



Convergent motifs of early olfactory processing are recapitulated by layer-wise efficient coding

Juan Carlos Fernández del Castillo^{a,b}, Farhad Pashakhanloo^a, Venkatesh N. Murthy^{a,c,d,1}, and Jacob A. Zavatone-Veth^{a,e,1}

Affiliations are included on p. 11.

Edited by Dmitry Rinberg, New York University School of Medicine, New York, NY; received September 4, 2025; accepted February 23, 2026 by Editorial Board Member John R. Carlson

The architecture of early olfactory processing is a striking example of convergent evolution. Typically, a panel of broadly tuned receptors is selectively expressed in sensory neurons (each neuron expressing only one receptor), and each glomerulus receives projections from just one neuron type. Taken together, these three motifs—broad receptors, selective expression, and glomerular convergence—constitute “canonical olfaction,” since a number of model organisms including mice and flies exhibit these features. The emergence of this distinctive architecture across evolutionary lineages suggests that it may be optimized for information processing, an idea known as efficient coding. In this work, we show that by maximizing mutual information one layer at a time, efficient coding recovers several features of canonical olfactory processing under realistic biophysical assumptions. We also explore the settings in which noncanonical olfaction may be advantageous. Along the way, we make several predictions relating olfactory circuits to features of receptor families and the olfactory environment.

olfaction | efficient coding | convergent evolution

Chemosensation is our oldest sensory modality, and for most animals it is the primary means of sensing the environment. The chemoreceptors underlying this process have a rich evolutionary history which mirrors the specialization for diverse habitats across species (1–5). These proteins evolve with dizzying speed, evincing their position at the interface between the organism and an ever-changing chemical landscape (6, 7). Surprisingly, however, the olfactory circuit in which these receptors are embedded has deep similarities across vertebrates and invertebrates (8, 9). In both lineages, organisms have evolved a large repertoire of receptors, many of them broadly tuned (though some receptors are specialists for a given ligand of high importance) (10, 11). Each primary sensory neuron typically expresses just one receptor, and only neurons expressing the same receptor converge onto a given olfactory glomerulus.

The transcriptional (12, 13) and wiring (14, 15) mechanisms for achieving this circuit organization vary widely across animals. This diversity suggests that these three motifs (broadly tuned receptors, one neuron-one receptor, and glomerular convergence) are the result of strong selective pressure, rather than evolutionary chance or molecular constraints. Why might this architecture be optimal?

To answer this question, we must specify an objective. Organisms rely on olfaction for many of their basic needs, including mating, feeding, and avoiding predation. From a computational perspective, this means that animals need to solve a wide array of tasks, including discriminating odor identity (11, 16, 17), segmenting odor landscapes (18–22), and matching single odorants to a known mixture (“pattern completion”) (23). Complicating the picture further are the rich temporal dynamics of olfactory stimuli, which induce correspondingly rich dynamics even in the earliest levels of sensory processing (17, 18, 24–27).

In light of this complexity, we consider a task-agnostic objective: maximizing mutual information. The idea that early sensory processing is organized to maximize mutual information between stimulus and neural representation is known as efficient coding (28, 29). Of course it is true that in practice, animals benefit from discarding irrelevant information. But critically, the relevance of information is often context- and task-dependent. Since higher areas of the animal brain can only access the olfactory stimulus through the glomeruli, it seems likely that up to the glomerular layer, information maximization is a plausible approximate objective. This reasoning has led many previous works to apply efficient coding principles to olfaction (30–35).

Significance

Animals from insects to rodents rely on smell as their primary means of sensing their environment. Strikingly, animals as different as ants and humans have convergently evolved similar sensory architectures to solve the common problem of recognizing smells. We show that the three core motifs of this convergently evolved architecture can be recovered using a minimal biophysically grounded model. This model is based on the efficient coding principle, which posits that early sensory circuits should be optimized to preserve as much information about the environmental stimulus as possible, subject to biophysical and anatomical constraints. Our work places the canonical olfactory circuit within a coherent normative framework, and makes predictions that may be tested in future experiments.

Author contributions: J.C.F.d.C., V.N.M., and J.A.Z.-V. designed research; J.C.F.d.C., F.P., and J.A.Z.-V. performed research; and J.C.F.d.C., F.P., V.N.M., and J.A.Z.-V. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. D.R. is a guest editor invited by the Editorial Board.

Copyright © 2026 the Author(s). Published by PNAS. This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: vnmurthy@fas.harvard.edu or jzavatoneveth@fas.harvard.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2524661123/-DCSupplemental>.

Published March 24, 2026.




Fig. 1. Overview of the model. (A) The structure of the stimulus and the network. (B) The frequency of each odorant across samples. (C) An example covariance matrix of the binarized odorants. The number of blocks (each representing a source) is a parameter we vary—see Section 2.2.

In this paper, we build on previous models for olfactory stimuli and receptor activity (30, 32) to capture the relevant biophysics (such as receptor saturation and widely distributed odorant concentrations) of the circuit. To allow tractable analysis, we ignore temporal response dynamics in this work (but see *Discussion*). Within this framework, we leverage recent technical advances from the domain of unsupervised representation learning to numerically maximize a variational lower bound for the mutual information between the stimulus and the corresponding neural response. We do this with respect to three matrices: the receptor by odorant sensing matrix W , the neuron by receptor expression matrix E , and the glomerulus by neuron connectivity matrix G (Fig. 1). In order to better respect the different timescales at play in the evolution and development of the circuit, we optimize one layer at a time, though joint optimizations yielded similar results. We find that efficient coding recovers the motifs of olfactory architecture that have now been described in several species separated by hundreds of millions of years of evolutionary time (9, 12, 13, 36, 37). Thus, this striking example of convergent evolution can be understood through the lens of optimality for information processing.

1. Setup

1.1. Stimulus and Encoding Model. We model the olfactory stimulus $c \in \mathbb{R}^N$ as a sparse vector in the space of concentrations of N monomolecular odorants (Fig. 1A). We first generate a binary vector $c_{bin} \in \{0, 1\}^N$ using a Gaussian copula. This allows us to tune the mean and covariance of c_{bin} . Means are drawn from a Gamma distribution, so that some odorants are frequent, but most are rare (Fig. 1B). We set the covariance of c_{bin} to be a block matrix, where each of the k blocks represents a mixture emitting

correlated odorants (Fig. 1C). A more realistic covariance matrix would have a nested, hierarchical structure (38), but it is less obvious how to parameterize such matrices.

We then assign an independent and identically distributed log-normal concentration with variance σ_c^2 to each odorant in the sample to generate $c \in \mathbb{R}^N$. A version of this sparse, log-normal model for c was used by Qin et al. (30), and we extend their model by incorporating varying frequencies, structured covariance, and a degree of sparsity that changes across samples. The circuit encounters one c vector at a time, and its task is to compute a useful representation of c over the distribution of possible stimuli $p(c)$. We include full details of the statistical model in *SI Appendix, section 1*.

Note that while we preserve the presence/absence statistics across samples (so that odorant α can occur more frequently than odorant β), we do not preserve odorant concentration ratios across samples (odorant α cannot occur at a typically higher concentration than odorant β). Such differences in concentration can be a critical aspect of identity coding (39). However, this structure is liable to disruption by the turbulent transport of volatile molecules to the sensory epithelium, as well as differences in molecular diffusivity (39, 40). Neglecting this structure greatly reduces the number of tunable parameters in our stimulus model.

