Environmental adaptation of olfactory receptor distributions

Tiberiu Teşileanu[1,2], Simona Cocco[3], Rémi Monasson[4], and Vijay Balasubramanian[1,2]

[1] *Initiative for the Theoretical Sciences, The Graduate Center, CUNY, New York, NY 10016*
[2] *David Rittenhouse Laboratories, University of Pennsylvania, Philadelphia, PA 19104*
[3] *Laboratoire de Physique Statistique, École Normale Supérieure and CNRS UMR 8550, PSL Research, UPMC Sorbonne Université, Paris, France*
[4] *Laboratoire de Physique Théorique, École Normale Supérieure and CNRS UMR 8550, PSL Research, UPMC Sorbonne Université, Paris, France*

January 28, 2018

Abstract

Olfactory receptor usage is highly non-uniform, with some receptor types being orders of magnitude more abundant than others. Intriguingly, in mammals the receptors are replaced regularly and the receptor distribution changes in response to olfactory experience. We propose a theoretical model that explains these experimental observations. Our model suggests that the receptor distribution is tuned to maximize the amount of information about the olfactory environment, in accordance with the efficient coding hypothesis. The optimal receptor distribution depends on natural odor statistics, implying that a change in environment should lead to changes in receptor abundances. The effect is more pronounced when olfactory receptors respond strongly to a smaller number of odors. Our model also predicts that, between closely-related species, those with more neurons in the olfactory epithelium should express more receptor types. We show that simple dynamical rules for neural birth and death processes converge to the predicted optimum.

1 Introduction

Olfaction is hard: according to some estimates, there may be tens of billions of odorants—volatile molecules that exhibit an odor—that can be encountered in nature [1]. The olfactory system has the daunting task of detecting these molecules in minute concentrations in order to guide an animal towards food or away from danger. To do so, the olfactory epithelium in mammals and the antennae in insects are populated with large numbers of olfactory sensory neurons (OSNs), each of which expresses a single kind of olfactory receptor. Each type of olfactory receptor binds to many different odorants, and each odorant activates many different receptors, leading to a complex encoding of olfactory scenes [2]. Olfactory receptors form the largest known gene family in mammalian genomes, with hundreds to thousands of members, owing perhaps to the importance that olfaction has for an animal's fitness [3, 4, 5]. Large olfactory receptor families can also be found in insects, although the receptor gene families in vertebrates and invertebrates appear to have evolved independently [6].

Although large, the number of different olfactory receptor types available to animals is tiny compared to the number of odorants that they potentially need to detect and identify in arbitrary mixtures at arbitrary concentrations. This implies that any attempt by the olfactory system to accurately encode natural odors must take advantage of redundancies in the structure of olfactory scenes. Put differently, the neural encoding used in the olfactory periphery should be adapted to the statistics of natural odors, in accordance with the efficient coding hypothesis [7, 8]. Indeed, this kind of adaptation has already been observed in relation to other sensory modalities, such as vision and audition [9, 10, 11, 12, 13, 14, 15]. It has also been suggested that olfactory sensing might be adapted to the sparsity of odor mixtures [16, 17, 18] and that receptor-odorant affinities might be specially adapted to odors in the environment [19, 20, 21].

In this paper, we take the receptor-odorant affinities as given and instead propose a model for the quantitative distribution of olfactory sensory neurons by receptor type. Our model is based on the hypothesis that the neural responses represent a maximally-informative picture of the olfactory environment of the animal. Qualitatively, this leads to the prediction that, in a noisy odor environment, the distribution of receptor types that is best adapted to natural olfactory scenes is highly non-uniform, but reproducible given the receptor affinities and the environmental odor statistics. The adaptation could happen during the lifetime of an individual, in response to the statistics of odors encountered by the organism, if such changes are allowed by the biology (*e.g.,* in mammals [22]); or it could happen over evolutionary times, in response to the typical odor environments encountered by organisms of the same species.

Short time-scale adaptation has already been observed in experiments [23, 24, 25, 26, 27, 28]. Furthermore, Ibarra-Soria et al. [28] also find that the distribution of receptor types in the mouse olfactory epithelium is highly non-uniform but strain-specific, suggesting a long-term component to the adaptation. Our model also predicts that the magnitude of the effect of a change in environment on the receptor distribution depends on how narrowly or broadly-tuned the receptors are, with a larger effect when each receptor tends to respond strongly to only a few odorants. The optimal receptor distribution is also affected by the total number of OSNs in the epithelium, with smaller numbers of neurons leading to some of the available receptor types not being used.

Information maximization in our model leads to a complex nonlinear dependence between receptor abundances and environmental statistics. We show, however, that the information maximum can be reached using a simple dynamical model based on logistic growth, in which the death rates of neurons depend on the covariance structure of receptor responses. This provides an avenue for a mechanistic implementation of our model.

2 Model

Before turning to the specifics of our model, we briefly introduce the anatomy of the olfactory periphery. In vertebrates, axons from olfactory neurons converge in the olfactory bulb on compact structures called glomeruli, where they form synapses with the dendrites of downstream neurons [29]; see Figure 1a. To a good approximation, each glomerulus receives axons from only one type of OSN, and all the OSNs expressing the same receptor type converge onto a small number of glomeruli, on average about 2 in mice to about 16 in humans [30]. A similar architecture can be found in insects [31].

The anatomy shows that in insects and vertebrates, the olfactory information that is passed to the brain can be summarized by activity at the level of the glomeruli. For the purposes of our model, we will treat this activity in a firing-rate approximation, ignoring the timing of individual spikes. There is evidence that at least in mammals, but perhaps also in insects, timing is important for odor discrimination [32, 33, 34, 35]. This could be incorporated into our model by using the timing of spikes instead of the average firing rate as a measure of receptor activity, but we not do so here. We will also assume that all neurons of a given type converge onto a single glomerulus (as in Figure 1a). This is without loss of generality because it turns out that within our model, information transfer only depends on the total number of neurons of a given type, and does not change if their outputs are split over several glomeruli (see Supplementary Information).

There is comparatively less known regarding the structure of odor environments. In fact, one of the difficulties specific to the study of the olfactory system when compared to other senses is the lack of a good canonical model for odor space. It is in general difficult to identify the common features shared by odorants that activate a given receptor type [36, 2], while attempts at defining a notion of distance in olfactory space have had only partial success [37], as have attempts to find reduced-dimensionality representations of odor space [38, 39]. In this work, we simply model the olfactory environment as a vector $\boldsymbol{c} = \{c_1, \ldots, c_N\}$ of concentrations, where c_i is the concentration in the environment of odorant i (Figure 1b). In fact, the mathematical analysis does not depend in any way on this interpretation—it could instead represent a different abstraction, such as the concentrations of classes of molecules clustered based on common chemical traits, or abstract coordinates in a low-dimensional representation of olfactory space.

With this parametrization, the statistics of natural scenes can be summarized by the joint probability distribution $P(c_1, \ldots, c_N)$. This ignores temporal correlations in olfactory cues, which is an approximation that has been used successfully in similar studies in the visual system [40, 41]. We aim here to obtain a

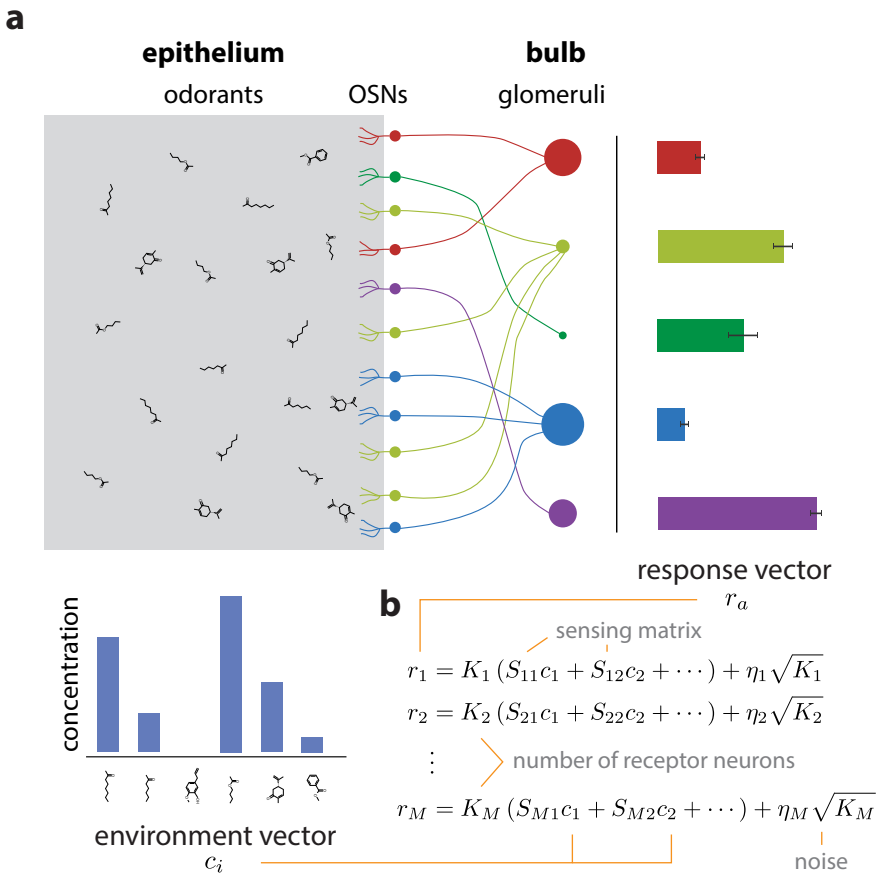

Figure 1: Sketch of the olfactory periphery as described in our model. **(a)** Brief overview of the anatomy of the olfactory epithelium and bulb in vertebrates. Different receptor types are represented by different colors. The architecture is similar in insects, with the olfactory sensory neurons located in the antennae instead of the epithelium, and the glomeruli located in the antennal lobes. The olfactory signal proceeds downstream to cortex in mammals and the mushroom body in insects. A given mixture of odorants in the environment (bar plot on the bottom left) leads to a response vector of activations for the glomeruli (bar plot on the top right). The noise in the response vector, shown by black error bars in the bar plot, depends on the number of receptors of each type, illustrated in the figure by the size of the corresponding glomerulus. Glomeruli receiving input from a small number of OSNs have higher variability due to receptor noise (*e.g.,* OSN, glomerulus, and activity bar in dark green), while those receiving input from many OSNs have smaller variability (*e.g.,* OSNs, glomerulus, and activity bar in blue). In contrast, the response magnitudes depend not only on the number of OSN inputs, but also on the odorants present in the medium and the affinity profile of the receptors (*e.g.,* blue bar is smaller than green bar despite blue glomerulus being larger). **(b)** The glomerular responses are calculated using a linear model based on a "sensing matrix" S, and include Gaussian noise.

tractable model that allows us to build intuition regarding the relation between natural odor statistics and olfactory receptor distributions. For this reason, we model the probability distribution of the environment vector as a multivariate Gaussian with mean $\boldsymbol{c}_0$ and covariance matrix Γ,

$$\text{environment } P(\boldsymbol{c}) \sim \mathcal{N}(\boldsymbol{c}_0, \Gamma)\,. \tag{1}$$

We only need to include in our information maximization procedure those odorants that are likely to be encountered at meaningful concentrations in the typical environment of a given animal. This number N, though still large, will be much smaller than the total number of possible odorants (which could be in the

billions [1]). In our numerical studies the main limitation on the number of odorants that we can use comes from the availability of affinity data. In this paper, we take N to be of order 100, matching the number of odorants available in the receptor affinity data that we use, 110 in fly [42] and 63 in mammals [43].

