

Structural spine plasticity: learning and forgetting of odor-specific subnetworks in the olfactory bulb

John Hongyu Meng^{1,□}, Hermann Riecke^{1,*}

¹ Engineering Sciences and Applied Mathematics, Northwestern University, Evanston, IL, 60208, USA

□Current Address: Center for Neural Science, New York University, New York, NY, 10003, USA

* h-riecke@northwestern.edu

Abstract

Learning to discriminate between different sensory stimuli is essential for survival. In rodents, the olfactory bulb, which contributes to odor discrimination via pattern separation, exhibits extensive structural synaptic plasticity involving the formation and removal of synaptic spines, even in adult animals. The network connectivity resulting from this plasticity is still poorly understood. To gain insight into this connectivity we present here a computational model for the structural plasticity of the reciprocal synapses between the dominant population of excitatory principal neurons and inhibitory interneurons. It incorporates the observed modulation of spine stability by odor exposure. The model captures the striking experimental observation that the exposure to odors does not always enhance their discriminability: while training with similar odors enhanced their discriminability, training with dissimilar odors actually *reduced* the discriminability of the training stimuli (Chu et al, 2016). Strikingly, this differential learning does not require the activity-dependence of the spine stability and occurs also in a model with purely random spine dynamics in which the spine density is changed homogeneously, e.g., due to a global signal. However, the experimentally observed odor-specific reduction in the response of principal cells as a result of extended odor exposure (Kato et al. 2012) and the concurrent disinhibition of a subset of principal cells arise only in the activity-dependent model. Moreover, this model predicts the experimentally testable recovery of odor response through weak but not through strong odor re-exposure and the forgetting of odors via exposure to interfering odors. Combined with the experimental observations, the computational model provides strong support for the prediction that odor exposure leads to the formation of odor-specific subnetworks in the olfactory bulb.

Author Summary

A key feature of the brain is its ability to learn through the plasticity of its network. The olfactory bulb in the olfactory system is a remarkable brain area whose anatomical structure evolves substantially still in adult animals by establishing new synaptic connections and removing existing ones. We present a computational model for this process and employ it to interpret recent experimental results. By comparing the results of our model with those of a random control model we identify various experimental observations that lend strong support to the notion that the network of the olfactory

SUBMISSION

bulb comprises learned, odor-specific subnetworks. Moreover, our model explains the recent observation that the learning of odors does not always improve their discriminability and provides testable predictions for the recovery of odor response after repeated odor exposure and for when the learning of new odors interferes with retaining the memory of familiar odors.

41
42
43
44
45

Introduction

Learning and long-term memory are to a large extent implemented through the plasticity of neuronal connectivity. In adult animals this plasticity is typically dominated by long-term potentiation and long-term depression, which change the strength of the synapses, while during development structural plasticity in the form of the addition and removal of neurons as well as the formation and pruning of synapses plays a central role. Strikingly, in olfaction structural plasticity appears to remain a key component of learning even in adult animals. Thus, new adult-born granule-cell interneurons are persistently integrated into the olfactory bulb, which is the first brain region to receive sensory input from the nose, and new synapses between granule cells and the principal mitral cells are formed. In parallel, granule cells undergo controlled apoptosis [1,2], the extent of which, however, has recently come under dispute ([3], see also discussion with the reviewers in [4]), and synapses are removed [5,6]. This makes the olfactory bulb an excellent system to study structural plasticity.

A widely assumed function of the olfactory bulb is to aid in the discrimination between odors by enhancing differences in the activity patterns that they evoke [7,8].

Computational modeling has identified a number of mechanism by which this can be achieved, which are not mutually exclusive. This includes a nonlinear response of the neurons without any lateral interactions among the neurons [9] or with random [10] or all-to-all connectivity [11]. Given the extensive plasticity of the olfactory bulb, adaptive connectivity with linear or nonlinear response of the neurons has also been identified as an efficient mechanism for pattern separation [6,12–15]. A common approach to probe this function of the olfactory system is to test whether animals learn to discriminate between two similar odors. In these tasks their correct choice may or may not be associated with a reward. It has been found that certain types of perceptual learning depend on the olfactory bulb [16], require adult neurogenesis [17], and are associated with increased densities of the spines on granule-cell (GC) dendrites that form the synapses with the mitral cells (MCs) [18]. The mechanisms controlling these plasticity mechanisms are not very well understood. The survival of GCs is enhanced by odor enrichment [4,19] and by GC activity [20]. Similarly, odor enrichment increases the stability of spines in bulbar regions that are activated by the enrichment odors [5,21]. Simultaneous measurements of the spine stability and the activity of the cells associated with that synapse are challenging. Thus, knowledge about the response of the bulbar spine dynamics to signals (like glutamate, BDNF) associated with neuronal activity is somewhat limited [22,23]. So far no direct connection between the activity of individual neurons, the stability of their spines, and the circuits emerging from that plasticity has been established experimentally. In fact, the organizational principle of the network of mitral and granule cells is still only poorly understood, reflecting also the lack of significant chemotopical organization of the bulb [24]. The overall dependence of the spine stability on odor enrichment suggests that the resulting connectivity would reflect the specific odors that the animal has experienced. Computational modeling predicts such a structure and suggests that it could underly the observed improvement in odor discrimination [1,2,6].

In this paper we pursue two goals. Motivated by a number of experimental observations [6,18,22,23], we first put forward a phenomenological model for the activity dependence of the spine dynamics that incorporates the memory of previous odor exposures. We then use this model to interpret two sets of experimental results [25,26], focusing particularly on the question to what extent these experiments support the computational prediction that the olfactory bulb exhibits a learned, odor-specific subnetwork structure. In [25] the repeated exposure to an odor over multiple days has been found to lead to a reduction of the mitral cell response that is

specific to that odor, i.e. the response of the mitral cell to a different odor changes only little during the same time. A different set of experiments [26] revealed that repeated odor exposure makes the bulbar mitral cell representations of two odors more discriminable if the two odors are very similar. However, if the two odors are very different from each other, the learning makes their representations actually less discriminable. Moreover, while the repeated odor exposure leads to an overall reduction of the mitral cell responses, individual mitral cells can exhibit an increase in activity and this increase can be odor-specific, i.e. for one of two very similar odors the mitral cell response can increase, while it remains low for the other odor, resembling differential disinhibition. Differential inhibition is also observed.

We show that our model is able to capture these observations. To disentangle the role of the activity-dependence of the spine dynamics we also consider a control model with purely random spine dynamics in which the statistically homogeneous spine density is modulated by a global signal. Somewhat surprisingly, the control model can capture the learning-induced improved discrimination of similar stimuli as well as the learning-induced reduced discrimination of dissimilar stimuli. However, the experimentally observed disinhibition of mitral cells is extremely unlikely in such a model. In combination with [25, 26] our model gives therefore strong support to the notion that the structural plasticity in the olfactory bulb leads to a learned, odor-specific network structure in which MCs preferentially inhibit each other disynaptically, if they have similar receptive fields [6, 12, 14, 15]. Moreover, the model makes testable predictions for the recovery of MC responses after repeated odor exposures, for the suppression of odor discrimination by odor exposure, and for the forgetting of odor memories through interference. An earlier account of some of these results is available at [27].

Results

Formulation of the model

Using computational modeling we investigated the ramifications of the structural plasticity of the reciprocal synapses between mitral cells (MCs) and granule cells (GCs). We focused on those two cell types and did not include the processing in the glomerular layer between the sensory neurons and the mitral cells. Instead, we considered the glomerular activation patterns as inputs to the MCs (Fig 1A). The MCs and GCs were described by firing-rate models. Reflecting the reciprocal nature of the MC-GC synapses the GCs inhibited those MCs that excited them as indicated in the inset of Fig 1A, which depicts a spine on a GC dendrite and an excitatory synapse (MC→GC) and an inhibitory synapse (GC→MC). Since the synapses are located on the secondary dendrites of the MCs, which reach across large portions of the olfactory bulb [28], we allowed synapses to be formed between all MCs and all GCs without any spatial limitations. We recognize that this ignores the sparse structure of the dendrites, which results in a decrease in the connection probability between widely separated MCs [29, 30].

The MC-GC synapses do exhibit some synaptic-weight plasticity. Its nature is, however, still quite unclear [31, 32]. Moreover, during perceptual learning the frequency of the inhibitory postsynaptic currents received by the MCs from the GCs has been found to increase, while their amplitude remained unchanged [18], suggesting that only the number but not the strength of the synapses changed. We therefore assume all connections to have the same weight and focus on the structural plasticity of the

synapses. Thus, in each time step we add and remove reciprocal synapses between MCs and GCs with probabilities that depend on the activity of the cells that are connected by that synapse. 144
145
146

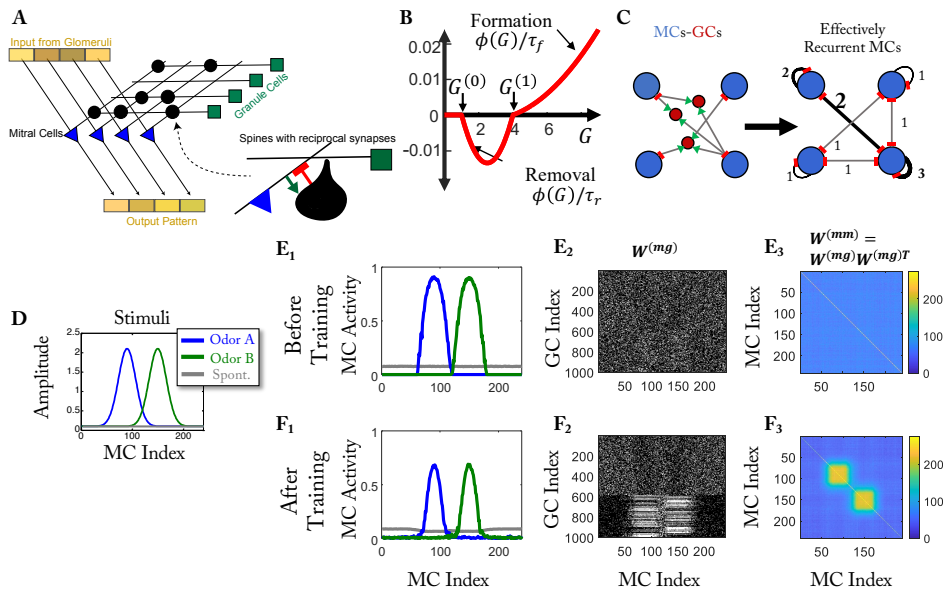


Fig 1. Computational model. (A) Sketch of the network. All spines provide reciprocal synapses that excite GCs (green arrow) and inhibit MCs (red bar). (B) Spine formation is controlled by $R = M\phi(G)$. The figure shows ϕ/τ_f for $\phi > 0$ and ϕ/τ_r for $\phi < 0$ (cf. Eq. 5). (C) Disynaptic recurrent inhibition of MCs via GCs. The numbers in the right panel indicate the strength of the effective inhibition. (D) Simplified training stimuli. (E₁, F₁) MC activity before and after training, respectively. (E₂, F₂) Connectivity between MCs and GCs before and after training, respectively. Each white dot represents a connection between an MC and a GC. (E₃, F₃) Effective recurrent connectivity among MCs. The color represents the number of GCs that mediate the mutual disynaptic inhibition of MCs via GCs (cf. C).

Our modeling is qualitatively motivated by a number of experimental studies. The overall stability of spines has been observed to be enhanced if animals experienced an enriched odor environment [5]. More precisely, across multiple days, spines in bulbar areas that were activated by the enrichment odors were stabilized, while those in other areas were not. Very recently, a positive correlation between the stability of individual spines and the GC-activity of the dendrite on which they are located has been observed [33]. Since filopodia often are precursors of spines, we also glean information from experiments addressing their dynamics. The frequency of formation and removal of filopodia has been found to increase with NMDA-receptor activation [22] and to decrease with Mg^{2+} -concentration. This suggests that the formation rate of spines increases with the activity of the GC-dendrite on which they are forming. The dynamics of filopodia that are located on spine heads are often indicative of subsequent dynamics of that spine. Specifically, the amplitude and direction of the movement of the spine head has been found to be correlated with the lifetime and orientation of such a filopodium, respectively [23]. Moreover, the dynamics of the filopodia have been found to be triggered by presynaptic glutamate release, suggesting a role of MC-activity in the spine dynamics [23]. Biophysical details about the mechanisms relating the spine dynamics with the temporal evolution of the neuronal activity during odor exposure are, however, not well known yet. 147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165

Based on these observations, we assume the formation and removal of synapses to follow Poisson processes with rates that depend on the activities of the neurons connected by the respective synapse (for details see Methods). We express the formation and removal rates in terms of a single rate function $R = M\phi(G)$ that depends on the activities M and G of the MC and the GC, respectively, that are connected by that synapse (Fig 1B). Similar to models for synaptic weight plasticity [34,35], we take this function to be non-monotonic. For GC activities above a threshold $G^{(1)}$ a new synapse is formed, while it is removed below that threshold as long as it is above a second, lower threshold, $G^{(0)}$. For GC activities below $G^{(0)}$ structural change is negligible. This lower threshold allows the synaptic connections of a GC to be remembered during times when the GC is not activated by any odor. Connections are removed, however, on GCs that are excited, but only weakly so. Since the change in the total number of synapses appears to be limited in experiments [18,36], we hypothesize that there exists a homeostatic mechanism that keeps the total number of synapses on a given GC within a limited range. For most of our results, we employ a top- k competition mechanism, in which only the k strongest synapses survive. We show that qualitatively the same results are obtained with a more biophysical mechanism that is based on a limited resource (see Realization of competition through competition for a limited resource and Methods). Both mechanisms operate on a relatively fast time scale, as is generally needed for the stability of networks with Hebbian plasticity [37].

We are particularly interested in the impact of the plasticity on the ability of the network to learn to discriminate between stimuli that the network is repeatedly exposed to. We train the network by using a sequence of alternating stimuli A and B as sensory inputs. We assess the ability of the network to discriminate between these stimuli by comparing the difference in the MC activity patterns in response to the stimuli with the variability of these patterns across repeated presentations of the same stimuli. The firing rates themselves do not fluctuate between successive presentations. However, the spike trains underlying those mean firing rates can fluctuate. Assuming Poisson-like spike trains, we take the mean firing rate of an MC as a proxy for its variance (for details see Methods section). To illustrate the mechanism by which the discriminability is modified, we use simplified stimuli (Fig 1D), while naturalistic stimuli are used to compare with experimental data (Fig 2A).

Plasticity induces mutual disynaptic inhibition of co-activated MCs

To illustrate the network evolution resulting in the structural-plasticity model, we use a pair of simple model stimuli A and B that excite partially overlapping sets of MCs (Fig 1D). In the absence of any odors ('air') the input is taken to be non-zero and homogeneous, reflecting the spontaneous activity of the MCs or of the neurons driving their input. Initially, the MC-GC connectivity is chosen to be random (Fig 1E₂). The learning process reflects the effectively Hebbian character of the structural plasticity rule through which connections are established between cells that fire together. If a GC responds to a stimulus by integrating excitatory inputs from activated MCs, other MCs that are activated but not yet connected to this GC are likely to form a connection with that GC. This is reflected by the two blocks in Fig 1F₂ involving GCs with indices above ~ 600 . Each of these GCs is connected with either the MCs responding to stimulus A (MC indices 60-110) or to stimulus B (MC indices 110-160). Conversely, if two MCs respond to the same stimulus, the number of GCs they both connect to increases, if the model is trained by that stimulus. Due to the reciprocal nature of these synapses, this induces enhanced mutual disynaptic inhibition (Fig 1C) between these

MCs. This is seen in the two blocks along the diagonal line of their effective connectivity matrix, which gives the number of GCs that mediate that inhibition (Fig 1F₃). As a result, the training preferentially reduces the activity of initially highly activated MCs (Fig 1E₁, F₁), consistent with observations [25,38].

Training enhances or reduces stimulus discriminability depending on their similarity

Intriguingly, in mice learning to discriminate between two odors the training enhanced the discriminability of the training stimuli *only* if they were similar, i.e., for hard discrimination tasks [26]. If the task was easy, i.e. if the stimuli were dissimilar, the training actually *reduced* the discriminability of the training stimuli. Here discriminability was measured using d' , which is given by the mean of the difference in MC-activity in response to the two odors relative to the square-root of the pooled variance. To provide insight into this surprising finding we employed the same training protocol as in the experiment, using naturalistic stimuli that were adapted from glomerular activation data of the Leon lab [39] (Fig 2A (top)) and down-sampled to reduce the computational effort (Fig 2A (bottom)). These stimuli were used for the easy discrimination task. As in [26], we used mixtures of the same stimuli for the hard discrimination task (60% : 40% *vs.* 40% : 60%, see Methods). The two odors of the hard task drove each MC to a very similar degree (Fig 2C).