Given the stimulus, the encoding model is

$$r = \varphi(EWc) + \xi, \quad [1]$$

where $r \in \mathbb{R}^L$ is the firing rate of the neurons,

$$\varphi(x) = \frac{x^n}{x^n + 1} \quad [2]$$

is a Hill function with coefficient n , applied element-wise, $E \in \mathbb{R}_{\geq 0}^{L \times M}$ is the neuron by receptor expression matrix, $W \in \mathbb{R}_{> 0}^{M \times N}$ is the receptor by odorant affinity matrix, and $\xi \sim \mathcal{N}(0, \sigma_\xi^2)$. This Hill function nonlinearity is consistent with measurements in the fly larva by Si et al. (41), in which the authors found that a Hill coefficient of $n = 1.46$ best fit their data. The coefficient can be higher in vertebrates (on average $n \approx 2$), where olfactory receptors are G protein-coupled receptors rather than ligand-gated ion channels (42, 43) (but see Grosmaître et al. (44) for $n \approx 1$ in vitro). We confirmed that the qualitative features of the optimal W and the optimal E did not depend on this coefficient (*SI Appendix, Fig. S1*). Note that, while receptor affinities can be negative (42, 45, 46), we have chosen not to model this inhibition in this work. We discuss this and other limitations of our model, as well as possible improvements, in *Discussion*. While both W (30, 32) and E (31) have been studied using efficient coding, most works have not considered the interplay between the two [with the exception of Lienkaemper et al. (47)].

Finally, the neural activity $r \in \mathbb{R}^L$ is integrated in regions of neuropil known as glomeruli $g \in \mathbb{R}^M$, with $M \ll L$. For simplicity, we match the number of glomeruli to the number of receptors, although in reality there could be more glomeruli [as in mice (48)] or fewer [as in mosquitoes (49)]. Here, the vector g consists of one representative dendrite per glomerulus. In paired recordings from pre- and postsynaptic neurons at the glomerulus in *Drosophila melanogaster*, Bhandawat et al. (50) found that the transfer function was sigmoid in shape, so we parameterize it using the hyperbolic tangent (again applied point-wise) with varying gain α :

$$g = G \tanh(\alpha r) \quad [3]$$

Here, G is a connectivity matrix describing the projections between neurons and glomeruli.

1.2. Layer-Wise Efficient Coding. Given this model, we perform three separate optimizations (see *SI Appendix, section 3* for details). The first is over receptor sensitivities:

$$\sup_{W|E_{can}} \hat{I}(r, c) \quad [4]$$

The notation $\sup$ indicates that we are finding the matrix W that maximizes the estimated mutual information $\hat{I}(r, c)$. When optimizing W , we constrain its elements to be positive. By E_{can} , we indicate that we are plugging in a canonical (one neuron-one receptor) E matrix into Eq. 1 and holding it fixed to optimize over W . While we will largely focus on this layer-wise approach, which allows us to flexibly test null models for each parameter matrix, we also tested joint optimization over W and E . This joint optimization yielded similar results to the separate optimizations in Eqs. 4 and 5 (*SI Appendix, Fig S2*).

Next, we optimize E given the solution W_{opt} to Eq. 4:

$$\sup_{E|W_{opt}} \hat{I}(r, c). \quad [5]$$

To ensure that our conclusions are robust to the details of the receptor matrix, we also perform this optimization with shuffled sensitivities. Varying the details of W and the parameters of the environment allows us to probe why canonical expression is so common, and when noncanonical expression might arise (49, 51). We constrain E to be positive and to sum row-wise to unity, since neurons must allocate a finite budget of receptor expression (49, 51–54).

Finally, we optimize the glomerular layer:

$$\sup_{G, \alpha|E_{opt}, W_{shuffle}} \hat{I}(g, c), \quad [6]$$

where $W_{shuffle}$ is a shuffled version of W_{opt} . This is likely to be a more realistic model of true receptors than W_{opt} due to the biophysical constraints of tuning in real receptors (something we discuss at length in Section 2.2). Again we constrain G to be positive and to sum row-wise to unity, since glomeruli receive excitatory inputs from finitely many olfactory receptor neurons (8).

We consider the trade-offs of this layer-wise approach, its alternatives, and its biological interpretation in the *Discussion*.

1.3. Maximizing Mutual Information by Proxy. The primary challenge in any efficient coding analysis is the estimation of mutual information. To enable analytical progress, strong assumptions on the stimulus and encoding are required (30–32, 47, 55, 56). Most simply, assuming linear processing of a Gaussian stimulus, the mutual information can be computed in terms of the log-determinants of the resulting covariance matrices. However, it is challenging to analytically compute mutual information in more realistic non-Gaussian models like the one we adopt here.

Numerical estimation of mutual information in high dimensions is also fraught with difficulties (56). Here, inspired by work on deep representation learning, we adopt a conservative approach based on the principle that finding an optimal encoding model does not require an estimate of the precise value of the mutual information. Instead, it is sufficient to maximize a reliably

estimable proxy objective function (57, 58). Committing to this approach limits our ability to compare different optimized models, but it allows for a much more robust and efficient optimization procedure.

As detailed in *SI Appendix, section 3*, we leverage a variational formulation for mutual information maximization first proposed by Nowozin et al. (59). Specifically, we maximize a bound on the Jensen–Shannon divergence (JSD) between the joint and marginal distributions of the stimulus and the encoding. This was inspired by previous works showing that the JSD can be estimated more reliably than the Kullback–Leibler divergence (KLD) that defines the mutual information (57, 58). We describe the trade-offs of this approach in *SI Appendix, section 3*. We handle constraints on W , E , and G using the framework of mirror descent; again, see *SI Appendix, section 3* for details.

2. Results

2.1. Receptors. We first optimized the mutual information over the sensing matrix W . Each row of this matrix represents a receptor, and each column an odorant. Thus W_{ij} is the sensitivity of the i -th receptor to the j -th odorant.

Chemoreceptors are both ancient (predating neurons, for example) and rapidly evolving (2, 6, 7, 60), given their role at the interface between the organism and the environment. Thus it seems that the matrix W is relatively tunable (notwithstanding biophysical constraints—see *Discussion*), and so we asked: To what extent can efficient coding recover the qualitative features of receptor affinities measured in biological circuits?

To guide our analysis, we turned to the well-characterized olfactory system of the *D. melanogaster* larva, which has just 21 olfactory receptor neurons, each identified by the expression of a unique receptor. The tractable size of this system permits comprehensive interrogation of the complete receptor repertoire, which was conducted by Si et al. (41). Plugging in a canonical expression pattern E as in Eq. 4, we optimized mutual information between neural activity r and the stimulus c , using parameters that matched the fly larva circuit.

The resulting W_{opt} had a nonnegligible degree of sparsity ($\approx 40\%$ entries functionally 0; see *SI Appendix, section 4*), and its nonzero elements were well fit by a log-normal distribution—see *SI Appendix, Figs. S3 and S4*. Both of these features have been derived analytically in previous theoretical work (30, 32) in the limit of vanishing neural noise. Singh et al. (34) further argue how sparsity and broad tuning can emerge from simple considerations about the dimensionality of odor and receptor space. Here, we confirmed that these results still approximately hold when neural noise is nonnegligible. We next sought to compare our optimized W to experimental results in a more detailed and biology-focused analysis, as shown in Fig. 2.

We found three qualitative matches between our optimized W and the fly larva W . First, the distribution of sensitivities across all receptors and odorants spans approximately six orders of magnitude (Fig. 2C). This heavy-tailed distribution is the object of much analysis by Si et al. (41), who argue for the computational advantages of such a code. In our results, it is difficult to confirm a precise power law to the exclusion of other heavy-tailed distributions. But the qualitative agreement shown in Fig. 2C may further support the optimality of this distribution for the entries of W .

Second, each odorant typically has at least several receptors tuned to it (Fig. 2B). This is necessary to transmit faithful information about the stimulus, since the two orders of magnitude

spanned by each receptor's Hill function activity do not cover the full range of odorant concentrations. (For Fig. 2, odorant concentrations were drawn from a log-normal distribution with $\sigma_c = 3$.)

