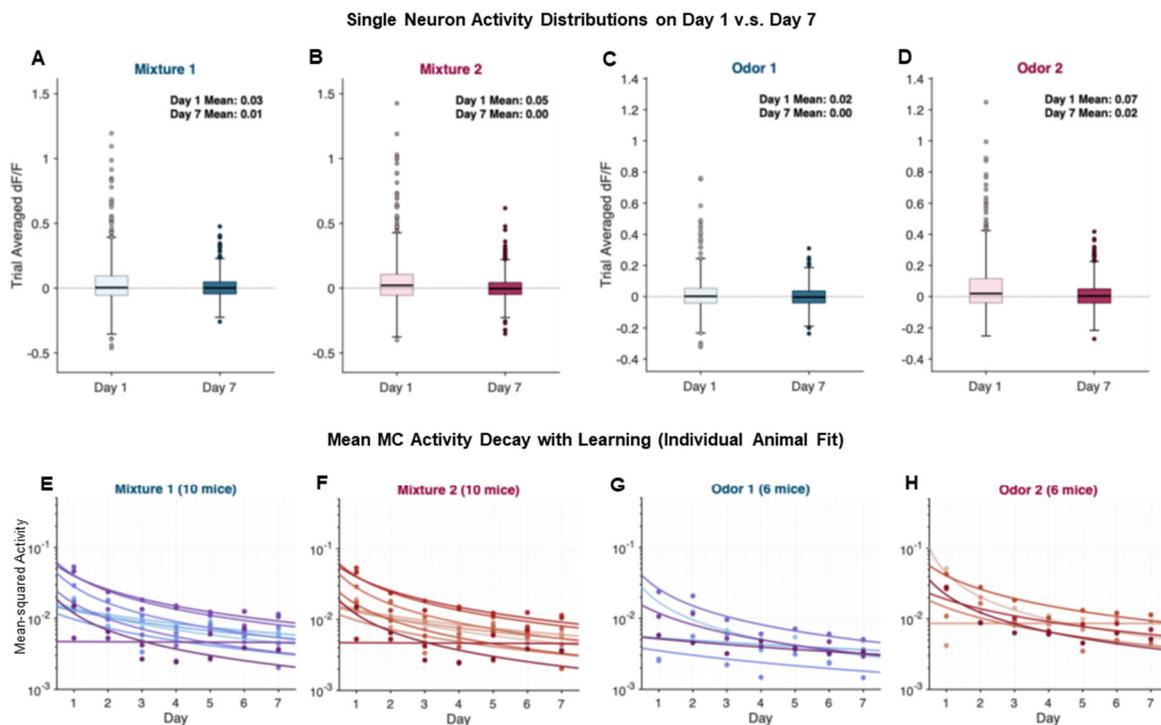
Trial-averaged dF/F traces are shown for individual neurons across seven days of training. Dotted lines indicate the 4-second odor presentation period.

(A-E) Similar odorant pair during passive exposure. Representative examples of different plasticity patterns observed in individual mitral cells from one mouse. Neurons were selected to illustrate the range of response changes for Odor 1 during the odorant delivery period. Note that responses to Odor 2 may follow different temporal dynamics than Odor 1 responses in the same cell. (A) Decreased excitation: progressive reduction in excitatory responses over training days. (B) Increased excitation: enhanced excitatory responses over training days. (C) Decreased inhibition: reduction in inhibitory responses over training days. (D) Increased inhibition: enhanced inhibitory responses over training days. (E) Sign reversal: neuron switching from inhibitory to excitatory responses over training days.

(F-J) Dissimilar odorant pair. Representative examples showing diverse activity pattern changes as in panels (A-E).

These examples demonstrate that while the population shows overall decreased activity (Fig. 2), individual neurons exhibit heterogeneous response dynamics that cannot be predicted solely from population-level changes.

Fig. S2

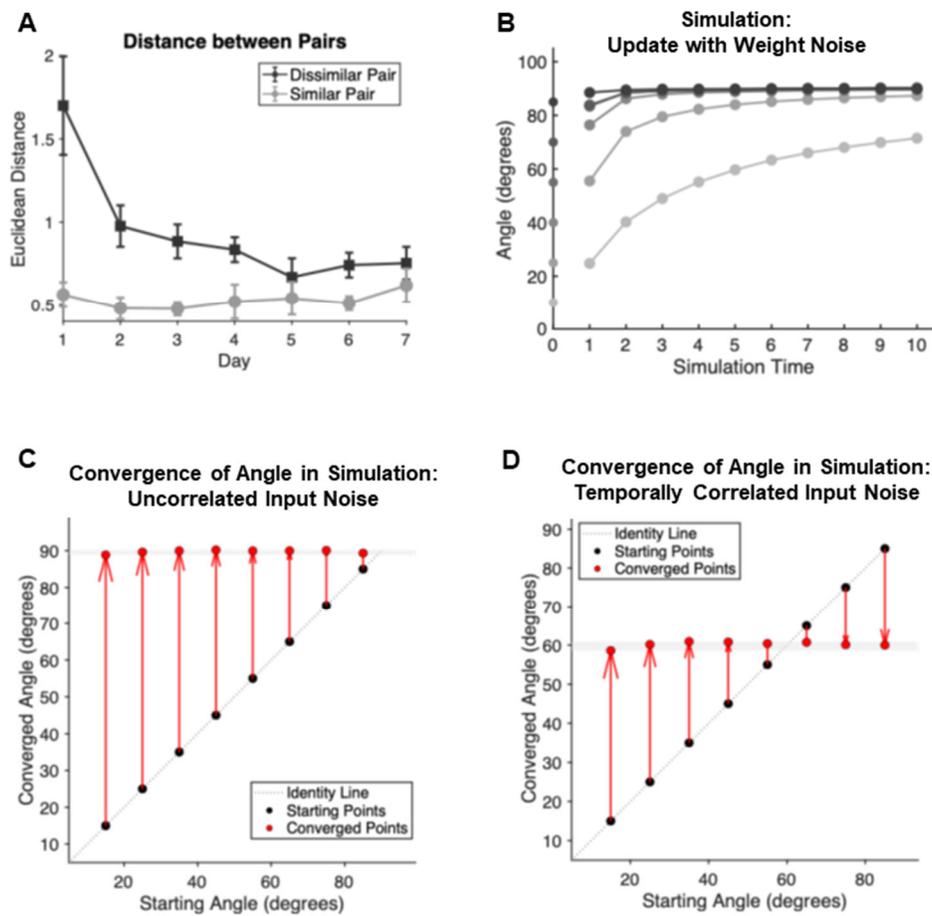


Supplementary Figure 2. Population activity decay manifests as reduced dynamics on a population level, and the trend is consistent over mice.

(A-D) Single-neuron activity distributions on Day 1 versus Day 7. Trial-averaged dF/F responses for individual neurons pooled across all animals exposed to each odor condition. Box plots show median, 25th and 75th percentiles; whiskers extend to 2 \times interquartile range. Note that mean activity remains near zero for all odorants, with neurons showing both excitation and inhibition relative to baseline, but the overall dynamic range decreases substantially by Day 7.

(E-H) Individual animal fits to mean-field decay model. Each color represents data from one animal. Points show daily mean-squared population activity; lines show best-fit power-law decay functions.

Fig. S3



Supplementary Figure 3. Angular dynamics depend on input similarity and noise correlation structure.