Our model successfully reproduced several experimental observations of [26]. There it was found that training with the easy task *reduced* the difference in the response of the MCs to the two odors used in the task; the majority of the MCs fell below the diagonal in Fig 2D₁. Here the response was defined as the difference between odor-evoked activity and air-evoked activity and the color of the dots in Fig 2D₁, E₁ denotes the average response of the respective MC across the two odors before training. In contrast, the difference *increased* with training in the hard task (Fig 2 E₁). In particular, a fraction of cells showed the same response to both odors before training, but significantly different responses after training (Fig 2E₁, dots along the y-axis). As in [26], we quantified these changes by classifying MCs into responsive cells, i.e., cells that showed a significant response to at least one of the two odors, and divergent cells, for which the two odors evoked significantly different responses (see Methods). In the easy task, not only the number of responsive cells decreased but also the number of divergent cells (Fig 2D₂). In the hard task the number of responsive cells also decreased. The number of divergent cells, however, actually increased by a small amount (Fig 2E₂). Thus, both results recapitulated the experimental observations [26]. These outcomes were not sensitive to the threshold θ that classified the cells (Fig S14).

The discriminability of two neuronal activity patterns depends not only on the difference of their mean firing rates but also on the trial-to-trial variability of their spike trains. The latter is not included in our firing-rate framework. It is quite common for the variability to increase with the firing rate; for instance, for Poisson spike trains the variance in the spike count is equal to the mean spike count, i.e. the variability increases like the square-root of the mean. As a proxy for the discriminability of the response patterns, we therefore defined d'_i of each divergent neuron as the difference of the mean activity in response to the two odors divided by the square-root of the sum of the means (see Eq.(12) in the Methods). As in [26], training increased the mean $\langle d'_i \rangle$ across the MCs in the hard task (Fig 2E₃), but it decreased $\langle d'_i \rangle$ in the easy task (Fig 2D₃), indicating reduced discriminability.

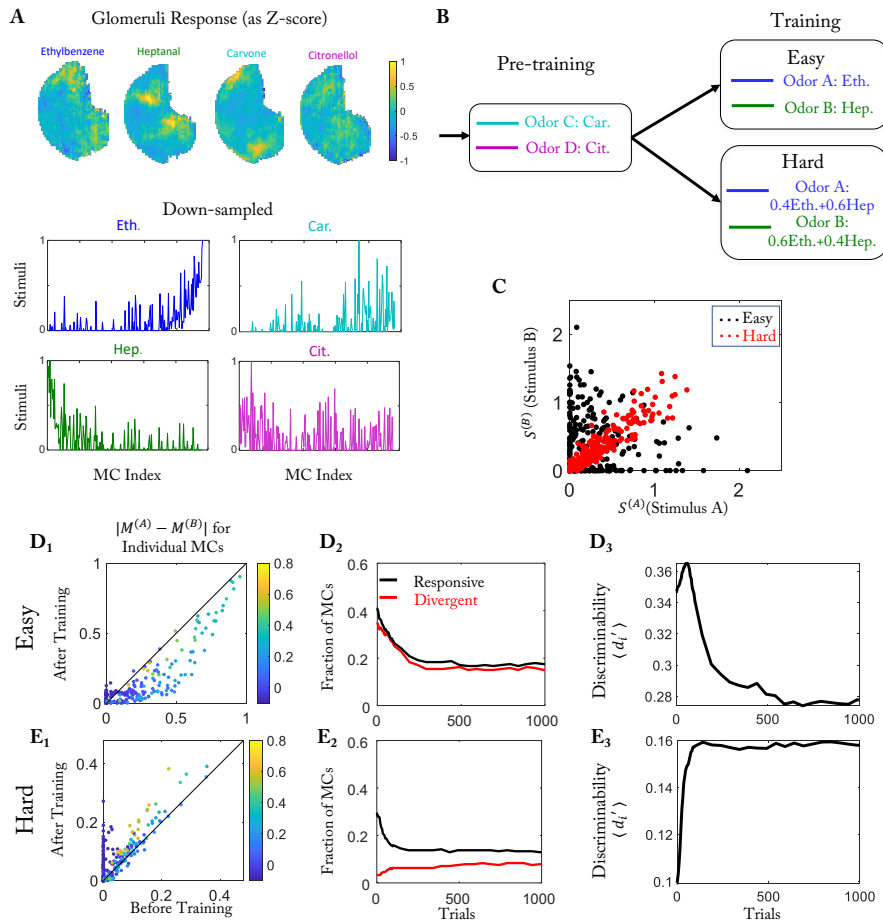


Fig 2. Easy and hard discrimination task using naturalistic stimuli. (A) Odors employed in the training. Top: Glomerular activation patterns (cf. [39]). Bottom: Activation patterns down-sampled to 240 points, serving as stimuli $S^{(i)}$. MCs are sorted based on the difference in activation by the stimuli of odor pair 1. (B) Computational protocol. (C) Activity of each MC for the two stimuli used in the easy (black) and hard (red) task. (D₁,E₁) Difference in the response to the two stimuli for each MC before and after training for the easy and the hard task, respectively. Color of the dots represents the mean response $|M^{(1)} + M^{(2)} - 2M^{(air)}|/2$ before training. (D₂,E₂) Temporal evolution of the number of responsive MCs and divergent MCs during the training in the easy task and hard task, respectively (cf. [26]). (D₃,E₃) Temporal evolution of $\langle d'_i \rangle$ for the easy task and the hard task, respectively. Parameters as in Table 1 except for $\gamma = 1.7 \cdot 10^{-4}$.

Visualization of reduction and enhancement of discriminability

262

To better understand and visualize the mechanism underlying the change in discriminability, we employed simplified stimuli (Fig 3A,B). As in the case of realistic stimuli, the stimuli for the hard task were mixtures of two easily discriminated stimuli. Again, we first exposed the network to a set of pre-training stimuli. They set up an initial connectivity that was independent of the task stimuli (Fig 3A). It affected the initial response of the network to the training stimuli (Fig 3C₁,D₁).

263

264

265

266

267

268

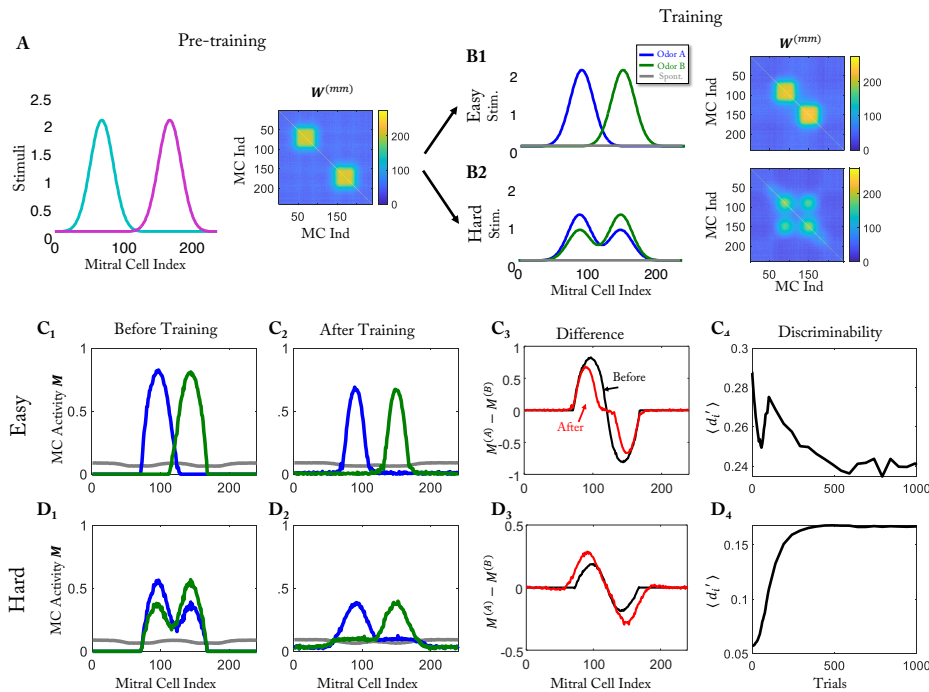


Fig 3. Visualization of the differential learning outcome for easy and hard tasks. (A) Stimuli for pre-training and effective connectivity matrix after pre-training. (B) Training stimuli for the easy task (B₁) and the hard task (B₂) and the resulting connectivity matrix after training. (C) Easy task: (C₁) MC Activity before training (after pre-training). (C₂) MC activity after training. (C₃) Activity difference before (black) and after training (red). (C₄) Temporal evolution of the discriminability during training. (D) As (C) but for hard task.

In the easy task, the effective connectivity resulting from the training had only two blocks, which were along the diagonal and corresponded separately to odors A and B (Fig 3B₁). This inhibition reduced the MC activity and with it the number of responsive cells. Since most MCs responded essentially to only one of the two stimuli, the response difference decreased along with the overall response (Fig 3C₃, red line vs. black line), decreasing the discriminability.

In the hard task, both stimuli activated the same set of MCs, but to a different degree, reflecting the difference in the concentrations of the two components. The discriminability of the stimuli was therefore not compromised if the MC activities were reduced, as long as the difference in the representations of the two stimuli was maintained. In fact, in that case the discriminability in terms of d' was enhanced, since we associate reduced overall activity with reduced trial-to-trial variability (cf. Eq.(12)). The effective connectivity resulting from the training achieved this through the two additional, off-diagonal blocks (Fig 3B₂). Each MC activated by component A received disinaptic inhibition from MCs that were activated by component A as well as from MCs activated by component B. With increasing concentration of component A the effective inhibition from the MCs activated by component A increased, while that from the MCs activated by component B decreased. Thus, the overall inhibition was relatively independent of the concentration of A and therefore about the same for both mixtures, preserving the response difference. In fact, for a large number of MCs the difference in the response even increased with training (Fig 3D₃). Assuming that for a

cell to be classified as divergent, this difference had to be above some threshold, the number of divergent cells decreased in the easy task, while it increased in the hard task (Fig 3C₄,D₄, cf. Fig S14C,D), as we found for the naturalistic stimuli (Fig 2D₂,E₂). Qualitatively similar results were obtained without pre-training phase (Fig S6).

Odor specific adaptation

Repeated odor exposure was seen in [26] to lead to the response of a smaller number of MCs. A change in the response amplitude through repeated exposure has already been found earlier [25,40]. Specifically, it was found that the repeated exposure to an odor over the course of a week substantially reduced the MC response specifically to that ‘familiar’ odor, but not to ‘novel’ odors, which the animal had experienced much less [25]. Such a differential response could be due to a modulation of the overall gain in the olfactory bulb by a signal that indicates the novelty or familiarity of the presented odors. Such a signal could arise some time after odor onset, since the odor has to be recognized as familiar first. A more parsimonious explanation is offered, however, by the activity-dependent spine dynamics of our plasticity model (Fig.4).

To mimic the procedure used in [25], we established a background connectivity by repeatedly exposing the network in the pre-training to two different naturalistic background odors B_1 and B_2 in an alternating fashion (Fig.4A). In the training period we included a third odor F in the odor ensemble as the ‘familiar’ odor and monitored the MC response to that odor as well as to 11 other, ‘novel’ odors $N_{1...11}$. For each MC-odor pair we quantified the change in the network response using a change index. It was defined as the response of the MC to that odor at the end of the training period minus that at the end of the pre-training divided by the sum of those two activities [25]. As in [25], for the majority of MC-odor pairs the change index for the familiar odor was negative (Fig.4B). This reduction in response resulted from an increase in the number of GCs mediating the disynaptic inhibition among the MCs responding to the familiar odor (MC indices mostly between 30 and 150 in Fig.4D).

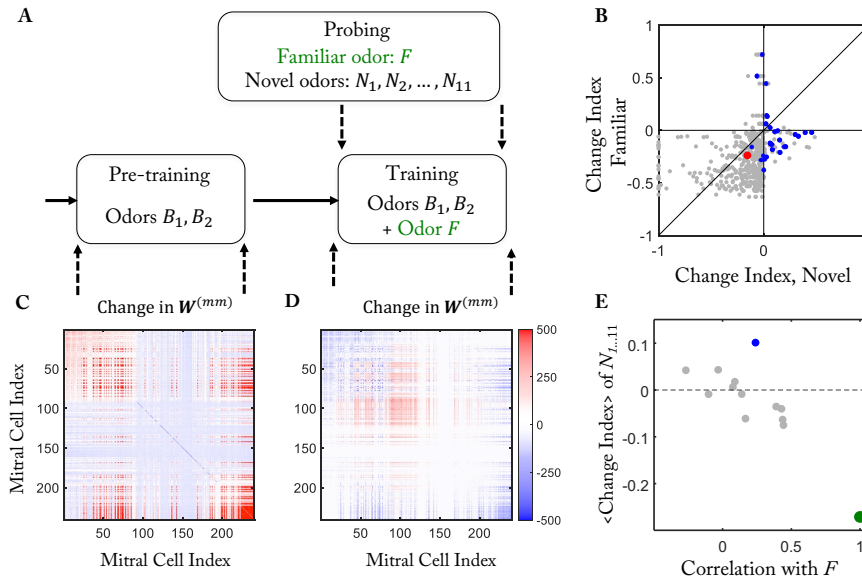


Fig 4. Odor specific adaptation (cf. [25]). (A) Simulation protocol. Pre-training with odors $B_{1,2}$ established a background connectivity. During training the network became familiar with an additional odor F . (B) The MC response to the familiar odor F is more reduced than to the novel odors $N_{1...11}$. (C) During pre-training connections are predominantly added from GCs to MCs that are activated by the background odors $B_{1,2}$ (MC indices below 90 and above 200). The color indicates the change in $\mathbf{W}^{(mm)}$, i.e. in the number of GCs connecting the respective MCs. (D) During training the number of connections of MCs responding to the familiar odor F increases, but that corresponding to odors $B_{1,2}$ decreases (cf. (C)). (E) The mean change index of the novel odors $N_{1...11}$ is correlated with their similarity to the familiar odor.

The network restructuring during the training generally also modified the MC responses to other, novel odors. For MCs that responded to a novel as well as to the familiar odor we compared the respective change indices. Since synapses that are not activated by the familiar odor do not become stabilized by the repeated odor exposure, the response of MCs to the novel odors was typically less reduced than that to the familiar odor, yielding a change index below the diagonal in Fig.4B, in agreement with [25].

The network restructuring during the training did not increase the inhibition for all MCs. For some MCs the inhibition was actually decreased leading to a positive change index (Fig.4B). Averaging across all responding MCs, the mean change index for a novel odor was negatively correlated with how similar it was to the familiar odor as measured in terms of the Pearson correlation of the stimuli (Fig.4E). A particularly strong reduction in the inhibition occurred among many of the MCs responding to novel odor N_4 (marked blue in Fig.4B,D). This odor was actually identical to odor B_2 , which was used in the pre-training, where it established mutual inhibition among MCs with indices in the range 30 to 90 and above 200 (left panel of Fig.4C). During the training period, which included odor F in addition to B_2 , many of these connections were removed (Fig.4D), leading to a strikingly positive mean change index for that odor (blue circle in Fig.4E). Positive change indices for familiar and particularly for novel odors were also observed experimentally for quite a few MC-odor pairs in [25].

Random spine dynamics can capture divergence and discrimination

336

337

So far we have seen that the model proposed here captured a variety of key experimental findings [25,26]. A key prediction emerging from the activity-dependence of the spine dynamics is the formation of subnetworks of highly connected MCs and GCs that specifically respond to one of the training odors. We therefore asked whether the activity dependence is necessary to capture the experimental results and to what extent the experiments therefore allow to infer the formation of an odor-specific, learned network structure in the olfactory bulb. Alternatively, the spine dynamics could be statistically *homogeneous* with a density that is controlled by a global signal reflecting, e.g., a neuromodulator indicating the novelty or familiarity of an odor [41, 42].