Third, some odorants are sensed by a greater number and range of receptors than others (compare the top and bottom lines of Fig. 2B). Within the context of our optimization, it is the more frequent odorants that are afforded greater dynamic range, as shown in Fig. 2D. While this result may hold in experimentally measured receptor sensitivities, testing it directly would require quantitative knowledge about natural olfactory landscapes, which are notoriously difficult to characterize.




Fig. 2. Optimizing over W for parameters typical of the fly larva, given canonical expression. (A) The trajectory of our mutual information proxy objective over the course of the optimization (see SI Appendix, section 3 for details). (B) Optimized sensitivity per odorant (blue) and measured sensitivities in the fly larva (red). Each tick denotes the sensitivity of one receptor for that odorant. In both cases, there are typically many receptors per odorant, and their sensitivities span several orders of magnitude. Odorants are sorted by their maximum sensitivity within each panel. (C) Distribution of W_{ij} . For the optimizations, 20 runs with different random seeds are shown. (D) More frequent odorants are afforded greater dynamic range. Error bars denote SD over 20 runs (only the upper error bars are shown). (E) For fixed $N = 1,000$, the number of “specialist receptors” grows with the number of receptors. “Fit” is an analytic control obtained by fitting a log-normal distribution to the optimized W matrix. Experimental data in panels (B and C) are from Si et al. (41). Environmental parameters were set to $\sigma_c = 3$ and 32 sources.

More fundamentally, however, this last finding underscores one limitation of our efficient coding framework. Mutual information is a statistical measure of the dependence between two distributions and does not privilege any component of the stimulus for reasons of ecology or behavior. Thus, when optimized, our model circuit simply prioritizes sensing of common odorants [although importantly, this is not a generic prediction of efficient coding (31, 61, 62)]. In natural settings, by contrast, animals may have evolved receptors to sense an infrequent but critically important odorant (such as a toxin or pheromone) with high sensitivity across a wide range of concentrations, as in Sakurai et al. (63). Thus, although the trend in Fig. 2D may hold on average, it is likely to admit important exceptions in biological circuits.

The latter observations consider the circuit's ability to report information about a given odorant (studying columns of W). The complementary viewpoint is to consider a single receptor's tuning across odorants (rows of W). For example, much of the recent experimental work on olfactory receptors has emphasized the promiscuity of their binding. Rather than a “lock and key” pairing, del Marmol et al. found that ligands fit loosely in the binding pocket of the bristletail receptor *MhOR5* (64). This strategy is partly due to the odorant-receptor bottleneck ($N \gg M$), and under our model, most receptors indeed exhibit broad tuning (SI Appendix, Fig. S3). However, this bottleneck argument does not explain why some receptors seem to have narrow tuning for one or a handful of odorants (10), especially in organisms with greater numbers of receptors (such as in mice, where $M \approx 1,300$).

To explore this question, we varied the degree of the bottleneck (keeping N fixed and increasing M) and inspected the resulting optimized receptor profiles. As we increased the number of receptors, the circuit acquired an increasing number of “specialist receptors” (Fig. 2E). We formalized this using the following criterion: If the maximum sensitivity of a receptor for an odorant was at least two orders of magnitude greater than the 99th percentile sensitivity, then we counted it as a specialist for that odorant. This corresponds to testing a panel of 100 odorants, and finding that a receptor has 100x greater sensitivity for its specialized odorant than for all other odorants in the panel.

In Fig. 2E, we plot two important controls. One is a shuffled W . Here, the trend still holds, which indicates that specialist receptors are developed by pushing up the sensitivity of the specialist receptor–ligand pair, rather than pushing down the sensitivities of the specialist receptor for other ligands. Therefore, a specialist in our model is a receptor with mostly typical sensitivities, and one very high outlier.

Our other control is an analytic fit. For this control, we fit a log-normal distribution to the nonzero values W_{ij} for each value of M , then sample from that distribution to generate a W , and count specialists in that W . While the log-normal is heavy-tailed and fits $p(W_{ij})$ well for the bulk of the distribution (SI Appendix, Fig. S4), it does not generate the extreme outliers needed to produce specialist receptors under our criterion. This demonstrates that the effect is not driven simply by drawing more samples from the same distribution.

These findings suggest that specialist receptors may “look normal” until their target ligand is found. Conversely, given knowledge about a specialist receptor, it may be worth testing other odorants against that receptor to see if they are detected at reasonable concentrations [such as in Meyerhof et al. (65)].

Interestingly, despite the allocation of greater dynamic range to higher frequency odorants (Fig. 2D), we did not find that they were more likely to be targeted by specialized receptors

(SI Appendix, Fig. S5). This may reflect the fact that, when specialist receptors are saturated in the presence of their target odorant, they effectively exacerbate the bottleneck faced by the rest of the circuit ($\frac{N-1}{M-1} > \frac{N}{M}$).

In this way, optimal sensing matrices W may leverage structure in the stimulus, as well as promiscuous binding, to get useful information from a specialist receptor when the target odorant is not present. Some support for this idea can be found in studies which showed that background odors can compromise the detection of a ligand by its cognate specialist receptor (66, 67). Conversely, host plant volatiles synergistically improve pheromone detection in the silk moth *Bombyx mori* at the receptor neuron level (68).

Finally, we confirmed that the distribution of the optimal W does not depend on initialization. In Fig. 2, the initial W is a scaled log-normal: $W_{init} = e^Z / \mathbb{E}[\|c\|]$, where $\mathbb{E}[\|c\|]$ is the mean magnitude of the stimulus vector (SI Appendix, section 3). Scaling the log-normal to the minimum sensitivity, below which W_{ij} is effectively zero, did not change the results (SI Appendix, Fig. S6). We derive this minimum sensitivity and discuss its implications in SI Appendix, section 4.

To summarize, the W_{opt} which results from our optimization procedure shares a number of properties with the experimental W measured in the fly larva. The distribution of values $p(W_{ij})$ spans six orders of magnitude and the nonzero values are well fit by a log-normal density. Receptor tuning is surprisingly broad, given that we do not explicitly model the imperfect binding of real ligand–receptor pairs. Lastly, increasing M permits the emergence of specialists in circuits with larger receptor families.

2.2. Expression. Perhaps the most striking feature of canonical olfaction is the “one neuron–one receptor” rule. This rule has emerged in diverse organisms including flies (37, 69), mice (12, 53, 54), and ants (13, 70). Notably, each of these animals has developed an entirely different molecular mechanism to exert this tight control over expression.

These and other disparate mechanisms (69) suggest that the one neuron–one receptor rule may confer a strong fitness advantage. Accordingly, when we numerically maximized our proxy mutual information with respect to E in Eq. 1 for numbers typical of the fly larva, we found that a clear pattern of canonical expression emerged, as shown in Fig. 3. This was despite random initialization (Fig. 3B).

Immediately, however, we are confronted with a “chicken-or-egg” problem. In Eq. 4, we optimize W given canonical E , then in Eq. 5, we optimize E given the resulting W_{opt} . This could in principle bias the optimization over E to converge on canonical expression. To account for this possibility, in our subsequent analysis we optimize E given five plausible alternative models for W (see analysis below, and SI Appendix, Fig. S13, for details). Each of these models favors canonical expression, which indicates that the result is not merely an artifact of our layer-wise approach.

