Neuron
Article

Sparse Incomplete Representations: A Potential Role of Olfactory Granule Cells

Alexei A. Koulakov[1,*] and Dmitry Rinberg[2]
[1]Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724
[2]Janelia Farm Research Campus, Howard Hughes Medical Institute, Ashburn, VA 20147
*Correspondence: akula@cshl.edu
DOI 10.1016/j.neuron.2011.07.031

SUMMARY

Mitral/tufted cells of the olfactory bulb receive odorant information from receptor neurons and transmit this information to the cortex. Studies in awake behaving animals have found that sustained responses of mitral cells to odorants are rare, suggesting sparse combinatorial representation of the odorants. Careful alignment of mitral cell firing with the phase of the respiration cycle revealed brief transient activity in the larger population of mitral cells, which respond to odorants during a small fraction of the respiration cycle. Responses of these cells are therefore temporally sparse. Here, we propose a mathematical model for the olfactory bulb network that can reproduce both combinatorially and temporally sparse mitral cell codes. We argue that sparse codes emerge as a result of the balance between mitral cells' excitatory inputs and inhibition provided by the granule cells. Our model suggests functional significance for the dendrodendritic synapses mediating interactions between mitral and granule cells.

INTRODUCTION

The olfactory bulb is the first processing center of information about odorants. In mammals, the olfactory system is the only sensory system in which peripheral information is sent directly to the cortex, bypassing the sensory thalamus. Therefore, it has been proposed that the bulb combines the function of peripheral sensory system and the thalamus (Kay and Sherman, 2007). Consistent with this proposal, several studies have demonstrated that activity in the olfactory bulb reflects not only sensory information but also the animal's internal state (Adrian, 1950; Rinberg et al., 2006) and task-dependent variables (Doucette and Restrepo, 2008; Fuentes et al., 2008; Kay and Laurent, 1999). The relative simplicity of the anatomy of the olfactory bulb and the combination of both sensory- and state-dependent activity in a single network make it an attractive model for the study of principles of sensory information processing.

The surface of the olfactory bulb is covered by ≈ 2000 glomeruli. Each glomerulus receives inputs from a set of receptor neurons expressing the same type of olfactory receptor protein. The inputs into individual glomerulus from receptor neurons are therefore substantially correlated (Koulakov et al., 2007; Lledo et al., 2005; Shepherd et al., 2004; Wachowiak et al., 2004). This glomerulus-based modularity is preserved further by the mitral cells (MCs), most of which receive direct excitatory inputs from a single glomerulus only. MCs are a major output class of the olfactory bulb. These cells transmit information about odorants to the olfactory cortex (Figure 1).

The representation of odorants by MCs is often described as combinatorial code (Firestein, 2004; Koulakov et al., 2007). For such a code, both odorant identity and concentration can be derived from the particular combination of active MCs. A large number of MCs provides enough combinatorial diversity to encode virtually any relevant stimulus. Early studies of the MC code have found that the sustained responses of MCs to odorants are sparse and state dependent (Adrian, 1950; Kay and Laurent, 1999; Rinberg et al., 2006). Sparseness of combinatorial representation implies that only a small fraction of cells displays detectable responses to odorants. In awake and behaving animals, the odor responses of most MCs vanish on the background of this high spontaneous activity (Adrian, 1950; Kay and Laurent, 1999; Rinberg et al., 2006). By contrast, in anesthetized animals, the responses are vigorous and dense, at least in the case of ketamine/xylazine anesthesia (Rinberg et al., 2006). Consequently, many MCs lose their responses to odorants when the effects of anesthesia are removed; this suggests that, in awake animals, these cells ignore their odorant-related inputs from the receptor neurons (Rinberg et al., 2006). Thus, in this paper, we ask how MCs can disregard their odorant-related inputs despite receiving substantial inputs from receptor neurons.

Another form of odorant representation by MCs is temporal code (Brody and Hopfield, 2003; Hopfield, 1995). In this coding scheme, MCs respond to odorants by forming ensembles of cells with transiently synchronized action potentials. The identities of the synchronized cells carry information about the stimulus. Recent observations by Cury and Uchida (2010) and Shusterman et al. (2011) demonstrate the potential importance of fine time scales in odor coding. A large fraction of MCs appeared to respond with sharp and temporally precise firing events. These transients are synchronized with the temporal phase of the respiration cycle and occur in a larger fraction of MCs than previously reported on the basis of sustained combinatorial code (Rinberg et al., 2006). It could be argued that the representation of odorants by these cells is temporally sparse; i.e., they respond with transient events that occupy a small

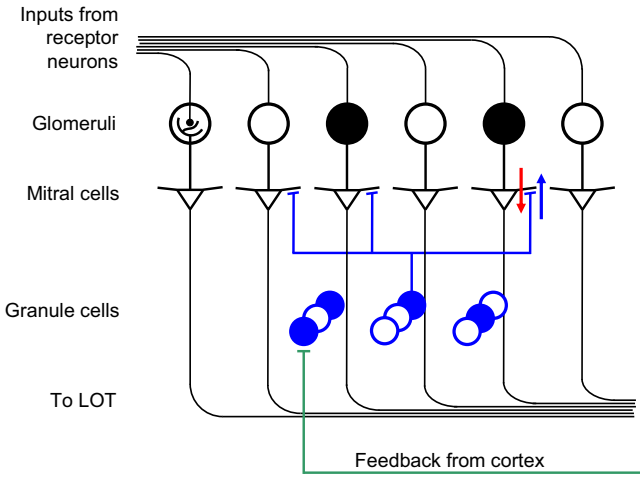

Figure 1. MC-GC Network Model

MCs (triangles) receive excitatory inputs from glomeruli (large circles). Active/inactive glomeruli are shown by the full/empty circles, respectively. The combinatorial pattern of glomerular activation represents the olfactory stimulus. The MC output is sent to the downstream parts of the brain for further processing via the lateral olfactory tract (LOT). MCs and GCs (blue circles) are connected by reciprocal dendrodendritic synapses, shown only for one GC. The GCs receive the excitatory inputs from the MCs (red arrow), and MCs are inhibited by GCs (blue arrow).

part of the respiration cycle. Although the relative importance of combinatorially and temporally sparse codes is unclear, the question that emerges from these studies is how two types of code can coexist in the same network.

In computational studies, progress has been achieved in understanding of how sparse codes can be generated by neural networks. It was shown that the recurrent network of inhibitory neurons can represent its inputs by sparse codes (Rozell et al., 2008). Understanding of these behaviors has become possible due to the approach based on Lyapunov function (Seung et al., 1998). While these models (Rozell et al., 2008) may provide neuronal mechanisms for sparse codes, they rely on the assumption that feedforward and feedback synaptic weights should satisfy a specific relationship. It is not clear how this condition is implemented biologically.

While MCs form the representation of odorants, their number is significantly smaller than the number of local inhibitory interneurons, granule cells (GCs), which play an important role in the network interactions. These cells are thought to implement lateral inhibition between MCs through a mechanism based on dendrodendritic reciprocal synapses (Figure 1) (Shepherd et al., 2004). Such interactions facilitate discrimination between similar stimuli and mediate competition between coactive neurons (Arevian et al., 2008). In agreement with this idea, facilitating inhibition between MCs and GCs improves performance in complex but not in simple discrimination tasks (Abraham et al., 2010).

In this paper, we study the mathematical model of olfactory bulb. Using this model, we address a series of questions about the responses of MCs to odorants. How can sparse combinatorial code emerge as a result of network activity in the olfactory bulb? That is, how can MCs disregard the inputs from receptor neurons? How can transient (i.e., temporally sparse) activity be generated by the same network? What is the role of network architecture of the olfactory bulb based on dendrodendritic synapses? How can olfactory code be state dependent, and is there a way to control the responses of MCs in a task-dependent manner?

To answer these questions, we propose a novel role for olfactory bulb GCs. We show that GCs can form representations of olfactory stimuli in the inhibitory inputs that they return to the MCs. MCs transmit to the olfactory cortex the errors of these representations. An exact balance between excitation from receptor neurons and inhibition from the GCs eliminates odorant responses for some MCs; however, other MCs retain the ability to respond to odors due to the incompleteness of the GCs' representations. This function is facilitated by the network architecture based on dendrodendritic reciprocal synapses between the MCs and the GCs. In this architecture, both feedforward and recurrent connections for the GCs are mediated by the same synapses, thus making biologically plausible the specific relationship between feedforward and feedback synaptic weights necessary for the existence of sparse coding in the current mathematical models (Rozell et al., 2008). Our model also predicts transient responses of silenced MCs, thus accounting for the recent data on fine temporal structure in responses to odorants.

RESULTS

Here, we describe the behavior of two ensembles of neurons: MCs and GCs (Figure 1). MCs receive inputs from olfactory receptor neurons through excitatory synapses located in glomeruli. The MC outputs are sent to olfactory cortices for further processing. Our purpose is therefore to understand the relationship between MCs' glomerular inputs and their outputs in the presence of GC inhibition. We first show several qualitative results for the model of the olfactory bulb with only a few neurons. Later, we analyze a more formal mathematical model.