338

339

340

341

342

343

344

345

346

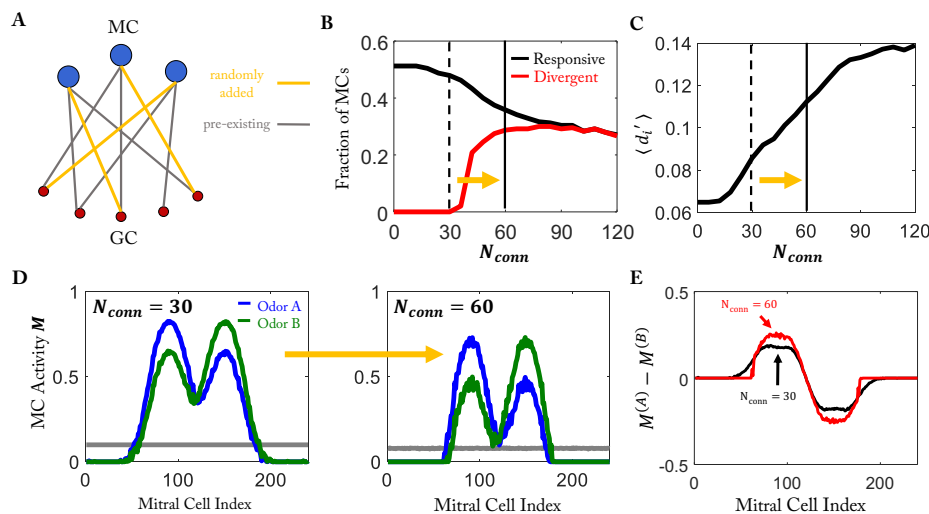


Fig 5. Random network model can capture enhanced differentiability, (A) Random addition of connections during training. (B) Increased connectivity decreases fraction of responsive cells but can increase that of divergent cells (cf. Fig.2H). (C) Discriminability of the odors in the hard task increases with increasing N_{conn} . (D,E) The difference in MC response does not decrease with increasing N_{conn} (cf. Fig.3I).

As a control model, we implemented purely random, activity-independent spine dynamics, where the training with the familiar odor lead to a change in the mean number of MC-GC connections. Since the experiments [25] and the simulations (4) revealed a reduction in the MC-activity in response to the familiar odor, we increased during the training the mean number of connections between MCs and GCs, which is also consistent with an increase in the number of spines during perceptual learning [18].

347

348

349

350

351

352

For the hard task using simplified stimuli (cf. Fig.3) the fraction of responsive MCs decreased when the number N_{conn} of connections was increased, while the fraction of divergent MCs increased substantially over some range of N_{conn} (Fig.5B, cf. Fig.2). In parallel, the discriminability measured in terms of d' increased (Fig.5C). This reflects the fact that increasing the random connectivity increased the inhibition for *all* MCs by roughly the same amount. In a linear firing-rate model this would preserve the difference between the MC responses; with the saturation of the MC response used in our model, the difference actually increased (Fig.5D,E).

353

354

355

356

357

358

359

360

Thus, the aspects considered in Fig.5 are not able to distinguish between the activity-dependent and the random model. We therefore considered the results of the

361

362

protocol of Fig.4 in more detail. 363

Random spine dynamics fail to induce observed disinhibition 364

As in Fig.4, we exposed the random model in the pre-training to 2 background odors, 365
added odor F as the ‘familiar’ odor to that ensemble during the training period, and 366
probed the response to additional 11 ‘novel’ odors. As expected, and in contrast to the 367
activity-dependent model, the random model showed no qualitative difference in the 368
change index distribution of the familiar and the novel odors (Fig.6D). Strikingly, even 369
though in the random model the mean change index for the familiar odor was less 370
negative than that in the activity-dependent model ($\langle CI \rangle_{random} = -0.14$ vs 371
 $\langle CI \rangle = -0.27$), for none of the MC-odor pairs the change index was positive in the 372
random model. In contrast, in the activity-dependent model about 8% of the MCs had 373
a positive change index for the familiar odor (Fig.6E), which is in rough agreement with 374
the experimental findings [25]. These MCs were disinhibited by the training (Fig.6A). 375
For visualization purposes we also compared the changes in MC activity for the two 376
models using simplified stimuli. Again, while in the random model the training led to a 377
decrease in the activity of all MCs, in the activity-dependent model there were quite a 378
few moderately activated MCs that were disinhibited by the training (Fig.6C). 379

A more detailed analysis of the results of the activity-dependent model showed that the 380
disinhibition in the response of a given MC to two odors can be quite different even if 381
the two odors are very similar to each other. Considering again the mixtures used in the 382
hard task of Fig.2, we found a large number of MCs for which the difference in the 383
response to the two stimuli was significantly increased by the network restructuring. For 384
some of them the enhancement arose from differential disinhibition, wherein the MCs 385
were inhibited by both odors before training, but only by one of them after training 386
(Fig.6A). For other MCs the increased difference was due to differential inhibition. 387
While these MCs were excited by both odors before training, they responded only to 388
one of the very similar odors after training (Fig.6B). Both, differential disinhibition and 389
inhibition were observed experimentally in [26] (their Supp. Fig.S1). 390

In contrast, in the random model no disinhibition was obtained. This is illustrated for 391
one odor in Fig.6C. It also manifested itself in the change index for novel and familiar 392
odors Fig.6D (cf. Fig.4B) and the cumulative distribution of the change indices (Fig.6E) 393

Thus, the improved discrimination of pairs of very similar stimuli and the reduced 394
discrimination of dissimilar stimuli does not require the emerging connectivity to be 395
odor-specific, but only an overall increase in inhibition. The observed disinhibition, 396
however, is not consistent with random spine dynamics. But it is recapitulated with 397
activity-dependent spine stability, which leads to odor-specific network structure. Thus, 398
combined with the experimental observations of [26] our computational model supports 399
the notion that odor exposure during perceptual learning leads to a network structure 400
that reflects the presented odor. 401

Removal of spines and forgetting by intermediate-amplitude 402 re-activation 403

A key feature of our model is the non-monotonic behavior of the resilience function ϕ 404
(Fig.1B), which renders spines stable not only if the associated GC is strongly active, 405
but also when it is inactive. This allows previously learned connectivities to persist even 406
if the corresponding odors are not presented any more. If MCs and GCs are weakly 407

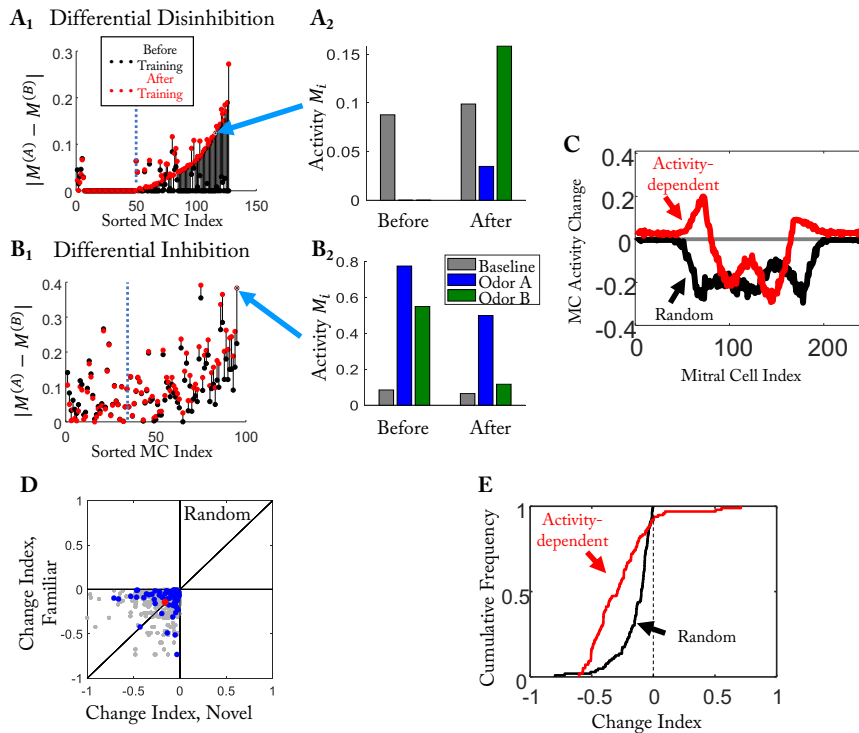


Fig 6. Disinhibition in the activity-dependent but not in the random model. (A) In the activity-dependent model training in the hard task with naturalistic stimuli (cf. Fig.2) disinhibited a large fraction of MCs differentially. The activity of a specific MC is shown in (A₂). (B) In the activity-dependent model training also inhibited many MCs differentially. (C) Training-induced change in MC activities for the random and the activity-dependent model for the simplified stimuli. (D) In the random model the change index for familiar and novel naturalistic odors was statistically the same (cf. Fig.4B); Odor N_4 is marked in blue. (E) The CDF of the change index shows disinhibition only for the activity-dependent, but not the random model.

driven, however, the corresponding spines are removed. While the non-monotonicity of the resilience function is reminiscent of models for synaptic weight plasticity [34,35], there is no explicit experimental evidence for it so far for the bulbar structural plasticity.

To illustrate the impact of the non-monotonicity of the resilience function we used simplified stimuli. During pre-training two dissimilar odors A,B were presented, which resulted in a connectivity comprised of two blocks on the diagonal (Fig.7A). In a subsequent training period we tested the persistence of the memory of the block corresponding to odor A when the network was re-exposed to that odor, but with a different amplitude (labeled $A_{1...4}$ in Fig.7). When that amplitude was small, the connectivity remained essentially unchanged by the training (Fig.7B₁). For intermediate activities, however, a fraction of the connections established in the pre-training were removed since the GC activity fell into the range in which the resilience function was negative (Fig.7B_{2,3}). For large amplitudes the resulting GC activity was large enough to render the resilience function positive and the connections were stable (Fig.7B₄).

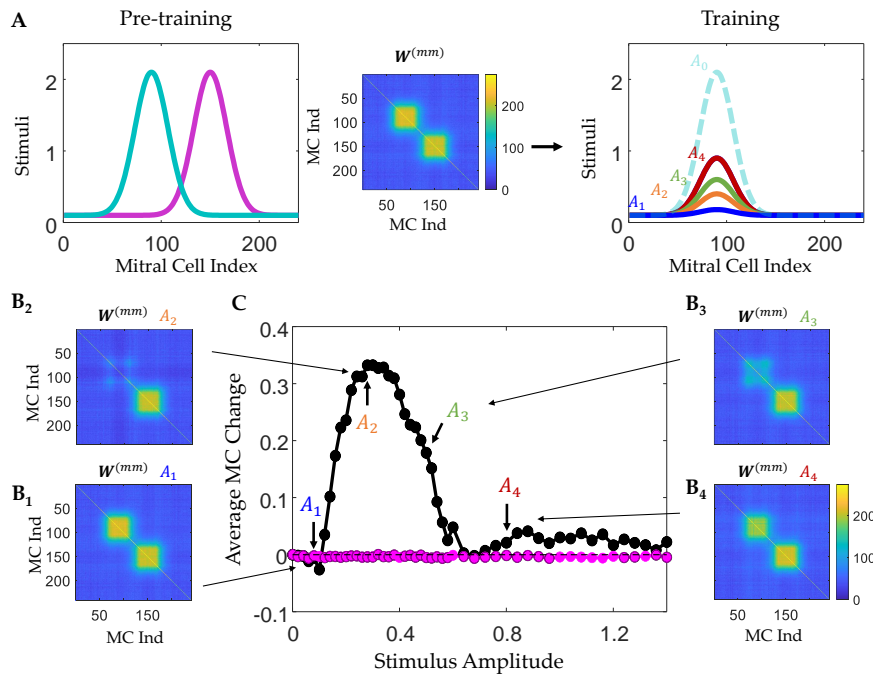


Fig 7. Removal of connections for intermediate GC activation. A) Pre-training with two odors, then training with only one of the two odors at different amplitudes $A_1 \dots A_4$. B) The connectivities $W^{(mm)}$ resulting from the training depend non-monotonically on the stimulus amplitude $A_1 \dots A_4$ during training. C) Non-monotonic mean change in the response to odor A after training as a function of the amplitudes $A_1 \dots A_4$ used in the training.

Experimentally the connectivity itself is only poorly accessible. One can, however, probe the predicted changes using the response of the MCs to odor A , similar to the experiments in [25]. When the connections are retained during the training, the response of the MCs to that odor will not change compared to their response before the training. However, if intermediate stimulus amplitudes are used in the training, the removal of the connections manifests itself in an increase of the mean response of the MCs to odor A after the training (Fig.7C). Correspondingly, the distribution of the differences between the MC activities after and before the training shifts and broadens in a non-monotonic fashion when the amplitude of the stimulus used in the training is increased (Fig.S2).

Thus, if the odor exposure experiments in [25] were to be continued with the same odor but at a suitably reduced concentration, the model predicts that the MC response to that odor would increase again within days rather than undergoing the slow recovery over a period of 1-2 months that was observed in the absence of that odor [25].

The protocol used in Fig.7 does not have a simple behavioral read-out. Moreover, it relies on a change in the overall activity of the MCs. Since the input to the MCs is pre-processed in the glomerular layer, which is thought to provide an overall gain control of the input [43], changes in the odor concentrations may not translate directly into corresponding changes in MC activation and the ensuing change in the connectivity predicted by the model. We therefore considered also a slightly modified protocol using mixtures composed of two dissimilar odor components. To reduce the impact of the overall gain control, only the concentration of one of the two components was varied

(Fig.8A). Exposing in pre-training 2 the network to two mixtures in which that component had similar but large values (Fig.8A₂) enhanced the discrimination between these two mixtures (teal-magenta line in Fig.8G), However, due to pre-training 2 the network developed strong mutual inhibition between the MCs responding to the two mixture components (Fig.8E). This allowed the strong component in the training stimulus (Fig.8A₃) to suppress the response to the weak component completely (Fig.8B). Subsequent training with the mixtures in which that component was weak, removed many of those connections (Fig.8A₃,F) and largely reversed that suppression (Fig.8C). This enhanced the discriminability of those mixtures above its initial value without any inhibition (blue-green line in Fig.8G).

444
445
446
447
448
449
450
451
452
453

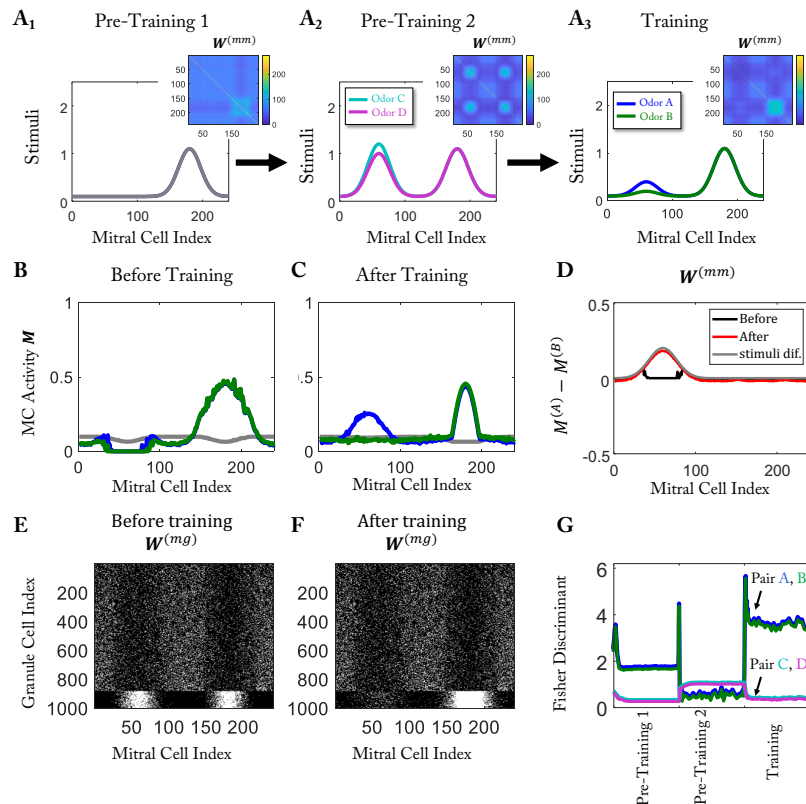


Fig 8. The discrimination of weak components of similar stimuli is abolished by training with strong components, but recovers through training with weak components. (A) Protocol: A₁: pre-training 1 with only 1 odor, A₂: pre-training 2 based on a mixture of that and a dissimilar odor, A₃: training with a mixture in which the second component is only weak. (B) Response to the training stimulus before training; due to pre-training 2 the response to the weaker odor component is suppressed completely. (C) Training with the weaker component recovers that component. (D) MC response difference between odors A and B before and after training. (E,F) MC-GC connectivity before and after training. (G) Fisher discriminant for odor pair (A,B) and (C,D).

The removal of spines on GCs that are only intermediately activated lead also to forgetting of a discrimination task by learning a subsequent discrimination task, if the odors of the second task overlapped to some extent with those of the first task and therefore interfered with the retention of the relevant connections (Fig.S3). Upon re-exposure to the odors of the first task, the network very quickly re-learned to discriminate between those odors. In fact, the relevant connectivity was re-established

454
455
456
457
458
459

faster than it had been learned in first odor exposure (Fig.S3G). In general re-learning occurred noticeably faster than the forgetting (Fig.S4).

Realization of competition through competition for a limited resource

In order to keep a Hebbian model stable, different compensatory processes can be imposed depending on the state of nearby synapses on the same dendrite [37]. In the model that we have discussed so far we used top- k competition because of its simplicity, without identifying a specific, biologically feasible mechanism. In this section, we discuss one such mechanism.

