LETTER **Communicated by Saket Navlakha**

Deconstructing Odorant Identity via Primacy in Dual Networks

Daniel R. Kepple
dkepple@cshl.edu
Hamza Giaffar
hgiaffar@cshl.edu
Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724, U.S.A.

Dmitry Rinberg
rinberg@nyu.edu
Neuroscience Institute, New York University School of Medicine, New York, NY 10016, U.S.A.

Alexei A. Koulakov
akula@cshl.edu
Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724, U.S.A.

In the olfactory system, odor percepts retain their identity despite substantial variations in concentration, timing, and background. We study a novel strategy for encoding intensity-invariant stimulus identity that is based on representing relative rather than absolute values of stimulus features. For example, in what is known as the primacy coding model, odorant identities are represented by the conditions that some odorant receptors are activated more strongly than others. Because, in this scheme, odorant identity depends only on the relative amplitudes of olfactory receptor responses, identity is invariant to changes in both intensity and monotonic nonlinear transformations of its neuronal responses. Here we show that sparse vectors representing odorant mixtures can be recovered in a compressed sensing framework via elastic net loss minimization. In the primacy model, this minimization is performed under the constraint that some receptors respond to a given odorant more strongly than others. Using duality transformation, we show that this constrained optimization problem can be solved by a neural network whose Lyapunov function represents the dual Lagrangian and whose neural responses represent the Lagrange coefficients of primacy and other constraints. The connectivity in such a dual network resembles known features of connectivity in olfactory circuits. We thus propose that networks in the piriform cortex implement dual computations to compute odorant identity with the sparse activities of individual neurons representing Lagrange coefficients. More generally, we propose that sparse neuronal firing rates may

Neural Computation **31, 710–737** (2019)
doi:10.1162/neco_a_01175

represent Lagrange multipliers, which we call the dual brain hypothesis. We show such a formulation is well suited to solve problems with multiple interacting relative constraints.

1 Introduction

Sensory systems face the problem of identifying stimulus features that are invariant to various transformations. The olfactory system, for example, has to identify stimuli despite substantial variations in odorant concentration. Computing concentration invariant odor identity is necessary to enable a stable perception in chemical gradients and turbulent odorant plumes. How can the olfactory system robustly represent odorant identity despite variable stimulus intensity? The first step of olfactory processing involves odorants binding to and activating a set of molecular sensors known as olfactory receptors (ORs). ORs are proteins expressed by olfactory sensory neurons (OSNs) located in the nose. Most mammalian olfactory systems contain about 1000 types of ORs, while humans rely on the responses of only 350 (Firestein, 2001; Koulakov, Gelperin, & Rinberg, 2007; Zhang & Firestein, 2002). Importantly, every OSN expresses only a single type of OR chosen randomly out of the large ensemble. Odorant identity is then interpreted from the patterns of activation of OR and, by extension, OSNs.

Here we examine the hypothesis that stimulus identity is inferred based on the relative amplitudes of OSN responses. In particular, we propose that odorant identity can be determined using only the information that a subset of receptors individually responds more strongly than all other receptors. We call this type of representation the primacy model. The primacy model is inspired by a recent experimental observation that odorant identity is recognized on the basis of inputs present within the first 100 milliseconds of the animal's sniff cycle (Wilson, Serrano, Koulakov, & Rinberg, 2017). We will show that the coding scheme based on the primacy model yields odorant representations that are independent of absolute odorant concentration. We formulate this identity decoding scheme using a dual Lagrange-Karush-Kuhn-Tucker problem and argue that this dual problem can be solved by a neural network that we call a dual network. Dual networks that implement the primacy model share many features with real olfactory networks. Our goal is therefore to derive the structure of olfactory circuits from first principles, based on the primacy model.

2 Results

2.1 Representing Odorants by Sparse Vectors. Ethologically important odorant stimuli are mixtures of monomolecular components. Such stimuli can be represented by a vector of concentrations $\vec{x}$. Each component of this vector x_j is equal to the concentration of an individual monomolecular

component numbered by index j. The number of potential monomolecular components, which is equal to the dimensionality of vector $\vec{x}$, will be denoted here by M. This number can be estimated to be around several million, $M \sim 10^6$, based on the count of potentially volatile molecules with molecular weight less than 300 Dalton in the popular database PubChem (Kim et al., 2016). However, ethological odorant mixtures do not contain all of these molecules at the same time and are therefore represented by sparse vectors $\vec{x}$. For the purposes of the olfactory system, sparseness of the concentration vector is further increased by the system's inability to detect or recognize individual components. Indeed, psychophysical studies suggest that the number of monomolecular components of the vector $\vec{x}$ detectable by a human observer K is close to 12 (Jinks & Laing, 1999). Overall, we suggest that ethologically relevant odorants can be defined by highly dimensional ($M \sim 10^6$), sparse concentration vectors $\vec{x}$ with very few nonzero elements ($K \sim 10$).

Odorant mixtures then enter the olfactory system through the responses of ORs $\vec{r}$. These responses, to the first approximation, can be represented by a linear nonlinear function of the concentration vector $\vec{x}$. In the simplest model, using receptors with only one binding site and no cooperativity, the law of mass action yields (see the appendix)

$$r_i = F(y_i), \tag{2.1}$$

$$y_i = \sum_j A_{ij} x_j. \tag{2.2}$$

Here index $i = 1, \ldots, N$ enumerates OR types. The total number of functional OR types, N, varies between 350 in humans and 1100 in rodents. Matrix element A_{ij} can be interpreted as the affinity of molecule type j to receptors of type i. $F(y) = y/(1 + y)$ is the nonlinear function describing activation of a receptor (see the appendix).

The problem solved by the olfactory system can then be formulated as follows: find the identity of the odorant stimulus x_j given the set of responses of olfactory sensory neurons r_i.

2.2 Sparse Olfactory Stimulus Recovery. In many concentration regimes, it is reasonable to approximate the receptor response r_i using only its linear input y_i. In this case, the problem solved by olfactory system (how to find $\vec{x}$ given $\vec{y}$) can be reduced to solving the system of linear equations (see equation 2.2). This problem is not entirely trivial because the number of unknowns (components of $\vec{x}$, $M \sim 10^6$) is substantially larger than the number of equations (components of $\vec{y}$, $N \sim 10^3$). Some help comes from the fact that the vector of unknowns is sparse. The problem of recovering a sparse vector from a system of linear equations has been addressed in compressed sensing (Baraniuk, 2007; Donoho & Tanner, 2005, 2006). In this

framework, vector $\vec{x}$ can be found exactly, despite the fact that it contains more components than equations. This is because the vector $\vec{x}$ is sparse and the number of nonzero components is small. To be able to find $\vec{x}$ exactly, a certain condition has to be met relating parameters M, N, and K. This condition can be understood by comparing the amounts of information contained in the input and output space. One can accurately determine the unknowns if the information capacity of $\vec{y}$ exceeds that of $\vec{x}$. These two amounts can be estimated by using the statistical physics definition of the amount of information:

$$H = \log \Gamma. \tag{2.3}$$

Here H and Γ are the amount of information and the number of states a variable can take, respectively. For example, a string of N binary numbers contains $H = N$ bits of information, which follows from equation 2.3 if the number of possible values of the binary string is $\Gamma = 2^N$. The number of values taken by the receptor response vector $\vec{y}$ can then be estimated as $\Gamma_y \sim 2^N$ (for the purposes of our order-of-magnitude estimate, we assume here that the elements of $\vec{y}$ are binary). The number of possible combinations contained in the concentration vector $\vec{x}$ can be estimated as the number of ways to place K nonzero elements into M bins—that is, $\Gamma_x \sim C_M^K \sim M^K$. Similar to the estimate for receptor responses, we assume here that concentrations are binary. This assumption is reasonable if the number of distinguishable concentrations for an element of $\vec{x}$ is much less than the total number of elements in $\vec{x}$, M. Thus, equation 2.2 can be solved exactly if the number of combinations contained in vector $\vec{y}$ is larger than the number of possible states of the concentration vector $\vec{x}$, $\Gamma_y > \Gamma_x$, which, given our estimates, yields the following condition:

$$N > K \log M. \tag{2.4}$$

This condition can be recognized as the necessary condition for sparse signal recovery using l_1 norm minimization obtained by Donoho and Tanner (2005, 2006). For $K \sim 10$ and $M \sim 10^6$ we obtain the following constraint for the number of OR types necessary for the recovery of the sparse stimulus: $N > 200$, which is satisfied for both humans ($N \approx 350$) and mice ($N \approx 1100$). Thus, equation 2.2 can in principle be solved with the existing machinery in the olfactory system. This means that with linear responses, the olfactory system can reconstruct $K \sim 10$ monomolecular components given the responses of $N \approx 1100$ ORs.