Given that the solution is robust, why is canonical expression optimal? In the limit of low neural noise, one simple way to increase mutual information is to decorrelate activity over the stimulus distribution $p(c)$ (29). Consistent with this classical prediction, we find that the distribution of variances explained by the principal components of activity is flatter under canonical expression than under random expression (Fig. 3D). The best way to decorrelate activity in this way is to place neurons at the corners of the simplex in gene expression space, since total receptor expression per neuron must sum to unity. This corresponds




Fig. 3. Optimization over E for numbers typical of the fly larva. (A) The objective function trajectory over the course of optimization. (B) The initial expression matrix. (C) The optimized expression matrix. The optimization is over E in Eq. 1, plugging in the optimal W from Fig. 2. (D) The fraction of variance in neural activity r explained by each of the top principal components, given the initial expression from panel (B) and the optimized expression from panel (C). Environmental parameters were set to $\sigma_c = 3$ and 32 sources. Statistics are computed over 1,000 samples.

precisely to the one neuron–one receptor rule. (See Lienkaemper et al. (47) for a related analysis in the case of one neuron.)

We next sought to understand if the same result would hold in a larger system. We chose a scale comparable to the adult fly, with $M = 60$ receptors and $L = 1,260$ olfactory neurons, since olfactory processing in the fly is well characterized (11). We first optimized W , and found a qualitative match to the measurements in Hallem and Carlson (11), with most receptors exhibiting broad tuning (SI Appendix, Fig. S7).

Surprisingly, when we optimized over E using this W matrix, a handful of receptors were typically coexpressed per neuron as shown in Fig. 4A. Here, and throughout this paper unless otherwise noted, we have initialized expression to be noncanonical, in order to test the robustness of canonical expression. When we initialize expression to be canonical instead, we obtain canonical expression as a solution (SI Appendix, Fig. S9). As a result, we can conclude that both canonical and noncanonical expressions are local solutions to the information-maximization problem in this setting.

Interestingly, recent work in the *Aedes aegypti* mosquito has discovered receptor coexpression in many olfactory receptor neurons (ORNs) (49, 51). Reanalysis of existing data by Adavi et al. (51) further characterized occasional exceptions to the one neuron–one receptor rule in *Drosophila*. The extent of this coexpression is controversial (69, 71) but some cases have been known for many years (72, 73).

In both species, the majority of coexpressed chemoreceptors are disproportionately close to each other in genomic and phylogenetic space (51). Adavi et al. (51) term this “coexpression by descent.” In only a minority of cases did the authors find that coexpressed receptors were far from each other (“coexpression by co-option.”) This suggests that much of the phenomenon may be a by-product of shared gene regulation. Since families of chemoreceptors frequently expand via tandem duplications




Fig. 4. Optimization over E for numbers typical of adult *Drosophila*. (A) The optimized expression E obtained from plugging in the optimized sensing matrix W_{opt} . (B) The optimized E obtained from plugging in a shuffled version of W_{opt} . (C) Covariance between W_{opt} receptor pairs which are coexpressed (purple), W_{opt} singly expressed pairs (blue), and $W_{shuffle}$ receptor pairs (orange). Covariances are computed on log sensitivities. (D) Probability of coexpression given binned covariances. For panels (A and B), environmental parameters were set to $\sigma_c = 2$ and 64 sources. Data in (C and D) are across 21 runs with varying environmental parameters σ_c and Σ_c , given by the Bottom three rows of the phase diagram in Fig. 5. The p -value for the difference in distribution between coexpressed pairs (purple, $n_c = 16,038$) and singly expressed pairs (blue, $n_s = 21,132$) is obtained using the Mann-Whitney U test. The error bars in panel (D) are computed using the Wilson score interval with $\alpha = 95\%$.

(the “birth and death model” of gene evolution) (7, 74), there are many examples of neighboring ORs with similar sequences whose coexpression might be driven by cis-regulatory elements (69).

These molecular mechanisms are unlikely to be fine-tuned for optimal information processing. However, the fact remains that different insects (like *Drosophila* and *A. aegypti*) exhibit strikingly different levels of coexpression. To understand the functional significance of these differences, it is helpful to consider the similarity between receptor affinity profiles, as shown in Fig. 4 C and D. Under this reasoning, coexpression is primarily driven by shared regulatory machinery, but once it emerges, its fitness effect depends on the similarity of the coexpressed receptors, with coexpression of more similar receptors being less detrimental. Since phylogenetically related receptors typically have similar affinity profiles (75, 76), these results are consistent with the analysis of Adavi et al. (51).

Our findings changed qualitatively between the scales corresponding to the fly larva (Fig. 3, $M = 21$) and to the adult fly (Fig. 4, $M = 60$). To better understand this puzzling result, we inspected the optimized matrix W_{opt} in the adult fly setting. Across environmental parameters, W_{opt} had a distinctive low-rank structure that was especially pronounced for higher values of σ_c (SI Appendix, Fig. S10). This is driven by our frequency model since, as indicated in Fig. 1B, only a small subset of odorants occur

with high frequencies. Flattening the frequency distribution abolished the low rank structure (SI Appendix, Fig. S11). This degeneracy in W_{opt} permits either canonical or noncanonical olfaction, suggesting that the precise details of expression do not matter when W is low rank.

Why then did we obtain such a clean canonical expression profile in our analysis of the larval fly (Fig. 3)? The answer may be in the number of receptors M , as we found that optimized receptor matrices with small M were closer to full rank (SI Appendix, Fig. S8), indicating that their receptors were more differentiated. These results suggest that organisms with only a handful of evolutionarily mature receptors cannot afford similar affinity profiles.

While potentially informative, these results also reflect a limitation of our mutual information approach. Under this objective, high-frequency odorants are allocated greater dynamic range (Fig. 2D), which induces low-rank structure in W_{opt} . In reality, organisms may need to detect common and rare odorants with equal sensitivity. Additionally, biophysical constraints limit the number of ligands a receptor can bind. Indeed, Si et al. (41) found that covariance in olfactory sensory neuron activity is primarily driven by the geometric structure of odorants. This geometry constitutes a significant constraint on receptor affinities that is not accounted for in our model, and thus real receptors are very likely to be less finely tuned than our W_{opt} matrices.

In light of these limitations, we next studied expression using a shuffled version ($W_{shuffle}$) of the optimized W matrix, as in Fig. 4B. Shuffling allowed us to preserve the scale, distribution, and sparsity of W_{opt} , while breaking the low rank structure. It also accounts, albeit coarsely, for the suboptimal tuning of real receptors. Plugging in $W_{shuffle}$ favored more canonical expression, as seen in Fig. 4B.

We next sought to understand when the one neuron-one receptor rule emerges at a scale closer to the number of receptor types found in mammals such as humans ($M = 400$) and mice ($M = 1,300$). Optimizing W and E for $M = 400$, we found very similar results compared with the $M = 60$ case; for W_{opt} , coexpression is tolerated, but for $W_{shuffle}$, near-perfect single expression emerges (see SI Appendix, Fig. S12; we did not optimize over $M = 1,300$ due to memory constraints). Critically, therefore, our results indicate that a larger number of receptors does not drive coexpression. Instead, it is the extremely low-rank structure of the optimized W matrix that permits (but does not favor) coexpression.

Real affinity matrices are most likely best modeled by something between W_{opt} and $W_{shuffle}$. They have some correlation structure, but are unlikely to be extremely low rank. We therefore tested a log-normal analytic W with block covariance structure, to model families of related receptors, and a log-normal W with Toeplitz covariance structure, to model a range of similarities across receptors in the adult fly. Regardless of receptor details, canonical expression was still largely favored (SI Appendix, Fig. S13).

