What the odor is not: Estimation by elimination

Vijay Singh[a,b,c], Martin Tchernookov[d], Vijay Balasubramanian[a,c]

[a]*Computational Neuroscience Initiative, University of Pennsylvania, Philadelphia, PA, 19104, USA*
[b]*Department of Physics, North Carolina A&T State University, Greensboro, NC, 27410, USA*
[c]*Department of Physics, University of Pennsylvania, Philadelphia, PA, 19104, USA*
[d]*Department of Physics, University of Wisconsin, Whitewater, WI, 53190, USA*

Abstract: The olfactory system uses a small number of broadly sensitive receptors to combinatorially encode a vast number of odors. We propose that the brain decodes this distributed representation by exploiting a statistical fact: receptors that do not respond to an odor carry more information than receptors that do because they signal the absence of all odorants that bind to them. Thus, it is easier to identify what the odor *is not*, rather than what the odor is. For biologically realistic numbers of receptors, response functions, and odor complexity, this method of elimination turns an underconstrained decoding problem into a solvable one, allowing accurate determination of odorants in a mixture and their concentrations. A neural network realization of our algorithm resembles circuit architecture in the piriform cortex. We propose several experimental tests and show an excellent match to data from olfactory discrimination experiments. Our theory also suggests why animals from the fly to the elephant have a few hundred olfactory receptor types, give or take a small factor.

1. Introduction

The olfactory system enables animals to sense, perceive, and respond to mixtures of volatile molecules that carry messages about the world. There are many monomolecular odorants, perhaps 10^4 or more [1, 2, 3], far more than the number of receptor types available to animals (~ 50 in fly, ~ 300 in human, ~ 1000 in rat, mouse and dog [4, 5, 6, 7]). The problem of representing such a high-dimensional chemical space in such a low-dimensional response space may be solved by the presence of many receptors that bind to numerous odorants [8, 9, 10, 11, 12, 13, 14], leading to a distributed, compressed, and combinatorial representation [12, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27].

We focus on the inverse problem: the estimation of odor composition from the response of olfactory receptors. We use a realistic competitive binding model of odor encoding by receptors [28, 29, 30, 31], and propose a scheme to estimate odor composition from such responses. The scheme works over a large range of biologically relevant parameters, does not require any special constraints on receptor-odorant interactions, and works for systems with few receptors, suggesting why the relatively small olfactory receptor repertoires of most organisms are sufficient for sensing complex natural odors. Our results provide performance bounds on behaviorally important tasks such as detecting the presence or absence of an odorant in a mixture [32], and discriminating between mixtures that differ in a some components [33, 34].

Our main idea is that a receptor that *does not* respond to an odor carries a lot more information about the odor than a receptor that *does* respond to it. This is because a receptor that does not respond to an odor signals that none of the odorants (individual chemicals) that could bind to this receptor are present. With just a few such non-responding receptors, most of the odorants that are not present can be identified and eliminated. Thus, it is easier to identify what the mixture *is not*, rather than what the

mixture is. For a large range of biologically relevant parameters, this elimination turns the estimation of odor concentration from an underdetermined problem to an overdetermined one. Thus, the concentration of the rest of the odorants can be estimated from how the remaining receptors respond. The specifics of the second step depend on the receptor encoding model.

We propose an experimental test of our model and show that it reproduces existing results on the olfactory cocktail party problem in mice [32]. We make additional predictions including for odor complexity threshold (number of odorants) at which human and mouse performance in this task will fall to chance levels. We then develop a neural network model based on the known structure of the olfactory system to decode odors from receptor responses. Stage by stage, the behavior of this network reflect the properties of our decoding model suggesting that olfactory system is implementing similar principles, subject to the constraints of neural circuitry. We developed our scheme with the olfactory system in mind but the method is general, and can be used for other types of chemical detection systems, such as electric noses [35, 36] for odor detection, and in medical tests.

2. Results

2.1. Identifying odorant presence

First, suppose we seek to simply identify the presence or absence of components of an odor and not their concentrations. This scenario can either arise because the behavioral task requires it, or because noise in receptor activity is high. In the latter case, binding dynamics is highly stochastic and the exact binding state is hard to predict; instead, receptor activation is determined by a threshold set by noise. In either case, if the concentration of an odorant is high enough for the activity of a receptor to be above threshold, the odorant is considered to be present and the receptor is said to respond. Otherwise, the odorant is not present and the receptor is considered inactive (no response). The main features of our decoding scheme can be explained in this simple model of receptor activity. A more realistic competitive binding (CB) model will be considered in the next section.

Consider a mixture of N_{L} odorants represented by the binary vector $\boldsymbol{c} = (c_1, c_2, \ldots, c_{N_{\mathrm{L}}})$, where $c_i = 1$ represents the presence of the i'th odorant. Let odors have complexity K: on average only K odorants are present in the mixture. These odorants bind to N_{R} receptors whose response is given by the binary vector $\boldsymbol{R} = (R_1, R_2, \ldots, R_{N_{\mathrm{R}}})$. Receptor sensitivity to the odorants is given by a matrix S. $S_{ij} = 1$ indicates that the odorant j can bind to receptor i and $S_{ij} = 0$ means it can not. Suppose that the probability that an odorant binds to a receptor is s, i.e., $P(S_{ij} = 1) = s$. Then, on average, each odorant binds to $s N_{\mathrm{R}}$ receptors and each receptor to $s N_{\mathrm{L}}$ odorants.

In this model a receptor will respond ($R_i = 1$) when stimulated by an odor that contains at least one odorant that binds to this receptor (i.e., $R_i = 1$ if, for some j, $c_j = 1$ & $S_{ij} = 1$). If no such odorant is present, the receptor is inactive ($R_i = 0$). The receptor thus acts as a binary 'OR' gate. This is an approximation of the biophysical competitive binding (CB) model [28, 29, 30, 31], which assumes that only one odorant molecule can bind to a receptor binding site at a time. The competition between the molecules to occupy the receptor binding site results in a sigmoidal "Hill" type receptor response function (see section: Identifying odorant concentration). The 'OR' gate can be considered as the response function of the sigmoidal CB model with a narrow transition region. This approximation can be used to model experiments where individual odorants have high concentrations when they are present (see section: Noisy decision making and comparison with experiments). We consider this simplified model as it captures the essential features of our decoder.

Odors encoded in this way can be decoded (odor estimate $\hat{\boldsymbol{c}}$) in two simple steps (Fig. 1):

1 First, all inactive receptors are identified and all odorants that bind to these inactive receptors are considered absent from the odor ($\hat{c}_j = 0$ for all j for which $S_{ij} = 1$ while $R_i = 0$).

2 The remaining odorants are considered to be present.

This simple decoder does not miss any odorant that is present in the mixture. This is because, assuming every odorant binds to at least one receptor, if an odorant is present ($c_j = 1$) all receptors that bind to it will respond. Hence, it will be identified as present.

However, false positives are possible. This is because an odorant could be absent, while all receptors that bind to it have a non-zero response because of other odorants present in the mixture. Thus, through lack of evidence, the odorant j will be declared present, giving a false positive. We can derive an approximate expression for the probability of such false positives (see SI: Identifying odorant presence):

$$P(\hat{c}_i = 1 | c_i = 0) \approx e^{-sN_\mathrm{R}e^{-sK}}. \tag{1}$$

We can also derive an approximate expression for the probability of correct estimation assuming that each odorant can be estimated independently of the others (SI: Identifying odorant presence Eq. S8):

$$P(\hat{\boldsymbol{c}} = \boldsymbol{c}) \approx \left(1 - N_\mathrm{L}e^{-sN_\mathrm{R}e^{-sK}}\right). \tag{2}$$

The second term is approximately the probability of false positives if there are N_L possible odorants in the environment.

These relations provide an intuitive understanding of the decoding process. For correct decoding the probability of false positives should be low. So, the term in the exponent of Eq. 1 should be large, i.e., sN_R should be large and sK small. This makes sense. sN_R is the average number of receptors that an odorant binds to. If the odorant does not bind to any receptor, then its concentration can not be estimated. Thus, sN_R should be large so that there are many receptors whose non-response can provide evidence of the absence of an odorant. Also, for successful decoding, sufficiently many receptors must be inactive to eliminate all odorants that are absent. For this to happen, the probability that any particular receptor responds to at least one of the K odorants in the mixture should be small. This probability is approximately sK when the likelihood s that a given odorant binds to a given receptor is small, and so we require that $sK < 1$.

The conditions $sN_\mathrm{R} > 1$ and $sK < 1$ will be necessary for any receptor response model and its decoder. This is because, in general, if an odorant does not bind to any receptor, its concentration cannot be estimated, while the process of converting an under-determined problem into a well determined one is going to require sufficiently many inactive receptors. Put otherwise, for a fixed number of receptors N_R, odorants in the environment N_L, and odor component complexity K, receptor sensitivity should be sufficiently high to ensure coverage of odorants, but small enough to avoid false positives.