To determine $\vec{x}$ given $\vec{y}$ within compressed sensing, one can use sparse signal recovery via minimization of the l_1 norm (Baraniuk, 2007; Donoho & Tanner, 2005, 2006):

$$\vec{x} = \underset{\substack{\vec{y}=A\vec{x} \\ x_j \geq 0}}{\arg\min} \sum_j x_j. \tag{2.5}$$

Here, because molecular concentrations $\vec{x}$ are constrained to be nonnegative, minimization of the l_1 norm is equivalent to minimization of the sum of the components of $\vec{x}$. Because the system is overcomplete, there are many vectors $\vec{x}$ for which $y_i = \sum_j A_{ij} x_j$ is satisfied exactly. Minimization of the l_1 norm ensures that the solution is sparse.

Solving the decoding problem using equation 2.5 is somewhat unrealistic because receptor responses are nonlinear. Furthermore, the solution of equation 2.5 is not concentration invariant. That is, doubling the mixture concentration results in doubling of each component of the concentration vector $\vec{x}$ and, as a result, the doubling of the vector of receptor responses $\vec{y} = A\vec{x}$. Solving equation 2.5 under the constraint of doubled vector $\vec{y}$ results in a doubled reconstructed vector of concentrations $\vec{x}$. Our goal, however, is to build a representation that is concentration invariant and not to be limited to regimes where receptor response is linear. More specifically, we want a representation of an odor that is invariant to multiplication of the concentration vector of that odor by a scalar. This means that we would like to obtain a framework in which doubling receptor responses $\vec{y}$ or the presence of nonlinearities, such as given by equation 2.1, does not affect the reconstructed stimulus. Such a model would yield a concentration-invariant odorant identity. Next, we show how a primacy-based decoding model can implement such a computation.

2.3 Concentration Invariant Decoding Algorithm via Primacy. To achieve an odorant representation that is invariant to both concentration fluctuations and monotonic nonlinearities of receptor responses, we reformulate the sparse recovery problem to use only relative constraints. Conceptually, the relative constraints we introduce isolate a set of receptor types, each of which responds more strongly than all other receptor types. We call this strongly responding set of receptors the primacy set P, while all other receptors are denoted by $\bar{P}$. The primacy set is expected to include $p \ll 1000$ receptor types. For these two sets we can state:

$$r_P \geq r_{\bar{P}}. \tag{2.6}$$

Here, r_P is the set of components of vector $\vec{r}$ that belong to the primacy group P of the given odorant. The expanded version of equation 2.6 is therefore $r_i \geq r_j$ for any $i \in P$ and $j \in \bar{P}$. Because the nonlinearity relating r_i and y_i is monotonic, equation 2.6 results in the same constraint on components of vector $\vec{y}$:

$$y_P \geq y_{\bar{P}}. \tag{2.7}$$

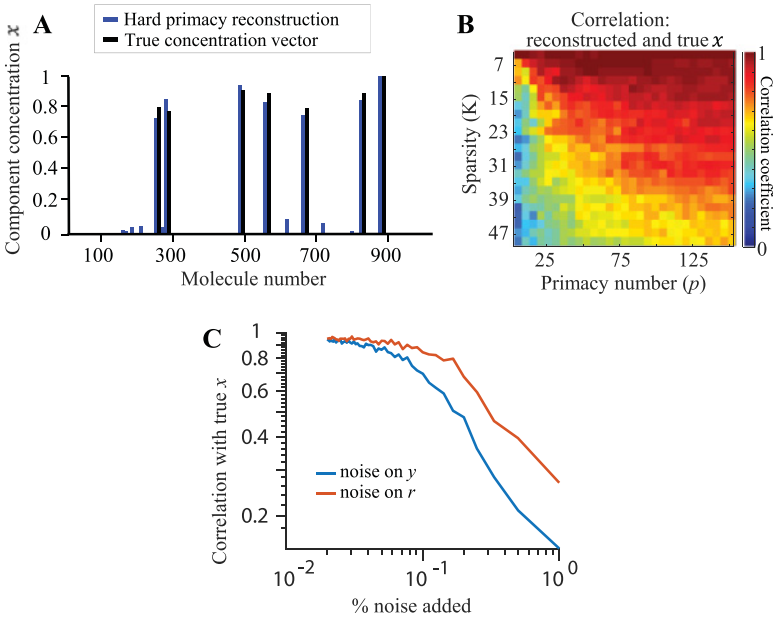

Figure 1: Using relative values of receptor responses can solve the problem of recovering sparse concentration vector $\vec{x}$. (A) An example concentration vector x alongside its reconstruction (blue) using only relative information that a group of receptors P respond more strongly than the rest. (B) Correlation between the reconstructed and true concentration vectors for different sparsity (K) values and number of primary receptors (p). (C) Correlation between reconstruction and stimulus with various levels of noise injected on input signal $\vec{y}$ (blue) and input neuron response $\vec{r}$ (red).

Importantly, inequality 2.7 is invariant to changes in concentration because concentration changes ($\vec{x} \to \alpha\vec{x}$) multiply both sides of equation 2.7 by the same factor. Also, the identities of the ORs forming the primacy set, P, are unaffected by the nonlinearity relating r to y, so long as this nonlinearity is monotonically increasing. Thus, reconstructions of $\vec{x}$ using these inequalities inherently have the desired invariances. We therefore reformulate the sparse recovery problem, equation 2.5, as follows:

$$\vec{x} = \underset{\substack{y_P \geq y_{\bar{P}} \\ x_j \geq 0}}{\arg\min} \sum_j x_j. \tag{2.8}$$

We implement this minimization and show that it converges to reconstruct the stimulus vector $\vec{x}$ in Figure 1. For this simulation and all

subsequent simulations in this letter, we use $N = 400$ receptors and $M = 1000$ molecules. We generate affinities A_{ij} from a zero-mean gaussian distribution truncated at zero such that all affinities are nonnegative. In Figure 2 in the online supplement, we show that this assumption on the distribution of A_{ij} is unnecessary and that our algorithm works for uniform or log-normally distributed A_{ij} matrices for a wide range of the number of receptors N and molecules M.

The nonlinearity in responses of each receptor, equation 2.1, is likely to be similar but not identical. The primacy model would not yield invariant odorant representations when nonlinearities are highly variant for each component of $\vec{y}$. In this case, large changes in concentration for certain odorants would elicit a different perception, as sometimes seen in human perception (Arctander, 1969).

Overall, using the relative rather than the absolute values of sensor responses to recover odorant identity results in invariance with respect to stimulus intensity. Using relative responses of sensors has a long history in neuroscience and includes color constancy (Foster, 2011), responses of differentiating (on/off-center) cells in retina, edge detection in the visual cortex (Rodieck, 1998), and relying on the relative responses of OSNs in the olfactory bulb implemented by normalization (Cleland et al., 2011; Cleland, Johnson, Leon, & Linster, 2007). Although we demonstrated this idea for the particular example of primacy coding, we propose that this neural relativity rule is used to produce perceptual invariance and invariant signal recovery with respect to other stimulus transformations.

In addition to invariance to changes in concentration and monotonic transformations of receptor signals, primacy codes are inherently robust to noise. Not only are the strongest responding receptors the most reliable, but many neighboring concentration vectors share the same primacy set. Furthermore, neighboring primacy sets can produce similar reconstructions. The result is a code unaltered by even high levels of noise, as shown in Figure 1.

Instead of minimizing the pure l_1 norm to find sparse solutions, one can minimize the elastic net functional:

$$\vec{x} = \underset{\substack{y_P \geq y_{\bar{P}} \\ x_j \geq 0}}{\arg\min} \sum_j \left(x_j + \varepsilon \frac{x_j^2}{2} \right) \tag{2.9}$$

Equation 2.9 is the same as the l_1 norm minimization, equation 2.8, when the parameter $\varepsilon \to 0$. For small parameter ε, equation 2.9 yields sparse solutions similar to equation 2.8. Problem 2.9, however allows for an easier dual space formulation. We will therefore use the elastic net minimization, equation 2.9, to recover olfactory stimuli in the rest of the letter with a sufficiently small parameter ε such that the elastic net and l_1 solutions are close.

2.4 Number of Receptors Needed to Implement Primacy Coding. Since we are using relative response information to decode the sensory input, we have to amend our estimate (see equation 2.4) of the number of receptors necessary to fully recover a sparse stimulus. In the primacy model, the amount of information in $\vec{y}$ is reduced as we are concerned only with the p strongest responding receptors. Therefore, the number of combinations in $\vec{y}$ is given by the binomial coefficient $\Gamma_y \sim C_N^p \sim N^p$. By keeping the estimate for the number of combinations in $\vec{x}$ the same, $\Gamma_x \sim M^K$, and assuming that for reconstructing $\vec{x}$, the number of combinations of $\vec{y}$ has to exceed the number of combinations in $\vec{x}$, we easily arrive at the following condition:

$$p > K\frac{\log M}{\log N}. \tag{2.10}$$

For a typical mammal, such as mouse, in which the number of ORs is approximately $N \approx 1000$, we obtain $p > 2K$. For humans, in which the number of functional ORs is $N \approx 350$, the primacy number (the size of the primacy receptor set) required is somewhat higher, $p > 2.4K$. In both estimates, we assumed that the number of molecule types available in the environment is close to $M \sim 10^6$. If one assumes that the number of discernable molecular components in each mixture is $K \sim 12$, as follows from human psychophysics (Jinks & Laing, 1999), we obtain the primacy number $p > 30$, which is substantially fewer than the total number of receptors in humans $N \approx 350$.