GCs are inhibitory interneurons that are much more abundant than MCs (Egger and Urban, 2006; Shepherd et al., 2004). GCs and MCs form reciprocal dendrodendritic bidirectional synapses (Shepherd et al., 2007). MC firing produces excitatory inputs into GCs, which provide feedback inhibition to the MCs. Because excitation and inhibition are localized to the same synapse, and to simplify our model, we assume that the synaptic strengths in both directions are proportional (see Experimental Procedures for further discussion of this approximation).

GCs Can Contribute to the Exact Balance between Excitation and Inhibition in the Input of the MCs

We first address the behavior of the bulbar network with only a single GC present (Figure 2). If the combined excitatory input received by the GC is not sufficient to drive it above the firing threshold, then the firing of MCs will be unaffected by the presence of the GC and will reflect the excitatory inputs received from the receptor neurons (Figure 2A). A more interesting regime occurs when the MCs drive the GC above the threshold for firing (Figure 2B). In this case, the GC will produce inhibitory inputs into the MCs that can substantially modify their odorant responses.

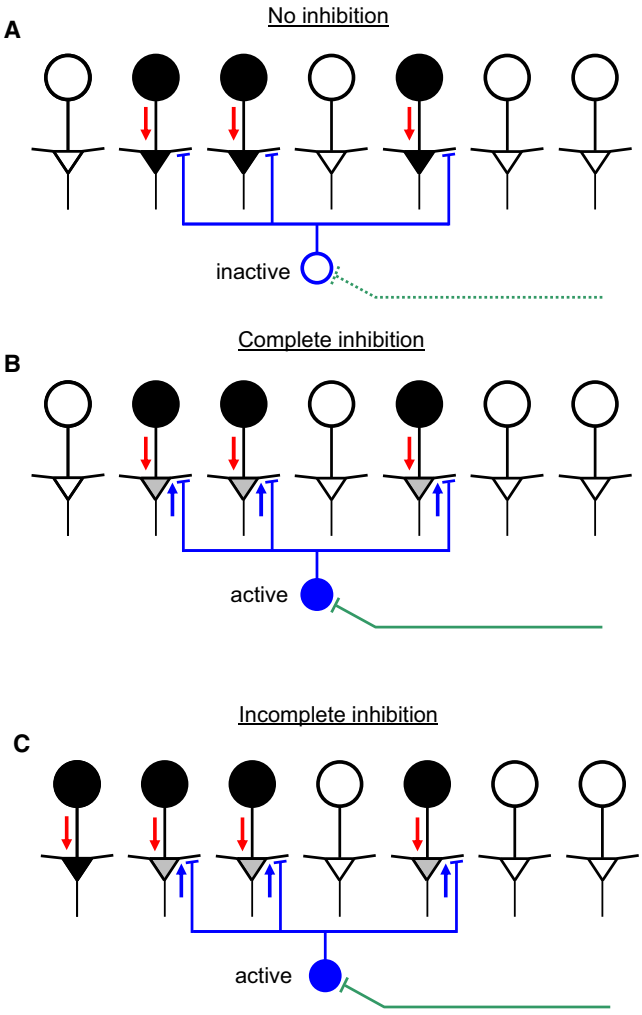

Figure 2. Balance between Excitation and Inhibition

The olfactory bulb network includes only one GC for simplicity.

(A) If activity of MCs is weak, they do not drive the GC above firing threshold. The inputs from receptor neurons (red arrows) determine the activity of MCs directly.

(B) If the inputs from MCs to the GC are strong enough to activate it, for example, due to additional inputs from the cortex (green), the GC substantially reduces the activation of the MCs through the inhibitory feedback (blue arrows) provided by the reciprocal dendrodendritic synapses, thus implementing a balance between excitation and inhibition for each MC. The activity of MCs (gray) is reduced to the level sufficient to drive the GC. In the limit of a large number of MCs, their response to odorants becomes weak.

(C) Strong responses of some MCs are possible if the pattern of inhibition is incomplete. Only the leftmost MC responds substantially, despite receptor inputs into the other cells. The MC responses are sparse, similar to those observed experimentally in awake behaving animals.

Indeed, according to our assumption, the synaptic strength between MCs and GCs is proportional. This means that the same subset of MCs that is excited by the receptor inputs may be inhibited by GCs (Figure 2B). The inhibitory feedback provided by the GC can substantially compensate for the excitatory inputs from receptor neurons, leading to a nearly exact balance between excitation and inhibition in the inputs of the MCs.

To understand the conditions for balance, consider the case when inhibitory weights from the GC to MCs are very strong. In this case, the GC will suppress any activity of the MCs that leads to the GC exceeding its firing threshold θ. The combined inputs to the GC from MCs will therefore barely exceed the GC firing threshold θ. If many MCs drive the GC, the increase in the firing rate of individual MCs needed to activate the GC is approximately given by θ / K (see Equation 16 in Experimental Procedures), where K is the number of MCs contributing to the excitatory input of the GC ($K = 3$ in Figure 2B). When more MCs are connected to a given GC (larger K), a smaller increase in activity of MCs is sufficient to activate the GC. For a large number of synapses K, the odorant-related increase in activity of a particular MC is practically undetectable. Thus, despite receiving substantial excitatory input from receptor neurons, the MCs have a very small response (defined here as the change in the MC firing rate with respect to spontaneous baseline activity). The reduction in the response is due to the strong inhibition provided by the GC. These inhibitory inputs almost completely cancel the excitation provided by the receptor neurons (see "The State Dependence of the MC Code" in Experimental Procedures for a more quantitative analysis).

Combinatorially Complete versus Incomplete Compensation

The balance between excitation and inhibition has implications for olfactory code carried by the MCs. For the MCs that receive inhibitory inputs from the GC, the odorant responses are substantially reduced. If all MCs receive these inhibitory inputs, only weak (i.e., undetectable) activity that is necessary to drive GC above the firing threshold remains (Figure 2B). Because the inputs to all of the MCs are substantially balanced by inhibition and none of the MCs displays strong odorant responses in this case, we call this complete combinatorial compensation. On the other hand, as shown in Figure 2C, the responses of a subset of MCs may accurately reproduce the inputs that they receive from the receptor neurons. This is because these cells do not receive the compensating inhibition from the GC, which is therefore incomplete. Inhibitory inputs from GCs selectively reduce the responses of some MCs, while leaving other MCs responsive. The sustained combinatorial representation carried by MCs becomes sparse. Therefore, our model can yield sparse sustained MC responses observed experimentally (Rinberg et al., 2006).

Sparsening of the responses of MCs reduces redundancy in the representation of odorants (Figure 3A). The role of GCs in this case is to remove overlaps between combinatorial receptor inputs. The removal of overlaps makes MC activation patterns more independent for different odorants. Redundancy reduction may occur in a task-dependent manner. This means that the particular overlap that is removed depends on the activation of the centrifugal cortico-bulbar projections (Figure 3A versus Figure 3B). By activating/deactivating the particular subsets of GCs, these projections may change the MC code to better discriminate the set of odorants relevant to specific behavior.

GCs Are Less Active in the Anesthetized State

Figure 2A illustrates the regime when GCs are never or rarely active. The MC code in this case is dense and reflects glomerular

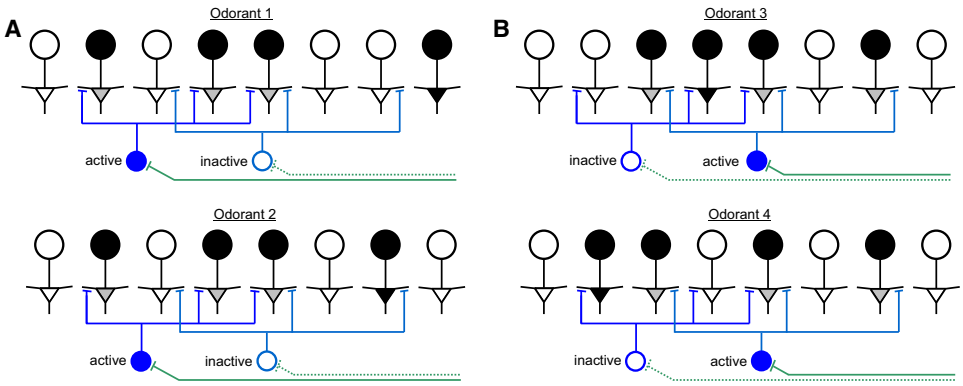

Figure 3. Redundancy Reduction in the MC Code

(A) Two strongly overlapping glomerular activations lead to dissimilar patterns of MC activity.

(B) If a different GC is brought close to firing threshold by the cortical inputs, the overlap is eliminated for a different pair of odorants.

inputs. Activation of GCs, as shown in Figure 2B, leads to sparse odorant representations. Because the transition between full and sparse codes occurs upon transition between anesthetized and awake states, we suggest that Figures 2A and 2B illustrate these two regimes of the bulbar network. The prediction of this model is therefore that GCs are less active in anesthetized animals than in awake and behaving animals.