In order to find the amount of information that the brain receives about an olfactory scene, we need to model the way in which glomerular responses depend on the odor environment. We assume here a linear model, in which increasing the concentration of an odorant by a certain factor leads to an increase in the activation of a receptor by the same factor, and the response to odor mixtures is the sum of the activations from each of the components. We define the response of a given glomerulus as simply the sum of the responses from all the olfactory neurons that project to it. These assumptions can be thought of as linearizing the full non-linear responses in the olfactory epithelium around a given operating point. Note that since we model all the responses in a linear way, we make no difference between changes in the number of receptor molecules per neuron, and changes in the number of olfactory neurons of a given type. In general these could have different effects due to nonlinear neuronal responses.

Mathematically, we write

$$r_a = K_a \sum_i S_{ai} c_i + \eta_a \sqrt{K_a}\,, \tag{2}$$

where S_{ai} is the expected response of a sensory neuron expressing receptor type a to a unit concentration of odorant i (see Figure 1b), and K_a is the number of neurons of type a. The second term models noise, with η_a being the noise for a single OSN. We use a Gaussian approximation for η_a and we normalize the responses by the standard deviation of the noise for a single neuron, so that η_a are drawn from independent standard normal distributions, $\eta_a \sim \mathcal{N}(0,1)$.

To quantify the amount of information that the response vector $\boldsymbol{r}$ contains about the environment vector $\boldsymbol{c}$, we use the mutual information $I(\boldsymbol{r},\boldsymbol{c})$ between these two quantities:

$$I(\boldsymbol{r},\boldsymbol{c}) = \int d^M r d^N c\, P(\boldsymbol{r},\boldsymbol{c}) \cdot \log\left[\frac{P(\boldsymbol{r}|\boldsymbol{c})}{P(\boldsymbol{r})}\right]\,, \tag{3}$$

where $P(\boldsymbol{r},\boldsymbol{c})$ is the joint probability distribution over response and concentration vectors, $P(\boldsymbol{r}|\boldsymbol{c})$ is the distribution of responses conditioned on the environment, and $P(\boldsymbol{r})$ is the marginal distribution of the responses alone. Given the assumptions we made, all of these distributions are Gaussian, and the integral can be evaluated analytically (see Supplementary Information). The result is

$$I(\boldsymbol{r},\boldsymbol{c}) = \frac{1}{2}\,\mathrm{Tr}\log(\mathbb{I} + \mathbb{K}Q)\,, \tag{4}$$

where the *overlap matrix* Q is related to the covariance matrix Γ of odorant concentrations in natural scenes (from eq. (1)),

$$Q = S\Gamma S^T\,, \tag{5}$$

and $\mathbb{K}$ is a diagonal matrix containing the receptor abundances K_a,

$$\mathbb{K} = \mathrm{diag}\big(K_1, K_2, \ldots, K_M\big)\,. \tag{6}$$

In the absence of noise, the overlap matrix Q is equal to the covariance matrix of receptor responses averaged over the ensemble of natural odors. Because of the way in which we normalized the receptor responses, the quantity $\mathbb{K}Q$ thus behaves as a signal-to-noise ratio (SNR), so that eq. (4) is essentially a generalization to multiple, correlated channels of the standard result for a single Gaussian channel, $I = \frac{1}{2}\log(1+\mathrm{SNR}^2)$ [44, 11, 8].

The receptor numbers K_a cannot grow without bound; they are constrained by the fact that the total number of neurons in the olfactory epithelium has to stay within realistic limits. Thus, to find the optimal distribution of receptor types, we maximize $I(\boldsymbol{r},\boldsymbol{c})$ with respect to $\{K_a\}$, subject to the constraints that: (1) the total number of receptor neurons is fixed; and (2) all neuron numbers are non-negative,

$$\{K_a\} = \arg\max\, I(\boldsymbol{r},\boldsymbol{c})\,, \qquad \text{subject to} \quad \begin{cases} K_{\mathrm{tot}} = \sum_a K_a = \mathrm{const}, \\ K_a \geq 0\,. \end{cases} \tag{7}$$

The optimization can be performed analytically, using the Karush-Kuhn-Tucker (KKT) conditions [45], which yields the following implicit equations for K_a (see Supplementary Information for a derivation):

$$\begin{cases} (Q^{-1} + \mathbb{K})^{-1}_{aa} = \lambda\,, & \text{if } K_a > 0, \text{ or} \\ (Q^{-1} + \mathbb{K})^{-1}_{aa} < \lambda\,, & \text{if } K_a = 0, \end{cases} \tag{8}$$

where λ is a Lagrange multiplier. The magnitude of λ is set by imposing the normalization condition in eq. (7).

3 Results

Our model predicts the optimal distribution of receptor types given receptor affinities, the statistics of natural odors, and the size of the olfactory epithelium. In the following sections, we analyze how the model results depend on these factors. The code that we used to generate all the results and figures in the paper is available on GitHub, at `https://github.com/ttesileanu/OlfactoryReceptorDistribution`.

3.1 Dependence on the total number of olfactory sensory neurons

Large number of sensory neurons In our model, the effect of receptor noise can be diminished by averaging over the responses from many sensory neurons. In the limit in which the total number of neurons is allowed to grow very large, $K_{\text{tot}} \to \infty$, the signal-to-noise ratio (SNR) becomes very large, and the averaged glomerular responses become effectively noiseless (see eq. (2)). When this happens, the KKT conditions from eq. (8) can be solved analytically (see Supplementary Information), and we find that the optimal receptor distribution is given by

$$K_a = \frac{K_{\text{tot}}}{M} - (A_{aa} - \bar{A}) + \mathcal{O}(A^2/K^2)\,, \tag{9}$$

where A is the inverse of the overlap matrix defined in eq. (5), $A = Q^{-1}$, and $\bar{A} = \sum A_{aa}/M$ is the average diagonal element of this matrix. In other words, the receptor distribution is essentially uniform, with each receptor type being expressed in a roughly equal fraction of the total population of sensory neurons, with small differences related to the diagonal elements of the inverse overlap matrix A (see Figure 2a). It turns out also that the information maximum in this regime is very shallow. This is because only a change in receptor numbers of order K_{tot}/M can have a significant effect on the noise level for the activity of each glomerulus. Put another way, when the receptor numbers K_a are very large, the responses are effectively noiseless, and in the noiseless case the number of receptors of each type has no effect on the reliability of the responses, and thus it has no effect on information transfer. The approximation in eq. (9) holds as long as the receptor abundances K_a are much larger than the elements of the inverse overlap matrix A, and it breaks down otherwise (see the points on the right in Figure 2f).

In natural olfactory systems, the distribution of receptor types is highly non-uniform, and is more conserved for individuals of the same species than between species [28]. This suggests that the high SNR regime in our model cannot adequately describe olfaction. Indeed, resource conservation in the organization of living systems would suggest that the number of olfactory sensory neurons (and, thus, the SNR) should instead be close to the minimum necessary.

Small number of sensory neurons When the number of neurons is very small, receptor noise can overwhelm the response to the environment. In this case, we find that the best strategy is to focus all the available neurons on a single receptor type, thus reducing noise as much as possible (Figure 2b). The receptor type that yields the highest contribution to the information turns out to be the one whose response is most variable in natural scenes (that is, the one that corresponds to the largest value of Q_{aa}; see Supplementary Information for a derivation). This is reminiscent of a result in vision where the variance of a stimulus predicted its perceptual salience [41]. As the total number of neurons increases, the added benefit from lowering the noise for a single receptor type diminishes, and at some critical value it becomes more advantageous in terms of maximizing information to start populating a second receptor type (Figure 2b).

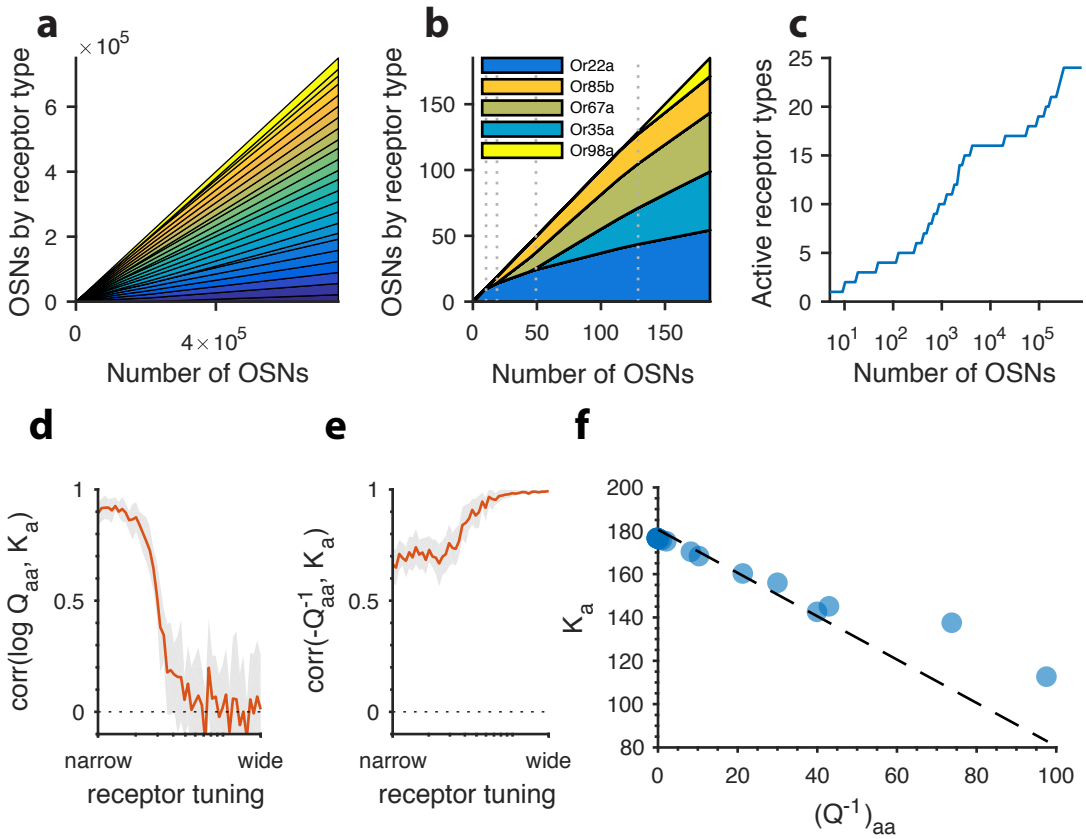

Figure 2: Building an intuition for the results from our model. **(a)** When the total number of neurons, K_{tot}, is large, the optimal distribution of receptors is approximately uniform, with each receptor type (represented here by strips of different colors) taking up a fraction $\approx 1/M$ of the neurons, where M is the number of receptor types. **(b)** When the number of neurons K_{tot} is low, only a small number of receptor types are expressed when information is maximized. **(c)** New receptor types become expressed in a series of step transitions as the total number of neurons increases. The results in panels **a**–**c** were obtained using a sensing matrix based on measurements in the fly [42] and environmental odor statistics following a random correlation matrix (see Supplementary Information for details). However, the qualitative aspects are generic. **(d)** Correlation between the abundance of a given receptor type, K_a, and the logarithm of the variance with which it is activated in olfactory scenes, $\log Q_{aa}$, shown here as a function of the tuning of the receptors. For every position along the x-axis, sensing matrices with a fixed receptor tuning width were generated from a random ensemble. Roughly-speaking, the tuning width on the x-axis indicates what fraction of all odorants elicit a strong response for the receptors; in the figure, this ranges from 0.01 to 0.2. See Supplementary Information for details. For narrow tuning (each receptor responding strongly to only a small number of odorants), response variance is a good predictor of abundance, while this is no longer true for wide tuning. **(e)** Receptor abundances correlate well with the diagonal elements of the inverse overlap matrix, $(Q^{-1})_{aa}$, for all tuning widths, with somewhat lower correlations at narrow tuning. In **d** and **e**, the red line is the mean obtained from 24 simulations, each performed using a different sensing matrix, and the light gray area shows the interval between the 20th and 80th percentiles of results. **(f)** Example showing the dependence between receptor abundances and the diagonal elements of the inverse overlap matrix in a simulation using a wide-tuning sensing matrix and an intermediate SNR. The dashed line corresponds to the approximation from eq. (9).