These considerations can be combined with the observed sensitivity of olfactory receptor responses to odorants ($s \sim 5\%$ for mammals [10]). Thus, for typical mammalian parameters ($\{N_\mathrm{L}, K, N_\mathrm{R}, s\} \sim \{10^4, 10, 500, 0.05\}$), the estimated probability of false positives is low ($P(\hat{c}_i = 1 | c_i = 0) \sim 10^{-7}$; Eq. 1) and the probability of correct estimate is high ($P(\hat{\boldsymbol{c}} = \boldsymbol{c}) \sim 0.998$; Eq. 2). These estimates should be regarded as upper bounds on performance that will be degraded by noise in the sensory and decision making circuitry. We will return to this below when we compare with experiments, by extracting the noise level directly from data and using it to discount the predicted accuracy.

We numerically estimated the accuracy of our decoding scheme. First, we randomly chose the elements of a sensitivity matrix so that they were non-zero with probability s, i.e. ($P(S_{ij} > 0) = s$). Fig. 1c shows the probability of a correct estimate ($P(\hat{\boldsymbol{c}} = \boldsymbol{c})$) as a function of the number of receptors (N_R) for odors that contain on average $K = 10$ odorants drawn from $N_L = 10000$ possibilities. When the number of receptors N_R is too low, the probability of correct decoding is zero. As the number of receptors increases, there is a transition to a region where recovery is nearly perfect, with probability approaching 1. The transition is sharp and occurs at N_R much smaller than the number of possible odorants (N_L). Over a wide range of odor complexities (K) and numbers of responsive receptors ($s * N_\mathrm{R}$), our decoding scheme shows excellent performance ($P(\hat{\boldsymbol{c}} = \boldsymbol{c}) \approx 1$) (Fig. 1d). Our analytical expression for the probability of correct decoding (SI Eq. S8) gives a good description of the numerical results, and an excellent estimate

(a) Encoding

(b) Decoding

Receptors   Random Sparse Sensing Matrix   Odor

(c) $N_{\mathrm{L}} = 10000$, $s = 0.05$, $K = 10$

(d) $N_{\mathrm{L}} = 100000$ $N_{\mathrm{R}} = 1000$

Binary Decoder

Analytical Eq. S8

Figure 1. Identifying odorant presence: (**a**) Encoding: For the odorant presence decoder, the odorant mixture and receptor activity are represented using binary vectors. If the concentration of an odorant is high enough for the activity of a receptor to be above threshold, the odorant is considered to be present and the receptor is said to respond; otherwise the odorant is not present and the receptor is non-responsive. The odorants that are present in the mixture and the receptors that are active are represented with filled (blue) elements of the corresponding binary vectors. In this example, odorant 1 and 8 are present in the mixture and receptors 3 and 4 respond to the mixture. The receptor response is obtained using the receptor-odorant interaction matrix whose filled (blue) elements indicate the odorants that binds to a receptor. In this example, receptor 1 binds to odorant 2 and 5; receptor 2 to odorant 3 and 7; and so on. A receptor is active if the odor mixture contains at least one odorant that binds to the receptor. Otherwise the receptor is inactive. (**b**) Decoding: (Step 1) Odorants that are absent from the mixture are identified from receptors that are inactive. (Step 2) All remaining odorants are considered to be present. (**c**) The probability of correct decoding ($P(\hat{\boldsymbol{c}} = \boldsymbol{c})$) plotted as a function of the number of receptors (N_{R}). Markers are numerical simulations, smooth curve is the approximate analytical result (Eq. S8). $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ was measured as the fraction of correct decodings over 1000 trials, each conducted with a random choice of odor mixture and sensitivity matrix (see SI). Mean and error bar (± 1 standard deviation) were computed over 10 replicate simulations, each with 1000 trials. (**d**) $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ plotted as a function of the average number of odorants present in the odor mixture (K) and the average number of receptors responding to an odorant ($s * N_{\mathrm{R}}$), with the number of receptors (N_{R}) and the number of odorants (N_L) held fixed. We are plotting $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ as a function of the product $s * N_{\mathrm{R}}$, since the parameters s and N_{R} always appear in this combination (Eq. 2 and Eq. S8). $s * N_{\mathrm{R}}$, the average number of receptors a typical odorant binds to, determines whether the odorant can be detected or not. $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ was calculated numerically over 10000 trials, each with a random choice of odor mixture and sensitivity matrix. The smooth white curve is the approximate boundary of the transition region of $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ estimated by setting the analytical expression Eq. 2 to half.

of the transition point between poor and excellent decoding as a function of odor complexity and the receptor sensitivity (solid lines in Fig. 1c,d; and SI Fig. S1). As we will see below, including noise will quantitatively reduce decoding accuracy, matching experiments, but the prediction of a sharp threshold between poor and good performance remains the same.

2.2. Identifying odorant concentrations

When noise is low, or integration times are long, the activity of a receptor can be considered as a numerical continuum, along with the odorant concentrations (c_i), receptor sensitivities (S_{ij}), and receptor responses (R_j). In this case, the decoder can be modified to estimate not just which odorants are present in the mixture, but also their concentrations. While the details of the decoding scheme depend on the encoding mechanism of the receptor, the main principle remains the same. First, inactive (below threshold) receptors are used to eliminate some odorants. Then, the active receptors are used to estimate the remaining concentrations.

Receptor responses can be realistically described by a competitive binding (CB) model in which odorant molecules compete to occupy the receptor binding site [28]. The response to a mixture of odorants with concentrations c_i is given by a Hill-type function [28, 31, 29, 30]:

$$R = \frac{\sum_{j=1}^{N_{\mathrm{L}}} x_j}{\left(1 + d * \sum_{j=1}^{N_{\mathrm{L}}} x_j\right)}. \tag{3}$$

Here $x_j = S_{ij}c_j$ and d parameterizes the affinity of odorants for the receptor. This CB model approximates to the binary model of the previous section for large values of d ($R \to 1/d$) and to the commonly used linear response model when $d \to 0$ [37, 38, 39, 40]. As discussed in [28], additional odor-receptor interactions such as synergy, suppression, antagonism, inhibition, etc., can readily be included in the receptor response model. The prevalence of such interaction is only beginning to be investigated on large scales [41]. When the statistics of such odor-receptor interactions are known, they can be included in our response model (3). We discuss this further below.

As discussed earlier, the odorant presence decoder works because an inactive receptor implies that all odorants that can bind to this receptor are absent. The concentrations of such odorants are set to zero, while concentrations of the remaining of the odorants are set to 1. Thus the success of the binary decoder depends on ensuring that for every odorant that is absent, there is at least one receptor that does not respond. In the continuous case, a weaker condition is sufficient. The starting point is an under-determined identification problem because the number of possible odorants exceeds the number of receptors ($N_L > N_R$). In the first step, we eliminate odorants that bind to receptors with below-threshold responses. This leaves $\tilde{N}_R$ active receptors responding to $\tilde{N}_L$ candidate odorants. If $\tilde{N}_{\mathrm{L}} \leq \tilde{N}_{\mathrm{R}}$ the problem is now over-determined and can be solved (Fig. 2), even if all the absent odorants have not been eliminated. Specifically, the odor encoding functions (Eq. 3) give a set of coupled equations that relate the $\tilde{N}_L$ odorant concentrations to the $\tilde{N}_R$ receptor responses. These equations can be inverted to get the unknown concentrations.

Our decoder will not eliminate any of the K odorants that are present in a mixture because all of them will evoke responses. To estimate the number of false positives from the remaining $N_L - K$ odorants, let s be the probability that a given receptor is sensitive to a given odorant ($P(S_{ij} > 0) = s$). Then, the number of active receptors will be about $\tilde{N}_R \sim sKN_{\mathrm{R}}$ while the number of inactive receptors will be approximately $(1 - sK)N_R$. The first inactive receptor eliminates approximately a fraction s of the remaining $N_L - K$ odorants. The second inactive receptor removes roughly another fraction s of the remaining $(1-s)(N_L-K)$ odorants. Summing over these eliminations for all $(1-sK)N_R$ inactive receptors leaves a total of $\tilde{N}_L \sim K + (N_L-K)(1-s)^{N_R(1-sK)-1}$ odorants under consideration. Typical parameters $\{N_{\mathrm{L}}, K, N_{\mathrm{R}}, s\} = \{10^4, 10, 500, 0.05\}$ give $\tilde{N}_L \sim K = 10$ which is less than $\tilde{N}_R \sim sKN_{\mathrm{R}} = 250$, showing that in the biologically relevant regime our elimination algorithm leads to an over-determined and hence solvable identification problem.

Figure 2. Identifying odorant concentrations: (**a**) Encoding: Odor mixture is represented as a vector with N_L components. The K non-zero components of the vector are represented by the filled (blue) elements, with the degree of saturation of the color representing the concentration levels. In this example, the odor has three non-zero components (2, 5 & 8). The elements of the sensing matrix indicate the strength of receptor-odor interaction. Each row corresponds to one receptor, each column an odorant. The colored elements indicate the odorants to which the receptor binds. For example, the first receptor responds to odorant 2, 4 and 8; the second to odorant 1, 6 and 9; and so on. The sensing matrix and the odor vector can be used to determine the response of the receptors using the response function. Color saturation levels represent concentrations of odorants. Colors in the sensing matrix represent receptor binding affinity to the corresponding odorant. Receptor activity is represented by colors, white = inactive. (**b**) Decoding (Step 1 Elimination): Inactive receptors are used to eliminate absent odorants, reducing an initially under-determined problem to a well-defined one. (Step 2 Estimation): Concentration of the remaining odorants can be estimated from responses of the remaining receptors. (**c**) Same as Fig. 1(c), now for the continuous decoder. Mean and error bar (± 1 standard deviation) were computed over 10 replicate simulations, each with 1000 trials. The free parameter γ in the analytical formula (Eq. 4) was chosen to minimize mean squared error between the probability obtained numerically and the formula. (**d**) $P(\hat{c} = c)$ plotted as a function of number of odorants in the odor mixture (K) and $s * N_R$ at fixed N_R. $P(\hat{c} = c)$ was calculated numerically over 1000 trials, each with a random choice of odor mixture and sensitivity matrix. The smooth curve in white is the approximate boundary of the perfect decoding region, estimated by setting Eq. 4 to 0.5 ($\gamma = 3$). (**e**) Error in concentration estimates, defined as Euclidean distance between actual and estimated odor concentration divided by the number of odorants ($(||\hat{c} - c||_2/K)$), plotted as a function of number of odorants (K) and $s * N_R$ at fixed N_R. The error is small even when recovery is not perfect. Other measures of error lead to similar results (SI Fig. S3: Other measures of estimation error).