To ensure that these results did not depend on specific parameter values, we next varied the two important parameters of our odorant model: the structure of the covariance matrix Σ_c and the log-normal noise parameter σ_c . We parameterized Σ_c by giving it a block structure and varying the number of blocks (Fig. 1 and SI Appendix, section 1), modeling the rough clustering of odorants into groups. Our results did not strongly depend on these environmental statistics (Fig. 5). For W_{opt} , we did not observe significant changes in E , irrespective of whether we initialized with noncanonical expression (Fig. 5A) or canonical




Fig. 5. The degree of canonical expression in optimized E matrices while varying environmental parameters σ_c and Σ_c . (A) The phase diagram using optimized W . (B) The phase diagram using shuffled W . (C–F): example expression matrices taken from the above phase diagrams. The score is $1 - \mathbb{E}[H(E_i)]$, where $H(E_i)$ measures the entropy of the i -th row of E , when the expression is viewed as a probability distribution over receptors, and the expectation is taken over the rows of E . Here, we initialize the expression matrices to be noncanonical: Each neuron expresses 3–7 receptors at roughly equal levels, for a “canonical score” of 0.74. In *SI Appendix, Fig. S9*, we show the corresponding results for canonical initialization.

expression (*SI Appendix, Fig. S9*). Thus, with the low-rank W_{opt} neither solution is strongly favored, regardless of environmental structure. For $W_{shuffle}$, however, near-canonical expression was favored across most combinations of environmental parameters (*Fig. 5B*). We found only a weak dependence on environmental parameters given $W_{shuffle}$, and the trends were reversed compared to W_{opt} . We tested two other unstructured models for W : shuffling W_{opt} within rows (in order to preserve mean receptor tuning), and fitting a log-normal distribution to W_{opt} . Canonical expression was favored under both of these models (*SI Appendix, Fig. S13*).

Finally, we sought to understand how the level of neural noise σ_0 affected these results. We swept across σ_0 for a representative pair of environmental parameters (blocks = 64, $\sigma_c = 2$) (*SI Appendix, Fig. S14*). We found that at extremely low noise levels ($\sigma_0 = 0.01$), both solutions were equally favored, but for more realistic levels ($\sigma_0 \in [0.1, 1.0]$), canonical olfaction was favored as shown above. All of the above results were run for $\sigma_0 = 0.1$, a value we chose because it sets the noise to be on the order of the mean activity (*SI Appendix, Fig. S16* and *Glomerular Convergence*). Increasing σ_0 beyond this point further favored canonical olfaction. This trend is consistent with the analytical theory of Lienkaemper et al. (47) (*Discussion*).

In sum, our analysis suggests that receptor cotuning, rather than environmental statistics, is decisive in determining optimal expression. In our model, only extremely low-rank W matrices permit significant levels of noncanonical olfaction. Biologically, such a set of receptors would have highly correlated tuning across many odorants. A number of more plausible models for W support largely canonical olfaction for realistic levels of neural and environmental noise. This may explain why the one neuron-one receptor rule has emerged in such distantly related organisms, despite the fact that these organisms sense different chemical environments with receptor families that are accordingly divergent. Conversely, our results suggest that coexpression may emerge when receptor affinities exhibit a large amount of redundancy.

2.3. Glomerular Convergence. Olfactory sensory neurons converge onto regions of neuropil known as glomeruli (77). In the canonical model, only neurons expressing the same receptor converge onto a given glomerulus (78). This “glomerular convergence” confers an intuitive advantage: Given canonical expression, the signals from each receptor can be averaged across neurons without mixing across receptors. We plugged in $W_{shuffle}$, the corresponding E_{opt} , and optimized our proxy for mutual information over G and α as in Eq. 6. We confirmed that glomerular convergence is recapitulated in our model across a wide range of environmental parameters (see *Fig. 6A–C*, and *SI Appendix, Fig. S16* for sweep).

But glomeruli do not just serve to denoise olfactory neuron responses. Careful recording of both the presynaptic ORN and the postsynaptic projection neuron (PN) in *Drosophila* has characterized other aspects of the transformation (50). Chief among these is “histogram equalization,” in which low ORN activity is amplified but higher activity is not. Since ORN activity is clustered near 0, this induces a more balanced distribution of activity in the postsynaptic neuron, which in turn enhances information transmission. This is a classical prediction from the early days of efficient coding (79).

We first confirmed that our activity clustered near 0 as in Bhandawat et al. (50) (see *Fig. 6D, Bottom axis*, and *SI Appendix, Fig. S16* for sweep). This is likely necessary due to the log-normally distributed stimulus; if the optimal W is to account for the rare, highest concentration odorants, then the bulk of stimuli will generate responses closer to baseline. Accordingly, we found that as we increased σ_c , the distribution of activity was more sharply peaked around 0 (*SI Appendix, Fig. S16*).

Optimizing over the gain parameter α in Eq. 6 resulted in a qualitative match to data from adult *Drosophila*, in which the postsynaptic activity was more evenly distributed across the dynamic range of the neuron (*Fig. 6D*). The resulting histograms were not perfectly equalized, but this may reflect the fact that increasing the entropy $H(g)$ need not increase the mutual information $I(g, c) = H(g) - H(g | c)$.




Fig. 6. An example optimization over the glomerular layer for 64 sources and $\sigma_c = 3$. (A) The course of the mutual information proxy estimate over the optimization. (B) The initial random connectivity. (C) The optimized connectivity. Neurons are sorted by their highest expressed receptor. (D) The transformation at the glomerular layer. A higher α_{opt} serves to flatten the distribution of postsynaptic activity.

In sum, our results are consistent with the idea that glomerular convergence allows for denoising and histogram equalization of olfactory receptor neuron responses when expression is canonical (although glomerular convergence can also account for variation in the abundance of different cell types—see *Discussion*). As we previously observed that canonical expression is largely favored across a range of environmental parameters, we did not investigate what connectivity emerges when multiple receptors are expressed in each sensory neuron.

3. Discussion

In this work, we set out to understand why odor coding exhibits deep similarities across vertebrates and invertebrates. Using efficient coding, we recovered three motifs of a widely shared olfactory logic: broad receptors, single receptor expression, and glomerular convergence.

3.1. Receptors. The nonzero sparsity of the optimal W and the broad distribution of its values have been derived in previous theoretical works that approximate the mutual information $I(r, c) = H(r) - H(r | c)$ as the entropy $H(r)$, an approximation

which becomes exact in the limit of no neural noise (30, 32, 34). Our analysis first recovers these results in the regime where neural noise is nonnegligible, then builds on them by studying more granular features of W_{opt} such as receptor tuning, odorant sensitivity across receptors, and specialists. These details allow us to make new predictions for experiment, particularly regarding specialist receptors.

In our model, receptor tuning is much more promiscuous (most receptors respond to most odorants) than strictly necessary to map all N odorants in an approximate labeled-line scheme. This correspondence was surprising, since broadly tuned receptors in biology might simply be a byproduct of ligand–receptor biophysics. If this were the case, then such receptors would not emerge in our unconstrained model when we initialize to zero. Additionally, recent work in the migratory locust *Locusta migratoria* has characterized strikingly narrow receptor tuning (80). This constitutes compelling evidence that broadly tuned receptor families are not inevitable.

Our results suggest instead that such tuning confers an information processing advantage. The rich theory of compressed sensing may shed light on this advantage, but it is difficult to apply classical results from this field directly to the problem at hand, since a linear measurement model is usually assumed (81). Recent progress in nonlinear compressed sensing may enable theoretical understanding of the bottleneck problem in olfaction, but much remains unknown (30, 33, 34, 82–84).

To the extent that specialists do emerge in our optimizations, they have increased their sensitivity for the target ligand rather than decreasing their sensitivity for off-target ligands, and are thus still available for sensing other odorants. One prediction of our analysis is therefore that, in the absence of the target ligand, specialist receptors may be co-opted for the sensing of other relevant molecules at realistic concentrations. This idea is indirectly supported by their diminished performance in the presence of background odorants (66, 67), which indicates that specialist receptors preserve some tuning for off-target ligands. This could be further explored simply by presenting specialist receptors with a typical panel of common odorants, ideally ones not naturally co-occurring with the target ligand.