General case More generally, our model predicts that the number of receptor types that are expressed should increase monotonically with the total number of sensory neurons, in a series of step transitions (see Figure 2c). This provides an interesting prediction that could be tested by looking at the relation between the

size of the olfactory epithelium and the number of functional receptor genes in various species. A caveat here is that we are assuming the sensing matrix S and environmental statistics Γ stay the same while varying only the total number of neurons K_{tot}. This restricts an empirical test of this prediction to groups of closely-related species.

3.2 Logic of adaptation

By maximizing the mutual information in eq. (4) while keeping the total number of receptors $K_{\text{tot}} = \sum_a K_a$ constant, we can predict the optimal receptor distribution K_a given the sensing matrix S and the environmental covariance matrix Γ. Note that the average concentration vector $\boldsymbol{c}_0$ has no impact on the optimal distribution because it corresponds to odors that are always present and thus offer no new information about the environment.

Is there an intuitive way to understand the results of the optimization? We know that receptors whose activations fluctuate extensively in natural olfactory scenes provide the brain with a lot of information about the odor environment (assuming noise is negligible), while receptors that are almost silent or are constantly active at about the same rate are not very informative. It thus seems advantageous to increase the number of neurons expressing receptors of the first kind, and not focus many resources on receptors of the second kind. In terms of our model, the variances of the glomerular responses in natural scenes are given (in the absence of noise) by the diagonal elements of the overlap matrix Q defined in eq. (5), so our intuition would dictate that the number of receptors of type a grows with the variance corresponding to that receptor type, $K_a \sim Q_{aa}$.

As mentioned above, this intuition is true in the very low SNR regime, when only one receptor type is active: in that case, the type that should be active to maximize information is indeed the one with the largest variance Q_{aa}. The intuition also holds when the sensing matrix is narrowly-tuned: if each receptor type responds to only a small number of odorants, the abundances of OSNs of each type correlate well with the amount of variability that the corresponding receptors experience in natural scenes (left side of Figure 2d). More generally, we can generate sensing matrices with varying tuning width by varying the number of odorants that exhibit strong responses for each receptor (detailed Methods in Supplementary Information). When we do this, we see that as each receptor type starts responding to more and more odorants, the correlation structure of the receptor responses becomes important in determining the optimal receptor distribution, and the intuition based on only the response variances no longer holds (Figure 2d). Importantly, we carried out this analysis in an intermediate-SNR regime in which noise is important, but does not dominate the responses of most receptors.

The optimal receptor distribution is also well approximated by a constant minus the diagonal elements of the *inverse* of the overlap matrix, $A = Q^{-1}$; Figure 2f. This is exact in the high-SNR limit, as described in eq. (9) and remains a good approximation even for lower signal-to-noise ratios, and even when the optimal receptor distribution doesn't correlate well with the receptor response variances Q_{aa}, *e.g.,* when the receptors are broadly tuned (compare Figure 2e and 2f). Note that, if the total number of neurons K_{tot} becomes too low, some of the values from eq. (9) become negative. To make sensible predictions for the corresponding abundances, the values can be clipped at 0. This provides an approximate method of determining which receptor abundances should vanish in the optimal distribution.

The correlation of the receptor distribution to the diagonal elements of the inverse overlap matrix Q_{aa}^{-1} may occur because these elements are the product of two factors (see Supplementary Information): (1) the reciprocals of the response variances, $1/Q_{aa}$; and (2) $1/\left(1-\rho_a^2\right)$, where ρ_a are goodness-of-fit parameters that show how well the responses of a given receptor are predicted by a linear combination of all the others. Thus, receptors that have a high Q_{aa}^{-1} either do not provide much olfactory information because their responses do not fluctuate much; or they provide information that is largely redundant given all the others. Either way, they should therefore not be very abundant in the optimal distribution. In contrast, receptors with low Q_{aa}^{-1} contain a lot of independent olfactory information, and so they are predicted to be abundant.

3.3 Dependence on the environment

General considerations To get some intuition for the dependence on the environment in our model, we used a sensing matrix based on measurements in the fly [42], and compared the predicted receptor distribution

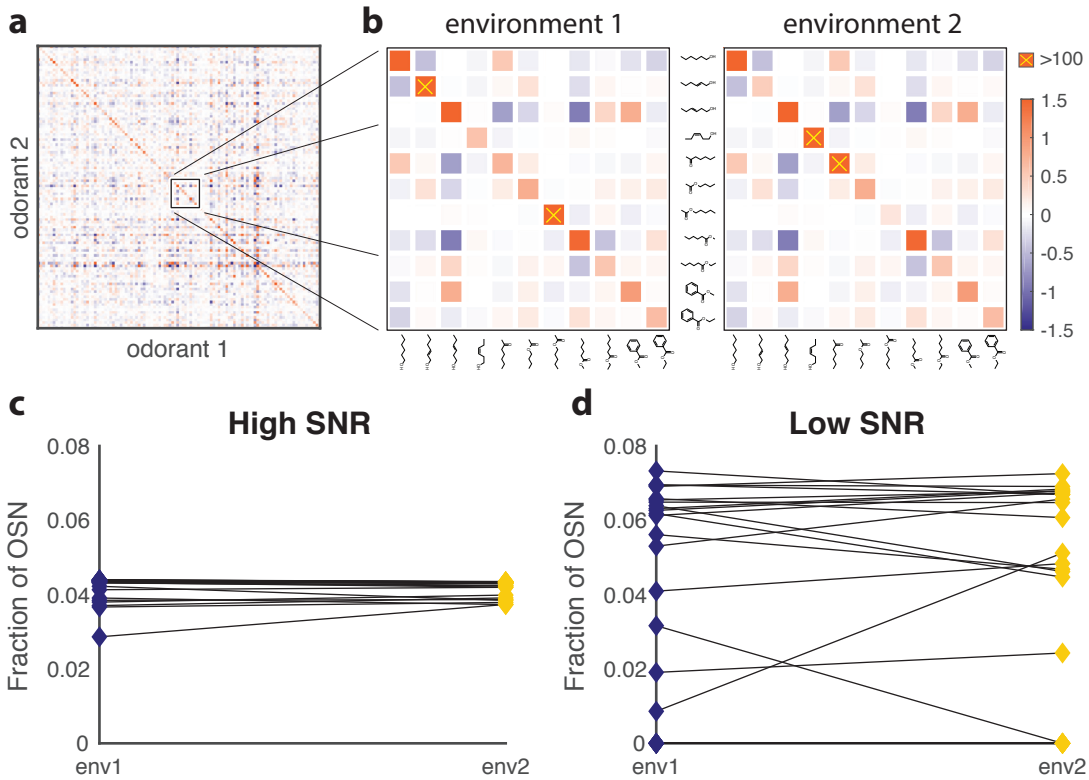

Figure 3: Effect of changing environment on the optimal receptor distribution in our model. **(a)** An example of an environment covariance matrix used in our model. We use an algorithm that generates random covariance matrices with a tunable amount of cross-correlation (see Supplementary Information for details). The variances are drawn from a lognormal distribution. **(b)** Close-ups showing some of the differences between the two environments used to generate the results from **c** and **d**. The two different covariance matrices are obtained by starting with the matrix from **a** and adding a large amount of variance to two different sets of 10 odorants (out of 110). The altered odorants are identified by yellow crosses; their corresponding variances go above the color scale on the plots by more than 60 times. **(c)** Change in receptor distribution when going from environment 1 to environment 2, in conditions where the total number of receptor neurons K_{tot} is large (the SNR is high). The blue diamonds on the left correspond to the optimal OSN fractions per receptor type in the first environment, while the yellow diamonds on the right correspond to the second environment. In this high-SNR regime, the effect of the environment is small, because in both environments the optimal receptor distribution is close to uniform. **(d)** When the total number of neurons K_{tot} is small (the SNR is low), changing the environment can have a dramatic effect on optimal receptor abundances, with some receptors that are almost vanishing in one setting becoming highly abundant in the other, and *vice versa*.

in two different environments, at two different noise levels (or, equivalently, two different values for the total number of OSNs K_{tot}); see Figure 3.

Unfortunately, there are no measurements available for the distributions of natural odors. As such, we have little information to guide our choice for the statistical structure of odorant concentrations, modeled here by a Gaussian $\mathcal{N}(\boldsymbol{c}_0, \Gamma)$. To get a qualitative understanding of the efficient coding of olfactory scenes, we therefore took the environment as having a random covariance matrix (see Figure 3a for an example; see Supplementary Information for details). Starting with a base environment, we then generated two different environments by adding a large amount of variance to 10 odorants in environment 1, and to 10 different odorants in environment 2 (Figure 3b).

Our results show that, when the number of olfactory sensory neurons K_{tot} is large, and thus the signal-to-noise ratio is high, the change in odor statistics has little effect on the distribution of receptors (Figure 3c).

This is because at high SNR, the distribution is essentially uniform regardless of environment, as explained earlier. When the number of neurons is smaller (or, equivalently, the signal to noise ratio is in a low or intermediate regime), the change in environment has a significant effect on the receptor distribution, with some receptor types becoming more abundant, others becoming less abundant, and yet others not changing much (see Figure 3d). This mimics the kinds of effects seen in experiments in mammals [23, 24, 25, 26, 27, 28].

Strikingly, and consistently with the findings in [23, 24, 25, 26, 27, 28], whether a certain receptor abundance goes up or down as a result of a given change in environment is hard to guess, as it is a global effect that depends on how the change affects the activation of all receptor types, as well as the correlations between these activations. As a rough approximation, the change in abundance ΔK_a is related to the change in the corresponding element of the inverse overlap matrix, ΔQ_{aa}^{-1}. However, this estimate does not work as well when receptors are narrowly-tuned, where the approximation $K_a \approx C - Q_{aa}^{-1}$ for some constant C is less accurate.

Factors that affect the sensitivity of the receptor distribution to the environment Above, we compared the effects of environments created by increasing the variances of some odorants in a background covariance matrix by a large amount—much larger than the typical element of the background matrix. Figure 4a shows what happens if we instead compare environments with two different randomly-generated covariance matrices, each generated in the same way as the background environment in Figure 3a. The resulting environmental covariance matrices (Figure 4a, top) are very different in detail (for example, the correlation coefficient between their entries is close to zero), although they look similar by eye. Nevertheless, the corresponding change in receptor distribution is very small, with just a few receptor types experiencing larger changes in abundance (Figure 4a, bottom).

In contrast, if we engineer two environments that are almost non-overlapping, meaning that each odorant is either common in environment 1, or in environment 2, but not in both (see top of Figure 4b), the changes in receptor abundances become much larger, with some receptors going from highly abundant to vanishing from the population, while others going in the opposite direction, from not being expressed at all, to high abundance (Figure 4b, bottom). This shows that the particular way in which the environment changes, and not only the magnitude of the change, can alter the effect on the receptor distribution. In biological terms, we would indeed expect animals that experience very different kinds of odors to have more striking differences in their receptor repertoires than those that merely experience the same odors with different frequencies. Our results quantify the expected difference.

The magnitude of the effect of environmental changes on the optimal olfactory receptor distribution is partly controlled by the tuning of the olfactory receptors. If receptors are narrowly-tuned, with each type responding to a small number of odorants, changes in environment tend to have more drastic effects on the receptor distribution than when the receptors are broadly-tuned (Figures 4c and 4d), an effect that could be experimentally tested.

Mammalian data Our study opens the exciting possibility of a causal test of the hypothesis of efficient coding in sensory systems, where a perturbation in the odor environment could lead to predictable adaptations of the olfactory receptor distribution during the lifetime of an individual. This does not happen in insects, but it can happen in mammals, since their receptor neurons regularly undergo apoptosis and are replaced.