The idea behind the competition is that within a single neuron, possibly within a single dendritic branch, resources used for the formation of synapses are limited. In hippocampal CA1 neurons, stimulating spines leads to the shrinking of neighboring unstimulated spines [44]. A plausible explanation of this phenomenon is that some resource is limited within the neuron. A computational model based on this idea has reproduced the fast multiplicative scaling behavior of synapses and explained a transient form of heterosynaptic plasticity [45]. Along similar lines, we assumed that the formation of each synapse depletes a resource pool by a fixed amount, and the formation rate is proportional to the fill-level of the resource pool. Further, we assumed that the resource for a synapse is returned to the pool when the synapse is removed and that the removal rate is independent of the resource pool size. We took the activation function (Fig 9A, red curve) to be comprised of a formation term, which depends on the fill-level of the resource pool (Fig 9A, blue curve), and a removal term, which is independent of the fill-level (Fig 9A, green curve).

It is easy to see how the limited resource pool provides a homeostatic mechanism. For low pool levels synapses are removed (open arrow in (Fig 9A)), which replenishes the pool, while for high pool levels new synapses are formed (solid arrow in Fig 9A)), decreasing the pool level (Fig 9B). The equilibrium pool level and with it the number of synapses in the steady state is regulated by the activity of the GCs, which in turn depends on the number of synapses *via* the activity of the MCs connected to the GCs.

With this biophysically plausible stabilization of the structural plasticity, our model yielded results that were qualitatively similar to those obtained with the simpler, top- k competition model (compare Fig 9C-F with Fig 2D,E and Fig S5 with Fig 3).

Discussion

In the rodent olfactory bulb structural plasticity in the form of the extensive formation and removal of synaptic spines persists into adulthood [6] and may be a key factor in establishing and maintaining the network structure of the olfactory bulb. Since it is operating on timescales that are comparable to some learning processes [6, 18] it may play an important role in these processes. In this study, we proposed a simple Hebbian-type model to investigate this role. Our model shows that a local unsupervised learning rule according to which the formation and removal of spines depends on GC and MC activity is consistent with and explains a host of experimental observations [18, 25, 26]. Among them is the intriguing observation that training can have opposite impact on the discriminability of very similar and of dissimilar odor pairs. As observed experimentally, in our model training enhanced the discriminability of similar odors, but it reduced the discriminability of the training odors if they were quite

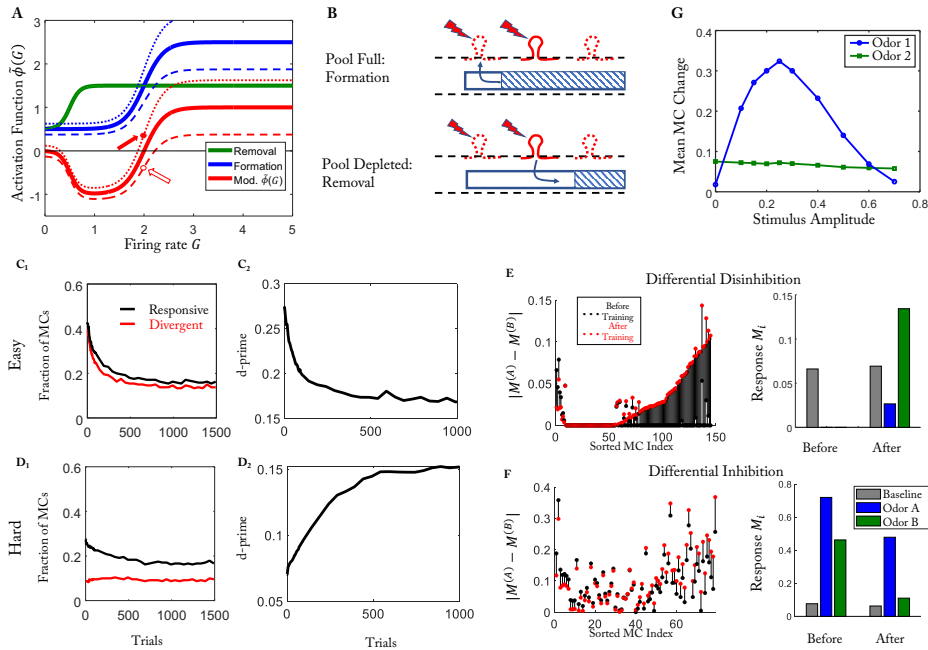


Fig 9. Synaptic competition via a common resource pool. The training protocol is the same protocol as in Fig 2. (A) The modified activation function $\tilde{\phi}(G)$ (red) is the difference between the formation function (blue) and the removal function (green). $P = P_0$ (thick), $P > P_0$ (dotted), $P < P_0$ (dashed). (B) (Top) For $P > P_0$ synapses tend to form (solid arrow in (A)). (Bottom) For $P < P_0$ synapses tend to be removed even though the GC activity is the same (open arrow in (A)). (C_{1,2}) Responsive and divergent cells and discriminability for the easy task (cf. Fig.2E,F). (D_{1,2}) Responsive and divergent cells and discriminability for the hard task (cf. Fig.2G,H). (E,F) Differential disinhibition and inhibition (cf. Fig.6). (G) Removal of spines for training with intermediate activities leads to an increase in MC activity (cf. Fig.7).

different from each other. In fact, for the dissimilar odors, the number of MCs that responded differently to the odors decreased with learning, deteriorating the performance.

Since a key prediction of our activity-dependent model is that upon repeated exposure with relevant odors the olfactory bulb develops subnetworks associated with these odors, we investigated whether the experimental results are also consistent with a model in which learning modifies the inhibition in a random, statistically homogeneous fashion, independent of the activity, and therefore without forming odor-specific subnetworks. Surprisingly, the random model also captured the differential outcome of training with similar and with dissimilar odor pairs. It failed, however, to recapitulate the observation that during that learning process, which led to an overall increase in the number of inhibitory connections, a sizable fraction of the MCs became more, rather than less active, i.e. they were disinhibited rather than inhibited [26]. Similarly, the activity-dependent, but not the random model, captured the odor-specific reduction in MC activity that is observed after an animal has been familiarized with an odor, which is much more pronounced for that familiar odor than for less experienced odors [25]. One could imagine that the reduced response to a familiar odor is the result of an inhibition-enhancing global modulatory signal that may be triggered by the recognition of the odor as familiar. This would, however, not explain the significant fraction of MCs

505
506
507

508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523

that are disinhibited in response to the familiar or the novel odor [25]. Thus, in combination with [25,26] our model provides strong indirect evidence that perceptual learning leads to the formation of MC-GC subnetworks comprised of MCs and GCs having significantly overlapping receptive fields.

Moreover, the model offers a foundation for further studies of learning-related functions in the olfactory bulb. For instance, in the current version of the model the novelty or salience of an odor is not considered explicitly. Only implicitly it is assumed that the presented odors have sufficient salience, since experiments suggest that the plastic changes occur predominantly for novel or otherwise salient odors [19,46].

In contrast to a previous model for structural spine plasticity [6], the model presented here describes a learning mechanism that allows the network to remember previously learned tasks in its connectivity. These memories are robust against subsequent learning of dissimilar tasks. The learned connectivity is only erased by exposure to interfering stimuli, which overlap significantly with previously learned ones. This prediction should be amenable to experimental tests using sequential perceptual learning [47] with two odor pairs, each comprised of two similar odors. Repeated exposure with one of the two odors of the first pair will enhance the spontaneous discriminability of the odors in that pair. The subsequent repeated exposure with an odor from the second pair is expected not to affect the discriminability of the odors in the first pair, if the odors in the second pair are dissimilar from the first pair. However, if there is a significant overlap between the pairs, the second training is expected to compromise the discriminability of the first pair. Note, however, that if the odors in the second pair are very similar to those in the first pair, this interference is not expected. The latter aspect would differentiate this forgetting from the widely observed retrieval-introduced forgetting [48].

When assessing the role of structural spine plasticity in forgetting, attention has to be paid to the time scales of the experiments. In recent experiments the successive perceptual learning and forgetting of different odor pairs has been investigated [4]. There it was found that learning a second pair of odors seven days after the training with the first pair strongly interfered with the memory of the first pair, while learning after 17 days did not. The experiments indicated that this interference arose from the enhanced apoptosis of adult-born GCs that was triggered if the time between the training periods was only seven days. Since the results of [25] suggest that significant changes in the connectivity arise already within the first day of exposure, performing the interference experiments on that shorter time scale may avoid complications arising from the apoptosis found in [4] and may focus on the role of the structural spine plasticity described by our model.

We discussed two fast homeostatic mechanisms that stabilize the network activity despite the Hebbian-type learning rule: an abstract top- k competition and a biologically plausible competition of synapses for a limited resource that operates via a sliding threshold for the formation of synapses, somewhat reminiscent of the BCM-model [34]. The overall behavior of both models did not depend much on which homeostatic mechanism was at work. It is worth noting that the specifics of the stabilization mechanism may vary across brain areas. Moreover, resource-dependent competition may be too local to lead to competition across different dendrites or the whole neuron. In addition, heterosynaptic plasticity may be controlled more by activity rather than a limited resource [44]. For these reasons, we leave open the question of how to implement the stabilization mechanism.

In principle, activity-dependent network restructuring could also arise from synaptic weight plasticity rather than structural plasticity. Recent experiments show, however, that in adult-born GCs the amplitudes of evoked IPSCs were not significantly changed

after learning [18]. This suggests that changes in the inhibitory synaptic weights are not likely to be the dominant mechanism for learning. At the same time, in the same experiment, the frequency of the evoked IPSCs was significantly up-regulated after learning, which highly suggests an increase in the number of synaptic spines [18].

What aspects of olfaction might favor structural plasticity even in adult animals, despite its seemingly high metabolic cost? A characteristic feature of olfaction is the high dimensionality of odor space. As a consequence, the layout of the olfactory bulb does not allow substantial chemotopy and MC dendrites have to reach across large portions of the bulb [28,49], where they establish only a very sparse connectivity with GCs. To wit, each of the about 5 million GCs connects with on the order of 50-100 of the 50,000 MCs [50]. Thus, only less than 1% of the possible connections are actually made. The sparseness of the connectivity has recently been identified as a key factor favoring structural plasticity over synaptic-weight plasticity in the context of feedforward networks [51]. This may also be the reason why structural plasticity seems to be essential in the adult motor system, as has been found recently [52].

A characteristic feature resulting from the structural plasticity implemented in our model is the emergence of MC-GC subnetworks that are specific to the learned stimuli. The interaction between the excitatory MCs and the inhibitory GCs is known to drive the γ -rhythms that are prominent in the olfactory bulb [53]. Since their power is enhanced with task demand [54] and since the optogenetic excitation of GCs in the γ -range enhances discrimination learning [55], they are presumed to play a role in odor processing. The different subnetworks that are predicted by our model to emerge from learning different odors are likely to generate each its own γ -rhythm with its own frequency. The question then arises whether these rhythms will synchronize if the bulbar network is driven by the corresponding odor mixtures. Depending on the temporal window over which the bulbar output is read out in downstream brain areas [56,57], the perception of such mixtures may then well vary with the degree of synchrony of the different rhythms. Interestingly, the synchrony of such interacting rhythms can be enhanced by uncorrelated noise and by neuronal heterogeneity *via* a paradoxical phase response [58,59].

A prominent feature of the olfactory bulb is its extensive centrifugal input *via* top-down projections that target particularly the GCs. Our model for the structural spine plasticity suggests that the GC develop receptive fields for specific odors of the training set and are preferentially connected with MCs with similar receptive fields. Due to the reciprocal nature of the MC-GC synapses, this implies that the activation of such a GC inhibits specifically MCs with the same response profile. If the top-down connectivity was able to target specific GCs, top-down inputs could modulate bulbar processing in a controlled fashion by inhibiting specific MCs. Computational modeling [15] suggests that the extensive adult neurogenesis of GCs observed in the bulb naturally leads to a network structure that allows just that. That model predicts that non-olfactory contexts that have been associated with a familiar odor allow to suppress that odor, enhancing the detection and discrimination of novel odors. The spine plasticity investigated here leads to a network structure that is similar to the bulbar component of the network obtained in the bulbar-cortical neurogenesis model. Structural spine plasticity is therefore expected to support the specific control of bulbar processing by higher brain areas.

Methods

620

The computational model comprises MCs and GCs (Fig 1A) that are described by the firing-rate models

621

622

$$\tau_M \frac{d\mathbf{M}}{dt} = -\mathbf{M} + F_M (\mathbf{S} - \gamma \mathbf{W}^{(mg)} \mathbf{G}), \quad (1)$$

$$\tau_G \frac{d\mathbf{G}}{dt} = -\mathbf{G} + F_G \left((\mathbf{W}^{(mg)})^T \mathbf{M} - g_{thr} \right). \quad (2)$$

Here the vectors \mathbf{M} and \mathbf{G} represent the firing rates of N_{MC} MCs and N_{GC} GCs, respectively. The MCs receive sensory inputs \mathbf{S} from the glomerular layer and inhibitory input from GCs, with γ denoting the inhibitory strength. The GCs receive excitation from MCs, with $g_{thr} > 0$ setting a non-zero firing threshold for the GCs. Reflecting the reciprocal nature of the MC-GC synapses the connectivity matrix $\mathbf{W}^{(gm)}$ is the transpose of $\mathbf{W}^{(mg)}$. Focusing on the structural plasticity of the synapses, we do not include synaptic weight plasticity, the characterization of which is still somewhat incomplete [31,32]. We therefore keep the weight γ fixed and set $\mathbf{W}^{(mg)}$ either to 1 or 0. Initially, each row of $\mathbf{W}^{(mg)}$ has N_{conn} non-zero entries, reflecting that each granule cell connects to N_{conn} mitral cells.

623

624

625

626

627

628

629

630

631

632

Reflecting the large spread of the secondary dendrites of MCs [28], we allow synapses to be formed between all MCs and all GCs without any spatial limitations. Thus, the ordering of the MCs and the GCs is arbitrary.

633

634

635

The firing rates depend on the neurons' inputs via the activation functions F_M and F_G given in (3,4), which are taken to be sigmoidal and piecewise linear, respectively. MCs within the olfactory bulb experience saturation behavior. In [60] the authors recorded responses of MCs when the concentration of the odorant changed 10-fold. Among the recorded cells, 38% of MCs responded linearly to stimuli, and 29% of the MCs experienced saturation behavior. This saturation may be the result of saturation in the inputs to the MCs or in the MCs themselves. The activity of the sensory neurons saturates on a log scale when the odor concentration changes 1000-fold [61]. For the concentration changes (0.6 : 0.4 mixture or pure odorant) relevant in our studies, we therefore assumed that saturation occurred mostly in the MC rather than their sensory inputs and took the activation functions in (1,2) to be

636

637

638

639

640

641

642

643

644

645

646

$$F_M(x) = [\tanh(x)]_+, \quad (3)$$

$$F_G(x) = [(x)]_+. \quad (4)$$

Based on the observations in [5, 22, 23, 33], we assume the rate of formation and removal of a synapse connecting MC i and GC j to depend on the activities of those neurons. As a proxy for those activities we take the steady-state solutions of Eqs.(1,2) for given stimulus \mathbf{S} . We express the formation and removal rates in terms of a single rate function,

647

648

649

650

651

$$R_{ij}(M_i, G_j) = M_i \phi(G_j), \quad (5)$$

where M_i and G_j are the respective steady-state activities. The activation function $\phi(G_j)$,

652

653

$$\phi(G_j) = [G_j - G^{(0)}]_+ (G_j - G^{(1)}), \quad (6)$$

changes sign at a threshold $G^{(1)}$, which controls whether a new synapse is formed or an existing synapse is removed (Fig 1B), giving the structural plasticity a bidirectional

654

655

dependence on the GC activity. In addition, there is a second threshold, $G^{(0)}$, below which the structural change is negligible. Specifically, in each trial of duration Δt , which is one step in our simulation, a new synapse is formed with probability P_{ij}^+ ,

$$P_{ij}^+ = 1 - \exp\left(-\lambda_f [R_{ij}]_+ \Delta t\right), \quad \text{if } W_{ij}^{(mg)} = 0. \quad (7)$$

Conversely, an existing synapse is removed with probability P_{ij}^- ,

$$P_{ij}^- = 1 - \exp\left(-\lambda_r [-R_{ij}]_+ \Delta t\right), \quad \text{if } W_{ij}^{(mg)} = 1. \quad (8)$$

Here, λ_f and λ_r are the formation and removal rates, respectively.