2.5 Hard Primacy Conditions. Problem 2.9 includes various inequality constraints, including the primacy constraints (see equation 2.6) $y_P \geq y_{\bar{P}}$. These conditions can be reformulated in a more convenient form. Indeed, conditions 2.6 have no scale in them. Therefore, minimization of the l_1 norm or elastic net functional of $\vec{x}$ connected to $\vec{y}$ via a set of positive coefficients A ($\vec{y} = A\vec{x}$) results in the trivial solution $\vec{x} = 0$. To obtain a nonzero solution, one has to introduce a finite scale into the conditions 2.6. A set of conditions with the constrained magnitude is

$$y_P \geq \gamma, \gamma \geq y_{\bar{P}} \tag{2.11}$$

where $\gamma \geq 0$ is the scale parameter. The value of this parameter is arbitrary, so we use $\gamma = 1$ in all simulations. In addition to providing a nonzero solution, introduction of a scale parameter also reduces the number of constraints in equation 2.6 from $p(N - p)$ to N. By introducing a sign variable $u_i = +1$ for $i \in P$ and $u_i = -1$ for $i \in \bar{P}$, equation 2.11 can be rewritten as a single set of conditions,

$$u_i(y_i - \gamma) \geq 0. \tag{2.12}$$

Here $y_i = \sum_j A_{ij} x_j$ where the affinity of receptor i to odorants j is given by a set of nonnegative numbers A_{ij}. Problem 2.9 combined with constraints 2.12 represents our primal optimization problem. Constraints defined by equation 2.12 will be called the hard primacy conditions.

2.6 Soft Primacy Conditions. As an alternative to equation 2.12, we propose a simpler set of conditions:

$$y_i - u_i \gamma \geq 0. \tag{2.13}$$

These conditions operate differently for primary versus nonprimary receptors. For the primary set, $i \in P$, they are equivalent to conditions in equation 2.12, that is, $y_i \geq \gamma$. Therefore, these inequalities ensure that the responses of primary receptors are larger than parameter γ. For nonprimary receptors, these inequalities become $y_i \geq -\gamma$. Since $y_i = \sum_j A_{ij} x_j$ and both A and x are nonnegative, this inequality is always satisfied. Therefore, the equation 2.13 conditions define the lower limit for the primary set of receptors, but not for the nonprimary receptors. We therefore call inequalities in equation 2.13 the soft primacy conditions.

Interestingly, however, the soft primacy conditions implicitly constrain the nonprimary receptors as well. Although the equation conditions 2.13 for nonprimary receptors become trivial, that is, $y_i \geq -\gamma$, minimization of vector $\vec{x}$ minimizes the unconstrained elements of vector $\vec{y}$, or nonprimary elements. In practice, we find that after minimizing the ℓ_1 norm of $\vec{x}$, nonprimary receptors obey $y_{\bar{p}} < \gamma$, as required by the hard primacy conditions. Thus, for practical purposes, the soft primacy conditions are equivalent to the hard ones. As we showed before (Kepple, Rinberg, & Koulakov, 2016), the hard primacy conditions result in dual networks with an interaction matrix dependent on odor stimuli, vector u_i, which is difficult to implement biologically. Below, we show that the soft primacy conditions can be implemented by a network with weights independent of the odorants presented. Thus, the network implementation based on the soft primacy conditions is more biologically plausible. We will therefore adopt the soft primacy conditions for the remainder of this letter. The approximate equivalence between soft and hard primacy conditions is illustrated in Figure 2. We therefore formulate the following primal problem that can potentially be solved by the olfactory system:

$$\vec{x} = \underset{\substack{y_i - u_i \gamma \geq 0 \\ x_j \geq 0}}{\arg\min} \sum_j \left(x_j + \varepsilon \frac{x_j^2}{2} \right). \tag{2.14}$$

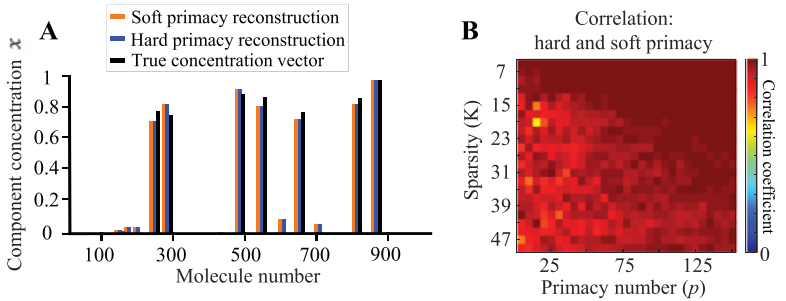

Figure 2: The solutions provided by soft and hard primacy conditions are similar. (A) An example concentration vector $\vec{x}$ alongside its hard primacy reconstruction (blue) and the soft primacy reconstruction (orange). (B) The correlation between the hard and soft primacy solutions for various sparsity K and primacy number p.

2.7 Formulation of the Dual Problem. The primal problem (see equation 2.14) cannot be easily solved by a neural network. To implement a neural network solution, we use its dual formulation. To transform the primal minimization problem (see equation 2.14) into its dual problem, we introduce a cost function (Lagrangian) with two sets of Lagrange multipliers, α_i and β_j. Each multiplier α_i, with $i = 1, .., N$, enforces an individual soft primacy constraint (see equation 2.13), $y_i - u_i\gamma \geq 0$. Each β_j enforces a constraint $x_j \geq 0$ for $j = 1..M$. The full Lagrangian for the elastic net minimization problem, equation 2.14, is

$$L(\vec{x}, \vec{\alpha}, \vec{\beta}) = \sum_j (x_j + \varepsilon x_j^2/2) - \sum_i \alpha_i \left(\sum_j A_{ij} x_j - u_i \gamma \right) - \sum_j \beta_j x_j. \tag{2.15}$$

In this equation, the Lagrange multipliers α_i and β_j represent the importance of different constraints. Because the cost function (see equation 2.15) will be minimized with respect to $\vec{x}$, Lagrange multipliers are constrained to have nonnegative values $\alpha_i, \beta_j \geq 0$ (Boyd & Vandenberghe, 2004). This ensures that at the minimum of the cost function, the target conditions, $y_i - u_i\gamma \geq 0$ and $x_j \geq 0$, are satisfied. If, for example, one of the coefficients β_j were negative in equation 2.15, decreasing the corresponding x_j below zero may be found to be advantageous from the point of view of minimizing the cost function.

To derive the dual formulation of the primal optimization problem, equation 2.14, we minimize the Lagrangian (see equation 2.15) with respect to $\vec{x}$ as if no constraints were present. We therefore find the optimal value

of $\vec{x}$ denoted here as $\vec{x}^*$. We then rewrite the Lagrangian with this optimal value of $\vec{x}$, which depends only on $\vec{\alpha}$ and $\vec{\beta}$. The resulting function is called the dual Lagrangian $\theta(\vec{\alpha}, \vec{\beta})$:

$$\frac{\partial L(\vec{x}, \vec{\alpha}, \vec{\beta})}{\partial \vec{x}} = 0 \rightarrow \quad x_j^* = \frac{1}{\varepsilon}\left(\sum_i \alpha_i A_{ij} + \beta_j - 1\right), \tag{2.16}$$

$$\theta(\vec{\alpha}, \vec{\beta}) \equiv L(\vec{x}^*, \vec{\alpha}, \vec{\beta}) = -\frac{1}{2\varepsilon}\sum_{ik}\alpha_i G_{ik}\alpha_k - \frac{1}{\varepsilon}\sum_{ij}\alpha_i A_{ij}\beta_j - \\ - \frac{1}{2\varepsilon}\sum_j(\beta_j^2 - 2\beta_j) + \gamma\sum_i \alpha_i u_i + \frac{1}{\varepsilon}\sum_i \alpha_i \sum_j A_{ij}. \tag{2.17}$$

Here $G_{ik} = \sum_j A_{ij}A_{kj}$ is the Gram matrix (the matrix of pairwise scalar products) for rows of affinity matrix $\hat{A}$.

