

Geometry and dynamics of representations in a precisely balanced memory network related to olfactory cortex

Reviewed Preprint

v2 • December 31, 2024

Revised by authors

Reviewed Preprint

v1 • May 9, 2024

Claire Meissner-Bernard, Friedemann Zenke, Rainer W Friedrich 

Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland • University of Basel, Basel, Switzerland

 https://en.wikipedia.org/wiki/Open_access Copyright information

eLife Assessment

This **important** study introduces a biologically constrained model of telencephalic area of adult zebrafish to highlight the significance of precisely balanced memory networks in olfactory processing. The authors provide **compelling** evidence that their model performs better in multiple situations (for e.g. in terms of network stability and shaping the geometry of representations), compared to traditional attractor networks and persistent activity. The work supports recent studies reporting functional E/I subnetworks in several sensory cortexes, and will be of interest to both theoretical and experimental neuroscientists studying network dynamics based on structured excitatory and inhibitory interactions.

<https://doi.org/10.7554/eLife.96303.2.sa4>

Abstract

Biological memory networks are thought to store information by experience-dependent changes in the synaptic connectivity between assemblies of neurons. Recent models suggest that these assemblies contain both excitatory and inhibitory neurons (E/I assemblies), resulting in co-tuning and precise balance of excitation and inhibition. To understand computational consequences of E/I assemblies under biologically realistic constraints we built a spiking network model based on experimental data from telencephalic area Dp of adult zebrafish, a precisely balanced recurrent network homologous to piriform cortex. We found that E/I assemblies stabilized firing rate distributions compared to networks with excitatory assemblies and global inhibition. Unlike classical memory models, networks with E/I assemblies did not show discrete attractor dynamics. Rather, responses to learned inputs were locally constrained onto manifolds that “focused” activity into neuronal subspaces. The covariance structure of these manifolds supported pattern classification when information was retrieved from selected neuronal subsets. Networks with E/I assemblies therefore transformed the geometry of neuronal coding space, resulting in continuous representations that reflected both relatedness of inputs and an individual’s experience. Such continuous representations enable fast pattern classification, can support continual learning, and may provide a basis for higher-order learning and cognitive computations.

Introduction

Autoassociative memory establishes internal representations of specific inputs that may serve as a basis for higher brain functions including classification and prediction. Representation learning in autoassociative memory networks is thought to involve experience-dependent synaptic plasticity and potentially other mechanisms that enhance connectivity among assemblies of excitatory neurons (Hebb, 1949 [↗](#); Ko et al., 2011 [↗](#); Miehl et al., 2022 [↗](#); Ryan et al., 2015 [↗](#)). Classical theories proposed that assemblies define discrete attractor states and map related inputs onto a common stable output pattern. Hence, neuronal assemblies are thought to encode internal representations, or memories, that classify inputs relative to previous experience via attractor dynamics (Amit and Tsodyks, 1991 [↗](#); Goldman-Rakic, 1995 [↗](#); Hopfield, 1982 [↗](#); Kohonen, 1984 [↗](#); Lagzi and Rotter, 2015 [↗](#); Mazzucato et al., 2015 [↗](#)). However, brain areas with memory functions such as the hippocampus or neocortex often exhibit dynamics that is atypical of attractor networks including irregular firing patterns, transient responses to inputs, and high trial-to-trial variability (Iurilli and Datta, 2017 [↗](#); Renart et al., 2010 [↗](#); Shadlen and Newsome, 1994 [↗](#)).

Irregular, fluctuation-driven firing reminiscent of cortical activity emerges in recurrent networks when neurons receive strong excitatory (E) and inhibitory (I) synaptic input (Brunel, 2000 [↗](#); Shadlen and Newsome, 1994 [↗](#); Van Vreeswijk and Sompolinsky, 1996 [↗](#)). In such “balanced state” networks, enhanced connectivity among assemblies of E neurons is prone to generate runaway activity unless matched I connectivity establishes co-tuning of E and I inputs in individual neurons. The resulting state of “precise” synaptic balance stabilizes firing rates because inhomogeneities in excitation across the population or temporal variations in excitation are tracked by correlated inhibition (Hennequin et al., 2017 [↗](#), 2014 [↗](#); Lagzi and Fairhall, 2024 [↗](#); Rost et al., 2018 [↗](#); Vogels et al., 2011 [↗](#)). E/I co-tuning has been observed experimentally in cortical brain areas (Bhatia et al., 2019 [↗](#); Froemke et al., 2007 [↗](#); Okun and Lampl, 2008 [↗](#); Rupprecht and Friedrich, 2018 [↗](#); Wehr and Zador, 2003 [↗](#)) and emerged in simulations that included spike-timing dependent plasticity at I synapses (Lagzi et al., 2021 [↗](#); Litwin-Kumar and Doiron, 2014 [↗](#); Vogels et al., 2011 [↗](#); Zenke et al., 2015 [↗](#)). In simulations, E/I co-tuning can be established by including I neurons in assemblies, resulting in “E/I assemblies” where I neurons track activity of E neurons (Barron et al., 2017 [↗](#); Eckmann et al., 2024 [↗](#); Lagzi and Fairhall, 2024 [↗](#); Mackwood et al., 2021 [↗](#)). Exploring the structural basis of E/I co-tuning in biological networks is challenging because it requires the dense reconstruction of large neuronal circuits at synaptic resolution (Friedrich and Wanner, 2021 [↗](#)).

Modeling studies started to investigate effects of E/I assemblies on network dynamics (Chenkov, 2017 [↗](#); Mackwood et al., 2021 [↗](#); Sadeh and Clopath, 2020a [↗](#); Schulz et al., 2021 [↗](#)) but the impact on neuronal computations in the brain remains unclear. Balanced state networks can exhibit a broad range of dynamical behaviors, including chaotic firing, transient responses and stable states (Festa et al., 2014 [↗](#); Hennequin et al., 2014 [↗](#); Litwin-Kumar and Doiron, 2012 [↗](#); Murphy and Miller, 2009 [↗](#); Roudi and Latham, 2007 [↗](#)), implying that computational consequences of E/I assemblies depend on network parameters. We therefore examined effects of E/I assemblies on autoassociative memory in a spiking network model that was constrained by experimental data from telencephalic area Dp of adult zebrafish, which is homologous to mammalian piriform cortex (Mueller et al., 2011 [↗](#)).

Dp and piriform cortex receive direct input from mitral cells in the olfactory bulb (OB) and have been proposed to function as autoassociative memory networks (Haberly, 2001 [↗](#); Wilson and Sullivan, 2011 [↗](#)). Consistent with this hypothesis, manipulations of neuronal activity in piriform cortex affected olfactory memory (Meissner-Bernard et al., 2018 [↗](#); Sacco and Sacchetti, 2010 [↗](#)). In both brain areas, odors evoke temporally structured, spatially distributed activity patterns (Blazing and Franks, 2020 [↗](#); Stettler and Axel, 2009 [↗](#); Yaksi et al., 2009 [↗](#)) that are dominated by

synaptic inputs from recurrent connections (Franks et al., 2011; Rupprecht and Friedrich, 2018) and modified by experience (Chapuis and Wilson, 2011; Frank et al., 2019; Jacobson et al., 2018; Pashkovski et al., 2020). Whole-cell voltage clamp recordings revealed that neurons in posterior Dp (pDp) received large E and I synaptic inputs during odor responses. These inputs were co-tuned in odor space and correlated on fast timescales, demonstrating that pDp enters a transient state of precise synaptic balance during odor stimulation (Rupprecht and Friedrich, 2018).

We found that network models of pDp with assemblies but without E/I co-tuning generated persistent attractor dynamics and exhibited a biologically unrealistic broadening of the firing rate distribution. Introducing E/I assemblies established E/I co-tuning, stabilized the firing rate distribution, and abolished persistent attractor states. In networks with E/I assemblies, population activity was locally constrained onto manifolds that represented learned and related inputs by “focusing” activity into neuronal subspaces. The covariance structure of manifolds supported pattern classification when information was retrieved from selected neuronal subsets. Furthermore, the continuity of the olfactory coding space provided a metric representing the similarity of inputs to learned stimuli. These results show that autoassociative memory networks constrained by biological data operate in a balanced regime where information is contained in the geometry of neural manifolds. Predictions derived from these analyses may be tested experimentally by measurements of neuronal population activity in zebrafish.

Results

A spiking network model based on pDp

To analyze memory-related computational functions of E/I assemblies under biologically realistic constraints we built a spiking neural network model, pDp_{sim}, based on experimental data from pDp (Figure 1A) (Blumhagen et al., 2011; Rupprecht and Friedrich, 2018; Yaksi et al., 2009). pDp_{sim} comprised 4000 E neurons and 1000 I neurons, consistent with the estimated total number of 4000 – 10000 neurons in adult zebrafish pDp (unpublished observations). The network received afferent input from 1500 mitral cells in the OB with a mean spontaneous firing rate of 6 Hz (Friedrich and Laurent, 2004, 2001; Tabor and Friedrich, 2008). Odors were modeled by increasing the firing rates of 10% of mitral cells to a mean of 15 Hz and decreasing the firing rates of 5% of mitral cells to a mean of 2 Hz (Methods, Figure 1B, Friedrich and Laurent, 2004, 2001). As a result, the mean activity increased only slightly while the variance of firing rates across the mitral cell population increased approximately 7-fold, consistent with experimental observations (Friedrich and Laurent, 2004; Wanner and Friedrich, 2020).

pDp_{sim} consisted of sparsely connected integrate-and-fire neurons with conductance-based synapses (connection probability ≤5%, Methods). Model parameters were taken from the literature when available and otherwise determined to reproduce measured observables (Figure 1F–H; Methods). The mean firing rate was <0.1 Hz in the absence of stimulation and increased to ~1 Hz during odor presentation (Figure 1C, F) (Blumhagen et al., 2011; Rupprecht et al., 2021; Rupprecht and Friedrich, 2018). Population activity was odor-specific and activity patterns evoked by uncorrelated OB inputs remained uncorrelated in pDp_{sim} (Figure 1H) (Yaksi et al., 2009). The synaptic conductance during odor presentation substantially exceeded the resting conductance and inputs from other E neurons (recurrent inputs) contributed >80% of the excitatory synaptic conductance (Figure 1G). Hence, pDp_{sim} entered a balanced state during odor stimulation (Figure 1D, E) with recurrent input dominating over afferent input, as observed in pDp (Rupprecht and Friedrich, 2018). Shuffling spike times of inhibitory neurons resulted in runaway activity with a probability of 0.79 ± 0.20, demonstrating that activity was inhibition-stabilized (Sadeh and Clopath, 2020b; Tsodyks et al., 1997). These results were robust against parameter variations (Methods). pDp_{sim} therefore reproduced key features of pDp.




Figure 1.

Spiking network model of pDp.

A. Schematic of pDpsim. OB, Olfactory bulb; E, excitatory; I, inhibitory neurons. **B.** Spike raster of a random subset of 50 mitral cells in the OB representing 2 odors. During odor stimulation, firing rates of 10% of mitral cells were increased and firing rates of 5% of mitral cells were decreased (baseline rate, 6 Hz). **C.** Spike raster of random subsets of 50 E and I neurons in response to 2 odors. **D.** Representative membrane potential trace (top) and excitatory (EPSC, black) and inhibitory (IPSC, red) currents (bottom) in one excitatory neuron in response to 2 odors. Purple trace shows net current (EPSC + IPSC). **E.** Odor-evoked inhibitory (red) and excitatory (black and blue) currents as measured in a hypothetical voltage clamp experiment (conductance multiplied by 70 mV, the absolute difference between holding potential and reversal potential; Rupprecht et al., 2018). Representative example of 1 network, averaged across neurons and odors. **F-H.** Measured values of the observables used to match pDpsim to experimental data. Each dot represents one network (average over 10 odors); $n = 20$ networks. Pink shading shows the experimentally observed range of values. **F.** Baseline and odor-evoked population firing rate. **G.** Left: g_{OE} is the synaptic conductance in E neurons contributed by afferents from the OB during odor stimulation. Middle: g_{syn} is the total odor-evoked synaptic conductance. Right: % recurrent input quantifies the percentage of E input contributed by recurrent connections during odor stimulation. **H.** Correlation coefficient between odor-evoked activity patterns in pDpsim. The dotted line indicates the mean correlation between odor patterns in the OB.

Co-tuning and stability of networks with E/I assemblies

To create networks with defined neuronal assemblies we rewired a small subset of the connections in randomly connected (*rand*) networks. An assembly was generated by identifying the 100 E neurons that received the most connections from the mitral cells activated by a given odor and increasing the probability of connections between these E neurons by a factor of κ (**Figure 2A** and **Supplementary Figure 1A**). The number of E input connections per neuron was maintained by randomly eliminating connections from neurons outside the assembly. We thus refer to an odor as “learned” when a network contains a corresponding assembly, and as “novel” when no such assembly is present. In each network, we created 15 assemblies representing memories of uncorrelated odors. As a result, ~30% of E neurons were part of an assembly, with few neurons participating in multiple assemblies. Odor-evoked activity within assemblies was higher than the population mean and increased with κ (**Figure 2B**). When κ reached a critical value of ~6, networks became unstable and generated runaway activity (**Figure 2B**).

We first set κ to 5, then uniformly scaled I-to-E connection weights by a factor of χ until E population firing rates in response to learned odors matched the corresponding firing rates in *rand* networks, i.e., 1 Hz (“*Scaled I*” networks; **Figure 2A, C, D, F**). Under these conditions, activity within assemblies was amplified substantially in comparison to the corresponding neurons in *rand* networks (“pseudo-assembly”) whereas activity outside assemblies was substantially reduced (**Figure 2E, G** and **Supplementary Figure 2A**). Hence, non-specific scaling of inhibition resulted in a large and biologically unrealistic divergence of firing rates (**Supplementary Figure 2B**) that nearly exhausted the dynamic range of individual neurons in the population, indicating that homeostatic global inhibition is insufficient to maintain a stable firing rate distribution. We further observed that neurons within activated assemblies produced regular spike trains (**Supplementary Figure 2A, C**), indicating that the balanced regime was no longer maintained.

In *rand* networks, correlations between E and I synaptic conductances in individual neurons were slightly above zero (**Figure 2H**), presumably because of stochastic inhomogeneities in the connectivity (Pehlevan and Sompolinsky, 2014). In *Scaled I* networks, correlations remained near zero, indicating that E assemblies by themselves did not enhance E/I co-tuning (**Figure 2H**). *Scaled I* networks with structured E but random I connectivity can therefore not account for the stability, synaptic balance, and E/I co-tuning observed experimentally (Rupprecht and Friedrich, 2018).

We next created structured networks with more precise E/I balance by including I neurons within assemblies. We first selected the 25 I neurons that received the largest number of connections from the 100 E neurons of an assembly. The connectivity between these two sets of neurons was then enhanced by two procedures to generate E/I assemblies: (1) in “*Tuned I*” networks, the probability of I-to-E connections was increased by a factor $\otimes$ while E-to-I connections remained unchanged. (2) In “*Tuned E+I*” networks, the probability of I-to-E connections was increased by $\otimes$ and the probability of E-to-I connections was increased by $\oslash$ (**Figure 2A**, **Supplementary Figure 1B**). As for “*Scaled I*” networks, $\otimes$ and $\oslash$ were adjusted to obtain mean population firing rates of ~1 Hz in response to learned odors (**Figure 2C, D, F**). The other observables used to constrain the *rand* networks remained unaffected (**Supplementary Figure 1 C-E**). Activity of networks with E assemblies could not be stabilized around 1 Hz by increasing connectivity from subsets of I neurons receiving dense input from activated mitral cells (**Supplementary Figure 1G-H**; Sadeh and Clopath, 2020).

In *Tuned* networks, correlations between E and I conductances in individual neurons were significantly higher than in *rand* or *Scaled I* networks (**Figure 2H**). To further analyze E/I co-tuning we projected synaptic conductances of each neuron onto a line representing the E/I ratio




Figure 2.

Neuronal assemblies (memories).

A. Schematic of assemblies. Each assembly contains the 100 E neurons that are most densely connected to the mitral cells activated by a given odor. Connection probability between these E neurons is increased by a factor α . In *Scaled I* networks, weights of all I-to-E synapses are increased by a factor χ . In *Tuned* networks, the 25 I neurons that are most densely connected to the 100 E neurons are included in each assembly. In *Tuned I* networks, the probability of I-to-E connections within the assembly is increased by a factor β . In *Tuned I+E* networks, probabilities of I-to-E and E-to-I connectivity within the assembly are increased by factors β and γ , respectively. $n = 20$ networks with 15 assemblies each were simulated for each group. **B.** Firing rates averaged over all E or I neurons (full lines) and over all assembly neurons (dashed lines) as a function of α (mean $\pm$ SD, across 20 networks). **C.** Mean E neuron firing rates of *Scaled* (left) and *Tuned* (right) networks in response to learned odors as a function of connection strength and probability, respectively. Squares depict parameters used in following figures unless stated otherwise. **D.** Spike raster plots showing responses of 50 E neurons to two odors (O1: novel odor; O2: learned odor) in a *Scaled I* and the corresponding *Tuned* networks (same neurons and odors in the corresponding *rand* network are shown in [Figure 1C](#)). **E.** Top: Mean firing rates in response to learned odors as a function of time, averaged across assembly or non-assembly neurons. Bottom: Mean correlation between activity patterns evoked by a learned odor at different time points during odor presentation. Correlation coefficients were calculated between pairs of activity vectors composed of the mean firing rates of E neurons in 100 ms time bins. Activity vectors were taken from the same or different trials, except for the diagonal, where only patterns from different trials were considered. The pink bar indicates odor presentation. **F.** Mean firing rate in response to learned odors or novel odors. Each data point represents one network-odor pair ($n = 20$ networks, 10 odors). **G.** Amplification within and outside assemblies, calculated as the ratio between the mean firing rates in response to learned odors averaged over the same populations of neurons in a given structured network (*Scaled I*, *Tuned I*, or *Tuned E+I*) and the corresponding *rand* network. Each data point represents one network-odor pair. **H.** Quantification of co-tuning by the correlation between time-averaged E and I conductances in response to different odors, average across neurons. Each data point corresponds to one network ($n = 20$). Mean $\pm$ SD. **I.** Quantification of co-tuning by the ratio of dispersion of joint conductances along balanced and counter-balanced axes (inset; Methods; $n = 20$ networks).

expected in a balanced network (“balanced axis”) and onto an orthogonal line (“counter-balanced axis”; **Figure 2I** [inset](#), Methods). The ratio between the standard deviations (SD.) along these axes has been used previously to quantify E/I co-tuning in experimental studies ([Rupprecht and Friedrich, 2018](#) [inset](#)). This ratio was close to 1 in *rand* and *Scaled I* networks but significantly higher in *Tuned I* and *Tuned E+I* networks (**Figure 2I** [inset](#)). Hence, *Tuned* networks exhibited significant co-tuning along the balanced axis, as observed in pDp ([Rupprecht and Friedrich, 2018](#) [inset](#)).

In *Tuned* networks, activity within assemblies was higher than the mean activity but substantially lower and more irregular than in *Scaled I* networks (**Figure 2E**, **G** [inset](#); **Supplementary Figure 2A-C** [inset](#)). Unlike in *Scaled I* networks, mean firing rates evoked by novel odors were indistinguishable from those evoked by learned odors and from mean firing rates in *rand* networks (**Figure 2F** [inset](#), **Supplementary Figure 1F** [inset](#)). Hence, E/I co-tuning prevented excessive amplification of activity in assemblies without affecting global network activity, consistent with experimental observations in Dp and piriform cortex ([Chapuis and Wilson, 2011](#) [inset](#); [Frank et al., 2019](#) [inset](#)).

Effects of E/I assemblies on attractor dynamics

We next explored effects of assemblies on network dynamics. In *rand* networks, firing rates increased after stimulus onset and rapidly returned to a low baseline after stimulus offset (**Figure 2E** [inset](#)). Correlations between activity patterns evoked by the same odor at different time points and in different trials were positive but substantially lower than unity, indicating high variability (**Figure 2E** [inset](#)). Hence, *rand* networks showed transient and variable responses to input patterns, consistent with the typical behavior of generic balanced state networks ([Shadlen and Newsome, 1994](#) [inset](#); [Van Vreeswijk and Sompolinsky, 1996](#) [inset](#)). *Scaled* networks responded to learned odors with persistent firing of assembly neurons and high pattern correlations across trials and time, implying attractor dynamics ([Hopfield, 1982](#) [inset](#); [Khona and Fiete, 2022](#) [inset](#)), whereas *Tuned* networks exhibited transient responses and modest pattern correlations similar to *rand* networks (**Figure 2E** [inset](#)). Hence, *Tuned* networks did not exhibit stable attractor states, presumably because precise synaptic balance prevented strong recurrent amplification within E/I assemblies.

In classical memory networks, attractor dynamics drives pattern completion, the retrieval of the whole memory from noisy or corrupted versions of the learned input. We therefore tested whether networks with E/I assemblies performed pattern completion using two different approaches. In computational studies, pattern completion is often assessed by testing whether targeted partial stimulation of an assembly can activate the assembly as a whole ([Sadeh and Clopath, 2021](#) [inset](#); [Vogels et al., 2011](#) [inset](#)). Stimulating subsets of E neurons in an assembly during baseline activity recruited the entire assembly in all structured pDp_{sim} networks (*Scaled* and *Tuned*) without a significant increase in the overall population activity (**Supplementary Figure 3A** [inset](#)).