Since the change in the total number of synapses appears to be limited in experiments [36, 62], we introduce a homeostatic mechanism that limits the maximal number of synapses on a given GC. For most of our results, we employ a top- k ($k > N_{conn}$) competition mechanism, in which only the k synapses with the largest values of R_{ij} survive if a GC has more than k synapses. A soft top- k competition can be realized by a resource-pool competition with additional assumptions, as discussed in the section Realization of competition through competition for a limited resource. However, other mechanisms like scaling [63] or the ABS (Artola, Bröcher and Singer) rule [64] may alternatively restrict the total number of synapses. Since the stability mechanism was not our focus, we mostly used the (hard) top- k competition without directly specifying its biological cause.

Summarizing the algorithm, in each step the steady-state firing rates \mathbf{M} and \mathbf{G} in response to a randomly chosen odor \mathbf{S} from the training set are computed. Based on the resulting formation/removal rates R_{ij} the connectivity matrix $\mathbf{W}^{(mg)}$ is updated in two parts. First, we implement homeostasis by top- k competition: if any GC has more than k synapses, the corresponding number with the smallest values of R_{ij} are removed. Then, in the learning step the connectivity is updated based on the Hebbian learning rules (7, 8), except that synapses that were just removed in the homeostasis step are not recreated in this time step. We simulate the model long enough for the connectivity to reach a statistically steady state as measured by the discriminability of the training odors (see Sec.Characterizing Discriminability). As initial condition we take a connectivity in which each GC is randomly connected to N_{conn} MCs. Unless specified otherwise, the model parameters are as listed in Table 1. The steady states are obtained by setting $\tau_G = 0$, solving Eq.(1) using ODE45 in Matlab, with \mathbf{G} solved from Eq.(2).

We found that the overall behavior of the model is robust with respect to changes in its parameters (Supplementary Materials).

Instead of the hard top- k competition the number of synapses can also be limited by a resource-pool competition. To this end we replace the activation function $\phi(G)$ in (5) by a function $\tilde{\phi}(G, P)$ that also depends on the fill-level P of that resource pool (Fig 9A, red curve) and assume that it is comprised of a formation term $\tilde{\phi}_{form}$ (Fig 9A, blue curve) and a removal term $\tilde{\phi}_{rem}$ (Fig 9A, green curve),

$$\tilde{\phi}(\mathbf{G}, P) = \tilde{\phi}_{form}(\mathbf{G}, P) - \tilde{\phi}_{remove}(\mathbf{G}), \quad (9)$$

which are, respectively, given by

$$\tilde{\phi}_{form}(\mathbf{G}, P) = (\tanh(\kappa_{form}(\mathbf{G} - r_{form})) + 1 + R_0) \frac{P}{P_0}, \quad (10)$$

$$\tilde{\phi}_{rem}(\mathbf{G}) = (\tanh(\kappa_{rem}(\mathbf{G} - r_{rem})) + 1) / 2 + R_0. \quad (11)$$

The fill-level P of the resource pool depends on the current number of synapses n , $P = P^{(all)} - n$, where $P^{(all)}$ is the total amount of the resource in the cell, which was assumed to have the same constant value for all GCs. Here P_0 is the equilibrium size of the resource pool for which no spines are formed or removed when the GC is not active. Thus, for low GC-activity P eventually goes to P_0 and the number of the synapses on that GC goes to the initial value $n = N_{conn} = P^{(all)} - P_0$.

In analogy to the original model, we define the two thresholds $\tilde{G}^{(0)} < \tilde{G}^{(1)}$ via $\tilde{\phi}(G^{(0,1)}, P) = 0$. The formation of synapses on significantly active GC is controlled by $\tilde{G}^{(1)}$. As more synapses are formed and the fill-level decreases, $\tilde{G}^{(1)}$ increases. This makes it more difficult for synapses to form on that GC, leading to a saturation of the number of its synapses. This dependence of the threshold $\tilde{G}^{(1)}$ is reminiscent of the sliding threshold of the BCM model for synaptic-weight plasticity [34]. However, the sliding is caused here by the changes in the limited resource instead of a temporal average over the previous activities. The change in $\tilde{G}^{(0)}$ is opposite to the change in $\tilde{G}^{(1)}$ when the size of the resource pool changes.

In [26], mice were exposed to the test odors only when they were on the training stage. When they were back in the cage, the training odors were absent, and the activities of GCs corresponding to those odors were likely to be low. In the original model, the activation function $\tilde{\phi}(G)$ vanishes when $G < G^{(0)}$, which means synapses are neither formed nor removed. Thus, the network is static. However, in the resource-pool model the activation function $\tilde{\phi}(G)$ is in general non-zero when $G < G^{(0)}$. As a result, periods of low input also lead to changes in the connectivity. We include such ‘air trials’, during which only spontaneous activity is presented, after every four training trials. During air trials, the GC activity is very low and synapses are formed and removed based particularly on the size of the respective resource pool, pushing each pool towards P_0 . Given the uniform input to the MCs during air trials and the dependence of spine formation on MC activity (cf. (5)), the formation is biased towards MCs that receive less inhibition, i.e., that have fewer connections with GCs, while the removal is biased towards MCs that receive more inhibition. We find that inserting such air trials increases the dynamics of the synapses and speeds up the convergence of the network evolution. It may compensate for the fact that in the resource-pool model an existing synapse is removed only if the activation function is negative, $\tilde{\phi}(G, P) < 0$, while in the top- k competition model, synapses can be removed not only when $\phi(G) < 0$, but also directly by that competition.

The top- k mechanism and the competition *via* a resource pool lead to similar results. We expect that this would also be the case for other mechanisms that may restrict the total number of synapses, like scaling [63] or the ABS (Artola, Bröcher and Singer) rule [64]. Since the stability mechanism was not our focus, we mostly used the (hard) top- k competition without directly specifying its biological cause.

N_{MC}	240	N_{GC}	1000
N_{conn}	60	k	66
τ_M	1	τ_G	0
Δt	1		
g_{thr}	4.4	γ	$5e - 4$
$G^{(0)}$	1	$G^{(1)}$	4
λ_f	$6 \cdot 10^{-4}$	λ_r	$6 \cdot 10^{-3}$

Table 1. Parameters used in the simulations of the model (unless stated otherwise).

The additional parameters of the pool-competition model are listed in Table 2.

κ_{form}	2.5	κ_{rem}	5
r_{form}	2	r_{rem}	1
R_0	0.8	P_0	20
Δt	1	γ	$4.2 \cdot 10^{-4}$
λ_f	1	λ_r	1

Table 2. Table of parameters for the resource-pool model (unless stated otherwise).

Generation of simplified stimuli and naturalistic stimuli

The naturalistic stimuli are based on glomerular activation data [39], which contains 2-d imaging z-scores of the glomeruli responses for a variety of odorants. In our simulations, we used carvone, citronellol (pre-training), ethylbenzene, heptanal (training, pure, and mixture). For the novel odors employed in Figs.4,6 we used in addition limonene, ethylvalerate, 2-heptanone, acetophenone, valeric acid, isoamylacetate, isoeugenol, 1-pentanol, p-anisaldehyde. Since in the original data, not all pixel values were available for all odorants, we kept only those that were common to all of them. We then downsampled the resulting 2074 data points to 240 sample points $\mathbf{S}^{(orig)}$.

The z-score data $\mathbf{S}^{(orig)}$ do not include the baseline activity of the glomeruli. Thus, a z-core of 0, which represents the mean activity across the whole population, does not imply that the activity was not affected by the odor presentation. To obtain a rough calibration of the mean activity, we used the observation that around 60% of the MCs are activated by any given strong stimulus [25]. In the absence of other information, we used this as a guide to re-calibrate the 40% percentile z-score $S^{(40\%)}$ as 0. Then, we normalized the data by their maximum, resulting in

$$\tilde{\mathbf{S}} = \frac{\mathbf{S}^{(orig)} - S^{(40\%)}}{\max(\mathbf{S}^{(orig)} - S^{(40\%)})}.$$

Further, MCs are activated by the airflow even without any odorants, and odor representations change mostly linearly with the concentration of odorants [60]. For any given mixture with concentration p of odor A and $(1 - p)$ of odor B we therefore modeled the stimulus as

$$\mathbf{S} = [p \tilde{\mathbf{S}}^{(A)} + (1 - p) \tilde{\mathbf{S}}^{(B)} + \tilde{\mathbf{S}}^{(air)}]_+.$$

Here $\tilde{\mathbf{S}}^{(air)} = 0.1$ is the ‘air’ stimulus and $[\cdot]_+$ the rectifier, $[x]_+ = \max(x, 0)$.

For the simplified model stimuli we used scaled Gaussian functions with different mean and standard deviation as $\tilde{\mathbf{S}}^{(A,B)}$ and set $\tilde{\mathbf{S}}^{(air)} = 0.2$.

Definition of responsive cells and divergent cells

We defined the response of an MC as the change in activity induced by an odor compared to the activity with air as stimulus. Specifically, if a MC had activity $\mathbf{M}^{(A,B)}$ in response to stimuli $\mathbf{S}^{(A,B)}$, and it had activity $\mathbf{M}^{(air)}$ when presenting air as a stimulus $\mathbf{S}^{(air)}$, we defined the response to odor A, B as $\mathbf{M}_i^{(A,B)} - \mathbf{M}_i^{(air)}$.

The definitions of responsive and divergent cells were adapted from [26]. A cell was classified as responsive, if its response to one odor in an odor pair was above a given threshold θ , i.e. $\max(\mathbf{M}_i^{(A)} - \mathbf{M}_i^{(air)}, \mathbf{M}_i^{(B)} - \mathbf{M}_i^{(air)}) > \theta$. A cell was classified as divergent, if the difference in its response to the two odors was above the same threshold, $|\mathbf{M}_i^{(A)} - \mathbf{M}_i^{(B)}| > \theta$. Our results were robust over a range of values of the threshold θ (Fig S14).

Characterizing Discriminability

We used d-prime and the Fisher discriminant to assess the discriminability of activity patterns. For each divergent neuron d-prime is given by the difference between the mean activities of the neuron for the two presented odors scaled by the combined variability of the activities. In our firing-rate model (1,2) the MC-activity exhibits fluctuations only on the time scale of network restructuring, which is much longer than the odor presentation. Assuming, as in [15], that the MC firing rates $M_i^{(A,B)}$ represent the mean values of independent Poisson spike trains, the variances $\Sigma^{(A,B)}$ of those spike trains are given by their respective means, $\Sigma_{ii}^{(A,B)} = M_i^{(A,B)}$ and $\Sigma_{ij}^{(A,B)} = 0, i \neq j$. Thus, for each MC we have

$$d'_i = \frac{|M_i^{(A)} - M_i^{(B)}|}{\sqrt{M_i^{(A)} + M_i^{(B)}}}. \quad (12)$$

One way to characterize the overall discriminability of the activity patterns is then given by the average $\langle d'_i \rangle$ of the d-prime values across all neurons, as was done in [26]. This average $\langle d'_i \rangle$ may decrease when additional, less-discriminating MCs are included in the ensemble. However, the information transmitted by the ensemble does not decrease. We therefore used also the Fisher discriminant to measure the discriminability. In general, it is given by

$$F = \frac{(\mathbf{w} \cdot (\mathbf{M}^{(A)} - \mathbf{M}^{(B)}))^2}{\mathbf{w}^T (\Sigma^{(A)} + \Sigma^{(B)}) \mathbf{w}},$$

where \mathbf{w} is an arbitrary vector, which can be interpreted as the weights with which the MC activity is fed into a linear read-out neuron. The weights \mathbf{w} that maximize F are given by $\mathbf{w}^{(opt)} = (\Sigma^{(A)} + \Sigma^{(B)})^{-1}(\mathbf{M}^{(A)} - \mathbf{M}^{(B)})$. The optimal Fisher discriminant used to assess discriminability is then given by

$$F_{opt} = \sum_{i=1}^{N_{MC}} \frac{(M_i^{(A)} - M_i^{(B)})^2}{M_i^{(A)} + M_i^{(B)}}. \quad (13)$$

It is related to the d-prime values of the individual neurons *via*

$$F_{opt} = \sum_{i=1}^{N_{MC}} (d'_i)^2.$$

Thus, F_{opt} always increases with the addition of MCs, reflecting the fact that even poorly discriminating MCs provide some additional information about the odors. The optimal read-out $\mathbf{w}^{(opt)}$ puts less weight on the poorly discriminating MCs to avoid that the trial-to-trial variability associated with those MCs deteriorates the overall discrimination.

Acknowledgments

We thank S. Saha, K.A. Sailor, and P.-M. Lledo for discussions and for sharing their data with us prior to publication. We gratefully acknowledge discussions with T. Komiyama.

This work was supported by funding from the National Institutes of Health (<https://www.nidcd.nih.gov/>) through grant DC015137 to HR.

Data availability

799

The authors declare that the data supporting the findings of this study are available within the paper and its Supplementary Information files or from the corresponding author on reasonable request. A sample code is available in the Supplementary Information files.

800

801

802

803

Author information 804

Affiliations 805

Engineering Sciences and Applied Mathematics, Northwestern University, 806
Evanston, IL, 60208, USA 807
John Hongyu Meng⁺, Hermann Riecke⁺ **current address: Center for Neural** 808
Science, New York University, New York, NY, 10003, USA 809
John Hongyu Meng 810

Contributions 811

J.M. and H.R. conceived the ideas of this study and designed the model. J.M. carried 812
out the simulation and analyzed the data. Both contributed to the writing of the 813
manuscript. 814

Corresponding authors 815

Correspondence to Hermann Riecke. 816

Ethics declarations 817

Competing interests 818

The authors declare no competing interests. 819

References 820

1. Yamaguchi M, Mori K. Critical period for sensory experience-dependent survival 821
of newly generated granule cells in the adult mouse olfactory bulb. *Proc Natl* 822
Acad Sci U S A. 2005;102(27):9697–9702. (document) 823
2. Lepousez G, Valley MT, Lledo PM. The impact of adult neurogenesis on 824
olfactory bulb circuits and computations. *Annual Review of Physiology*. 825
2013;75:339–363. (document) 826
3. Platel JC, Angelova A, Bugeon S, Wallace J, Ganay T, Chudotvorova I, et al. 827
Neuronal integration in the adult mouse olfactory bulb is a non-selective addition 828
process. *Elife*. 2019;8:e44830. (document) 829
4. Forest J, Moreno M, Cavelius M, Chalenccon L, Ziessel A, Sacquet J, et al. 830
Short-term availability of adult-born neurons for memory encoding. *Nature* 831
communications. 2019;10:5609. doi:10.1038/s41467-019-13521-7. (document) 832
5. Livneh Y, Mizrahi A. Experience-dependent plasticity of mature adult-born 833
neurons. *Nature Neuroscience*. 2012;15(1):26. (document) 834

SUBMISSION

6. Sailor KA, Valley MT, Wiechert MT, Riecke H, Sun GJ, Adams W, et al. 835
Persistent Structural Plasticity Optimizes Sensory Information Processing in the 836
Olfactory Bulb. *Neuron*. 2016;91(2):384–396. doi:10.1016/j.neuron.2016.06.004. 837
(document) 838
7. Friedrich RW, Laurent G. Dynamic Optimization of Odor Representations by 839
Slow Temporal Patterning of Mitral Cell Activity. *Science*. 2001;291:889. 840
(document) 841
8. Gschwend O, Abraham NM, Lagier S, Begnaud F, Rodriguez I, Carleton A. 842
Neuronal pattern separation in the olfactory bulb improves odor discrimination 843
learning. *Nature Neuroscience*. 2015;18:1474–1482. doi:10.1038/nn.4089. 844
(document) 845
9. Cleland TA, Sethupathy P. Non-Topographical Contrast Enhancement in the 846
Olfactory Bulb. *BMC Neurosci*. 2006;7:7. (document) 847
10. Wiechert MT, Judkewitz B, Riecke H, Friedrich RW. Mechanisms of pattern 848
decorrelation by recurrent neuronal circuits. *Nat Neurosci*. 2010;13(8):1003–1010. 849
(document) 850
11. Arevian AC, Kapoor V, Urban NN. Activity-dependent gating of lateral 851
inhibition in the mouse olfactory bulb. *Nat Neurosci*. 2008;11(1):80–87. 852
(document) 853
12. Linster C, Sachse S, Galizia CG. Computational Modeling Suggests That 854
Response Properties Rather Than Spatial Position Determine Connectivity 855
Between Olfactory Glomeruli. *J Neurophysiol*. 2005;93(6):3410–3417. (document) 856
13. Wick S, Wiechert M, Friedrich R, Riecke H. Pattern orthogonalization via 857
channel decorrelation by adaptive networks. *J Comput Neurosci*. 2010;28:29–45. 858
(document) 859
14. Chow SF, Wick SD, Riecke H. Neurogenesis Drives Stimulus Decorrelation in a 860
Model of the Olfactory Bulb. *PLoS Comp Biol*. 2012;8:e1002398. (document) 861
15. Adams W, Graham JN, Han X, Riecke H. Top-down inputs drive neuronal 862
network rewiring and context-enhanced sensory processing in olfaction. *PLoS 863
Comp Biol*. 2019;15(1):e1006611. (document) 864
16. McNamara AM, Magidson PD, Linster C, Wilson DA, Cleland TA. Distinct 865
Neural Mechanisms Mediate Olfactory Memory Formation At Different 866
Timescales. *Learning & Memory*. 2008;15(3):117–125. (document) 867
17. Moreno MM, Linster C, Escanilla O, Sacquet J, Didier A, Mandairon N. 868
Olfactory perceptual learning requires adult neurogenesis. *Proceedings of the 869
National Academy of Sciences*. 2009;106(42):17980. (document) 870
18. Mandairon N, Kuczewski N, Kermen F, Forest J, Midroit M, Richard M, et al. 871
Opposite regulation of inhibition by adult-born granule cells during implicit 872
versus explicit olfactory learning. *eLife*. 2018;7:e34976. doi:10.7554/eLife.34976. 873
(document) 874
19. Rochefort C, Gheusi G, Vincent JD, Lledo PM. Enriched Odor Exposure 875
Increases the Number of Newborn Neurons in the Adult Olfactory Bulb and 876
Improves Odor Memory. *J Neurosci*. 2002;22(7):2679–2689. (document) 877

