What the odor is not: Estimation by elimination

Vijay Singh^{a,b}, Martin Tchernookov^c, and Vijay Balasubramanian^{a,b,1}

^aComputational Neuroscience Initiative, University of Pennsylvania, Philadelphia, PA, 19104, USA; ^bDepartment of Physics, University of Pennsylvania, Philadelphia, PA, 19104, USA; ^cDepartment of Physics, University of Wisconsin, Whitewater, WI, 53190, USA

The olfactory system uses responses of a small number of broadly sensitive receptors to combinatorially encode a vast number of odors. Here, we propose a method for decoding such a distributed representation. Our main idea is that a receptor that does not respond to an odor carries more information than a receptor that does, because a typical receptor binds to many odorants. So a response below threshold signals absence of all such odorants. As a result, it is easier to identify what the odor *is not*, rather than what the odor *is*. We demonstrate that, for biologically realistic numbers of receptors, response functions, and odor mixture complexity, this remarkably simple method of elimination turns an underdetermined decoding problem into an overdetermined one, allowing accurate determination of odorants in a mixture and their concentrations. We give a simple neural network realization of our algorithm resembling the circuit architecture in piriform cortex, and propose experimental tests.

Olfaction | Receptors | Odor Decoding

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 (1–3), far more than the number of receptor types available to animals (~ 50 in fly, ~ 300 in human, ~ 1000 in rat, mouse and dog (4–7)). 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 (8–14), leading to a distributed, compressed, and combinatorial representation (12, 15–27).

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 (28–31), and propose a scheme to decode 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 detecting complex natural odors.

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) while making additional predictions. We then propose a neural network to implement our decoding scheme. Remarkably, the statistics of olfactory receptor binding to natural odors (10, 33), and circuits in the olfactory epithelium, the olfactory bulb (34) and the olfactory cortex (35–37), all reflect properties of this network. 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 (38, 39) for odor detection, and in medical tests.

Results

Decoder of odor identity. First, suppose we seek to simply identify the 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 binary model of receptor activity.

In detail, consider a mixture of N_L odorants represented by the binary vector $\mathbf{c} = (c_1, c_2, \dots, c_{N_L})$, where $c_i = 1$ represents the presence of the i 'th odorant. Let odors have complexity K : on average only K odorants are present. These odorants bind to N_R receptors whose response is given by the binary vector $\mathbf{R} = (R_1, R_2, \dots, R_{N_R})$. 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 sN_R receptors and each receptor to sN_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.

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

1. First, all inactive receptors are identified and all odorants that bind to them 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.

¹To whom correspondence should be addressed. E-mail: vijay@physics.upenn.edu

Fig. 1. Odor identity decoder: 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. Thus, the odorant concentration and receptor response are represented as binary vectors. Odorants present in the mixture and active receptors are depicted as filled blue squares. Elements of the sensing matrix represent receptor-odorant interactions: a filled blue matrix element indicates an odorant that binds to a receptor. (a) Encoding: 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 ($P(\hat{c} = c)$) plotted as a function of the number of receptors (N_R). Markers are numerical simulations, smooth curve is the approximate analytical result (Eq. S8). $P(\hat{c} = c)$ 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) $P(\hat{c} = c)$ 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 ($s * N_R$), with the number of receptors (N_R) and the number of odorants (N_L) held fixed. We are plotting $P(\hat{c} = c)$ as a function of the product $s * N_R$, since the parameters s and N_R always appear in this combination (Eq. 2 and Eq. S8). $s * N_R$, the average number of receptors a typical odorant binds to, determines whether the odorant can be detected or not. $P(\hat{c} = c)$ 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 $P(\hat{c} = c)$ estimated by setting the analytical expression Eq. 2 to half.

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, the decoded concentration $\hat{c}_j$ will not be set to zero.

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, $\hat{c}_j$ will be set to 1. We can derive an approximate expression for the probability of false positives (see SI: Odor identity decoder):

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

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: Odor identity decoder Eq. S8):

$$P(\hat{c} = c) \approx \left(1 - N_L e^{-sN_R} e^{-sK}\right). \quad [2]$$

The second term is approximately the probability of false positives if there are N_L 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 $sN_R > 1$ 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 compare well with the observed sensitivity of olfactory receptor responses to odorants ($s \sim 14\%$ for *Drosophila* (33), $s \sim 4\%$ for humans (10)). Thus, for typical values ($\{N_L, K, N_R, s\} = \{10^4, 10, 500, 0.05\}$), 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 ($P(\hat{c} = c) \sim 0.998$; Eq. 2). These estimates suggest that our proposed decoding scheme will be effective at decoding complex odors in the biologically relevant regime.