A recent study demonstrated that reproducible changes in olfactory receptor distributions of the sort that we predict can occur in mice [28]. Specifically, Ibarra-Soria et al. [28] raised two groups of mice in similar conditions, exposing one group to a mixture of four odorants (acetophenone, eugenol, heptanal, and R-carvone) either continuously or intermittently (by adding the mixture to their water supply). Continuous exposure to the odorants had no effect on the receptor distribution, in agreement with the predictions of our model, in which the average concentration vector $\boldsymbol{c}_0$ drops out of the mutual information calculations. In contrast, intermittent exposure did lead to systematic changes in the receptor distribution (see Figure 5a).

We used our model to run an experiment similar to that of Ibarra-Soria et al. [28] *in silico*. Using a sensing matrix based on odor response curves for mouse and human receptors [43], we calculated the predicted change in receptor abundances between two different environments with random covariance matrices. The results (see Figure 5b) show that most receptor abundances do not change much, especially for those receptors that were very abundant to start with; there is more change in the middle of the abundance range; while the abundances of rare receptors change significantly, but are highly sensitive to small perturbations in the environment.

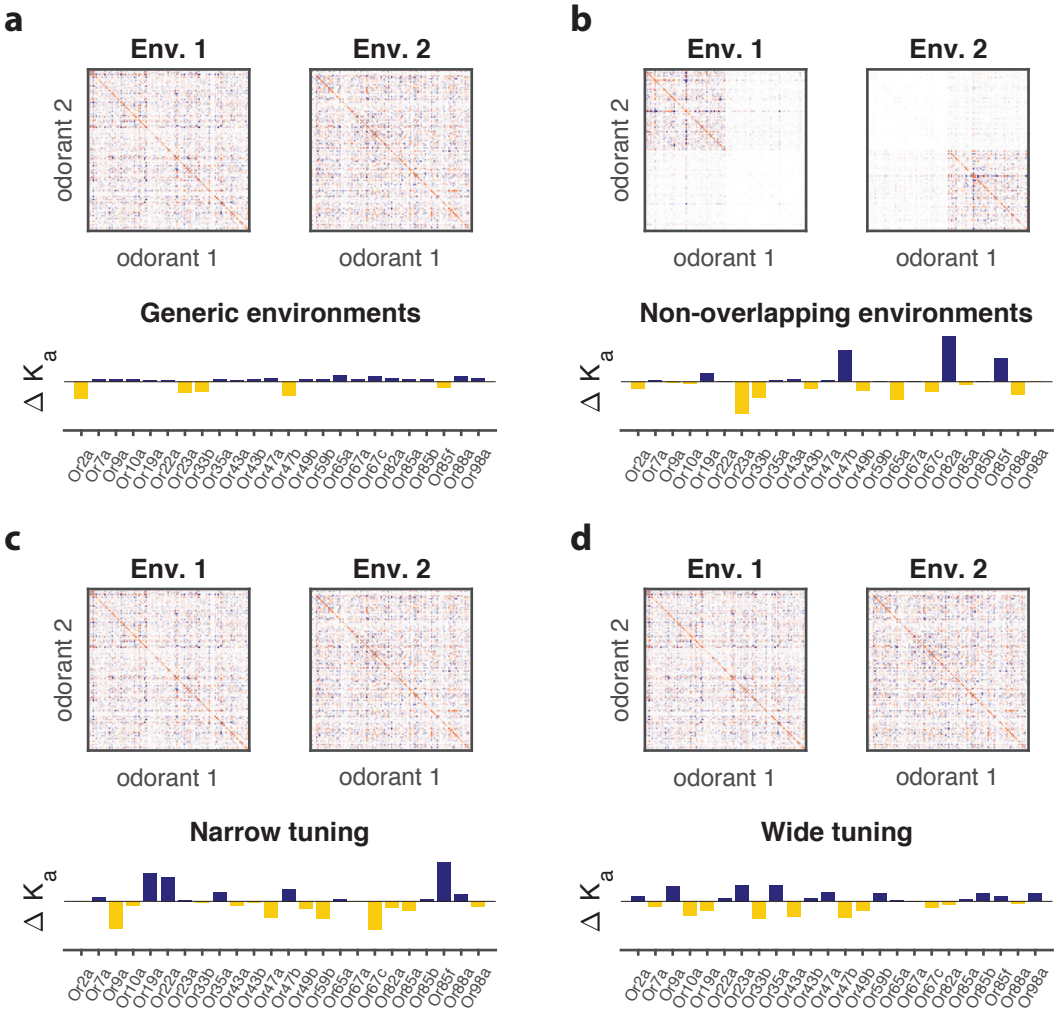

Figure 4: The effect of environmental statistics on the optimal receptor distribution as a function of overlap in the odor content of the two environments, and the tuning properties of the olfactory receptors. In all panels, the heat maps on top show the random covariance matrices used to model the two environments (red is positive [co-]variance, blue is negative, white is zero), while the bar plots show the difference in the optimal receptor distributions calculated in the two environments (blue is positive change, yellow is negative change). The scaling of the y-axis is arbitrary, but was kept the same across the four bar plots. **(a)** The change in receptor distribution tends to be small when both environments span a similar set of odors, even when the detailed statistics are very different. **(b)** The change in receptor distribution tends to be larger when the two environments contain largely non-overlapping sets of odors. The results in **a** and **b** are based on receptor affinity data from the fly [42]. **(c)** Narrowly-tuned receptors (this corresponds to a tuning parameter of 0.05—see Supplementary Information), which respond to small numbers of odorants, result in optimal receptor distributions that are more sensitive to changes in environment. **(d)** In contrast, broadly-tuned (corresponding to a tuning parameter of 0.8—see Supplementary Information) receptors tend to average out differences in environmental statistics of odors.

The sensitivity to perturbations was estimated by running the simulation 24 times, each time modifying the covariance matrices for the initial and final environments. This was done by adding a small amount of Gaussian random noise to the square roots of these covariance matrices (see Supplementary Information).

The simulation results qualitatively match the experimental results (see Figure 5a), where we see the same pattern of the largest reproducible changes occurring for receptors with intermediate abundances. Note

that in Figure 5a we replotted the raw data from Figure 5B from Ibarra-Soria et al. [28] without a Bayesian correction that was used in that paper which assumed no *a priori* change between exposed and control animals (personal communication with Professor Darren Logan, June 2017). Such a correction seemed inappropriate for our setting, since we know that the control and exposed groups have experienced different environments.

In principle this comparison could be made quantitative by using the same perturbation in our simulations as that used in the experiments from Ibarra-Soria et al. [28]. This is not yet possible for several reasons. Most importantly, sensing data is missing for one of the four odorants used in the experiments (eugenol). Furthermore, affinity data is only available for a small fraction of receptor types (52 mouse and 10 human receptors, compared to more than 1000 total receptors in the mouse), and a small number of odorants (63, although even for these 62×63 pairs, most responses were below the measurement threshold; see Supplementary Information). Since the reorganization of the receptor distribution is a global phenomenon that depends on the responses from all glomeruli, this makes it hard to interpret any results obtained with such a small subset of receptors. Precise predictions are also made difficult because the statistics of natural odors to which the mice are exposed are unknown, and even the statistics with which the mice encounter the odorant mixture used in the experiment are hard to guess. All these issues make it impossible to produce specific predictions at this point, although many of them can be addressed in future experiments.

3.4 Dynamical model

We have explored the structure of olfactory receptor distributions that code odors efficiently, *i.e.,* are adapted to maximize the amount of odor information that the brain gets about the environment. The full solution to the corresponding optimization problem, (8), depends in a complicated nonlinear way on the receptor affinities S and natural distribution of odorants Γ. The distribution of olfactory receptors in the mammalian epithelium, however, must arise dynamically, from interactions between neurons expressing various receptor types mediated by activity generated by the odor environment [46]. By starting with a gradient ascent argument, we worked out a dynamical model that obeys these constraints, and converges to the desired optimum (details in Supplementary Information). The resulting model is described by the equation

$$\frac{dK_a}{dt} = \alpha \left\{ K_a - \lambda K_a^2 - (R^{-1})_{aa} K_a^2 \right\} , \tag{10}$$

where α is a learning rate and R is the covariance matrix of glomerular responses,

$$R_{ab} = \langle r_a r_b \rangle - \langle r_a \rangle \langle r_b \rangle , \tag{11}$$

with the angle brackets denoting ensemble averaging over natural odors and averaging over receptor noise. In the absence of the experience-related term R_{aa}^{-1}, the dynamics from eq. (10) is simply logistic growth: the population of OSNs of type a grows at a rate α, but saturates when $K_a = 1/\lambda$ because of a population-dependent death rate λK_a. The last term in eq. (10) is a death rate that is influenced by olfactory experience. Such experience-dependent apoptosis in the olfactory epithelium has been observed in experiments, for example Santoro and Dulac [24]. Since a neural circuit could estimate the elements of the response covariance matrix R_{ab} by taking temporal averages of products of glomerular activations (assuming ergodicity), the dynamics proposed here, or something like it, could be biologically implemented. A caveat is that our dynamical model involves taking the inverse of the response covariance matrix R^{-1} rather than using the covariance matrix itself. This means that an additional layer of processing is needed to calculate the required diagonal elements, R_{aa}^{-1}. Working out how this processing can be implemented in a neural system is interesting, but beyond the scope of this paper.

The total number of neurons $\sum_a K_a$ must be constrained for information to be bounded (and thus, for a maximum to exist). That constraint is enforced here by the experience-independent death rate λ (corresponding to the Lagrange multiplier from eq. (8)). If this is a constant, the final number of neurons will depend on both λ and the environmental statistics R. We can instead promote λ itself to a dynamical variable that is adjusted until the total number of neurons reaches a certain value K_{tot}:

$$\frac{d\lambda}{dt} = \beta \left[\sum_a K_a - K_{\text{tot}} \right] , \tag{12}$$