We can derive an approximate analytical expression for the probability of correct estimation (SI: Identifying odorant concentrations). This derivation assumes that the typical number of receptors that respond to a mixture is larger than the average odor complexity ($\tilde{N}_{\mathrm{R}} > K$), while, at the same time, enough receptors are inactive to eliminate absent odorants. This requires $s(N_{\mathrm{R}} - \tilde{N}_{\mathrm{R}}) > \gamma$, where $\gamma > 1$ is a parameter that depends on the receptor response model (details in SI: Identifying odorant concentrations). With these assumptions,

$$P(\hat{\boldsymbol{c}} = \boldsymbol{c}) \sim P(\tilde{N}_{\mathrm{R}} > K) * P(N_{\mathrm{R}} - \tilde{N}_{\mathrm{R}} > (\gamma/s)) = \left[1 - \Phi\left(\frac{K - \tilde{N}_{\mathrm{R}}}{\sqrt{\tilde{N}_{\mathrm{R}}}} \right) \right] \Phi\left(\frac{N_{\mathrm{R}} - \tilde{N}_{\mathrm{R}} - \frac{\gamma}{s}}{\sqrt{\tilde{N}_{\mathrm{R}}}} \right) . \quad (4)$$

Φ is the standard normal cumulative distribution function.

To numerically estimate the probability of correct decoding ($P(\hat{\boldsymbol{c}} = \boldsymbol{c})$), we generated sparse odor vectors with K odorants on average. Concentrations were drawn from a uniform distribution on the interval $[0, 1)$. Elements of the sensitivity matrix were chosen to be non-zero with probability s, i.e., ($P(S_{ij} > 0) = s$). The values of these non-zero elements were chosen from a log uniform distribution (SI: Numerical Simulations; similar results with other distributions in SI Fig. S2). When there are very few receptors (N_{R}), the probability of correct decoding is zero. But the probability transitions sharply to finite values at a threshold N_{R} which is much smaller than the number of possible odorants (N_{L}) (Fig. 2c). Odor compositions are recovered well for a wide range of parameters (Fig. 2d,e), so long as receptors are sufficiently sensitive $s * N_{\mathrm{R}} > 6$. Odors with the highest complexity are decoded when $s * N_{\mathrm{R}} \sim 10 - 15$. We quantified the error in odor estimates and found that even when decoding is not perfect there is a large parameter space where the error is small (Fig. 2e). These results have a weak dependence on the total number of odorants (SI Fig. S4).

Since humans have about 300 receptors, our model predicts that odors with most components can be decoded with $s \sim 3 - 5\%$ so that $sN_{\mathrm{R}} = 10 - 15$. This is consistent with observation – human receptors have $s \sim 4\%$ [10]. For *Drosophila*, which has ~ 50 receptors, the observed sensitivity of $s \sim 14\%$ [42] gives $sN_{\mathrm{R}} \sim 7$, in the expected range for successful decoding.

While we have assumed a CB model of odor encoding, the decoder can be modified to incorporate other biophysical interactions between odorants and receptors. If non-competitive interactions are known for a receptor-odor pair, the corresponding response model can be used instead of the CB model (Eq. 3). Odorant suppression of receptor responses is easily included in the algorithm. If responses are attenuated but not completely silenced, the algorithm goes through as before, albeit with a response model that includes the attenuation. So long as odor-receptor interactions as sparse as experiments suggest [9, 10], receptor silencing can be accounted for before the elimination step by removing silent receptors that bind to suppressive odorants for which we have evidence from other receptor responses.

2.3. Noisy decision making and comparison with behavioral experiment

We can compare the predictions of our decoder with animal performance in behavioral studies. In one such experiment [32], mice were presented with an odor (a mixture of odorants) and were trained to report the presence of a target odorant (details in SI: Comparison to experiment). The odors were formed by mixing up to 14 odorants chosen from a panel of 16. Any one of the 16 odorants could be either present or absent in the mixture. Every odorant was added at the same concentration independently of the other odorants. Performance of mice at identifying a target odorant in the mixture was tested for each of 16 odorants. Irrespective of the target, mice reported presence or absence of the odorant with a high probability (over 80%). Such high performance in reporting the presence or absence of the 16 possible constituent odorants suggests that mice could identify all components of the mixture. We can compare the predictions of our decoder with mouse performance on this task. Since the concentration levels and the decision variable were both binary, we will use the binary decoder of odor identity, modified to just report presence or absence of a target odorant, for this comparison.

Figure 3. Experimental Test: We predict performance of mice in the the olfactory cocktail-party task of detecting odorants in K-component mixtures [32] using a version of our odorant presence decoder (details in text). (a) Schematic of the two step noisy decoder. (b) We assume a normally distributed decision variable with a mean of 0 when the target is absent, and a mean equal to the probability of correct detection for the binary decoder, $P(\hat{c} = c)$, when the target is present. In both conditions the standard deviation is σ. The ideal observer detection threshold is indicated. The probability of correct response equals the probability of correct rejection plus the probability of true positives. (c) 1/d-prime estimated from experiment (see text) follows a linear trend with K. (d) Probability of correct detection of presence or absence of an odorant in a K-component mixture. Blue markers = fraction of correct responses (true positive + correct rejection) by mice. Continuous lines = prediction in the absence of noise for the binary decoder (red) and linear classifier (black). Dashed lines = prediction including noise determined from the data in panel (c) for the binary decoder (red) and the linear classifier (black). For the binary decoder, the noisy decision model fits the data well and predicts a striking cliff in performance for odors with ~ 37 components. The linear classifier predicts a drop much earlier. Parameters of the theoretical model: number of odorants $N_{\mathrm{L}} = 10^4$, number of receptors $N_{\mathrm{R}} = 1000$, receptor response sensitivity $s = 0.05$ [10].

We model the decision making process in this olfactory "cocktail-party problem" as consisting of two steps: (1) Internally representing the components of the mixture using the decoder of odor identity described above, and (2) Using noisy higher level processes to decide on the presence or absence of the target odorant based on the output of the estimation step (Fig. 3a).

Decision noise will degrade performance relative to the ideal decoder. We model this process in the brain in terms of a noisy decision variable derived from the activity of neurons in the decision circuit [43, 44]. If the target is absent, the baseline-subtracted decision variable should take the value 0. If the target is present, the variable should take a value that is proportional to the probability $P(\hat{c} = c) = p$ of correct detection. However, in both cases the decision variable is actually distributed around the desired value with a standard deviation determined by the noise. An ideal observer then simply asks whether the target odorant is more likely to be present or absent, given the observed value of the decision variable and its distribution in the two cases (Fig. 3b).

Decision noise in this picture can be directly estimated from the data in [32] using signal detection theory [45]. Briefly, assuming for simplicity that the decision variable in (Fig. 3b) is distributed as a Gaussian, we can estimate the standard deviation from the rate of hits (fraction of instances when the target is correctly reported to be present) and false alarms (fraction of instances where the target is falsely reported to be present). Signal detection theory [45] relates the the signal to noise ratio (SNR; also called d-prime) of the go/no-go task to the true/false positive rates as: SNR = d-prime = z(true positive) - z(false positives), where z is the z-score. This analysis gives the SNR (d-prime) for mice as a function of K, the number of components in the mixture ([32]; Fig. 3c). We estimate SNR at other values of K by extrapolating the experimentally observed relationship (red line in Fig. 3c). For a Gaussian decision variable with the same standard deviation σ for both conditions, and a difference in means of μ, standard theory [45] gives SNR $= \mu/\sigma$. We took the noise standard deviation in our model to be σ (estimated as above from the data for each K) times a constant chosen to minimize the mean squared difference between theory and experiment.

In the absence of noise ($\sigma = 0$) our binary decoder predicts essentially perfect performance for mice identifying missing odorants in odors with up to ~ 27 components, and a sharp fall-off thereafter (red line in Fig. 3d). Passing this through the noisy decision making process in (Fig. 3a,b) with noise estimated as described above leads to the dashed red line in Fig. 3d. There is an excellent match to the data in [32] and a new prediction: the performance of mice in this olfactory cocktail-party problem will continue to decline linearly as the complexity of odors increases, until there are about 27 component odorants. At that point there will be a sharp fall-off in the probability of correct detection, which will approach chance for odors with about 37 components. These predictions have a weak dependence on N_{L} (see SI: Comparison to experiment), and we do not predict a substantial change in the fall-off as N_{L} changes. However, there is a stronger dependence of the fall-off on the number of receptors (N_{R}). Since humans have ~ 300 receptors, assuming similar noise in the decision making stage, our model predicts that human performance at the "olfactory cocktail-party problem" would be much worse. We predict that this performance will decline linearly till about $5 - 8$ components ($s = 0.05 - 0.10$), and then fall sharply to chance level at about 14 components (details in SI: Comparison to experiment).