During partial assembly stimulation, most pDp_{sim} neurons do not receive any input, which is unlikely to happen during odor presentation. Therefore, we also assessed pattern completion with more realistic variations of inputs, building on experimental studies in the olfactory system and hippocampus ([Wilson and Sullivan, 2011](#) [inset](#); [Yassa and Stark, 2011](#) [inset](#)). We morphed a novel odor into a learned odor (**Figure 3A** [inset](#)), or a learned odor into another learned odor (**Supplementary Figure 3B** [inset](#)), and quantified the similarity between morphed and learned odors by the Pearson correlation of the OB activity patterns (input correlation). We then compared input correlations to the corresponding pattern correlations among E neurons in pDp_{sim} (output correlation). In *rand* networks, output correlations increased linearly with input correlations but did not exceed them (**Figure 3B** [inset](#) and **Supplementary Figure 3B** [inset](#)), consistent with the absence of pattern completion in generic random networks ([Babadi and Sompolinsky, 2014](#) [inset](#); [Marr, 1969](#) [inset](#); [Schaffer et al., 2018](#) [inset](#); [Wiechert et al., 2010](#) [inset](#)). *Scaled I* networks, in contrast, showed typical signatures of pattern completion: output correlations increased abruptly as a morphed odor approached a learned odor and eventually exceeded the corresponding input correlations (**Figure 3B** [inset](#) and **Supplementary Figure 3B** [inset](#)). In *Tuned* networks, output correlations were higher than in *rand*

networks but remained lower than input correlations and did not increase steeply near the learned odor (**Figure 3B** and **Supplementary Figure 3B**). Hence, networks with E/I assemblies did not show signatures of pattern completion in response to naturalistic stimuli, consistent with the absence of strong recurrent amplification within assemblies.

To further understand the effect of E/I assemblies on network dynamics we examined the synaptic currents underlying the modest amplification of activity in assemblies of *Tuned* networks that occurred as a learned odor was approached (**Figure 3C**). We measured the relative contribution of synaptic inputs from afferent and recurrent sources received by assembly and non-assembly neurons. In *rand* networks, activity in pseudo-assembly neurons increased linearly as the corresponding odor was approached (**Figure 3C**). This linear increase in activity can be attributed entirely to an increase in afferent input (**Figure 3D**) because the net recurrent input remained constant due to E/I balance (**Figure 3E**). In *Scaled* networks, assembly firing rates increased massively as the learned odor was approached (**Figure 3C**), reflecting a large increase in recurrent excitation (**Figure 3D-F**). In *Tuned* networks, in contrast, assembly firing rates increased only modestly due to a slightly larger increase in E than I currents near the learned odor (**Figure 3D-F**). Hence, the modest amplification by assemblies of *Tuned* networks involves an “imperfect” synaptic balance.

Geometry of activity patterns in networks with E/I assemblies

We next examined how E/I assemblies transform the geometry of neuronal representations, i.e. their organization in a state space where each axis represents the activity of one neuron or one pattern of neural covariance (Chung and Abbott, 2021; Gallego et al., 2017; Langdon et al., 2023). To address this general question, we created an odor subspace and examined its transformation by pDp_{sim} . The subspace consisted of a set of OB activity patterns representing four uncorrelated pure odors and mixtures of these pure odors. Pure odors were assigned to the corners of a square and mixtures were generated by selecting active mitral cells from each of the pure odors with probabilities depending on the relative distances from the corners (**Figure 4A**, Methods). Correlations between OB activity patterns representing pure odors and mixtures decreased approximately linearly as a function of distance in the subspace (**Figure 4B**). The odor subspace therefore represented a hypothetical olfactory environment with four odor sources at the corners of a square arena and concentration gradients within the arena. Locations in the odor subspace were visualized by the color code depicted in **Figure 4A**.

To examine how pDp_{sim} transforms this odor subspace we projected time-averaged activity patterns onto the first two principal components (PCs). As expected, the distribution of OB activity patterns in PC space closely reflected the square geometry of the subspace (**Figure 4C**). This geometry was largely preserved in the output of *rand* networks, consistent with the notion that random networks tend to maintain similarity relationships between input patterns (**Figure 4D**) (Babadi and Sompolsky, 2014; Marr, 1969; Schaffer et al., 2018). We next examined outputs of *Scaled* or *Tuned* networks containing 15 assemblies, two of which were aligned with pure odors (**Supplementary Figure 4A**). The four odors delineating the odor subspace therefore consisted of two learned and two novel odors. In *Scaled* networks, odor inputs were discretely distributed between three locations in state space representing the two learned odors and residual odors, consistent with the expected discrete attractor states (**Figure 4D, E**). *Tuned* networks, in contrast, generated continuous representations of the odor subspace (**Figure 4D**). The geometry of these representations was distorted in the vicinity of learned odors, which were more distant from most mixtures than novel odors. These geometric transformations were less obvious when activity patterns of *Tuned* networks were projected onto the first two PCs extracted from *rand* networks (**Supplementary Figure 4B**). Hence, E/I assemblies introduced local curvature into the coding space that partially separated learned from novel odors without disrupting the continuity of the subspace representation.




Figure 3.

Changes of output activity to gradual modifications of inputs.

A. Morphing of a novel odor N into a learned odor L. Morphed odors were generated by gradually changing the fractions of activated mitral cells defining odors N or L. The x-axis indicates the similarity of a morphed odor to odor L, here quantified by the correlation between the OB activity patterns representing the morphed and the learned odors (same x-axis in all panels). **B.** Pearson correlation between activity patterns evoked by a morphed odor and a learned odor in pDpsim as a function of the corresponding correlations in the OB (full line). The dashed line shows the correlation to the novel odor instead of the learned odor. A signature of pattern completion is a steep increase in the correlation between activity patterns representing a morphed and a learned odor in pDpsim, often exceeding the OB input correlation (gray shaded area). **C.** Firing rates in response to morphed odors averaged across assembly neurons (learned odor, full line) or pseudo-assembly neurons (novel odor, dashed line) as a function of similarity between the presented odor and learned odor L (see A). **D.** Excitatory input from the OB averaged over assembly neurons or non-assembly neurons. **E.** Sum of inhibitory and excitatory currents evoked by other pDpsim neurons (recurrent input), averaged over assembly or non-assembly neurons. **F.** Absolute inhibitory and excitatory recurrent currents averaged over assembly neurons. **B-F:** averages over 8 networks.




Figure 4.

Geometry of odor representations in pDpsim.

A. Odor subspace delineated by 2 learned (L, M) and 2 novel (N, O) pure odors at the vertices of a virtual square. Pixels within the square represent odor mixtures. **B.** Left: Pearson correlations between OB activity patterns defining pure odors. Right: Correlation between one pure odor (L; top left vertex) and all other odors. The odor from one vertex gradually morphs into the odor from another vertex. **C.** Projection of activity patterns in the OB onto the first 2 principal components (PCs). Colors represent patterns in the odor subspace as shown in **A**. **D.** Projection of activity patterns in pDpsim in response to the odor subspace onto the first 2 PCs. Representative examples of different pDp networks. **E.** Density plot showing distribution of data points and demonstrating clustering at distinct locations in PC space for *Scaled I* networks. **F.** Quantification of dimensionality of neural activity by the participation ratio: activity evoked by novel odors and related mixtures (left), activity evoked by learned odors and related mixtures (center), and activity evoked by all stimuli (right). Each data point represents one network (n = 8, mean ± SD.); dotted line represents the participation ratio of OB activity. **G.** Variance along the first 40 PCs extracted from activity patterns in different networks. Insets: variance along PCs 200 - 400. **H.** Angles between edges connecting a pure odor response and related versions thereof (inset). The analysis was performed using the first 400 PCs, which explained >75% of the variance in all networks. n = 168 angles per pure odor in each of 8 networks (Methods). Similar results were obtained in the full-dimensional space.

The curvature of the activity manifold in *Tuned* networks suggests that E/I assemblies confine population activity along specific dimensions of the state space. To test this hypothesis, we first quantified the dimensionality of odor-evoked activity by the participation ratio, a measure that depends on the eigenvalue distribution of the pairwise neural covariance matrix (Altan et al., 2021) (Methods). As expected, dimensionality was highest in *rand* networks and very low in *Scaled I* networks, reflecting the discrete attractor states (Figure 4F). In *Tuned* networks, dimensionality was high compared to *Scaled I* networks but lower than in *rand* networks (Figure 4F). The same trend was observed when we sampled data from a limited number of neurons to mimic experimental conditions and when we used other measures of dimensionality (Supplementary Figure 4C-E). Furthermore, when restraining the analysis to activity evoked by novel odors and related mixtures, dimensionality was similar between *rand* and *Tuned* networks (Figure 4F). These observations support the hypothesis that E/I assemblies locally constrain neuronal activity onto lower-dimensional subspaces (manifolds), consistent with recent findings showing effects of specific circuit motifs on the dimensionality of neural activity (Dahmen et al., 2023; Recanatesi et al., 2019).

We further tested this hypothesis by examining the local geometry of activity patterns around representations of learned and novel odors. If E/I assemblies locally confine activity onto manifolds, small changes of input patterns should modify output patterns along preferred dimensions near representations of learned but not novel odors. To test this prediction, we selected sets of input patterns including each pure odor and the seven most closely related mixtures. We then quantified the variance of the projections of their corresponding output patterns onto the first 40 PCs (Figure 4G). In *Tuned* networks, this variance decreased only slightly as a function of PC rank for activity patterns related to novel odors, indicating that patterns varied similarly in all directions. For patterns related to learned odors, in contrast, the variance was substantially higher in the direction of the first few PCs, implying variation along preferred dimensions. In addition, we measured the distribution of angles between edges connecting activity patterns representing pure odors and their corresponding related mixtures in PC space (Figure 4H, inset; Methods; Schoonover et al., 2021). Angles were narrowly distributed around 1 rad in *rand* networks but smaller in the vicinity of learned patterns in *Tuned* networks (Figure 4H).

Activity may be constrained non-isotropically by amplification along a subset of dimensions, by inhibition along other dimensions, or both. E neurons participating in E/I assemblies had large loadings on the first two PCs (Supplementary Figure 4 F, G) and responded to learned odors with increased firing rates as compared to the mean rates in *Tuned E+I* and *rand* networks. Firing rates of the remaining neurons, in contrast, were slightly lower than the corresponding mean rates in *rand* networks (Figure 2G). Consistent with these observations, the variance of activity projected onto the first few PCs was higher in *Tuned E+I* than in *rand* networks (Figure 4G) while the variance along higher-order PCs was lower (Figure 4G, inset). These observations indicate that local activity manifolds are delineated both by amplification of activity along preferred directions and by suppression of activity along other dimensions.

To obtain first insights into network mechanisms underlying effects of E/I assemblies on manifold geometry we simulated additional networks. Equalizing biophysical parameters of E and I neurons had little impact on network behavior (Supplementary Figure 5; Methods), indicating that observed effects do not depend on specific biophysical differences between E and I neurons. We next explored the hypothesis that the modest amplification of activity within E/I assemblies contributes to the geometry of neural manifolds in pDp_{sim} . First, we adjusted the number of I-to-E connections ($\otimes$) within assemblies of *Tuned E+I* networks to prevent amplification of activity within assemblies during presentation of learned odors (*Tuned[adjust]*, Supplementary Figure 6A-C, E). In contrast to *Tuned E+I* networks with control parameters, we observed neither a curvature of activity manifolds in PC space nor an asymmetry of d_M in these modified networks (Supplementary Figure 6F-H). Second, we reduced recurrent amplification in assemblies of

Scaled I networks by decreasing γ and χ (*Scaled[adjust]*, **Supplementary Figure 6A-C, E** [↗](#)). This resulted in distorted activity manifolds and an asymmetric increase in d_M , similar to observations in *Tuned* networks (**Supplementary Figure 6F-H** [↗](#)). Third, we artificially amplified activity in pseudo-assemblies of *rand* networks through concentration-dependent modifications of excitability instead of connectivity changes (**Supplementary Figure 6A** [↗](#)). This manipulation increased the mean population activity as compared to control *rand* networks, but it also induced geometric modifications similar to *Tuned* networks, although the manifold curvature in PC space appeared less pronounced (**Supplementary Figure 6B-H** [↗](#)). Hence, different manipulations had consistent effects on manifold geometry despite different effects on mean firing rates: signatures of manifold geometry observed in *Tuned* networks were mimicked by manipulations producing a modest, input-dependent amplification of firing rates in (pseudo-)assemblies but abolished by manipulations that suppressed this amplification. These results indicate that the modest amplification of activity within assemblies contributes to the geometry of odor representations in pDP_{sim} . However, further analyses are required to explore potential contributions of other processes and to understand the underlying mechanisms in more detail.

Pattern classification by networks with E/I assemblies

The lack of discrete attractor states raises the question how transformations of activity patterns by *Tuned* networks affect pattern classification. To quantify the association between an activity pattern and a class of patterns representing a given odor we computed the Mahalanobis distance (d_M). This measure quantifies the distance between the pattern and the class center, taking into account covariation of neuronal activity within the class. In bidirectional comparisons between patterns from different classes, the mean d_M may be asymmetric if neural covariance differs between classes. We first quantified d_M between representations of pure odors based on activity patterns across 80 E neurons drawn from the corresponding (pseudo-) assemblies. d_M was asymmetrically increased in *Tuned E+I* networks as compared to *rand* networks. Large increases were observed for distances between patterns related to learned odors and reference classes representing novel odors (**Figure 5A, B** [↗](#)). In the other direction, increases in d_M were smaller. Moreover, distances between patterns related to novel odors were almost unchanged (**Figure 5B** [↗](#)). Further analyses showed that increases in d_M in *Tuned E+I* networks involved both increases in the Euclidean distance between class centers and non-isotropic scaling of intra-class variability (**Supplementary Figure 7** [↗](#)). The geometric transformation of odor representations by E/I assemblies therefore particularly enhanced the discriminability of patterns representing learned odors.

To further analyze pattern classification, we performed multi-class quadratic discriminant analysis (QDA), an extension of linear discriminant analysis for classes with unequal covariances. Using QDA, we determined the probabilities that an activity pattern evoked by a mixture is classified as being a member of each of four classes representing the pure odors, thus mimicking a 4-way forced choice odor classification task. The four classes were defined by the distributions of activity patterns evoked by the pure and closely related odors. We then examined the classification probability of patterns evoked by mixtures with respect to a given class as a function of the similarity between the mixture and the corresponding pure odor (“target”) in the OB. As expected, the classification probability increased with similarity. Furthermore, in *Tuned E+I* networks, the classification probability of mixtures similar to a pure odor was significantly higher when the pure odor was learned (**Figure 5C** [↗](#)). Hence, E/I assemblies facilitated the classification of inputs related to learned patterns. Pattern classification by QDA could not be applied to outputs of *Scaled* networks because intra-class variability was too low (Methods).

E/I assemblies also improved pattern classification when activity patterns were classified based on correlations to template vectors representing the four odor classes (Methods; **Supplementary Figure 8** [↗](#)). This analysis further revealed hysteresis effects in *Scaled* but not in *Tuned* networks that impaired pattern classification (**Supplementary Figure 8** [↗](#)).




Figure 5.

Distance relationships and classification of odor representations.

A. Odor-evoked activity patterns used as class distributions and vectors (**B**) or training and test sets (**C**). Same odor subspace as in **Figure 4**. **B.** Top: schematic illustration of Mahalanobis distance d_M . Bottom: d_M between one activity vector (v) and reference classes (Q) in *rand* and *Tuned E+I* networks. d_M was computed based on the activity patterns of 80 E neurons drawn from the four (pseudo-) assemblies. Average of 50 draws. Note that d_M between patterns related to a learned odor and non-matching reference classes was higher in *Tuned E+I* networks. **C.** Pattern classification probability quantified by QDA. P_{Target} quantifies the probability that an activity pattern from the test set (odor mixtures, see **A**) is assigned to a target class from the training set (pure or closely related odor; see **A**). Left: classification probability as a function of the similarity (Pearson correlation) between the test and target odors in the OB (input patterns). Note enhanced classification probability for patterns evoked by odors similar to learned odors in *Tuned E+I* networks. Right: Classification probability for patterns similar to the training set (see **A**). Each data point represents one network ($n = 8$, mean $\pm$ SD; Wilcoxon signed-rank test, $p < 0.01$). **D.** The odor subspace of **Figure 4** was subdivided in classes reflecting different concentrations of one of the pure odors (Methods). Odor concentration is defined here as the percentage of the 150 mitral cells representing a given odor that are activated. **E.** The concentration of a given odor was regressed against the mean firing rates of a subset of neurons (Methods). Mean square error (squared difference between the actual concentration and the estimated concentration). **F.** Accuracy of a linear support vector machine (SVM) in predicting the concentration of a novel (left) or learned odor (right) in a mixture. Each data point represents one network ($n = 8$, mean $\pm$ SD.).

When neuronal subsets were randomly drawn not from assemblies but from the entire population, d_M was generally lower (**Figure 5B**). These results indicate that assembly neurons convey higher-than-average information about learned odors. Together, these observations imply that pDp_{sim} did not function as a categorical classifier but nonetheless supported the classification of learned odors, particularly when the readout focused on assembly neurons. Conceivably, classification may be further enhanced by optimizing the readout strategy, for example, by a learning-based approach. However, modeling biologically realistic readout mechanisms requires further experimental insights into the underlying circuitry.

Additional computational properties of networks with E/I assemblies

E/I assemblies may not only support discrete classification but also other tasks. We hypothesized that continuous manifolds provide a distance metric to quantitatively evaluate the relatedness between an input and other stimuli. The geometric modifications induced by E/I assemblies may specifically provide a metric between a new sensory input and previously learned stimuli. To explore these hypotheses, we first used multilinear regression to predict the concentration of a pure odor in subspace mixtures (**Figure 4A**) based on population activity. Highest prediction errors were observed for *Scaled I* networks, presumably because their discontinuous outputs poorly represent gradual concentration differences (**Figure 5D, E**). Lowest errors were observed for concentration estimates of learned odors in *Tuned* networks. We further trained support vector machines to estimate the concentrations of pure odors in the subspace mixtures and again found highest accuracy for learned odors in *Tuned* networks (**Figure 5F**). These results show that networks with E/I assemblies enable the quantitative analysis of selected stimulus components in the presence of distractors, presumably due to the geometric properties of the underlying neural manifold.

We next examined whether E/I assemblies can stabilize activity when multiple memories are successively added into a network, as in a continual learning process. In the absence of precise E/I balance, the divergence of firing rates and, thus, the risk of network instability are expected to increase as assemblies accumulate. *Tuned* networks may be more resistant against these risks, particularly when assemblies overlap, because activity is controlled more precisely within each additional assembly. To test this hypothesis, we examined responses of *Tuned E+I* and *Scaled I* networks to an additional odor subspace where four of the six pairwise correlations between the pure odors were clearly positive (range, 0.27 – 0.44; **Figure 6A**). We then compared networks with 15 randomly created assemblies to networks with two additional assemblies representing two of the correlated pure odors. These two additional assemblies had, on average, 16% of neurons in common due to the similarity of the odors. In *Tuned* networks, introducing additional overlapping assemblies selectively increased the discriminability of the corresponding odor representations in PC space (**Figure 6B, C**), as observed for uncorrelated assemblies (**Figure 4D**), but population firing rates remained almost unchanged (**Figure 6B, C**). In *Scaled I* networks, in contrast, creating two additional memories resulted in a convergence of the corresponding representations in PC space and in a substantial increase in firing rates, particularly in response to the learned and related odors (**Figure 6C**). These observations are consistent with the assumption that precise balance in E/I assemblies protects networks against instabilities during continual learning, even when memories overlap. We further observed that, in this regime of higher pattern similarity, d_M was again increased upon learning, particularly between learned odors and reference classes representing other odors (**Supplementary Figure 9**). E/I assemblies therefore consistently increased d_M in a directional manner under different conditions.




Figure 6.

Representation of correlated patterns and resilience against additional memories.

A. Subspace delineated by four positively correlated odors (see Methods). Top: Correlations between pure odors. Bottom: Projection of OB activity patterns onto the first two PCs. **B.** Firing rates (top) and PC projection of output activity of a *Tuned E+I* network with 15 E/I assemblies that did not represent any of the four pure odors of the subspace. **C.** Firing rates (top) and PC projection of output activity after creation of two additional assemblies representing two of the pure odors (Y and Z). Left: *Tuned E+I* network. Right: *Scaled I* network. Note that in the *Scaled I* network, but not in the *Tuned E+I* network, firing rates evoked by newly learned odors were increased and patterns evoked by these odors were not well separated in PC space.