GCs Form Representations of MC Inputs by Balancing Excitation and Inhibition

We now consider the case of several GCs. The simplest case is when the dendritic fields of the GCs (i.e., synaptic connection patterns) do not overlap, as shown in Figure 4A. Two GCs in this figure provide inhibitory feedback to two nonoverlapping sets of MCs. This feedback can balance excitation with inhibition for the subset of MCs, similarly to that of the single GC case (Figure 2). By providing inhibitory inputs to the MCs, GCs represent the combinatorial glomerular inputs by decomposing them into a set of simpler patterns contained in the dendrodendritic synaptic weights. The result of such a representation is contained in the pattern of inhibitory inputs returned to the MCs by the dendrodendritic synapses. The accurate representation of odorant-related inputs by the GCs leads to the reduction of activity of MCs due to the balance between excitation and inhibition.

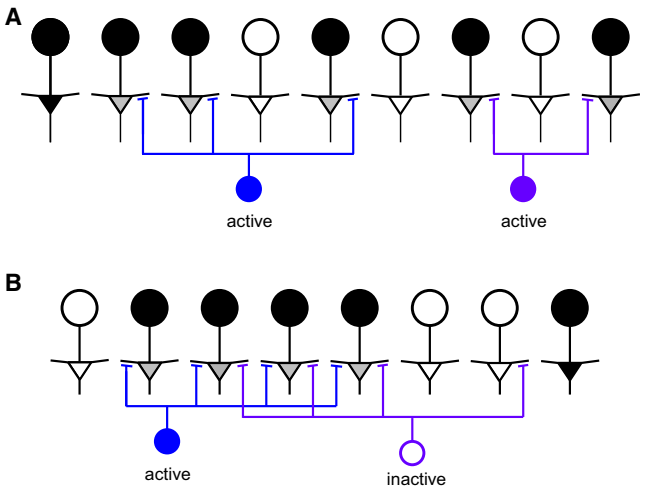

Figure 4. GCs Form Representations of MC Glomerular Inputs

(A) Several GCs may form a more complete representation.

(B) GCs compete for better representation. The cell with a larger overlap with the glomerular inputs (left) can remove the excitation from the inputs of the GC with smaller overlap (right), rendering this cell inactive.

GCs Compete for a Better Representation of the MC Glomerular Inputs

If the dendritic fields of two GCs overlap, only the GC whose pattern of connectivity better matches the pattern of MC activation becomes active, suggesting that GCs compete with each other for inputs from MCs. The nature of the GC competition is in their second-order inhibitory connectivity. Indeed, because GCs inhibit MCs, while the latter excite other GCs through the dendrodendritic synapses, the GCs, in effect, inhibit each other. This leads to GCs competing for the most complete representation of the MC inputs. Thus, the GC with the largest overlap with the glomerular input cancels the excitatory inputs into GCs with smaller overlap, rendering them inactive (Figure 4B). In Experimental Procedures, we prove that in the stationary state, i.e., after all activity patterns have stabilized, the number of coactive GCs cannot exceed the number of MCs (theorem 2; see "The Number of Coactive GCs" in Experimental Procedures). Because the number of GCs substantially exceeds the number of MCs, this statement implies that only a small fraction of GCs is coactive. This means that the GC code is also sparse.

Because of pressure to reduce the number of coactive GCs and their tendency to produce the most accurate representation (with the largest overlap), GCs form representations of the odorants that are parsimonious, i.e., the most simple and accurate. However, even the most accurate representations may be imprecise or incomplete, which is necessary for the observation of substantial MC responses (Figure 2C). In the case of many GCs, the conditions for incompleteness can be examined quantitatively with the use of the approach based on the Lyapunov function, which is described in the next section.

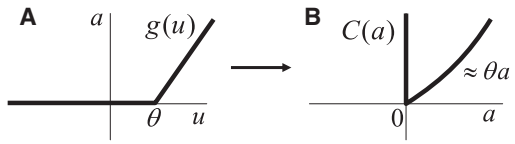

Figure 5. GC Individual Cost Function *C*(*a*)

(A) The input-output relationship is shown (*u* and *a* are the GC's input current and the resulting firing rate).

(B) The cost function prohibits negative firing rates. For positive firing rates, the cost function is approximately linear and proportional to the firing threshold θ.

The System of MCs and GCs Connected by Reciprocal Dendrodendritic Synapses Can Optimize Representation in Terms of Accuracy and Simplicity

In Experimental Procedures, we show that the dynamics of bulbar network can be viewed as a gradient descent (minimization) of the cost function called the Lyapunov function. Minimization of the Lyapunov function describes optimization of GC representations in terms of both their accuracy and their simplicity. The Lyapunov function is a standard construct in neural network theory that has been extensively used to study the properties of complex networks (Hertz et al., 1991). The dynamics of neural responses can be viewed as a descent on the landscape described by the Lyapunov function. As a corollary of this property, the neural activities' steady states can be obtained as minima of this function. The Lyapunov function, however, does not represent any form of metabolic cost or chemical binding energy. While minimization of this function may sometimes lead to sparse neural codes (Rozell et al., 2008), which often obey the principle of energy minimization (Olshausen and Field, 2004), the Lyapunov function itself cannot be related to any physical form of energy. Instead, the only definition of the Lyapunov function is that it is minimized by the dynamics of the neural network. By studying the Lyapunov function, it is possible to understand both the steady-state responses of MCs and GCs to odorants and the transitional nonstationary regime.

In our model, the inputs and outputs of MCs are defined by variables x_m and y_m, where m is the MC index that runs between 1 and M, the total number of MCs. The activities of GCs are represented by a_i, where index i is restricted between 1 and N, the total number of GCs. The synaptic weights for dendrodendritic synapses are W_{mi} and $\tilde{W}_{im}$ for GC-to-MC and MC-to-GC connectivity, respectively. This network can be described by the Lyapunov function if the two weights describing the same dendrodendritic synapse are proportional to each other:

$$W_{mi} = \varepsilon \tilde{W}_{im} \qquad \text{(Equation 1)}$$

This condition means that for different dendrodendritic synapses within the same olfactory bulb, the ratio of two opposite synaptic strengths is the same and equal to ε. Equation 1 is sufficient but not necessary for the existence of the Lyapunov function. A more general condition is described in Experimental Procedures.

The Lyapunov function for the MC-GC network has the following form:

$$\mathcal{L} = \frac{1}{2\varepsilon}\sum_{m=1}^{M}\left(x_m - \sum_i W_{mi}a_i\right)^2 + \sum_{i=1}^{N} C(a_i). \qquad \text{(Equation 2)}$$

The first term in the cost function contains the sum of the squared differences between the excitatory inputs to the MCs from receptor neurons x_m and the inhibitory inputs from the GCs. The inhibitory inputs are proportional to the activity of GCs a_i weighted by the synaptic matrix W_{mi}. The first term therefore reflects the balance between excitation and inhibition on the inputs to MCs. Minimization of this term leads to the establishment of an exact balance between excitation from receptors x_m and inhibition from GCs $\sum_i W_{mi}a_i$.

The first term in the Lyapunov function also describes the error in the representation of the odorant-related inputs by the GCs. Indeed, the inhibitory inputs returned to the MCs by the GC, $\tilde{x}_m = \sum_i W_{mi}a_i$, contain GC representations of the inputs that MCs receive from receptor neurons x_m. The first term in the Lyapunov function is proportional to the sum of the squared differences between actual glomerular inputs, x_m, and the GC approximation of these inputs, $\tilde{x}_m$. Minimization of this difference leads to the most accurate approximation of the receptor neuron activity by the GCs (i.e., to $\tilde{x}_m$ approaching x_m).

The second term of the Lyapunov function represents a cost imposed on the GC firing; i.e., the simplicity or parsimony constraint. The cost is imposed individually on each neuron and is independent of network connectivity, as follows from the form of Equation 2. The individual cost $C(a)$ is dependent on the GC input-output relationship (Figure 5) and is defined by Equation 9 in Experimental Procedures. Here, we mention two important features of this cost. First, the cost function becomes infinitely large for negative values of firing rates, thus prohibiting the firing rate to fall below zero (Figure 5B). Second, for the positive levels of activity, the cost raises approximately linearly (Figure 5B). This behavior can be traced to the membrane leak current, which, in the absence of other factors, forces the firing rates of neurons to zero. Thus, minimization of the second term in the Lyapunov function leads to minimization of the activity of GCs and can be viewed as the implementation of the simplicity constraint. The GC representation of the odorants, as described by the Lyapunov function, is subject to two conflicting constraints: those of accuracy (first term) and those of simplicity (second term).

GCs Represent MC Inputs as a Linear Sum of Dendrodendritic Weights by Balancing Excitation and Inhibition for the MCs

The set of numbers MC glomelular inputs x_m can be combined into an M-dimensional column-vector $\vec{x}$. The error $\vec{r}$ in the GC representation is given by

$$\vec{r} = \vec{x} - \sum_i \vec{W}_i a_i. \qquad \text{(Equation 3)}$$

Here, $\vec{W}_i$ is the vector containing the synaptic weights of a GC number i onto all of the MCs. Minimization of the Lyapunov function means minimization of the length of vector $\vec{r}$ with the constraints. Therefore, the olfactory bulb, through the dynamics

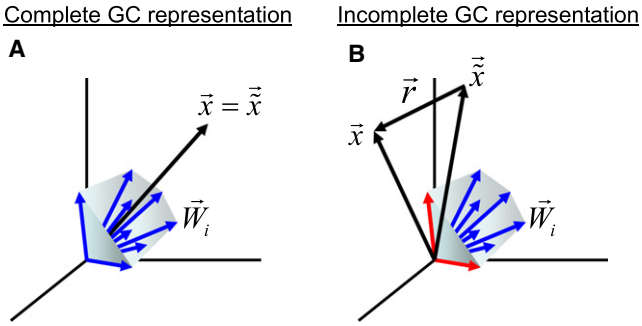

Figure 6. Nonnegativity of the GC Firing Rates Leads to Incomplete Odorant Representations

The olfactory bulb includes three MCs (3D input space) and eight GCs. The inputs into MCs from receptor neurons are represented by the black arrows. The synaptic weights from eight GCs to the MCs are shown by the blue arrows.