We compared this decision model with a linear classifier trained on the receptor responses to report whether a target odorant is present in a mixture or not. Such linear classifiers have been used previously to model the performance of mice in the odorant identification task [46]. We used the CB model to calculate the response of N_{R} (=1000) receptors to random mixtures with K components drawn from a possible N_{L} (=10,000) odorants. The linear classifier was trained on 1000 random mixtures. Half of the mixtures contained a target odorant and the other half did not. After training, the performance of the classifier was estimated over 1000 test mixtures, half of which contained the target odorant. The performance was averaged over 10 random sensitivity matrices, each with a different set of training and test mixtures. Overall, the linear classifier shows a trend similar to the binary decoder (black line and dashed black line in Fig. 3d), but it predicts a drop in performance much earlier than our model.

2.4. Olfactory pathways and limits on odor decoding

The reconstruction of a high-dimensional sparse odor signal from the response of a set of olfactory receptors involves the solution of an underdetermined set of equations. We have demonstrated an efficient algorithm in which receptor responses that are below threshold are used to eliminate the vast majority of odorants, converting an underdetermined problem into an overdetermined one. The olfactory system must implicitly solve a similar problem of solving a set of under-determined equations in order to extract odor information from receptor responses. Thus it should show similar limitations on the ability to decode complex odors. To analyze these limitations, we constructed a neural network with an architecture that resembles the olfactory pathway and studied its computational logic and performance.

In the first stage of the olfactory pathway in mammals, each type of receptor is individually expressed in thousands of Olfactory Sensory Neurons (OSNs) (Olfactory Receptor Neurons in insects). The responses of each type of OSN (OSR in insects) are then aggregated in glomeruli of the Olfactory Bulb (Antennal Lobe in insects) (Fig. 4a). Such replication of response and averaging is the simplest way to mitigate stochastic noise and produce reliable responses, and non-responses, of receptors to odor mixtures.

The second stage of olfactory processing (Olfactory Bulb/Antennal Lobe) involves inhibitory interneurons that implement a recurrent circuit to suppress the weakest responses. These interneurons shut down the output neurons (mitral cells in mammals; projection neurons in insects) of weakly active glomeruli [47, 48], thus silencing the output neurons firing below threshold. As we showed, reliable estimates of these below-threshold responses can be used to turn the underdetermined problem of odor decoding into a well-determined one.

The responses coming out of the first two stages of the olfactory system are projected in a feed-forward manner to the third stage of olfactory system (Piriform Cortex in mammals; Mushroom Body in insects). The neurons in the third stage receive inputs from many glomeruli in the second stage. But a response in these neurons occurs only when most of the inputs to the neuron are active simultaneously [49, 50, 51, 52]. This resembles the elimination step of our algorithm: the non-response of a glomerulus in the second stage is likely to lead to inactivity of the neurons that it projects to in the third stage.

Consider a population of neurons in the third stage whose activity reflects the presence of a particular odorant. If the odorant j is not present in the mixture ($c_j = 0$), the probability that a receptor that binds to this odorant is inactive when responding to the mixture is $P(R_i = 0|c_j = 0) \sim e^{-sK}$ (SI: Eq. S15). So, of the roughly sN_R receptor types that bind to this odorant, nearly $sN_R e^{-sK}$ will be inactive. Taking typical numbers $\{K, N_R, s\} = \{10, 500, 0.05\}$, about 25 receptor types (and their corresponding glomeruli) will respond to a particular odorant, and about 15 these will be silent if the odorant is absent. The other 10 will remain responsive because of other odorants in the mixture, along with other receptor types that are responsive to these odorants.

Consider the units in the third stage (Piriform Cortex in mammals; Mushroom Body in insects) that are not silenced. These units receive excitatory feedforward input from the second stage and inhibitory recurrent inputs from other neurons in the third stage. To model this circuit, suppose that the responses R_i corresponding to the ith receptor type are conveyed to the jth readout unit with feedforward weight $\hat{S}_{ij}$. Also suppose that the jth readout unit provides recurrent input to the kth readout unit with a weight p_{jk}. A standard linearized neural network with these connections will satisfy

$$\frac{d\hat{c}_j}{dt} = -\hat{c}_j + \sum_{i=1}^{N_R} \hat{S}_{ji} R_i + \sum_{k=1, k \neq j}^{N_L} p_{jk} \hat{c}_k \,, \tag{5}$$

where $\hat{c}_j$ is the response of the jth unit. The first term on the right side represents decay of activity in the absence of inputs. In steady state, the activity of the readout units is:

$$\hat{c} = (1-p)^{-1} \hat{S} R. \tag{6}$$

Thus, the steady state response $\hat{c}$ represents a linear transformation of the response of the receptors (R). Suitable choice of feedforward and recurrent connections can ensure that the steady state response

Figure 4. Network implementation resembles piriform cortex architecture: (a) An odorant c_j binds many types of receptors (colors). Responses are reliably estimated by averaging multiple receptors of the same type in a second layer, where axons of each type converge onto glomeruli. In the second layer inhibitory interneurons shut down the outputs of weakly activated glomeruli. The surviving above-threshold average response is relayed to a readout layer. Units in the readout layer also receive recurrent inhibitory inputs from other readout units. Connections for one odorant and final readout unit are shown. (b) Probability of correct decoding ($P(\hat{\boldsymbol{c}} = \boldsymbol{c})$) as a function of odor complexity K and $s * N_R$ for N_R receptors, with $s =$ the probability that an odorant binds to a receptor. $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ was calculated numerically over 100 trials, each with random choice of odor mixture and sensitivity matrix (SI: Numerical simulations). The criterion for correct decoding was that the Euclidean distance between the odor vector $\boldsymbol{c}$ and the decoded vector $\hat{\boldsymbol{c}}$ was less than 0.01.

$\hat{c}$ contains the information required to estimate the signal c. This choice depends on the functional form of the receptor response function which, in general, will be a non-linear function of the odorant concentrations. While these non-linearities are not fully characerized in experiments, the properties and limitations of a decoder with the architecture discussed above can be understood using a linearized receptor response function, i.e., $R_i = \sum_j S_{ij} c_j$, where S_{ij} is the sensitivity matrix and c_j's are odorant concentrations. In this context we get:

$$\hat{c} = \left[(1-p)^{-1} \hat{S} S \right] c = Uc, \tag{7}$$

where the readout population represents a linear transform of the odor $\left(U = \left[(1-p)^{-1} \hat{S} S \right] \right)$. The transform U depends on the feedforward connection matrix $\hat{S}$, which is believed to be a random projection of the glomerular outputs to the neurons in the third stage [53]. The linear nature of the transform ($\hat{c} = Uc$) ensures that $\hat{c}$ is a faithful representaton of the odor concentration vector (c).

To understand the limits of this decoder, let us consider the simple case where the readout unit activity $\hat{c}$ directly represents the odorant concentrations c. This requires $\left[(1-p)^{-1} \hat{S} S \right] = 1$. In a feedforward network ($p = 0$) [23], this can be achieved if the matrix representing feedforward connection strengths ($\hat{S}$) is the inverse of the receptor-odorant sensitivity matrix (S), i.e., $\hat{S} = S^{-1}$. In such a feedforward network, each feedforward connection $\hat{S}_{ji}$ will depend on every element of the sensitivity matrix S. These strengths could only be learnt through a "global" learning rule. Instead, in the presence of recurrent activation ($p \neq 0$), the condition $\left[(1-p)^{-1} \hat{S} S \right] = 1$ can be satisfied if the feedforward and recurrent weights are chosen such that:

$$\sum_{i=1}^{N_R} \hat{S}_{ji} S_{ij} = 1 \ \text{ and } \ p_{jk} + \sum_{i=1}^{N_R} \hat{S}_{ji} S_{ik} = 0 \tag{8}$$

for all j and every $k \neq j$. The first criterion relates the feedforward weights $\hat{S}_{ji}$ to the sensing matrix S_{ij}. The second criterion balances the network – feedforward excitation ($\hat{S}_{ji}$) is compensated by recurrent inhibition (p_{jk}). This recurrent balanced inhibition recalls circuits of the Piriform Cortex in mammals where long-range inhibition arises via large-scale distance-independent random projections from pyramidal cells to locally inhibitory interneurons [51, 52, 54]. In insects similar recurrent inhibition is provided by a giant interneuron.

The N_L constraints from the first criterion in (8) can be solved along with $N_L(N_L - 1)$ constraints from the second criterion because we have about $s * N_R N_L$ parameters in the feedforward matrix $\hat{S}$ and $N_L(N_L - 1)$ parameters in the recurrent matrix (p). This gives more free parameters than constraints if sN_R, the typical number of receptors responding to an odorant, is larger than one. These network parameters can be acquired through local learning rules because the feedforward weights for readout unit j are only related to the sensitivities of receptors to the corresponding odorant j, while the excitatory-inhibitory balance is local (unit by unit). If the response R_i is a nonlinear function of its inputs, e.g. Eq. 3, there will still be enough parameters in a recurrent network to decode odor concentrations. However, units in the network may need nonlinear responses, or be organized in a deep network with multiple feedforward layers.