According to optimization theory (Boyd & Vandenberghe, 2004), the dual Lagrangian $\theta(\vec{\alpha}, \vec{\beta})$ is to be maximized to find the optimal values of $\vec{\alpha}$ and $\vec{\beta}$. These values can then be used to find the solution of the primal problem using equation 2.16. One can think of the dual Lagrangian, a function of the constraint multipliers $\vec{\alpha}$ and $\vec{\beta}$, as a lower bound for the primal problem. Thus, we maximize the dual Lagrangian because we want to find the greatest lower bound for our minimization problem. The dual problem can therefore be formulated as

$$(\vec{\alpha}^*, \vec{\beta}^*) = \underset{\substack{\alpha_i \geq 0 \\ \beta_j \geq 0}}{\arg\max}\, \theta(\vec{\alpha}, \vec{\beta}). \tag{2.18}$$

The dual problem is formulated in terms of Lagrange multipliers $\vec{\alpha}$ and $\vec{\beta}$. This makes it different from the primal optimization problem, equation 2.14, which is formulated in terms of $\vec{x}$. The solution to the dual problem, however, is also the solution to the primal problem, since they are connected via equation 2.16.

The reason the dual problem, equation 2.18, is sometimes preferred to the primal problem, equation 2.14, is in the simplicity of its constraints. In many cases, the constraints $\alpha_i, \beta_j \geq 0$ are easier to implement than the primal problem's inequalities. The important observation that we make here is that neural systems are well suited for implementing dual optimizations. For example, it is especially easy to impose nonnegativity constraints because neural responses are described by firing rates that cannot fall below zero. Motivated by this observation, we argue that the Lagrange multipliers $\vec{\alpha}$ and $\vec{\beta}$ could be represented by responses of different types of olfactory neurons that solve the dual rather than the primal representation problem.

Another motivation for linking neural activity with Lagrange multipliers can be derived from the Karush-Kuhn-Tucker theorem (KKTT) (Boyd & Vandenberghe, 2004; Kuhn & Tucker, 1951). According to KKTT, at the maximum of a dual Lagrangian, such as equation 2.17, the Lagrangian contributions in equation 2.15 vanish:

$$\alpha_i \left(\sum_j A_{ij} x_j - u_i \gamma \right) = 0, \tag{2.19}$$

$$\beta_j x_j = 0. \tag{2.20}$$

These equations are valid for all values of i and j. The former equation implies that either $\alpha_i = 0$, in which case $\sum_j A_{ij} x_j - u_i \gamma$ can assume any nonnegative value (inactive constraint), or $\sum_j A_{ij} x_j - u_i \gamma = 0$, which allows α_i to be nonzero (active constraint). The Lagrange coefficients $\alpha_i = 0$ describe constraints that have no impact on the solution of the optimization problem, explaining why they are called inactive. This observation implies that many α_i are expected to be zero, making the vector of responses $\vec{\alpha}$ sparse. This observation could provide a rationale to the observed sparsity of neural activities in the olfactory system (Kay & Laurent, 1999; Koulakov & Rinberg, 2011; Rinberg, Koulakov, & Gelperin, 2006; Stettler & Axel, 2009) and beyond (DeWeese & Zador, 2006; Hromadka, DeWeese, & Zador, 2008; Lehky, Sejnowski, & Desimone, 2005; Vinje & Gallant, 2000). A mapping between the dual problem and neural responses could thus explain sparsity of neural responses as a corollary of KKTT.

2.8 Dual Networks. We now describe neural networks that can solve the dual problem, equation 2.18. We associate vectors $\vec{\alpha}$ and $\vec{\beta}$ with the firing rates of two groups of neurons (cell types). The conditions $\alpha_i \geq 0$ and $\beta_j \geq 0$ are then satisfied automatically as the firing rates of neurons cannot be negative. We assume that the neurons are connected to a network and the Lyapunov function of that network $H(\vec{\alpha}, \vec{\beta})$ is proportional to the negative of the dual Lagrangian. We use the negative dual Lagrangian in order to follow the convention of constructing networks that minimize a Lyapunov function (whereas the dual Lagrangian is to be maximized). For simplicity, we assume that the Lyapunov function of the network is $H(\vec{\alpha}, \vec{\beta}) = -\varepsilon\theta(\vec{\alpha}, \vec{\beta})$. For the Lyapunov function, from equation 2.17, we obtain

$$\begin{aligned} H(\vec{\alpha}, \vec{\beta}) = {} & \frac{1}{2} \sum_{ik} \alpha_i G_{ik} \alpha_k + \sum_{ij} \alpha_i A_{ij} \beta_j + \\ & + \sum_j (\beta_j^2/2 - \beta_j) - \varepsilon\gamma \sum_i \alpha_i u_i - \sum_i \alpha_i \sum_j A_{ij} \end{aligned} \tag{2.21}$$

To be able to interpret equation 2.21 as a Lyapunov function, we have to propose a network whose dynamics minimizes it. To generate network equations for each neuron in the network, we define internal variables that can be viewed as the total synaptic input current for each neuron. For α_i and β_j neurons, these currents will be denoted a_i and b_j, respectively. Consider the following equations for a_i and b_j:

$$\dot{a}_i + a_i = -\sum_k W_{ik}\alpha_k - \sum_j A_{ij}\beta_j + \varepsilon\gamma u_i + \sum_j A_{ij}, \tag{2.22}$$

$$\dot{b}_j + b_j = -\sum_i \alpha_i A_{ij} + 1, \tag{2.23}$$

$$W_{ik} = G_{ik} - \delta_{ik}, \tag{2.24}$$

$$\alpha_i = [a_i]_+, \beta_j = [b_j]_+. \tag{2.25}$$

Here and throughout this letter, we use $\dot{x} \equiv dx/dt$. Equation 2.22 describes inputs into neurons with firing rates α_i connected to each other by weights $-W_{ik}$. Connectivity between α-neurons is symmetric. These cells are also connected to β_j neurons with synaptic weights—A_{ij}, the affinity matrix. α-neurons receive an external excitatory drive equal to $\varepsilon\gamma u_i + \sum_j A_{ij}$. Equation 2.23 describes β_j neurons that are connected symmetrically to α-neurons. For both cell types, their firing rates (α_i and β_j) are related to their inputs (a_i and b_j) by rectifying threshold-linear relationships (ReLU), equation 2.25 ($[x]_+ = x$ for $x \geq 0$ and $[x]_+ = 0$ for $x < 0$). The circuit diagram for the network described here is presented in Figure 3.

It is straightforward to verify that these equations can be rewritten as a gradient descent:

$$\dot{a}_i = -\frac{\partial H(\vec{\alpha}, \vec{\beta})}{\partial \alpha_i}, \tag{2.26}$$

$$\dot{b}_j = -\frac{\partial H(\vec{\alpha}, \vec{\beta})}{\partial \beta_j}. \tag{2.27}$$

Here $H(\vec{\alpha}, \vec{\beta})$ is the Lyapunov function given by equation 2.21. To show that $H(\vec{\alpha}, \vec{\beta})$ is indeed a Lyapunov function, that is, a function that is not increasing and bounded, we follow conventional methods. To prove monotonicity of time evolution, we observe that

$$\begin{aligned}\frac{dH(\vec{\alpha}, \vec{\beta})}{dt} &= \sum_i \frac{\partial H(\vec{\alpha}, \vec{\beta})}{\partial \alpha_i}\dot{\alpha}_i + \sum_j \frac{\partial H(\vec{\alpha}, \vec{\beta})}{\partial \beta_j}\dot{\beta}_j = \\ &= -\sum_i \dot{a}_i\dot{\alpha}_i - \sum_j \dot{b}_j\dot{\beta}_j = -\sum_i f'(a_i)\dot{a}_i^2 - \sum_j f'(b_j)\dot{b}_j^2 \leq 0.\end{aligned} \tag{2.28}$$

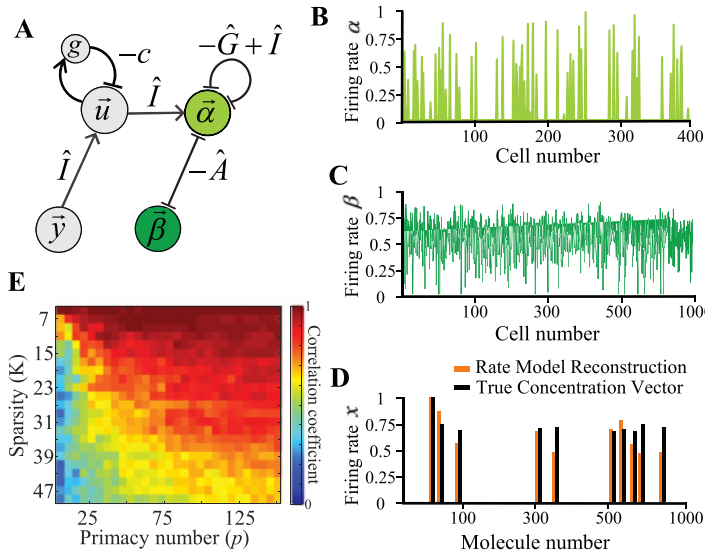

Figure 3: The dual network model described by equations 2.22 through 2.25. (A) The structure of the network. α cells (light green) implement the dual representation of the concentration vector. β cells (dark green) implement the nonnegativity constraints on the concentration vector. (B) Example firing rates of alpha cells. (C) Example firing rates of beta cells. (D) Example of the firing rate model's reconstruction (orange) of the concentration vector compared with the original concentration vector (black). (E) Correlation between the firing rate reconstruction of the concentration vector and the true concentration vector for a range of stimulus and network parameters, K and p.