Discussion

A precisely balanced memory network constrained by pDp

Autoassociative memory networks map inputs onto output patterns representing learned information. Classical models proposed this mapping to be accomplished by discrete attractor states that are defined by assemblies of E neurons and stabilized by global homeostatic inhibition. However, as seen in *Scaled I* networks, global inhibition is insufficient to maintain a stable, biologically plausible firing rate distribution. This problem can be overcome by including I neurons in assemblies, which leads to more precise synaptic balance. To explore network behavior in this regime under biologically relevant conditions we built a spiking network model constrained by experimental data from pDp. The resulting *Tuned* networks reproduced additional experimental observations that were not used as constraints including irregular firing patterns, lower output than input correlations, and the absence of persistent activity. Hence, pDp_{sim} recapitulated characteristic properties of a biological memory network with precise synaptic balance.

Neuronal dynamics and representations in precisely balanced memory networks

Simulated networks with global inhibition showed attractor dynamics and pattern completion, consistent with classical attractor memory. However, the distribution of firing rates broadened as connection density within assemblies increased, resulting in unrealistically high (low) rates inside (outside) assemblies and, consequently, in a loss of synaptic balance. Hence, global inhibition was insufficient to stabilize population activity. In networks with E/I assemblies, in contrast, firing rates remained within a realistic range and the balanced state was maintained. Such *Tuned* networks showed no discrete attractor states but transformed the geometry of the coding space by confining activity to continuous manifolds near representations of learned inputs. This observation is consistent with the hypothesis that E/I co-tuning decreases the probability of multistable attractor states as compared to balanced networks with global inhibition or lateral inhibition between assemblies (Boerlin et al., 2013 [DOI](#); Hennequin et al., 2018 [DOI](#); Wu and Zenke, 2021 [DOI](#)).

Geometric transformations in *Tuned* networks may be considered as intermediate between two extremes: (1) geometry-preserving transformations as, for example, performed by many random networks (Babadi and Sompolinsky, 2014 [DOI](#); Marr, 1969 [DOI](#); Schaffer et al., 2018 [DOI](#)), and (2) discrete maps as, for example, generated by discrete attractor networks (Freeman and Skarda, 1985 [DOI](#); Hopfield, 1982 [DOI](#); Khona and Fiete, 2022 [DOI](#)) (Figure 7 [DOI](#)). We found that transformations became more discrete map-like when amplification within assemblies was increased and precision of synaptic balance was reduced. Likewise, decreasing amplification in assemblies of *Scaled* networks changed transformations towards the intermediate behavior, albeit with broader firing rate distributions than in *Tuned* networks (Supplementary Figure 6B [DOI](#)). These observations indicate that a modest amplification of activity within assemblies contributes to geometric modifications of activity manifolds in *Tuned* networks, but other factors such as structured inhibition may also contribute. Hence, further analyses are required to obtain a deeper mechanistic understanding of manifold geometry in networks with E/I assemblies. Nonetheless, our results suggest that precise synaptic balance may generally favor intermediate over discrete transformations because this regime tends to linearize the relationship between the mean input and output firing rates of neuronal populations (Baker et al., 2020 [DOI](#); Denève and Machens, 2016 [DOI](#)).




Figure 7.

Schematic of geometric transformations.

A. Randomly connected networks tend to preserve the geometry of coding space. Such networks can support neuronal computations, e.g., by projecting activity patterns in a higher-dimensional coding space for pattern classification. **B.** We found that balanced networks with E/I assemblies transform the geometry of representations by locally restricting activity onto manifolds. These networks stored information about learned inputs while preserving continuity of the coding space. Such a geometry may support fast classification, continual learning and cognitive computations. Note that the true manifold geometry cannot be visualized appropriately in 2D because activity was “focused” in different subsets of dimensions at different locations of coding space. As a consequence, the dimensionality of activity remained substantial. **C.** Neuronal assemblies without precise balance established discrete attractor states, as observed in memory networks that store information as discrete items. Networks establishing locally defined activity manifolds (**B**) may thus be considered as intermediates between networks generating continuous representations without memories (**A**) and classical memory networks with discrete attractor dynamics (**C**).

E/I assemblies increased variability of population activity along preferred directions of state space and reduced its dimensionality in comparison to *rand* networks. Nonetheless, dimensionality remained high compared to *Scaled* networks with discrete attractor states. These observations indicate that geometric transformations in *Tuned* networks involved (1) a modest amplification of activity in one or a few directions aligned to the assembly, and (2) a modest reduction of activity in other directions. E/I assemblies therefore created a local curvature of coding space that “focused” activity in a subset of dimensions and, thus, stored information in the geometry of coding space.

As E/I assemblies were small relative to the total size of the E neuron population, stored information may be represented predominantly by small neuronal subsets. Consistent with this hypothesis, d_M was increased and the classification of learned inputs by QDA was enhanced when activity was read out from subsets of assembly neurons as compared to random neuronal subsets. Moreover, signatures of pattern completion were found in the activity of assemblies but not in global pattern correlations. The retrieval of information from networks with small E/I assemblies therefore depends on the selection of informative neurons for readout. Unlike in networks with global attractor states, signatures of memory storage may thus be difficult to detect experimentally without specific knowledge of assembly memberships.

Computational functions of networks with E/I assemblies

In theory, precisely balanced networks with E/I assemblies may support pattern classification despite high variability of spike trains and the absence of discrete attractor states (Denève and Machens, 2016 [DOI](#)). Indeed, we found in *Tuned E+I* networks that input patterns were classified successfully by generic classifiers, particularly relative to learned inputs. Analyses based on the Mahalanobis distance d_M indicate that classification of learned inputs was enhanced by two effects: (1) local manifolds representing learned odors became more distant from representations of other odors due to a modest increase in firing rates within E/I assemblies, and (2) the concomitant increase in variability was not isotropic, remaining sufficiently low in directions that separated novel from learned patterns. Hence, information contained in the geometry of coding space can be retrieved by readout mechanisms aligned to activity manifolds. Efficient readout mechanisms may thus integrate activity primarily from assembly neurons, as mimicked in our QDA-based pattern classification. This notion is consistent with the finding that the integrated activity of E/I assemblies can be highly informative despite variable firing of individual neurons (Boerlin et al., 2013 [DOI](#); Denève et al., 2017 [DOI](#); Denève and Machens, 2016 [DOI](#)). It will thus be interesting to explore how the readout of information from local manifolds could be further optimized.

Networks with E/I assemblies may also perform computations other than pattern classification. A toy model of continual learning indicates that precise balance can stabilize firing rate distributions when new memories are added, helping to prevent catastrophic failures. Furthermore, we found that continuous manifolds in *Tuned* networks support metric analyses, for example, to determine the relatedness of a (potentially complex) input to previously learned stimuli. Such tasks are not supported by discrete attractor states because information about gradual differences between inputs is lost by pattern completion. E/I assemblies can thus enhance the quantitative analysis of relevant information through changes in manifold geometry.

Conceivably, E/I assemblies may have further consequences for neuronal computation. Unlike discrete attractor networks, *Tuned* networks do not exhibit persistent activity, suggesting that they mediate fast computations rather than short-term memory functions. Fast classification may, for example, be important to interpret dynamical sensory inputs on a moment-to-moment basis. Moreover, the representation of learned inputs by small neuronal subsets, rather than global activity states, raises the possibility that multiple inputs can be classified simultaneously. Generally, the absence of discrete attractor states indicates that information is not stored in the form of distinct items. Rather, E/I assemblies cause geometric modifications of a continuous coding space that result in an overrepresentation of learned (relevant) inputs at the expense of other

stimuli. These pattern transformations may thus contribute to different classification and learning processes by re-formatting and extracting relevant information for further processing by distributed networks. Such a function would be loosely related to computations within layers of artificial neuronal networks and consistent with the notion that piriform cortex is embedded in a larger network comprising multiple telencephalic brain areas (Haberly, 2001 [↗](#)).

Balanced state networks with E/I assemblies as models for olfactory cortex

Piriform cortex and Dp have been proposed to function as attractor-based memory networks for odors. Consistent with this hypothesis, pattern completion and its modulation by learning has been observed in piriform cortex of rodents (Barnes et al., 2008 [↗](#); Chapuis and Wilson, 2011 [↗](#)). However, odor-evoked firing patterns in piriform cortex and Dp are typically irregular, variable, transient and less reproducible than in the OB even after learning (Jacobson et al., 2018 [↗](#); Pashkovski et al., 2020 [↗](#); Schoonover et al., 2021 [↗](#); Yaksi et al., 2009 [↗](#)), indicating that activity does not converge onto stable attractor states. Balanced networks with E/I assemblies, in contrast, are generally consistent with these experimental observations. Alternative models for pattern classification in the balanced state include networks endowed with short-term plasticity, which respond to stimuli with an initial amplification phase followed by a tonic inhibition-stabilized state (Wu and Zenke, 2021 [↗](#)), or mechanisms related to “balanced amplification”, which typically generate pronounced activity transients (Ahmadian and Miller, 2021 [↗](#); Murphy and Miller, 2009 [↗](#)). However, it has not been explored whether these models can be adapted to reproduce characteristic features of Dp or piriform cortex.

Our results generate predictions to test the hypothesis that E/I assemblies establish local manifolds in Dp: (1) odor-evoked population activity should be constrained onto manifolds, particularly in response to learned odors. (2) Learning should increase the magnitude and asymmetry of d_M between odor representations. (3) Activity evoked by learned and related odors should exhibit lower dimensionality and more directional variability than activity evoked by novel odors. (4) Careful manipulations of inhibition may unmask assemblies by increasing amplification. These predictions may be addressed experimentally by large-scale measurements of odor-evoked activity after learning. The direct detection of E/I assemblies will ultimately require dense reconstructions of large neuronal networks at synaptic resolution. Given the small size of Dp, this challenge may be addressed in zebrafish by connectomics approaches based on volume electron microscopy (Denk et al., 2012 [↗](#); Friedrich and Wanner, 2021 [↗](#); Kornfeld and Denk, 2018 [↗](#)).

The hypothesis that memory networks contain E/I assemblies and operate in a state of precise synaptic balance can be derived from the basic assumptions that (1) synaptic plasticity establishes assemblies and (2) that firing rate distributions remain stable as network structure is modified by experience (Barron et al., 2017 [↗](#); Hennequin et al., 2017 [↗](#)). Hence, *Tuned* networks based on Dp may also reproduce features of other recurrently connected brain areas such as hippocampus and neocortex, which also operate in a balanced state (Renart et al., 2010 [↗](#); Sadeh and Clopath, 2020b [↗](#); Shadlen and Newsome, 1994 [↗](#); Znamenskiy et al., 2024 [↗](#)). Future experiments may therefore explore representations of learned information by local manifolds also in cortical brain areas.

Methods

pDp spiking network model

pDp_{sim} consisted of 4000 excitatory (E) and 1000 inhibitory (I) neurons which were modeled as adaptive leaky integrate-and-fire units with conductance-based synapses of strength w .

A spike emitted at time t by the presynaptic neuron y from population Y triggered an increase in the conductance g_{YX} in the postsynaptic neuron x , connected to y with strength w_{yx} :

$$\tau_{syn,Y} \frac{dg_{YX}}{dt} = -g_{YX} + \tau_{syn,Y} \sum_y w_{yx} \delta(t - t_{spike,y}) \quad (\text{equation 1})$$

The synaptic time constant $\tau_{syn,Y}$ was set to 10 ms for inhibitory synapses and 30 ms for excitatory synapses, except for the simple model where $\tau_{syn,Y}$ was 10 ms for both synapses. Neuron x received synaptic inputs from the olfactory bulb (OB) as well as from the different local neuronal populations P . Its membrane potential V_x evolved according to:

$$g_{rest,X} \tau_X \frac{dV_x}{dt} = g_{rest,X} (E_{rest,X} - V_x) + g_{OBX} (E_{exc} - V_x) + \sum_P g_{PX} (E_P - V_x) - z_x \quad (\text{equation 2})$$

When the membrane potential reached a threshold V_{th} , the neuron emitted a spike and its membrane potential was reset to E_{rest} and clamped to this value during a refractory period τ_{ref} . Excitatory neurons were endowed with adaptation with the following dynamics (Brette and Gerstner, 2005 [↗](#)):

$$\tau_a \frac{dz}{dt} = a(V_x - E_{rest,E}) - z, \text{ with } z \rightarrow z + b \text{ after each spike} \quad (\text{equation 3}).$$

In inhibitory neurons, z was set to 0 for simplicity.

The neuronal parameters of the model are summarized in **Table 1** [↗](#). The values of the membrane time constant τ_X , the resting conductance g_{rest} and resting potential E_{rest} of excitatory neurons, and the inhibitory and excitatory reversal potential E_{inh} and E_{exc} are in the range of experimentally measured values (Rupprecht and Friedrich, 2018 [↗](#) and Blumhagen et al., 2011 [↗](#)). The remaining parameters were then fitted to fulfill two conditions (derived from unpublished experimental observations): (1) the neuron should not generate action potentials in response to a step current injection of duration 500 ms and amplitude 15 pA, and (2) the mean firing rate should be on the order of tens of Hz when the amplitude of the step current is 100 pA. Furthermore, the firing rates of inhibitory neurons should be higher than the firing rates of excitatory neurons, as observed experimentally (unpublished data).

To verify that the behavior of pDp_{sim} was robust, we simulated 20 networks with different connection probabilities p_{YX} and synaptic strengths w_{YX} (**Table 2** [↗](#)). The connections between neurons were drawn from a Bernoulli distribution with a predefined $p_{YX} \leq 0.05$. As a consequence, each neuron received the same number of input connections. Care was also taken to ensure that the variation in the number of output connections across neurons was low in pDp_{sim} : connections of a presynaptic neuron y to postsynaptic neurons x were randomly deleted when their total number exceeded the average number of output connections by $\geq 5\%$, or added when they were lower by $\geq 5\%$.

The connection strengths w_{YX} were then fitted to reproduce experimental observations in pDp (five observables in total, see below and **Figure 1** [↗](#)). For this purpose, a lower and an upper bound for w_{YX} were set such that the amplitudes of single EPSPs and IPSPs were in the biologically plausible range of 0.2 to 2 mV. w_{OE} was then further constrained to maintain the odor-evoked, time-averaged g_{OE} in the range of experimental values (Rupprecht and Friedrich, 2018 [↗](#)). Once w_{OE} was fixed, the lower bound of w_{EE} was increased to obtain a network rate >10 Hz in the absence of inhibition. A grid search was then used to refine the remaining w_{YX} .

Olfactory bulb input

Each pDp_{sim} neuron received external input from the OB, which consisted of 1500 mitral cells spontaneously active at 6 Hz. Odors were simulated by increasing the firing rate of 150 randomly selected mitral cells. Firing rates of these “activated” mitral cells were drawn from a discrete

Table 1

Values of the neuronal parameters. The superscripts indicate the reference where the experimental measurements can be found. ¹ Rupperecht and Friedrich, 2018 [↗](#) ² Blumhagen et al., 2011 [↗](#)

Neuronal parameters	Symbol	Value		
		Excitatory neuron	Inhibitory neuron	Simple model
Membrane time constant	τ_X	85 ms ^{1,2}	50 ms ²	68 ms
Resting conductance	g_{rest}	1.35 nS ¹	0.9 nS	1.1 nS
Resting potential	E_{rest}	-60 mV ^{1,2}	-65 mV	-62 mV
Excitatory reversal potential	E_{exc}	0 mV ¹	0 mV	0 mV
Inhibitory reversal potential	E_{inh}	-70 mV ¹	-70 mV	-70 mV
Spiking threshold	V_{th}	-38 mV ¹	-45 mV	-41 mV
Reset potential	E_{reset}	-60 mV	-65 mV	-62 mV
Refractory period	τ_{ref}	8 ms	8 ms	8 ms
Adaptation time constant	τ_a	40 ms	/	20 ms
Subthreshold adaptation	a	1 nS	/	0.5 nS
Spike-triggered adaptation	b	10 pA	/	5 pA

Table 2

Values of the connectivity parameters of different networks {probability p_{YX} and synaptic strength w_{YX} in pS}. In the simple model, each neuron received the same number of inputs from each population.

Network structure	OB → E	OB → I	E → E	E → I	I → E	I → I
A (n=5)	{0.02, 128}	{0.01, 68}	{0.05, 128}	{0.04, 68}	{0.05, 480}	{0.04, 250}
B (n=5)	{0.02, 128}	{0.01, 66}	{0.05, 128}	{0.04, 66}	{0.05, 450}	{0.04, 210}
C (n=5)	{0.02, 128}	{0.01, 68}	{0.05, 108}	{0.02, 80}	{0.05, 520}	{0.02, 310}
D (n=5)	{0.03, 95}	{0.02, 42}	{0.05, 128}	{0.04, 58}	{0.05, 590}	{0.04, 270}
Simple (n=5)	{0.03, 160}	{0.03, 160}	{0.025, 370}	{0.025, 370}	{0.1, 1010}	{0.1, 1010}

uniform distribution ranging from 8 to 32 Hz and their onset latencies were drawn from a discrete uniform distribution ranging from 0 to 200 ms. An additional 75 randomly selected mitral cells were inhibited. Firing rates and latencies of these neurons were drawn from discrete uniform distributions ranging from 0 to 5 Hz and from 0 to 200 ms, respectively. After odor onset, firing rates decreased exponentially with a time constant of 1, 2 or 4 s (equally distributed). Spikes were then generated from a Poisson distribution, and this process was repeated to create trial-to-trial variability. As all odors had almost identical firing patterns, the total OB input did not vary much across odors. In **Figures 1** – 3, the odor set consisted of 10 novel and/or 10 learned odors, all of which were uncorrelated (pattern correlations close to zero). Odors were presented for 2 s and separated by 1 s of baseline activity.

Olfactory subspaces comprised 121 OB activity patterns. Each pattern was represented by a pixel in a 11 x 11 square. The pixel at each vertex corresponded to one pure odor with 150 activated and 75 inhibited mitral cells as described above, and the remaining pixels corresponded to mixtures. When pure odors were correlated (**Figure 6**), adjacent pure odors shared half of their activated and half of their inhibited mitral cells. We generated 8 different trajectories within the square, each visiting all possible virtual odor locations for 1s. Each trajectory thus consisted of 121s of odor presentation, and trajectories were separated by 2 s of baseline activity. The dataset for analysis therefore comprised 968 activity patterns (8 trajectories x 121 odors). The total number of activated and inhibited cells at each location in the virtual square remained within the range of $150 \pm 10\%$ and $75 \pm 10\%$, respectively. The fraction of activated and inhibited mitral cells from a given pure odor (its concentration) decreased with the distance from the corresponding vertex as shown in **Table 3**. At each location within the square and for each trajectory, mitral cells that remained activated in the mixture were randomly selected from the pool of C mitral cells with the shortest latencies from each pure odor. C decreased with the distance from the vertex as shown in **Table 4**. The firing rate of each selected mitral cell varied ± 1 Hz around its rate in response to the pure odor. The identity, but not the firing rate, of the activated mitral cells therefore changed gradually within the odor subspace. This procedure reflects the experimental observation that responses of zebrafish mitral cells to binary odor mixtures often resemble responses to one of the pure components (Tabor et al., 2004).

Assemblies

Unless noted otherwise, *Scaled* and *Tuned* networks contained 15 assemblies (“memories”). An assembly representing a given odor contained the 100 E neurons that received the highest density of inputs from the corresponding active mitral cells. Hence, the size of assemblies was substantially smaller than the total population, consistent with the observation that only a minority of neurons in Dp or piriform cortex are activated during odor stimulation (Miura et al., 2012; Stettler and Axel, 2009; Yaksi et al., 2009) and upregulate *cfos* during olfactory learning (Meissner-Bernard et al., 2018). We then rewired assembly neurons: additional connections were created between assembly neurons, and a matching number of existing connections between non-assembly and assembly neurons were eliminated. The number of input connections per neuron therefore remained unchanged. A new connection between two assembly neurons doubled the synaptic strength w_{EE} if it added to an existing connection. As a result of this rewiring, the connection probability within the assembly increased by a factor γ relative to the baseline connection probability.

In *Scaled* networks, w_{IE} was increased globally by a constant factor χ . In *Tuned* networks, connections were modified between the 100 E neurons of an assembly and the 25 I neurons that were most densely connected to these E neurons, using the same procedure as for E-to-E connections. In *Tuned I* networks, only I-to-E connections were rewired, while in *Tuned E+I networks*, both I-to-E and E-to-I connections were rewired (**Table 5**). Whenever possible,

Table 3

Percentage of mitral cells out of the 150 mitral cells defining one pure odor (top left corner) that are activated in mixtures. Percentages may be interpreted as relative odor concentration. Values given here are for uncorrelated pure odors.

100	90	80	70	60	50	40	30	20	10	0
90	90	80	70	60	50	40	30	15	5	0
80	80	70	60	50	40	35	25	15	5	0
70	70	60	50	50	35	30	20	15	5	0
60	60	50	50	40	30	25	20	15	5	0
50	50	40	35	30	25	25	15	10	5	0
40	40	35	30	25	25	15	10	10	5	0
30	30	25	20	20	15	10	10	5	5	0
20	15	15	15	15	10	10	5	5	5	0
10	5	5	5	5	5	5	5	5	0	0
0	0	0	0	0	0	0	0	0	0	0

Table 4

Percentage of activated cells available for selection for uncorrelated pure odors. C is obtained by multiplying the values by 1.5.