We characterized the limitations of such a neural network at decoding complex odors (details in SI: Numerical simulations). We selected odor sensitivity matrices S_{ij} such that each odorant binds randomly to a fraction s of the receptors, and assumed a response function that is linear in the odorant concentrations (Eq. 3 with $d = 0$). This linear response represented the statistically stable average over many stochastic receptors. We then selected feedforward projection matrices $\hat{S}_{ij}$ to readout neurons with recurrent weight matrices p_{jk} satisfying the constraints in Eq. 8. Imitating the projections to the olfactory cortex [49, 50, 51, 52], we set a threshold so that the readout units only respond if at least 95% of their feedforward inputs were active. Finally, we allowed the network to decode odor concentrations

as the steady state of the network in Eq. 5. Figure 4b shows that the probability of correct decoding by the network. This is similar to results shown for the abstract decoders discussed in previous sections.

3. Discussion

Our central idea is that receptors which do not respond to an odor convey far more information than receptors that do. This is because the olfactory code is combinatorial – each receptor binds to many different odorants and each odorant binds to many receptors. Hence, an inactive receptor indicates that all the odorants that could have bound to it must be absent. Natural odors are mixtures of perhaps 10-40 components drawn from the more than 10^4 volatile molecules in nature [1, 2, 3]. We showed that if most of these molecules bind to a fraction of the receptors that is neither too small nor too large, odorants that are absent from a mixture can be eliminated from consideration with nearly perfect accuracy by a system with just a few dozen to a few hundred receptors types. The response of the active receptors can then be used to accurately decode the concentrations of the molecules that are present. Our results show that odors of natural complexity can be faithfully encoded in, and decoded from, signals produced by a relatively small number of receptor types that each bind to 5-15% of odorants. Perhaps this explains why all animals express ~ 300 receptor types, give or take a small factor, although receptor diversity does increase with body size [40]. Even at the extremes, the fruitfly and the billion-fold heavier African elephant have $\sim 300/6$ and $\sim 300 \times 6$ receptor types respectively.

We developed a neural network implementation based on the architecture of the early olfactory pathway in the brain. Our neural network model shows properties similar to the abstract algorithms proposed in our work. First, with a few dozen to a few hundred receptor types, our algorithm and the neural network, both show the best performance if each receptor type binds to $\sim 5 - 15\%$ of odorants. This requirement, which recalls ideas from "primacy coding" [55, 56], is consistent with observations from *Drosophila* to human [42, 10]. The next step in our network, similar to the second stage of olfactory processing [47, 48], is to pool signals from multiple receptors of the same type into "glomeruli" and then use lateral inhibition to suppress spurious responses in noisy receptors. This step is crucial to get reliable responses and non-responses of receptors to odors, as required in the elimination step of our proposed decoding algorithms. The third stage of our neural network has readout units that pool from many glomeruli, most of which must be active to produce a response. In addition, the readout units have large-scale, recurrent, balanced inhibition, based on the architecture observed in projections from the second to the third stage of the animal olfactory pathway and in the recurrent circuits of the third stage [49, 50, 51, 52, 54]. Previous work has highlighted that the elements of this architecture could enable robust feedforward reconstruction of compressed odor codes [23], and supports both similarity search [24] and novelty detection [25]. A comparison to our decoding algorithm suggests a new role for the complete circuit, to build a distributed representation of the composition of the odor.

In our network, implementation of the activation function of each unit was linear in the activities of other units. Real neurons have nonlinear activation functions with a threshold, saturation, and sometimes nonlinear summation of inputs. Our model, which can be regarded as a linearization of such neurons around an operating point, can be generalized to nonlinear units, which still have a high threshold for activation to implement feedforward elimination of absent odorants, and recurrent inhibitory balance for concentration decoding.

In the brain, exposure to an odorant activates a sparse, distributed collection of cortical neurons. We modeled this in our network model by representing the readout units as a population of neurons that determine the presence or absence of odorants. It has also been observed that individual cortical neurons may participate in the representation of multiple odorants. This can be incorporated through a simple extension of our network. For a given odorant, we can construct the readout population that samples only from a sub-group of glomeruli that are activated by the odorant. An odorant would then be represented by the collective activation of all neurons in the readout population, some of which may also participate

in the representation of other odorants, as seen in the brain. Such distributed representations might have resulted due to the constraints during development. During development the brain does not know which odorants are present in the world and which receptors they activate. Thus, a good wiring strategy would be to project small groups of glomeruli to target readout neurons. Each such target would be a guess for a subset of receptors that will be co-activated by some odorant. The odorant is then represented by activity in every readout neuron that samples from a proper subset of the activated receptors. Finally, our feedforward weight matrix should be related to the odor sensitivity matrix in order to decode the odorants concentrations. As we discussed, these weights could be acquired through a local learning rule.

According to our decoding model, performance at estimating and discriminating complex odors should increase with the size of the olfactory receptor repertoire. So, while performance might be similar at low odor complexity, say between humans and dogs or mice, the latter animals should be better than humans at discriminating more complex odors as they have 2.5 times more olfactory receptor types. Thus, while comparing performance between species, one should account for odor complexity as well as the total number of receptors. The quantitative predictions of our theory for these differences can be tested by studying odor discrimination thresholds as a function of odor complexity for animals with olfactory receptor repertoires of different sizes.

Our results suggest that the brain may be able to discriminate the detailed composition of odors, contrary to our usual experience of olfaction as a synthetic sense. In fact behavioral experiments do show that it is possible to discriminate complex odors that differ by just a few components [33, 34]. If some version of our decoding algorithm is realized in the brain, as suggested by the analogy with the neural network of the olfactory pathway that we constructed, odors that bind to an inactive receptor type should be largely eliminated. A way of testing this prediction would be to block a specific receptor type pharmacologically, or via optogenetic suppression of the associated glomerulus. Our theory predicts that animals will then tend to behave as if odorants that bind to this receptor are absent, even if other receptors do bind them. Finally, in our model odors can be very well decoded (yellow regions in Figs. 1,2) if they are composed of fewer than $K_{\max}$ components, where $K_{\max}$ is determined by the number of receptor types (N_{R}) and the fraction of them that bind on average to the typical odorant (s). This prediction can be tested by measuring N_{R} and s for different species and then characterizing discrimination performance between odors of complexity bigger and smaller than $K_{\max}$. We illustrated this for mouse in Fig. 3.

Our algorithm can decode complex natural odors detected by chemosensing devices like electric noses [35, 36]. In this engineered setting, the target odorants and response functions are explicitly known so that our method of "Estimation by Elimination" can be precisely implemented.

4. Acknowledgements

VS was supported by a University of Pennsylvania Computational Neuroscience Initiative fellowship. VB was supported by Simons Foundation MMLS grant 400425, and NSF grants PHY-160761 and PHY-1734030. VB thanks the Kavli IPMU for hospitality as this work was completed.

5. Code availability

Code will be made available upon request to the authors.

References

1. M. Dunkel, U. Schmidt, S. Struck, L. Berger, B. Gruening, J. Hossbach, I. S. Jaeger, U. Effmert, B. Piechulla, R. Eriksson, et al., SuperScent- a database of flavors and scents, Nucleic Acids Research 37 (suppl_1) (2008) D291–D294.
2. K. Touhara, L. B. Vosshall, Sensing odorants and pheromones with chemosensory receptors, Annual Review of Physiology 71 (2009) 307–332.

3. C. W. Yu, K. A. Prokop-Prigge, L. A. Warrenburg, J. D. Mainland, Drawing the borders of olfactory space, in: Chemical Senses, Vol. 40, Oxford Univ. Press, 2015, pp. 565–565.

4. L. B. Vosshall, A. M. Wong, R. Axel, An olfactory sensory map in the fly brain, Cell 102 (2) (2000) 147–159.

5. S. Zozulya, F. Echeverri, T. Nguyen, The human olfactory receptor repertoire, Genome Biology 2 (6) (2001) research0018–1.

6. X. Zhang, S. Firestein, The olfactory receptor gene superfamily of the mouse, Nature Neuroscience 5 (2) (2002) 124.

7. P. Quignon, M. Giraud, M. Rimbault, P. Lavigne, S. Tacher, E. Morin, E. Retout, A.-S. Valin, K. Lindblad-Toh, J. Nicolas, et al., The dog and rat olfactory receptor repertoires, Genome Biology 6 (10) (2005) R83.1–9.

8. K. Audouze, A. Tromelin, A. M. Le Bon, C. Belloir, R. K. Petersen, K. Kristiansen, S. Brunak, O. Taboureau, Identification of odorant-receptor interactions by global mapping of the human odorome, PLoS One 9 (4) (2014) e93037.

9. H. Saito, Q. Chi, H. Zhuang, H. Matsunami, J. D. Mainland, Odor coding by a mammalian receptor repertoire, Science Signaling 2 (60) (2009) ra9–ra9.

10. J. D. Mainland, Y. R. Li, T. Zhou, W. L. L. Liu, H. Matsunami, Human olfactory receptor responses to odorants, Scientific Data 2 (2015) 150002.