Here $f(x) = [x]_+$, $f'(x) \geq 0$. Because, according to equation 2.28, $dH(\vec{\alpha}, \vec{\beta})/dt \leq 0$, and because equation 2.21 is bounded from below, our network will minimize the Lyapunov function, equation 2.21. Due to the physically imposed nonnegativity of firing rates, the variables in equation 2.25 will stay nonnegative throughout the course of this optimization, automatically satisfying the inequalities needed to solve the dual problem. We conclude therefore that our two-cell-type network can solve the dual optimization problem, equation 2.17, and thereby compute accurately molecular composition of a mixture in dual space.

2.9 Implementing the Primacy Variables u_i. The purpose of the variables u_i is to identify the set of primary variables y_i, that is, the set of p components of vector $\vec{y}$ that are larger than all others. For primary/nonprimary y_i, variables u_i are expected to be equal to $+1/-1$, respectively. To compute these variables, we introduce a p-winners-take-it-all network that identifies

p strongest inputs and suppresses the representations of the remainder of inputs. This network contains a population of inhibitory neurons, g, which are activated by the firing of p of the primacy neurons u. The equations specifying the dynamics of such a network are

$$g = h\left(\sum_k h(u_k) - p\right), \tag{2.29}$$

$$u_i = s\,(r_i - T - cg). \tag{2.30}$$

Here, $h(x)$ is the Heaviside function, while $s(x)$ is a hysteretic sign function. The latter function is activated (changes activity from -1 to $+1$) when $x > 0$ and is deactivated at very low levels of input. Other parameters include c, the strength of inhibition from neurons of type g, and T, a detection threshold of the u neurons. Odorants that do not activate receptors above T will not be perceived. Because the activities of each u_i are controlled by its input from receptor r_i, u-cells with the strongest inputs (primary) win over the cells with weaker inputs (nonprimary). This dynamic will select a certain number of cells that have the strongest receptor input (see Figure 1 in the online supplement).

We found that this network finds robust solutions if $s(x)$ function, the activation function for the u-cells in equation 2.29, is hysteretic, similar to previous studies (Sanders, Berends, Major, Goldman, & Lisman, 2013; Sanders et al., 2014; Wilson et al., 2017). This is because activities of receptors are often transient, and therefore the sustained activity of hysteretic neurons following an initial activation makes representations stable until the end of the sniff cycle. Such behavior ensures that primary cells identified early in a sniff cycle remain primary even though receptor responses undergo adaptation or are affected by other network dynamics. This behavior is consistent with the presence of nonlinear voltage-dependent synaptic currents in the piriform cortex (Poo & Isaacson, 2011).

2.10 Networks Implementing Primacy. Figure 3 displays the network implementing the dual problem, equation 2.17. It contains five cell types, each designated by a letter. Receptor neurons (y-cells) are connected to u-cells via feedforward excitatory connections. u-cells then connect to g cells and α neurons. g cells enforce the primacy conditions, and α neurons compute the dual representations of the stimulus. α-neurons interact with each other by structured inhibitory connections as indicated. They also form inhibitory connections to β-neurons described by the affinity matrix $\hat{A}$. In turn, β-neurons inhibit α-cells with the same strength. α-cells enforce the soft primacy constraints, equation 2.13. β-cells represent Lagrange-KKT coefficients that enforce nonnegativity of concentrations of individual monomolecular components of the stimulus vector $\vec{x}$.

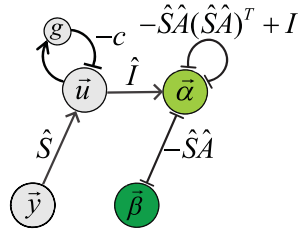

Figure 4: Network diagram of the simplicial dual network.

Overall, we propose a network that can solve dual constrained optimization problems with time-dependent firing rates. These networks capitalize on the ease with which neural firing rates can implement the dual problem's nonnegativity conditions compared to much more complex conditions in the primal problem. We therefore call these circuits *dual networks*.

2.11 Simplicial Dual Networks. In the previous network architecture (see Figure 3), the dual variables u_i receive inputs from the receptor neurons y_i via an identity weight matrix. This implies that the number of dual variables $u_i = \pm 1$ is equal to the number of ORs. This assumption is not biologically justified and is unnecessary. Instead, we could amend our soft primacy conditions, equation 2.3, to include weighted combinations of activities of ORs, that is, $y_i \to \sum_j S_{ij} y_j$. In this case, the soft primacy conditions read:

$$\sum_j S_{ij} y_j - u_i \gamma \geq 0. \tag{2.31}$$

The number of soft primacy conditions, indexed by i, is then arbitrary and is not limited by the number of receptors. Matrix S_{ij} determines the weights with which each receptor contributes to a given condition. For example, if $\hat{S} = \hat{I}$, the identity matrix, the previous primacy conditions, equation 2.13, are recovered. If $\hat{S}$ is a sparse matrix of zeros and ones, each primacy condition will constrain the sum of various receptor responses as opposed to individual receptor responses.

In the simplest case, each row of the matrix $\hat{S}$ contains exactly s nonzero values that are all equal to one. Condition 2.30 then enforces, for each $u_i = +1$, that a sum of s receptor responses is larger than γ. Thus, matrix $\hat{S}$ generalizes primacy from relating the responses of individual receptors to relating the responses of groups of receptors. We now explore the consequences of such a generalization on the architecture of our dual network. Because the subsets of s receptors form simplexes in the $\vec{y}$-space, we call this network simplicial (see Figure 4).

Since we changed the structure of our soft primacy conditions, equation 2.30, we must also update the information-theoretic argument for mixture recovery, equation 2.10. Assume that matrix $\hat{S}$ contains q rows, that is, the total number of simplexes is q. In this case, the primacy conditions imply that p sums of receptor responses are larger than the $q - p$ remaining sums. The number of distinct configurations of receptor activities can therefore be described by $\Gamma_y \propto q^p$. Since $\Gamma_x \propto M^K$, as before, we conclude that stimulus recovery in the simplicial model is possible if

$$p \geq K \frac{\log M}{\log q}. \tag{2.32}$$

This condition is less restrictive than in the case of individual receptor-based primacy (see equation 2.10). This is because the number of primacy conditions q can substantially exceed the number of receptors N.

Including the generalized soft primacy conditions, equation 2.31, into our approach is straightforward. Before, $y_i = \sum_k A_{ik} x_k$ and equation 2.13 lead to $\sum_k A_{ik} x_k \geq \gamma u_i$. Equation 2.30 states instead that $\sum_{jk} S_{ij} A_{jk} x_k \geq \gamma u_i$. To include the generalized condition into our approach, one should replace the affinity matrix A_{ik} with the new matrix $V_{ik} \equiv \sum_j S_{ij} A_{jk}$ throughout. We thus obtain, instead of equation 2.21, the following equation for the dual Lyapunov function of the network:

$$H(\vec{\alpha}, \vec{\beta}) = \frac{1}{2} \sum_{ik} \alpha_i \Gamma_{ik} \alpha_k + \sum_{ij} \alpha_i V_{ij} \beta_j + \\ + \sum_j (\beta_j^2/2 - \beta_j) - \varepsilon\gamma \sum_i \alpha_i u_i - \sum_i \alpha_i \sum_j V_{ij}. \tag{2.33}$$

Here $\Gamma_{ik} = \sum_j V_{ij} V_{kj}$ is the interaction matrix between the coefficients α, which is also a Gram matrix for the rows of matrix $\hat{V}$. The network that minimizes function 2.32 can also be derived as before (see equations 2.22 to 2.25). For completeness, we list these equations here:

$$\dot{a}_i + a_i = -\sum_k W_{ik} \alpha_k - \sum_j V_{ij} \beta_j + \varepsilon\gamma u_i + \sum_j V_{ij}, \tag{2.34}$$

$$\dot{b}_j + b_j = -\sum_i \alpha_i V_{ij} + 1, \tag{2.35}$$

$$W_{ik} = \Gamma_{ik} - \delta_{ik}, \tag{2.36}$$

$$\alpha_i = [a_i]_+, \ \beta_j = [b_j]_+. \tag{2.37}$$

These equations can be obtained from equations 2.22 to 2.25 by replacing $\hat{G} \to \hat{\Gamma}$ and $\hat{A} \to \hat{V}$.

Equations for the primacy variables u_i can be obtained from equation 2.37 by replacing r_i with $\sum_i S_{ij} r_i$.