100	100	100	90	80	70	60	50	40	30	0
100	100	100	90	80	70	60	50	35	25	0
100	100	90	60	70	60	55	45	35	25	0
90	90	80	70	70	55	50	40	35	25	0
80	80	70	70	60	50	45	40	35	25	0
70	70	60	55	50	45	45	35	30	25	0
60	60	55	50	45	45	35	30	30	25	0
50	50	45	40	40	35	30	30	25	25	0
40	35	35	35	35	30	30	25	25	25	0
30	25	25	25	25	25	25	25	25	0	0
0	0	0	0	0	0	0	0	0	0	0

parameters ζ , θ , ϕ , χ that generated networks with less than 15% change in population firing rates compared to the corresponding *rand* network were selected. In [Figure 6](#), two additional assemblies were created in *Scaled* or *Tuned* networks without adjusting any parameters.

Observables

All variables were measured in the E population and time-averaged over the first 1.5 seconds of odor presentations, unless otherwise stated.

1. The *firing rate* is the number of spikes in a time interval T divided by T.
2. g_{OE} is the mean conductance change in E neurons triggered by spikes from the OB.
3. g_{syn} is the total synaptic conductance change due to odor stimulation, calculated as the sum of g_{OE} , g_{EE} and g_{IE} . g_{EE} and g_{IE} are the conductance changes contributed by E synapses and I synapses, respectively.
4. The *percentage of recurrent input* quantifies the average contribution of the recurrent excitatory input to the total excitatory input in E neurons. It was defined for each excitatory neuron as the ratio of the time-averaged g_{EE} to the time-averaged total excitatory conductance ($g_{EE} + g_{OE}$) multiplied by 100.
In (2,3,4), the time-averaged E and I synaptic conductances during the 500 ms before odor presentation were subtracted from the E and I conductances measured during odor presentation for each neuron.
5. In addition, we required the Pearson correlation between activity patterns to be close to zero in response to uncorrelated inputs. The Pearson correlation between pairs of activity vectors composed of the firing rate of E neurons was averaged over all possible odor pairs.

Co-tuning

Co-tuning was quantified by two different procedures: (1) For each neuron, we calculated the Pearson correlation between the time-averaged E and I conductances in response to 10 learned odors. (2) As described in ([Rupprecht and Friedrich, 2018](#)), we projected observed pairs of E and I conductances onto a “balanced” and “counter-balanced” axis. The balanced axis was obtained by fitting a linear model without constant to the E and I conductances of 4000*10 neuron-learned odor pairs. The resulting model was a constant I/E ratio (~1.2) that defined a membrane potential close to spike threshold. The counter-balanced axis was orthogonal to the balanced axis. For each neuron, synaptic conductances were projected onto these axes and their dispersions quantified by the standard deviations.

Characterization of population activity in state space

Principal Component Analysis (PCA) was applied to the OB activity patterns of the odor subspace or to the corresponding activity patterns across E neurons in pDp_{sim} ($8 \times 121 = 968$ patterns, each averaged over 1s).

The *participation ratio* PR provides an estimate of the maximal number of principal components (PC) required to recapitulate the observed neuronal activity. It is defined as $PR = \frac{(\sum \lambda_i)^2}{(\sum \lambda_i^2)}$, where λ_i are the eigenvalues obtained from PCA (variance along each PC).

For angular analyses we projected activity patterns onto the first 400 PCs, which was the number of PCs required to explain at least 75% of the variance in all networks. We measured the *angle* θ between the edges connecting the pattern p evoked by a pure odor to two patterns s_y and s_z . s_y and s_z were patterns evoked by 2 out of the 7 odors that were most similar to the pure odor (168 angles in total). θ was defined as $\theta_{yz} = \cos^{-1} \frac{(s_y - p) \cdot (s_z - p)}{\|s_y - p\| \|s_z - p\|}$. This metric is sensitive to non-uniform expansion and other non-linear transformations.

	Figure 2-6	Supplementary Figure 6 (mean)
Scaled I	$\alpha=5, \chi=1.4$	Scaled[adjust]: $\alpha=3.7, \chi=1.07$
Tuned I	$\alpha=5, \beta=18$	/
Tuned E+I	$\alpha=5, \beta=4, \gamma=3$	Tuned[adjust]: $\alpha=5, \beta=11.5, \gamma=3;$

Table 5

Values of $\alpha, \beta, \gamma, \chi$ used in the simulations.

α : increase in E-E connection probability within assemblies. β : increase in I to E connection probability. γ : increase in E to I connection probability.

The *Mahalanobis distance* d_M is defined as $d_M(v, Q) = \sqrt{(v - \mu)^T S^{-1} (v - \mu)}$, where v is a vector representing an activity pattern, and Q a reference class consisting of a distribution of activity patterns with mean μ and covariance matrix S .

Classification: four-way forced choice task

To assess the assignment of odor-evoked patterns to representations of pure odors in the odor subspace, we used *Quadratic Linear Discriminant* (QDA) analysis, a non-linear classifier that takes into account the separation and covariance patterns of different classes (Ghojogh and Crowley, 2019). The training set consisted of the population response to multiple trials of each of the 4 pure odors, averaged over the first and second half of the 1 s odor presentation. To increase the number of training data in each of the 4 odor classes, the training set also included population response to odors that were closely related to the pure odor (Pearson correlation between OB response patterns > 0.6). Analyses were performed using subsets of 80 neurons (similar results were obtained using 50 – 100 neurons). These neurons were randomly selected (50 iterations) either from the (pseudo-) assemblies representing the pure odors (400 E neurons; pseudo-assemblies in *rand* networks and for novel odors) or from the entire population of E neurons. We verified that the data came from a Gaussian mixture model. The trained classifier was then applied to activity patterns evoked by the remaining odors of the subspace (test set; correlation with pure odors < 0.6). Each pattern x from the test set was assigned to the class k that maximized the discriminant function $\delta_k(x) = -\frac{1}{2} \log |S_k| - \frac{1}{2} (x - \mu_k)^T S_k^{-1} (x - \mu_k) + \log \pi_k$, where S_k is the neuronal covariance matrix of each odor class k and π_k is the prior probability of class k . This discriminant function is closely related to d_M . The subsampling of neurons ensured the invertibility of S_k in *Tuned* and *rand* networks. In *Scaled* networks, S_k is low-rank and therefore not invertible, and QDA cannot be applied.

Template matching (Supplementary Figure 8) was performed in a similar way except that each test pattern x was assigned to the class k that maximized the correlation between pattern x and a randomly selected pattern in each class (10 iterations).

Classification: odor concentration

As mentioned above, an odor is represented by 150 activated mitral cells, a subset of which remain activated in mixtures. We define the percentage of remaining activated cells as the concentration of the odor. We performed multiple linear regression to predict the concentration of a given odor based on the firing rate of 80 E neurons averaged over the first second of odor presentation. As for QDA, these 80 neurons were randomly selected (50 iterations) from the (pseudo-) assemblies representing the pure odors (400 E neurons; pseudo-assemblies in *rand* networks and for novel odors). We included 4 or 8 trials to equilibrate the number of population responses across the different concentrations ($n=352$).

Using the Matlab function *fitcecoc*, we also trained a linear Support Vector Machine (SVM) model to assess the accuracy of different networks in predicting the concentration of an odor in a mixture. Mixtures were divided into 6 different concentration classes, with 15% increments (see Figure 5D). We included 2 to 6 trials to equilibrate the number of population responses within the different classes. The dataset consisted of 162 population responses of 80 randomly selected neurons (20 iterations) and was randomly partitioned in training and testing set (2-fold cross-validation, 20 iterations).


Simulations

Simulations were performed using Matlab and Python. Differential equations were solved using the forward Euler method and an integration time step of $dt = 0.1$ ms.

Data and code availability

Matlab code to simulate pDp_{sim} is available at <https://github.com/clairemb90/pDp-model>.


Supplementary Figure Legends



Supplementary Figure 1.

Structured networks reproduce key features of pDp.


A. Connection probability between activated (activ.) or non-activated (non-activ.) mitral cells and assembly or non-assembly neurons. **B.** Connection probability between classes of E and I neurons in a *rand*, a *Tuned I*, and a *Tuned E+I* network. **C.** EPSCs and IPSCs in a *Tuned I* network as observed in a hypothetical voltage clamp recording, averaged across neurons and odors. An equivalent plot for a *rand* network is shown in **Fig. 1E**. **D, E.** Values of observables for *rand* networks (same as **Figure 1G-H**) and different structured networks (*Scaled I*, *Tuned I*, *Tuned E+I*). **F.** Spike raster plots showing responses of 50 I neurons to two odors (O1: novel odor; O2: learned odor) in a *Scaled I* and the corresponding *Tuned* networks (same neurons and odors in the corresponding *rand* network are shown in **Fig. 1C**). **G.** Network with increased connectivity between the E assembly neurons ($\alpha = 5$) and the 25 I neurons that are most densely connected to the mitral cells activated by a given odor. **H.** Mean firing rate of the network in **E** in response to learned odors as a function of connection probability. Activity could not be stabilized efficiently by selecting I neurons based on their afferent connectivity.



Supplementary Figure 2.

Structured networks: additional results.


A. Raster plots showing responses of assembly or non-assembly neurons to a learned odor. **B.** Ratio of the average firing rate of (pseudo-)assembly neurons to non-assembly neurons in response to learned odors. **C.** Coefficients of variation of the inter-spike interval (ISI) in assembly neurons. **D.** Correlation between activity patterns across E neurons evoked by the same novel odor in different trials as a function of time. Pink bar indicates odor presentation. Note that correlations in response to novel odors are similar across networks and different from responses to learned odors in *Scaled* and *Tuned* networks (**Fig. 2E**).



Supplementary Figure 3.

Pattern completion: additional results.


A. Artificial reactivation of E assemblies. During 6 Hz baseline activity of the olfactory bulb, a subset of the assembly neurons was artificially reactivated by current injection (500 ms, 28 pA). Mean firing rates were quantified in the injected assembly neurons (i), in the remaining, non-injected assembly or pseudo-assembly neurons (ii), and in the non-assembly neurons (iii) as a function of time. The orange bar indicates duration of current injection. **B.** Correlation between activity patterns across E neurons (output patterns) evoked by a series of input patterns representing a morph of one learned odor into another learned odor. Correlations between output patterns are plotted as a function of the correlation between the corresponding OB patterns. **C.** Absolute value of the inhibitory and excitatory recurrent currents averaged over non-assembly neurons (see Fig. 3F for the average over assembly neurons).



Supplementary Figure 4.

Transformations and dimensionality of activity patterns: additional results.


A. Spike raster of 80 E neurons in response to a trajectory within the odor subspace (Fig. 4A) including one pure learned and one pure novel odor. **B.** Projection of activity patterns representing the odor subspace (Fig. 4A) onto the first 2 PCs of the corresponding *rand* networks (representative examples of one network each). **C.** Scree plot for the PCA results shown in Fig. 4D. **D.** Error in the reconstruction of odor-evoked activity patterns as a function of the number of PCs. Euclidean distances between all pairs of activity patterns were calculated in the full-dimensional state space (Df) and in reduced-dimensional embedding spaces (DI). The reconstruction error was defined as $1 - (\text{correlation between Df and DI})$. **E.** Participation ratio of the neural activity sampled from different numbers of neurons (50 iterations). **F.** Loading plot: contribution of neurons to the first two PCs of a *rand* and a *Tuned E+I* network (Fig. 4D). Each line represents one neuron and neurons that are part of the assemblies representing the two learned odors are color-coded in magenta. **G.** For each network and PC, the 100 E neurons with the highest absolute loadings were selected and grouped into 3 categories: neurons part of the two assemblies representing the learned odors, neurons part of the two “pseudo-assemblies” representing the two novel odors, and the remaining neurons (non-assembly).



Supplementary Figure 5.

Geometry and dynamics in a simple model with equal parameters for excitation and inhibition.


A. Firing rates averaged over all E neurons as a function of α (mean $\pm$ SD across 5 networks). **B.** Mean firing rate in response to learned odors or novel odors. Each data point represents one network-odor pair ($n = 5$ networks, 10 odors). **C.** Spike raster plots showing responses of 50 E neurons to two odors (O1: novel odor; O2: learned odor). **D.** Left: synaptic conductance contributed by afferents from the OB during odor stimulation (g_{OE}). Middle: total odor-evoked synaptic conductance g_{syn} . Right: percentage of E input contributed by recurrent connections during odor stimulation. Each dot represents one network (average over 10 odors); $n = 5$ networks. **E.** Co-tuning: correlation between time-averaged E and I conductances in response to different odors, average across neurons. **F.** Amplification within assemblies, calculated as the ratio between mean firing rates in response to learned odors averaged across assembly neurons in a structured network (*Scaled I* or *Tuned E+I*) and the same neurons in the corresponding *rand* network. **G.** Projection of activity patterns in pDpsim in response to the odor subspace described in Fig. 4 onto the first 2 PCs. Representative examples of different networks. **H.** Participation ratio (similar to Fig. 4F): activity evoked by novel odors and related mixtures (left), activity evoked by learned odors and related mixtures (center), and activity evoked by all stimuli (right).



Supplementary Figure 6.

Manipulating amplification within assemblies.

A. Overview of modified pDpsim networks that were generated to manipulate amplification of activity within assemblies. (1) *Tuned[adjust]*: in existing *Tuned E+I* networks, the probability of I-to-E connections (β) is further increased until the mean activity within assemblies is similar to the activity of the corresponding pseudo-assemblies in *rand* networks. (2) *Scaled[adjust]*: Parameters α and χ are set such that the mean activity within assemblies matches the mean activity in assemblies of *Tuned E+I* networks. (3) *Excitability*: in *rand* networks, during presentation of odor L or M, the resting conductance g_{rest} of the corresponding pseudo-assembly neurons is decreased until their mean activity matches the mean activity of the *Tuned E+I* assembly. During mixture presentation, g_{rest} decreases linearly with the concentration of the odor. **B.** Mean firing rate in response to learned or novel odors (see Fig. 2F). Each data point represents one network-odor pair ($n = 8$ networks, 10 odors). In *Scaled[adjust]*, mean firing rates for novel odors were lower than for learned odors. **C.** Amplification within assemblies, calculated as the ratio between mean firing rates in response to learned odors averaged across assembly neurons (structured network) and the corresponding pseudo-assembly neurons (*rand* network); see Fig. 2G. **D.** Quantification of E/I co-tuning. Left: correlation between time-averaged E and I conductances in response to different odors, average across neurons. Right: ratio of dispersion of joint conductances along balanced and counter-balanced axes (see Fig. 2H, I). Note that E/I co-tuning remained weak in *Scaled[adjust]* networks, inconsistent with experimental observations. **E.** Sum of recurrent inhibitory and excitatory currents averaged over assembly or non-assembly neurons (see Fig. 3E). In *Tuned E+I* and *Scaled[adjust]* networks, recurrent E inputs in assemblies slightly exceeded I inputs in the vicinity of learned odors. In *Tuned[adjust]*, recurrent E and I inputs cancelled out. **F.** Quantification of dimensionality of neural activity by the participation ratio (see Fig. 4F): activity evoked by the odor subspace. Each data point represents one network ($n = 8$). **G.** Projection of activity patterns in response to the odor subspace onto the first 2 PCs (representative examples of different networks; see Fig. 4D). **H.** d_M between one activity vector (v) and reference classes (Q). d_M was computed based on activity across subsets of 80 E neurons drawn from the four (pseudo-) assemblies with equal probability (see Fig. 5B).



Supplementary Figure 7.

Further analyses of pattern distances: uncorrelated odors.

Differences in dM between *rand* and *Tuned* networks may involve differences in the distance between class centers and/or differences in intra-class variability. To dissect the contributions of these effects we compared three distance measures:

1. The Euclidean distance between class centers:

$$d_E = \sqrt{(\mu_V - \mu_Q)^T (\mu_V - \mu_Q)}$$

2. The mean Euclidean distance between patterns of one class and the center of another class ($d_{\hat{E}}$):


$$d_{\hat{E}} = \frac{1}{n} \sum_{i=1,n} \sqrt{(v_i - \mu_Q)^T (v_i - \mu_Q)}$$

3. The Mahalanobis distance dM:

$$d_M = \frac{1}{n} \sum_{i=1,n} \sqrt{(v_i - \mu_Q)^T S^{-1} (v_i - \mu_Q)}$$

Here, μ_V and μ_Q are the centers (average patterns) of classes V and Q, respectively; v is a vector from $V = [v_1, v_2, \dots, v_n]$, and S^{-1} is the inverse of the covariance matrix of the neuronal population within reference class Q. Note that $d_{\hat{E}}$ corresponds to d_M without normalization by variability (covariance).


A. Analysis of dE. Left: dE in *rand* and *Tuned E+I* networks based on 80 E neurons drawn from (pseudo-) assemblies. Right: same based on 80 E neurons drawn from the entire population. Note that dE between learned and other odors was increased in *Tuned E+I* networks as compared to *rand* networks, particularly when neurons were drawn from assemblies. **B.** Equivalent plots for $d_{\hat{E}}$. Note that distances were increased nearly symmetrically, similar to d_E . **C.** Equivalent plots for dM (same plots as in [Fig. 5B](#)). Note that dM was increased asymmetrically. These observations show that the changes in dM relative to *rand* networks involved an increase in the distance between class centers (dE) and a non-isotropic change in intra-class variability (comparison between $d_{\hat{E}}$ and d_M). These effects were prominent when E neurons were drawn from assemblies. An important contribution to the increase in dM in the direction from learned odors to reference classes representing novel odors was made by the increased distance between class centers. In the other direction, dM was smaller, implying that variability in the reference class was higher. Nonetheless, variability in the relevant direction did not fully counteract the increased distance between class centers in *Tuned E+I* networks. As a consequence, dM was still increased slightly relative to the corresponding *rand* networks. Most of these effects were still observed, albeit weakly, when E neurons were drawn from the whole population.



Supplementary Figure 8.

Classification of odor representations (template matching).

A. Pattern classification probability quantified by template matching (see Methods). P_{Target} quantifies the probability that an activity pattern from the test set is assigned to a target class from the training set (see Fig. 5A) and is plotted as a function of the similarity (Pearson correlation) between the test and target odor patterns in the OB. Unlike QDA, template matching can also be applied to outputs of *Scaled* networks. Note enhanced classification probability for patterns evoked by odors similar to learned odors in *Tuned E+I* networks, and to a lesser extent in *Scaled* networks. **B.** The dataset was divided into 2 categories, depending on the trajectory in the odor subspace (insets; $n = 484$ activity patterns each), and template matching was performed again on each of the generated datasets. Note that learned odors were presented further apart in time in the top panels than in the bottom panels. Left: P_{Target} as a function of similarity in the OB, see panel A. Right: Classification probability for patterns similar to the training set (see A). Note that odor classification by *Scaled* networks depends on the trajectories in the odor subspace. **C.** Projection of activity patterns in response to the odor subspace onto the first 2 PCs in *Scaled* networks. Odor representations depend on stimulus history: note the reduced separation between representations of different odor classes during trajectories morphing one learned odor into another learned odor (arrows). This effect arises from the attractor dynamics of the networks.



Supplementary Figure 9.

Further analyses of pattern distances: correlated odors.

A-C. Pattern distances in *Tuned E+I* networks with (right) and without (left) two additional assemblies representing Y and Z (see Fig. 6). Similar analysis as in Supplementary Fig. 7, but for the odor subspace described in Fig. 6. **A.** d_E . **B.** d_E . **C.** d_M . **D.** Pattern classification probability quantified by QDA. P_{Target} quantifies the probability that an activity pattern from the test set is assigned to a target class from the training set (pure or closely related odor; see also Fig. 5 A & C).

Acknowledgements

We thank the Friedrich lab for insightful discussions. This work was supported by the Novartis Research Foundation, by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement no. 742576), and by the Swiss National Science Foundation (grants no. 31003A_172925/1, PCEFP3_202981).