SUBMISSION

20. Lin CW, Sim S, Ainsworth A, Okada M, Kelsch W, Lois C. Genetically increased cell-intrinsic excitability enhances neuronal integration into adult brain circuits. *Neuron*. 2010;65:32–9. (document) 878
879
880
21. Wu A, Yu B, Chen Q, Matthews GA, Lu C, Campbell E, et al. Context-dependent plasticity of adult-born neurons regulated by cortical feedback. *Science Advances*. 2020;6(42):eabc8319. (document) 881
882
883
22. Breton-Provencher V, Coté D, Saghatelian A. Activity of the principal cells of the olfactory bulb promotes a structural dynamic on the distal dendrites of immature adult-born granule cells via activation of NMDA receptors. *Journal of Neuroscience*. 2014;34(5):1748–1759. (document) 884
885
886
887
23. Breton-Provencher V, Bakhshetyan K, Hardy D, Bammann RR, Cavarretta F, Snappyan M, et al. Principal cell activity induces spine relocation of adult-born interneurons in the olfactory bulb. *Nature Communications*. 2016;7:12659. (document) 888
889
890
891
24. Soucy ER, Albeanu DF, Fantana AL, Murthy VN, Meister M. Precision and diversity in an odor map on the olfactory bulb. *Nat Neurosci*. 2009;12(2):210–220. (document) 892
893
894
25. Kato HK, Chu MW, Isaacson JS, Komiyama T. Dynamic sensory representations in the olfactory bulb: modulation by wakefulness and experience. *Neuron*. 2012;76(5):962–975. (document), 4 895
896
897
26. Chu MW, Li WL, Komiyama T. Balancing the robustness and efficiency of odor representations during learning. *Neuron*. 2016;92(1):174–186. (document), 2 898
899
27. Meng JH, Riecke H. Structural Spine Plasticity in Olfaction: Memory and Forgetting, Enhanced vs. Reduced Discriminability after Learning. *bioRxiv* 20201204411629. 2020;doi:10.1101/2020.12.04.411629. (document) 900
901
902
28. Orona E, Rainer EC, Scott JW. Dendritic and axonal organization of mitral and tufted cells in the rat olfactory bulb. *Journal of Comparative Neurology*. 1984;226(3):346–356. (document) 903
904
905
29. Egger V, Urban NN. Dynamic connectivity in the mitral cell-granule cell microcircuit. *Sem Cell Devel Biol*. 2006;17(4):424–432. (document) 906
907
30. Kersen DEC, Tavoni G, Balasubramanian V. Connectivity and dynamics in the olfactory bulb. *PLoS computational biology*. 2022;18:e1009856. (document) 908
909
31. Gao Y, Strowbridge BW. Long-term plasticity of excitatory inputs to granule cells in the rat olfactory bulb. *Nature Neuroscience*. 2009;12(6):731–733. (document) 910
911
912
32. Chatterjee M, Perez de los Cobos Pallares F, Loebel A, Lukas M, Egger V. Sniff-like patterned input results in long-term plasticity at the rat olfactory bulb mitral and tufted cell to granule cell synapse. *Neural Plasticity*. 2016;2016. (document) 913
914
915
916
33. Saha S, Meng JH, Riecke H, Agoranos G, Sailor KA, Lledo PM. Interneuron activity-structural plasticity association is driven by context-dependent odor experience. *bioRxiv* 20210128428612. 2021;. (document) 917
918
919

34. Bienenstock EL, Cooper LN, Munro PW. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*. 1982;2(1):32–48. (document)
35. Graupner M, Brunel N. Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proceedings of the National Academy of Sciences of the United States of America*. 2012;109(10):3991–3996. (document)
36. Lepousez G, Nissant A, Bryant AK, Gheusi G, Greer CA, Lledo PM. Olfactory learning promotes input-specific synaptic plasticity in adult-born neurons. *Proc Nat Acad Sci USA*. 2014;111(38):13984–13989. (document)
37. Zenke F, Gerstner W. Hebbian plasticity requires compensatory processes on multiple timescales. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2017;372(1715):20160259. (document)
38. Moreno MM, Linster C, Escanilla O, Sacquet J, Didier A, Mandairon N. Olfactory perceptual learning requires adult neurogenesis. *Proc Nat Acad Sci USA*. 2009;106(42):17980. (document)
39. Johnson BA, Farahbod H, Saber S, Leon M. Effects of functional group position on spatial representations of aliphatic odorants in the rat olfactory bulb. *Journal of Comparative Neurology*. 2005;483(2):192–204. (document), 2
40. Buonviso N, Chaput M. Olfactory experience decreases responsiveness of the olfactory bulb in the adult rat. *Neuroscience*. 2000;95:325. (document)
41. Linster C, Nai Q, Ennis M. Non-linear effects of noradrenergic modulation of olfactory bulb function in adult rodents. *J Neurophysiol*. 2011;. (document)
42. Poe GR, Foote S, Eschenko O, Johansen JP, Bouret S, Aston-Jones G, et al. Locus coeruleus: a new look at the blue spot. *NATURE REVIEWS NEUROSCIENCE*. 2020;21(WOS:000570488300001):644–659. doi:10.1038/s41583-020-0360-9. (document)
43. Cleland TA, Johnson BA, Leon M, Linster C. Relational representation in the olfactory system. *Proc Natl Acad Sci U S A*. 2007;104(6):1953–1958. (document)
44. Oh WC, Parajuli LK, Zito K. Heterosynaptic structural plasticity on local dendritic segments of hippocampal CA1 neurons. *Cell Reports*. 2015;10(2):162–169. (document)
45. Triesch J, Vo AD, Hafner AS. Competition for synaptic building blocks shapes synaptic plasticity. *Elife*. 2018;7:e37836. (document)
46. Veyrac A, Sacquet J, Nguyen V, Marien M, Jourdan F, Didier A. Novelty determines the effects of olfactory enrichment on memory and neurogenesis through noradrenergic mechanisms. *Neuropsychopharmacology*. 2009;34(3):786–795. (document)
47. Mandairon N, Stack C, Kiselycznyk C, Linster C. Enrichment to odors improves olfactory discrimination in adult rats. *Behav Neurosci*. 2006;120(1):173–179. (document)
48. Anderson MC, Bjork EL, Bjork RA. Retrieval-induced forgetting: Evidence for a recall-specific mechanism. *Psychonomic Bulletin & Review*. 2000;7(3):522–530. (document)

49. Urban NN. Lateral inhibition in the olfactory bulb and in olfaction. *Physiology & Behavior*. 2002;77(4-5):607–612. (document) 964
965
50. Greer CA. Golgi Analyses of Dendritic Organization among Denervated Olfactory-bulb Granule Cells. *Journal of Comparative Neurology*. 1987;257(3):442–452. doi:10.1002/cne.902570311. (document) 966
967
968
51. Hiratani N, Fukai T. Hebbian wiring plasticity generates efficient network structures for robust inference with synaptic weight plasticity. *Frontiers in Neural Circuits*. 2016;10:41. (document) 969
970
971
52. Hedrick NG, Lu Z, Bushong E, Singhi S, Nguyen P, Magana Y, et al. Learning binds new inputs into functional synaptic clusters via spinogenesis. *Nature neuroscience*. 2022;. (document) 972
973
974
53. Buonviso N, Amat C, Litaudon P, Roux S, Royet JP, Farget V, et al. Rhythm sequence through the olfactory bulb layers during the time window of a respiratory cycle. *European Journal of Neuroscience*. 2003;17(9):1811–1819. (document) 975
976
977
978
54. Beshel J, Kopell N, Kay LM. Olfactory Bulb Gamma Oscillations Are Enhanced with Task Demands. *Journal of Neuroscience*. 2007;27(31):8358–8365. doi:10.1523/JNEUROSCI.1199-07.2007. (document) 979
980
981
55. Grelat A, Benoit L, Wagner S, Moigneu C, Lledo PM, Alonso M. Adult-born neurons boost odor-reward association. *Proc Nat Acad Sci USA*. 2018;115:2514–2519. doi:10.1073/pnas.1716400115. (document) 982
983
984
56. Luna VM, Schoppa NE. GABAergic circuits control input-spike coupling in the piriform cortex. *J Neurosci*. 2008;28(35):8851–8859. (document) 985
986
57. Bolding KA, Franks KM. Recurrent cortical circuits implement concentration-invariant odor coding. *Science (New York, NY)*. 2018;361. doi:10.1126/science.aat6904. (document) 987
988
989
58. Meng JH, Riecke H. Synchronization by uncorrelated noise: interacting rhythms in interconnected oscillator networks. *Scientific Reports*. 2018;8:6949. doi:10.1038/s41598-018-24670-y. (document) 990
991
992
59. Xu X, Riecke H. Paradoxical phase response of gamma rhythms facilitates their entrainment in heterogeneous networks. *PLoS Comp Bio*. 2021;17(6):e1008575. (document) 993
994
995
60. Khan AG, Thattai M, Bhalla US. Odor representations in the rat olfactory bulb change smoothly with morphing stimuli. *Neuron*. 2008;57(4):571–585. (document) 996
997
61. Singh V, Murphy NR, Balasubramanian V, Mainland JD. Competitive binding predicts nonlinear responses of olfactory receptors to complex mixtures. *Proc Nat Acad Sci USA*. 2019;116:9598–9603. (document) 998
999
1000
62. Mandairon N, Kuczewski N, Kermen F, Forest J, Midroit M, Richard M, et al. Opposite regulation of inhibition by adult-born granule cells during implicit versus explicit olfactory learning. *Elife*. 2018;7:e34976. (document) 1001
1002
1003
63. Turrigiano G. Homeostatic synaptic plasticity: local and global mechanisms for stabilizing neuronal function. *Cold Spring Harbor Perspectives in Biology*. 2012;4(1):a005736. (document) 1004
1005
1006

SUBMISSION

64. Artola A, Singer W. Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends in Neurosciences*. 1993;16(11):480–487. (document) 1007
1008
1009

Supplementary Materials

1010

Randomized Connectivity

1011

To further illustrate that the network learned not by simply changing the number of synapses, but by developing a specific, stimulus-dependent connectivity, we assessed the performance of the network when the connections of each GC were rewired to random MCs with probability p_{rewire} while keeping the number of synapses on each GC the same. For $p_{rewire} = 0$ (Fig S1A), the connection was the originally learned network; for $p_{rewire} = 1$, the connection for each GC was totally random (Fig S1C). Indeed, for the highly similar training odors, the discriminability predominantly decreased as the randomness is increased (Fig S1D). For the network trained on the dissimilar odors (easy task), the discriminability increased, as expected, when the randomness was increased (Fig S1F). Interestingly, when P_{rewire} was increased beyond ≈ 0.7 these trends reversed (Fig S1D, F) somewhat.

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

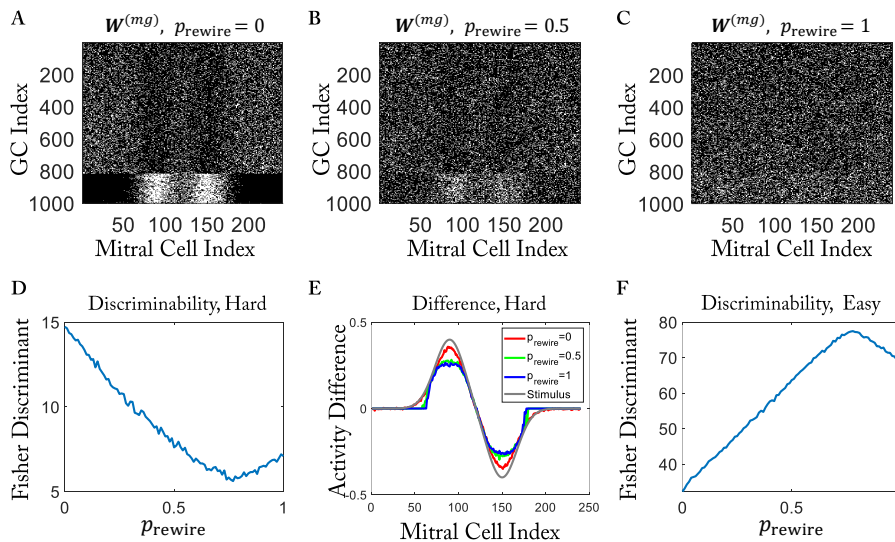


Fig S1. Randomly rewiring trained networks. (A to C) Connectivity matrices after rewiring a network that was trained with the hard task (Fig 3C, bottom). Rewiring probability $p_{rewire} = 0, 0.5, 1$, respectively. (D) For the hard task discriminability predominantly decreases with increasing rewiring probability p_{rewire} . (E) Activity differences are reduced after rewiring. (F) For the easy task discriminability predominantly increases with increasing p_{rewire} (cf. Fig 3C, top).

Removal of Spines for Intermediate GC Activation (cf. Fig.7)

1023

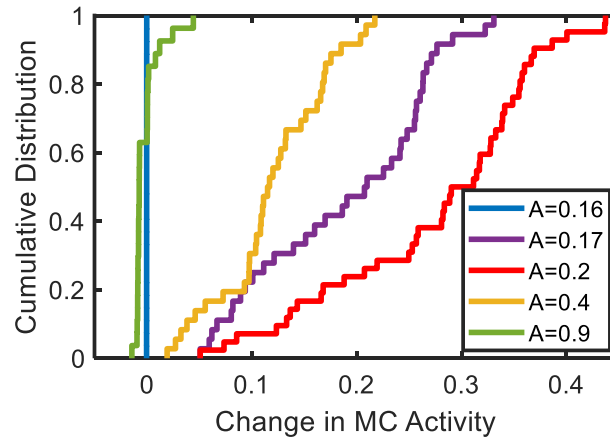


Fig S2. CDF for the increase in the MC response to odor A due to training with that odor at a reduced amplitude $2A$ after the network has been pre-trained with that odor at amplitude 2. For protocol see Fig.7.

1024

Forgetting due to Interference

1025

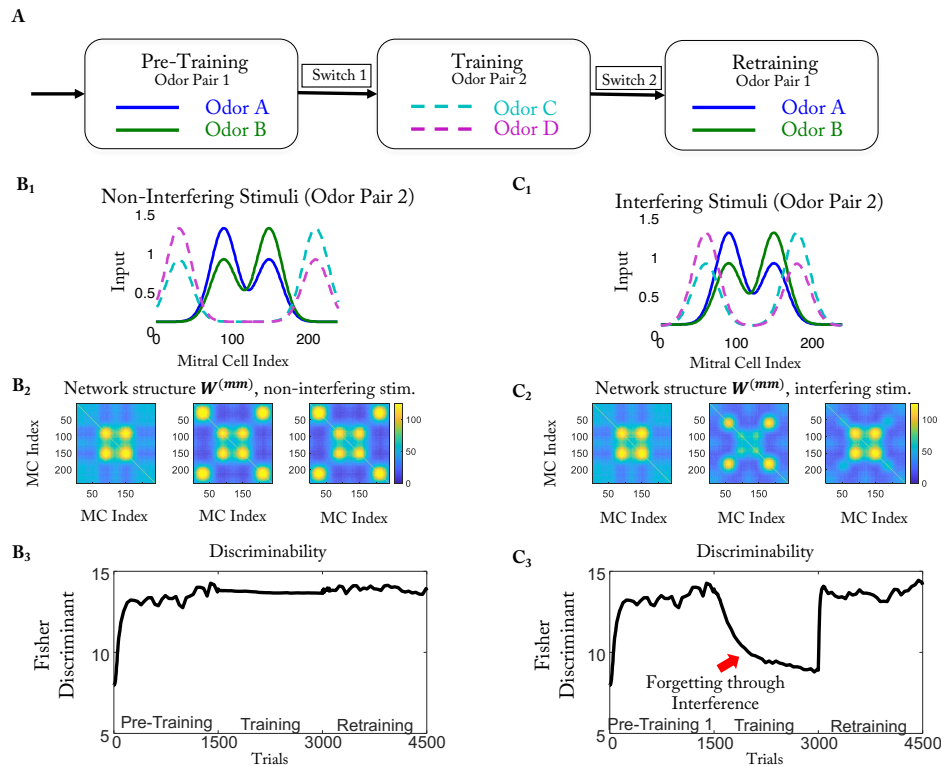


Fig S3. Interference leads to forgetting. (A) Training protocol. (B₁) Non-interfering stimuli for odor pair 2. (B₂) Connectivities after pre-training, training, and retraining, respectively. (B₃) Discriminability is retained throughout the training. (C) as (B) for interfering stimuli. During the training the network forgets most of the previously learned structure. Retraining recovers the previous performance.