11. R. C. Araneda, A. D. Kini, S. Firestein, The molecular receptive range of an odorant receptor, Nature Neuroscience 3 (12) (2000) 1248.

12. B. Malnic, J. Hirono, T. Sato, L. B. Buck, Combinatorial receptor codes for odors, Cell 96 (5) (1999) 713–723.

13. C.-Y. Su, K. Menuz, J. R. Carlson, Olfactory perception: receptors, cells, and circuits, Cell 139 (1) (2009) 45–59.

14. J. D. Mainland, J. N. Lundström, J. Reisert, G. Lowe, From molecule to mind: an integrative perspective on odor intensity, Trends in Neurosciences 37 (8) (2014) 443–454.

15. J. Hopfield, Odor space and olfactory processing: collective algorithms and neural implementation, Proceedings of the National Academy of Sciences 96 (22) (1999) 12506–12511.

16. G. Laurent, Olfactory network dynamics and the coding of multidimensional signals, Nature Reviews Neuroscience 3 (11) (2002) 884.

17. M. Stopfer, V. Jayaraman, G. Laurent, Intensity versus identity coding in an olfactory system, Neuron 39 (6) (2003) 991–1004.

18. A. Koulakov, A. Gelperin, D. Rinberg, Olfactory coding with all-or-nothing glomeruli, Journal of Neurophysiology 98 (6) (2007) 3134–3142.

19. C. F. Stevens, What the fly's nose tells the fly's brain, Proceedings of the National Academy of Sciences 112 (30) (2015) 9460–9465.

20. D. Zwicker, A. Murugan, M. P. Brenner, Receptor arrays optimized for natural odor statistics, Proceedings of the National Academy of Sciences 113 (20) (2016) 5570–5575.

21. K. Krishnamurthy, A. M. Hermundstad, T. Mora, A. M. Walczak, V. Balasubramanian, Disorder and the neural representation of complex odors: smelling in the real world, arXiv preprint arXiv:1707.01962.

22. C. Pehlevan, A. Genkin, D. B. Chklovskii, A clustering neural network model of insect olfaction, in: 2017 51st Asilomar Conference on Signals, Systems, and Computers, IEEE, 2017, pp. 593–600.

23. Y. Zhang, T. O. Sharpee, A robust feedforward model of the olfactory system, PLoS Computational Biology 12 (4) (2016) e1004850.

24. S. Dasgupta, C. F. Stevens, S. Navlakha, A neural algorithm for a fundamental computing problem, Science 358 (6364) (2017) 793–796.

25. S. Dasgupta, T. C. Sheehan, C. F. Stevens, S. Navlakha, A neural data structure for novelty detection, Proceedings of the National Academy of Sciences 115 (51) (2018) 13093–13098.

26. S. Qin, Q. Li, C. Tang, Y. Tu, The optimal odor-receptor interaction network is sparse in olfactory systems: Compressed sensing by nonlinear neurons with a finite dynamic range, bioRxiv (2019) 464875.

27. N. Kadakia, T. Emonet, Front-end weber-fechner gain control enhances the fidelity of combinatorial odor coding, eLife 8 (2019) e45293.

28. V. Singh, N. R. Murphy, V. Balasubramanian, J. D. Mainland, Competitive binding predicts nonlinear responses of olfactory receptors to complex mixtures, Proceedings of the National Academy of Sciences (2019) 201813230.

29. J.-P. Rospars, P. Lansky, M. Chaput, P. Duchamp-Viret, Competitive and noncompetitive odorant interactions in the early neural coding of odorant mixtures, Journal of Neuroscience 28 (10) (2008) 2659–2666.

30. G. Cruz, G. Lowe, Neural coding of binary mixtures in a structurally related odorant pair, Scientific Reports 3 (1220) (2013) 1–11.

31. G. Reddy, J. D. Zak, M. Vergassola, V. N. Murthy, Antagonism in olfactory receptor neurons and its implications for the perception of odor mixtures, eLife 7 (2018) e34958.

32. D. Rokni, V. Hemmelder, V. Kapoor, V. N. Murthy, An olfactory cocktail party: figure-ground segregation of odorants in rodents, Nature Neuroscience 17 (9) (2014) 1225.

33. A. Jinks, D. G. Laing, A limit in the processing of components in odour mixtures, Perception 28 (3) (1999) 395–404.

34. C. Bushdid, M. O. Magnasco, L. B. Vosshall, A. Keller, Humans can discriminate more than 1 trillion olfactory stimuli, Science 343 (6177) (2014) 1370–1372.

35. A. C. Johnson, C. Staii, M. Chen, S. Khamis, R. Johnson, M. Klein, A. Gelperin, Dna-decorated carbon nanotubes for chemical sensing, Semiconductor Science and Technology 21 (11) (2006) S17.

36. B. R. Goldsmith, J. J. Mitala Jr, J. Josue, A. Castro, M. B. Lerner, T. H. Bayburt, S. M. Khamis, R. A. Jones, J. G. Brand, S. G. Sligar, et al., Biomimetic chemical sensors using nanoelectronic readout of olfactory receptor proteins, ACS Nano 5 (7) (2011) 5408–5416.

37. P. Gupta, D. F. Albeanu, U. S. Bhalla, Olfactory bulb coding of odors, mixtures and sniffs is a linear sum of odor time profiles, Nature Neuroscience 18 (2) (2015) 272.

38. C. Martelli, J. R. Carlson, T. Emonet, Intensity invariant dynamics and odor-specific latencies in olfactory receptor neuron response, Journal of Neuroscience 33 (15) (2013) 6285–6297.

39. V. Ferreira, Revisiting psychophysical work on the quantitative and qualitative odour properties of simple odour mixtures: a flavour chemistry view. Part 1: intensity and detectability. A review., Flavour and Fragrance Journal 27 (2) (2012) 124–140.

40. T. Tesileanu, S. Cocco, R. Monasson, V. Balasubramanian, Adaptation of olfactory receptor abundances for efficient coding, eLife 8 (2019) e39279. doi:10.7554/eLife.39279.

41. P. Pfister, B. C. Smith, B. J. Evans, J. H. Brann, C. Trimmer, M. Sheikh, R. Arroyave, G. Reddy, H.-Y. Jeong, D. A. Raps, et al., Odorant receptor inhibition is fundamental to odor encoding, Current Biology 30 (2020) 1–14. doi:10.1016/j.cub.2020.04.086.

42. D. Münch, C. G. Galizia, Door 2.0-comprehensive mapping of drosophila melanogaster odorant responses, Scientific Reports 6 (2016) 21841.

43. A. J. Parker, W. T. Newsome, Sense and the single neuron: probing the physiology of perception, Annual Review of Neuroscience 21 (1) (1998) 227–277.

44. J. I. Gold, M. N. Shadlen, The neural basis of decision making, Annual Review of Neuroscience 30 (2007) 535–574.

45. D. M. Green, J. A. Swets, et al., Signal detection theory and psychophysics, Vol. 1, Wiley New York, 1966.

46. A. Mathis, D. Rokni, V. Kapoor, M. Bethge, V. N. Murthy, Reading out olfactory receptors: feedforward circuits detect odors in mixtures without demixing, Neuron 91 (5) (2016) 1110–1123.

47. S. R. Olsen, V. Bhandawat, R. I. Wilson, Divisive normalization in olfactory population codes, Neuron 66 (2) (2010) 287–299.

48. B. Roland, R. Jordan, D. L. Sosulski, A. Diodato, I. Fukunaga, I. Wickersham, K. M. Franks, A. T. Schaefer, A. Fleischmann, Massive normalization of olfactory bulb output in mice with a 'monoclonal nose', eLife 5 (2016) e16335.

49. K. Miyamichi, F. Amat, F. Moussavi, C. Wang, I. Wickersham, N. R. Wall, H. Taniguchi, B. Tasic, Z. J. Huang, Z. He, et al., Cortical representations of olfactory input by trans-synaptic tracing, Nature 472 (7342) (2011) 191.

50. I. G. Davison, M. D. Ehlers, Neural circuit mechanisms for pattern detection and feature combination in olfactory cortex, Neuron 70 (1) (2011) 82–94.

51. D. M. Johnson, K. R. Illig, M. Behan, L. B. Haberly, New features of connectivity in piriform cortex visualized by intracellular injection of pyramidal cells suggest that "primary" olfactory cortex functions like "association" cortex in other sensory systems, Journal of Neuroscience 20 (18) (2000) 6974–6982.

52. K. M. Franks, M. J. Russo, D. L. Sosulski, A. A. Mulligan, S. A. Siegelbaum, R. Axel, Recurrent circuitry dynamically shapes the activation of piriform cortex, Neuron 72 (1) (2011) 49–56.

53. S. J. Caron, V. Ruta, L. Abbott, R. Axel, Random convergence of olfactory inputs in the drosophila mushroom body, Nature 497 (7447) (2013) 113–117.

54. B. Bathellier, T. W. Margrie, M. E. Larkum, Properties of piriform cortex pyramidal cell dendrites: implications for olfactory circuit design, Journal of Neuroscience 29 (40) (2009) 12641–12652.

55. C. D. Wilson, G. O. Serrano, A. A. Koulakov, D. Rinberg, A primacy code for odor identity, Nature Communications 8 (1) (2017) 1477.