References

- Ahmadian Y, Miller KD (2021) **What is the dynamical regime of cerebral cortex?** *Neuron* **109**:3373–3391 <https://doi.org/10.1016/j.neuron.2021.07.031>
- Altan E, Solla SA, Miller LE, Perreault EJ (2021) **Estimating the dimensionality of the manifold underlying multi-electrode neural recordings** *Plos Comput Biol* **17** <https://doi.org/10.1371/journal.pcbi.1008591>
- Amit DJ, Tsodyks MV (1991) **Quantitative study of attractor neural network retrieving at low spike rates: I. substrate—spikes, rates and neuronal gain** *Netw: Comput Neural Syst* **2**:259–273 https://doi.org/10.1088/0954-898x_2_3_003
- Babadi B, Sompolinsky H (2014) **Sparseness and Expansion in Sensory Representations** *Neuron* **83**:1213–1226 <https://doi.org/10.1016/j.neuron.2014.07.035>
- Baker C, Zhu V, Rosenbaum R (2020) **Nonlinear stimulus representations in neural circuits with approximate excitatory-inhibitory balance** *Plos Comput Biol* **16** <https://doi.org/10.1371/journal.pcbi.1008192>
- Barnes DC, Hofacer RD, Zaman AR, Rennaker RL, Wilson DA (2008) **Olfactory perceptual stability and discrimination** *Nat Neurosci* **11**:1378–1380 <https://doi.org/10.1038/nn.2217>
- Barron HC, Vogels TP, Behrens TE, Ramaswami M (2017) **Inhibitory engrams in perception and memory** *Proc National Acad Sci* **114**:6666–6674 <https://doi.org/10.1073/pnas.1701812114>
- Bhatia A, Moza S, Bhalla US (2019) **Precise excitation-inhibition balance controls gain and timing in the hippocampus** *eLife* **8** <https://doi.org/10.7554/elife.43415>
- Blazing RM, Franks KM (2020) **Odor coding in piriform cortex: mechanistic insights into distributed coding** *Elsevier Ltd* :1–7 <https://doi.org/10.1016/j.conb.2020.03.001>
- Blumhagen F, Zhu P, Shum J, Schärer Y-PZ, Yaksi E, Deisseroth K, Friedrich RW (2011) **Neuronal filtering of multiplexed odour representations** *Nature* **479**:493–498 <https://doi.org/10.1038/nature10633>
- Boerlin M, Machens CK, Denève S (2013) **Predictive Coding of Dynamical Variables in Balanced Spiking Networks** *PLoS Comput Biol* **9** <https://doi.org/10.1371/journal.pcbi.1003258>
- Brette R, Gerstner W (2005) **Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity** *J Neurophysiol* :1–6 <https://doi.org/10.1152/jn.00686.2005>
- Brunel N (2000) **Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons** *J Comput Neurosci* **8**:183–208 <https://doi.org/10.1023/a:1008925309027>
- Chapuis J, Wilson DA (2011) **Bidirectional plasticity of cortical pattern recognition and behavioral sensory acuity** *Nat Neurosci* **15**:155–161 <https://doi.org/10.1038/nn.2966>
- Chenkov HSRKN (2017) **Memory replay in balanced recurrent networks** *PLOS Computational Biology* **13** <https://doi.org/10.1371/journal.pcbi.1005359>

- Chung S, Abbott LF (2021) **Neural population geometry: An approach for understanding biological and artificial neural networks** *Curr Opin Neurobiol* **70**:137–144 <https://doi.org/10.1016/j.conb.2021.10.010>
- Dahmen D, Recanatesi S, Jia X, Ocker GK, Campagnola L, Seeman S, Jarsky T, Helias M, Shea-Brown E (2023) **Strong and localized recurrence controls dimensionality of neural activity across brain areas** *bioRxiv* <https://doi.org/10.1101/2020.11.02.365072>
- Denève S, Alemi A, Bourdoukan R (2017) **The Brain as an Efficient and Robust Adaptive Learner** *Neuron* **94**:969–977 <https://doi.org/10.1016/j.neuron.2017.05.016>
- Denève S, Machens CK (2016) **Efficient codes and balanced networks** *Nat Neurosci* **19**:375–382 <https://doi.org/10.1038/nn.4243>
- Denk W, Briggman KL, Helmstaedter M (2012) **Structural neurobiology: missing link to a mechanistic understanding of neural computation** *Nat Rev Neurosci* **13**:351–358 <https://doi.org/10.1038/nrn3169>
- Eckmann S, Young EJ, Gjorgjieva J (2024) **Synapse-type-specific competitive Hebbian learning forms functional recurrent networks** *PNAS*
- Festa D, Hennequin G, Lengyel M (2014) **Analog Memories in a Balanced Rate-Based Network of E-I Neurons** *Advances in Neural Information Processing Systems* **27**:1–9
- Frank T, Mönigl NR, Satoul C, Friedrich SH, Rainer W, Friedrich RW (2019) **Associative conditioning remaps odor representations and modifies inhibition in a higher olfactory brain area** *Nat Neurosci* :1–12 <https://doi.org/10.1038/s41593-019-0495-z>
- Franks KM, Russo MJ, Sosulski DL, Mulligan AA, Siegelbaum SA, Axel R (2011) **Recurrent Circuitry Dynamically Shapes the Activation of Piriform Cortex** *Neuron* **72**:49–56 <https://doi.org/10.1016/j.neuron.2011.08.020>
- Freeman WJ, Skarda CA (1985) **Spatial EEG patterns, non-linear dynamics and perception: the neo-sherringtonian view** *Brain Res Rev* **10**:147–175 [https://doi.org/10.1016/0165-0173\(85\)90022-0](https://doi.org/10.1016/0165-0173(85)90022-0)
- Friedrich RW, Laurent G (2004) **Dynamics of Olfactory Bulb Input and Output Activity During Odor Stimulation in Zebrafish** *Journal of Neurophysiology* **91**:2658–2669 <https://doi.org/10.1152/jn.01143.2003>
- Friedrich RW, Laurent G (2001) **Dynamic Optimization of Odor Representations by Slow Temporal Patterning of Mitral Cell Activity** *Science* **291**:1–7
- Friedrich RW, Wanner A (2021) **Dense Circuit Reconstruction to Understand Neuronal Computation: Focus on Zebrafish** *Annual Review of Neuroscience* **44**:1–19 <https://doi.org/10.1146/annurev-neuro-110220-013050>
- Frome RC, Merzenich MM, Schreiner CE (2007) **A synaptic memory trace for cortical receptive field plasticity** *Nature* **450**:425–429 <https://doi.org/10.1038/nature06289>
- Gallego JA, Perich MG, Miller LE, Solla SA (2017) **Neural Manifolds for the Control of Movement** *Neuron* **94**:978–984 <https://doi.org/10.1016/j.neuron.2017.05.025>

- Ghojogh B, Crowley M. (2019) **Linear and Quadratic Discriminant Analysis: Tutorial** *arXiv* <https://doi.org/10.48550/arxiv.1906.02590>
- Goldman-Rakic PS (1995) **Cellular basis of working memory** *Neuron* **14**:477–485 [https://doi.org/10.1016/0896-6273\(95\)90304-6](https://doi.org/10.1016/0896-6273(95)90304-6)
- Haberly LB (2001) **Parallel-distributed Processing in Olfactory Cortex: New Insights from Morphological and Physiological Analysis of Neuronal Circuitry** *Chem Senses* **26**:551–576 <https://doi.org/10.1093/chemse/26.5.551>
- Hebb DO (1949) **The organization of behavior; a neuropsychological theory** Wiley
- Hennequin G, Agnes EJ, Vogels TP (2017) **Inhibitory Plasticity: Balance, Control, and Codependence** *Annu Rev Neurosci* **40**:557–579 <https://doi.org/10.1146/annurev-neuro-072116-031005>
- Hennequin G, Ahmadian Y, Rubin DB, Lengyel M, Miller KD (2018) **The Dynamical Regime of Sensory Cortex: Stable Dynamics around a Single Stimulus-Tuned Attractor Account for Patterns of Noise Variability** *Neuron* **98**:846–860 <https://doi.org/10.1016/j.neuron.2018.04.017>
- Hennequin G, Vogels TP, Gerstner W (2014) **Optimal Control of Transient Dynamics in Balanced Networks Supports Generation of Complex Movements** *Neuron* **82**:1394–1406 <https://doi.org/10.1016/j.neuron.2014.04.045>
- Hopfield JJ (1982) **Neural networks and physical systems with emergent collective computational abilities** *Proc Natl Acad Sci* :2554–2558
- Iurilli G, Datta SR (2017) **Population Coding in an Innately Relevant Olfactory Area** *Neuron* **93**:1180–1197 <https://doi.org/10.1016/j.neuron.2017.02.010>
- Jacobson G, Rupprecht P, Friedrich RW (2018) **Experience-Dependent Plasticity of Odor Representations in the Telencephalon of Zebrafish** *Curr Biol* **28**:1–14 <https://doi.org/10.1016/j.cub.2017.11.007>
- Khona M, Fiete IR (2022) **Attractor and integrator networks in the brain** *Nat Rev Neurosci* **23**:744–766 <https://doi.org/10.1038/s41583-022-00642-0>
- Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, Mrcic-Flogel TD (2011) **Functional specificity of local synaptic connections in neocortical networks** *Nature* **473**:87–91 <https://doi.org/10.1038/nature09880>
- Kohonen T (1984) **Self-Organization and Associative Memory** Springer Berlin, Heidelberg
- Kornfeld J, Denk W (2018) **Progress and remaining challenges in high-throughput volume electron microscopy** *Curr Opin Neurobiol* **50**:261–267 <https://doi.org/10.1016/j.conb.2018.04.030>
- Lagzi F, Bustos MC, Oswald A-M, Doiron B. (2021) **Assembly formation is stabilized by Parvalbumin neurons and accelerated by Somatostatin neurons** *Biorxiv* <https://doi.org/10.1101/2021.09.06.459211>

- Lagzi F, Fairhall AL (2024) **Emergence of co-tuning in inhibitory neurons as a network phenomenon mediated by randomness, correlations, and homeostatic plasticity** *Sci Adv* **10** <https://doi.org/10.1126/sciadv.adi4350>
- Lagzi F, Rotter S (2015) **Dynamics of Competition between Subnetworks of Spiking Neuronal Networks in the Balanced State** *PLoS ONE* **10** <https://doi.org/10.1371/journal.pone.0138947>
- Langdon C, Genkin M, Engel TA (2023) **A unifying perspective on neural manifolds and circuits for cognition** *Nat Rev Neurosci* **24**:363–377 <https://doi.org/10.1038/s41583-023-00693-x>
- Litwin-Kumar A, Doiron B (2014) **Formation and maintenance of neuronal assemblies through synaptic plasticity** *Nature Communications* **5**:1–12 <https://doi.org/10.1038/ncomms6319>
- Litwin-Kumar A, Doiron B (2012) **Slow dynamics and high variability in balanced cortical networks with clustered connections** *Nat Neurosci* **15**:1498–1505 <https://doi.org/10.1038/nn.3220>
- Mackwood O, Naumann LB, Sprekeler H (2021) **Learning excitatory-inhibitory neuronal assemblies in recurrent networks** *Elife* **10** <https://doi.org/10.7554/elife.59715>
- Marr D (1969) **A theory of cerebellar cortex** *J Physiol* **202**:437–470
- Mazzucato L, Fontanini A, Camera GL (2015) **Dynamics of Multistable States during Ongoing and Evoked Cortical Activity** *J Neurosci* **35**:8214–8231 <https://doi.org/10.1523/jneurosci.4819-14.2015>
- Meissner-Bernard C, Dembitskaya Y, Venance L, Fleischmann A (2018) **Encoding of Odor Fear Memories in the Mouse Olfactory Cortex** *Current Biology* **29**:1–19 <https://doi.org/10.1016/j.cub.2018.12.003>
- Miehl C, Onasch S, Festa D, Gjorgjieva J (2022) **Formation and computational implications of assemblies in neural circuits** *J Physiology* <https://doi.org/10.1113/jp282750>
- Miura K, Mainen ZF, Uchida N (2012) **Odor Representations in Olfactory Cortex: Distributed Rate Coding and Decorrelated Population Activity** *Neuron* **74**:1087–1098 <https://doi.org/10.1016/j.neuron.2012.04.021>
- Mueller T, Dong Z, Berberoglu MA, Guo S (2011) **The dorsal pallium in zebrafish, *Danio rerio* (Cyprinidae Teleostei).** *Brain Res* **1381**:95–105 <https://doi.org/10.1016/j.brainres.2010.12.089>
- Murphy BK, Miller KD (2009) **Balanced Amplification: A New Mechanism of Selective Amplification of Neural Activity Patterns** *Neuron* **61**:635–648 <https://doi.org/10.1016/j.neuron.2009.02.005>
- Okun M, Lampl I (2008) **Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities** *Nat Neurosci* **11**:535–537 <https://doi.org/10.1038/nn.2105>
- Pashkovski SL, Iurilli G, Brann D, Chicharro D, Drummey K, Franks K, Panzeri S, Datta SR (2020) **Structure and flexibility in cortical representations of odour space** *Nature* :1–28 <https://doi.org/10.1038/s41586-020-2451-1>

- Pehlevan C, Sompolinsky H (2014) **Selectivity and Sparseness in Randomly Connected Balanced Networks** *PLOS ONE* 9
- Recanatesi S, Ocker GK, Buice MA, Shea-Brown E (2019) **Dimensionality in recurrent spiking networks: Global trends in activity and local origins in connectivity** *PLoS Comput Biol* 15 <https://doi.org/10.1371/journal.pcbi.1006446>
- Renart A, Rocha J de la, Bartho P, Hollender L, Parga N, Reyes A, Harris KD (2010) **The Asynchronous State in Cortical Circuits** *Science* 327:587–590 <https://doi.org/10.1126/science.1179850>
- Rost T, Deger M, Nawrot MP (2018) **Winnerless competition in clustered balanced networks: inhibitory assemblies do the trick** *Biol Cybern* 112:81–98 <https://doi.org/10.1007/s00422-017-0737-7>
- Roudi Y, Latham PE (2007) **A Balanced Memory Network** *PLoS Comput Biol* 3:e141–22 <https://doi.org/10.1371/journal.pcbi.0030141>
- Rupprecht P *et al.* (2021) **A database and deep learning toolbox for noise-optimized, generalized spike inference from calcium imaging** *Nat Neurosci* 24:1324–1337 <https://doi.org/10.1038/s41593-021-00895-5>
- Rupprecht P, Friedrich RW (2018) **Precise Synaptic Balance in the Zebrafish Homolog of Olfactory Cortex** *Neuron* 100:1–29 <https://doi.org/10.1016/j.neuron.2018.09.013>
- Ryan TJ, Roy DS, Pignatelli M, Arons A, Tonegawa S (2015) **Engram Cells Retain Memory Under Retrograde Amnesia** *Science* :1007–1013
- Sacco T, Sacchetti B (2010) **Role of Secondary Sensory Cortices in Emotional Memory Storage and Retrieval in Rats** *Science* 329:649–656 <https://doi.org/10.1126/science.1183165>
- Sadeh S, Clopath C (2021) **Excitatory-inhibitory balance modulates the formation and dynamics of neuronal assemblies in cortical networks** *Sci Adv* 7 <https://doi.org/10.1126/sciadv.abg8411>
- Sadeh S, Clopath C (2020) **Patterned perturbation of inhibition can reveal the dynamical structure of neural processing** *eLife* 9 <https://doi.org/10.7554/elife.52757>
- Sadeh S, Clopath C (2020) **Inhibitory stabilization and cortical computation** *Nat Rev Neurosci* :1–17 <https://doi.org/10.1038/s41583-020-00390-z>
- Schaffer ES, Stettler DD, Kato D, Choi GB, Axel R, Abbott LF (2018) **Odor Perception on the Two Sides of the Brain: Consistency Despite Randomness** *Neuron* 98:736–742 <https://doi.org/10.1016/j.neuron.2018.04.004>
- Schoonover CE, Ohashi SN, Axel R, Fink AJP (2021) **Representational drift in primary olfactory cortex** *Nature* :1–34 <https://doi.org/10.1038/s41586-021-03628-7>
- Schulz A, Miehl C II MJB, Gjorgjieva J (2021) **The generation of cortical novelty responses through inhibitory plasticity** *eLife* 10:1–28 <https://doi.org/10.7554/eLife.65309>
- Shadlen MN, Newsome WT (1994) **Noise, neural codes and cortical organization** *Curr Opin Neurobiol* 4:569–579 [https://doi.org/10.1016/0959-4388\(94\)90059-0](https://doi.org/10.1016/0959-4388(94)90059-0)

- Stettler DD, Axel R (2009) **Representations of Odor in the Piriform Cortex** *Neuron* **63**:854–864 <https://doi.org/10.1016/j.neuron.2009.09.005>
- Tabor R, Friedrich RW (2008) **Pharmacological Analysis of Ionotropic Glutamate Receptor Function in Neuronal Circuits of the Zebrafish Olfactory Bulb** *PLoS ONE* **3** <https://doi.org/10.1371/journal.pone.0001416>
- Tabor R, Yaksi E, Weislogel J-M, Friedrich RW (2004) **Processing of Odor Mixtures in the Zebrafish Olfactory Bulb** *The Journal of Neuroscience* **24**:6611–6620 <https://doi.org/10.1523/jneurosci.1834-04.2004>
- Tsodyks MV, Skaggs WE, Sejnowski TJ, McNaughton BL (1997) **Paradoxical Effects of External Modulation of Inhibitory Interneurons** *J Neurosci* **17**:4382–4388 <https://doi.org/10.1523/jneurosci.17-11-04382.1997>
- Van Vreeswijk C, Sompolinsky H (1996) **Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity** *Science* **274**:1724–1726 <https://doi.org/10.1126/science.274.5293.1724>
- Vogels TP, Sprekeler H, Zenke F, Clopath C, Gerstner W (2011) **Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks** *Science* **334**:1–7 <https://doi.org/10.1126/science.1212991>
- Wanner AA, Friedrich RW (2020) **Whitening of odor representations by the wiring diagram of the olfactory bulb** *Nat Neurosci* :1–24 <https://doi.org/10.1038/s41593-019-0576-z>
- Wehr M, Zador AM (2003) **Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex** *Nature* **426**:1–5
- Wiechert MT, Judkewitz B, Riecke H, Friedrich RW (2010) **Mechanisms of pattern decorrelation by recurrent neuronal circuits** *Nat Neurosci* **13**:1003–1010 <https://doi.org/10.1038/nn.2591>
- Wilson DA, Sullivan RM (2011) **Cortical Processing of Odor Objects** *Neuron* **72**:506–519 <https://doi.org/10.1016/j.neuron.2011.10.027>
- Wu YK, Zenke F (2021) **Nonlinear transient amplification in recurrent neural networks with short-term plasticity** *eLife* **10** <https://doi.org/10.7554/eLife.71263>
- Yaksi E, Paul F, von S, Niessing J, Bundschuh ST, Friedrich RW (2009) **Transformation of odor representations in target areas of the olfactory bulb** *Nat Neurosci* **12**:474–482 <https://doi.org/10.1038/nn.2288>
- Yassa MA, Stark CEL (2011) **Pattern separation in the hippocampus** *Trends Neurosci* **34**:515–525 <https://doi.org/10.1016/j.tins.2011.06.006>
- Zenke F, Agnes EJ, Gerstner W (2015) **Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks** *Nature Communications* **6**:1–13 <https://doi.org/10.1038/ncomms7922>

Znamenskiy P, Kim M-H, Muir DR, Iacaruso MF, Hofer SB, Mrsic-Flogel TD (2024) **Functional specificity of recurrent inhibition in visual cortex** *Neuron* <https://doi.org/10.1016/j.neuron.2023.12.013>

Author information

Claire Meissner-Bernard

Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland

ORCID iD: [0009-0007-2038-8398](https://orcid.org/0009-0007-2038-8398)

Friedemann Zenke

Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland, University of Basel, Basel, Switzerland

ORCID iD: [0000-0003-1883-644X](https://orcid.org/0000-0003-1883-644X)

Rainer W Friedrich

Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland, University of Basel, Basel, Switzerland

ORCID iD: [0000-0001-9107-0482](https://orcid.org/0000-0001-9107-0482)

For correspondence: rainer.friedrich@fmi.ch

Editors

Reviewing Editor

Julijana Gjorgjieva

Technical University of Munich, Freising, Germany

Senior Editor

Panayiota Poirazi

FORTH Institute of Molecular Biology and Biotechnology, Heraklion, Greece

Reviewer #1 (Public review):

Summary:

Meissner-Bernard et al present a biologically constrained model of telencephalic area of adult zebrafish, a homologous area to the piriform cortex, and argue for the role of precisely balanced memory networks in olfactory processing.

This is interesting as it can add to recent evidence on the presence of functional subnetworks in multiple sensory cortices. It is also important in deviating from traditional accounts of memory systems as attractor networks. Evidence for attractor networks has been found in some systems, like in the head direction circuits in the flies. However, the presence of attractor dynamics in other modalities, like sensory systems, and their role in computation has been more contentious. This work contributes to this active line of research in experimental and computational neuroscience by suggesting that, rather than being represented in attractor networks and persistent activity, olfactory memories might be coded by balanced excitation-inhibitory subnetworks.

Strengths:

The main strength of the work is in: (1) direct link to biological parameters and measurements, (2) good controls and quantification of the results, and (3) comparison across multiple models.

(1) The authors have done a good job of gathering the current experimental information to inform a biological-constrained spiking model of the telencephalic area of adult zebrafish. The results are compared to previous experimental measurements to choose the right regimes of operation.

(2) Multiple quantification metrics and controls are used to support the main conclusions, and to ensure that the key parameters are controlled for - e.g. when comparing across multiple models.

(3) Four specific models (random, scaled I / attractor, and two variant of specific E-I networks - tuned I and tuned E+I) are compared with different metrics, helping to pinpoint which features emerge in which model.

In the revised manuscript, the authors have also:

(a) made a good effort to provide a mechanistic explanation of their results (especially on the mechanism underlying medium amplification in specific E/I network models);

(b) performed a systematic analysis of the parameter space by changing different parameters of E and I neurons (specifically showing that different time constants of E and I neurons do not change the results and therefore the main effects result from connectivity);

(c) added further analysis and discussion on the potential functional and computational significance of balanced specific E-I subnetworks.

These additions substantially strengthen the study, presenting compelling evidence for how networks with specific E-I structure can underpin olfactory processing and memory representations. The findings have potential implications that extend beyond the olfactory system and may be applicable to other neural systems and species.