The question of the relevant range of concentrations for odorant stimuli is critical. Wachowiak et al. (85) have recently argued that the concentrations used in typical experimental studies may exceed those encountered in the environment by several orders of magnitude. If this is true, then much of the broad tuning of receptors as currently characterized by experiment would be functionally inaccessible to the organism. On the other hand, careful experimental work (86–88), supported by models of turbulent transport (39, 40), suggests that odorant concentrations fluctuate wildly at the sensory epithelium. Animals are thus exposed to high-concentration bursts, and may not average over long timescales (16, 21).

Granting that experimental concentrations are unnaturally high, the problem is mitigated by the consideration that all studies necessarily use a limited panel of odorants. As more odorants are tested, more extremely high sensitivities will be filled in for each receptor, and the current picture of combinatorial coding might survive, even if all sensitivities are shifted upward. Our model, which only considers the relative values of sensitivity and concentration, can say little to resolve this question, which amounts to fixing the mean of the stimulus distribution $p(c)$. But our results do suggest that the variance of $p(c)$ is likely to be high, since this variance is what drives the spread in the optimized

sensitivities, and these in turn are a close match to experimental data.

There are several limitations of our receptor-level analysis. First, we have only accounted for excitation of sensory neurons by odors. This was primarily because inhibition was not detected by calcium imaging in Si et al. (41), which was our primary experimental comparison. In reality, however, antagonistic interactions play a crucial role in mammalian odor coding (42, 45, 46). Inhibitory responses also emerged when Qin et al. (30) optimized W given nonzero baseline neural activity. Mechanistically, a minimal two step model in which the odorant first binds the receptor with some affinity and then activates it with a different, potentially uncoupled affinity, explained a number of classic observations in psychophysical experiments such as synergy and overshadowing between pairs of odors (89) (see also ref. 90). In future work, we hope to revisit the above analyses after optimizing the parameters of this more realistic biophysical model (initial attempts to do so proved numerically unstable).

Second, lurking in any biophysical model is the question of where to place the nonlinearity. In Eq. 1, we have placed the expression matrix E inside of the Hill function. Therefore we are using the Hill function as an empirical fit to ORN activity in the presence of multiple odors, not as a mechanistic model for cooperativity in odorant receptor signaling. A plausible extension of our model would place E outside of the nonlinearity, then include another nonlinearity for neural activity. To our knowledge, this issue has not been considered in previous works on efficient coding in olfaction, either because they assumed linear processing of the stimulus (31, 47) or canonical expression (30, 32).

Third, on the theoretical side, the compressed sensing perspective suggests that it may be sufficient to consider bulk statistical properties of W , rather than the detailed arrangement of particular elements (91). This was part of our motivation for considering alternative, shuffled models to the optimized W when analyzing downstream processing. However, biological receptor affinity matrices are subject to odorant-specific evolutionary pressure, as evidenced by the emergence of specialist receptors, and sense different odors with different dynamic ranges (Fig. 2). Such details can be captured by optimizing W , but not by assuming independent random affinities.

Last, and most fundamentally, biological evolution is a much more constrained and local process than our numerical optimization procedure. For example, real receptors likely cannot change their affinity for one odorant without changing their affinities for others, especially since odorant geometry is a primary determinant of receptor specificity (41, 64). Furthermore, organisms develop new receptors from existing ones through a birth-and-death model (7, 74), rather than optimizing a set of randomly initialized receptors as we have done. Any attempt to encode these evolutionary dynamics in our optimization would be computationally challenging, but a principled approach could be fruitful.

3.2. Expression. One key takeaway from our work is that approximate one neuron-one receptor expression emerges across a broad range of environmental parameters and for all but one of the model receptor matrices which we considered. In any normative model, a chief concern is the extent to which such results depend on specific details of the setup. This is why we have swept over a broad range of parameter values for Σ_c , σ_c , and σ_0 , and inserted multiple models for W . The fairly general

emergence of canonical expression in these analyses is likely due in part to the decorrelation of activity (Fig. 3D) that occurs when neurons “spread out” in gene expression space by choosing one receptor. A complete answer, however, would require a more developed theory for the nonlinear setting, which we leave for future work.

Critically, we do not find evidence that receptor coexpression should be finely tuned to either the statistics of receptors or the environment. On the contrary, we find that one neuron-one receptor is typically preferred, and that exceptions to this rule should have varying effects depending on receptor similarity. Stronger claims would need to be weighed against the recent work suggesting that many coexpressed pairs are recent duplicates whose coexpression is driven simply by shared regulatory factors (51, 69). Ramdya and Benton (92) raise the possibility that such instances may reflect a transient evolutionary state, in which recent duplicates are coexpressed only until regulatory machinery has “caught up” to the duplication event by creating new cell types.

In this respect, our findings differ from a very recent theoretical analysis by Lienkaemper et al. (47), who characterized the dependence of optimal expression on environmental noise and the “signal correlation” $W\Sigma_c W^T$ using a linear-Gaussian theory. Discrepancies may be due to the nonlinearity we include in the receptor activity or the log-normality of our odor model. Our formulations also differ slightly: Whereas they add noise to a Gaussian stimulus $c = c' + \xi$ and compute $I(r, c')$, we compute $I(r, c)$, where c is log-normal with variance σ_c^2 . Thus, any trends that depend on environmental noise cannot be directly compared. On the other hand, in accordance with their theory, we do find that increasing neural noise favors canonical olfaction (SI Appendix, Fig. S14). Another important difference is that Lienkaemper et al. (47) focus on the setting where the number of neurons L is less than the number of receptors M , forcing neurons into coexpression if no receptors are to be neglected. In our setting, $L > M$, which permits neurons to spread out their activity by singly expressing without losing information.

3.3. Glomeruli. At the glomerular level, the optimized circuit averages across neurons expressing the same receptor, and uses the transfer function to spread out postsynaptic activity. This has already been understood by Bhandawat et al. (50) in terms of efficient coding, and it is the most intuitive of the three layers. In this way, it serves as a useful sanity check of our computational approach.

Although the simple denoising picture is intuitive, glomerular convergence can also mirror differences in overall levels of receptor abundance (the total number of ORNs which each singly express a given receptor, and the number of copies of that receptor each such ORN expresses). How these differences reflect the environment and experience of the animal has been well documented experimentally (52, 93), and theoretically analyzed by Teşileanu et al. (31). As it stands, our numerical approach does not support reliable conclusions about optimal allocation of receptor abundance, but this is an area for future development. For the same reason, we did not attempt to make predictions about glomerular convergence in cases of noncanonical expression.

3.4. Layer-Wise Efficient Coding. The key structural choice we made in this work was to adopt a layer-wise optimization approach: We first optimized receptor affinities W assuming canonical expression E , then optimized E given fixed W , and

finally optimized glomerular connectivity G given fixed E and W . This is both computationally convenient, and—as we will argue below—biologically interpretable.

The risk of such a procedure is that the results could depend on the order of the optimizations. As mentioned before, the analysis of canonical expression represents a chicken-or-egg problem. If we optimize W given canonical E , then optimize E given the resulting W , it might be expected that canonical expression would emerge, if this structure is somehow encoded in the optimal W . We therefore sought to check that the finding of generally canonical expression was robust to the details of the optimization procedure by optimizing expression for different plausible models of the receptor affinity matrix. For these alternative models, optimization resulted in strictly greater levels of canonical expression than the optimized W . As an additional check, jointly optimizing over W and E yielded similar results (SI Appendix, Fig. S2). Together, these controls suggest that the key takeaway—that canonical expression is typically optimal—is not an artifact of our optimization procedure.