2.12 Sparse Incomplete Representations in Simplicial Dual Networks. The Lyapunov function (see equation 2.32) is a relatively simple function on the activities of β-cells. This function could be explicitly optimized with respect to β cell activity, resulting in a substantial simplification of the network's dynamics. Indeed, the optimal value of β_j in equation 2.32 is

$$\beta_j^* = \left[1 - \sum_i \alpha_i V_{ij}\right]_+ . \tag{2.38}$$

Instead of integrating the time-dependent equations 2.34, we can evaluate the instantaneous optimal β-cell membrane voltage $b_j = 1 - \sum_i \alpha_i V_{ij}$ and equation 2.38 to obtain β-cell firing rates. Variables given by equation 2.38 are not sparse. Indeed, when the values of α-cell responses are small, during the early stages of network dynamics, β-cell firing rates would be close to one. It would only be after some time that values of β_j^* approach zero. The dense β-cell representation comes as a consequence of the KKT theorem. Because the concentration vector $\vec{x}$ is sparse, $\vec{\beta}$, which enforces nonnegativity on the zero components of $\vec{x}$, is dense. To obtain a more biologically plausible network, we introduce a variable $\bar{\beta}_j \equiv \beta_j - 1 + \sum_i \alpha_i V_{ij}$, which provides a sparse representation of the nonnegativity constraints. The optimal values of this variable are

$$\bar{\beta}_j^* = \beta_j^* - 1 + \sum_i \alpha_i V_{ij} = \left[\sum_i \alpha_i V_{ij} - 1\right]_+ \tag{2.39}$$

By comparing this equation with equation 2.16, we observe that each $\bar{\beta}_j^*$ represents a component of the reconstructed concentration vector $\vec{x}$:

$$\bar{\beta}_j^* = \varepsilon x_j. \tag{2.40}$$

As such, $\bar{\beta}_j^*$ are expected to inherit the sparsity of the concentration vector. In previous work, we suggested by independent logic that granule cells could function directly as the components of the concentration vector (Kepple et al., 2018; Koulakov & Rinberg, 2011). There, we proposed that granule cells in the olfactory bulb form a sparse representation of components in odor mixtures. In that work, we assumed that representations were either temporarily or spatially incomplete. Therefore, we called it the sparse incomplete representation (SIR) model. Here, we associate the variables $\bar{\beta}_j^*$, which represent the concentration vector $\vec{x}$, with responses of the

olfactory bulb's granule cells, which are also found to be sparse (Cazakoff, Lau, Crump, Demmer, & Shea, 2014). The dual network relying on responses $\bar{\beta}_j^*$ can therefore implement the previously proposed SIR model (Kepple et al., 2018; Koulakov & Rinberg, 2011).

That $\bar{\beta}_j^*$ represent components of the reconstructed concentration vector also has implications for later olfactory processing. In particular, as seen in equations 2.39 and 2.40, the stimulus can be reconstructed by a relatively simple operation on α cells. That is, α-cells alone form a full dual representation of the stimulus; no information about the responses of β-cells is required. α-cell representation contains a unique identifier for each odorant, which is concentration invariant and can be used by other brain regions.

Network equations describing the dual SIR network are obtained by expressing the values of β_j from equation 2.39 and plugging these values in equations 2.33 through 2.36:

$$\dot{a}_i + a_i = \varepsilon\gamma u_i - \sum_j V_{ij}\bar{\beta}_j^*, \tag{2.41}$$

$$\alpha_i = [a_i]_+ \tag{2.42}$$

$$\bar{\beta}_j^* = \left[\sum_i \alpha_i V_{ij} - 1\right]_+ . \tag{2.43}$$

These equations describe a much simpler network than 2.33 through 2.36. Indeed, as follows from equation 2.41, this new network lacks structured recurrent connectivity between α-cells (see Figure 5). Interestingly, connectivity between α- and $\bar{\beta}^*$-cells is antisymmetric: the connectivity is described by the same matrix $\hat{V} = \hat{S}\hat{A}$ in both directions but has an opposite sign.

Equation 2.40 suggests that $\bar{\beta}^*$-cells build representations of olfactory mixtures, with activities of individual neurons encoding the concentrations of individual mixture components. As we mentioned, a similar suggestion was made for the activities of granule cell neurons of the olfactory bulb, within the SIR model. We thus can identify $\bar{\beta}^*$-cells in our model with granule cells of the olfactory bulb. In agreement with this suggestion, the number of granule cells (a few million) matches our estimate of the number of monomolecular chemical compounds.

3 Discussion

Here we propose a novel model for intensity-invariant encoding of olfactory stimuli. According to the primacy model, an odor can be identified from the identities of the p strongest responding ORs. Importantly, this does not mean that our model uses only p receptors to identify the stimulus. Instead, weakly responding receptors are still informative of which odors are not present. Because the primacy model relies only on the relative

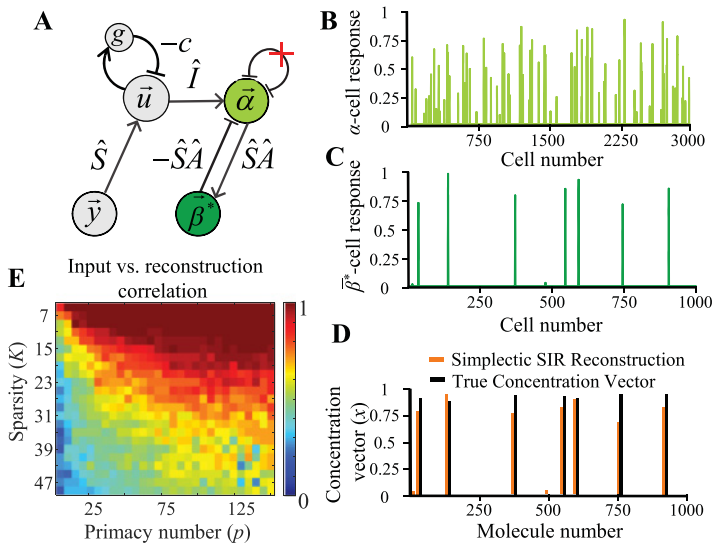

Figure 5: (A) The structure of the dual network implementing SIR model. α-cells (light green) implement the dual representation of the concentration vector. $\bar{\beta}^*$-cells (dark green) represent the reconstruction of the concentration vector. (B) Example firing rates of α-cells. (C) Example of the sparse firing rates of $\bar{\beta}^*$-cells. (D) Comparison of the simplicial SIR model's reconstruction (orange) and the concentration vector. (E) Correlation between simplicial SIR model's reconstruction of the concentration vector and the true concentration vector for a range of stimulus and network parameters, K and p.

rather than absolute strengths of receptor responses, the recovered stimulus is independent of the absolute molecular concentrations of the stimulus. Although we demonstrated this idea for the particular example of primacy coding in olfaction, we suggest that this principle could be used to produce intensity-invariant signal recovery for more complex conditions and in other modalities.

We formulated a solution of the olfactory decoding problem from first principles, using optimization theory with inspiration from compressed sensing. Since, within the primacy model, the constraints for inferring a stimulus use relative relationships between the responses of sensors (ORs), we attempted to formulate the problem using a Lagrangian approach to optimize under inequality constraints. Our solution thus involves minimizing a Lagrangian cost function containing two sets of Lagrange coefficients, $\vec{\alpha}$ and $\vec{\beta}$. Lagrange coefficients describe the importance of individual constraints—which constraints were actually used in a particular inference problem. The first set of Lagrange coefficients, $\vec{\alpha}$, described the

importance of individual primacy conditions, while the second set, $\vec{\beta}$, described nonnegativity constraints imposed on individual molecular concentrations. Solving the sparse optimization problem under inequality constraints is often performed in the dual representation via optimizing the dual Lagrangian. Dual optimization is performed in the dual space—using only the Lagrange coefficients and not the primal variables (i.e., values of molecular concentrations).

The driving observation that we make in this study is that neural systems are well suited for implementing dual optimizations. The nonnegativity of Lagrange coefficients is especially easy to enforce in neural responses described by firing rates that cannot physically fall below zero. Furthermore, a well-known theorem in optimization theory states that the Lagrange coefficient acting on an unused constraint is zero. For a dual neural network, we showed that this can result in sparse neural activity. These observations led us to propose the dual brain theory—that firing rates of individual neurons represent the Lagrange coefficients for a set of conditions of varying complexity.

It has been known for a while that recurrent neural nets can minimize cost functions of firing rates, known as Lyapunov functions (Hertz, Krogh, & Palmer, 1991). Mapping a dual Lagrangian onto a Lyapunov function of a recurrent neural network shows that this network solves the dual problem. Using this approach, we found the structure of the neural network whose Lyapunov function matches the dual Lagrangian for the problem of olfactory stimulus recovery.