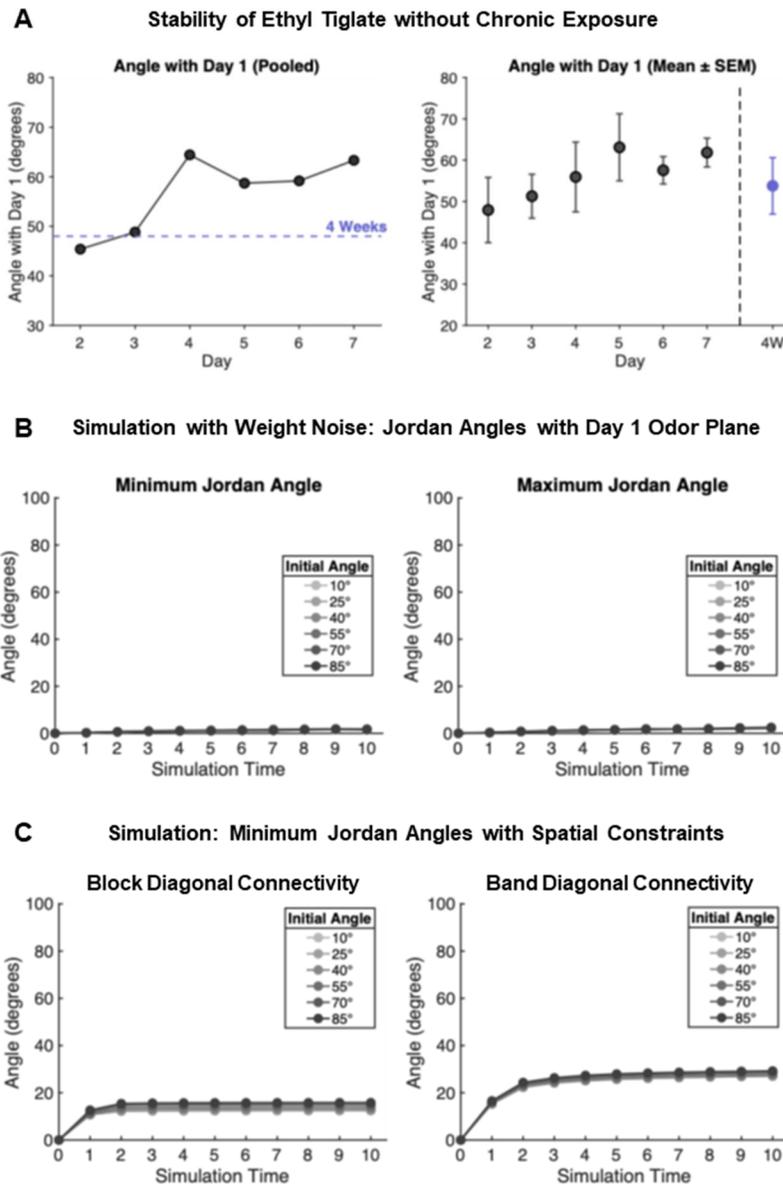
(A) Between-odor Euclidean distance changes reveal similarity-dependent dynamics. During passive exposure, Euclidean distance between population vectors for dissimilar odor pairs (dark) decreases substantially over days, while distance between similar odor pairs (light) remains relatively stable despite overall activity decay (see Fig. 2, Fig. S2). Data show mean \pm SEM across animals.

(B) Weight update noise does not affect decorrelation dynamics. Simulation with random normal noise added to Hebbian weight updates (magnitude comparable) shows that decorrelation toward complete orthogonality remains unchanged across all initial angles.

(C-D) Input noise correlation determines convergence angle. Simulations comparing uncorrelated versus temporally correlated Gaussian input noise (5% of input variance). (C) With uncorrelated noise, all odor pairs converge to complete orthogonality (90°) regardless of starting angle. (D) With temporally correlated noise, pairs converge to approximately 60° , matching

experimental observations. Eight different starting angles tested (15°-85°). Final angles measured as average of last 50 trials.

Fig. S4



Supplementary Figure 4. Representational drift requires chronic exposure and emerges from structural connectivity constraints.

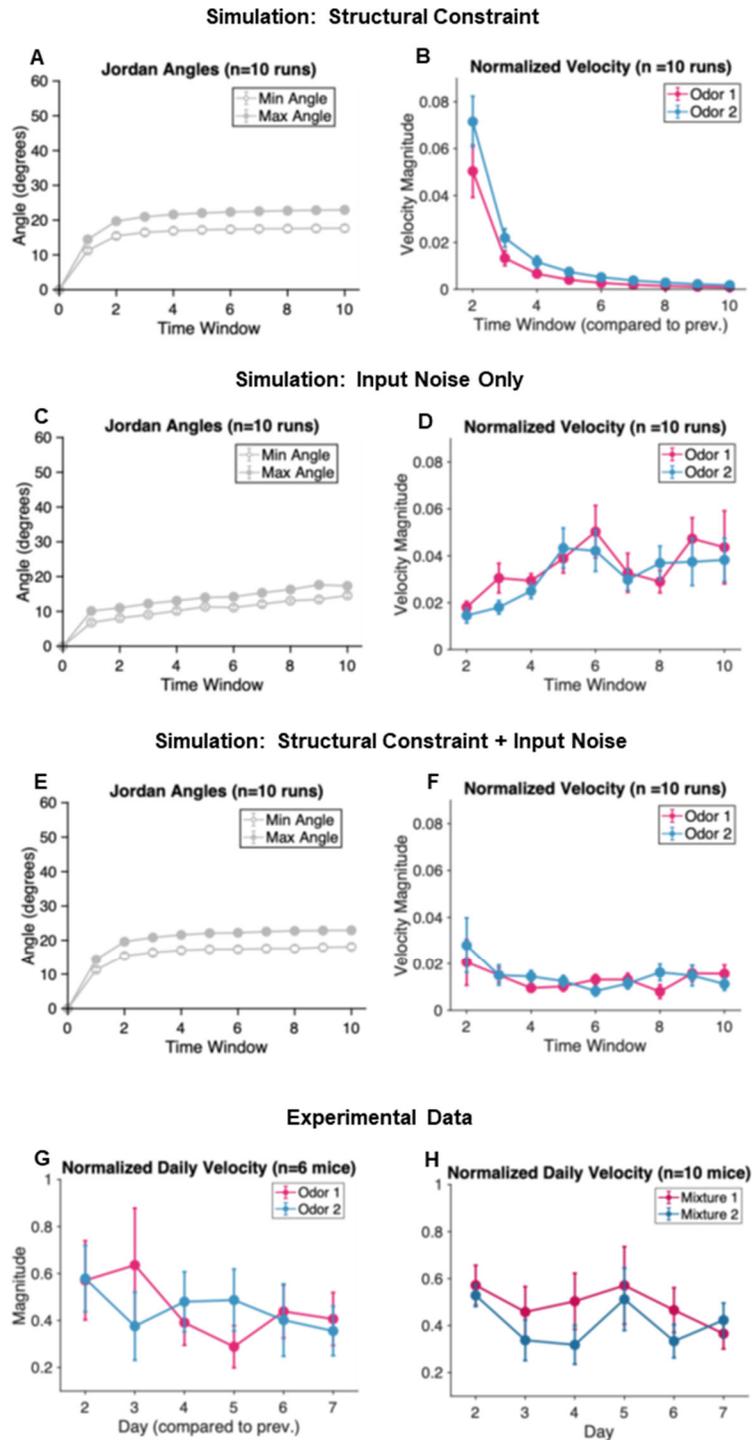
(A) Drift magnitude from chronic exposures exceeds baseline variability from minimal exposure. Angular change of Ethyl Tiglate (referred to as odor 2) representations relative to Day 1, compared with data from an independent study that also exposed animals to this odor at 4-week intervals with no intervening exposure (Shani-Narkiss et al., 2023; Shani-Narkiss et al., 2020). Left: pooled population vectors (N = 217 neurons, 6 mice for chronic exposure; N = 294 neurons, 7 mice for 4-week study). Right: daily angles of odor representation vectors of individual animals to Day 1 (Mean \pm SEM). The comparison shows that chronic exposure

produces substantially larger angles to initial exposure than the 4-week baseline, demonstrating that drift is exposure-dependent rather than spontaneous.

(B) Weight update noise alone does not produce subspace rotation. Simulation results for fully connected MC-GC networks with random normal weight noise (magnitude comparable to learning-driven changes). No Jordan angle rotation is observed.

(C) Structured connectivity produces drift in both minimum and maximum Jordan angles. Minimum Jordan angles for block-diagonal (left) and band-diagonal (right) connectivity patterns, using identical parameters as Fig. 4G and 4J. Both connectivity structures produce changes in minimum Jordan angles, demonstrating that the drift phenomenon affects multiple dimensions of the representational subspace.

Fig. S5



Supplementary Figure 5. Distinguishing contributions of structural connectivity constraints versus input noise to representational drift.

(A-F) Simulation analysis of drift mechanisms using normalized odor velocity. To isolate odor-specific drift effects from global subspace changes, we analyzed normalized velocity, the magnitude of change between consecutive time windows (normalized to remove decay effects from gain adaptation). Simulations used identical parameters except for the varied factors. Trial averages were computed by averaging over 100-trial windows; $n=10$ runs with different random seeds. Simulation data show mean \pm SEM over random seeds.