(A) If inputs lie within the cone restricted by the weight vectors, the input vector can be represented as an exact superposition (linear sum) of the weight vectors with positive coefficients. The GC representation of MC inputs is exact, and the MCs are expected to respond weakly due to an exact balance between inhibition and excitation.

(B) If the input vector is *outside* the cone of the GC weights, the GCs cannot represent the inputs exactly. This is because the GC firing rates (coefficients of expansion) cannot be negative. The best approximation of the inputs $\vec{x}$ is the nearest vector to the input vector $\vec{x}$ on the surface of the cone. $\vec{x}$ is formed by two GCs (red). Because $\vec{x}$ is different from $\vec{x}$, the GC code is *incomplete*. MC response is $\vec{r} = \vec{x} - \vec{x}$.

of GCs, attempts to represent vector $\vec{x}$ as a superposition (linear sum) of weight-vectors:

$$\vec{x} \approx \sum_i \vec{W}_i a_i. \qquad \text{(Equation 4)}$$

The activities of GCs a_i represent the coefficients with which each weight vector contributes to the representation. If the representation is perfect, MCs receive no odorant-related inputs; i.e., inhibition and excitation for each MC are perfectly balanced.

MCs Transmit to the Olfactory Cortex the Error of Representation by the GCs of the Odorant-Related Sensory Neuron Activation

The odorant-dependent MC response is the difference between excitatory inputs x_m received by an MC number m and the inhibition from GCs; i.e., $x_m - \sum_i W_{mi} a_i$. The same difference defines the error in the representation of MCs' glomerular inputs by the GCs given by Equation 3. Thus, the MCs transmit to the olfactory cortex the error of representation of the olfactory inputs by the GCs. The dynamics of the bulbar network, by minimizing the error of representation, minimizes the odorant-dependent responses of MCs. In our model, MC odor responses, defined as the deviation from the baseline firing rate, can be both negative and positive. Negative errors occur when, for example, a MC does not receive any input from receptor neurons but is inhibited by a GC (Figure 7A, blue triangle). The total firing rate of each MC is therefore equal to $y_m = \bar{y}_m + r_m$, where $\bar{y}_m$ is the spontaneous firing rate, while r_m is the error of GC representation defined by Equation 3, which is also the response of the MC to odorant.

To satisfy Equation 4, vector $\vec{x}$ should be represented as a sum of N basis vectors $\vec{W}_i$ with coefficients a_i. The latter are unknown coefficients, which correspond to the activity of GCs. The number of unknowns is therefore equal to the number of GCs. The number of conditions that the unknowns have to satisfy is equal to the number of MCs for which the balance between excitation and inhibition is sought. Because there are more unknowns (GCs) than constraints (MCs), the problem specified by Equation 4 has many solutions; i.e., it is overcomplete. This means that several combinations of firing patterns of GCs are consistent with any given set of glomerular inputs $\vec{x}$. This also means that, in the absence of constraints, GCs can accurately represent any MC input. Thus, constraints imposed by the second term in the Lyapunov function lead to inaccurate representations of the glomerular inputs by the GCs and, consequently, nonvanishing responses of the MCs to odorants.

Constraints on GC Firing Rates Lead to Incomplete GC Representations and Detectable Responses of the MCs

If GCs represent the excitatory inputs of the MCs exactly, then MCs are unresponsive as a result of the exact balance between excitation and inhibition. Therefore, GCs must fail to represent glomerular inputs; i.e., the GC code must be incomplete. Among several reasons for the incompleteness of the GC code, the nonnegativity of the GC firing rates is the most straightforward. Indeed, any M-dimensional vector (MC receptor input) can be accurately represented as a linear sum of M independent vectors (GC-to-MC synaptic weights) if the coefficients in this representation (GC firing rates) are allowed to be both positive and negative. This is certainly true if more than M basis vectors are available (i.e., N >> M). However, because the GC firing rates cannot be negative, the representation cannot be always performed accurately, which leads to substantial MC responses.

The impact of nonnegative GC firing rates is illustrated for the olfactory bulb network with three MCs and eight GCs (Figure 6). A particular odorant input is shown as a vector $\vec{x}$ in three-dimensional space, where each dimension corresponds to the receptor input to one of the MCs. Each GC is shown by the blue basis vector (Figure 6A). The number of basis vectors is given by the number of GCs; e.g., eight in this example. The components of each basis vector determine the strength of the inhibitory dendrodendritic synapse from a given GC to all of the MCs. The olfactory bulb is therefore expected to represent the input vector $\vec{x}$ as a superposition of eight synaptic weight vectors $\vec{W}_i$ $(i = 1..8)$. The input vectors within the convex cone enveloping the weight vectors can be obtained from the weight vectors by mixing them with positive coefficients (GC firing rates). The input vectors outside the cone cannot be represented by the GC with positive coefficients. When the input vector $\vec{x}$ is outside of this cone (Figure 6B), the network attempts to minimize the error of representation by finding the nearest vector within the boundaries of the cone $\vec{x}$. This nearest vector forms the imperfect representation of the odorant by the GC. The difference between GC representation and real vector of inputs $\vec{x} - \vec{x}$ is the error of the representation; i.e., MC odorant response $\vec{r}$ that is transmitted to the olfactory cortex.

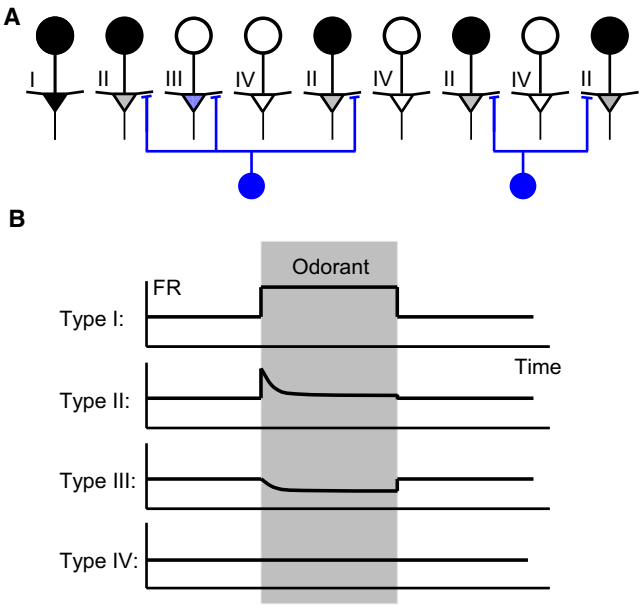

Figure 7. The Transient Responses of MCs

(A) Four classes of MCs can be identified depending on their relationship to the GC connectivity and glomerular inputs. These classes include cells that receive glomerular inputs but are not inhibited by the GC set corresponding to given odorant (I), cells in which excitation is balanced by inhibition almost precisely (II), cells that receive inhibitory inputs only (III), and cells that receive no inputs (IV). Because responses of MCs are defined with respect to their spontaneous activity, they can be negative, as indicated by the color of the type III cell.

(B) Firing rates (FR) of cells in four classes as a function of time after a sharp odorant onset (gray region). Type II cells exhibit sharp activity transients that decay to undetectable levels as a function of time.

The Number of Coactive GCs Is Less than the Number of MCs

In the case of incomplete representations (Figure 6B), the GCs encode a point $\tilde{x}$ on the boundary of the enveloping cone. Not all GCs are simultaneously active. Indeed, in Figure 6B, only two GCs on the boundary (red weight vectors) are active, while the others contribute to the representation with zero coefficients (firing rates). The number of coactive GCs is one less than the dimensionality of the input space determined by the number of the MCs M. Thus, in Figure 6B, two GCs are contributing to the representation. In Experimental Procedures, we prove that the number of coactive GCs in the model described is less than the number of MCs. Because the number of GCs in the olfactory bulb is substantially larger than the number of MCs, only a small fraction of the GCs is coactive. Therefore, our model predicts sparse responses of GCs.

Incomplete GC Representations Are Typical for Random Network Weights

For a large number of MCs and a random set of network weights, the representations of odorants by GCs are typically incomplete (see Supplemental Information available online). Hence, for a large network, the region inside of the cone of completeness (see Figure 6) is expected to shrink. This implies that it becomes almost impossible to expand a random input vector to the basis containing vectors with positive components by using only nonnegative coefficients. In the Supplemental Information, we show that the number of coactive GCs for random binary inputs with M MCs is $\sim\sqrt{M}$. Because an exact representation of the M-dimensional random input requires M vectors, this result implies that the representation of odorants by GCs is typically imprecise. The GC code is therefore incomplete. We also show that for sparse GC-to-MC connectivity, when only $K<<M$ weights are nonzero, the number of coactive GCs is somewhat larger, $\sim M/\sqrt{K}$. This number is smaller than what is required for a complete GC code; i.e., M. Therefore, GCs cannot represent MC inputs precisely in the case of random connectivity, which implies ubiquity of incomplete representations.