Although we presented a conceptual model in which our initial approach was to implement a mathematical concept rather than to describe the biology of the olfactory system, we found that some features of our network bear resemblance to the real olfactory networks. In the dual network, the sets of Lagrange multipliers $\vec{\alpha}$ and $\vec{\beta}$ are represented by the firing rates of two cell types, α and β cells. The firing rates of β cells represent individual molecular components of the olfactory stimulus mixture. In our previous work, we proposed that granule cells form a representation of odorant components (Kepple et al., 2018; Koulakov & Rinberg, 2011). Since in this study, β cells perform a similar function, we argued that the β cells of our network (more precisely, $\bar{\beta}^*$ neurons, equation 2.39) are analogous to the granule cells in the real olfactory bulb. The number of β cells in our network is equal to the number of potential molecular components in the environment. It is encouraging that the number of granule cells (several million) (Shepherd, 1972) is also similar to the number of potential volatile monomolecular compounds present in PubChem (Kim et al., 2016).

Because of the downstream position of α neurons from receptor input $\vec{y}$ and their excitatory effect on putative granule cells β, we suggest that α cells could reside in the piriform cortex, which is known to feed back to the olfactory bulb. The inhibition from putative granule cells β to putative cortical

cells α is then possible through an indirect pathway, via inhibition of mitral cells projecting to the piriform cortex. For this reason, we suggest that receptor input instead travels through the distinct tufted cell channel. As a result, we place u cells in the anterior olfactory nucleus (AON), which receives most projections from tufted cells (Haberly & Price, 1977). Due to the *ab initio* nature of our approach, the mapping of our dual network onto the real biological network needs further refinement. Overall, we suggest that within the dual network approach, each neuronal cell type can be associated with a set of Lagrange coefficients implementing individual constraints.

In our study, we presented networks that are capable of inferring concentration-invariant odorant representations by performing a sparse norm minimization (l_1 norm or elastic net). We proposed therefore that ORs implement compression of the odorant concentration vector and relied on elements of compressed sensing in olfactory decoding. Several recent studies have argued that a decoding algorithm in the olfactory system may use compressed sensing (Grabska-Barwinska et al., 2017; Krishnamurthy, Hermundstad, Mora, Walczak, & Balasubramanian, 2017; Tootoonian & Lengyel, 2014; Zhang & Sharpee, 2016). These studies usually assume a linear encoding scheme similar to equation 2.2 and a decoding mechanism based on a sparse norm minimization. Some of these studies place the representation of individual concentration components into the piriform cortex or, equivalently, insect mushroom body (Tootoonian & Lengyel, 2014; Zhang & Sharpee, 2016). Similarly, a study of Grabska-Barwinska et al. (2017) suggests a Bayesian inference-based approach that relies on representing odors in the activity of cortical cells.

A significant distinction of our study from the previous approaches is that we assume that the activity of cortical neurons (α cells) is representative of, but distinct from, the vector of molecular concentrations. Thus, the vector of molecular concentrations can be decoded uniquely from the activities of α cells, but activities of individual neurons cannot be interpreted as molecular concentrations. In this regard, our model is similar to the recent study by Krishnamurthy et al. (2017). We place the vector of molecular concentrations into the olfactory bulb (β cells; see Figure 6). In doing so, we propose that the information of molecular composition of the odorant mixture is retained within the olfactory bulb, while cortical neurons contain a dual representation of the mixture that is synthetic (i.e., constructed rather than reconstructed). Thus, in our model, the activity of cortical cells displays similar statistics in case of both mixtures and monomolecular stimuli. The reconstruction of the stimulus into monomolecular components is performed by the granule cells in the olfactory bulb (β cells) that receive inputs from the cortex (α cells; see Figure 6). The purpose of this reconstruction is to check whether the synthetic cortical representation is consistent with any possible mixture of monomolecular compounds. In a separate work, we suggest that this reconstruction could be useful during olfactory learning (Kepple et al., 2018).

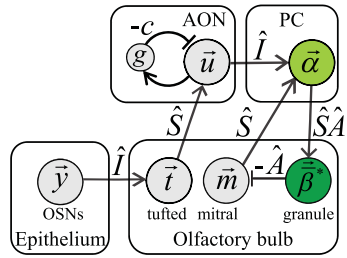

Figure 6: Possible mapping of our simplicial dual network implementing the primacy model to the known olfactory circuitry. The neurons of our network are depicted with circles, and the corresponding brain region we suggest for the neurons are shown in the surrounding box. The suggested analogy with specific cell types is given outside each circle. AON, PC, and OSNs stand for anterior olfactory nucleus, piriform cortex, and olfactory sensory neurons, respectively.

In formally linking the neurons of our dual network to biological neuronal cell types, we are able to make specific predictions about connectivity in the olfactory system. In particular, our theory suggests that feedback and feedforward connections between the bulb and cortex are dependent. Thus, the feedback connection matrix from cortex ($\hat{S}\hat{A}$; see Figure 6) is a product of feedforward connectivity ($\hat{S}$) and granule-to-mitral-cell connectivity ($\hat{A}$). We also suggest that granule cells represent the olfactory system's reconstruction of the stimulus, consequently predicting that more complex stimuli (mixtures) should evoke more complex activity in granule cells than monomolecular odorants. More generally, using our dual brain theory to connect the activity of specific cell types to a class of Lagrange coefficients enables one to determine connectivity on both mesoscopic and single-neuron scales.

Biologically it is known that OSN responses to an odor are roughly exponentially distributed. This fact has been used to motivate maximum entropy models of olfactory coding (Stevens, 2016). For these models, the exponential distribution of OSN firing rates is important as it provides a maximal entropy code under certain assumptions. For our primacy-based model, however, there is very little constraint on the distribution of firing rates. Indeed, we show that we can construct the affinity matrix A with elements that have gaussian, log-normal, or even uniform distributions (see Figure 2 in the online supplement). This finding results from the primacy model's invariance to any uniform, monotonic nonlinearity relating r to y.

Our model uses relative rather than absolute responses of ORs to solve the decoding problem. This allowed us to encode the stimulus in a concertation-invariant manner, even if the responses of receptors are nonlinear. One mechanism for concentration invariance that has been proposed previously is based on normalization of bulbar output (Banerjee et al., 2015;

Cleland et al., 2011, 2007; Kato, Gillet, Peters, Isaacson, & Komiyama, 2013; Miyamichi et al., 2013; Olsen, Bhandawat, & Wilson, 2010). This mechanism uses global inhibition to change the gain of receptor responses. The primacy model proposed here is distinct from the normalization model, although they share similar circuit features, such as broadly projecting inhibitory elements. Although we also implement inhibition to achieve concentration invariance, we place the inhibitory circuits in cortex with the goal of identifying the strongest responding receptors. Our mechanism may therefore account for fast odor-guided decisions that are dependent on receptors with high affinity to odorants. We thus propose that normalization and primacy models may operate in series on different levels of olfactory system.

Tootoonian and Lengyel (2014) have proposed that the maximum a posteriori (MAP) solution to the inference problem of recovering a sparse N-dimensional odor vector $\vec{x}$ can be achieved in a low-dimensional measurement space, reflecting the known biology of olfactory processing. In their study, the compressed sensing problem was formulated as an l_1 norm minimization of the odor concentration vector $\vec{x}$ subject to the linear equality encoding constraint $\vec{y} = \hat{A}\vec{x}$ and solved by considering the problem in dual space. The resulting generalized energy function for their network reflects the equality constraints, and, as such, their proposed network implementation differs significantly from our solution. Our networks rely on Karush-Kuhn-Tucker-type inequality conditions that implement primacy. Our formulation of the decoding problem is therefore distinct from the work of Tootoonian and Lengyel. In particular, our formulation allows decoding an odorant composition even if the encoding problem is nonlinear (see equation 2.6 and the following discussion). In addition, due to our inclusion of inequality constraints and the Karush-Kuhn-Tucker theorem, the responses of neurons in our dual network are expected to be sparse.

Two features of dual networks are worth mentioning. First, according to the Karush-Kuhn-Tucker theorem (KKTT), dual Lagrangians are optimized under the constraint of the nonnegativity of Lagrange coefficients. Neuronal responses can naturally enforce these constraints because they are described by firing rates that cannot fall below zero. Thus, dual neural networks could be viewed as analog computers that convert complex conditions, such as those imposed by primacy theory, into nonnegativity constraints. The latter constraints can be relatively easily implemented by neuronal firing rates. Second, due to KKTT (Boyd & Vandenberghe, 2004), a large number of the Lagrange coefficients are likely to be zero. This observation is consistent with the observed sparsity of neuronal responses in olfaction (Kay & Laurent, 1999; Koulakov & Rinberg, 2011; Rinberg et al., 2006; Stettler & Axel, 2009) and beyond (DeWeese & Zador, 2006; Hromadka et al., 2008; Lehky et al., 2005; Vinje & Gallant, 2000), further strengthening the possible association between neuronal responses and Lagrange coefficients. We therefore propose that neural networks may implement

optimization of dual Lagrangians with the responses of individual neurons representing Lagrange multipliers corresponding to a set of individual constraints.