56. A. Dewan, A. Cichy, J. Zhang, K. Miguel, P. Feinstein, D. Rinberg, T. Bozza, Single olfactory receptors set odor detection thresholds, Nature Communications 9 (1) (2018) 2887.

Supplementary Information

6. Analytic estimate of the probability of correct decoding

6.1. Identifying odorant presence

We want the probability $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ that the decoded vector $\hat{\boldsymbol{c}}$ equals the input vector $\boldsymbol{c}$, i.e., the corresponding elements of the vectors $\hat{\boldsymbol{c}}$ and $\boldsymbol{c}$ are equal. Assuming statistical independence of the decoding of each odorant, we can write

$$P(\hat{\boldsymbol{c}} = \boldsymbol{c}) = \prod_{i=1}^{N_{\mathrm{L}}} P(\hat{c}_i = c_i) = [P(\hat{c}_i = c_i)]^{N_{\mathrm{L}}} . \tag{S1}$$

The assumption of independence is an approximation that we will validate by comparing with the full numerical results.

The decoded concentration $\hat{c}_i$ could be equal to c_i, if either both of them equal 1 or both of them equal zero. Thus, the term in the square bracket in Eq. S1 can be written as:

$$P(\hat{c}_i = c_i) = P(\hat{c}_i = 1|c_i = 1)P(c_i = 1) + P(\hat{c}_i = 0|c_i = 0)P(c_i = 0). \tag{S2}$$

where $P(c_i = 1) = K/N_{\mathrm{L}} = \alpha$ is the probability that an odorant is present in the mixture, and $P(c_i = 0) = (1 - \alpha)$.

The decoder guarantees that if an odorant c_i is present and there is a receptor R_j that is sensitive to it ($S_{ji}=1$), then the receptor will respond, and the decoded vector will set the corresponding element $\hat{c}_i$ to 1. If no receptor is sensitive to this odorant (i.e, $\forall j \; : \; j \in [1, N_{\mathrm{R}}], S_{ji} = 0$), the decoded element will still be set to 1 by default. So, $P(\hat{c}_i = 1|c_i = 1) = 1$.

To calculate $P(\hat{c}_i = 0|c_i = 0)$, recall that in our decoding scheme, $\hat{c}_i = 0$ if there exists at least one receptor such that $R_j = 0$ for which $S_{ji} = 1$. Thus,

$$P(\hat{c}_i = 0|c_i = 0) = P(\exists \, j : R_j = 0 \cap S_{ji} = 1|c_i = 0) \tag{S3}$$

where $\cap$ is the binary AND operation. The probability on the right is 1 minus the probability that for all receptors either $R_j = 1$ or $R_j = 0 \cap S_{ji} = 0$. So,

$$P(\hat{c}_i = 0|c_i = 0)$$
$$= 1 - P(\forall \, j : R_j = 1 \cup (R_j = 0 \cap S_{ji} = 0)|c_i = 0) = 1 - [P(R_j = 1 \cup (R_j = 0 \cap S_{ji} = 0)|c_i = 0)]^{N_{\mathrm{R}}}, \tag{S4}$$

where in the second step we have again made the assumption that the receptors are independent conditional on the response of c_i. The quantity in the bracket in Eq. S4 can be written as:

$$P(R_j = 1 \cup (R_j = 0 \cap S_{ji} = 0)|c_i = 0)$$
$$= P((R_j = 1 \cup R_j = 0) \cap (R_j = 1 \cup S_{ji} = 0)|c_i = 0)$$
$$= 1 \cap (R_j = 1 \cup S_{ji} = 0)|c_i = 0)$$
$$= P(R_j = 1 \cup S_{ji} = 0|c_i = 0)$$
$$= 1 - P(R_j = 0 \cap S_{ji} = 1|c_i = 0)$$
$$= 1 - P(R_j = 0|c_i = 0)P(S_{ji} = 1|c_i = 0) \tag{S5}$$

Now, $P(S_{ji} = 1|c_i = 0) = P(S_{ji} = 1) = s$, where entries of the sensing matrix are chosen to be non-zero independently and with probability s.

To calculate $P(R_j = 0|c_i = 0)$ recall that the receptors are OR gates with inputs $S_{jk}c_k$. Thus, for $R_j = 0$ all terms $S_{jk}c_k$ should be zero. The probability that any one such term is zero is $(1 - s\alpha)$. Since we already have $c_i = 0$, there are $(N_L - 1)$ additional terms that need to be zero. Hence,

$$P(R_j = 0|c_i = 0) = (1 - s\alpha)^{(N_L - 1)}, \tag{S6}$$

and

$$P(\hat{c}_i = 0|c_i = 0) = \left(1 - \left[1 - s(1 - s\alpha)^{(N_L - 1)}\right]^{N_R}\right) \tag{S7}$$

Putting this all together (using Eq. S7 in Eq. S2), we get:

$$P(\hat{\boldsymbol{c}} = \boldsymbol{c}) = \left[\alpha + (1 - \alpha)\left(1 - \left[1 - s(1 - s\alpha)^{(N_L - 1)}\right]^{N_R}\right)\right]^{N_L} \tag{S8}$$

Using Eq. S7, we can also get the (approximate) probability of a false detection as $P(\hat{c}_i = 1|c_i = 0) = 1 - P(\hat{c}_i = 0|c_i = 0)$:

$$P(\hat{c}_i = 1|c_i = 0) \approx \left[1 - s(1 - s\alpha)^{(N_L - 1)}\right]^{N_R}. \tag{S9}$$

This expression is approximate due to our independence assumptions.

6.1.1. Approximation:

Since the average number of odorants present in the mixture ($K = \alpha N_L$) is small compared to N_L and $N_L \gg 1$, we can approximate:

$$(1 - s\alpha)^{(N_L - 1)} = \left(1 - \frac{sK}{N_L}\right)^{(N_L - 1)} \approx e^{-sK}. \tag{S10}$$

Now, since the odor sensitivity (s) is small, so that se^{-sK} is also small, while $N_R \gg 1$, we further approximate

$$\left[1 - se^{-sK}\right]^{N_R} = \left[1 - \frac{sN_R e^{-sK}}{N_R}\right]^{N_R} \approx e^{-sN_R e^{-sK}}. \tag{S11}$$

This results in:

$$P(\hat{\boldsymbol{c}} = \boldsymbol{c}) = \left[\alpha + (1 - \alpha)\left(1 - e^{-sN_R e^{-sK}}\right)\right]^{N_L} \tag{S12}$$

which simplifies to:

$$P(\hat{\boldsymbol{c}} = \boldsymbol{c}) = \left[1 - e^{-sN_R e^{-sK}} + \alpha e^{-sN_R e^{-sK}}\right]^{N_L} \tag{S13}$$

This expression approximates to Eq. 2 in the main text:

$$P(\hat{\boldsymbol{c}} = \boldsymbol{c}) = \left[1 - N_L e^{-sN_R e^{-sK}}\right] \tag{S14}$$

Similarly, Eq. S6 approximates to

$$P(R_j = 0|c_i = 0) = (1 - s\alpha)^{(N_L - 1)} \approx e^{-s\alpha N_L} = e^{-sK}, \tag{S15}$$

and Eq. S9 approximates to

$$P(\hat{c}_i = 1|c_i = 0) \approx e^{-sN_R e^{-sK}}. \tag{S16}$$

6.2. Identifying odorant concentrations

For the continuous decoder to give a unique solution, the number of receptors that respond to the mixture should be larger than the number of odorants with non-zero concentrations ($K = \alpha N_{\mathrm{L}}$). This ensures that the system of equations is over-determined and can, in principle, be solved.

Additionally, the number of receptors that do not respond should be such that the absent odorants can be set to zero. Since every receptor binds to sN_{L} odorants on average, we need at least $1/s$ receptors to cover all the odorants. In general, as the entries of the sensitivity matrix are statistically distributed, the number of receptors that do not respond should be larger than γ/s for correct odor estimation, where γ is a small number greater than 1.

Putting this all together, if $P(\tilde{N}_{\mathrm{R}})$ is the probability of the number of receptors with non-zero response, we are interested in the probability that $P(\tilde{N}_{\mathrm{R}} > K = \alpha N_{\mathrm{L}}) * P(N_{\mathrm{R}} - \tilde{N}_{\mathrm{R}} > (\gamma/s))$. The probability that a receptor responds is:

$$P(R > 0) = (1 - P(R = 0)) = \left(1 - (1 - s\alpha)^{N_{\mathrm{L}}}\right). \tag{S17}$$

Taking the number of receptors that respond to be a Poisson variable with rate $\langle \tilde{N}_{\mathrm{R}} \rangle = N_{\mathrm{R}} * P(R > 0)$, we can estimate the typical number of receptors that respond. For biologically appropriate parameters $\{N_{\mathrm{L}}, N_{\mathrm{R}}, K, s\} \sim \{10^4, 500, 10, 0.05\}$, the mean number of receptors that respond is $\langle \tilde{N}_{\mathrm{R}} \rangle \sim 200$. The standard deviation is $\sqrt{\langle \tilde{N}_{\mathrm{R}} \rangle} \sim 14$. For these values of the mean and variance, we can approximate the Poisson distribution with a Gaussian $P(\tilde{N}_{\mathrm{R}}) = \mathcal{N}\left(\tilde{N}_{\mathrm{R}}, \sqrt{\tilde{N}_{\mathrm{R}}}\right)$. Thus,

$$P(\hat{\boldsymbol{c}} = \boldsymbol{c}) \sim P(\tilde{N}_{\mathrm{R}} > \alpha N_{\mathrm{L}}) * P(N_{\mathrm{R}} - \tilde{N}_{\mathrm{R}} > (\gamma/s)) = \left[1 - \Phi\left(\frac{\alpha N_{\mathrm{L}} - \tilde{N}_{\mathrm{R}}}{\sqrt{\tilde{N}_{\mathrm{R}}}}\right)\right] \Phi\left(\frac{N_{\mathrm{R}} - \tilde{N}_{\mathrm{R}} - \frac{\gamma}{s}}{\sqrt{\tilde{N}_{\mathrm{R}}}}\right) \tag{S18}$$

where Φ is the cumulative distribution function of the standard normal distribution.