Combinatorial versus Temporal Sparseness

So far, we have discussed the responses of MCs and GCs in the stationary state established after odorant onset. We found that the responses of MCs may be *spatially* or *combinatorially* sparse in the steady state. This means that a small fraction of MCs carries sustained responses to odorants. Here, we address the responses during the transitional period immediately following the odorant onset. Within this model, many MCs should display sharp activity transients that are followed by exponentially decaying responses. The responses of most of the MCs are *temporally* sparse. This means that the MCs display vigorous responses only within a small fraction of time during odor presentation.

The Model Predicts Vigorous Transient Responses of MCs during the Transitional Regime

Because GC inhibition arrives at MCs with a delay, the spatially sparse MC responses are not expected to form immediately after odorant onset. During a brief initial period, the receptor neuron inputs affect MC responses directly, without strong inhibition from the GC. This observation leads to two conclusions. First, the initial responses of MCs during odor presentation are not sparse. Until inhibition from the GC arrives, MC responses reflect the pattern of receptor neuron inputs directly and are less sparse and more vigorous, as in the anesthetized state. Second, due to the small time constant of inhibition, the initial vigorous responses are suppressed quickly by the GC. As a result, for some MCs, the odorant responses display transients synchronized with the odorant onset (Figure 7B, Type II cells). Within this model, the transients have an exponential shape with the time constant $\tau=\tau_0/K$, where K is the number of synapses per GC and τ_0 is the time constant related to the synaptic delay ($\tau_0=\tau_d/g'W\tilde{W}$, where τ_d is the synaptic delay; see Supplemental Information for a full description of the transient regime). When the number of synapses K is large, the shape of transients becomes very sharp and is controlled mostly by the precision of odorant delivery to the receptor neurons. Our model therefore predicts temporarily sparse responses for most MCs. To be observed, the sharp transient responses have to be aligned precisely with the odorant onset.

DISCUSSION

Sparse Overcomplete Representations

Sparseness in neural networks emerges in the theory of sparse overcomplete representations (Olshausen and Field, 1996;

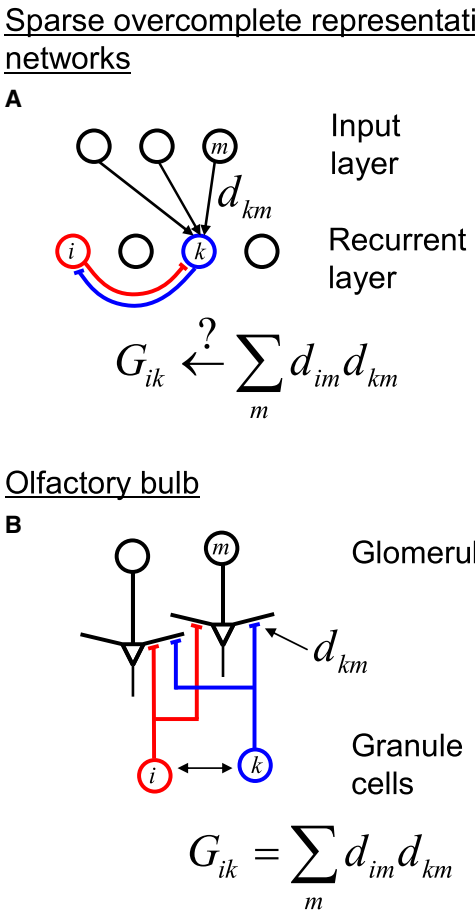

Figure 8. Possible Functional Significance of Dendrodendritic Synapses

(A) Networks of inhibitory neurons implementing sparse overcomplete representations decompose the input pattern into the linear sum of dictionary elements (Rozell et al., 2008). The response of each cell in the recurrent layer indicates whether the given dictionary element is present in the input. The feedforward excitatory weights of a recurrent layer neuron number k contain the dictionary element represented by this neuron d_{km}. The recurrent inhibitory weights are proportional to the overlap between dictionary elements, making similar input patterns compete more strongly. The biological implementation of this constraint is unclear.

(B) In networks with dendrodendritic synapses, the connectivity between GCs is inhibitory. Because every two GCs are two synapses away from each other, the strength of pairwise inhibition contains an overlap between dendrodendritic weights, thus automatically satisfying the constraint on the recurrent inhibitory weights shown in (A).

Olshausen and Field, 2004; Rozell et al., 2008). According to these models, a sensory input can be decomposed into a linear sum of primitives called dictionary elements. The decomposition is sought in the form of a set of coefficients with which different dictionary elements contribute to the input. These coefficients represent the responses of neurons in a high-level sensory area, such as the visual cortex, that indicate whether a given feature is present in the stimulus. Because the number of dictionary elements available is usually quite large, several decompositions are consistent with the given input. That is why this representation is called overcomplete. To make representation unambiguous, some form of the parsimony principle is added to the model in the form of a cost function on the coefficients/responses. The solution that yields the minimum of the cost function is assumed to be chosen by the nervous system. The decomposition is found to be dependent on the cost function. The general form of a cost function is a sum of firing rates in power α: $L_\alpha = \sum_i |a_i|^\alpha$. For L_2, the simple sum of squares of the coefficients, all neurons generally respond to any stimulus, and, therefore, the code is not sparse. For L_1, when the cost is equal to the sum of absolute values of the responses, the solution is found to be sparse, which means that most coefficients/responses are exactly zero. Thus, the L_1 norm is called sparse, and the corresponding neural representation is sparse overcomplete.

It was shown that the recurrent network of inhibitory neurons can implement sparse overcomplete representations (Rozell et al., 2008). To show this, the network dynamics are represented as a minimization of a cost function called the Lyapunov function, similarly to the representation of Hopfield networks (Hertz et al., 1991; Hopfield, 1982). Hopfield networks have attractor states that contain memory of activation patterns. In contrast to Hopfield networks, in purely inhibitory networks, the recurrent weights enter the Lyapunov function with a minus sign, which abolishes the attractor memory states and makes the network purely sensory (Rozell et al., 2008). Minimization of Lyapunov function in realistic recurrent networks with inhibition was suggested as a means to implement the parsimony constraint (L_1) mentioned above.

To implement sparse overcomplete representations with realistic networks of neurons, two requirements have to be met (Rozell et al., 2008). First, the feedforward weights between the input layer of the network and the inhibitory neurons have to contain the dictionary elements (Figure 8A). This ensures that inhibitory neurons representing a particular dictionary element will be driven strongly when it is present in the input, due to a high overlap between the stimulus and the feedforward weights. Second, the recurrent inhibitory weight between any pair of neurons has to be proportional to the overlap between their dictionary elements (Figure 8A). This feature implies that similarly tuned inhibitory neurons compete more strongly. Therefore, the two types of network weights, feedforward and recurrent, have to closely match each other, one of them constructed as the overlap of the other. Here, we suggest that the olfactory bulb network architecture based on dendrodendritic synapses can ensure that the feedforward and recurrent connectivity are closely matched. In the architecture based on dendrodendritic synapses, both the feedforward weights received by the GCs and their recurrent connections are dependent on the same set of synapses. Similar architectures have been proposed for analysis-synthesis networks (Mumford, 1994; Olshausen and Field, 1997).

Dendrodendritic Synapses Facilitate the Representation of Odorants by the GCs

The GCs of the olfactory bulb receive excitatory inputs from the MCs through dendrodendritic synapses (Shepherd et al., 2004). These synapses encode patterns that can strongly drive

individual GCs. The effective connectivity between GCs is inhibitory (GC-to-MC and MC-to-GC synapses are inhibitory and excitatory, respectively). To calculate the strength of mutual inhibition, one has to calculate the sum over intermediate synapses, which leads to the evaluation of a convolution or overlap between GC input weights (Figure 8B). Therefore, GCs inhibit each other more strongly if their inputs from MCs are more similar. Thus, both requirements necessary to implement sparse overcomplete representations are met. This implies that the function of GCs is to detect specific patterns of activity in the inputs that MCs receive. The GCs then are capable of building representation of MC inputs. The parsimony of representation is ensured by the mutual inhibition between GCs, with more similar GCs inhibiting each other more strongly. The latter condition is facilitated by the network architecture based on dendrodendritic synapses. This observation provides a potential explanation for the existence of these synapses (Shepherd et al., 2004).

MCs Respond to Odorants, because the GC Code Is Incomplete

Two problems emerge if we assume that GCs implement sparse overcomplete codes. First, GCs are interneurons and, as such, cannot directly transmit their representation to the downstream network. The significance of these representations becomes unclear. Second, if GCs indeed establish an absolutely accurate representation of their inputs, the MCs will respond to odorants very weakly. This is because GCs can eliminate these responses from the MCs' firing by balancing receptor neuron inputs with inhibition. These considerations suggest that the representations by the GCs are incomplete; i.e., that GCs cannot find accurate representations of their inputs, for example, because this would require their firing rates to become negative. If GCs' codes are incomplete, the MCs transmit only the unfinished portion of the representation to the downstream olfactory networks. As a consequence, the MCs' odorant representations become sparse. The redundancies in the MC codes are reduced, and the overlaps in representations of similar odorants are erased, yielding more distinguishable responses to similar odorants.