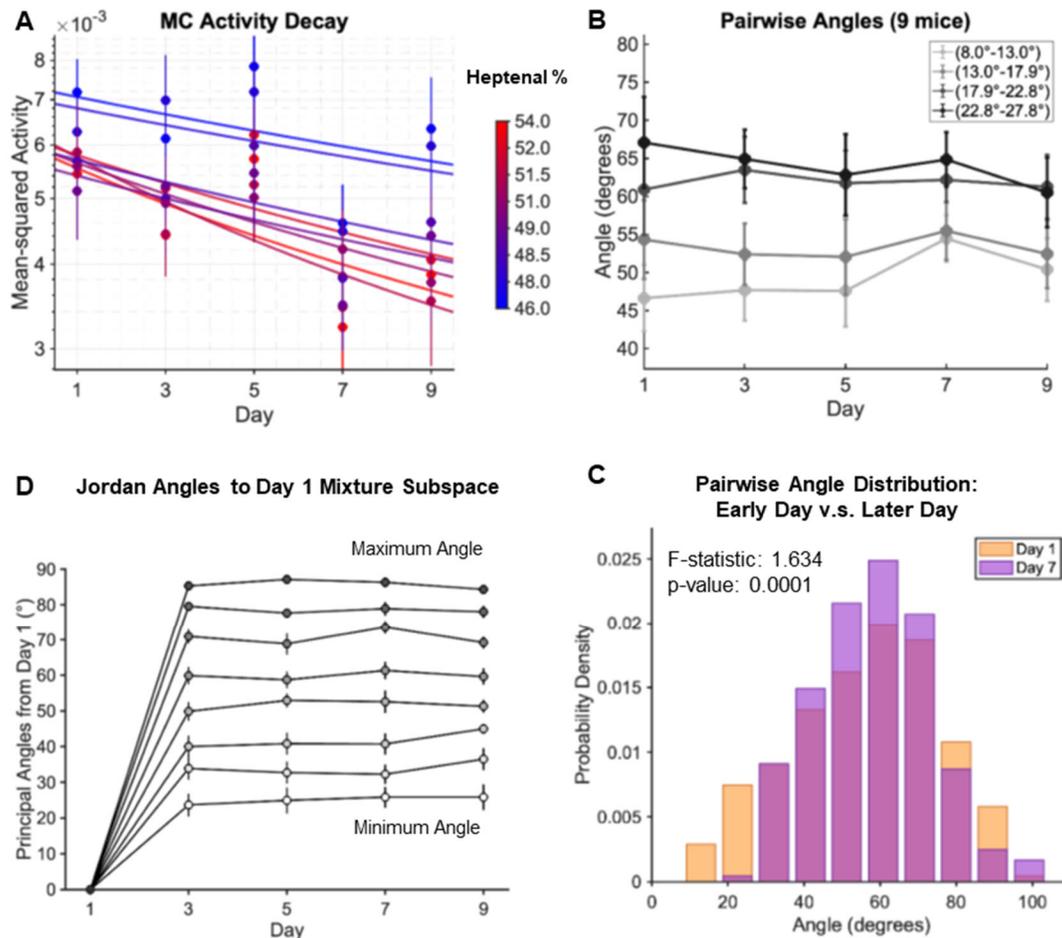
(A-B) Structural constraint alone. Block-diagonal connectivity produces substantial Jordan angle drift (A) with high initial velocity that rapidly decreases (B).

(C-D) Input noise alone. Input noise (10% variance) produces modest Jordan angle changes (C) but increasing velocity over time (D), as noise overwhelms diminishing residual changes due to input dependent plasticity.

(E-F) Combined structural constraint and input noise. Both factors together produce Jordan angle drift (E) with velocity dynamics (F) that show initial higher velocity followed by gradual decrease.

(G-H) Experimental validation. Normalized daily velocity from experimental data shows patterns most consistent with combined mechanisms. (G) Dissimilar odor pairs. (H) Similar odor pairs. Data show mean \pm SEM across animals.

Fig. S6



Supplementary Figure 6. Neural representations for binary mixture panel exhibits similar representational changes as odor pairs during passive exposure.

(A) Mean-field decay model fits for individual mixture responses. Mean-squared population activity (\pm SEM) across all mice for each of the 8 binary mixtures, color-coded by Heptanal percentage. Solid lines show best-fit power-law decay functions from our mean-field model, demonstrating that gain adaptation occurs consistently across the concentration gradient.

(B) Pairwise angle changes depend on initial separation. Angular changes between mixture representations, binned by initial Day 1 angles into 4 equally spaced intervals (containing 4, 7, 8, and 9 pairs respectively). Data show mean \pm SEM across animals. Pairs with larger initial angles tend to decrease (convergence), while pairs with smaller initial angles tend to increase (separation), consistent with convergence toward an intermediate angle.

(C) Distribution of pairwise angles narrows over time. Histogram comparing all pairwise angles between the 8 mixtures on Day 1 (orange) versus Day 7 (purple) across all animals (28 pairs \times 9 animals = 252 data points). The distribution becomes more concentrated around 50-60°, supporting the theoretical prediction of angle convergence. F-statistic and p-value from two-sample F-test for equal variances.

(D) Jordan angles quantify subspace drift for mixture representations.

Supplemental Information

Contents

1	Learning in the GC-MC circuit	2
1.1	Separation of Time Scale	3
2	Mean-field Decay	4
3	Learning on Single Odor Input	7
3.1	Residual m Aligns to r	7
3.2	Presence of Weight Noise	9
3.2.1	Approximating $\mathbb{E}[\delta m]$	11
3.3	Presence of Input Noise	12
3.4	Presence of Structural Connectivity	13
3.4.1	Block-diagonal Connectivity	14
3.4.2	Band-diagonal Connectivity	16
4	Learning on Odor Pair	18
4.1	The Decorrelation Motif	18
4.2	Temporally correlated noise	20

1 Learning in the GC-MC circuit

Our model is based on the following equations describing the responses of MCs $m \in \mathbb{R}^M$, and responses of GCs $a \in \mathbb{R}^N$:

$$\dot{u} = -u + W^T m \quad (1)$$

$$m = r - W a \quad (2)$$

with

$$a = f(u) \quad (3)$$

where $u \in \mathbb{R}^N$ is the membrane voltage of the GCs. Here, $W \in \mathbb{R}^{M \times N}$ is the connectivity matrix from GCs to MCs, mediating local inhibitory feedback, r is input from the olfactory sensory neurons (OSNs). Since the connection is via dendrodendritic synapses, excitation from MCs to GCs occurs through W^T , and we assume equal synaptic strengths for simplicity. For tractability, we consider linear activation $f(u) = \theta u$, θ being a shared scalar parameter for all GCs, and we assume instantaneous MC responses.

The differential equations (2)-(3) admit a Lyapunov function:

$$\mathcal{L} = \frac{1}{2} \|r - W a\|_2^2 + \frac{1}{2\theta} \|a\|_2^2. \quad (4)$$

as the temporal behavior of the system can be viewed as gradient descent, i.e., $\dot{u} \propto -\partial\mathcal{L}/\partial a$, and the time derivative of the Lyapunov function is always negative:

$$\dot{\mathcal{L}} = \frac{\partial\mathcal{L}}{\partial a} \cdot \dot{a} = -\frac{1}{\theta} \|\dot{a}\|_2^2 \leq 0.$$

Inspecting the Lyapunov function, the first term describes the *error* in the representation of the inputs r by the GCs, and minimization of this difference leads to the most accurate approximation of the receptor neuron activity by the GCs. The second term of the Lyapunov function represents a regularization cost imposed on the GC firing, which can be seen as a constraint favoring small activation magnitudes. Therefore, the dynamical system can be thought of as performing regularized linear reconstruction, where the GCs learn a representation a that reconstructs the input r using the fixed basis W while penalizing large activations, and the output m to the cortex is the residual error between the input and its reconstruction.

The same Lyapunov function also provides a natural framework for synaptic plasticity. Taking the gradient of \mathcal{L} with respect to the connectivity matrix W yields:

$$-\frac{\partial\mathcal{L}}{\partial W} = m a^T$$

This gradient gives rise to the classic Hebbian learning rule, where synaptic strength increases when both pre-synaptic GC activity (a) and post-synaptic MC activity (m) are simultaneously active.