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 ($P(\hat{c} = c)$) 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 ($s * N_R$), our decoding scheme shows excellent performance ($P(\hat{c} = c) \approx 1$) (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 of the transition point between poor and excellent decoding (solid lines in Fig. 1c,d; and SI Fig. S1).

Decoder of odor composition. If noise is low, or integration times are long, the activity of a receptor is more appropriately represented by 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

Fig. 2. Odor composition decoder: (a) Encoding: 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): Silent 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).

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):

$$R = \frac{\sum_{j=1}^{N_L} x_j}{\left(1 + d * \sum_{j=1}^{N_L} x_j\right)}. \quad [3]$$

Here $x_j = S_{ij}c_j$ and d parameterizes the affinity of odorants for the receptor. The CB model approximates to the binary model when ($d \rightarrow \infty$) and to the commonly used linear response model when $d \rightarrow 0$ (40–43).

As discussed earlier, the binary 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 $\tilde{N}_L \leq \tilde{N}_R$ 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 $\tilde{N}_R \sim sKN_R$ 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 $\{N_L, K, N_R, s\} = \{10^4, 10, 500, 0.05\}$ give $\tilde{N}_L \sim K = 10$ which is less than $\tilde{N}_R \sim sKN_R = 250$, showing that in the biologically relevant regime our elimination algorithm leads to an over-determined and hence solvable identification problem.

We can derive an approximate analytical expression for the probability of correct estimation (SI: Odor composition

decoder). This derivation assumes that the typical number of receptors that respond to a mixture is larger than the average odor complexity ($\tilde{N}_R > K$), while, at the same time, enough receptors are inactive to eliminate absent odorants. This requires $s(N_R - \tilde{N}_R) > \gamma$, where $\gamma > 1$ is a parameter that depends on the receptor response model (details in SI: Odor composition decoder). With these assumptions,

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

Φ is the standard normal cumulative distribution function.

To numerically estimate the probability of correct decoding ($P(\hat{c} = c)$), 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). The probability of correct decoding is zero when there are very few receptors (N_R). 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 perfectly for a wide range of parameters (Fig. 2d,e), so long as receptors are sufficiently sensitive $s * N_R > 6$. Odors with the highest complexity are decoded when $s * N_R \sim 10 - 15$. The dependence on the total number of odorants (N_L) is weak (SI Fig. S4). 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).

Since humans have about 300 receptors, our model predicts that odors with most components can be decoded with $s \sim 3 - 5\%$ so that $sN_R = 10 - 15$. This is consistent with observation – human receptors have $s \sim 4\%$ (10). For *Drosophila*, which has ~ 50 receptors, the observed sensitivity of $s \sim 14\%$ (33) gives $sN_R \sim 7$, in the expected range for successful decoding.

Noisy decision making and comparison with experiment. We can compare the predictions of the odor identity decoder to the performance of mice (32) in behavioral studies where the animal is presented with an odor (a mixture of odorants) and is asked to report whether a target odorant is present or not (details in SI: Comparison to experiment). 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 odor identity decoder 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 will model this process in the brain in terms of a noisy decision variable derived from the activity of neurons in the decision circuit (44, 45). 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

Fig. 3. Experimental Test: We predict performance of mice in the olfactory cocktail-party task of detecting missing odorants in K -component mixtures (32). (a) Schematic of the two step noisy decoder. (b) Normally distributed decision variable: mean when target is absent = 0; mean when target is present = $P(\hat{c} = c)$, the probability of correct detection of the binary decoder; standard deviations σ in both conditions. The ideal observer detection threshold is indicated. Probability of correct response = probability of correct rejection + 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. Red line = theoretical prediction in the absence of noise. Black line = theoretical prediction including noise determined from the data in panel (c). The noisy decision model fits the data well and predicts a striking cliff in performance for odors with ~ 37 components. Parameters of the theoretical model: number of odorants $N_L = 10^4$, number of receptors $N_R = 1000$, receptor response sensitivity $s = 0.05$ (10).