From a biological perspective, layer-wise optimization may be a reasonable model for the evolution of the convergent motifs we consider. Organisms are not presented with the opportunity to optimize an entire sensory circuit ab initio. Instead, W , E , and G are tuned by different mechanisms, on different timescales, and in different contexts. We therefore had to approximate the constrained setting in which each motif evolved given limited experimental knowledge.

Chemoreceptor families are evolutionarily ancient and continually evolving (2, 6, 7, 36, 60), so we chose to optimize over W first. Since approximate canonical expression is widespread, we plugged in canonical E for this optimization. This also permitted direct comparison to the experimental results presented by Si et al. (41). Next, it seems reasonable to ask which gene expression programs are optimal given a broad, mature set of receptors. In the mouse, for example, the relative abundance of olfactory receptors can be adapted over the timescale of just a few hours (52, 94), although this process does not generate coordinated noncanonical expression. This framing also addresses the emergence of different expression programs across organisms which each possess mature receptor families. Hence we optimized over E given fixed W .

Glomerular convergence, which constitutes an exquisite example of wiring specificity, emerges by pruning during development (14, 15). Since examples of highly adaptive connectivity abound in organisms that learn, we chose to study glomerular convergence in the setting of a mature receptor array and largely canonical expression. Hence we optimized over G given fixed, optimized W and E .

These conceptual distinctions are naturally much cleaner than the underlying biology. For example, receptors continue to evolve after a program of noncanonical expression is established (92). However, our breakdown enables a tractable model and represents a best guess at the constraints which govern the circuit's evolution. It is of course possible that layer-wise optimization could yield sensitivities and expression patterns that do not lend themselves to glomerular convergence (95). However, given access to at least as many neurons as receptor types and canonical expression, the simple intuition that glomerular convergence allows denoising suggests that not much is being lost.

One alternative to our approach is an unconstrained end-to-end optimization of the entire circuit. Wang et al. (96) recovered largely canonical expression and glomerular convergence using unconstrained optimization on a match-to-prototype classification task. This is qualitatively consistent with our

results. Their work, however, differs in assuming a particular classification task, and starts with a very different assumption on the stimulus: They directly model W by assuming that the activity of each receptor is independent and uniformly distributed. This differs substantially from the statistics resulting from our optimized sensitivity matrices, and from those measured by Si et al. (41). This kind of unconstrained optimization can shed light on the computation instantiated by the circuit, but cannot capture the path-dependent quirks of biological evolution (97).

Another alternative would be to formalize our context-specific model as a joint dynamical system whose dynamics are decomposed across different timescales. Most simply, this would lead to a setting where we optimize W given the expression E that is optimized for each sensitivity matrix, i.e., to study the nested optimization $\sup_W \{\sup_E \hat{I}(r, c)\}$. This however could be too flexible—it seems that many organisms have “committed” to canonical expression at the level of gene regulatory mechanisms, for example, and do not have the ability to smoothly modulate receptor coexpression as receptors evolve. Correspondingly, there may be some biological merit to freezing E while W is optimized.

3.5. Dynamics. Our work does not address the dynamics of the odor landscape and its neural representation in the olfactory bulb. As mentioned above, turbulent transport of air to the sensory epithelium induces high-frequency fluctuations in concentration that are richly informative about identity and location (18, 39). Here, we have ignored these dynamics in order to focus on a statistical problem of compression: namely, the representation of a high-dimensional chemical space by a finite set of receptors. The temporal challenge of encoding rapidly fluctuating odor signals is equally daunting. Recent work has begun to bridge these statistical and dynamical pictures of olfactory sensing (26, 83, 98–101), but many questions remain regarding how to efficiently encode the temporal statistics of the olfactory world. In this vein, it could be interesting to consider minimal extensions of our model that include temporal filtering (55).

3.6. Conclusion. We reiterate in closing that mutual information is only a rough proxy for performance in olfactory processing. For example, organisms may privilege the sensing of rare but critical odors, even at the expense of more common ones—this would not be captured by our statistical model (63, 67, 102). Therefore some of our results are unlikely to transfer perfectly to biological circuits. That said, the great variety of olfactory tasks faced by the typical organism suggests mutual information as a reasonable starting point (29, 30, 102).

The correspondence between our results and experimental measurements suggests that these motifs may constitute uniquely accessible solutions to the problem of processing olfactory stimuli. As chemosensation is further explored in nonmodel organisms (2, 3, 80), the extent to which canonical principles hold will be an interesting point of focus. While receptors always reflect the organism's chemical environment, the other layers we consider (receptor expression and the downstream routing of receptor activity) are not obviously bound to any stereotypy. Especially interesting will be cases where departures from canonical olfaction are unambiguous, tunable, and clearly advantageous. Since our simplified modeling suggests that canonical olfaction is optimal, such departures should serve as a flag that something interesting is afoot.

4. Materials and Methods

Our optimization setup and the structure of the models we consider are as presented in Section 1. We defer details of the odorant scene model to *SI Appendix, section 1*, a detailed comparison of our model for OSN responses to those used in previous works to *SI Appendix, section 2*, and the full details of our optimization procedure to *SI Appendix, section 3*. We make extensive use of the automatic differentiation capabilities provided by the JAX library (103). Code to perform all optimizations and reproduce all figures is available on GitHub under an MIT License at <https://github.com/VNMurthyLab/olfactory-ec>.

The dataset of estimated *Drosophila* larva receptor affinities from Si et al. (41) that we analyzed in Fig. 2 was downloaded from A. D. T. Samuel's lab GitHub, where it is available under an MIT License: <https://github.com/samuellab/Larval-ORN/blob/master/Figure3/results/MLEFit.mat>.

Data, Materials, and Software Availability. All code for this project has been deposited in GitHub (<https://github.com/VNMurthyLab/olfactory-ec>). All other data are included in the manuscript and/or *SI Appendix*. Previously published data were used for this work (The dataset of estimated *Drosophila* larva receptor affinities from Si et al. (41) that we analyzed in Fig. 2 was downloaded from A. D. T. Samuel's lab GitHub, where it is available under an MIT License: <https://github.com/samuellab/Larval-ORN/blob/master/Figure3/results/MLEFit.mat>).