Several factors may contribute to the incompleteness of GC representations. Here, we analyzed the nonnegativity of the GC firing rates as one possibility. In addition, we argue in Experimental Procedures that the high threshold for GC activation can hinder accurate representation of odorants by these cells and suggest that the increase in GC activation threshold may contribute to less-sparse responses of MCs in the anesthetized state. In addition, if the ensemble of GCs available is small, the set of combinations represented by them may be limited, leading to incomplete representations. Finally, inhibitory inputs to MCs cannot be represented exactly by GCs without invoking a more complex network mechanism. Such inputs may arise from inhibition of the receptor neurons by some odorants (Ukhanov et al., 2010) or inhibition in the glomerular layer network (Aungst et al., 2003).

The Lyapunov Function

Lyapunov functions are standard tools in neural network theory (Hertz et al., 1991). Seung et al., 1998 have shown that the network containing two populations of neurons, inhibitory and excitatory, can be described by the Lyapunov function. This model can be related to the system of MCs and GCs. Here, we eliminate the dynamics of MCs from consideration and reduce the description to GCs only. By doing so, we show that GC could be used to extract features from MC inputs. Analysis-synthesis networks have architectures similar to our model (Mumford, 1994; Olshausen and Field, 1997). We propose that the network based on dendrodendritic synapses provides a mechanism for balancing feedforward and feedback weights that can be easily implemented biologically. Lee and Seung, 1997 studied the nonlinear network mechanism to implement sensory network with nonnegativity constraints; i.e., conic networks. Although their nonlinear network mechanism is somewhat different from the one used here, the set of encoding/error neurons described in Lee and Seung can be viewed as analogs of GCs/MCs, respectively, in our model. In our study, we propose a mapping of the conic networks onto the olfactory bulb network and study the implications of this mapping for the olfactory code.

Several Recent Studies Suggest a Complex Role for GCs in Olfactory Discrimination

First, Arevian et al., 2008 showed that inhibition between MCs is gated by postsynaptic activity so that a minimum threshold is required before lateral inhibition is engaged. The inhibition between MCs is nonlinear, which can facilitate discrimination of correlated patterns. This behavior is fully compatible with our model, in which the active GCs are selected from the large population on the basis of competitive interactions. In comparison to Arevian et al., 2008, we suggest that behavior of the bulbar network cannot be viewed as pairwise nonlinear interaction between MCs, and the description on the basis of GCs approximates the network nonlinearities with better accuracy. Fantana et al., 2008 show that MCs do not have a center-surround inhibitory receptive field. Instead, the MCs are inhibited by a small number of spatially dispersed glomeruli. Here, we suggest that the identities of the interacting MCs may be odorant or state dependent. We thus propose that the description of bulbar interactions as lateral inhibition is not fully adequate and that a model in which the network weights vary with odorants may fit the data more accurately. Our model predicts that in awake animals, effective bulbar connectivity may be less sparse. The role of bulbar inhibition in the discrimination of similar stimuli was recently studied by Abraham et al., 2010. Remarkably, it was shown that enhancement of GC inhibition affected complex but not simple discrimination tasks. Here, we propose that GCs remove overlaps in glomerular representation of similar odorants (Figure 3), leading to their larger impact on complex discriminations, which is in agreement with Abraham et al., 2010.

Several theoretical studies have proposed that GCs implement orthogonalization of stimuli on the level of the olfactory bulb (Cecchi et al., 2001). Recently, Wick et al., 2010 showed that GCs can orthogonalize the responses of MCs. Their study reduces the dynamics of the olfactory bulb network to the pairwise interactions between MCs by eliminating GCs. Here, we propose that GCs cannot be eliminated, because they build representations of MC inputs. This implies that GCs represent MC inputs in the inhibitory current returned to the MCs. As a

result, MCs transmit to the cortex errors of GC representations. The responses of GCs in our model are highly nonlinear, with most of them remaining silent. Because MCs play the role of error neurons, their sustained responses are sparse, which is a form of orthogonalization that is alternative to Wick et al., 2010.

Overlap reduction in the olfactory bulb network was previously proposed theoretically on the basis of a periglomerular network implementing surround inhibition (Linster and Hasselmo, 1997). This hypothesis was supported by enhanced generalization between chemically similar odorants by rats with strengthened periglomerular inhibition (Linster et al., 2001). We suggest a mechanism for redundancy reduction by GC inhibition that is organized functionally rather than spatially in a task-dependent manner. This proposal is consistent with nonlocal interglomerular connectivity (Fantana et al., 2008).

Feedback from Higher Brain Areas and State Dependence of the Olfactory Code

The sparseness of the MC responses depends on the nonlinearity of the GCs and, specifically, on the GC activation threshold θ. In this study, we assumed that all GCs have similar activation thresholds that are small enough for GCs to be easily activated by low levels of activity in MCs. If the thresholds for activation of individual GCs are different, it is possible to envision a mechanism by which the olfactory code carried by both MCs and GCs can be controlled to adapt to a particular task. Thus, if the threshold for activation is raised for a subset of GCs, these cells will be no longer active; therefore, their activity will not be extracted from the firing of MCs. If, for example, the threshold for all of the GCs is increased, thus making them unresponsive, then the olfactory code carried by the MCs replicates their inputs from receptor neurons. If the activation threshold is lowered for a subset of GCs, these cells will efficiently extract their activity from the MCs' responses. Thus, a particular redundancy among similar odorants can be excluded in a task-dependent manner. Therefore, the thresholds for GC activation may regulate both an overall sparseness of MC responses and the fine structure of the bulbar olfactory code.

GC excitability depends on cellular properties but can also be effectively modulated by additional input into these cells. The GCs in the mammalian olfactory bulb are recipients of the efferent projections from the cortex and other brain areas (Davis and Macrides, 1981; Luskin and Price, 1983). These signals to GCs can change their effective threshold values. If a GC receives excitatory inputs from the cortex, then the MC signal is closer to the threshold value, and the GC is more readily excited by the odorant-related inputs. These mechanisms can dynamically control the sparseness and information content of the MC signal by regulating the input to the GCs from the cortex.

Implicit experimental evidence for the presence of such regulation already exists. In the anesthetized state, when the efferent signal coming to the bulb from the cortex is minimal, the MCs respond strongly to the odor stimulation. In the awake state, when the cortex is active, the MC code becomes sparser (Adrian, 1950; Kay and Laurent, 1999; Rinberg et al., 2006). Cutting the lateral olfactory tract and eliminating feedback from the brain in an awake rabbit led to MC responses that were similar to those in an anesthetized rabbit (Moulton, 1963). Therefore, centrifugal projection may indeed regulate the sparseness of the olfactory code.

Some evidence also points toward the possibility of finer network tuning by specific activation or deactivation of subsets of GCs to enhance extraction of relevant information. First, Doucette and Restrepo, 2008 demonstrated that the MC responses to odorants change as animals learn the task. Second, Fuentes et al., 2008 showed that the number of responding MCs depends on the task. When a rat is involved in an odor discrimination task, the average number of MC excitatory responses is less than that in the rat passively smelling an odor. The assumption is that when an animal discriminates odorants, it may be advantageous to suppress redundant MC responses and enhance those that carry the most behaviorally relevant information. Our model proposes the network mechanism for this phenomenon.

Our Model Predicts Sharp Transient Responses of MCs

If the signal first appears in the inputs to MCs, it causes an elevation of the MC firing rates, which, in turn, causes activation of GCs and suppression of MCs by feedback inhibition. The sensory inputs of the MCs are therefore represented initially by transients of activity followed by decay to the steady state as described in this study (Figure 7B). This observation raises several questions. First, what information about the stimulus is sent to the cortex by transients and the consequent steady state responses? Second, what are the experimental conditions for the observations of such bursts?

While the roles of different modes of information transmission are unclear, we can make some predictions about the second question. The MC activity transients are short and stand on top of high levels of spontaneous activity. In order to observe such transients reliably, one needs to synchronize spike trains with stimulus delivery. In mammals, stimulus delivery is controlled by sniffing. In previous studies (Doucette and Restrepo, 2008; Fuentes et al., 2008; Kay and Laurent, 1999; Rinberg et al., 2006), the authors synchronized their recordings with stimulus onset but not with sniffing/breathing. This approach may lead to smearing of the short bursts and emphasize the steady state responses of the network. In such a regime, the odor responses should be combinatorially sparse as predicted by the model. New evidence by Shusterman et al., 2011 and Cury and Uchida, 2010 suggests the presence of fast bursts of MC activity synchronized with sniffing, which is reminiscent of the transient responses proposed here.

While our model predicts temporarily precise, strong, and brief MC response to odors, more complex temporal responses observed in Shusterman et al., 2011 and Cury and Uchida, 2010 are beyond accurate description by our model, mainly due to oversimplified temporal profile of the stimulus. In a more realistic situation, the sniff dynamics controls the raise of the odorant concentration at the epithelium, and different receptors are activated at different phases of the sniffing cycle.