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 (46). 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 (46) 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: $SNR = d\text{-prime} = z(\text{true positive}) - z(\text{false positives})$, 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 (46) 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 a 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 black 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 with continue 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.

Network implementation. We have demonstrated an efficient algorithm for decoding odor identity from a combinatorial code in which receptors that are below threshold are used to eliminate the vast majority of odorants, converting an underdetermined problem into an overdetermined one. Here, we develop a neural network implementation of the algorithm.

To instantiate our algorithm mechanistically it is important to have reliable responses and non-responses in receptors, reflecting the actual concentrations of odorants. However, receptor-odorant binding is inherently stochastic. So the first step is to mitigate sensing noise. The simplest way to achieve this is to have many copies of each receptor and to average their responses. Indeed, 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) and, subsequently, responses of each type are aggregated in glomeruli of the Olfactory Bulb (Antennal Lobe in insects) (Fig. 4a).

Below-threshold responses are especially important for our algorithm. To further ensure their reliability we can arrange for receptor types to compete to suppress each other, thereby muting the weakest responses. In the presence of a firing threshold such a suppression will cause units firing at very low rates to fall silent, as we require. Well-known computational principles show that recurrent inhibitory circuits can achieve this effect. Indeed, in the second stage of olfactory processing, inhibitory interneurons implement a circuit that shuts down the output neurons (mitral cells in mammals; projection neurons in insects) of weakly active glomeruli (47, 48).

Next, we need a mechanism to eliminate absent odorants. To achieve this, we organize projections from glomeruli of receptors binding a given odorant to a readout unit whose activity $\hat{c}_j$ represents the odorant concentration (Figure 4a). We can then implement the elimination step of our decoding algorithm by setting the readout unit threshold so that most of its inputs must be active to trigger a response. If odorant j is not present in the mixture ($c_j = 0$), the probability that a receptor which 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 receptors 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 receptors will respond to a ligand, and about 15 these will be silent if the ligand is absent. Hence, the corresponding readout unit will be silent ($\hat{c}_j = 0$). A similar architecture is seen in the feedforward projections from the Olfactory Bulb to the Piriform Cortex in mammals (Antennal Lobe to Mushroom Body in insects). Specifically, each neuron in the third stage of olfactory processing receives inputs from many glomeruli in the second stage, but simultaneous activation from most of these is necessary for a response (36, 37, 49, 50).

Finally we need a mechanism to set the activity of the readout units that have not been eliminated to represent concentrations of odorants. This can be achieved through a network

Fig. 4. Network implementation: (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. This 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 ($P(\hat{c} = c)$) 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. $P(\hat{c} = c)$ 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 c and the decoded vector $\hat{c}$ was less than 0.01.

of recurrent connections between readout units (Figure 4). To illustrate, suppose that the responses R_i corresponding to the i th receptor are conveyed to the j th readout unit with a feedforward weight $\hat{S}_{ij}$. Also suppose that the j th readout unit provides recurrent input to the k th readout unit with a weight p_{jk} . A standard linearized neural network with these connections satisfies the equation

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

where $\hat{c}_j$ is the response of the j th unit. The first term on the right side represents decay of activity in the absence of inputs. In this context, we also linearize the responses so that $R_i = \sum_j S_{ij} c_j$, where S_{ij} is a sensitivity matrix and c_j are odorant concentrations. The steady state occurs when

$$0 = \left(-\hat{c}_j + c_j \sum_{i=1}^{N_R} \hat{S}_{ji} S_{ij} \right) + \sum_{k=1, k \neq j}^{N_L} \left(p_{jk} \hat{c}_k + c_k \sum_{i=1}^{N_R} \hat{S}_{ji} S_{ik} \right). \quad [6]$$

In the steady state $\hat{c}_j = c_j$, i.e. the activity of readout unit j equals the concentration of odorant j , provided the feedforward and recurrent weights are adjusted to obey

$$\sum_{i=1}^{N_R} \hat{S}_{ji} S_{ij} = 1 \quad \text{and} \quad p_{jk} + \sum_{i=1}^{N_R} \hat{S}_{ji} S_{ik} = 0 \quad [7]$$

for all j and every $k \neq j$. The first criterion relates the feedforward weights $\hat{S}_{ji}$ to the sensing matrix S_{ij} (see (23) for a similar relation in a related context). 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 (35–37). In insects similar recurrent inhibition is provided by a giant interneuron.