The behavior of the model depended strongly on the similarity of the two odor pairs (Fig S3B,C). When the MCs that were activated during the training did not overlap with those activated in the pre-training, the network preserved the previously learned structure (Fig S3B₂). However, if there was significant overlap, the learning of the new stimuli interfered with the previously learned structure and that odor pair was forgotten during the training (Fig S3C₂, middle panel). It was re-learned by re-training (Fig S3C₂, right panel). The different evolution for non-interfering and for interfering stimuli is reflected in the Fisher discriminant $\mathcal{F}^{(1,2)}$, which is given by the sum over the squares of the d_i^2 of the individual MCs in response to stimulus pair 1 and 2, respectively (see Methods). The Fisher discriminant $\mathcal{F}^{(1)}$ remained high during the training with odor pair 2 if that pair was non-interfering (Fig S3B₃), while it decreased if odor pair 2 was interfering (Fig S3C₃). We expect that animals will spontaneously discriminate the odors in a pair only if their bulbar representations are sufficiently different, i.e., if \mathcal{F} is sufficiently large [38]. Thus, the model predicts that at the end of the training phase of a multi-phase perceptual learning task an animal will not spontaneously discriminate any more the odors it learned during pre-training if the training odors interfere with those used in the pre-training. However, for sufficiently different odor pairs, the spontaneous discrimination should remain intact.

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043

Rapid Relearning

1044

Repeatedly switching between odor pairs did not impair the learning ability (Fig S4). In addition, the model predicts that re-learning during re-training is faster than the initial learning during pre-training and forgetting is slower than either of the learning processes.

1045

1046

1047

1048

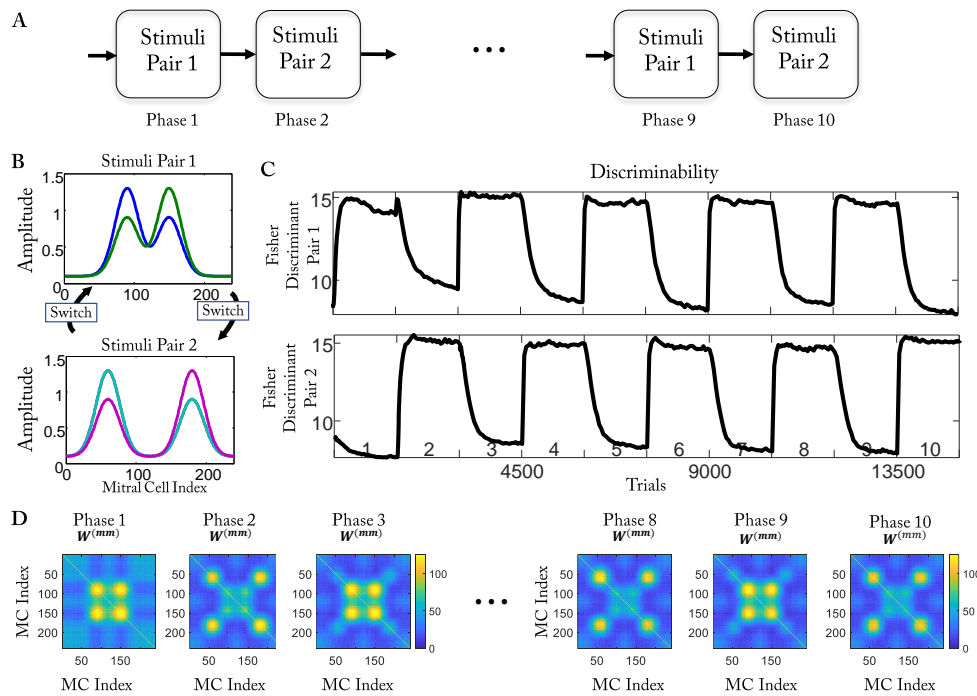


Fig S4. Alternating training does not impair learning ability. (A) Expanding the training protocol of S3A to 10 phases. (B) (Top) Stimuli of odor pair 1. (Bottom) Interfering stimuli of odor pair 2. (C) The Fisher discriminant of odor pair 1 (top) and pair 2 (bottom). Learning (increasing Fisher discriminant) proceeds faster than forgetting (decreasing Fisher discriminant). The numbers above the x-axis indicate the learning phase. (D) Effective connectivity $W^{(mm)}$ alternates.

1049

Resource Pool Model

1050

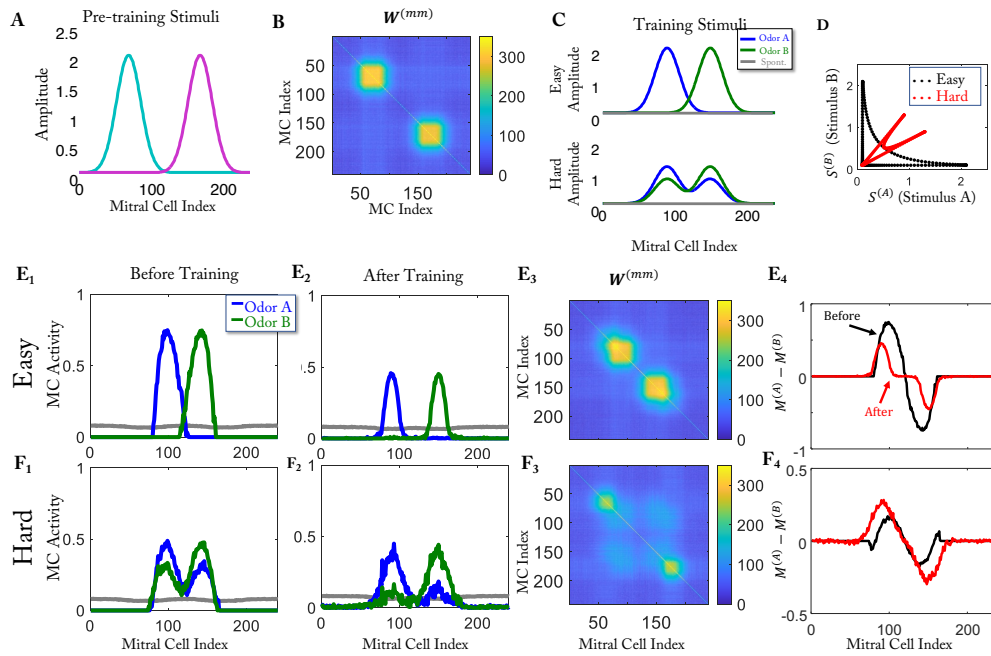


Fig S5. Results of the resource-pool model when trained with simplified stimuli. Here $\gamma = 1.7 \cdot 10^{-3}$. Training with dissimilar stimuli reduces their discriminability (E), whereas training with very similar enhances their discriminability (F) (cf. Fig.3)

1051

Impact of the Pre-Training

1052

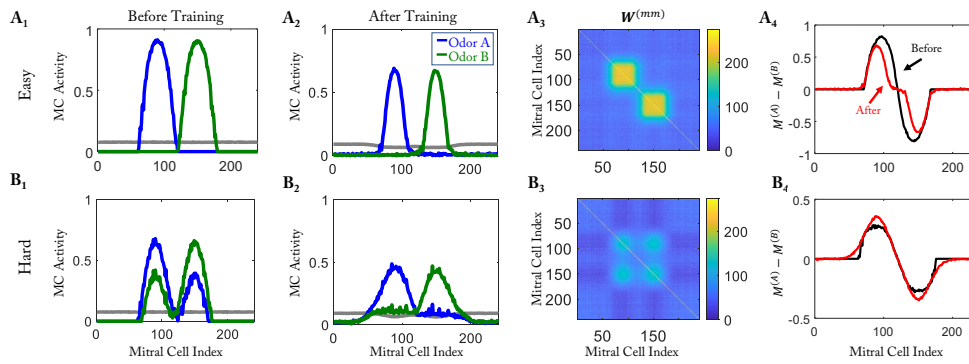


Fig S6. Training results without pre-training phase. The training starts with a random homogeneous network (cf. Fig 3)

1053

1054

Robustness of the Model

1055

To demonstrate the robustness of the model, we first assessed the impact of some modifications of the model.

1056

1057

For the formation of the reciprocal synapses little is known about the timescales of the maturation of the excitatory synapses relative to that of the inhibitory synapses. In the main part of the paper we assumed that inhibition becomes functional at the same as the excitation. Here, we tested if the model behaves differently if inhibitory synapses mature later than the excitatory synapses. If the inhibitory synapses become only functional 2 computational time steps after the excitatory synapses, the model performs similarly (Fig S7 A-C). Since 100 timesteps in our model correspond very roughly to about 1 day training in the experiments, a lag of 2 timesteps would correspond to a lag on the order of 30 minutes. However, the performance deteriorates if the lag is increased and the model fails for lags of 10 or larger. A detailed analysis of this break-down is beyond the scope of this study.

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

While action potentials can propagate along the secondary dendrites of MCs and can drive excitatory synapses on spines even far from the soma, the GC-driven inhibition originating from those spines is likely to affect the MC somata only within a quite limited spatial range. To assess the impact of this spatial limitation, we blocked the inhibition in half of the MC-GC pairs: while in all spines the MC→GC connections excited the GCs, only half of the MC←GC connections in those spines were effective in inhibiting the MCs. To make the total amount of inhibition comparable, we doubled the number of granule cells. With this modification, the results are similar to the original model (Fig S7 D, E).

1069

1070

1071

1072

1073

1074

1075

1076

1077

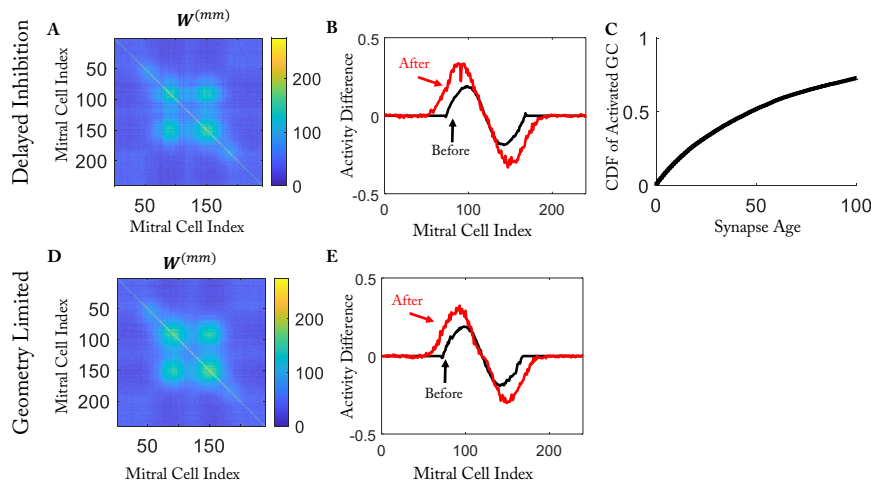


Fig S7. Modifications with asymmetric inhibitory and excitatory synapses show similar results in a hard task with simplified stimuli. (A - C) The results of setting maturation of the inhibitory synapses lagged 2 timesteps. (A) The effective connectivity after training. (B) Activity difference before and after training. (C) The cumulative distribution function of synapse age. When the lag is 2, only a limited fraction of connections does not inhibit MCs. (D, E) The results when the inhibitory synapses have a geometry limitation. Organized as (A, B).

Further, we tested the influence of various parameters on the model performance.

1078

Changing the total number N_{GC} of GCs affected the performance very little, since the number of GCs that were strongly activated and connected to specific MCs did not

1079

1080

change significantly, which influenced the effective connectivity $W^{(mm)}$ only mildly (Fig S8). Analogously, changes in the inhibitory strength γ were largely compensated by a change in the number of activated GCs (Fig S9). This feedback stabilizing the overall inhibition arises from the reciprocal character of the MC-GC connections: new connections increase the inhibition of the MCs, which decreases the excitation of the GC and in turn inhibits the formation of new connections.

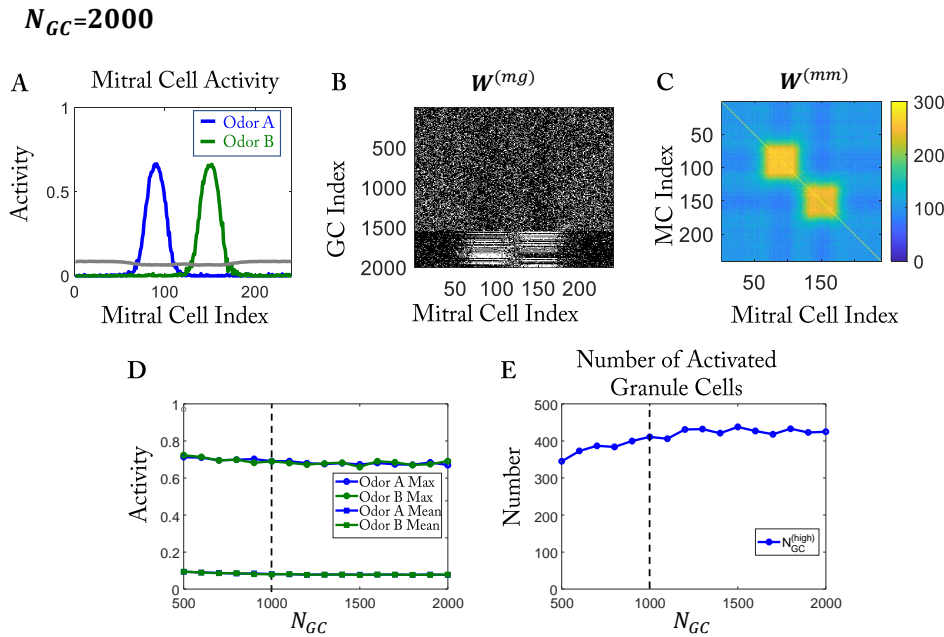


Fig S8. The overall inhibition of the model depends only weakly on the number N_{GC} of GCs. (A-C) Training with simplified easy stimuli as in Fig 1D and random initial network as in Fig 1F, but with twice as many GCs. (A) MC activity after training (cf. Fig 1H). (B) Connectivity $W^{(mg)}$ after training. (C) Effective connectivity $W^{(mm)}$. Compared to Fig 1J the connectivity is not quite as selective. (D) The maximal and mean MC activity after training as a function of the number of GCs. (E) The number of activated GCs ($G > G^{(1)}$) depends only weakly on the total number of GCs. In (D, E) the vertical dashed line marks the value used in the rest of the paper.

What does control the overall inhibition, is the threshold $G^{(1)}$ in the activation function $\phi(G)$ (Fig S10). Only GCs with activity larger than $G^{(1)}$ form more synapses and generate stronger selective inhibition. Thus, for larger $G^{(1)}$ fewer GCs had sufficiently large activity to sustain their synapses, which reduced the overall inhibition and increased the MC activity level. The change in overall inhibition could be associated with a change in the selectivity. To quantify this, we called a GC responsive to odor A (B), if after training its activity surpassed $G^{(1)}$ for odor A (B). Despite the increase in the inhibition, the number of GCs that responded to both odors remained low when decreasing $G^{(1)}$ (Fig S10H, red line), indicating that the selectivity depended very little on $G^{(1)}$.