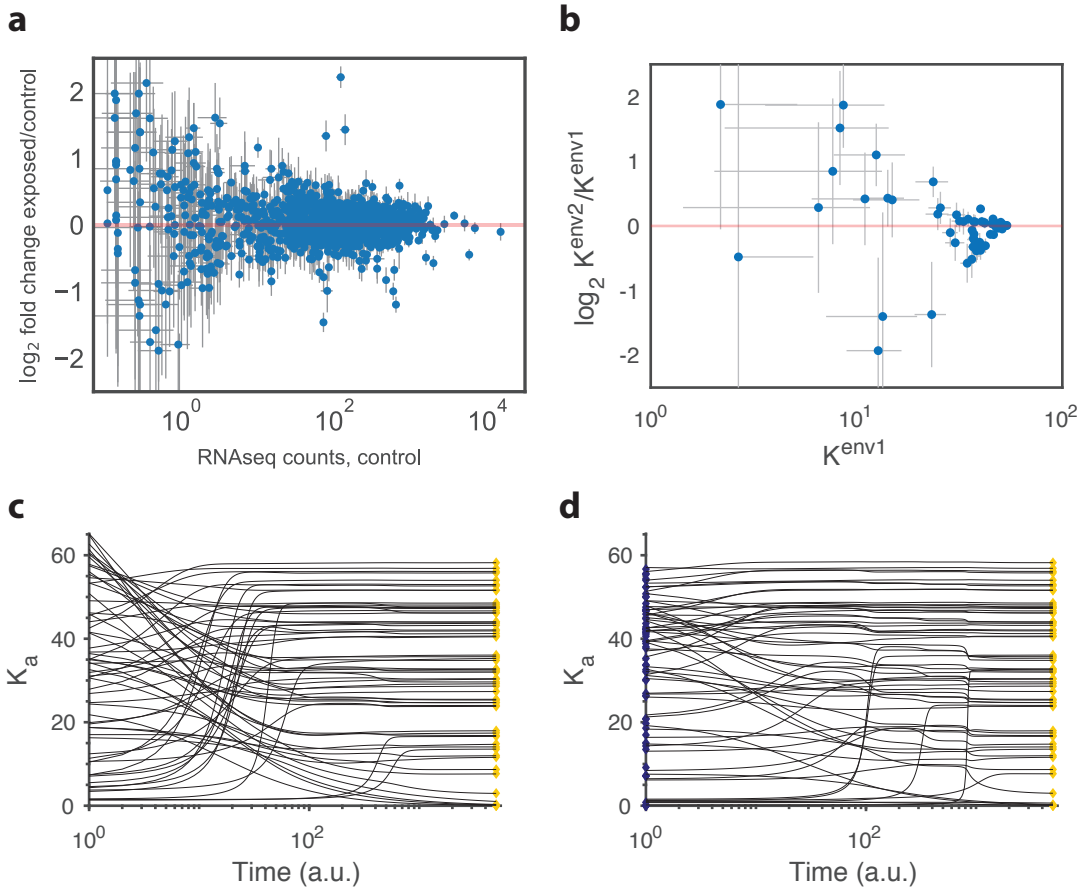

Figure 5: Qualitative comparison between experiment and theory regarding the changes in receptor distribution that occur with a change in environment; and an example of how our dynamical model converges to the optimal receptor distribution. **(a)** Reproduced using the raw data from Ibarra-Soria et al. [28], this shows the experimental data for the log-ratio between the receptor abundances in the test environment (where four odorants were added to the water supply) and those in the control environment, plotted against the values in the control conditions (on a log scale). The error bars show standard deviation across six individuals. Compared to Figure 5B in Ibarra-Soria et al. [28], this plot does not use a Bayesian estimation technique that shrinks ratios of abundances of rare receptors towards 1. **(b)** A similar plot produced in our model using mouse and human receptor response curves [43]. The error bars show the range of variation found in the optimal receptor distribution when slightly perturbing the initial and final environments (see the text). The simulation includes 62 receptor types for which response curves were measured [43], compared to 1115 receptor types assayed in Ibarra-Soria et al. [28]. **(c)** Example convergence curves in our dynamical model showing how the optimal receptor distribution (yellow diamonds) is reached from a random initial distribution of receptors. Note that the time axis is logarithmic. **(d)** Convergence curves when starting close to the optimal distribution from one environment (blue diamonds) but optimizing for another. A small, random deviation from the optimal receptor abundance in the initial environment was added to avoid the simulation getting trapped in a local optimum.

with β another learning rate. This picture envisions that the rate of apoptosis of olfactory sensory neurons is self-regulated by the population—one can readily imagine a a mechanism for this through the secretion of signaling factors that allow a sort of quorum sensing.

Figure 5c shows example convergence curves in our dynamical model when starting from a random initial distribution. The sensing matrix used here is based on mammalian data [43] and the environment covariance matrix is generated using the random procedure mentioned earlier (and described in more detail in the

Supplementary Information). We see that some receptor types take much longer than others to converge (the time axis is logarithmic, which helps visualize the whole range of convergence behaviors). Roughly-speaking, convergence is slower when the final receptor abundance is small, which is related to the fact that the rate of change dK_a/dt in eq. (10) vanishes in the limit $K_a \to 0$.

In Figure 5d, we show the convergence to the same final state, but this time starting from a distribution that is not random, but was optimized for a different environment. The initial and final environments are the same as those used in the previous section to compare the simulations to experimental findings, Figure 5b. Interestingly, many receptor types actually take longer to converge in this case compared to the random starting point, showing that getting trapped in local optima is a potential danger. This can be alleviated by adding a small amount of stochasticity to the dynamics, which in realistic situations would enter naturally since the response covariance matrix R would have to be estimated based on stochastic odor encounters and noisy receptor readings. In fact, in Figure 5d, we needed to add a small amount of noise (corresponding to $\pm 0.05 K_{\text{tot}}$) to the initial distribution to improve convergence and avoid getting trapped.

4 Discussion

We built a model for the distribution of receptor types in the olfactory epithelium that is based on efficient coding, and assumes that the abundances of different receptor types are adapted to the statistics of natural odors in a way that maximizes the amount of information conveyed to the brain by glomerular responses. We showed how this model predicts a non-uniform distribution of receptor types in the olfactory epithelium, as well as reproducible qualitative changes in receptor distribution as a result of perturbations in the odor environment. Both of these striking phenomena have been observed in recent experiments, and our model provides a potential conceptual framework for understanding the data. In our model, the sensitivity of the receptor distribution to changes in odor statistics is affected by the tuning of the olfactory receptors, with narrowly-tuned receptors being more readily affected by such changes than broadly-tuned ones. The model also predicts that environments that differ in the identity of the odors that are present will lead to larger differences in optimal receptor distribution than environments that differ only in the statistics with which these odors are encountered. Another prediction is that there is a monotonic relationship between the number of receptor types found in the epithelium and the total number of olfactory sensory neurons. All of these predictions can be tested in future experiments.

A quantitative comparison of the model with data was difficult because of the scarcity of the sorts of experiments needed to inform some of the parameters. More complete surveys of the response curves of olfactory receptors to larger panels of odorants in more organisms would help to fix the structure of odor-receptor interactions that is necessary for our model. More high-throughput measurements of the distribution of receptor types in the olfactory epithelium, as in Ibarra-Soria et al. [28], would enable precise tests of the predictions. However, the biggest missing piece at this point is a knowledge of the statistics of natural odors. Perhaps imaging studies focusing on glomeruli in behaving animals could provide a novel way to make such measurements. Or perhaps a mix of theory and experiment could provide a way to estimate naturalistic odor statistics even if these are not precisely matched to any particular environment.

Our model can be extended in several ways, with the caveat that each extension introduces more parameters, which require even more experimental data to fix accurately. The OSN responses do not depend linearly on odorant concentrations, but respond in more complex, nonlinear ways. Using a more realistic response function can have important effects on the way in which increasing neuron numbers relates to changes in the number of receptor types. We could also use an improved noise model, in which external sources of noise could be included, as opposed to just receptor noise. Also, the distribution of odorants in natural scenes might be more appropriately modeled using a Gaussian mixture instead of a normal distribution. This would imply the existence of many odor objects, each of which would be modeled by a Gaussian distribution with a certain concentration profile, and would more closely approximate the sparse nature of olfactory scenes [18]. Finally, the importance of a given glomerular response to an animal might be related not only to how much information this gives about the odor environment, but might be modulated by a value assignment [47]. Both theoretical and experimental work is needed to figure out how and if the adaptation of olfactory receptor distributions depends on value. All of these extensions we leave for future work.

Our model suggests the exciting possibility of a first causal test of the efficient coding hypothesis for

sensory coding. Given sufficiently-detailed information regarding receptor affinities and natural odor statistics, experiments could be designed that perturb the environment in specified ways, and then measure the change in olfactory receptor distribution. This could then be compared to the theoretically-predicted changes to provide a strong test of efficient coding ideas. This is in contrast to previous applications of the efficient coding hypothesis in vision and audition, which provided elegant *post-hoc* explanations for observed neural receptive fields but lacked an avenue for testing causality [9, 10, 11, 12, 13, 14, 15, 48, 40, 49, 41, 50, 51].

5 Acknowledgments

We would like to thank Joel Mainland for helpful discussions, and Elissa Hallem, Joel Mainland, and Darren Logan for olfactory receptor affinity data. During the completion of this project, VB was supported by Simons Foundation Mathematical Modeling in Living Systems grant 400425, and TT was supported by the Swartz Foundation.

References

[1] Chung Wen Yu, Lindsay A. Warrenburg, Katharine A. Prokop-Prigge, and Joel D. Mainland. Drawing the border of olfactory space. *Chemical Senses*, 40(7):565–565, 2015.

[2] Bettina Malnic, Junzo Hirono, Takaaki Sato, and Linda B. Buck. Combinatorial receptor codes for odors. *Cell*, 96(5):713–23, 1999.

[3] Linda Buck and Richard Axel. A novel multigene family may encode odorant receptors: A molecular basis for odor recognition. *Cell*, 65(1):175–187, 1991.

[4] Andrew Chess, Itamar Simon, Howard Cedar, and Richard Axel. Allelic inactivation regulates olfactory receptor gene expression. *Cell*, 78(5):823–834, 1994.

[5] Longzhi Tan, Qian Li, and X Sunney Xie. Olfactory sensory neurons transiently express multiple olfactory receptors during development. *Molecular systems biology*, 11(12):844, 2015.

[6] Christine Missbach, Hany KM Dweck, Heiko Vogel, Andreas Vilcinskas, Marcus C. Stensmyr, Bill S. Hansson, and Ewald Grosse-Wilde. Evolution of insect olfactory receptors. *eLife*, 2014(3):e02115, 2014.

[7] Horace B. Barlow. Possible principles underlying the transformations of sensory messages. *Sensory Communication*, pages 217–234, 1961.

[8] Johannes H. van Hateren. A theory of maximizing sensory information. *Biological Cybernetics*, 68(1): 23–29, 1992.

[9] Simon Laughlin. A Simple Coding Procedure Enhances a Neuron's Information Capacity. *Z. Naturforsch.*, 36(c):910–912, 1981.

[10] Joseph J. Atick and A. Norman Redlich. Towards a Theory of Early Visual Processing. *Neural Computation*, 2:308–320, 1990.

[11] Johannes H. van Hateren. Theoretical predictions of spatiotemporal receptive fields of fly LMCs, and experimental validation. *Journal of Comparative Physiology A*, 171:157–170, 1992.

[12] Bruno A. Olshausen and David J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(June):607–609, 1996.

[13] Eero P. Simoncelli and Bruno A. Olshausen. Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24:1193–1216, 2001.

[14] Adrienne L. Fairhall, Geoffrey D. Lewen, and William Bialek. Efficiency and ambiguity in an adaptive neural code. *Nature*, 412(August):787–792, 2001.

[15] Michael S. Lewicki. Efficient coding of natural sounds. *Nature neuroscience*, 5(4):356–363, 2002.

[16] Charles F. Stevens. What the fly's nose tells the fly's brain. *PNAS*, 112(30):9460–5, 2015.

[17] Yilun Zhang and Tatyana O. Sharpee. A Robust Feedforward Model of the Olfactory System. *PLoS computational biology*, 12(4):e1004850, 2016.

[18] Kamesh Krishnamurthy, Ann M. Hermundstad, Thierry Mora, Aleksandra M. Walczak, and Vijay Balasubramanian. Disorder and the neural representation of complex odors: smelling in the real world. *arXiv*, 2017.

[19] Andreas Keller and Leslie B. Vosshall. Influence of odorant receptor repertoire on odor perception in humans and fruit flies. *Proceedings of the National Academy of Sciences of the United States of America*, 104(13):5614–9, 2007.

[20] Carolyn S. McBride, Felix Baier, Aman B. Omondi, Sarabeth A. Spitzer, Joel Lutomiah, Rosemary Sang, Rickard Ignell, and Leslie B. Vosshall. Evolution of mosquito preference for humans linked to an odorant receptor. *Nature*, 515(7526):222–227, 2014.

[21] David Zwicker, Arvind Murugan, and Michael P. Brenner. Receptor arrays optimized for natural odor statistics. *Proceedings of the National Academy of Sciences*, 113(20):5570–5575, 2016.

[22] G. A. Monti Graziadei and P. P. C. Graziadei. Neurogenesis and neuron regeneration in the olfactory system of mammals. II. Degeneration and reconstitution of the olfactory sensory neurons after axotomy. *Journal of Neurocytology*, 8(2):197–213, 1979.

[23] James E. Schwob, Karen E. Mieleszko Szumowski, and Audrey A. Stasky. Olfactory sensory neurons are trophically dependent on the olfactory bulb for their prolonged survival. *The Journal of Neuroscience*, 12(10):3896–3919, 1992.

[24] Stephen W. Santoro and Catherine Dulac. The activity-dependent histone variant H2BE modulates the life span of olfactory neurons. *eLife*, 2012(1):1–32, 2012.

[25] Shaohua Zhao, Huikai Tian, Limei Ma, Ying Yuan, C. Ron Yu, and Minghong Ma. Activity-Dependent Modulation of Odorant Receptor Gene Expression in the Mouse Olfactory Epithelium. *PLoS ONE*, 8(7): e69862, 2013.