Inhibitory Neurons in Other Systems

The mechanism suggested here could be implemented in other parts of the nervous system and in other species. Thus, the inhibitory interneurons of the insect antennal lobe (Assisi et al., 2011) could form representations of odorant-dependent inputs of

the projection neurons in a similar way. Our mechanisms could also be implemented on the basis of axon-dendrite synapses if the symmetry in the synaptic strength between excitatory and inhibitory neurons (Equation 1) is approximately satisfied. This model could therefore be implemented by cortical networks.

Conclusions

We theoretically address the experimental evidence of spatially sparse odor codes carried by MCs in awake rodents. We propose a novel role for the GCs of the olfactory bulb in which they collectively build an incomplete representation of the odorants in the inhibitory currents returned to the MCs. These inhibitory currents lead to the balance between excitation and inhibition and suppression of responses for most of the MCs. Because the representation formed by GCs is incomplete, a small number of MCs can carry information to the cortex, leading to the sparse olfactory codes. Our model predicts sharp transient responses in a large population of MCs. This function is facilitated by the network architecture that includes bidirectional dendrodendritic MC-to-GC synapses.

EXPERIMENTAL PROCEDURES

Derivation of the Lyapunov Function

Our model is based on the following equations describing the responses of MCs and GCs, r_m and a_i:

$$r_m = x_m - \sum_{i=1}^{N} W_{mi} a_i, \qquad \text{(Equation 5)}$$

$$\dot{u}_i = -u_i + \sum_{m=1}^{M} \tilde{W}_{im} r_m, \qquad \text{(Equation 6)}$$

and

$$a_i = g(u_i), \qquad \text{(Equation 7)}$$

where u_i is the membrane voltage of the GC and $r_m = y_m - \overline{y}_m$. Here, y_m and $\overline{y}_m$ are instantaneous, and average activity of the MC number m, $\dot{u} = du/dt$. We assume therefore that the dynamics of MC responses is fast enough to reflect the values of inputs instantly.

The description in terms of differential Equations 5–7 is equivalent to the minimization of the Lyapunov function:

$$\mathcal{L}(\vec{a}) = \frac{1}{2\varepsilon} \sum_{m=1}^{M} \left(x_m - \sum_i W_{mi} a_i \right)^2 + \sum_{i=1}^{N} C(a_i), \qquad \text{(Equation 8)}$$

where $C(a)$ the cost function is defined as

$$C(a) = \int_0^a g^{-1}(a')\,da'. \qquad \text{(Equation 9)}$$

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

$$\dot{\mathcal{L}} = \sum_i \frac{\partial \mathcal{L}}{\partial a_i} \dot{a}_i = \sum_i (-\dot{u}_i) \dot{a}_i = -\sum_i \left[g^{-1}(a_i) \right]' \left[\dot{a}_i \right]^2 \le 0. \qquad \text{(Equation 10)}$$

The last inequality holds if $g(u)$ is a monotonically increasing function. Therefore, the Lyapunov function always decreases if the system behaves according to Equations 5–7. Since the function is limited from below, the stable states of the system are described by the local minima of the Lyapunov function.

The cost function for the threshold-linear neuronal activation function $g(u) = [u - \theta]_+$ shown in Figure 5A can be approximated by a linear function when the firing rates of the GC are not too large:

$$C(a) \approx \theta a, \quad a \ge 0. \qquad \text{(Equation 11)}$$

For negative values of firing rates a, the cost function is infinitely big, reflecting the fact that negative firing rates are not available.

The Importance of Condition $W_{mi} = \varepsilon \tilde{W}_{im}$

For our model to have a Lyapunov function, a more general condition than $W_{mi} = \varepsilon \tilde{W}_{im}$ may hold. Indeed, the sufficient condition for the existence of the Lyapunov function is that the network weight matrix $G_{ik} = \sum_i \tilde{W}_{im} W_{mk}$ is symmetrical (Hertz et al., 1991). This is true if, for example, $W_{mi} = \tilde{W}_{in} E_{nm}$, where E_{nm} is an arbitrary symmetrical M-by-M matrix. Thus, condition $W_{mi} = \varepsilon \tilde{W}_{im}$ is sufficient but not necessary for the network to have a Lyapunov function. However, we argue that this condition is necessary for the system to be described by the Lyapunov function in the form of Equation 2.

The Number of Coactive GCs

We will distinguish two types of Lyapunov functions. First, we consider the homogeneous case, when the Lyapunov function has the following form:

$$\mathcal{L}_0(\vec{a}) = \frac{1}{2\varepsilon} \sum_{m=1}^{M} \left(x_m - \sum_i W_{mi} a_i \right)^2. \qquad \text{(Equation 12)}$$

The homogeneous case corresponds to a vanishingly small threshold for the activation of the GCs. In this case, we still constrain the firing rates to be nonnegative. The inhomogeneous Lyapunov function is

$$\mathcal{L}(\vec{a}) = \mathcal{L}_0(\vec{a}) + \sum_i \theta_i a_i, \qquad \text{(Equation 13)}$$

with the same constraint $a_i \ge 0$. We prove here the two theorems that limit the number of coactive GCs (i.e., the ones for which $a_i \ne 0$). The response of MCs is $r_m = x_m - \sum_i W_{mi} a_i$.

Theorem 1: Sparseness of the Homogeneous Solution

Assume that the M-dimensional receptive fields of the GCs $\vec{W}_i$ are the vectors of the general position; i.e., every subset of M vectors $\vec{W}_i$ are linearly independent. Then, in the minimum of the homogeneous Lyapunov function (Equation 12), either $r_m = 0$ for all m (i.e., MCs do not respond, and GC representation is complete) or fewer than M GCs are active. In the former case ($r_m = 0$), all of the GCs may be active.

Proof: Assume that M or more GCs are simultaneously active. Let us vary slightly the activity of only one active GC: $\Delta a_k = \varepsilon$. The corresponding variation in the Lyapunov function is $\Delta \mathcal{L}_0 = (\vec{r} \cdot \vec{W}_k)\varepsilon + O(\varepsilon^2)$. Because we are considering the minimum of the Lyapunov function, all of the scalar products $(\vec{r} \cdot \vec{W}_k)$ must be zero, which is possible only if $\vec{r} = 0$ or the number of vectors $\vec{W}_k$ is less than M.

Theorem 2: Sparseness of the Inhomogeneous Solution

Consider the set of N $(M + 1)$-dimensional vectors $\vec{\Omega}_i = (\vec{W}_i, \theta_i)$. Assume that these vectors are of the general position; i.e., any subset of $M + 1$ of these vectors is linearly independent. For the minimum of the inhomogeneous Lyapunov function (Equation 13), no more than M GCs can be simultaneously active.

Proof: For any active GC, the following equation is valid:

$$\frac{\partial \mathcal{L}}{\partial a_i} = -\sum_m W_{mi} r_m + \theta_i = 0. \qquad \text{(Equation 14)}$$

Assume that more than M GCs are active. Then, we have at least $M + 1$ such equations for M unknowns r_m. Such a system in the general case (if $M + 1$ corresponding vectors $\vec{\Omega}_i$ are independent) is inconsistent and has no solution. Thus, the number of coactive GCs cannot exceed M.

Note: Consider the case of small but nonzero firing thresholds of GCs θ. In this case, two regimes can be distinguished. If vector $\vec{x}$ can be expanded in terms of vectors $\vec{W}_i$ with positive coefficients, the firing rates of M GCs are generally ~ 1, but the responses of MCs are small ($\sim \theta$). This is the regime of sparse overcomplete representations. If the glomerular input vector $\vec{x}$

cannot be represented as a superposition of GCs weights $\vec{W}_i$ with positive coefficients (incomplete representation), the responses of cells are essentially (ignoring contributions $\sim\theta$) given by the solution of homogeneous problem (Equation 12), which explains our attention to this problem. In this case, according to theorem 1, fewer than M GCs have large firing rates, and only one has a small firing rate ($\sim\theta$).

State Dependence of the MC Code

Here, we suggest that the presence of a large threshold for GC firing (Figure 5A) can lead to inaccurate representations of odorants, similar to the nonnegativity of the GC firing rates (Figure 6). This observation allows us to explain the transition between the awake and anesthetized responses. We use a simplified model of a bulbar network containing only one GC (Figure 2). This network has the advantage that an exact solution can be found even when a finite threshold for firing is present for the activation of the GCs. Consider the input configuration shown in Figure 2C. Assume for simplicity that all of the nonzero weights and MC inputs have unit strengths. Then, the Lyapunov function for the activity of the single GC a is

$$\mathcal{L}(a) = \frac{K}{2}(1-a)^2 + \theta a. \qquad \text{(Equation 15)}$$

Here, we have to assume that $a \geq 0$; K is the number of nonzero weights for GCs ($K = 3$ in Figure 2). By minimizing the Lyapunov function, we obtain $a = 1-\theta / K$, for $\theta \leq K$ and zero otherwise. The activity of the first MC (Figure 2C, Figure S1) cannot be affected by the GC, because the GC makes no synapses onto this cell. The activities of MCs 2, 3, and 5 are given by

$$r_{2,3,5} = 1 - a = \frac{\theta}{K} \qquad \text{(Equation 16)}$$

for $\theta \leq K$. MCs increase their firing rate to activate the GC. The amount of increase is equal to the threshold for activation of the GC divided by the number of MCs contributing to the input current; i.e., K. The activity of the GC is assumed to rise fast above the threshold so that it suppresses all significant increases of inputs above the value given by Equation 16. The responses of MCs as functions of the threshold θ are shown in Figure S1. For large firing thresholds θ, all MCs that receive receptor inputs respond to the odorant. In this case, the GC is silent due to the high threshold and, therefore, the MCs relay the receptor neuron inputs directly. The MC code is not sparse. The case of large thresholds corresponds to the network in the anesthetized animal. In the opposite case of low GC firing threshold, the MC firing becomes sparse. This regime corresponds to the awake animal. According to this model, the transition from the awake to the anesthetized state is accomplished by an increase in the thresholds of GC firing, which could be mediated by the centrifugal cortico-bulbar projections or decrease in the spontaneous activity of MCs.