The N_L constraints from the first criterion in (7) 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 bigger 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 to have nonlinear responses, or be organized in a deep network with multiple feedforward layers.

To test our network decoder (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. 7. Imitating the projections to the olfactory cortex (36, 37, 49, 50), we set a threshold so that the readout units only responded 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 is similar to results shown for the abstract decoders discussed in previous sections.

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 (1-3). 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 receptor 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 fully 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 (43). 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 proposed a network implementation of our algorithm that resembles the architecture of the early olfactory pathway in the brain. First, with a few dozen to a few hundred receptor types, our algorithm requires each receptor type to bind to $\sim 5 - 15\%$ of odorants. This requirement, which recalls ideas from “primacy coding” (51, 52), is consistent with observations

from *Drosophila* to human (10, 33). If receptors are noisy, the next step in our decoding network is to pool signals from multiple receptors of the same type into “glomeruli”, and to then allow lateral inhibition to suppress spurious responses due to noise. This pooling and inhibition motif is realized in the second stage of olfactory processing (47, 48). The third stage of our decoding network has readout units that pool from many glomeruli, most of which must be active to produce a response. In addition, the readout units must have large-scale, recurrent, balanced inhibition. A similar architecture is visible in projections from the second to the third stage of the animal olfactory pathway, and in the recurrent circuits of the third stage (35-37, 49, 50). Previous work has highlighted that this architecture could enable robust feedforward reconstruction of compressed odor codes (23), and supports both similarity search (24) and novelty detection (25).

In our network implementation 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.

Our network readout units individually represented the presence or absence of odorants. By contrast, in the brain, exposure to an odorant activates a sparse, distributed collection of cortical neurons. A simple extension of our network produces such a representation. Instead of collecting all glomeruli that respond to a given odorant, we can construct readout units that sample groups of these glomeruli. An odorant would then be represented by the collective activation of a set of readout neurons, some of which may also participate in the representation of other odorants, as seen in the brain. We did not pursue this approach because we assumed knowledge of the olfactory environment and the receptor sensitivity matrix. But 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 was related to the odor sensitivity matrix in order to decode the actual concentrations of odorants. As we discussed these weights could be acquired through a local learning rule. Our theory could be tested by sampling sensitivities of receptors for a particular odorant (9, 10), along with feedforward projection strengths from glomeruli to their targets, perhaps measured by optogenetically activating individual glomeruli while imaging the strength of downstream responses.

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, these 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. This prediction 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 indeed 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 (53, 54). If our decoding algorithm is realized in the brain, all odors that bind to an inactive receptor type should be 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 ligands 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 (38, 39). 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.

ACKNOWLEDGMENTS. 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.