[26] Brian G. Dias and Kerry J. Ressler. Parental olfactory experience influences behavior and neural structure in subsequent generations. *Nature neuroscience*, 17(1):89–96, 2014.

[27] Hervé Cadiou, Imad Aoudé, Bassim Tazir, Adrien Molinas, Claire Fenech, Nicolas Meunier, and Xavier Grosmaitre. Postnatal Odorant Exposure Induces Peripheral Olfactory Plasticity at the Cellular Level. *Journal of Neuroscience*, 34(14):4857–4870, 2014.

[28] Ximena Ibarra-Soria, Thiago S. Nakahara, Jingtao Lilue, Yue Jiang, Casey Trimmer, Mateus A. A. Souza, Paulo H. M. Netto, Kentaro Ikegami, Nicolle R. Murphy, Mairi Kusma, Andrea Kirton, Luis R. Saraiva, Thomas M. Keane, Hiroaki Matsunami, Joel D. Mainland, Fabio Papes, and Darren W. Logan. Variation in olfactory neuron repertoires is genetically controlled and environmentally modulated. *eLife*, 6:e21476, 2017.

[29] John G. Hildebrand and Gordon M. Shepherd. Mechanisms of olfactory discrimination: converging evidence for common principles across phyla. *Annual Review of Neuroscience*, 20:595–631, 1997.

[30] Alison Maresh, Diego Rodriguez Gil, Mary C. Whitman, and Charles A. Greer. Principles of glomerular organization in the human olfactory bulb - Implications for odor processing. *PLoS ONE*, 3(7), 2008.

[31] Leslie B. Vosshall, Allan M. Wong, and Richard Axel. An olfactory sensory map in the fly brain. *Cell*, 102(2):147–159, 2000.

[32] Arbora Resulaj and Dmitry Rinberg. Novel Behavioral Paradigm Reveals Lower Temporal Limits on Mouse Olfactory Decisions. *Journal of Neuroscience*, 35(33):11667–11673, 2015.

[33] Shamik DasGupta and Scott Waddell. Learned odor discrimination in Drosophila without distinct combinatorial maps in the antennal lobe. *Current Biology*, 18(21):1668–1674, 2008.

[34] Michael Wehr and Gilles Laurent. Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature*, 384(November):162–166, 1996.

[35] Stephen J. Huston, Mark Stopfer, Stijn Cassenaer, Zane N. Aldworth, and Gilles Laurent. Neural Encoding of Odors during Active Sampling and in Turbulent Plumes. *Neuron*, 88:1–16, 2015.

[36] Karen J. Rossiter. Structure-odor relationships. *Chemical Reviews*, 96(8):3201–3240, 1996.

[37] Kobi Snitz, Adi Yablonka, Tali Weiss, Idan Frumin, Rehan M. Khan, and Noam Sobel. Predicting Odor Perceptual Similarity from Odor Structure. *PLoS Computational Biology*, 9(9), 2013.

[38] Manuel Zarzo and David T. Stanton. Identification of latent variables in a semantic odor profile database using principal component analysis. *Chemical Senses*, 31(8):713–724, 2006.

[39] Alexei A. Koulakov, Brian E. Kolterman, Armen G. Enikolopov, and Dmitry Rinberg. In search of the structure of human olfactory space. *Frontiers in Systems Neuroscience*, 5(September):1–8, 2011.

[40] Gasper Tkacik, Jason S. Prentice, Jonathan D. Victor, and Vijay Balasubramanian. Local statistics in natural scenes predict the saliency of synthetic textures. *Proceedings of the National Academy of Sciences of the United States of America*, 107(42):18149–18154, 2010.

[41] Ann M. Hermundstad, John J. Briguglio, Mary M. Conte, Jonathan D. Victor, Vijay Balasubramanian, and Gasper Tkacik. Variance predicts salience in central sensory processing. *eLife*, 10.7554:e03722, 2014.

[42] Elissa A. Hallem and John R. Carlson. Coding of odors by a receptor repertoire. *Cell*, 125(1):143–160, 2006.

[43] Harumi Saito, Qiuyi Chi, Hanyi Zhuang, Hiroaki Matsunami, and Joel D. Mainland. Odor Coding by a Mammalian Receptor Repertoire. *Science Signaling*, 2(60):ra9–ra9, mar 2009.

[44] Claude E. Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and*, 27(July), 1948.

[45] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[46] Anne L. Calof, Nobuko Hagiwara, J. David Holcomb, Jeffrey S. Mumm, and Jianyong Shou. Neurogenesis and cell death in olfactory epithelium. *Journal of Neurobiology*, 30(1):67–81, 1996.

[47] Olivier Rivoire and Stanislas Leibler. The Value of Information for Populations in Varying Environments. *Journal of Statistical Physics*, 142:1124–1166, 2011.

[48] Patrick Garrigan, Charles P. Ratliff, Jennifer M. Klein, Peter Sterling, David H. Brainard, and Vijay Balasubramanian. Design of a trichromatic cone array. *PLoS Computational Biology*, 6(2):e1000677, 2010.

[49] Charles P. Ratliff, Bart G. Borghuis, Yen-Hong Kao, Peter Sterling, and Vijay Balasubramanian. Retina is structured to process an excess of darkness in natural scenes. *PNAS*, 107(40):17368–17373, 2010.

[50] Stephanie E. Palmer, Olivier Marre, Michael J. II Berry, and William Bialek. Predictive information in a sensory population. *PNAS*, 112(22):6908–6913, 2015.

[51] Jared M. Salisbury and Stephanie E. Palmer. Optimal Prediction in the Retina and Natural Motion Statistics. *Journal of Statistical Physics*, 162(5):1309–1323, 2016.

[52] Daniel Lewandowski, Dorota Kurowicka, and Harry Joe. Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9):1989–2001, 2009.

A Supplementary Information

A.1 Choice of sensing matrices

We used three types of sensing matrices in this study. Two were based on experimental data, one using fly receptors [42], and one using mouse and human receptors [43]; and another sensing matrix was based on randomly-generated receptor affinity profiles. These can all be either directly downloaded from our repository on GitHub, `https://github.com/ttesileanu/OlfactoryReceptorDistribution`, or generated using the code available there.

Fly sensing matrix Some of our simulations used a sensing matrix based on *Drosophila* receptor affinities, as measured by Hallem and Carlson [42]. This includes the responses of 24 of the 60 receptor types in the fly against a panel of 110 odorants, measured using single-unit electrophysiology in a mutant antennal neuron. We directly used the values from Table S1 in Hallem and Carlson [42], normalized to obtain the desired SNR (high for Figure 2a, low for Figure 2b, and intermediate values for Figures 2c, 4a, and 4b).

The fly data has the advantage of being more complete than the equivalent datasets in mammals, but since olfactory sensory neurons don't regenerate in the fly, it is less useful for understanding the effects of changing environments.

Mammalian sensing matrix When comparing the results from our simulation to the experimental findings from Ibarra-Soria et al. [28], we used a sensing matrix based on mouse and human receptor affinity data from Saito et al. [43]. This was measured using heterologous expression of olfactory genes, and in total it tested 219 mouse and 245 human receptor types against 93 different odorants. However, only 52 mouse receptors and 10 human receptors exhibited detectable responses against any of the odorants, while only 63 odorants activated any receptors. From the remaining $62 \times 63 = 3906$ receptor–odorant pairs, only 335 (about 9%) showed a response, and were assayed at 11 different concentration points. In this paper, we used the values obtained for the highest concentration (3 mM).

The sparsity of the sensing matrix obtained from the Saito et al. [43] data makes convergence slow in our simulations. For this reason, we replaced the vanishing entries from this matrix with normally-distributed random variables, with a standard deviation of about 12% of the median value of the non-zero entries.

Random sensing matrices We generated random sensing matrices by first treating the columns of the sensing matrix (*i.e.,* the odorants) like equally-spaced points along a circle, with positions parametrized by a normalized angle $x = \theta/2\pi$. We then generated receptive fields with a Gaussian profile,

$$\phi(x) = \exp\left[-\frac{1}{2}\left(\frac{2\sin\pi(x-x_0)}{\sigma}\right)^2\right], \tag{S1}$$

where the expression in the exponent is the planar distance between a point with coordinate x and the center point of the Gaussian, with coordinate x_0. The centers x_0 for the Gaussian profiles for each of the receptors were chosen uniformly at random, and the standard deviation σ was a fixed parameter. As a final step, we shuffled the columns of the resulting sensing matrix to remove the circular ordering of the odorants.

The parameter σ is the tuning width which appears on the x-axis in Figures 2d and 2e, where it ranges from 0.01 to 0.2. In Figures 4c and 4d, the tuning parameters are 0.05 and 0.8, respectively.

A.2 Deriving the expression for the mutual information

Given that we assume a Gaussian distribution for odorant concentrations and that we approximate receptor responses as linear with additive Gaussian noise, eq. (2), it follows that the marginal distribution of receptor responses in natural olfactory scenes is also Gaussian. Taking averages of the responses, $\langle r_a \rangle$, and of products of responses, $\langle r_a r_b \rangle$, over both the noise distribution and the odorant distribution, and using eq. (2), we get a normal distribution of responses:

$$\boldsymbol{r} \sim \mathcal{N}(\boldsymbol{r}_0, R)\,, \tag{S2}$$

where the mean response vector $\boldsymbol{r}_0$ and the response covariance matrix R are given by

$$\begin{aligned} \boldsymbol{r}_0 &= \mathbb{K} S \boldsymbol{c}_0 \,, \\ R &= \big[\mathbb{I} + \mathbb{K} Q\big]\, \mathbb{K} \,. \end{aligned} \tag{S3}$$

Here, as in eq. (1) in the main text, $\boldsymbol{c}_0$ is the mean concentration vector, Γ is the covariance matrix of concentrations, and we use the overlap matrix from eq. (5), $Q = S\Gamma S^T$. Note that in the absence of noise, the response matrix is simply the overlap matrix Q modulated by the number of OSNs of each type, $R_{\text{noiseless}} = \mathbb{K} Q \mathbb{K}$.

The joint probability distribution over responses and concentrations, $P(\boldsymbol{r}, \boldsymbol{c})$, is itself Gaussian. To calculate the corresponding covariance matrix, we need the covariances between responses, $\langle r_a r_b \rangle - \langle r_a \rangle \langle r_b \rangle$, which are just the elements of the response matrix R from eq. (S3) above; and between concentrations, $\langle c_i c_j \rangle - \langle c_i \rangle \langle c_j \rangle$, which are the elements of the environment covariance matrix Γ, eq. (1). In addition, we need the covariances between responses and concentrations, $\langle r_a c_i \rangle - \langle r_a \rangle \langle c_i \rangle$, which can be calculated using eq. (2). We get:

$$(\boldsymbol{r}, \boldsymbol{c}) \sim \mathcal{N}\big((\boldsymbol{r}_0, \boldsymbol{c}_0), \Lambda\big) \,, \tag{S4}$$

with

$$\Lambda = \begin{pmatrix} R & \mathbb{K} S \Gamma \\ \Gamma S^T \mathbb{K} & \Gamma \end{pmatrix} . \tag{S5}$$

The mutual information between responses and odors is then given by (see the next section for a derivation):

$$I(\boldsymbol{r}, \boldsymbol{c}) = \frac{1}{2} \log \frac{\det \Gamma \det R}{\det \Lambda} \,. \tag{S6}$$

From eq. (S3) we have

$$\det R = \det(\mathbb{I} + \mathbb{K} Q) \det \mathbb{K} \,, \tag{S7}$$

and from eq. (S5),

$$\begin{aligned} \det \Lambda &= \det \begin{pmatrix} R & \mathbb{K} S \Gamma \\ \Gamma S^T \mathbb{K} & \Gamma \end{pmatrix} = \det \Gamma \cdot \det(R - \mathbb{K} S \Gamma\, \Gamma^{-1}\, \Gamma S^T \mathbb{K}) \\ &= \det \Gamma \cdot \det(\mathbb{K} + \mathbb{K} Q \mathbb{K} - \mathbb{K} S \Gamma S^T \mathbb{K}) \\ &= \det \Gamma \cdot \det \mathbb{K} \,, \end{aligned} \tag{S8}$$

where we used eq. (S3) again, and employed Schur's determinant identity (which we also derive in the next section). Thus,

$$I(\boldsymbol{r}, \boldsymbol{c}) = \frac{1}{2} \log \frac{\det \Gamma \cdot \det(\mathbb{I} + \mathbb{K} Q) \cdot \det \mathbb{K}}{\det \Gamma \cdot \det \mathbb{K}} = \frac{1}{2} \log \det(\mathbb{I} + \mathbb{K} Q) \,, \tag{S9}$$

recovering the result quoted in the main text, eq. (4).