SUPPLEMENTAL INFORMATION

Supplemental Information includes one figure and Supplemental Experimental Procedures and can be found with this article online at doi:10.1016/j.neuron.2011.07.031.

ACKNOWLEDGMENTS

We thank Dmitry Chklovskii, Venkerakesh Murty, Barak Pearlmutter, Sebastian Seung, and Anthony Zador for useful discussions; Henry Greenside and Joshua Dudman for comments on the manuscript; and Aspen Center for Physics for support. A.A.K. was supported by NIH R01EY018068.

Accepted: July 28, 2011
Published: October 5, 2011

REFERENCES

Abraham, N.M., Egger, V., Shimshek, D.R., Renden, R., Fukunaga, I., Sprengel, R., Seeburg, P.H., Klugmann, M., Margrie, T.W., Schaefer, A.T., and Kuner, T. (2010). Synaptic inhibition in the olfactory bulb accelerates odor discrimination in mice. Neuron *65*, 399–411.

Adrian, E.D. (1950). The electrical activity of the mammalian olfactory bulb. Electroencephalogr. Clin. Neurophysiol. *2*, 377–388.

Arevian, A.C., Kapoor, V., and Urban, N.N. (2008). Activity-dependent gating of lateral inhibition in the mouse olfactory bulb. Nat. Neurosci. *11*, 80–87.

Assisi, C., Stopfer, M., and Bazhenov, M. (2011). Using the structure of inhibitory networks to unravel mechanisms of spatiotemporal patterning. Neuron *69*, 373–386.

Aungst, J.L., Heyward, P.M., Puche, A.C., Karnup, S.V., Hayar, A., Szabo, G., and Shipley, M.T. (2003). Centre-surround inhibition among olfactory bulb glomeruli. Nature *426*, 623–629.

Brody, C.D., and Hopfield, J.J. (2003). Simple networks for spike-timing-based computation, with application to olfactory processing. Neuron *37*, 843–852.

Cecchi, G.A., Petreanu, L.T., Alvarez-Buylla, A., and Magnasco, M.O. (2001). Unsupervised learning and adaptation in a model of adult neurogenesis. J. Comput. Neurosci. *11*, 175–182.

Cury, K.M., and Uchida, N. (2010). Robust odor coding via inhalation-coupled transient activity in the mammalian olfactory bulb. Neuron *68*, 570–585.

Davis, B.J., and Macrides, F. (1981). The organization of centrifugal projections from the anterior olfactory nucleus, ventral hippocampal rudiment, and piriform cortex to the main olfactory bulb in the hamster: an autoradiographic study. J. Comp. Neurol. *203*, 475–493.

Doucette, W., and Restrepo, D. (2008). Profound context-dependent plasticity of mitral cell responses in olfactory bulb. PLoS Biol. *6*, e258.

Egger, V., and Urban, N.N. (2006). Dynamic connectivity in the mitral cell-granule cell microcircuit. Semin. Cell Dev. Biol. *17*, 424–432.

Fantana, A.L., Soucy, E.R., and Meister, M. (2008). Rat olfactory bulb mitral cells receive sparse glomerular inputs. Neuron *59*, 802–814.

Firestein, S. (2004). A code in the nose. Sci. STKE *2004*, pe15.

Fuentes, R.A., Aguilar, M.I., Aylwin, M.L., and Maldonado, P.E. (2008). Neuronal activity of mitral-tufted cells in awake rats during passive and active odorant stimulation. J. Neurophysiol. *100*, 422–430.

Hertz, J., Krogh, A., and Palmer, R.G. (1991). Introduction to the theory of neural computation (Reading, Mass: Addison-Wesley Pub. Co.).

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. P Natl Acad Sci-Biol *79*, 2554–2558.

Hopfield, J.J. (1995). Pattern recognition computation using action potential timing for stimulus representation. Nature *376*, 33–36.

Kay, L.M., and Laurent, G. (1999). Odor- and context-dependent modulation of mitral cell activity in behaving rats. Nat. Neurosci. *2*, 1003–1009.

Kay, L.M., and Sherman, S.M. (2007). An argument for an olfactory thalamus. Trends Neurosci. *30*, 47–53.

Koulakov, A., Gelperin, A., and Rinberg, D. (2007). Olfactory coding with all-or-nothing glomeruli. J. Neurophysiol. *98*, 3134–3142.

Lee, D.D., and Seung, H.S. (1997). Unsupervised learning by convex and conic coding. In Advances in Neural Information Processing Systems, Michael I. Jordan, Michael J. Kearns, and Sara A. Solla, eds. (Cambridge, MA, USA: MIT Press), pp. 515–521.

Linster, C., and Hasselmo, M. (1997). Modulation of inhibition in a model of olfactory bulb reduces overlap in the neural representation of olfactory stimuli. Behav. Brain Res. *84*, 117–127.

Linster, C., Garcia, P.A., Hasselmo, M.E., and Baxter, M.G. (2001). Selective loss of cholinergic neurons projecting to the olfactory system increases perceptual generalization between similar, but not dissimilar, odorants. Behav. Neurosci. *115*, 826–833.

Lledo, P.M., Gheusi, G., and Vincent, J.D. (2005). Information processing in the mammalian olfactory system. Physiol. Rev. *85*, 281–317.

Luskin, M.B., and Price, J.L. (1983). The topographic organization of associational fibers of the olfactory system in the rat, including centrifugal fibers to the olfactory bulb. J. Comp. Neurol. *216*, 264–291.

Moulton, D.G. (1963). Electrical activity in the olfactory system of rabbits with indwelling electrodes. In Olfaction and Taste I, Y. Zotterman, ed. (Oxford: Pergamon Press), pp. 71–84.

Mumford, D. (1994). In Neuronal architectures for pattern-theoretic problems. In large scale neuronal theories of the brain, C. Koch and J.L. Davis, eds. (Cambridge, MA: MIT Press), pp. 125–152.

Olshausen, B.A., and Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature *381*, 607–609.

Olshausen, B.A., and Field, D.J. (1997). Sparse coding with an overcomplete basis set: a strategy employed by V1? Vision Res. *37*, 3311–3325.

Olshausen, B.A., and Field, D.J. (2004). Sparse coding of sensory inputs. Curr. Opin. Neurobiol. *14*, 481–487.

Rinberg, D., Koulakov, A., and Gelperin, A. (2006). Sparse odor coding in awake behaving mice. J. Neurosci. *26*, 8857–8865.

Rozell, C.J., Johnson, D.H., Baraniuk, R.G., and Olshausen, B.A. (2008). Sparse coding via thresholding and local competition in neural circuits. Neural Comput. *20*, 2526–2563.

Seung, H.S., Richardson, T.J., Lagarias, J.C., and Hopfield, J.J. (1998). Minimax and Hamiltonian dynamics of excitatory-inhibitory networks. In Advances in neural information processing systems, Michael J. Kearns, Sara A. Solla, and David A. Cohn, eds. (Cambridge, MA, USA: MIT Press), pp. 329–335.

Shepherd, G.M., Chen, W.R., and Greer, C.A. (2004). Olfactory Bulb. In The synaptic organization of the brain, G.M. Shepherd, ed. (New York: Oxford University Press), pp. 165–216.

Shepherd, G.M., Chen, W.R., Willhite, D., Migliore, M., and Greer, C.A. (2007). The olfactory granule cell: from classical enigma to central role in olfactory processing. Brain Res. Brain Res. Rev. *55*, 373–382.

Shusterman, R., Smear, M.C., Koulakov, A.A., and Rinberg, D. (2011). Nat. Neurosci. *14*, 1039–1044.

Ukhanov, K., Corey, E.A., Brunert, D., Klasen, K., and Ache, B.W. (2010). Inhibitory odorant signaling in Mammalian olfactory receptor neurons. J. Neurophysiol. *103*, 1114–1122.

Wachowiak, M., Denk, W., and Friedrich, R.W. (2004). Functional organization of sensory input to the olfactory bulb glomerulus analyzed by two-photon calcium imaging. Proc. Natl. Acad. Sci. USA *101*, 9097–9102.

Wick, S.D., Wiechert, M.T., Friedrich, R.W., and Riecke, H. (2010). Pattern orthogonalization via channel decorrelation by adaptive networks. J. Comput. Neurosci. *28*, 29–45.