Remarkably, this learning rule emerges naturally from the same optimization principle that governs neural dynamics, suggesting that both inference (finding appropriate representations) and learning (adapting synaptic weights) can be understood as minimizing the same objective function. The resulting plasticity rule is both i) local, requiring only information available at individual synapses, and ii) unsupervised, making it biologically plausible while providing a principled approach to dictionary learning in the GC-MC circuit.

1.1 Separation of Time Scale

We study this system assuming fast dynamics in the bulb and slow synaptic plasticity, which occurs when the animal is repeatedly exposed to odor inputs.

First, with every inhalation, the dynamics quickly reach equilibrium, and a^* and m^* can be solved to take the form:

$$a^* = \theta(I + \theta W^T W)^{-1} W^T r \quad (5)$$

$$m^* = (I + \theta W W^T)^{-1} r \quad (6)$$

Then, the synaptic weights update incrementally according to $\Delta W = \eta m^* a^{*T} - \lambda W$, where η is learning rate, and λ is weight decay (turnover).

Note that we do not constrain the sign of W 's elements. While biological synapses are non-negative, our analysis focuses on the mathematical structure where $W W^T$ symmetry is preserved, which holds regardless of W 's positivity constraint.

2 Mean-field Decay

We consider a simplified network model to understand the decay of mitral cell population activity over learning. The model consists of M MCs with input vector $\mathbf{r} = [r_1, r_2, \dots, r_M]^T$. The MCs are connected to a single granule cell to capture the inhibitory dynamics while maintaining analytical tractability. We used vector bold font in this section for clarity of vectorized analysis.

The simplified network is governed by the following equations:

1. **GC activities:** $a = \sum_{i=1}^M w_i m_i$
2. **MC activities:** $m_i = r_i - w_i a$
3. **Hebbian weight update:** $\tau \dot{w}_i = \frac{1}{M} m_i a$

where τ is the synaptic time constant, and w_i are the dendrodendritic synaptic weights. Scaling the learning rate by $1/M$ ensures that larger populations do not learn faster just because they have more neurons contributing to the weight updates, and the normalization maintains consistent dynamics across population sizes.

Substituting the GC equation into the MC equation:

$$m_i = r_i - w_i \sum_{j=1}^M w_j m_j$$

We can solve for the self-consistent solution for a by summing over i for $m_i = r_i - w_i a$ and get:

$$a = \sum_{j=1}^M w_j (r_j - w_j a)$$

Solving it for a we obtain:

$$a = \frac{\sum_{j=1}^M w_j r_j}{1 + \sum_{j=1}^M w_j^2} \quad (7)$$

Weight dynamics decomposition. We can decompose the weight vector as:

$$\mathbf{w}(t) = \alpha(t) \mathbf{r} + \mathbf{w}_\perp(t) \quad (8)$$

where $\alpha(t) = \frac{\mathbf{w} \cdot \mathbf{r}}{\|\mathbf{r}\|^2}$ is the projection coefficient onto the input direction, and $\mathbf{w}_\perp \cdot \mathbf{r} = 0$. We evolve the evolution for the parallel and perpendicular components next. The goal is to show that the orthogonal direction would diminish over time, and then weight dynamics, and MC activity can be expressed using scalar dimension $\alpha(t)$.

Parallel component evolution. Since $\dot{\mathbf{w}} = \frac{1}{\tau M} \mathbf{m} a = \frac{a}{\tau M} (\mathbf{r} - a \mathbf{w})$,

$$\mathbf{r} \cdot \dot{\mathbf{w}} = \frac{a}{\tau M} (\|\mathbf{r}\|^2 - a \mathbf{r} \cdot \mathbf{w})$$

With orthogonality constraint, we also have:

$$\mathbf{r} \cdot \dot{\mathbf{w}} = \mathbf{r} \cdot (\dot{\alpha}\mathbf{r} + \mathbf{w}) = \|\mathbf{r}\|^2 \dot{\alpha} + 0 = \dot{\alpha} \|\mathbf{r}\|^2$$

With the above two equations for $\mathbf{r} \cdot \dot{\mathbf{w}}$, we can readily derive:

$$\dot{\alpha} = \frac{a}{\tau M} (1 - a\alpha) \quad (9)$$

Perpendicular component evolution. From $\dot{\mathbf{w}}_{\perp} = \dot{\mathbf{w}} - \dot{\alpha}\mathbf{r}$, and the derived dynamics for $\dot{\mathbf{w}}$ and $\dot{\alpha}$:

$$\dot{\mathbf{w}}_{\perp} = \frac{a}{\tau M} (\mathbf{r} - a\mathbf{w}) - \frac{a}{\tau M} (1 - a\alpha)\mathbf{r} = -\frac{a^2}{\tau M} \mathbf{w}_{\perp}$$

Since $-a^2/\tau < 0$, the perpendicular component decays near exponentially:

$$\mathbf{w}_{\perp}(t) = \mathbf{w}_{\perp}(0) \exp\left(-\int_0^t \frac{a^2(s)}{\tau M} ds\right) \rightarrow \mathbf{0}$$

After the perpendicular component decays, $\mathbf{w}(t) \approx \alpha(t)\mathbf{r}$, which simplifies the collective variables:

$$\sum_{j=1}^M w_j r_j = \alpha \|\mathbf{r}\|^2, \quad \sum_{j=1}^M w_j^2 = \alpha^2 \|\mathbf{r}\|^2$$

and therefore we can write a in terms of α and \mathbf{r} :

$$a = \frac{\alpha \|\mathbf{r}\|^2}{1 + \alpha^2 \|\mathbf{r}\|^2} \quad (10)$$

The evolution $\dot{\alpha}$ becomes:

$$\dot{\alpha} = \frac{1}{\tau M} a (1 - a\alpha) = \frac{1}{\tau M} \frac{\alpha \|\mathbf{r}\|^2}{(1 + \alpha^2 \|\mathbf{r}\|^2)^2}$$

For $\alpha^2 \|\mathbf{r}\|^2 \gg 1$, which is equivalent to $\sum_{j=1}^M w_j^2 \gg 1$ (valid when the synaptic weights have grown substantially large),

$$\dot{\alpha} \approx \frac{1}{\tau M \alpha^3 \|\mathbf{r}\|^2}$$

This approximation captures the asymptotic regime where the weights have strengthened considerably through Hebbian learning.

This separable equation integrates to:

$$\alpha^4 = \frac{4t}{\tau M \|\mathbf{r}\|^2} + C_1$$

On the other hand, we can derive for individual MC activity with expression of α too:

$$m_i(t) = r_i - \alpha r_i a = r_i \left(1 - \frac{\alpha^2 \|\mathbf{r}\|^2}{1 + \alpha^2 \|\mathbf{r}\|^2} \right) = \frac{r_i}{1 + \alpha^2 \|\mathbf{r}\|^2}$$

Therefore, population second moment:

$$\langle m^2 \rangle = \frac{1}{M} \sum_i m_i^2 = \frac{1}{M} \sum_i \frac{r_i^2}{(1 + \alpha^2 \|\mathbf{r}\|^2)^2} = \frac{\langle r^2 \rangle}{(1 + \alpha^2 \|\mathbf{r}\|^2)^2}$$

Substituting expression for α into the second moment expression and using $\|\mathbf{r}\|^2 = M \langle r^2 \rangle$:

$$\langle m^2 \rangle \approx \frac{\langle r^2 \rangle}{\alpha^4 \|\mathbf{r}\|^4} = \frac{\langle r^2 \rangle \tau M \|\mathbf{r}\|^2}{(4t + C_2 \tau) \|\mathbf{r}\|^4} = \frac{\tau}{4t + C_2 \tau}$$

This provides an explicit functional form for fitting experimental data, with parameter τ representing the effective synaptic time constant, and C_2 a shifting constant. In the main text, we use $C = C_2$ to denote this constant, and therefore we have:

$$\boxed{\langle m^2 \rangle \approx \frac{\tau}{4t + C\tau}} \quad (11)$$

The mean-field approach gives us the inverse-time decay, which provides mechanistic insight for how weight alignment with input patterns drives the reductions of mitral cell population activities.

3 Learning on Single Odor Input

3.1 Residual m Aligns to r

In this section, we show that the residual for a single input during learning with Hebbian plasticity aligns with the input. By “alignment,” we mean that m is parallel to r , achieving $\cos(\psi) = 1$ where ψ is the angle between m and r . We assume the network has been pre-trained on samples drawn from the same distribution as the input, a centered Gaussian distribution. This initialization controls for distributional novelty effects, ensuring that observed alignment dynamics reflect learning of the specific input pattern r rather than adaptation to an unfamiliar input statistics.