<https://doi.org/10.7554/eLife.96303.2.sa3>

Reviewer #2 (Public review):

Summary:

The authors conducted a comparative analysis of four networks, varying in the presence of excitatory assemblies and the architecture of inhibitory cell assembly connectivity. They found that co-tuned E-I assemblies provide network stability and a continuous representation of input patterns (on locally constrained manifolds), contrasting with networks with global inhibition that result in attractor networks.

Strengths:

The findings presented in this paper are very interesting and cutting-edge. The manuscript effectively conveys the message and presents a creative way to represent high-dimensional inputs and network responses. Particularly, the result regarding the projection of input patterns onto local manifolds and continuous representation of input/memory is very Intriguing and novel. Both computational and experimental neuroscientists would find value in reading the paper.

Weaknesses:

Intuitively, classification (decodability) in discrete attractor networks is much better than in networks with continuous representations. This could also be shown in Figure 5B, along with the performance of the random and tuned E-I networks. The latter networks have the

advantage of providing network stability compared to the Scaled I network, but at the cost of reduced network salience and, therefore, reduced input decodability. Thus, tuned E-I networks cannot always perform better than any other network.

<https://doi.org/10.7554/eLife.96303.2.sa2>

Reviewer #3 (Public review):

Summary:

This work investigates computational consequences of assemblies containing both excitatory and inhibitory neurons (E/I assembly) in a model with parameters constrained by experimental data from the telencephalic area Dp of zebrafish. The authors show how this precise E/I balance shapes the geometry of neuronal dynamics in comparison to unstructured networks and networks with more global inhibitory balance. Specifically, E/I assemblies lead to the activity being locally restricted onto manifolds - a dynamical structure in-between high-dimensional representations in unstructured networks and discrete attractors in networks with global inhibitory balance. Furthermore, E/I assemblies lead to smoother representations of mixtures of stimuli while those stimuli can still be reliably classified, and allows for more robust learning of additional stimuli.

Strengths:

Since experimental studies do suggest that E/I balance is very precise and E/I assemblies exist, it is important to study the consequences of those connectivity structures on network dynamics. The authors convincingly show that E/I assemblies lead to different geometries of stimulus representation compared to unstructured networks and networks with global inhibition. This finding might open the door for future studies for exploring the functional advantage of these locally defined manifolds, and how other network properties allow to shape those manifolds.

The authors also make sure that their spiking model is well-constrained by experimental data from the zebrafish pDp. Both, spontaneous and odor stimulus triggered spiking activity is within the range of experimental measurements. But the model is also general enough to be potentially applied to findings in other animal models and brain regions.

Weaknesses:

All my previous points have been addressed.

<https://doi.org/10.7554/eLife.96303.2.sa1>

Author response:

The following is the authors' response to the original reviews.

The revised manuscript contains new results and additional text. Major revisions:

(1) Additional simulations and analyses of networks with different biophysical parameters and with identical time constants for E and I neurons (Methods, Supplementary Fig. 5).

(2) Additional simulations and analyses of networks with modifications of connectivity parameters to further analyze effects of E/I assemblies on manifold geometry (Supplementary Fig. 6).

- (3) Analysis of synaptic current components (Figure 3 D-F; to analyze mechanism of modest amplification in *Tuned* networks).
- (4) More detailed explanation of pattern completion analysis (Results).
- (5) Analysis of classification performance of *Scaled* networks (Supplementary Fig.8).
- (6) Additional analysis (Figure 5D-F) and discussion (particularly section “Computational functions of networks with E/I assemblies”) of functional benefits of continuous representations in networks with E-I assemblies.

Public Reviews:

Reviewer #1 (Public Review):

Summary:

Meissner-Bernard et al present a biologically constrained model of telencephalic area of adult zebrafish, a homologous area to the piriform cortex, and argue for the role of precisely balanced memory networks in olfactory processing.

This is interesting as it can add to recent evidence on the presence of functional subnetworks in multiple sensory cortices. It is also important in deviating from traditional accounts of memory systems as attractor networks. Evidence for attractor networks has been found in some systems, like in the head direction circuits in the flies. However, the presence of attractor dynamics in other modalities, like sensory systems, and their role in computation has been more contentious. This work contributes to this active line of research in experimental and computational neuroscience by suggesting that, rather than being represented in attractor networks and persistent activity, olfactory memories might be coded by balanced excitation-inhibitory subnetworks.

Strengths:

The main strength of the work is in: (1) direct link to biological parameters and measurements, (2) good controls and quantification of the results, and (3) comparison across multiple models.

(1) The authors have done a good job of gathering the current experimental information to inform a biological-constrained spiking model of the telencephalic area of adult zebrafish. The results are compared to previous experimental measurements to choose the right regimes of operation.

(2) Multiple quantification metrics and controls are used to support the main conclusions and to ensure that the key parameters are controlled for - e.g. when comparing across multiple models. (3) Four specific models (random, scaled I / attractor, and two variant of specific E-I networks - tuned I and tuned E+I) are compared with different metrics, helping to pinpoint which features emerge in which model.

Weaknesses:

Major problems with the work are: (1) mechanistic explanation of the results in specific E-I networks, (2) parameter exploration, and (3) the functional significance of the specific E-I model.

(1) The main problem with the paper is a lack of mechanistic analysis of the models. The models are treated like biological entities and only tested with different assays and metrics to describe their different features (e.g. different geometry of representation in Fig. 4). Given that all the key parameters of the models are known and can be changed

(unlike biological networks), it is expected to provide a more analytical account of why specific networks show the reported results. For instance, what is the key mechanism for medium amplification in specific E/I network models (Fig. 3)? How does the specific geometry of representation/manifolds (in Fig. 4) emerge in terms of excitatory-inhibitory interactions, and what are the main mechanisms/parameters? Mechanistic account and analysis of these results are missing in the current version of the paper.

We agree that further mechanistic insights would be of interest and addressed this issue at different levels:

(1) Biophysical parameters: to determine whether network behavior depends on specific choices of biophysical parameters in E and I neurons we equalized biophysical parameters across neuron types. The main observations are unchanged, suggesting that the observed effects depend primarily on network connectivity (see also response to comment [2]).

(2) Mechanism of modest amplification in E/I assemblies: analyzing the different components of the synaptic currents demonstrate that the modest amplification of activity in *Tuned* networks results from an “imperfect” balance of *recurrent* excitation and inhibition within assemblies (see new Figures 3D-F and text p.7). Hence, E/I co-tuning substantially reduces the net amplification in *Tuned* networks as compared to *Scaled* networks, thus preventing discrete attractor dynamics and stabilizing network activity, but a modest amplification still occurs, consistent with biological observations.

(3) Representational geometry: to obtain insights into the network mechanisms underlying effects of E/I assemblies on the geometry of population activity we tested the hypothesis that geometrical changes depend, at least in part, on the modest amplification of activity within E/I assemblies (see Supplementary Figure 6). We changed model parameters to either prevent the modest amplification in *Tuned* networks (increasing I-to-E connectivity within assemblies) or introduce a modest amplification in subsets of neurons by other mechanisms (concentration-dependent increase in the excitability of pseudo-assembly neurons; *Scaled* I networks with reduced connectivity within assemblies). Manipulations that introduced a modest, input-dependent amplification in neuronal subsets had geometrical effects similar to those observed in *Tuned* networks, whereas manipulations that prevented a modest amplification abolished these effects (Supplementary Figure 6). Note however that these manipulations generated different firing rate distributions. These results provide a starting point for more detailed analyses of the relationship between network connectivity and representational geometry (see p.12).

In summary, our additional analyses indicate that effects of E/I assemblies on representational geometry depend primarily on network connectivity, rather than specific biophysical parameters, and that the resulting modest amplification of activity within assemblies makes an important contribution. Further analyses may reveal more specific relationships between E/I assemblies and representational geometry, but such analyses are beyond the scope of this study.

(2) The second major issue with the study is a lack of systematic exploration and analysis of the parameter space. Some parameters are biologically constrained, but not all the parameters. For instance, it is not clear what the justification for the choice of synaptic time scales are (with E synaptic time constants being larger than inhibition: $\tau_{\text{syn}_i} = 10 \text{ ms}$, $\tau_{\text{syn}_E} = 30 \text{ ms}$). How would the results change if they are varying these - and other unconstrained - parameters? It is important to show how the main results, especially the manifold localisation, would change by doing a systematic exploration of the key parameters and performing some sensitivity analysis. This would also help to see how robust the results are, which parameters are more important and which parameters are less relevant, and to shed light on the key mechanisms.

We thank the reviewer for raising this point. We chose a relatively slow time constant for excitatory synapses because experimental data indicate that excitatory synaptic currents in Dp and piriform cortex contain a prominent NMDA component. Nevertheless, to assess whether network behavior depends on specific choices of biophysical parameters in E and I neurons, we have performed additional simulations with equal synaptic time constants and equal biophysical parameters for all neurons. Each neuron also received the same number of inputs from each population (see revised Methods). Results were similar to those observed previously (Supplementary Fig.5 and p.9 of main text). We therefore conclude that the main effects observed in *Tuned* networks cannot be explained by differences in biophysical parameters between E and I neurons but is primarily a consequence of network connectivity.

(3) It is not clear what the main functional advantage of the specific E-I network model is compared to random networks. In terms of activity, they show that specific E-I networks amplify the input more than random networks (Fig. 3). But when it comes to classification, the effect seems to be very small (Fig. 5c). Description of different geometry of representation and manifold localization in specific networks compared to random networks is good, but it is more of an illustration of different activity patterns than proving a functional benefit for the network. The reader is still left with the question of what major functional benefits (in terms of computational/biological processing) should be expected from these networks, if they are to be a good model for olfactory processing and learning.

One possibility for instance might be that the tasks used here are too easy to reveal the main benefits of the specific models - and more complex tasks would be needed to assess the functional enhancement (e.g. more noisy conditions or more combination of odours). It would be good to show this more clearly - or at least discuss it in relation to computation and function.

In the previous manuscript, the analysis of potential computational benefits other than pattern classification was limited and the discussion of this issue was condensed into a single itemized paragraph to avoid excessive speculation. Although a thorough analysis of potential computational benefits exceeds the scope of a single paper, we agree with the reviewer that this issue is of interest and therefore added additional analyses and discussion.

In the initial manuscript we analyzed pattern classification primarily to investigate whether *Tuned* networks can support this function at all, given that they do not exhibit discrete attractor states. We found this to be the case, which we consider a first important result.

Furthermore, we found that precise balance of E/I assemblies can protect networks against catastrophic firing rate instabilities when assemblies are added sequentially, as in continual learning. Results from these simulations are now described and discussed in more detail (see Results p.11 and Discussion p.13).

In the revised manuscript, we now also examine additional potential benefits of *Tuned* networks and discuss them in more detail (see new Figure 5D-F and text p.11). One hypothesis is that continuous representations provide a distance metric between a given input and relevant (learned) stimuli. To address this hypothesis, we (1) performed regression analysis and (2) trained support vector machines (SVMs) to predict the concentration of a given odor in a mixture based on population activity. In both cases, *Tuned E+I* networks outperformed *Scaled* and *_rand n_* networks in predicting the concentration of learned odors across a wide range mixtures (Figure 5D-F). E/I assemblies therefore support the quantification of learned odors within mixtures or, more generally, assessments of how strongly a (potentially complex) input is related to relevant odors stored in memory. Such a metric assessment of stimulus quality is not well supported by discrete attractor networks because inputs are mapped onto discrete network states.

The observation that *Tuned* networks do not map inputs onto discrete outputs indicates that such networks do not classify inputs as distinct items. Nonetheless, the observed geometrical modifications of continuous representations support the classification of learned inputs or the assessment of metric relationships by hypothetical readout neurons. Geometrical modifications of odor representations may therefore serve as one of multiple steps in multi-layer computations for pattern classification (and/or other computations). In this scenario, the transformation of odor representations in Dp may be seen as related to transformations of representations between different layers in artificial networks, which collectively perform a given task (notwithstanding obvious structural and mechanistic differences between artificial and biological networks). In other words, geometrical transformations of representations in *Tuned* networks may overrepresent learned (relevant) information at the expense of other information and thereby support further learning processes in other brain areas. An obvious corollary of this scenario is that Dp does not perform odor classification *per se* based on inputs from the olfactory bulb but reformats representations of odor space based on experience to support computational tasks as part of a larger system. This scenario is now explicitly discussed (p.14).

Reviewer #2 (Public Review):

Summary:

The authors conducted a comparative analysis of four networks, varying in the presence of excitatory assemblies and the architecture of inhibitory cell assembly connectivity. They found that co-tuned E-I assemblies provide network stability and a continuous representation of input patterns (on locally constrained manifolds), contrasting with networks with global inhibition that result in attractor networks.

Strengths:

The findings presented in this paper are very interesting and cutting-edge. The manuscript effectively conveys the message and presents a creative way to represent high-dimensional inputs and network responses. Particularly, the result regarding the projection of input patterns onto local manifolds and continuous representation of input/memory is very Intriguing and novel. Both computational and experimental neuroscientists would find value in reading the paper.

Weaknesses:

that have continuous representations. This could also be shown in Figure 5B, along with the performance of the random and tuned E-I networks. The latter networks have the advantage of providing network stability compared to the Scaled I network, but at the cost of reduced network salience and, therefore, reduced input decodability. The authors may consider designing a decoder to quantify and compare the classification performance of all four networks.

We have now quantified classification by networks with discrete attractor dynamics (*Scaled*) along with other networks. However, because the neuronal covariance matrix for such networks is low rank and not invertible, pattern classification cannot be analyzed by QDA as in Figure 5B. We therefore classified patterns from the odor subspace by template matching, assigning test patterns to one of the four classes based on correlations (see Supplementary Figure 8). As expected, *Scaled* networks performed well, but they did not outperform *Tuned* networks. Moreover, the performance of *Scaled* networks, but not *Tuned* networks, depended on the order in which odors were presented to the network. This hysteresis effect is a direct consequence of persistent attractor states and decreased the general classification performance of *Scaled* networks (see Supplementary Figure 8 for details). These results

confirm the prediction that networks with discrete attractor states can efficiently classify inputs, but also reveal disadvantages arising from attractor dynamics. Moreover, the results indicate that the classification performance of *Tuned* networks is also high under the given task conditions, which simulate a biologically realistic scenario.

We would also like to emphasize that classification may not be the only task, and perhaps not even a main task, of Dp/piriform cortex or other memory networks with E/I assemblies. Conceivably, other computations could include metric assessments of inputs relative to learned inputs or additional learning-related computations. Please see our response to comment (3) of reviewer 1 for a further discussion of this issue.

Networks featuring E/I assemblies could potentially represent multistable attractors by exploring the parameter space for their reciprocal connectivity and connectivity with the rest of the network. However, for co-tuned E-I networks, the scope for achieving multistability is relatively constrained compared to networks employing global or lateral inhibition between assemblies. It would be good if the authors mentioned this in the discussion. Also, the fact that reciprocal inhibition increases network stability has been shown before and should be cited in the statements addressing network stability (e.g., some of the citations in the manuscript, including Rost et al. 2018, Lagzi & Fairhall 2022, and Vogels et al. 2011 have shown this).

We thank the reviewer for this comment. We now explicitly discuss multistability (see p. 12) and refer to additional references in the statements addressing network stability.

Providing raster plots of the pDp network for familiar and novel inputs would help with understanding the claims regarding continuous versus discrete representation of inputs, allowing readers to visualize the activity patterns of the four different networks. (similar to Figure 1B).

We thank the reviewer for this suggestion. We have added raster plots of responses to both familiar and novel inputs in the revised manuscript (Figure 2D and Supplementary Figure 4A).

Reviewer #3 (Public Review):

Summary:

This work investigates the computational consequences of assemblies containing both excitatory and inhibitory neurons (E/I assembly) in a model with parameters constrained by experimental data from the telencephalic area Dp of zebrafish. The authors show how this precise E/I balance shapes the geometry of neuronal dynamics in comparison to unstructured networks and networks with more global inhibitory balance. Specifically, E/I assemblies lead to the activity being locally restricted onto manifolds - a dynamical structure in between high-dimensional representations in unstructured networks and discrete attractors in networks with global inhibitory balance. Furthermore, E/I assemblies lead to smoother representations of mixtures of stimuli while those stimuli can still be reliably classified, and allow for more robust learning of additional stimuli.

Strengths:

Since experimental studies do suggest that E/I balance is very precise and E/I assemblies exist, it is important to study the consequences of those connectivity structures on network dynamics. The authors convincingly show that E/I assemblies lead to different geometries of stimulus representation compared to unstructured networks and networks with global inhibition. This finding might open the door for future studies for exploring

the functional advantage of these locally defined manifolds, and how other network properties allow to shape those manifolds.

The authors also make sure that their spiking model is well-constrained by experimental data from the zebrafish pDp. Both spontaneous and odor stimulus triggered spiking activity is within the range of experimental measurements. But the model is also general enough to be potentially applied to findings in other animal models and brain regions.

Weaknesses:

I find the point about pattern completion a bit confusing. In Fig. 3 the authors argue that only the Scaled I network can lead to pattern completion for morphed inputs since the output correlations are higher than the input correlations. For me, this sounds less like the network can perform pattern completion but it can nonlinearly increase the output correlations. Furthermore, in Suppl. Fig. 3 the authors show that activating half the assembly does lead to pattern completion in the sense that also non-activated assembly cells become highly active and that this pattern completion can be seen for Scaled I, Tuned E+I, and Tuned I networks. These two results seem a bit contradictory to me and require further clarification, and the authors might want to clarify how exactly they define pattern completion.

We believe that this comment concerns a semantic misunderstanding and apologize for any lack of clarity. We added a definition of pattern completion in the text: “...the retrieval of the whole memory from noisy or corrupted versions of the learned input.”. Pattern completion may be assessed using different procedures. In computational studies, it is often analyzed by delivering input to a subset of the assembly neurons which store a given memory (partial activation). Under these conditions, we find recruitment of the entire assembly in all structured networks, as demonstrated in Supplementary Figure 3. However, these conditions are unlikely to occur during odor presentation because the majority of neurons do not receive any input.

Another more biologically motivated approach to assess pattern completion is to gradually modify a realistic odor input into a learned input, thereby gradually increasing the overlap between the two inputs. This approach had been used previously in experimental studies (references added to the text p.6). In the presence of assemblies, recurrent connectivity is expected to recruit assembly neurons (and thus retrieve the stored pattern) more efficiently as the learned pattern is approached. This should result in a nonlinear increase in the similarity between the evoked and the learned activity pattern. This signature was prominent in *Scaled* networks but not in *Tuned* or *rand* networks. Obviously, the underlying procedure is different from the partial activation of the assembly described above because input patterns target many neurons (including neurons outside assemblies) and exhibit a biologically realistic distribution of activity. However, this approach has also been referred to as “pattern completion” in the neuroscience literature, which may be the source of semantic confusion here. To clarify the difference between these approaches we have now revised the text and explicitly described each procedure in more detail (see p.6).

The authors argue that Tuned E+I networks have several advantages over Scaled I networks. While I agree with the authors that in some cases adding this localized E/I balance is beneficial, I believe that a more rigorous comparison between Tuned E+I networks and Scaled I networks is needed: quantification of variance (Fig. 4G) and angle distributions (Fig. 4H) should also be shown for the Scaled I network. Similarly in Fig. 5, what is the Mahalanobis distance for Scaled I networks and how well can the Scaled I network be classified compared to the Tuned E+I network? I suspect that the Scaled I network will actually be better at classifying odors compared to the E+I network. The authors might want to speculate about the benefit of having networks with both sources

of inhibition (local and global) and hence being able to switch between locally defined manifolds and discrete attractor states.

We agree that a more rigorous comparison of Tuned and Scaled networks would be of interest. We have added the variance analysis (Fig 4G) and angle distributions (Fig. 4H) for both Tuned I and Scaled networks. However, the Mahalanobis distances and Quadratic Discriminant Analysis cannot be applied to Scaled networks because their neuronal covariance matrix is low rank and not invertible. To nevertheless compare these networks, we performed template matching by assigning test patterns to one of the four odor classes based on correlations to template patterns (Supplementary Figure 8; see also response to the first comment of reviewer 2). Interestingly, *Scaled* networks performed well at classification but did not outperform *Tuned* networks, and exhibited disadvantages arising from attractor dynamics (Supplementary Figure 8; see also response to the first comment of reviewer 2). Furthermore, in further analyses we found that continuous representational manifolds support metric assessments of inputs relative to learned odors, which cannot be achieved by discrete representations. These results are now shown in Figure 5D-E and discussed explicitly in the text on p.11 (see also response to comment 3 of reviewer 1).

We preferred not to add a sentence in the Discussion about benefits of networks having both sources of inhibition, as we find this a bit too speculative.

At a few points in the manuscript, the authors use statements without actually providing evidence in terms of a Figure. Often the authors themselves acknowledge this, by adding the term "not shown" to the end of the sentence. I believe it will be helpful to the reader to be provided with figures or panels in support of the statements.