A.3 Details regarding the derivation of the mutual information

Schur's determinant identity The identity for the determinant of a 2×2 block matrix that we used in eq. (S8) above can be derived in the following way. First, note that

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} \mathbb{I} & B \\ 0 & D \end{pmatrix} \begin{pmatrix} A - BD^{-1}C & 0 \\ D^{-1}C & \mathbb{I} \end{pmatrix} . \tag{S10}$$

Now, from the definition of the determinant it can be seen that

$$\det \begin{pmatrix} A & B \\ 0 & \mathbb{I} \end{pmatrix} = \det \begin{pmatrix} A & 0 \\ C & \mathbb{I} \end{pmatrix} = \det A \,, \tag{S11}$$

since all the products involving elements from the off-diagonal blocks must necessarily also involve elements from the 0 matrix. Thus, taking the determinant of eq. (S10), we get the desired identity

$$\det \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \det D \cdot \det(A - BD^{-1}C) \,. \tag{S12}$$

Mutual information for Gaussian distributions The expression (S6) for the mutual information $I(\boldsymbol{r}, \boldsymbol{c})$ can be derived by starting with the fact that I is equal to the Kullback-Leibler divergence from the joint distribution $P(\boldsymbol{r}, \boldsymbol{c})$ to the product distribution $P(\boldsymbol{r})P(\boldsymbol{c})$. As a first step, let us calculate the Kullback-Leibler divergence between two multivariate normals in n dimensions:

$$D = D_{\mathrm{KL}}(p\|q) = \int p(\boldsymbol{x}) \log \frac{p(\boldsymbol{x})}{q(\boldsymbol{x})}\, d\boldsymbol{x}\,, \tag{S13}$$

where

$$\begin{aligned} p(\boldsymbol{x}) &= \frac{1}{\sqrt{(2\pi)^n \det A}} \exp\left[-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu}_A)^T A^{-1}(\boldsymbol{x}-\boldsymbol{\mu}_A)\right]\,, \\ q(\boldsymbol{x}) &= \frac{1}{\sqrt{(2\pi)^n \det B}} \exp\left[-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu}_B)^T B^{-1}(\boldsymbol{x}-\boldsymbol{\mu}_B)\right]\,. \end{aligned} \tag{S14}$$

Plugging the distribution functions into the logarithm, we have

$$D = \frac{1}{2}\log\frac{\det B}{\det A} + \frac{1}{2}\int p(\boldsymbol{x})\Big[(\boldsymbol{x}-\boldsymbol{\mu}_B)^T B^{-1}(\boldsymbol{x}-\boldsymbol{\mu}_B) - (\boldsymbol{x}-\boldsymbol{\mu}_A)^T A^{-1}(\boldsymbol{x}-\boldsymbol{\mu}_A)\Big]\, d\boldsymbol{x}\,, \tag{S15}$$

where the normalization property of $p(\boldsymbol{x})$ was used. Using also the definition of the mean and of the covariance matrix, we have

$$\int p(\boldsymbol{x}) x_i\, d\boldsymbol{x} = \mu_{A,i}\,, \tag{S16a}$$

$$\int p(\boldsymbol{x}) x_i x_j\, d\boldsymbol{x} = A_{ij}\,, \tag{S16b}$$

which implies

$$\int p(\boldsymbol{x})\,(\boldsymbol{x}-\boldsymbol{\mu})^T C^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\, d\boldsymbol{x} = \mathrm{Tr}(AC^{-1}) + (\boldsymbol{\mu}_A-\boldsymbol{\mu})^T C^{-1}(\boldsymbol{\mu}_A-\boldsymbol{\mu}) \tag{S17}$$

for any vector $\boldsymbol{\mu}$ and matrix C. Plugging this into eq. (S15), we get

$$D = \frac{1}{2}\log\frac{\det B}{\det A} + \frac{1}{2}\left[\mathrm{Tr}(AB^{-1}) - n\right] + \frac{1}{2}(\boldsymbol{\mu}_A-\boldsymbol{\mu}_B)^T B^{-1}(\boldsymbol{\mu}_A-\boldsymbol{\mu}_B)\,. \tag{S18}$$

We can now return to calculating the KL divergence from $P(\boldsymbol{r}, \boldsymbol{c})$ to $P(\boldsymbol{r})P(\boldsymbol{c})$. Note that, since $P(\boldsymbol{r})$ and $P(\boldsymbol{c})$ are just the marginals of the joint distribution, the means of the variables are the same in the joint and in the product, so that the last term in the KL divergence vanishes. The covariance matrix for the product distribution is

$$\Lambda_{\mathrm{prod}} = \begin{pmatrix} R & 0 \\ 0 & \Gamma \end{pmatrix}\,, \tag{S19}$$

so the product inside the trace becomes

$$\Lambda\Lambda_{\mathrm{prod}}^{-1} = \begin{pmatrix} R & \dots \\ \dots & \Gamma \end{pmatrix}\begin{pmatrix} R^{-1} & 0 \\ 0 & \Gamma^{-1} \end{pmatrix} = \begin{pmatrix} \mathbb{I} & \dots \\ \dots & \mathbb{I} \end{pmatrix}\,, \tag{S20}$$

where the entries replaced by "…" don't need to be calculated because they drop out when the trace is taken. The sum of the dimensions of R and Γ is equal to the dimension, n, of Λ, so that the term involving the trace from eq. (S18) also drops out, leaving us with the final result:

$$I = D_{\mathrm{KL}}\left(p(\boldsymbol{r}, \boldsymbol{c})\|p(\boldsymbol{r})p(\boldsymbol{c})\right) = \frac{1}{2}\log\frac{\det R \det \Gamma}{\det \Lambda}\,, \tag{S21}$$

which is the same as eq. (S6), which we used in the previous section.

A.4 Deriving the KKT conditions for the information optimum

In order to find the optimal distribution of olfactory receptors, we need to maximize the mutual information from eq. (4), subject to constraints. Let us first calculate the gradient of the mutual information with respect to the receptor numbers:

$$\frac{\partial I}{\partial K_a} = \frac{1}{2}\frac{\partial}{\partial K_a}\log\det(\mathbb{I}+\mathbb{K}Q) = \frac{1}{2}\frac{\partial}{\partial K_a}\operatorname{Tr}\log(\mathbb{I}+\mathbb{K}Q)\,. \tag{S22}$$

The cyclic property of the trace allows us to use the usual rules to differentiate under the trace operator, so we get

$$\begin{aligned}\frac{\partial I}{\partial K_a} &= \frac{1}{2}\operatorname{Tr}\left[\frac{\partial \mathbb{K}}{\partial K_a}\left(Q^{-1}+\mathbb{K}\right)^{-1}\right] = \frac{1}{2}\sum_{b,c}\frac{\partial\left(K_b\delta_{bc}\right)}{\partial K_a}\left(Q^{-1}+\mathbb{K}\right)^{-1}_{ca} \\ &= \frac{1}{2}\left(Q^{-1}+\mathbb{K}\right)^{-1}_{aa}\,.\end{aligned} \tag{S23}$$

We now have to address the constraints. We have two kinds of constraints: an equality constraint that sets the total number of neurons, $\sum K_a = K_{\text{tot}}$; and inequality constraints that ensure that all receptor abundances are non-negative, $K_a \geq 0$. This can be done using the Karush-Tucker-Kuhn (KKT) conditions, which require the introduction of Lagrange multipliers, λ for the equality constraint, and μ_a for the inequality constraints. At the optimum, we must have:

$$\begin{aligned}\frac{\partial I}{\partial K_a} &= \frac{1}{2}\lambda\frac{\partial}{\partial K_a}\left(\sum_b K_b - K_{\text{tot}}\right) - \sum_b \mu_b\frac{\partial}{\partial K_a}K_b \\ &= \lambda - \mu_a\,,\end{aligned} \tag{S24}$$

where the Lagrange multipliers for the inequality constraints, μ_a, must be non-negative, and must vanish unless the inequality is saturated:

$$\begin{aligned}\mu_a &\geq 0\,, \\ \mu_a K_a &= 0\,.\end{aligned} \tag{S25}$$

Put differently, if $K_a > 0$, then $\mu_a = 0$ and $\partial I/\partial K_a = \lambda/2$; while if $K_a = 0$, then $\partial I/\partial K_a = \lambda/2 - \mu_a \leq \lambda/2$. Combined with eq. (S23), this recovers the conditions from the main text, eq. (8).

A.5 Deriving the high total neuron number approximation

Suppose we are in the regime in which the total number of neurons is large, and in particular, each of the abundances K_a is large. Then we can perform an expansion of the expression appearing in the KKT equations from eq. (8):

$$(A+\mathbb{K})^{-1} = \mathbb{K}^{-1}(\mathbb{I}+A\mathbb{K}^{-1})^{-1} \approx \mathbb{K}^{-1}(\mathbb{I}-A\mathbb{K}^{-1})\,, \tag{S26}$$

where we used the notation $A = Q^{-1}$. This implies

$$\lambda \approx \frac{1}{K_a} - \frac{A_{aa}}{K_a^2}\,. \tag{S27}$$

This quadratic equation has only one large solution, and it is given approximately by

$$K_a \approx \frac{1}{\lambda} - A_{aa}\,. \tag{S28}$$

Combined with the normalization constraint, $\sum_a K_a = K_{\text{tot}}$, this recovers eq. (9) from the main text.

A.6 First receptor type to be activated

When there is only one active receptor, $K_x = K_{\text{tot}}$, $K_{a\neq x} = 0$, the KKT conditions (8) are automatically satisfied. The receptor that is activated first can be found in this case by calculating the information $I(\boldsymbol{r}, \boldsymbol{c})$ using eq. (4) while assuming an arbitrary index x for the active receptor, and then finding $x = x^*$ that yields the maximum value. Without loss of generality, we can permute the receptor indices such that $x = 1$. We have, when $K_1 = K_{\text{tot}}$,

$$I_1(\boldsymbol{r}, \boldsymbol{c}) = \frac{1}{2}\operatorname{Tr}\log(\mathbb{I} + \mathbb{K}Q) = \frac{1}{2}\log\det(\mathbb{I} + \mathbb{K}Q) = \frac{1}{2}\log\begin{vmatrix} 1 + K_{\text{tot}}Q_{11} & K_{\text{tot}}Q_{12} & K_{\text{tot}}Q_{12} & \cdots & K_{\text{tot}}Q_{1M} \\ 0 & 1 & 0 & & 0 \\ \vdots & & & \ddots & \\ 0 & 0 & 0 & \cdots & 1 \end{vmatrix}$$
$$= \frac{1}{2}\log\bigl(1 + K_{\text{tot}}Q_{11}\bigr)\,. \tag{S29}$$

Thus, in general, the information when only receptor type x is activated is given by

$$I_x(\boldsymbol{r}, \boldsymbol{c}) = \frac{1}{2}\log\bigl(1 + K_{\text{tot}}Q_{xx}\bigr)\,, \tag{S30}$$

which implies that information is maximized when x matches the receptor corresponding to the highest diagonal value of the overlap matrix Q,

$$x^* = \arg\max Q_{xx}\,. \tag{S31}$$

Another way to think of this result is by employing the usual expression for the capacity of a single Gaussian channel, and then finding the channel that maximizes this capacity.