Appendix: Derivation of Equations 2.1 and 2.2

Consider the mass-action law for an OR number i binding odorant number j. For the number of receptor molecules bound by the odorant, R_{ij}, we have

$$dR_{ij}/dt = -\gamma_{ij}R_{ij} + \alpha_{ij}x_j R_i^*. \tag{A.1}$$

Here α_{ij} and γ_{ij} are binding and unbinding rates, and the number of available (unbound) receptors, R_i^*, is given by

$$R_i^* = R_{i0} - \sum_j R_{ij}. \tag{A.2}$$

Here R_{i0} is the total number of OR molecules of type i exposed to odorant binding. Because in the equilibrium $dR_{ij}/dt = 0$ and $R_{ij} = \alpha_{ij}x_j R_i^*/\gamma_{ij} \equiv A_{ij}x_j R_i^*$, the total number of unbound receptors can be found from

$$R_i^* = R_{i0} - R_i^* \sum_j A_{ij}x_j.$$

We thus obtain $R_i^* = R_{i0}/(1 + \sum_j A_{ij}x_j)$ and the number of active receptors

$$R_i = \sum_j R_{ij} = R_{i0} \sum_j A_{ij}x_j / \left(1 + \sum_j A_{ij}x_j\right).$$

Assuming that the activity of the cell reflects its relative OR activation, we obtain

$$r_i \equiv R_i/R_{i0} = y_i/(1 + y_i) \tag{A.3}$$

with $y_i = \sum_j A_{ij}x_j$.

Acknowledgments

We are grateful for discussions with Petr Sulc, Venkatesh Murthy, and Dmitry Chklovskii. This work was made possible by the Swartz Foundation and NIH R01DC014366. We are sincerely grateful to the Kavli Institute for Theoretical Physics at UC Santa Barbara (grants NSF PHY-1748958,

NSF PHY-1748958, NIH grant R25GM067110, and the Gordon and Betty Moore Foundation, grant 2919.01) for support during the development of this project. This work was also supported by Aspen Center for Physics (NSF PHY-1066293).

References

Arctander, S. (1969). *Perfume and flavor chemicals (aroma chemicals)*. Montclair, NJ: Author.

Banerjee, A., Marbach, F., Anselmi, F., Koh, M. S., Davis, M. B., Garcia da Silva, P., . . . Albeaunu, D. F. (2015). An interglomerular circuit gates glomerular output and implements gain control in the mouse olfactory bulb. *Neuron*, *87*, 193–207.

Baraniuk, R. G. (2007). Compressive sensing. *IEEE Signal. Proc. Mag.*, *24*, 118–121.

Boyd, S. P., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge: Cambridge University Press.

Cazakoff, B. N., Lau, B. Y. B., Crump, K. L., Demmer, H. S., & Shea, S. D. (2014). Broadly tuned and respiration-independent inhibition in the olfactory bulb of awake mice. *Nature Neuroscience*, *17*, 569–576.

Cleland, T. A., Chen, S. Y., Hozer, K. W., Ukatu, H. N., Wong, K. J., & Zheng, F. (2011). Sequential mechanisms underlying concentration invariance in biological olfaction. *Front. Neuroeng.*, *4*, 21.

Cleland, T. A., Johnson, B. A., Leon, M., & Linster, C. (2007). Relational representation in the olfactory system. *Proc. Natl. Acad. Sci. USA*, *104*, 1953–1958.

DeWeese, M. R., & Zador, A. M. (2006). Non-gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *J. Neurosci.*, *26*, 12206–12218.

Donoho, D. L., & Tanner, J. (2005). Sparse nonnegative solution of underdetermined linear equations by linear programming. *P. Natl. Acad. Sci. USA*, *102*, 9446–9451.

Donoho, D. L., & Tanner, J. (2006). Thresholds for the recovery of sparse solutions via L1 minimization. In *Proceedings of the 40th Annual Conference on Information Sciences and Systems* (pp. 202–206). Piscataway, NJ: IEEE.

Firestein, S. (2001). How the olfactory system makes sense of scents. *Nature*, *413*, 211–218.

Foster, D. H. (2011). Color constancy. *Vision Res.*, *51*, 674–700.

Grabska-Barwinska, A., Barthelme, S., Beck, J., Mainen, Z. F., Pouget, A., & Latham, P. E. (2017). A probabilistic approach to demixing odors. *Nat. Neurosci.*, *20*, 98–106.

Haberly, L. B., & Price, J. L. (1977). Axonal projection patterns of mitral and tufted cells of olfactory-bulb in rat. *Brain Res.*, *129*, 152–157.

Hertz, J., Krogh, A., & Palmer, R. G. (1991). *Introduction to the theory of neural computation*. Redwood City, CA: Addison-Wesley.

Hromadka, T., DeWeese, M. R., & Zador, A. M. (2008). Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biology*, *6*, e16.

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

Kato, H. K., Gillet, S. N., Peters, A. J., Isaacson, J. S., & Komiyama, T. (2013). Parvalbumin-expressing interneurons linearly control olfactory bulb output. *Neuron*, *80*, 1218–1231.

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

Kepple, D. R., Cazakoff, B. N., Demmer, H. S., Eckmeier, D., Shea, S. D., & Koulakov, A. A. (2018). *Computational algorithms and neural circuitry for compressed sensing in the mammalian main olfactory bulb*. doi: 10.1101/339689

Kepple, D., Giaffar, H., Rinberg, D., & Koulakov, A. (2016). *Deconstructing odorant identity via primacy in dual networks*. arXiv:1609.02202.

Kim, S., Thiessen, P. A., Bolton, E. E., Chen, J., Fu, G., Gindulyte, A., Han, L., . . . Bryant, S. H. (2016). PubChem Substance and Compound databases. *Nucleic Acids Res.*, *44*, D1202–1213.

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

Koulakov, A. A., & Rinberg, D. (2011). Sparse incomplete representations: A potential role of olfactory granule cells. *Neuron*, *72*, 124–136.

Krishnamurthy, K., Hermundstad, A. M., Mora, T., Walczak, A. M., & Balasubramanian, V. (2017). *Disorder and the neural representation of complex odors: smelling in the real world*. arXiv:1707.01962v1.

Kuhn, H. W., & Tucker, A. W. (1951). Nonlinear programming. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley: University of California Press.

Lehky, S. R., Sejnowski, T. J., & Desimone, R. (2005). Selectivity and sparseness in the responses of striate complex cells. *Vision Res.*, *45*, 57–73.

Miyamichi, K., Shlomai-Fuchs, Y., Shu, M., Weissbourd, B. C., Luo, L., & Mizrahi, A. (2013). Dissecting local circuits: Parvalbumin interneurons underlie broad feedback control of olfactory bulb output. *Neuron*, *80*, 1232–1245.

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

Poo, C., & Isaacson, J. S. (2011). A major role for intracortical circuits in the strength and tuning of odor-evoked excitation in olfactory cortex. *Neuron*, *72*, 41–48.

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

Rodieck, R. W. (1998). *The first steps in seeing*. Sunderland, MA: Sinauer Associates.

Sanders, H., Berends, M., Major, G., Goldman, M. S., & Lisman, J. E. (2013). NMDA and GABAB (KIR) conductances: The "perfect couple" for bistability. *J. Neurosci.*, *33*, 424–429.

Sanders, H., Kolterman, B. E., Shusterman, R., Rinberg, D., Koulakov, A., & Lisman, J. (2014). A network that performs brute-force conversion of a temporal sequence to a spatial pattern: Relevance to odor recognition. *Front. Comput. Neurosci.*, *8*, 108.

Shepherd, G. M. (1972). Synaptic organization of mammalian olfactory bulb. *Physiol. Rev.*, *52*, 864–917.

Stettler, D. D., & Axel, R. (2009). Representations of odor in the piriform cortex. *Neuron*, *63*, 854–864.

Stevens, C. F. (2016). A statistical property of fly odor responses is conserved across odors. *Proc. Natl. Acad. Sci. USA*, *113*, 6737–6742.

Tootoonian, S., & Lengyel, M. (2014). A dual algorithm for olfactory computation in the locust brain. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, *27*. Red Hook, NY: Curran.

Vinje, W. E., & Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, *287*, 1273–1276.

Wilson, C. D., Serrano, G. O., Koulakov, A. A., & Rinberg, D. (2017). A primacy code for odor identity. *Nat. Commun.*, *8*, 1477.

Zhang, X., & Firestein, S. (2002). The olfactory receptor gene superfamily of the mouse. *Nat. Neurosci.*, *5*, 124–133.

Zhang, Y., & Sharpee, T. O. (2016). A robust feedforward model of the olfactory system. *PLoS Comput. Biol.*, *12*, e1004850.

Received September 22, 2017; accepted December 8, 2018.