6.3. Numerical Simulations

6.3.1. Identifying odorant presence

For the binary case, the elements of the odor vector were chosen to be non-zero with probability $P(c_i > 0) = K/N_{\mathrm{L}}$. The entries of the sensitivity matrix S_{ij} were chosen to be non-zero with a probability s, $(P(S_{ij}) = s)$. The receptor response was calculated using the binary 'OR' function. The decoded concentration $\hat{c}$ was estimated using the two steps described in the main paper. First, the decoded concentration of any odorant to which an inactive receptor is sensitive, was set to zero. All remaining concentrations were set to 1.

6.3.2. Identifying odorant concentrations

For the continuous case, the elements of the odor vector were chosen to be non-zero with probability $P(c_i > 0) = K/N_{\mathrm{L}}$, and the elements of the sensitivity matrix were chosen to be non-zero with probability s, $(P(S_{ij} > 0) = s)$. The values of the non-zero elements in the odor vector were chosen from a uniform distribution on the interval $[0, 1)$, and for the sensitivity matrix from a log-uniform distribution between 10^{-1} and 10^1. The activity of each receptor was determined using Eq. 3 of the main text (d = 1).

The concentration of any odorant to which an inactive receptor is sensitive was set to zero. After this elimination, let $\tilde{\boldsymbol{R}}$ be the vector representing the response of the set of active receptors, $\tilde{\boldsymbol{c}}$ be the vector representing the concentration of the odorants that have not been set to zero, and $\tilde{S} = \{\tilde{S}_{ij}\}$ be the $\tilde{N}_{\mathrm{R}} \times \tilde{N}_{\mathrm{L}}$ sensitivity submatrix over active receptors $\tilde{\boldsymbol{R}}$ and the remaining odorants $\tilde{\boldsymbol{c}}$. Then, if $\tilde{N}_{\mathrm{R}} < \tilde{N}_{\mathrm{L}}$ (non-invertible case), all decoded concentrations were set to zero. Otherwise, the decoded concentrations were given by the vector that minimized the L_2 distance $||\tilde{\boldsymbol{R}} - \tilde{S} \cdot \tilde{\boldsymbol{c}}||_2$. The Levenberg-Marquardt solver with geodesic acceleration from the *GNU GSL* library was used to find the minimum.

$N_L = 100000$ $N_R = 1000$

Figure S1. Plot of the difference between $P(\hat{c} = c)$ as given by Eq. S8 and as estimated numerically for the binary decoder. For most parameters the analytical results match the simulations.

(a) $N_L = 10000$, $s = 0.05$, $K = 10$ (b) $N_L = 10000$, $s = 0.05$, $K = 10$

Figure S2. $P(\hat{c} = c)$ **as a function of** N_R **for the continuous decoder and alternative choices of the sensitivity matrix**. Results for binary encoding are the same as in Fig. 1c and are plotted here for comparison. (a) Uniform distribution: Similar to Fig. 2c except that the non-zero elements of the sensitivity matrix were chosen uniformly at random between [0,1]. (b) Log-normal distribution: Similar to Fig. 2c except that the non-zero elements of the sensitivity matrix were chosen at random from a log normal distribution with the corresponding normal distribution having mean zero and standard deviation 1.

Figure S3. Other measures of estimation error: (a) Total error: L_2 norm (or the root mean square) of the difference between actual and estimated concentrations for the continuous decoder with compressive binding encoding model. (b) L_2 norm of the difference between actual and estimated concentrations divided by the L_2 norm of the actual concentration.

Multiple trials were run for each choice of parameters. At the end of each trial, the L_2 norm of the difference between actual and decoded concentration vectors was reported. The trial was considered a success if this norm was less than a threshold of 0.01.

Simulations were performed in C++. The sensitivity matrix S and the odorant concentrations were generated from streams of (pseudo)random numbers drawn by the *Xoroshiro128+* random number generator. Each stream is seeded with a 2^{64} forward jump from the seed of the previous stream. The first stream is seeded from the output of the *SplitMix64* generator initialized by current system time. Random number production as well as vector operation code were optimized using *Intel*'s SIMD instruction set.

6.3.3. Neural network

To simulate the neural network, we generated random sparse odor vectors and sensitivity matrices. The elements of the odor vector were chosen to be non-zero with probability $P(c_i > 0) = K/N_{\mathrm{L}}$, and the elements of the sensitivity matrix were chosen to be non-zero with probability $P(S_{ij}) > 0 = s$. The value of the non-zero elements were chosen from a uniform distribution on the interval $[0, 1)$. The receptor response was calculated using a linear response model (d = 0 in Eq. 3). To get the feed forward connections $\hat{S}_{ji}$, we first made a matrix $\bar{S}$ defined as: $\bar{S}_{ji} = 1/(S_{ji})$ if S_{ji} is non-zero, and $\bar{S}_{ji} = 0$ otherwise. The matrix $\hat{S}$ was then chosen as: $\hat{S}_{ji} = \bar{S}_{ji}/\left(\sum_i \bar{S}_{ji} S_{ij} \right)$. The elements of the recurrent connectivity matrix were obtained as $p_{jk} = -\sum_i \hat{S}_{ji} S_{ik}$.

If more than 5% of the receptors connected to a readout unit c_j were inactive, the decoded concentration $\hat{c}_j$ was set to zero. The feedforward input to the remaining readout units were calculated as $c_j^{\mathrm{init}} = \hat{S}_{ji} R_i$. The remaining concentrations were computed as $(1 - \tilde{p})^{-1} \mathbf{c}^{\mathrm{init}}$, where $\mathbf{c}^{\mathrm{init}}$ is the vector representing the total feedforward input to neurons that have more than 95% of their receptors active, and $\tilde{p}$ represents the sub-matrix of connection weights between these neurons.

6.4. Comparison to experiment

In the main text the predictions of the binary decoder are compared to the performance of mice [32] in go/no-go experiments where the animal is presented with an odor (a mixture of odorants) and is asked to report the presence or absence of a target odorant. The experiments of [32] were conducted

Figure S4. Dependence of $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ on N_{L}: $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ plotted as a function of odor complexity K and sN_{R} at a fixed value of number of receptors $N_{\mathrm{R}} = 500$. Each panel gives $P(\hat{\boldsymbol{c}} = \boldsymbol{c})$ for a different value of the total number of possible odorants N_{L}. The minimum value of sN_R for successful decoding and the optimal value where the most complex odors can be decoded are both relatively independent of N_{L}.

Figure S5. Prediction for human and mouse performance at olfactory cocktail party problem: Probability of correct detection of presence or absence of an odorant in a K-component mixture. Blue markers = fraction of correct responses (true positive + correct rejection) by mice. Continuous lines = prediction in the absence of noise for the binary decoder. Dashed lines = prediction including noise determined from the data in [32] (Fig. 3c main text). Human performance (number of receptors $N_R \sim 300$) at the task is predicted to be lower than of mouse ($N_R \sim 1000$), falling to chance level at $K \sim 14$. We have shown the prediction for humans at two values of the receptor-odor sensitivity ($s = 0.05$ (green lines) and $s = 0.10$ (blue lines)).

as follows. Thirteen mice were trained to report the presence of a target odorant by licking at a water spout, and absence by abstaining from licking. Feedback was provided by giving a water drop on correct lick, and punishing an incorrect lick by a 5 second timeout. Each mouse was trained to identify 2 target odorants from a set of 16. A total of 8 different sets of target odorant pairs were used. The odor mixtures in the experiment contained between 1-14 odorants. In each trial, the target odorant was present with probability 0.5, and the two target odorants for a particular mouse were never presented in the same trial. The mice were first trained to identify targets in mixtures with few components, and allowed to reach a performance level of 80%. Once mice reached this criterion, the complexity of the mixtures (number of component odorants) was gradually increased such that the distribution of mixture complexity presented in any trial gradually became uniform. Mice took typically 1000 trials to reach the 80% criterion with uniform distribution of odor complexity. Performance was measured after such learning had occurred.

Figure S6. Effect of the number of odorants N_{L} and receptor sensitivity s on performance in olfactory cocktail party problem: Probability of correct detection of presence or absence of an odorant in a K-component mixture. Compare to Figure 3d. (a) $N_{\mathrm{R}} = 10^3, s = 0.05$. (b) $N_{\mathrm{R}} = 10^3, N_{\mathrm{L}} = 10^4$.