Thank you for this comment. We have provided additional data figures to support the following statements:

“ d_M was again increased upon learning, particularly between learned odors and reference classes representing other odors (Supplementary Figure 9)”

“decreasing amplification in assemblies of Scaled networks changed transformations towards the intermediate behavior, albeit with broader firing rate distributions than in Tuned networks (Supplementary Figure 6 B)”

Recommendations for the authors:

Reviewer #1 (Recommendations For The Authors):

Meissner-Bernard et al present a biologically constrained model of telencephalic area of adult zebrafish, a homologous area to the piriform cortex, and argue for the role of precisely balanced memory networks in olfactory processing.

This is interesting as it can add to recent evidence on the presence of functional subnetworks in multiple sensory cortices. It is also important in deviating from traditional accounts of memory systems as attractor networks. Evidence for attractor networks has been found in some systems, like in the head direction circuits in the flies. However, the presence of attractor dynamics in other modalities, like sensory systems, and their role in computation has been more contentious. This work contributes to this active line of research in experimental and computational neuroscience by suggesting that, rather than being represented in attractor networks and persistent activity, olfactory memories might be coded by balanced excitation-inhibitory subnetworks.

The paper is generally well-written, the figures are informative and of good quality, and multiple approaches and metrics have been used to test and support the main results of

the paper.

The main strength of the work is in: (1) direct link to biological parameters and measurements, (2) good controls and quantification of the results, and (3) comparison across multiple models.

(1) The authors have done a good job of gathering the current experimental information to inform a biological-constrained spiking model of the telencephalic area of adult zebrafish. The results are compared to previous experimental measurements to choose the right regimes of operation.

(2) Multiple quantification metrics and controls are used to support the main conclusions and to ensure that the key parameters are controlled for - e.g. when comparing across multiple models. (3) Four specific models (random, scaled I / attractor, and two variant of specific E-I networks - tuned I and tuned E+I) are compared with different metrics, helping to pinpoint which features emerge in which model.

Major problems with the work are: (1) mechanistic explanation of the results in specific E-I networks, (2) parameter exploration, and (3) the functional significance of the specific E-I model.

(1) The main problem with the paper is a lack of mechanistic analysis of the models. The models are treated like biological entities and only tested with different assays and metrics to describe their different features (e.g. different geometry of representation in Fig. 4). Given that all the key parameters of the models are known and can be changed (unlike biological networks), it is expected to provide a more analytical account of why specific networks show the reported results. For instance, what is the key mechanism for medium amplification in specific E/I network models (Fig. 3)? How does the specific geometry of representation/manifolds (in Fig. 4) emerge in terms of excitatory-inhibitory interactions, and what are the main mechanisms/parameters? Mechanistic account and analysis of these results are missing in the current version of the paper.

Precise balancing of excitation and inhibition in subnetworks would lead to the cancellation of specific dynamical modes responsible for the amplification of responses (hence, deviating from the attractor dynamics with an unstable specific mode). What is the key difference in the specific E/I networks here (tuned I or/and tuned E+I) which make them stand between random and attractor networks? Excitatory and inhibitory neurons have different parameters in the model (Table 1). Time constants of inhibitory and excitatory synapses are also different (P. 13). Are these parameters causing networks to be effectively more excitation dominated (hence deviating from a random spectrum which would be expected from a precisely balanced E/I network, with exactly the same parameters of E and I neurons)? It is necessary to analyse the network models, describe the key mechanism for their amplification, and pinpoint the key differences between E and I neurons which are crucial for this.

To address these comments we performed additional simulations and analyses at different levels. Please see our reply to comment (1) of the public review (reviewer 1) for a detailed description. We thank the reviewer for these constructive comments.

(2) The second major issue with the study is a lack of systematic exploration and analysis of the parameter space. Some parameters are biologically constrained, but not all the parameters. For instance, it is not clear what the justification for the choice of synaptic time scales are (with E synaptic time constants being larger than inhibition: $\tau_{\text{syn}_i} = 10 \text{ ms}$, $\tau_{\text{syn}_E} = 30 \text{ ms}$). How would the results change if they are varying these - and other unconstrained - parameters? It is important to show how the main results, especially the manifold localisation, would change by doing a systematic exploration of

the key parameters and performing some sensitivity analysis. This would also help to see how robust the results are, which parameters are more important and which parameters are less relevant, and to shed light on the key mechanisms.

We thank the reviewer for this comment. We have now carried out additional simulations with equal time constants for all neurons. Please see our reply to the public review for more details (comment 2 of reviewer 1).

(3) It is not clear what the main functional advantage of the specific E-I network model is compared to random networks. In terms of activity, they show that specific E-I networks amplify the input more than random networks (Fig. 3). But when it comes to classification, the effect seems to be very small (Fig. 5c). Description of different geometry of representation and manifold localization in specific networks compared to random networks is good, but it is more of an illustration of different activity patterns than proving a functional benefit for the network. The reader is still left with the question of what major functional benefits (in terms of computational/biological processing) should be expected from these networks, if they are to be a good model for olfactory processing and learning.

One possibility for instance might be that the tasks used here are too easy to reveal the main benefits of the specific models - and more complex tasks would be needed to assess the functional enhancement (e.g. more noisy conditions or more combination of odours). It would be good to show this more clearly - or at least discuss it in relation to computation and function.

Please see our reply to the public review (comment 3 of reviewer 1).

Specific comments:

*Abstract: "resulting in continuous representations that reflected both relatedness of inputs and *an individual's experience*"*

It didn't become apparent from the text or the model where the role of "individual's experience" component (or "internal representations" - in the next line) was introduced or shown (apart from a couple of lines in the Discussion)

We consider the scenario that that assemblies are the outcome of an experience-dependent plasticity process. To clarify this, we have now made a small addition to the text: "Biological memory networks are thought to store information by experience-dependent changes in the synaptic connectivity between assemblies of neurons."

P. 2: "The resulting state of "precise" synaptic balance stabilizes firing rates because inhomogeneities or fluctuations in excitation are tracked by correlated inhibition"

It is not clear what the "inhomogeneities" specifically refers to - they can be temporal, or they can refer to the quenched noise of connectivity, for instance. Please clarify what you mean.

The statement has been modified to be more precise: "...“precise” synaptic balance stabilizes firing rates because inhomogeneities in excitation across the population or temporal variations in excitation are tracked by correlated inhibition...".

P. 3 (and Methods): When odour stimulus is simulated in the OB, the activity of a fraction of mitral cells is increased (10% to 15 Hz) - but also a fraction of mitral cells is suppressed (5% to 2 Hz). What is the biological motivation or reference for this? It is not provided. Is

it needed for the results? Also, it is not explained how the suppressed 5% are chosen (e.g. randomly, without any relation to the increased cells?).

We thank the reviewer for this comment. These changes in activity directly reflect experimental observations. We apologize that we forgot to include the references reporting these observations (Friedrich and Laurent, 2001 and 2004); this is now fixed.

In our simulation, OB neurons do not interact with each other, and the suppressed 5% were indeed randomly selected. We changed the text in Methods accordingly to read: “An additional 75 randomly selected mitral cells were inhibited”

P. 4, L. 1-2: "... sparsely connected integrate-and-fire neurons with conductance-based synapses (connection probability {less than or equal to}5%)."

Specify the connection probability of specific subtypes (EE, EI, IE, II).

We now refer to the Methods section, where this information can be found.

“... conductance-based synapses (connection probability $\leq 5\%$, Methods)”

P. 4, L. 6-7: "Population activity was odor-specific and activity patterns evoked by uncorrelated OB inputs remained uncorrelated in Dp (Figure 1H)"

What would happen to correlated OB inputs (e.g. as a result of mixture of two overlapping odours) in this baseline state of the network (before memories being introduced to it)? It would be good to know this, as it sheds light on the initial operating regime of the network in terms of E/I balance and decorrelation of inputs.

This information was present in the original manuscript at (Figure 3) but we improved the writing to further clarify this issue: “ (...) we morphed a novel odor into a learned odor (Figure 3A), or a learned odor into another learned odor (Supplementary Figure 3B), and quantified the similarity between morphed and learned odors by the Pearson correlation of the OB activity patterns (input correlation). We then compared input correlations to the corresponding pattern correlations among E neurons in Dp (output correlation). In *rand* networks, output correlations increased linearly with input correlations but did not exceed them (Figure 3B and Supplementary Figure 3B)”

P. 4, L. 12-13: "Shuffling spike times of inhibitory neurons resulted in runaway activity with a probability of ~80%, .." Where is this shown?

(There are other occasions too in the paper where references to the supporting figures are missing).

We now provide the statistics: “Shuffling spike times of inhibitory neurons resulted in runaway activity with a probability of 0.79 ± 0.20 ”

P. 4: "In each network, we created 15 assemblies representing uncorrelated odors. As a consequence, ~30% of E neurons were part of an assembly ..."

$15 \times 100 / 4000 = 37.5\%$ - so it's closer to 40% than 30%. Unless there is some overlap?

Yes: despite odors being uncorrelated and connectivity being random, some neurons (6 % of E neurons) belong to more than one assembly.

P. 4: "When a reached a critical value of ~6, networks became unstable and generated runaway activity (Figure 2B)."

Can this transition point be calculated or estimated from the network parameters, and linked to the underlying mechanisms causing it?

We thank the reviewer for this interesting question. The instability arises when inhibitions fails to counterbalance efficiently the increased recurrent excitation within Dp. The transition point is difficult to estimate, as it can depend on several parameters, including the probability of E to E connections, their strength, assembly size, and others. We have therefore not attempted to estimate it analytically.

P. 4: "Hence, non-specific scaling of inhibition resulted in a divergence of firing rates that exhausted the dynamic range of individual neurons in the population, implying that homeostatic global inhibition is insufficient to maintain a stable firing rate distribution."

I don't think this is justified based on the results and figures presented here (Fig. 2E) - the interpretation is a bit strong and biased towards the conclusions the authors want to draw.

To more clearly illustrate the finding that in *Scaled* networks, assembly neurons are highly active (close to maximal realistic firing rates) whereas non-assembly neurons are nearly silent we have now added Supplementary Fig. 2B. Moreover, we have toned down the text: "Hence, non-specific scaling of inhibition resulted in a large and biologically unrealistic divergence of firing rates (Supplementary Figure 2B) that nearly exhausted the dynamic range of individual neurons in the population, indicating that homeostatic global inhibition is insufficient to maintain a stable firing rate distribution"

P. 5, third paragraph: Description of Figure 2I, inset is needed, either in the text or caption.

The inset is now referred to in the text: "we projected synaptic conductances of each neuron onto a line representing the E/I ratio expected in a balanced network ("balanced axis") and onto an orthogonal line ("counter-balanced axis"; Figure 2I inset, Methods)."

P. 5, last paragraph: another example of writing about results without showing/referring to the corresponding figures:

"In rand networks, firing rates increased after stimulus onset and rapidly returned to a low baseline after stimulus offset. Correlations between activity patterns evoked by the same odor at different time points and in different trials were positive but substantially lower than unity, indicating high variability ..."

And the continuation with similar lack of references on P. 6:

"Scaled networks responded to learned odors with persistent firing of assembly neurons and high pattern correlations across trials and time, implying attractor dynamics (Hopfield, 1982; Khona and Fiete, 2022), whereas Tuned networks exhibited transient responses and modest pattern correlations similar to rand networks."

Please go through the Results and fix the references to the corresponding figures on all instances.

We thank the reviewer for pointing out these overlooked figure references, which are now fixed.

P. 8: "These observations further support the conclusion that E/I assemblies locally constrain neuronal dynamics onto manifolds."

As discussed in the general major points, mechanistic explanation in terms of how the interaction of E/I dynamics leads to this is missing.

As discussed in the reply to the public review (comment 3 of reviewer 1), we have now provided more mechanistic analyses of our observations.

P. 9: "Hence, E/I assemblies enhanced the classification of inputs related to learned patterns." The effect seems to be very small. Also, any explanation for why for low test-target correlation the effect is negative (random doing better than tuned E/I)?

The size of the effect ($p_{\text{learned}} - p_{\text{novel}} = 0.074$; difference of means; Figure 5C) may appear small in terms of absolute probability, but it is substantial relative to the maximum possible increase ($1 - p_{\text{novel}} = 0.133$; Figure 5C). The fact that for low test-target correlations the effect is negative is a direct consequence of the positive effect for high test-target correlations and the presence of 2 learned odors in the 4-way forced choice task.

P. 9: "In Scaled I networks, creating two additional memories resulted in a substantial increase in firing rates, particularly in response to the learned and related odors" Where is this shown? Please refer to the figure.

We thank the reviewer again for pointing this out. We forgot to include a reference to the relevant figure which has now been added in the revised manuscript (Figure 6C).

P. 10: "The resulting Tuned networks reproduced additional experimental observations that were not used as constraints including irregular firing patterns, lower output than input correlations, and the absence of persistent activity"

It is difficult to present these as "additional experimental observations", as all of them are negative, and can exist in random networks too - hence cannot be used as biological evidence in favour of specific E/I networks when compared to random networks.

We agree with the reviewer that these additional experimental observations cannot be used as biological evidence favouring Tuned E+I networks over random networks. We here just wanted to point out that additional observations which we did not take into account to fit the model are not invalidating the existence of E-I assemblies in biological networks. As assemblies tend to result in persistent activity in other types of networks, we feel that this observation is worth pointing out.

Methods:

P. 13: Describe the parameters of Eq. 2 after the equation.

Done.

P. 13: "The time constants of inhibitory and excitatory synapses were 10 ms and 30 ms, respectively."

What is the (biological) justification for the choice of these parameters?

How would varying them affect the main results (e.g. local manifolds)?

We chose a relatively slow time constant for excitatory synapses because experimental data indicate that excitatory synaptic currents in Dp and piriform cortex contain a prominent NMDA component. We have now also simulated networks with equal time constants for excitatory and inhibitory synapses and equal biophysical parameters for excitatory and

inhibitory neurons, which did not affect the main results (see also reply to the public review: comment 2 of reviewer 1).

P. 14: "Care was also taken to ensure that the variation in the number of output connections was low across neurons" How exactly?

More detailed explanations have now been added in the Methods section: "connections of a presynaptic neuron y to postsynaptic neurons x were randomly deleted when their total number exceeded the average number of output connections by $\geq 5\%$, or added when they were lower by $\geq 5\%$."

Reviewer #2 (Recommendations For The Authors):

Congratulations on the great and interesting work! The results were nicely presented and the idea of continuous encoding on manifolds is very interesting. To improve the quality of the paper, in addition to the major points raised in the public review, here are some more detailed comments for the paper:

(1) Generally, citations have to improve. Spiking networks with excitatory assemblies and different architectures of inhibitory populations have been studied before, and the claim about improved network stability in co-tuned E-I networks has been made in the following papers that need to be correctly cited:

- Vogels TP, Sprekeler H, Zenke F, Clopath C, Gerstner W. 2011. Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. *Science* 334:1-7. doi:10.1126/science.1212991 (mentions that emerging precise balance on the synaptic weights can result in the overall network stability)
- Lagzi F, Bustos MC, Oswald AM, Doiron B. 2021. Assembly formation is stabilized by Parvalbumin neurons and accelerated by Somatostatin neurons. bioRxiv doi: <https://doi.org/10.1101/2021.09.06.459211> (among other things, contrasts stability and competition which arises from multistable networks with global inhibition and reciprocal inhibition)
- Rost T, Deger M, Nawrot MP. 2018. Winnerless competition in clustered balanced networks: inhibitory assemblies do the trick. *Biol Cybern* 112:81-98. doi:10.1007/s00422-017-0737-7 (compares different architectures of inhibition and their effects on network dynamics)
- Lagzi F, Fairhall A. 2022. Tuned inhibitory firing rate and connection weights as emergent network properties. bioRxiv 2022.04.12.488114. doi:10.1101/2022.04.12.488114 (here, see the eigenvalue and UMAP analysis for a network with global inhibition and E/I assemblies)

Additionally, there are lots of pioneering work about tracking of excitatory synaptic inputs by inhibitory populations, that are missing in references. Also, experimental work that show existence of cell assemblies in the brain are largely missing. On the other hand, some references that do not fit the focus of the statements have been incorrectly cited.

The authors may consider referencing the following more pertinent studies on spiking networks to support the statement regarding attractor dynamics in the first paragraph in the Introduction (the current citations of Hopfield and Kohonen are for rate-based networks):

- Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience*, 26(4), 1314-1328. <https://doi.org/10.1523/JNEUROSCI.3733-05.2006>

- Wang, X.-J. (2008). *Decision making in recurrent neuronal circuits*. *Neuron*, 60(2), 215-234. <https://doi.org/10.1016/j.neuron.2008.09.034>
- F. Lagzi, & S. Rotter. (2015). *Dynamics of competition between subnetworks of spiking neuronal networks in the balanced state*. *PloS One*.
- Goldman-Rakic, P. S. (1995). *Cellular basis of working memory*. *Neuron*, 14(3), 477-485.
- Rost T, Deger M, Nawrot MP. 2018. *Winnerless competition in clustered balanced networks: inhibitory assemblies do the trick*. *Biol Cybern* 112:81-98. doi:10.1007/s00422-017-0737-7.
- Amit DJ, Tsodyks M (1991) *Quantitative study of attractor neural network retrieving at low spike rates: I. substrate-spikes, rates and neuronal gain*. *Network* 2:259-273.
- Mazzucato, L., Fontanini, A., & La Camera, G. (2015). *Dynamics of Multistable States during Ongoing and Evoked Cortical Activity*. *Journal of Neuroscience*, 35(21), 8214-8231.

We thank the reviewer for the references suggestions. We have carefully reviewed the reference list and made the following changes, which we hope address the reviewer's concerns:

- (1) We adjusted References about network stability in co-tuned E-I networks.
- (2) We added the Lagzi & Rotter (2015), Amit et al. (1991), Mazzucato et al. (2015) and GoldmanRakic (1995) papers in the Introduction as studies on attractor dynamics in spiking neural networks. We preferred to omit the two X.J Wang papers, as they describe attractors in decision making rather than memory processes.
- (3) We added the Ko et al. 2011 paper as experimental evidence for assemblies in the brain. In our view, there are few experimental studies showing the existence of cell assemblies in the brain, which we distinguish from cell ensembles, group of coactive neurons.
- (4) We also included Hennequin 2018, Brunel 2000, Lagzi et al. 2021 and Eckmann et al. 2024, which we had not cited in the initial manuscript.
- (5) We removed the Wiechert et al. 2010 reference as it does not support the statement about geometry-preserving transformation by random networks.

(2) *The gist of the paper is about how the architecture of inhibition (reciprocal vs. global in this case) can determine network stability and salient responses (related to multistable attractors and variations) for classification purposes. It would improve the narrative of the paper if this point is raised in the Introduction and Discussion section. Also see a relevant paper that addresses this point here:*

Lagzi F, Bustos MC, Oswald AM, Doiron B. 2021. *Assembly formation is stabilized by Parvalbumin neurons and accelerated by Somatostatin neurons*. *bioRxiv* doi: <https://doi.org/10.1101/2021.09.06.459211>

Classification has long been proposed to be a function of piriform cortex and autoassociative memory networks in general, and we consider it important. However, the computational function of Dp or piriform cortex is still poorly understood, and we do not focus only on odor classification as a possibility. In fact, continuous representational manifolds also support other functions such as the quantification of distance relationships of an input to previously memorized stimuli, or multi-layer network computations (including classification). In the revised manuscript, we have performed additional analyses to explore these notions in more detail, as explained above (response to public reviews, comment 3 of reviewer 1).

Furthermore, we have now expanded the discussion of potential computational functions of *Tuned* networks and explicitly discuss classification but also other potential functions.

(3) A plot for the values of the inhibitory conductances in Figure 1 would complete the analysis for that section.

In Figure 1, we decided to only show the conductances that we use to fit our model, namely the afferent and total synaptic conductances. As the values of the inhibitory conductances can be derived from panel E, we refrained from plotting them separately for the sake of simplicity.

(4) How did the authors calculate correlations between activity patterns as a function of time in Figure 2E, bottom row? Does the color represent correlation coefficient (which should not be time dependent) or is it a correlation function? This should be explained in the Methods section.

The color represents the Pearson correlation coefficient between activity patterns within a narrow time window (100 ms). We updated the Figure legend to clarify this: “Mean correlation between activity patterns evoked by a learned odor at different time points during odor presentation. Correlation coefficients were calculated between pairs of activity vectors composed of the mean firing rates of E neurons in 100 ms time bins. Activity vectors were taken from the same or different trials, except for the diagonal, where only patterns from different trials were considered.”

(5) Figure 3 needs more clarification (both in the main text and the figure caption). It is not clear what the axes are exactly, and why the network responses for familiar and novel inputs are different. The gray shaded area in panel B needs more explanation as well.

We thank the reviewer for the comment. We have improved Figure 3A, the figure caption, as well as the text (see p.6). We hope that the figure is now clearer.

(6) The "scaled I" network, known for representing input patterns in discrete attractors, should exhibit clear separation between network responses in the 2D PC space in the PCA plots. However, Figure 4D and Figure 6D do not reflect this, as all network responses are overlapped. Can the authors explain the overlap in Figure 4D?