- Dunkel M, et al. (2008) SuperScent- a database of flavors and scents. *Nucleic Acids Research* 37(suppl_1):D291–D294.
- Touhara K, Vosshall LB (2009) Sensing odorants and pheromones with chemosensory receptors. *Annual Review of Physiology* 71:307–332.
- Yu CW, Prokop-Prigge KA, Warrenburg LA, Mainland JD (2015) Drawing the borders of olfactory space in *Chemical Senses*. (Oxford Univ. Press), Vol. 40, pp. 565–565.
- Vosshall LB, Wong AM, Axel R (2000) An olfactory sensory map in the fly brain. *Cell* 102(2):147–159.
- Zozulya S, Echeverri F, Nguyen T (2001) The human olfactory receptor repertoire. *Genome Biology* 2(6):research0018–1.
- Zhang X, Firestein S (2002) The olfactory receptor gene superfamily of the mouse. *Nature Neuroscience* 5(2):124.
- Quignon P, et al. (2005) The dog and rat olfactory receptor repertoires. *Genome Biology* 6(10):R83.1–9.
- Audouze K, et al. (2014) Identification of odorant-receptor interactions by global mapping of the human odorome. *Plos One* 9(4):e93037.
- Saito H, Chi Q, Zhuang H, Matsunami H, Mainland JD (2009) Odor coding by a mammalian receptor repertoire. *Science Signaling* 2(60):ra9–ra9.
- Mainland JD, Li YR, Zhou T, Liu WLL, Matsunami H (2015) Human olfactory receptor responses to odorants. *Scientific Data* 2:150002.
- Araneda RC, Kini AD, Firestein S (2000) The molecular receptive range of an odorant receptor. *Nature Neuroscience* 3(12):1248.
- Mainic B, Hirono J, Sato T, Buck LB (1999) Combinatorial receptor codes for odors. *Cell* 96(5):713–723.
- Su CY, Menzies K, Carlson JR (2009) Olfactory perception: receptors, cells, and circuits. *Cell* 139(1):45–59.
- Mainland JD, Lundström JN, Reiser J, Lowe G (2014) From molecule to mind: an integrative perspective on odor intensity. *Trends in Neurosciences* 37(8):443–454.
- Hopfield J (1999) Odor space and olfactory processing: collective algorithms and neural implementation. *Proceedings of the National Academy of Sciences* 96(22):12506–12511.
- Laurent G (2002) Olfactory network dynamics and the coding of multidimensional signals. *Nature reviews neuroscience* 3(11):884.
- Stopfer M, Jayaraman V, Laurent G (2003) Intensity versus identity coding in an olfactory system. *Neuron* 39(6):991–1004.
- Koulakov A, Gelperin A, Rinberg D (2007) Olfactory coding with all-or-nothing glomeruli. *Journal of Neurophysiology* 98(6):3134–3142.
- Stevens CF (2015) What the fly's nose tells the fly's brain. *Proceedings of the National Academy of Sciences* 112(30):9460–9465.
- Zwicker D, Murugan A, Brenner MP (2016) Receptor arrays optimized for natural odor statistics. *Proceedings of the National Academy of Sciences* 113(20):5570–5575.
- Krishnamurthy K, Hermundstad AM, Mora T, Walczak AM, Balasubramanian V (2017) Disorder and the neural representation of complex odors: smelling in the real world. *arXiv preprint arXiv:1707.01962*.
- Pehlevan C, Genkin A, Chklovskii DB (2017) A clustering neural network model of insect olfaction in 2017 51st Asilomar Conference on Signals, Systems, and Computers. (IEEE), pp. 593–600.
- Zhang Y, Sharpee TO (2016) A robust feedforward model of the olfactory system. *PLoS Computational Biology* 12(4):e1004850.
- Dasgupta S, Stevens CF, Navlakha S (2017) A neural algorithm for a fundamental computing problem. *Science* 358(6364):793–796.
- Dasgupta S, Sheehan TC, Stevens CF, Navlakha S (2018) A neural data structure for novelty detection. *Proceedings of the National Academy of Sciences* 115(51):13093–13098.
- Kadukia N, Emonet T (2019) Front-end weber-fechner gain control enhances the fidelity of combinatorial odor coding. *eLife* 8:e45293.
- Qin S, Li Q, Tang C, Tu Y (2018) Statistical properties of the optimal sensitivity matrix for compressed sensing with nonlinear sensors. *bioRxiv* p. 464875.