Setup: Before exposing the network to a specific input r , we assume it learns from samples $r' \sim \mathcal{N}(0, \sigma I)$ in dimension M .

Lemma 1: After pre-training on centered inputs, the residual m aligns with any new input r at first exposure.

Proof. During pre-training, the network learns from samples $r' \sim \mathcal{N}(0, \sigma^2 I)$. In the limit of many samples, the sample covariance becomes:

$$\frac{1}{N} RR^T \rightarrow \sigma^2 I$$

At steady state with weight decay λ , the dynamics satisfy:

$$\dot{W} = \eta\theta(I + \theta WW^T)^{-1} RR^T W (I + \theta W^T W)^{-1} - \lambda W = 0$$

With SVD $W = USV^T$, this becomes:

$$\eta\theta(I + \theta S^2)^{-1} U^T (RR^T) US (I + \theta S^2)^{-1} = \lambda S$$

Since $U^T(\sigma^2 I)U = \sigma^2 I$, we get:

$$\eta\theta(I + \theta S^2)^{-1} \sigma^2 S (I + \theta S^2)^{-1} = \lambda S$$

For each singular value s_i :

$$\eta\theta\sigma^2 \frac{s_i}{(1 + \theta s_i^2)^2} = \lambda s_i$$

Solving for s_i :

$$\frac{\eta\theta\sigma^2}{(1 + \theta s_i^2)^2} = \lambda$$

This gives the same solution for all i :

$$s_i^2 = \frac{1}{\theta} \left(\sqrt{\frac{\eta\sigma^2}{\theta\lambda}} - 1 \right) \equiv s^2$$

For the pre-trained network, all singular values are equal: $S = sI$. The residual for any new input r is:

$$m = (I + \theta WW^T)^{-1}r = (I + \theta s^2 UU^T)^{-1}r = (I + \theta s^2 I)^{-1}r = \frac{1}{1 + \theta s^2}r$$

Therefore, m is perfectly aligned with r during first exposure. □

Lemma 2: After W update on r after first exposure, m aligns with r , and same applies to all future updates.

Proof. We call the weight matrix after pretraining W_0 , with $W_0 = U_0 S_0 V_0^T$, $s_i \equiv s$, and that $W_1 = W_0 + \eta m a^T$.

Therefore W_1 can be written as:

$$W_1 = U_0(S_0 + C_{(S_0)} U_0^T r r^T U_0) V_0^T = U_0(S_0 + C_{(S_0)} z z^T) V_0^T$$

where $C_{(S_0)}$ is diagonal with all entries of same values, and we also simplify by defining $z = U_0^T r$.

Let $M = S_0 + C_{(S_0)} z z^T$. Since $S_0 = sI$ and $C_{(S_0)} = cI$ for some scalar c , we have:

$$M = sI + c z z^T$$

The eigenvalues of M are: $\lambda_1 = s + c\|z\|^2$ with eigenvector $q_1 = \frac{z}{\|z\|}$; $\lambda_i = s$ for $i = 2, \dots, M$ with eigenvectors orthogonal to z .

Let $M = Q\Lambda Q^T$,

$$W_1 = U_0 M V_0^T = U_0 Q \Lambda Q^T V_0^T$$

Then we know singular vectors u_i of W_1 , $u_i = U_0 q_i$. Since $q_1 = \frac{z}{\|z\|}$, $u_1 = U_0 q_1 = U_0 U_0^T \frac{r}{\|r\|} \propto r$.

Now, for the residual with updated weights, since $W_1^T W_1 = V_0 M^T M V_0^T$, they share singular values:

$$m = (I + \theta W_1 W_1^T)^{-1}r = (I + \theta U_0 \Lambda^2 U_0^T)^{-1}r$$

Since $r \propto u_1$, we can write $r = \alpha u_1$ for some scalar $\alpha > 0$. Then:

$$\begin{aligned} m &= (I + \theta U_0 \Lambda^2 U_0^T)^{-1} (\alpha u_1) = \alpha (I + \theta U_0 \Lambda^2 U_0^T)^{-1} u_1 \\ &= \frac{\alpha}{1 + \theta \lambda_1^2} u_1 \propto u_1 \propto r \end{aligned}$$

Therefore, m remains aligned with r after the first weight update.

Inductive argument: By induction, this alignment is preserved for all future updates. At each exposure to r , the Hebbian update $\Delta W = \eta m a^T$ further strengthens the connection in the r direction since both m and a are aligned with r . This increases the singular value λ_1 corresponding to the r direction while leaving other singular values unchanged. Since the leading eigenvector q_1 remains proportional to $U_0^T r$ (and hence $u_1 \propto r$), the alignment is preserved through all subsequent updates. \square

3.2 Presence of Weight Noise

In this section we show that there is no systematic rotation in expectation with stochasticity in weight update. For systematic rotation to exist in the expected trajectory, we would need the infinitesimal expected change in m , δm , to contain a part $\propto Jm$, where J is an antisymmetric rotation matrix. Our goal is to inspect $\mathbb{E}[\delta m] = \mathbb{E}[m(t+1)] - \mathbb{E}[m(t)]$ when δm becomes very small.

Recall that we have arrived at a succinct form of m , expressing it with W and r :

$$m = (\theta W W^T + I_M)^{-1} r$$

Let's analyze how a single noisy update affects m . At time step t :

$$W(t+1) = W(t) + \eta(m(t)a(t)^T + \xi^W(t))$$

where $\xi^W(t)$ has i.i.d. entries with variance $\sigma^2 \leq \eta^2 \|ma^T\|^2$.

Some key properties of the independent noise ξ^W :

$$\begin{aligned} \mathbb{E}[\xi_{ij}^W] &= 0 \quad \text{for all } i, j \\ \mathbb{E}[\xi_{ij}^W \xi_{kl}^W] &= \sigma^2 \quad \text{if } i = k \text{ and } j = l \\ &= 0 \quad \text{otherwise} \end{aligned}$$

Let's write the update as a perturbation:

$$W(t+1) = W(t) + \Delta W_{\text{det}} + \Delta W_{\text{noise}} \tag{12}$$

where:

$$\begin{aligned}\Delta W_{\text{det}} &= \eta m a^T \\ \Delta W_{\text{noise}} &= \eta \xi^W\end{aligned}$$

Using the relationship for $m(t+1)$:

$$m(t+1) = (\theta W(t+1)W(t+1)^T + I_M)^{-1} r$$

Using the expression of eq (12), multiply out $(W + \Delta W_{\text{det}} + \Delta W_{\text{noise}})(W + \Delta W_{\text{det}} + \Delta W_{\text{noise}})^T$, we get:

$$\begin{aligned}W(t+1)W(t+1)^T &= (W + \Delta W_{\text{det}} + \Delta W_{\text{noise}})(W + \Delta W_{\text{det}} + \Delta W_{\text{noise}})^T \\ &= (W + \Delta W_{\text{det}} + \Delta W_{\text{noise}})(W^T + \Delta W_{\text{det}}^T + \Delta W_{\text{noise}}^T) \\ &= WW^T \text{ (zeroth order terms)} \\ &+ W\Delta W_{\text{det}}^T + \Delta W_{\text{det}}W^T \text{ (first order in det. update)} \\ &+ W\Delta W_{\text{noise}}^T + \Delta W_{\text{noise}}W^T \text{ (first order in noise)} \\ &+ \Delta W_{\text{det}}\Delta W_{\text{det}}^T \text{ (second order in det. update)} \\ &+ \Delta W_{\text{det}}\Delta W_{\text{noise}}^T + \Delta W_{\text{noise}}\Delta W_{\text{det}}^T \text{ (mixed second order)} \\ &+ \Delta W_{\text{noise}}\Delta W_{\text{noise}}^T \text{ (second order in noise)}\end{aligned}$$

Now, substituting the $\Delta W_{\text{det}} = \eta m a^T$ and $\Delta W_{\text{noise}} = \eta \xi^W$ back, we get:

$$\begin{aligned}W(t+1)W(t+1)^T &= WW^T \\ &+ \eta(Wam^T + ma^TW^T) \\ &+ \eta(W\xi^W{}^T + \xi^WW^T) \\ &+ \eta^2 ma^T am^T \\ &+ \eta^2 (ma^T \xi^W{}^T + \xi^W am^T) \\ &+ \eta^2 \xi^W \xi^W{}^T\end{aligned}$$

First-order noise terms. When we multiply $W(t+1)W(t+1)^T$, the first-order terms in noise are:

$$\eta W \xi^W{}^T + \eta \xi^W W^T$$

For each entry (i, j) in this matrix:

$$[\eta W \xi^{W^T}]_{ij} = \eta \sum_k W_{ik} \xi_{jk}^W$$

$$[\eta \xi^W W^T]_{ij} = \eta \sum_k \xi_{ik}^W W_{jk}$$

Therefore:

$$\mathbb{E}[W \xi^{W^T}]_{ij} = \eta \sum_k W_{ik} \mathbb{E}[\xi_{jk}^W] = 0$$

$$\mathbb{E}[\xi^W W^T]_{ij} = \eta \sum_k \mathbb{E}[\xi_{ik}^W] W_{jk} = 0$$

This shows that first-order noise terms vanish in expectation, and therefore any directional changes from noise must come from higher-order terms. This helps explain why when noise is comparable to the input ($\sigma \sim \eta \|ma^T\|$) has minimal effect on the direction of m .