In Figure 4D, activity of Scaled networks is distributed between three subregions in state space that are separated by the first 2 PCs. Two of them indeed correspond to attractor states representing the two learned odors while the third represents inputs that are not associated with these attractor states. To clarify this, please see also the density plot in Figure 4E. The few datapoints between these three subregions are likely outliers generated by the sequential change in inputs, as described in Supplementary Figure 8C.

(7) The reason for writing about the ISN networks is not clear. Co-tuned E-I assemblies do not necessarily have to operate in this regime. Also, the results of the paper do not rely on any of the properties of ISNs, but they are more general. Authors should either show the paradoxical effect associated with ISN (i.e., if increasing input to I neurons decreases their responses) or show ISN properties using stability analysis (See computational research conducted at the Allen Institute, namely Millman et al. 2020, eLife). Currently, the paper reads as if being in the ISN regime is a necessary requirement, which is not true. Also, the arguments do not connect with the rest of the paper and never show up again. Since we know it is not a requirement, there is no need to have those few

sentences in the Results section. Also, the choice of $\alpha=5.0$ is extreme, and therefore, it would help to judge the biological realism if the raster plots for Figs 2-6 are shown.

We have toned down the part on ISN and reduced it to one sentence for readers who might be interested in knowing whether activity is inhibition-stabilized or not. We have also added the reference to the Tsodyks et al. 1997 paper from which we derive our stability analysis. The text now reads “Hence, pDp_{sim} entered a balanced state during odor stimulation (Figure 1D, E) with recurrent input dominating over afferent input, as observed in pDp (Rupprecht and Friedrich, 2018). Shuffling spike times of inhibitory neurons resulted in runaway activity with a probability of 0.79 ± 0.20 , demonstrating that activity was inhibition-stabilized (Sadeh and Clopath, 2020b, Tsodyks et al., 1997).”

We have now also added the raster plots as suggested by the reviewer (see Figure 2D, Supplementary Figure 1 G, Supplementary Figure 4). We thank the reviewer for this comment.

(8) In the abstract, authors mention "fast pattern classification" and "continual learning," but in the paper, those issues have not been addressed. The study does not include any synaptic plasticity.

Concerning “continual learning” we agree that we do not simulate the learning process itself. However, Figure 6 show results of a simulation where two additional patterns were stored in a network that already contained assemblies representing other odors. We consider this a crude way of exploring the end result of a “continual learning” process. “Fast pattern classification” is mentioned because activity in balanced networks can follow fluctuating inputs with high temporal resolution, while networks with stable attractor states tend to be slow. This is likely to account for the occurrence of hysteresis effects in *Scaled* but not *Tuned* networks as shown in Supplementary

Fig. 8.

(9) In the Introduction, the first sentence in the second paragraph reads: "... when neurons receive strong excitatory and inhibitory synaptic input ...". The word strong should be changed to "weak".

Also, see the pioneering work of Brunel 2000.

In classical balanced networks, strong excitatory inputs are counterbalanced by strong inhibitory inputs, leading to a fluctuation-driven regime. We have added Brunel 2000.

(10) In the second paragraph of the introduction, the authors refer to studies about structural co-tuning (e.g., where "precise" synaptic balance is mentioned, and Vogels et al. 2011 should be cited there) and functional co-tuning (which is, in fact, different than tracking of excitation by inhibition, but the authors refer to that as co-tuning). It makes it easier to understand which studies talk about structural co-tuning and which ones are about functional co-tuning. The paper by Znamenski 2018, which showed both structural and functional tuning in experiments, is missing here.

We added the citation to the now published paper by Znamenskyi et al. (2024).

(11) The third paragraph in the Introduction misses some references that address network dynamics that are shaped by the inhibitory architecture in E/I assemblies in spiking networks, like Rost et al 2018 and Lagzi et al 2021.

These references have been added.

(12) The last sentence of the fourth paragraph in the Introduction implies that functional co-tuning is due to structural co-tuning, which is not necessarily true. While structural co-tuning results in functional co-tuning, functional co-tuning does not require structural co-tuning because it could arise from shared correlated input or heterogeneity in synaptic connections from E to I cells.

We generally agree with the reviewer, but we are uncertain which sentence the reviewer refers to.

We assume the reviewer refers to the last sentence of the second (rather than the fourth paragraph), which explicitly mentions the “...structural basis of E/I co-tuning...”. If so, we consider this sentence still correct because the “structural basis” refers not specifically to E/I assemblies, but also includes any other connectivity that may produce co-tuning, including the connectivity underlying the alternative possibilities mentioned by the reviewer (shared correlated input or heterogeneity of synaptic connections).

(13) In order to ensure that the comparison between network dynamics is legit, authors should mention up front that for all networks, the average firing rates for the excitatory cells were kept at 1 Hz, and the background input was identical for all E and I cells across different networks.

We slightly revised the text to make this more clear “We (...) uniformly scaled I-to-E connection weights by a factor of χ until E population firing rates in response to learned odors matched the corresponding firing rates in *rand* networks, i.e., 1 Hz”

(14) In the last paragraph on page 5, my understanding was that an individual odor could target different cells within an assembly in different trials to generate trial to trial variability. If this is correct, this needs to be mentioned clearly.

This is not correct, an odor consists of 150 activated mitral cells with defined firing rates. As now mentioned in the Methods, “Spikes were then generated from a Poisson distribution, and this process was repeated to create trial-to-trial variability.”

(15) The last paragraph on page 6 mentions that the four OB activity patterns were uncorrelated but if they were designed as in Figure 4A, due to the existing overlap between the patterns, they cannot be uncorrelated.

This appears to be a misunderstanding. We mention in the text (and show in Figure 4B) that the four odors which “... were assigned to the corners of a square...” are uncorrelated. The intermediate odors are of course not uncorrelated. We slightly modified the corresponding paragraph (now on page 7) to clarify this: “The subspace consisted of a set of OB activity patterns representing four uncorrelated pure odors and mixtures of these pure odors. Pure odors were assigned to the corners of a square and mixtures were generated by selecting active mitral cells from each of the pure odors with probabilities depending on the relative distances from the corners (Figure 4A, Methods).”

(16) The notion of "learned" and "novel" odors may be misleading as there was no plasticity in the network to acquire an input representation. It would be beneficial for the authors to clarify that by "learned," they imply the presence of the corresponding E assembly for the odor in the network, with the input solely impacting that assembly. Conversely, for "novel" inputs, the input does not target a predefined assembly. In Figure 2 and Figure 4, it would be especially helpful to have the spiking raster plots of some sample E and I cells.

As suggested by the reviewer, we have modified the existing spiking raster plots in Figure 2, such that they include examples of responses to both learned and novel odors. We added spiking raster plots showing responses of I neurons to the same odors in Supplementary Figure 1F, as well as spiking raster plots of E neurons in Supplementary Figure 4A. To clarify the usage of “learned” and “novel”, we have added a sentence in the Results section: “We thus refer to an odor as “learned” when a network contains a corresponding assembly, and as “novel” when no such assembly is present.”.

(17) In the last paragraph of page 8, can the authors explain where the asymmetry comes from?

As mentioned in the text, the asymmetry comes from the difference in the covariance structure of different classes. To clarify, we have rephrased the sentence defining the Mahalanobis distance:

“This measure quantifies the distance between the pattern and the class center, taking into account covariation of neuronal activity within the class. In bidirectional comparisons between patterns from different classes, the mean dM may be asymmetric if neural covariance differs between classes.”

(18) The first paragraph of page 9: random networks are not expected to perform pattern classification, but just pattern representation. It would have been better if the authors compared Scaled I network with E/I co-tuned network. Regardless of the expected poorer performance of the E/I co-tuned networks, the result would have been interesting.

Please see our reply to the public review (reviewer 2).

(19) Second paragraph on page 9, the authors should provide statistical significance test analysis for the statement "... was significantly higher ...".

We have performed a Wilcoxon signed-rank test, and reported the p-value in the revised manuscript ($p < 0.01$).

(20) The last sentence in the first paragraph on page 11 is not clear. What do the authors mean by "linearize input-output functions", and how does it support their claim?

We have now amended this sentence to clarify what we mean: “...linearize the relationship between the mean input and output firing rates of neuronal populations...”.

(21) In the first sentence of the last paragraph on page 11, the authors mentioned “high variability”, but it is not clear compared with which of the other 3 networks they observed high variability.

Structurally co-tuned E/I networks are expected to diminish network-level variability.

“High variability” refers to the variability of spike trains, which is now mentioned explicitly in the text. We hope this more precise statement clarifies this point.

(22) Methods section, page 14: "firing rates decreased with a time constant of 1, 2 or 4 s". How did they decrease? Was it an implementation algorithm? The time scale of input presentation is 2 s and it overlaps with the decay time constant (particularly with the one with 4 s decrease).

Firing rates decreased exponentially. We have added this information in the Methods section.

Reviewer #3 (Recommendations For The Authors):

In the following, I suggest minor corrections to each section which I believe can improve the manuscript.

- There was no github link to the code in the manuscript. The code should be made available with a link to github in the final manuscript.

The code can be found here: <https://github.com/clairemb90/pDp-model>. The link has been added in the Methods section.

Figure 1:

- Fig. 1A: call it pDp not Dp. Please check if this name is consistent in every figure and the text.

Thank you for catching this. Now corrected in Figure 1, Figure 2 and in the text.

- The authors write: "Hence, pDpsim entered an inhibition-stabilized balanced state (Sadeh and Clopath, 2020b) during odor stimulation (Figure 1D, E)." and then later "Shuffling spike times of inhibitory neurons resulted in runaway activity with a probability of ~80%, demonstrating that activity was indeed inhibition-stabilized. These results were robust against parameter variations (Methods)." I would suggest moving the second sentence before the first sentence, because the fact that the network is in the ISN regime follows from the shuffled spike timing result.

Also, I'd suggest showing this as a supplementary figure.

We thank the reviewer for this comment. We have removed “inhibition-stabilized” in the first sentence as there is no strong evidence of this in Rupprecht and Friedrich, 2018. And removed “indeed” in the second sentence. We also provided more detailed statistics. The text now reads “Hence, pDpsim entered a balanced state during odor stimulation (Figure 1D, E) with recurrent input dominating over afferent input, as observed in pDp (Rupprecht and Friedrich, 2018). Shuffling spike times of inhibitory neurons resulted in runaway activity with a probability of 0.79 ± 0.20 , demonstrating that activity was inhibition-stabilized (Sadeh and Clopath, 2020b).”

Figure 2:

- "... Scaled I networks (Figure 2H." Missing)

Corrected.

- The authors write "Unlike in Scaled I networks, mean firing rates evoked by novel odors were indistinguishable from those evoked by learned odors and from mean firing rates in rand networks (Figure 2F)."

Why is this something you want to see? Isn't it that novel stimuli usually lead to high responses? Eg in the paper Schulz et al., 2021 (eLife) which is also cited by the authors it is shown that novel responses have high onset firing rates. I suggest clarifying this (same in the context of Fig. 3C).

In Dp and piriform cortex, firing rates evoked by learned odors are not substantially different from firing rates evoked by novel odors. While small differences between responses to learned versus novel odors cannot be excluded, substantial learning-related differences in firing rates, as observed in other brain areas, have not been described in Dp or piriform cortex. We added references in the last paragraph of p.5. Note that the paper by Schulz et al. (2021) models a different type of circuit.

- Fig. 2B: Indicate in figure caption that this is the case "Scaled I"

This is not exactly the case "Scaled I", as the parameter χ (increased I to E strength) is set to 1.

- Suppl Fig. 2I: Is E&F ever used in the manuscript? I couldn't find a reference. I suggest removing it if not needed.

Suppl. Fig 2I E&F is now Suppl Fig.1G&H. We now refer to it in the text: "Activity of networks with E assemblies could not be stabilized around 1 Hz by increasing connectivity from subsets of I neurons receiving dense feed-forward input from activated mitral cells (Supplementary Figure 1GH; Sadeh and Clopath, 2020)."

Figure 3:

- As mentioned in my comment in the public review section, I find the arguments about pattern completion a little bit confusing. For me it's not clear why an increase of output correlations over input correlations is considered "pattern completion" (this is not to say that I don't find the nonlinear increase of output correlations interesting). For me, to test pattern completion with second-order statistics one would need to do a similar separation as in Suppl Fig. 3, ie measuring the pairwise correlation at cells in the assembly L that get direct input from L OB with cells in the assembly L that do not get direct input from OB. If the pairwise correlations of assembly cells which do not get direct input from OB increase in correlations, I would consider this as pattern completion (similar to the argument that increase in firing rate in cells which are not directly driven by OB are considered a sign of pattern completion).

Also, for me it now seems like that there are contradictory results, in Fig. 3 only Scaled I can lead to pattern completion while in the context of Suppl. Fig. 3 the authors write "We found that assemblies were recruited by partial inputs in all structured pDpsim networks (Scaled and Tuned) without a significant increase in the overall population activity (Supplementary Figure 3A)." I suggest clarifying what the authors exactly mean by pattern completion, why the increase of output correlations above input correlations can be considered as pattern completion, and why the results differs when looking at firing rates versus correlations.

Please see our reply to the public review (reviewer 3).

- I actually would suggest adding Suppl. Fig. 3 to the main figure. It shows a more intuitive form of pattern completion and in the text there is a lot of back and forth between Fig. 3 and Suppl. Fig. 3

We feel that the additional explanations and panels in Fig.3 should clarify this issue and therefore prefer to keep Supplementary Figure 3 as part of the Supplementary Figures for simplicity.

- In the whole section "We next explored effects of assemblies ... prevented strong recurrent amplification within E/I assemblies." the authors could provide a link to the respective panel in Fig. 2 after each statement. This would help the reader follow your arguments.

We thank the reviewer for pointing this out. The references to the appropriate panels have been added.

- Fig. 3A: I guess the x-axis has been shifted upwards? Should be at zero.

We have modified the x-axis to make it consistent with panels B and C.

- Fig. 3B: In the figure caption, the dotted line is described as the novel odor but it is actually the unit line. The dashed lines represent the reference to the novel odor.

Fixed.

- Fig. 3C: The " is missing for Pseudo-Assembly N

Fixed.

- "...or a learned odor into another learned odor." Have here a ref to the Supplementary Figure 3B.

Added.

Figure 4:

- "This geometry was largely maintained in the output of rand networks, consistent with the notion that random networks tend to preserve similarity relationships between input patterns (Babadi and Sompolinsky, 2014; Marr, 1969; Schaffer et al., 2018; Wiechert et al., 2010)." I suggest adding here reference to Fig. 4D (left).

Added.

- Please add a definition of E/I assemblies. How do the authors define E/I assemblies? I think they consider both, Tuned I and Tuned E+I as E/I assemblies? In Suppl. Fig. 2I it looks like tuned feedforward input is defined as E/I assemblies.

We thank the reviewer for pointing this out. E/I assemblies are groups of E and I neurons with enhanced connectivity. In other words, in E/I assemblies, connectivity is enhanced not only between subsets of E neurons, but also between these E neurons and a subset of I neurons. This is now clarified in the text: "We first selected the 25 I neurons that received the largest number of connections from the 100 E neurons of an assembly. To generate E/I assemblies, the connectivity between these two sets of neurons was then enhanced by two procedures." We removed "E/I assemblies" in Suppl. Fig. 2, where the term was not used correctly, and apologize for the confusion.

- Suppl. Fig. 4: Could the authors please define what they mean by "Loadings"

The loadings indicate the contribution of each neuron to each principal component, see adjusted legend of Suppl. Fig. 4: "G. Loading plot: contribution of neurons to the first two PCs of a rand and a Tuned E+I network (Figure 4D)."

- Fig. 4F: The authors might want to normalize the participation ratio by the number of neurons (see e.g. Dahmen et al., 2023 bioRxiv, "relative PR"), so the PR is bound between 0 and 1 and the dependence on N is removed.

We thank the reviewer for the suggestion, but we prefer to use the non-normalized PR as we find it more easily interpretable (e.g. number of attractor states in Scaled networks).

- Fig. 4G&H: as mentioned in the public review, I'd add the case of Scaled I to be able to compare it to the Tuned E+I case.

As already mentioned in the public review, we thank the reviewer for this suggestion, which we have implemented.

- Figure caption Fig. 4H "Similar results were obtained in the full-dimensional space." I suggest showing this as a supplemental panel.

Since this only adds little information, we have chosen not to include it as a supplemental panel to avoid overloading the paper with figures.

Figure 5:

- As mentioned in the public review, I suggest that the authors add the Scaled I case to Fig. 5 (it's shown in all figures and also in Fig. 6 again). I guess for Scaled I the separation between L and M will be very good?

Please see our reply to the public review (reviewer 3).

- Fig. 5A&B: I am a bit confused about which neurons are drawn to calculate the Mahalanobis distance. In Fig. 5A, the schematic indicates that the vector B from which the neurons are drawn is distinct from the distribution Q. For the example of odor L, the distribution Q consists of pure odor L with odors that have little mixtures with the other odors. But the vector v for odor L seems to be drawn only from odors that have slightly higher mixtures (as shown in the schematic in Fig. 5A). Is there a reason to choose the vector v from different odors than the distribution Q?

The distribution Q and the vector v consist of activity patterns across the same neurons in response to different odors. The reason to choose a different odor for v was to avoid having this test datapoint being included in the distribution Q. We also wanted Q to be the same for all test datapoints.

What does "drawn from whole population" mean? Does this mean that the vectors are drawn from any neuron in pDp? If yes, then I don't understand how the authors can distinguish between different odors (L,M,O,N) on the y-axis. Or does "whole population" mean that the vector is drawn across all assemblies as shown in the schematic in Fig. 5A and the case "neurons drawn from (pseudo-) assembly" means that the authors choose only one specific assembly? In any case, the description here is a bit confusing, I think it would help the reader to clarify those terms better.

Yes, "drawn from whole population" means that we randomly draw 80 neurons from the 4000 E neurons in pDp. The y-axis means that we use the activity patterns of these neurons evoked by one of the 4 odors (L, M, N, O) as reference. We have modified the Figure legend to clarify this: " d_M was computed based on the activity patterns of 80 E neurons drawn from the

four (pseudo-) assemblies (top) or from the whole population of 4000 E neurons (bottom). Average of 50 draws.”

- Suppl Fig. 5A: In the schematic the distance is called $d_E(\bar{Q}, \bar{V})$ while the colorbar has $d_E(\bar{Q}, \bar{Q})$ with the Qs in different color. The green Q should be a V.

We thank the reviewer for spotting this mistake, it is now fixed.

- Fig. 5: Could the authors comment on the fact that a random network seems to be very good in classifying patterns on it's own. Maybe in the Discussion?

The task shown in Figure 5 is a relatively easy one, a forced-choice between four classes which are uncorrelated. In Supplementary Figure 9, we now show classification for correlated classes, which is already much harder.

Figure 6:

- Is the correlation induced by creating mixtures like in the other Figures? Please clarify how the correlations were induced.

We clarified this point in the Methods section: “The pixel at each vertex corresponded to one pure odor with 150 activated and 75 inhibited mitral cells (...) and the remaining pixels corresponded to mixtures. In the case of correlated pure odors (Figure 6), adjacent pure odors shared half of their activated and half of their inhibited cells.”. An explicit reference to the Methods section has also been added to the figure legend.

- Fig. 6C (right): why don't we see the clear separation in PC space as shown in Fig. 4? Is this related to the existence of correlations? Please clarify.

Yes. The assemblies corresponding to the correlated odors X and Y overlap significantly, and therefore responses to these odors cannot be well separated, especially for Scaled networks. We added the overlap quantification in the Results section to make this clear. “These two additional assemblies had on average 16% of neurons in common due to the similarity of the odors.”

- "Furthermore, in this regime of higher pattern similarity, dM was again increased upon learning, particularly between learned odors and reference classes representing other odors (not shown)." Please show this (maybe as a supplemental figure).

We now show the data in Supplementary Figure 9.

Discussion:

- The authors write: "We found that transformations became more discrete map-like when amplification within assemblies was increased and precision of synaptic balance was reduced. Likewise, decreasing amplification in assemblies of Scaled networks changed transformations towards the intermediate behavior, albeit with broader firing rate distributions than in Tuned networks (not shown)."

Where do I see the first point? I guess when I compare in Fig. 4D the case of Scaled I vs Tuned E+I, but the sentence above sounds like the authors showed this in a more step-wise way eg by changing the strength of α or β (as defined in Fig. 1).

Also I think if the authors want to make the point that decreasing amplification in assemblies changes transformation with a different rate distribution in scaled vs tuned networks, the authors should show it (eg adding a supplemental figure).

The first point is indeed supported by data from different figures. Please note that the revised manuscript now contains further simulations that reinforce this statement, particularly those shown in Supplementary Figure 6, and that this point is now discussed more extensively in the Discussion. We hope that these revisions clarify this general point.

The data showing effects of decreasing amplification in assemblies is now shown in Supplementary Figure 6 (Scaled[adjust])

- I suggest adding the citation Znamenskiy et al., 2024 (Neuron; <https://doi.org/10.1016/j.neuron.2023.12.013>), which shows that excitatory and inhibitory (PV) neurons with functional similarities are indeed strongly connected in mouse V1, suggesting the existence of E/I assembly structure also in mammals.

Done.

<https://doi.org/10.7554/eLife.96303.2.sa0>