A.7 Random environment matrices

Generating plausible olfactory environments is difficult because so little is known about natural odor scenes. However, it is reasonable to expect that there will be some strong correlations. This could, for instance, be due to the fact that an animal's odor is composed of several different odorants in fixed proportions, and thus concentrations of these odorants will be correlated.

To generate covariance matrices that have a non-trivial correlation structure, we used a modified form of an algorithm based on partial correlations [52]. The basic idea is that there is a recurrence relation that allows one to eliminate the conditioning variables one by one, so that a set of unconditioned correlations can be obtained by recursively processing the partial correlations.

The distributions from which the partial correlations are drawn in Lewandowski et al. [52] are beta distributions with parameters that depend on the order of the correlations. These are chosen such that the resulting matrices provide a uniform sampling in the space of correlation matrices. However, this leads to samples that are close to diagonal and thus not particularly suitable for modeling olfactory environments. By simply keeping the beta distribution parameters equal to a given fixed value, $\alpha = \beta = \text{const}$, instead of changing depending on the order of the partial correlations, we can instead obtain matrices with a tunable amount of correlation. More specifically, large values of β lead to almost uncorrelated environments, while small values lead to environments with strong correlations between odorant concentrations. We found this modified algorithm in an entry on the question-and-answer website `Stack Exchange`, at `https://stats.stackexchange.com/q/125020`. This method generates a correlation matrix, in which all the variances are equal to 1. Random variances are then drawn from a lognormal distribution and the rows and columns are multiplied by their square roots, generating a final covariance matrix. The detailed code is available on GitHub, at `https://github.com/ttesileanu/OlfactoryReceptorDistribution`.

When comparing the qualitative results from our model against experiments in which the odor environment changes [28], we used small perturbations of the initial and final environments to estimate error bars on receptor abundances. To generate a perturbed covariance matrix, $\tilde{\Gamma}$, from a starting matrix Γ, we first took the matrix square root: a symmetric matrix M, which obeys

$$\Gamma = MM^T \equiv M^2\,. \tag{S32}$$

We then perturbed M by adding normally-distributed *i.i.d.* values to its elements,

$$\tilde{M}_{ij} = M_{ij} + \sigma\eta_{ij}\,, \tag{S33}$$

and recreated a covariance matrix by multiplying the perturbed square root with its transpose,

$$\tilde{\Gamma} = \tilde{M}\tilde{M}^T\,. \tag{S34}$$

This approach ensures that the perturbed matrix $\tilde{\Gamma}$ remains a valid covariance matrix—symmetric and positive-definite—which wouldn't be guaranteed if the random perturbation was added directly to Γ. We chose the magnitude σ of the perturbation so that the error bars in our simulations are of comparable magnitude to those in the experiments.

A.8 Deriving the dynamical model

To turn the maximization requirement into a dynamical model, we turn to a gradient ascent argument. Given the current abundances K_a, we demand that they change in proportion to the corresponding components of the information gradient, plus a Lagrange multiplier to impose the constraint on the total number of neurons:

$$\dot{K}_a = 2\alpha\left(\frac{\partial I}{\partial K_a} - \lambda\right) = \alpha\big[\big(Q^{-1} + \mathbb{K}\big)^{-1}_{aa} - \lambda\big]\,. \tag{S35}$$

The brain does not have direct access to the overlap matrix Q, but it could measure the response covariance matrix R from eq. (S3). Thus, we can write the dynamics as

$$\begin{aligned}\dot{K}_a &= \alpha\Big\{\big[Q(\mathbb{I} + \mathbb{K}Q)^{-1}\big]_{aa} - \lambda\Big\} = \alpha\Big\{\big[\mathbb{K}^{-1}(R\mathbb{K}^{-1} - \mathbb{I})\mathbb{K}R^{-1}\big]_{aa} - \lambda\Big\} \\ &= \alpha\Big\{K_a^{-1} - \lambda - R_{aa}^{-1}\Big\}\,.\end{aligned} \tag{S36}$$

These equations do not yet obey the non-negativity constraint on the receptor abundances. The divergence in the K_a^{-1} term would superficially appear to ensure that positive abundances stay positive, but there is a hidden quadratic divergence in the response covariance term, R_{aa}^{-1}; see eq. (S3). To ensure that all constraints are satisfied while avoiding divergences, we multiply the right-hand-side of eq. (S36) by K_a^2, yielding

$$\dot{K}_a = \alpha\big[K_a - K_a^2\big(\lambda + R_{aa}^{-1}\big)\big]\,, \tag{S37}$$

which is the same as eq. (10) which we used in the main text.

A.9 Multiple glomeruli with the same affinity profile

In mammals, the axons from neurons expressing a given receptor type can project to anywhere from 2 to 16 different glomeruli. Here we show that in our setup, information transfer only depends on the total number of neurons of a given type, and not on the number of glomeruli to which they project.

The key observation is that mutual information, eq. (3), is unchanged when the responses and/or concentrations are modified by monotonic transformations. In particular, linear transformations of the responses do not affect the information values. Suppose that we have a case in which two receptors p and q have identical affinities, so that $S_{pi} = S_{qi}$ for all odorants i. We can then form linear combinations of the corresponding glomerular responses,

$$\begin{aligned}r_+ &= r_p + r_q = (K_p + K_q)\sum_i S_{pi}c_i + \eta_p\sqrt{K}_p + \eta_q\sqrt{K}_q\,, \\ r_- &= K_q r_p - K_p r_q = \eta_p K_q\sqrt{K}_p - \eta_q K_p\sqrt{K}_q\,,\end{aligned} \tag{S38}$$

and consider a transformation that replaces (r_p, r_q) with (r_+, r_-). Since r_- is pure noise, *i.e.,* it does not depend on the concentration vector $\boldsymbol{c}$ in any way, it has no effect on the mutual information.

We have thus shown that the amount of information that M receptor types contain about the environment when two of the receptors have identical affinity profiles is the same as if there were only $M-1$ receptor

types. The two redundant receptors can be replaced by a single one with an abundance equal to the sum of the abundances of the two original receptors. (Remember that in eq. (2) we had assumed that the responses are normalized such that the variance of η is unity. This is automatically satisfied for the r_+ response from eq. (S38) because the sum of two Gaussian variables with the same mean is Gaussian itself and has a variance equal to the sum of the variances of the two variables. This means that the noise term in eq. (S38) is equivalent to $\eta_+\sqrt{K_p + K_q}$, with $\eta_+ \sim \mathcal{N}(0,1)$.)

A.10 Interpretation of diagonal elements of inverse overlap matrix

In the main text we saw that the diagonal elements of the inverse overlap matrix Q^{-1}_{aa} were well-correlated with the abundances of OSNs K_a. Recall that the overlap matrix was defined as $Q = S\Gamma S^T$, where S is the sensing matrix, and Γ is the environment covariance matrix (see eqns. (1) and (2)). In the main text, we suggested that a possible explanation for this correlation is based on the relation between inverse covariance matrices (also called precision matrices) and partial correlations.

More specifically, as noted around eq. (S3) above, the overlap matrix Q and the response covariance matrix are related: in particular, Q is equal to R when there is a single receptor of each type ($K_a = 1$) and there is no noise ($\eta = 0$ in eq. (2)). Thus, the inverse overlap matrix $A = Q^{-1}$ is effectively a precision matrix. Diagonal elements of a precision matrix are inversely related to the corresponding diagonal elements of the corresponding covariance matrix, but they are also monotonically-related to goodness-of-fit parameters that show how well each variable can be linearly predicted from all the others. Since receptor responses that either do not fluctuate much or whose values can be guessed based on the responses from other receptors are not very informative, we would expect that abundances K_a are low when the corresponding diagonal elements of the inverse overlap matrix A_{aa} are high, which is what we see. In the following we give a short derivation of the connection between the diagonal elements of precision matrices and goodness-of-fit parameters.

Let us thus work in the particular case in which there is one copy of each receptor and there is no noise, so that $Q = R$, *i.e.*, $Q_{ij} = \langle r_i r_j\rangle - \langle r_i\rangle\langle r_j\rangle$. Without loss of generality, we focus on calculating the first diagonal element of the inverse overlap matrix, A_{11}, where $A = Q^{-1}$. For notational convenience, we will also denote the mean-centered first response variable by $y \equiv r_1 - \langle r_1\rangle$, and the subsequent ones by $x_a \equiv r_{a+1} - \langle r_{a+1}\rangle$. Then the covariance matrix Q can be written in block form

$$Q = \begin{pmatrix} \langle y^2\rangle & \langle y\boldsymbol{x}^T\rangle \\ \langle y\boldsymbol{x}\rangle & M \end{pmatrix}, \tag{S39}$$

where M is

$$M = \langle \boldsymbol{x}\boldsymbol{x}^T\rangle\,. \tag{S40}$$

Using the definition of the inverse together with Laplace's formula for determinants, we get

$$A_{11} = \frac{\det M}{\det Q}\,. \tag{S41}$$

Using the Schur determinant identity (derived above) on the block form (eq. (S39)) of the matrix Q,

$$\begin{aligned} A_{11} &= \frac{\det M}{\det M \cdot \det\big[\langle y^2\rangle - \langle y\boldsymbol{x}^T\rangle M^{-1}\langle y\boldsymbol{x}\rangle\big]} \\ &= \frac{1}{\langle y^2\rangle - \langle y\boldsymbol{x}^T\rangle M^{-1}\langle y\boldsymbol{x}\rangle}\,, \end{aligned} \tag{S42}$$

where we used the fact that the argument of the second determinant is a scalar.

Now, consider approximating the first response variable y by a linear function of all the others:

$$y = \boldsymbol{a}^T\boldsymbol{x} + q\,, \tag{S43}$$

where q is the residual. Note that we do not need an intercept term because we mean-centered our variables, $\langle y\rangle = \langle x\rangle = 0$. Finding the coefficients $\boldsymbol{a}$ that lead to a best fit (in the least-squares sense) requires minimizing the variance of the residual, and a short calculation yields

$$\boldsymbol{a}^* = \arg\min_{\boldsymbol{a}} \langle q\rangle^2 = \arg\min_{\boldsymbol{a}} (y - \boldsymbol{a}^T\boldsymbol{x})^2 = M^{-1}\langle y\boldsymbol{x}\rangle\,, \tag{S44}$$

where M is the same as the matrix defined in eq. (S40).

The coefficient of determination ρ^2 is defined as the ratio of explained variance to total variance of the variable y,

$$\begin{aligned}\rho^2 &= \frac{\langle(\boldsymbol{a}^{*T}\boldsymbol{x})^2\rangle}{\langle y^2\rangle} = \frac{\boldsymbol{a}^{*T}\langle\boldsymbol{x}\boldsymbol{x}^T\rangle\boldsymbol{a}^*}{\langle y^2\rangle} = \frac{\langle y\boldsymbol{x}^T\rangle M^{-1}MM^{-1}\langle y\boldsymbol{x}\rangle}{\langle y^2\rangle} \\ &= \frac{\langle y\boldsymbol{x}^T\rangle M^{-1}\langle y\boldsymbol{x}\rangle}{\langle y^2\rangle}\,.\end{aligned} \tag{S45}$$

Comparing this to eq. (S42), we see that

$$A_{11} = \frac{1}{\langle y^2\rangle}\frac{1}{1-\rho^2}\,, \tag{S46}$$

showing that indeed the diagonal elements of the precision matrix are monotonically-related to the goodness-of-fit parameter ρ^2 that indicates how well the corresponding variable can be linearly predicted by all the other variables. In addition, the inverse dependence on the variance of the response $\langle y\rangle^2$ shows that variables that do not fluctuate much (low $\langle y\rangle^2$ and thus high A_{11}) tend to have low abundances.