Second-order terms. The second-order terms in $W(t+1)W(t+1)^T$ are:

$$\begin{aligned} \eta^2 \xi^W \xi^{W^T} & \quad \text{(pure noise)} \\ + \eta^2 ma^T am^T & \quad \text{(pure deterministic)} \\ + \eta^2 (ma^T \xi^{W^T} + \xi^W am^T) & \quad \text{(mixed)} \end{aligned}$$

Taking expectations, same as in the pretraining session:

$$\begin{aligned} \mathbb{E}[\xi^W \xi^{W^T}]_{ij} &= \begin{cases} \sigma^2 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} \\ &= \sigma^2 I_M \end{aligned}$$

The mixed terms vanish in expectation since $\mathbb{E}[\xi^W] = 0$.

3.2.1 Approximating $\mathbb{E}[\delta m]$

Recall that $m = (\theta W W^T + I_M)^{-1} r$. For the current state, we leave out the time subscripts and let's write:

$$\theta W(t+1)W(t+1)^T + I_M = (\theta W W^T + I_M) + \theta \Delta$$

where Δ includes all update terms expanding $W(t+1)W(t+1)^T$. Taking the expectation, then many terms zeroed out, and we are left with:

$$\mathbb{E}[\theta \Delta] = \theta \eta (W am^T + ma^T W^T) + \theta \eta^2 \sigma^2 I_M + \theta \eta^2 ma^T am^T$$

Now, treat $\theta WW^T + I_M$ as A , and the aggregated perturbation term $\theta\Delta$ as ϵB , especially identify η with ϵ , we can use the below expansion to approximate $(\theta W(t+1)W(t+1)^T + I_M)^{-1}$:

$$(A + \epsilon B)^{-1} = A^{-1} - \epsilon A^{-1} B A^{-1} + \epsilon^2 A^{-1} B A^{-1} B A^{-1} + O(\epsilon^3) + \dots$$

Applying this to our case, keeping only first order term,

$$\begin{aligned} m(t+1) &= [\theta WW^T + I_M + \theta\Delta]^{-1} r \\ &\approx (\theta WW^T + I_M)^{-1} r \\ &\quad - (\theta WW^T + I_M)^{-1} \theta\Delta (\theta WW^T + I_M)^{-1} r \end{aligned}$$

Taking expectation and use $m(t) = (\theta WW^T + I_M)^{-1} r$

$$\begin{aligned} \mathbb{E}[\delta m] &= \mathbb{E}[m(t+1)] - \mathbb{E}[m(t)] \\ &= -\eta \theta A^{-1} (W a m^T + m a^T W^T) A^{-1} r \end{aligned}$$

Note that A^{-1} is symmetric, and $W a m^T + m a^T W^T$ is a matrix summed with its transpose so always symmetric, thus there is no antisymmetric component. We conclude that when update W with isotropic Gaussian noise, there is no systematic rotation being introduced.

3.3 Presence of Input Noise

Assume that input is $\tilde{r} = r + \varepsilon$, ε being random Gaussian with zero mean and a small variance, compared to the input variability. We find the expected converged \tilde{m} 's expression will give us information about the deviation from r 's direction.

Consider training with noisy input $\tilde{r} = r + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma^2 I)$. The residual becomes:

$$\tilde{m} = (\theta WW^T + I_M)^{-1} (r + \varepsilon) = m + m_{\text{noise}} \quad (13)$$

where $m = (\theta WW^T + I_M)^{-1} r$ and $m_{\text{noise}} = (\theta WW^T + I_M)^{-1} \varepsilon$.

Steady-state analysis with input noise. The weight update becomes:

$$\dot{W} = \eta \tilde{m} \tilde{a}^T - \lambda W = \eta (\theta WW^T + I_M)^{-1} (r + \varepsilon) (r + \varepsilon)^T W (\theta W^T W + I_N)^{-1} - \lambda W$$

At steady state, $\mathbb{E}[\dot{W}] = 0$. Since $\mathbb{E}[\varepsilon] = 0$ and $\mathbb{E}[\varepsilon \varepsilon^T] = \sigma^2 I$:

$$\mathbb{E}[(r + \varepsilon)(r + \varepsilon)^T] = r r^T + \sigma^2 I$$

Using our standard SVD approach $W = USV^T$, the steady-state condition becomes:

$$\eta\theta(\theta S^2 + I)^{-1}U^T(rr^T + \sigma^2 I)US(\theta S^2 + I)^{-1} = \lambda S$$

For this to be diagonal, we need $U^T(rr^T + \sigma^2 I)U$ to be diagonal. This requires $U^T r$ to have only one non-zero component, giving:

$$U^T(rr^T + \sigma^2 I)U = \begin{bmatrix} \|r\|^2 + \sigma^2 & 0 & \dots \\ 0 & \sigma^2 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

Therefore, $u_1 = \frac{r}{\|r\|}$ remains aligned with r , and the singular values satisfy:

$$s_1 = \sqrt{\frac{1}{\theta} \left(\sqrt{\frac{\eta\theta(\|r\|^2 + \sigma^2)}{\lambda}} - 1 \right)}$$

$$s_i = \sqrt{\frac{1}{\theta} \left(\sqrt{\frac{\eta\theta\sigma^2}{\lambda}} - 1 \right)} \quad \text{for } i > 1$$

The residual thus decomposes as:

$$\tilde{m} = \frac{1}{1 + \theta s_1^2} \frac{rr^T}{\|r\|^2} r + \sum_{i=2}^M \frac{1}{1 + \theta s_i^2} u_i u_i^T \varepsilon \quad (14)$$

which shows that, the deterministic component $m \propto r$ maintains alignment with the input, but noise is filtered differently across directions: strongly suppressed along r (due to large s_1), less suppressed in orthogonal directions (due to smaller s_i). The noise magnitude is larger in directions orthogonal to r , creating bounded deviations from perfect alignment.

Moreover, with continued learning, the deterministic component of \tilde{m} decays faster than the noise component. Since $s_1^2 > s_i^2$ for $i > 1$, the deterministic term $m_{\text{det}} \propto \frac{1}{1 + \theta s_1^2}$ is suppressed more strongly than the noise terms $m_{\text{noise}} \propto \frac{1}{1 + \theta s_i^2}$. This leads to random walk-like drift in the orthogonal subspace, even though $\mathbb{E}[\tilde{m}]$ remains aligned with r in expectation.

3.4 Presence of Structural Connectivity

In realistic neural circuits, synaptic connectivity is constrained by anatomy. While analyzing arbitrary sparse connectivity matrices is intractable, we can gain insight by studying simplified cases that capture key features of olfactory bulb organization.

The olfactory bulb exhibits structured connectivity: MCs and GCs are arranged in layers with dense local connections due to physical proximity and limited dendritic spread, supplemented by sparse

long-range connections. To understand how such connectivity patterns affect learning dynamics, we analyze two tractable cases that preserve essential structural features:

1) Block diagonal connectivity: A is block diagonal with all-to-all connections within blocks, modeling distinct local circuits.

2) Banded connectivity: A has a band matrix structure, modeling distance-dependent connection probabilities.

Both cases constrain the weight updates $W \leftarrow W + \eta(ma^T \odot A)$, where \odot denotes element-wise multiplication, fundamentally altering the learning dynamics from the all-to-all case.