The other threshold in the activation function, $G^{(0)}$, affects the removal but not the formation of synapses. For low values of $G^{(0)}$ weakly activated GCs removed many of their synapses, particularly those connecting them to active MCs (Fig S11C, D). This lead to higher selectivity of the connectivity and with it to a weaker inhibition of the spontaneously active MCs that were not odor-driven (Fig S11B, E). It had, however,

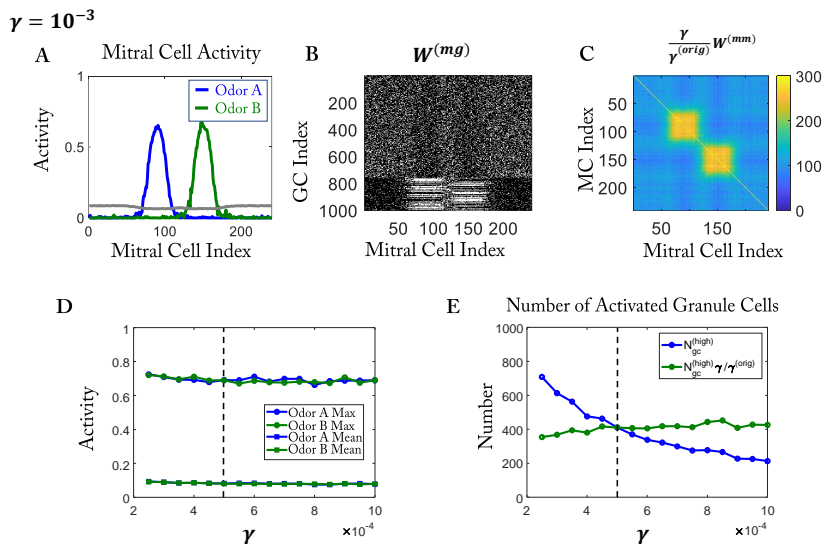


Fig S9. The overall inhibition of the model depends only weakly on the inhibitory strength γ . The results are organized as in Fig S8. (C) The connectivity is slightly less selective than in Fig 1J. (D,E) The vertical dashed line indicates the value used in the rest of the paper. (D, E) To allow a direct comparison, $W^{(mm)}$ and $N_{GC}^{(high)}$ have been re-scaled by $\frac{\gamma}{\gamma^{(orig)}}$ with $\gamma^{(orig)} = 5 \times 10^{-4}$.

little impact on the maximal MC amplitudes (Fig S11H). The removal of synapses played a larger role in retaining and forgetting connections. To assess its impact, we trained the model in a second phase with a new pair of stimuli that overlapped with the stimuli learned in phase 1 (Fig S11I). For $G^{(0)} = 0$ any weak activity ($0 < G < G^{(1)}$) triggered the removal of spines and lead in phase 2 to an almost complete forgetting of the connectivity learned in phase 1 (Fig S11D, K) and with it to poor discrimination of the previously learned odors (Fig S11N). However, for large $G^{(0)}$ the previously learned connectivity was maintained (Fig S11G, M) and with it the discriminability of the previously learned stimuli. The discrimination of the newly learned stimulus pair 2 was, however, not significantly affected by $G^{(0)}$ (Fig S11O).

The maximal number of connections k that each GC can make had a significant impact on MC amplitudes and on the selectivity of the connections. With increasing k , each GC integrated input from more MCs, and more easily surpassed the threshold $G^{(1)}$. Each of these more activated GCs inhibited a larger number of MCs, increasing the overall inhibition (Fig S12C). When the maximal number of connections was larger than the number of MCs activated by one of the odors, not all connections of a GC could be made to MCs that were activated by odor A, say. The remaining connections could then be made to MCs corresponding to odor B without destabilizing the other synapses. As a result, GCs that responded to both odors emerged (Fig S12D).

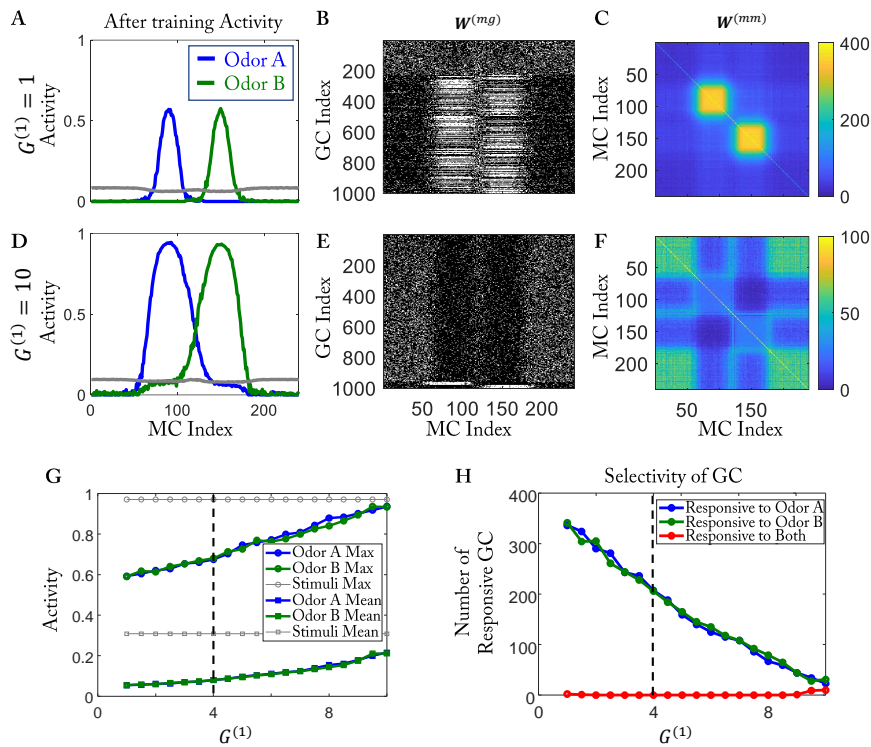


Fig S10. Increasing $G^{(1)}$ impairs the ability to form connections between activated MCs and GCs. (A,B,C) Results for $G^{(1)} = 1$. (A) MC activity after training. (B) Connectivity $W^{(mg)}$ between MCs and GCs. (C) Effective connectivity $W^{(mm)}$. (D,E,G) as (A,B,C) except for $G^{(1)} = 10$. Activated MCs are connected with fewer GCs (compare E with B), resulting in weaker disynaptic inhibition (compare F with C) and higher MC activity (compare D to A). In (F) most of the GCs cannot reach the high threshold. As a result, the synapses that connect to strongly activated MCs are removed faster than those connecting to weakly activated MCs. Thus, the effective connectivity among the activated MCs is lower than the background. (G) Maximal and mean MC activity increase with increasing $G^{(1)}$. The gray lines indicate the corresponding values without inhibition from GCs (by setting $\gamma = 0$). (H) The number of responsive GC decreases with increasing $G^{(1)}$. (G,H) The vertical dashed line indicates the value used in the rest of the paper.

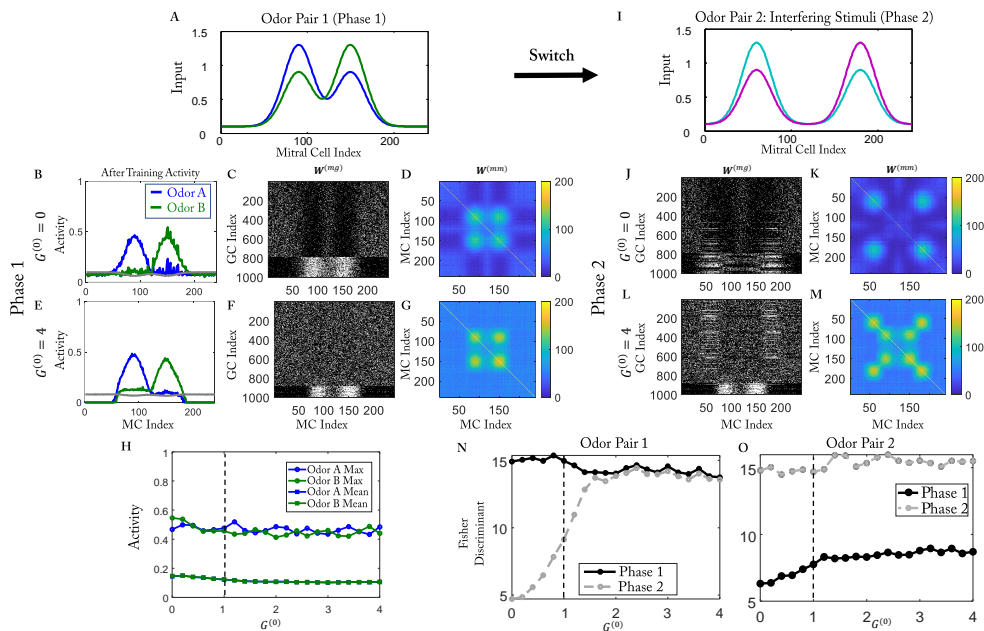


Fig S11. Learning ability is not affected by $G^{(0)}$, but the memory is. (A-G) Phase 1. (A) Training stimuli in phase 1. (B) MC activity after training. (C) Connectivity $W^{(mg)}$. (D) Effective connectivity $W^{(mm)}$. (E to G) as (B to D) except for $G^{(0)} = 4$. (H) Maximal and mean MC activity after training in phase 1 as a function of $G^{(1)}$. (I to M) as (A to G) but after Phase 2. (I) Interfering training stimuli. (J-M) For $G^{(0)} = 0$ the network forgets the previously learned connectivity, but not for $G^{(0)} = 4$. (N) The Fisher discriminant for odor pair 1 is unaffected by $G^{(0)}$ in phase 1, but substantially reduced in phase 2 for small $G^{(0)}$. (O) For odor pair 2 the Fisher discriminant depends only little on $G^{(0)}$. The vertical dashed lines in (H,N,O) indicate the value used in the rest of the paper.

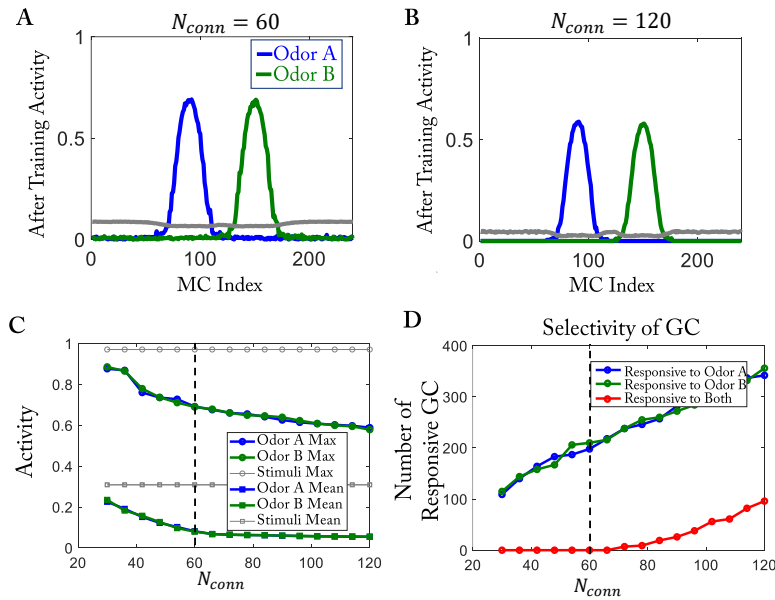


Fig S12. The selectivity of GCs is influenced by the number of connections each GC makes. Easy stimuli as in Fig 1D with random initial network as in Fig 1F. The ratio of maximal to initial number of connections is kept fixed at $k/N_{conn} \equiv 1.1$. (A, B) The MC activity is very similar for $N_{conn} = 60$ and $N_{conn} = 120$. (C) Maximal and mean MC activity decrease with increasing N_{conn} . The gray lines indicate the corresponding values without inhibition from GCs (by setting $\gamma = 0$). (D) The selectivity of the GCs is impaired for $N_{conn} > 60$: a population of cells emerges that respond to both odors (red line).

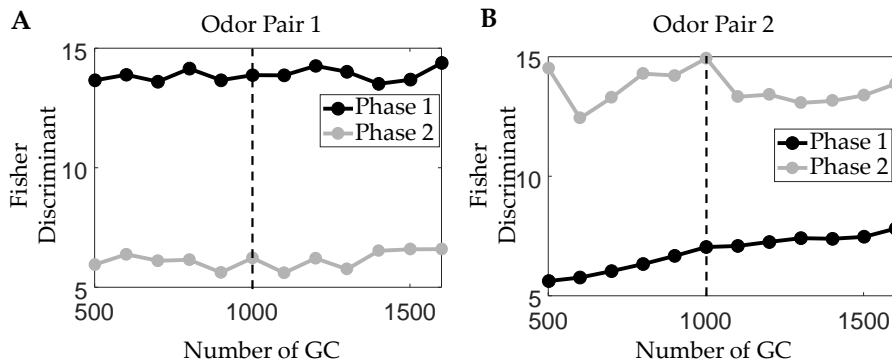


Fig S13. Increasing the number of GCs does not reduce the interference-induced forgetting of odor pair 1 during training phase 2. Independent of the number of GCs, the discriminability of odor pair 1, which is learned during phase 1 (black line in panel A), is substantially reduced by the subsequent training during phase 2 (gray line in panel A). The discriminability of odor pair 2 that is reached after training with those odors during phase 2 also does not depend on the number of GCs. Only the learning of odor pair 2 during the training with odor pair 1 is slightly improved in the larger network. ($G^{(0)} = 1$, cf. Fig S11N,O).

Dependence on the Threshold θ

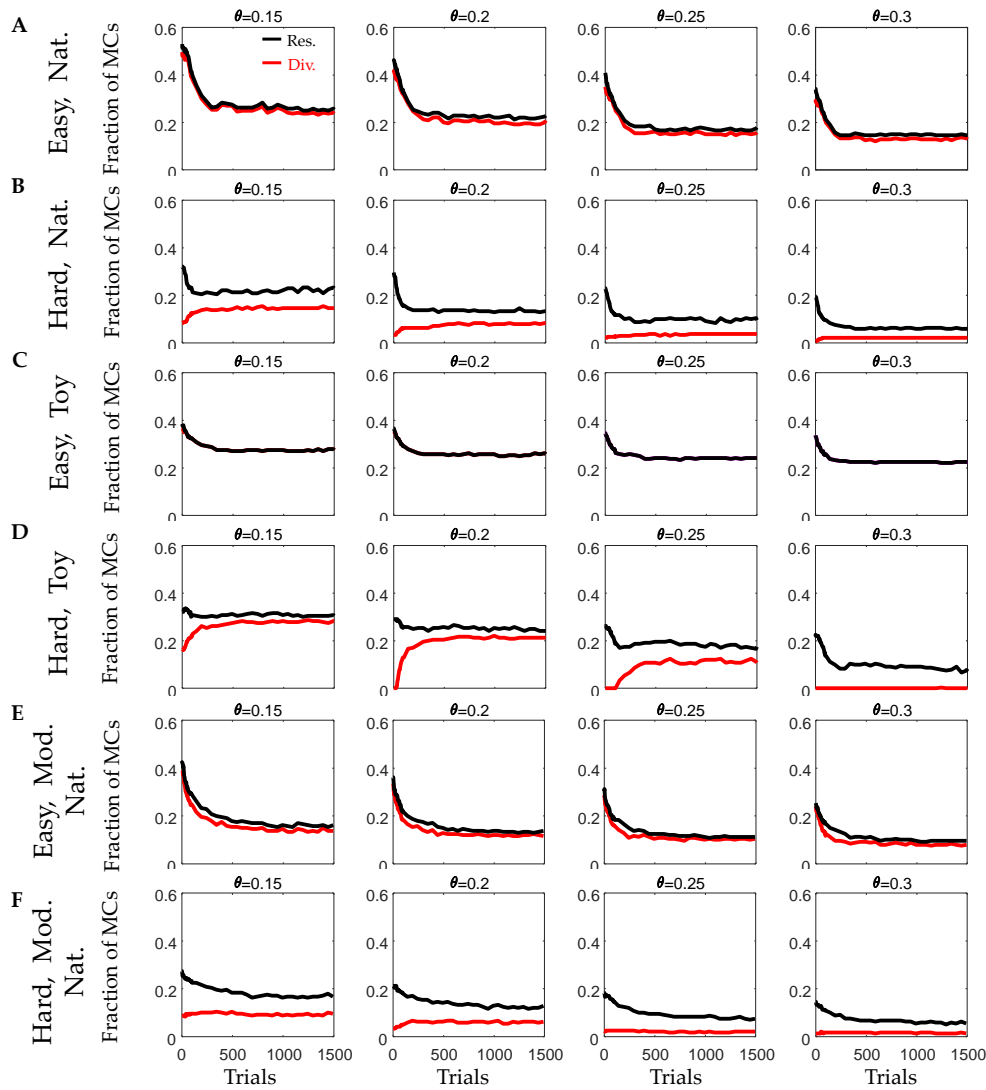


Fig S14. Qualitatively, the results for the fraction of responsive and divergent cells do not depend on the threshold θ for their classification. (A, B) Training with naturalistic stimuli for different thresholds θ . The results in Fig 2 and Fig 9 are based on $\theta = 0.2$. (C, D) as (A, B) but for training with simplified stimuli. (E, F) as (A, B) but in the resource-pool model.