- Singh V, Murphy NR, Balasubramanian V, Mainland JD (2019) Competitive binding predicts nonlinear responses of olfactory receptors to complex mixtures. *Proceedings of the National Academy of Sciences* p. 201813230.
- Rospars JP, Lansky P, Chaput M, Duchamp-Viret P (2008) Competitive and noncompetitive odorant interactions in the early neural coding of odorant mixtures. *Journal of Neuroscience* 28(10):2659–2666.
- Cruz G, Lowe G (2013) Neural coding of binary mixtures in a structurally related odorant pair. *Scientific Reports* 3(1220):1–11.
- Reddy G, Zak JD, Vergassola M, Murthy VN (2018) Antagonism in olfactory receptor neurons and its implications for the perception of odor mixtures. *Elife* 7:e34958.
- Rokni D, Hemmelder V, Kapoor V, Murthy VN (2014) An olfactory cocktail party: figure-ground segregation of odorants in rodents. *Nature neuroscience* 17(9):1225.
- Münch D, Galizia CG (2016) Door 2.0-comprehensive mapping of drosophila melanogaster odorant responses. *Scientific Reports* 6:21841.
- Nagayama S, Homma R, Imamura F (2014) Neuronal organization of olfactory bulb circuits. *Frontiers in Neural Circuits* 8:98.
- Bathellier B, Margrie TW, Larkum ME (2009) Properties of piriform cortex pyramidal cell dendrites: implications for olfactory circuit design. *Journal of Neuroscience* 29(40):12641–12652.
- Johnson DM, Illig KR, Behan M, Haberly LB (2000) 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):6974–6982.
- Franks KM, et al. (2011) Recurrent circuitry dynamically shapes the activation of piriform cortex. *Neuron* 72(1):49–56.
- Johnson AC, et al. (2006) Dna-decorated carbon nanotubes for chemical sensing. *Semiconductor science and technology* 21(11):S17.
- Goldsmith BR, et al. (2011) Biomimetic chemical sensors using nanoelectronic readout of olfactory receptor proteins. *ACS nano* 5(7):5408–5416.
- Gupta P, Albeanu DF, Bhalla US (2015) Olfactory bulb coding of odors, mixtures and sniffs is a linear sum of odor time profiles. *Nature neuroscience* 18(2):272.
- Martelli C, Carlson JR, Emonet T (2013) Intensity invariant dynamics and odor-specific latencies in olfactory receptor neuron response. *Journal of Neuroscience* 33(15):6285–6297.
- Ferreira V (2012) 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):124–140.
- Tesileanu T, Cocco S, Monasson R, Balasubramanian V (2019) Adaptation of olfactory receptor abundances for efficient coding. *Elife* 8:e39279.
- Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. *Annual review of neuroscience* 21(1):227–277.
- Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu. Rev. Neurosci.* 30:535–574.
- Green DM, Swets JA, et al. (1966) *Signal detection theory and psychophysics*. (Wiley New York) Vol. 1.
- Olsen SR, Bhandawat V, Wilson RI (2010) Divisive normalization in olfactory population codes. *Neuron* 66(2):287–299.
- Roland B, et al. (2016) Massive normalization of olfactory bulb output in mice with a “monoclonal nose”. *Elife* 5:e16335.
- Miyamichi K, et al. (2011) Cortical representations of olfactory input by trans-synaptic tracing. *Nature* 472(7342):191.
- Davison IG, Ehlers MD (2011) Neural circuit mechanisms for pattern detection and feature combination in olfactory cortex. *Neuron* 70(1):82–94.
- Wilson CD, Serrano GO, Koulakov AA, Rinberg D (2017) A primacy code for odor identity. *Nature communications* 8(1):1477.
- Dewan A, et al. (2018) Single olfactory receptors set odor detection thresholds. *Nature communications* 9(1):2887.
- Jinks A, Laing DG (1999) A limit in the processing of components in odour mixtures. *Perception* 28(3):395–404.
- Bushdid C, Magnasco MO, Vosshall LB, Keller A (2014) Humans can discriminate more than 1 trillion olfactory stimuli. *Science* 343(6177):1370–1372.