3.4.1 Block-diagonal Connectivity

Intuition for the block diagonal case: consider A with two blocks:

$$A = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix}, \quad W = \begin{bmatrix} W_1 & 0 \\ 0 & W_2 \end{bmatrix}$$

The key insight is that $(\theta WW^T + I_M)^{-1}$ inherits the block structure:

$$(\theta WW^T + I_M)^{-1} = \begin{bmatrix} (\theta W_1 W_1^T + I_{M_1})^{-1} & 0 \\ 0 & (\theta W_2 W_2^T + I_{M_2})^{-1} \end{bmatrix}$$

Therefore, for input $r = [r_1^T, r_2^T]^T$:

$$m = (\theta WW^T + I_M)^{-1} r = \begin{pmatrix} (\theta W_1 W_1^T + I_{M_1})^{-1} r_1 \\ (\theta W_2 W_2^T + I_{M_2})^{-1} r_2 \end{pmatrix}$$

Key consequences for block constraint is that each block processes its input independently, and each block's suppression is scaled by its own input, and this differential suppression by block creates an angle between the overall m and r .

For constant input r and block sizes that are approximately Gaussian-distributed $x \sim \mathcal{N}(\mu, \sigma^2)$ (which approximates the binomial distribution under the assumption of large M and many blocks) with $K = M/\mu$ total blocks, the angle between m and r is:

$$\cos(\psi) = \frac{\sum_k r_k^T (\theta W_k W_k^T + I_{M_k})^{-1} r_k}{\sqrt{\sum_k \|(\theta W_k W_k^T + I_{M_k})^{-1} r_k\|^2} \sqrt{\sum_k \|r_k\|^2}}$$

Steady state angle for varying block sizes

At steady state with weight decay λ , each block k satisfies:

$$\eta \theta (\theta S_k^2 + I)^{-1} U_k^T r_k r_k^T U_k S_k (\theta S_k^2 + I)^{-1} = \lambda S_k$$

Following our standard analysis, the first singular vector $u_{k,1}$ aligns with r_k , giving:

$$s_{k,1} = \sqrt{\sqrt{\frac{\eta \|r_k\|^2}{\theta \lambda}} - \frac{1}{\theta}}$$

For each block, the residual becomes:

$$m_k = (\theta W_k W_k^T + I_{M_k})^{-1} r_k = \frac{1}{\theta s_{k,1}^2 + 1} r_k$$

The overall angle between m and r is:

$$\cos(\psi_\infty) = \frac{\sum_k \frac{1}{\theta s_{k,1}^2 + 1} \|r_k\|^2}{\sqrt{\sum_k \frac{1}{(\theta s_{k,1}^2 + 1)^2} \|r_k\|^2} \sqrt{\sum_k \|r_k\|^2}}$$

For constant input density, $\|r_k\|^2 = \frac{x_k}{M} \|r\|^2$ where x_k is the size of block k .

Defining $c = \sqrt{\frac{\lambda}{\theta \eta}}$ and substituting:

$$\frac{1}{\theta s_{k,1}^2 + 1} = \frac{cM}{x_k \|r\|}$$

This yields:

$$\cos(\psi_\infty) = \frac{\sum_k \frac{c}{\|r\|}}{\sqrt{\sum_k \frac{c^2 M}{\|r\|^2 x_k}}} = \sqrt{\frac{K}{\sum_k \frac{M}{x_k}}}$$

where K is the number of blocks.

Expected angle with block size variance: For $x_k \sim \mathcal{N}(\mu, \sigma^2)$, using the approximation $\mathbb{E}[1/x_k] \approx \frac{1}{\mu} (1 + \frac{\sigma^2}{\mu^2})$:

$$\mathbb{E}[\cos(\psi_\infty)] \approx \sqrt{\frac{1}{1 + \frac{\sigma^2}{\mu^2}}} \quad (15)$$

If we have uniform blocks ($\sigma = 0$), we would have perfect alignment ($\cos(\psi_\infty) = 1$); on the other hand, if we have high variance, ($\sigma^2/\mu^2 \gg 1$): alignment would be poor ($\cos(\psi_\infty) \rightarrow 0$).

Analysis of input distribution and final angle

Now we fix the block size d with K total blocks, and analyze how input distribution affects alignment.

From our steady-state analysis:

$$m_k = \frac{1}{\theta s_{k,1}^2 + 1} r_k = \frac{r_k}{c' \|r_k\|}$$

where $c' = \sqrt{\frac{\theta\eta}{\lambda}}$.

The cosine of the angle between m and r becomes:

$$\cos(\psi) = \frac{\sum_k m_k^T r_k}{\|m\| \|r\|} = \frac{\sum_k \frac{\|r_k\|^2}{c' \|r_k\|}}{\sqrt{\sum_k \frac{\|r_k\|^2}{(c' \|r_k\|)^2}} \sqrt{\sum_k \|r_k\|^2}} = \frac{\sum_k \|r_k\|}{\sqrt{K} \sqrt{\sum_k \|r_k\|^2}}$$

Gaussian input analysis. For each component $r_{k,i} \sim \mathcal{N}(\mu, \sigma^2)$ independently, we have:

$$\mathbb{E}[\|r_k\|^2] = \mathbb{E}\left[\sum_{i=1}^d r_{k,i}^2\right] = d(\sigma^2 + \mu^2)$$

For the expected norm $\mathbb{E}[\|r_k\|]$, since $\|r_k\|^2$ follows a non-central chi-squared distribution, we can compute numerically:

$$\mathbb{E}[\|r_k\|] = \int_0^\infty \sqrt{t} f_{\|r_k\|^2}(t) dt$$

Therefore:

$$\mathbb{E}[\cos(\psi)] = \frac{K \cdot \mathbb{E}[\|r_k\|]}{\sqrt{K} \sqrt{K \cdot \mathbb{E}[\|r_k\|^2]}} = \frac{\mathbb{E}[\|r_k\|]}{\sqrt{d(\sigma^2 + \mu^2)}} \quad (16)$$

When we have uniform inputs (all $\|r_k\|$ equal), $\cos(\psi) = 1$. When not, alignment depends on both mean μ and variance σ^2 of inputs. The larger block count K , the smaller the block-wise variance, the greater the alignment (simulation see Fig. 4H in main).

3.4.2 Band-diagonal Connectivity

Intuition with band structure: consider connectivity matrix A with fixed bandwidth k :

$$A_{ij} = \begin{cases} 1 & \text{if } |i - j| \leq k \\ 0 & \text{otherwise} \end{cases}$$

This constrains W to have the same band structure: $W_{ij} = 0$ if $|i - j| > k$.

Let's inspect the band structure effects. For $B = \theta W^T W + I_N$, we can use the geometric series expansion:

$$B^{-1} = (I_N + \theta W^T W)^{-1} = I_N - \theta W^T W + \theta^2 (W^T W)^2 - \theta^3 (W^T W)^3 + \dots$$

With power n , $(W^T W)^n$ has bandwidth $2nk$, but with decreasing magnitude due to the θ^n factor (assuming $\theta < 1$). The residual becomes:

$$m_i = r_i - \sum_{|j-i| \leq 4k} Q_{ij} r_j + (\text{higher-order terms})$$

This shows that each m_i is primarily influenced by r_j within distance $\sim 4k$, creating a “sliding window” like effect, where adjacent regions share input components, therefore still creating local suppression patterns, similar to the block diagonal case.

4 Learning on Odor Pair

4.1 The Decorrelation Motif

Theorem: When the system learns two input vectors r_1 and r_2 , their corresponding residuals m_1 and m_2 will be orthogonalized at steady state.

Proof. We will show that if r_1 and r_2 initially form angle ϕ , then m_1 and m_2 will each move away from the other input by angle $\psi = \frac{1}{2}|\phi - \pi/2|$, achieving orthogonality when $\phi \neq \pi/2$.

Coordinate system and eigendecomposition. Without loss of generality, consider two input vectors r_1 and r_2 with equal magnitude α . We choose coordinates such that:

$$r_1 = \alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad r_2 = \alpha \begin{pmatrix} \cos \phi \\ \sin \phi \end{pmatrix}$$

where ϕ is the angle between r_1 and r_2 .

The covariance matrix becomes:

$$RR^T = \alpha^2 \begin{bmatrix} 1 + \cos^2 \phi & \cos \phi \sin \phi \\ \cos \phi \sin \phi & \sin^2 \phi \end{bmatrix}$$

Solving the characteristic equation $\det(RR^T - \delta I) = 0$ yields eigenvalues:

$$\delta_{\pm} = \alpha^2(1 \pm \cos \phi)$$

with corresponding normalized eigenvectors:

$$v_+ = \frac{1}{\sqrt{2(1 + \cos \phi)}} \begin{pmatrix} 1 + \cos \phi \\ \sin \phi \end{pmatrix}$$

$$v_- = \frac{1}{\sqrt{2(1 - \cos \phi)}} \begin{pmatrix} 1 - \cos \phi \\ -\sin \phi \end{pmatrix}$$

Steady-state weight matrix. At steady state with weight decay λ :

$$\eta\theta(\theta S^2 + I)^{-1}U^T Q \Delta Q^T U S(\theta S^2 + I)^{-1} = \lambda S$$

For this equation to hold with diagonal right-hand side, we require $Q = U$. Therefore, the left singular vectors of the steady-state W are $U = [v_+ | v_-]$:

$$U = \begin{bmatrix} \frac{\sqrt{1+\cos \phi}}{\sqrt{2}} & \frac{\sqrt{1-\cos \phi}}{\sqrt{2}} \\ \frac{\sin \phi}{\sqrt{2(1+\cos \phi)}} & \frac{-\sin \phi}{\sqrt{2(1-\cos \phi)}} \end{bmatrix}$$

with singular values:

$$s_i = \sqrt{\sqrt{\frac{\eta\delta_i}{\theta\lambda}} - \frac{1}{\theta}}$$

Residual computation and angle derivation. The residual for r_1 is:

$$m_1 = (I + \theta WW^T)^{-1} r_1 = U \begin{bmatrix} \frac{1}{\theta s_1^2 + 1} & 0 \\ 0 & \frac{1}{\theta s_2^2 + 1} \end{bmatrix} U^T r_1$$

Defining $c = \alpha\sqrt{\theta\eta/\lambda}$ and substituting the expressions for s_i and U , we obtain:

$$m_1 = \frac{\alpha}{c\sqrt{2}} \begin{pmatrix} \sqrt{1 + \cos\phi} + \sqrt{1 - \cos\phi} \\ \sqrt{1 - \cos\phi} - \sqrt{1 + \cos\phi} \end{pmatrix}$$

The angle ψ between m_1 and r_1 satisfies:

$$\cos\psi = \frac{m_1^T r_1}{\|m_1\| \|r_1\|} = \frac{\sqrt{1 + \cos\phi} + \sqrt{1 - \cos\phi}}{2} \quad (17)$$

Now we show that this implies $\phi = \frac{1}{2}|\psi - \pi/2|$:

$$\begin{aligned} \frac{\sqrt{1 + \cos\psi} + \sqrt{1 - \cos\psi}}{2} &= \frac{\sqrt{2}}{2} \sqrt{\frac{1 + \cos\psi}{2}} + \frac{\sqrt{2}}{2} \sqrt{\frac{1 - \cos\psi}{2}} \\ &= \frac{\sqrt{2}}{2} \cos(\psi/2) + \frac{\sqrt{2}}{2} \sin(\psi/2) \\ &= \cos(\pi/4) \cos(\psi/2) + \sin(\pi/4) \sin(\psi/2) = \cos(\pi/4 - \psi/2) \\ &= \cos\left(\frac{\pi/2 - \psi}{2}\right) \end{aligned}$$

Since $\cos\phi = \cos\left(\frac{\pi/2 - \psi}{2}\right)$, we conclude that $\phi = \frac{|\pi/2 - \psi|}{2} = \frac{1}{2}|\psi - \pi/2|$.

□

Prediction for more than two inputs

Our prediction is that for n input odors, the system will achieve pairwise orthogonalization of all residuals m^i . Each residual will move in the direction that maximizes its distance from the subspace spanned by all other inputs, ultimately converging to a configuration where all residuals are mutually orthogonal. Complete orthogonalization can be achieved as long as the input dimension satisfies $M \geq n$.

Prediction for learning input sets with lower rank

Consider a set of n_1 input mixtures composed of n_2 underlying odorants, where $n_2 < n_1$. Although there are n_1 distinct inputs, the covariance matrix RR^T has rank n_2 , constraining all inputs to lie within an n_2 -dimensional subspace. We predict that the system will learn this covariance structure and optimally separate the inputs within the subspace they span. The residuals will achieve maximal separation subject to the constraint that they remain orthogonal to the null space of the input covariance.

4.2 Temporally correlated noise

In Section 3.3, we showed that with uncorrelated input noise, the residual m aligns with r in expectation for a single input. Therefore, with two inputs and independent noise, m_1 and m_2 will follow their expected decorrelation trajectory as derived in the noiseless case.

However, motivated by trial-to-trial noise correlations observed in neural responses, we now consider correlated noise between inputs: $\mathbb{E}[\varepsilon_1 \varepsilon_2^T] = \sigma'^2 I$ where $\sigma' < \sigma$. These shared fluctuations across trials prevent perfect decorrelation.

For two noisy inputs $\tilde{r}_1 = r_1 + \varepsilon_1$ and $\tilde{r}_2 = r_2 + \varepsilon_2$, we analyze the residual correlation:

$$\cos(\phi) = \frac{\tilde{m}_1^T \tilde{m}_2}{\|\tilde{m}_1\| \|\tilde{m}_2\|}$$

The numerator:

$$\begin{aligned} \tilde{m}_1^T \tilde{m}_2 &= [(\theta WW^T + I)^{-1}(r_1 + \varepsilon_1)]^T [(\theta WW^T + I)^{-1}(r_2 + \varepsilon_2)] \\ &= m_1^T m_2 + r_1^T (\theta WW^T + I)^{-2} \varepsilon_2 + \varepsilon_1^T (\theta WW^T + I)^{-2} r_2 + \varepsilon_1^T (\theta WW^T + I)^{-2} \varepsilon_2 \end{aligned}$$

Taking expectations and using $\mathbb{E}[\varepsilon_i] = 0$:

$$\mathbb{E}[\tilde{m}_1^T \tilde{m}_2] = m_1^T m_2 + \mathbb{E}[\varepsilon_1^T (\theta WW^T + I)^{-2} \varepsilon_2] = 0 + \sigma'^2 \text{Tr}[(\theta WW^T + I)^{-2}]$$

where the first term vanishes at convergence due to perfect decorrelation in the noiseless case.

For denominator similarly:

$$\mathbb{E}[\|\tilde{m}_i\|^2] = \|m_i\|^2 + \sigma'^2 \text{Tr}[(\theta WW^T + I)^{-2}]$$

Defining $T = \text{Tr}[(\theta WW^T + I)^{-2}]$:

$$\mathbb{E}[\cos(\phi)] = \mathbb{E} \left[\frac{\tilde{m}_1^T \tilde{m}_2}{\|\tilde{m}_1\| \|\tilde{m}_2\|} \right] \approx \frac{\mathbb{E}[\tilde{m}_1^T \tilde{m}_2]}{\sqrt{\mathbb{E}[\|\tilde{m}_1\|^2] \mathbb{E}[\|\tilde{m}_2\|^2]}} \quad (18)$$

$$= \frac{\sigma'^2 T}{\sqrt{(\|m_1\|^2 + \sigma'^2 T)(\|m_2\|^2 + \sigma'^2 T)}} \quad (19)$$

Inspecting this equation, we can see that perfect decorrelation is impossible: with correlated noise ($\sigma' > 0$), $\mathbb{E}[\cos(\phi)] > 0$, residuals cannot achieve perfect orthogonality.

From our following two-input analysis, we can relatively easily show that $\|m_1\|^2 = \|m_2\|^2$ regardless of the initial angle ϕ between r_1 and r_2 .

With $W = USV^T$, we have at steady state:

$$T = \sum_{i=1}^M \frac{1}{(1 + \theta s_i^2)^2} \approx \frac{M - 2}{(1 + \theta \sigma^2)^2} \quad \text{for } M \gg 2 \quad (20)$$

Since T is dominated by the noise subspace dimensions and is similar across all pairs, all odor pairs should converge to approximately the same residual correlation angle and the angle is controlled by mostly the noise magnitude.

This prediction is confirmed in stimulation (see Fig. 3G in main).