Supplementary information

Fig. 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. Analytic estimate of the probability of correct decoding.

A.1. Odor identity decoder. We want the probability $P(\hat{c} = c)$ that the decoded vector $\hat{c}$ equals the input vector c , i.e., the corresponding elements of the vectors $\hat{c}$ and c are equal. Assuming statistical independence of the decoding of each odorant, we can write

$$P(\hat{c} = c) = \prod_{i=1}^{N_L} P(\hat{c}_i = c_i) = [P(\hat{c}_i = c_i)]^{N_L}. \quad [S1]$$

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:

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

where $P(c_i = 1) = K/N_L = \alpha$ 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., $\forall j : j \in [1, N_R], S_{ji} = 0$), 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,

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

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,

$$\begin{aligned} P(\hat{c}_i = 0 | c_i = 0) &= 1 - P(\forall j : R_j = 1 \cup (R_j = 0 \cap S_{ji} = 0) | c_i = 0) \\ &= 1 - [P(R_j = 1 \cup (R_j = 0 \cap S_{ji} = 0) | c_i = 0)]^{N_R}, \end{aligned} \quad [S4]$$

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:

$$\begin{aligned} P(R_j = 1 \cup (R_j = 0 \cap S_{ji} = 0) | c_i = 0) &= P((R_j = 1 \cup R_j = 0) \cap (R_j = 1 \cup S_{ji} = 0) | c_i = 0) \\ &= 1 \cap (R_j = 1 \cup S_{ji} = 0) | c_i = 0 \\ &= P(R_j = 1 \cup S_{ji} = 0 | c_i = 0) \\ &= 1 - P(R_j = 0 \cap S_{ji} = 1 | c_i = 0) \\ &= 1 - P(R_j = 0 | c_i = 0)P(S_{ji} = 1 | c_i = 0) \end{aligned} \quad [S5]$$

Fig. 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.

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,

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

and

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

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

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

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

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

This expression is approximate due to our independence assumptions.

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:

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

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

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

This results in:

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

which simplifies to:

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

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

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

Similarly, Eq. S6 approximates to

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

Fig. 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.

and Eq. S9 approximates to

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

A.2. Odor composition decoder. 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 ($K = \alpha N_L$). 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 $P(\tilde{N}_R)$ is the probability of the number of receptors with non-zero response, we are interested in the probability that $P(\tilde{N}_R > K = \alpha N_L) * P(N_R - \tilde{N}_R > (\gamma/s))$. The probability that a receptor responds is:

$$P(R > 0) = (1 - P(R = 0)) = (1 - (1 - s\alpha)^{N_L}). \quad [S17]$$

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

$$\begin{aligned} P(\hat{c} = c) &\sim P(\tilde{N}_R > \alpha N_L) * P(N_R - \tilde{N}_R > (\gamma/s)) \\ &= \left[1 - \Phi\left(\frac{\alpha N_L - \tilde{N}_R}{\sqrt{\tilde{N}_R}}\right) \right] \Phi\left(\frac{N_R - \tilde{N}_R - \frac{\gamma}{s}}{\sqrt{\tilde{N}_R}}\right) \end{aligned} \quad [S18]$$

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

B. Numerical Simulations.

B.1. Odor identity decoder. For the binary case, the elements of the odor vector were chosen to be non-zero with probability $P(c_i > 0) = K/N_L$. The entries of the sensitivity matrix S_{ij} were chosen to be non-zero with a probability s , ($P(S_{ij} > 0 = s)$). 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.

Fig. S4. Dependence of $P(\hat{c} = c)$ on N_L : $P(\hat{c} = c)$ plotted as a function of odor complexity K and sN_R at a fixed value of $N_R = 500$. Each panel gives $P(\hat{c} = c)$ for 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 .

B.2. Odor composition decoder. For the continuous case, the elements of the odor vector were chosen to be non-zero with probability $P(c_i > 0) = K/N_L$, 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 $\tilde{\mathbf{R}}$ be the vector representing the response of the set of active receptors, $\tilde{\mathbf{c}}$ be the vector representing the concentration of the odorants that have not been set to zero, and $\tilde{\mathbf{S}} = \{\tilde{S}_{ij}\}$ be the $\tilde{N}_R \times \tilde{N}_L$ sensitivity submatrix over active receptors $\tilde{\mathbf{R}}$ and the remaining odorants $\tilde{\mathbf{c}}$. Then, if $\tilde{N}_R < \tilde{N}_L$ (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 $\|\tilde{\mathbf{R}} - \tilde{\mathbf{S}} \cdot \tilde{\mathbf{c}}\|_2$. The Levenberg-Marquardt solver with geodesic acceleration from the GNU GSL library was used to find the minimum.

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 $\mathbf{S}$ 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.

B.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 $P(c_i > 0) = K/N_L$, 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 $\bar{S}$ 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 $c_j^{\text{init}} = \hat{S}_{ji} R_i$. The remaining concentrations were computed as $(1 - \bar{p})^{-1} \mathbf{c}^{\text{init}}$, where $\mathbf{c}^{\text{init}}$ is the vector representing the total feedforward input to neurons that have more than 95% of their receptors active, and $\bar{p}$ represents the sub-matrix of connection weights between these neurons.

C. 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 subject is presented with an odor (a mixture of odorants) and is asked to report whether a target odorant is present in the odor or not. The experiments of (32) were conducted as follows. 13 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.