1. S. Hayden *et al.*, Ecological adaptation determines functional mammalian olfactory subgenomes. *Genome Res.* **20**, 1–9 (2009).
2. W. A. Valencia-Montoya, N. E. Pierce, N. W. Bellono, Evolution of sensory receptors. *Annu. Rev. Cell Dev. Biol.* **40**, 353–379 (2024).
3. L. R. Yohe, P. Brand, Evolutionary ecology of chemosensation and its role in sensory drive. *Curr. Zool.* **64**, 525–533 (2018).
4. T. O. Auer *et al.*, Olfactory receptor and circuit evolution promote host specialization. *Nature* **579**, 402–408 (2020).
5. H. M. Robertson, Molecular evolution of the major arthropod chemoreceptor gene families. *Annu. Rev. Entomol.* **64**, 227–242 (2019).
6. Y. Nimura, M. Nei, Extensive gains and losses of olfactory receptor genes in mammalian evolution. *PLoS One* **2**, e708 (2007).
7. M. Nei, Y. Nimura, M. Nozawa, The evolution of animal chemosensory receptor gene repertoires: Roles of chance and necessity. *Nat. Rev. Genet.* **9**, 951–963 (2008).
8. J. G. Hildebrand, G. M. Shepherd, Mechanisms of olfactory discrimination: Converging evidence for common principles across phyla. *Annu. Rev. Neurosci.* **20**, 595–631 (1997).
9. K. A. Fulton, D. Zimmerman, A. Samuel, K. Vogt, S. R. Datta, Common principles for odour coding across vertebrates and invertebrates. *Nat. Rev. Neurosci.* **25**, 453–472 (2024).
10. H. Saito, O. Chi, H. Zhuang, H. Matsunami, J. D. Mainland, Odor coding by a mammalian receptor repertoire. *Sci. Signal.* **2**, ra9 (2009).
11. E. A. Hallem, J. R. Carlson, Coding of odors by a receptor repertoire. *Cell* **125**, 143–160 (2006).
12. E. Bashkirova, S. Lomvardas, Olfactory receptor genes make the case for inter-chromosomal interactions. *Curr. Opin. Genet. Dev.* **55**, 106–113 (2019).
13. A. Brahma *et al.*, Transcriptional and post-transcriptional control of odorant receptor choice in ants. *Curr. Biol.* **33**, 5456–5466.e5 (2023).
14. W. Hong, L. Luo, Genetic control of wiring specificity in the fly olfactory system. *Genet.* **196**, 17–29 (2014).
15. D. J. Zou *et al.*, Postnatal refinement of peripheral olfactory projections. *Science* **304**, 1976–1979 (2004).
16. N. Uchida, Z. F. Mainen, Speed and accuracy of olfactory discrimination in the rat. *Nat. Neurosci.* **6**, 1224–1229 (2003).
17. C. D. Wilson, G. O. Serrano, A. A. Koulakov, D. Rinberg, A primacy code for odor identity. *Nat. Commun.* **8**, 1477 (2017).
18. T. Ackels *et al.*, Fast odour dynamics are encoded in the olfactory system and guide behaviour. *Nature* **593**, 558–563 (2021).
19. D. Rokni, V. Hemmelder, V. Kapoor, V. N. Murthy, An olfactory cocktail party: Figure-ground segregation of odorants in rodents. *Nat. Neurosci.* **17**, 1225–1232 (2014).
20. S. Kundu, A. Ganguly, T. S. Chakraborty, A. Kumar, O. Siddiqi, Synergism and combinatorial coding for binary odor mixture perception in *Drosophila*. *ENEURO* **3**, ENEURO.0056-14.2016 (2016).
21. J. A. Riffell *et al.*, Flower discrimination by pollinators in a dynamic chemical environment. *Science* **344**, 1515–1518 (2014).
22. S. Jayakumar *et al.*, Mice navigate scent trails using predictive policies. bioRxiv [Preprint] (2025). <https://doi.org/10.1101/2025.08.27.672631> (Accessed 3 September 2025).
23. J. Chapuis, D. A. Wilson, Bidirectional plasticity of cortical pattern recognition and behavioral sensory acuity. *Nat. Neurosci.* **15**, 155–161 (2011).
24. C. Y. Su, C. Martelli, T. Emonet, J. R. Carlson, Temporal coding of odor mixtures in an olfactory receptor neuron. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 5075–5080 (2011).
25. C. Martelli, J. R. Carlson, T. Emonet, Intensity invariant dynamics and odor-specific latencies in olfactory receptor neuron response. *J. Neurosci.* **33**, 6285–6297 (2013).
26. G. Laurent, Olfactory network dynamics and the coding of multidimensional signals. *Nat. Rev. Neurosci.* **3**, 884–895 (2002).
27. M. Stopfer, V. Jayaraman, G. Laurent, Intensity versus identity coding in an olfactory system. *Neuron* **39**, 991–1004 (2003).
28. R. Linsker, Self-organization in a perceptual network. *Computer* **21**, 105–117 (1988).

ACKNOWLEDGMENTS. We are grateful to G. Si and A. D. T. Samuel for making their data publicly available. J.C.F.d.C. thanks D. M. Zimmerman and I. Chandok for helpful discussions regarding the evolution and biophysics of receptors. We thank C. Lienkaemper, M. A. Younger, and G. K. Ocker for inspiring discussions regarding noncanonical expression. Moreover, we are indebted to A. D. T. Samuel and the members of his group, and to G. K. Ocker and M. A. Younger, for their thoughtful comments, which have helped us improve our manuscript. J.C.F.d.C. was supported by the Harvard Biophysics Graduate Program through training grant T32GM158477 from the NIH. F.P. was supported by the Harvard Center for Brain Science (CBS)-NTT Fellowship Program on the Physics of Intelligence. V.N.M. was supported by grant number A47994 from NTT Research, Inc. J.A.Z.-V. was supported by the Office of the Director of the NIH under Award Number DP5OD037354, and by a Junior Fellowship from the Harvard Society of Fellows. This work has been made possible in part by a gift from the Chan Zuckerberg Initiative Foundation to establish the Kempner Institute for the Study of Natural and Artificial Intelligence at Harvard University.

Author affiliations: ^aCenter for Brain Science, Harvard University, Cambridge, MA 02138; ^bProgram in Biophysics, Harvard University, Cambridge, MA 02138; ^cDepartment of Molecular and Cellular Biology, Harvard University, Cambridge, MA 02138; ^dKempner Institute for the Study of Natural and Artificial Intelligence, Harvard University, Cambridge, MA 02138; and ^eSociety of Fellows, Harvard University, Cambridge, MA 02138

29. E. P. Simoncelli, B. A. Olshausen, Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
30. S. Qin, Q. Li, C. Tang, Y. Tu, Optimal compressed sensing strategies for an array of nonlinear olfactory receptor neurons with and without spontaneous activity. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 20286–20295 (2019).
31. T. Teşileanu, S. Cocco, R. Monasson, V. Balasubramanian, Adaptation of olfactory receptor abundances for efficient coding. *eLife* **8**, e39279 (2019).
32. D. Zwicker, A. Murugan, M. P. Brenner, Receptor arrays optimized for natural odor statistics. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 5570–5575 (2016).
33. K. Krishnamurthy, A. M. Hermundstad, T. Mora, A. M. Walczak, V. Balasubramanian, Disorder and the neural representation of complex odors. *Front. Comput. Neurosci.* **16**, 917786 (2022).
34. V. Singh, M. Tchernookov, V. Balasubramanian, What the odor is not: Estimation by elimination. *Phys. Rev. E* **104**, 024415 (2021).
35. J. D. Victor *et al.*, Olfactory navigation and the receptor nonlinearity. *J. Neurosci.* **39**, 3713–3727 (2019).
36. R. Benton, *Drosophila* olfaction: Past, present and future. *Proc. R. Soc. B Biol. Sci.* **289**, 20222054 (2022).
37. L. B. Vosshall, A. M. Wong, R. Axel, An olfactory sensory map in the fly brain. *Cell* **102**, 147–159 (2000).
38. J. Y. Yang *et al.*, Restructuring of olfactory representations in the fly brain around odor relationships in natural sources. bioRxiv [Preprint] (2023). <https://doi.org/10.1101/2023.02.15.528627> (Accessed 3 September 2025).
39. A. Celani, E. Villermaux, M. Vergassola, Odor landscapes in turbulent environments. *Phys. Rev. X* **4**, 041015 (2014).
40. T. Nowotny, P. Szyszka, *Dynamics of Odor-Evoked Activity Patterns in the Olfactory System* (Springer International Publishing, 2017), pp. 243–261.
41. G. Si *et al.*, Structured odorant response patterns across a complete olfactory receptor neuron population. *Neuron* **101**, 950–962.e7 (2019).
42. J. D. Zak, G. Reddy, M. Vergassola, V. N. Murthy, Antagonistic odor interactions in olfactory sensory neurons are widespread in freely breathing mice. *Nat. Commun.* **11**, 3350 (2020).
43. A. Dewan *et al.*, Single olfactory receptors set odor detection thresholds. *Nat. Commun.* **9**, 2887 (2018).
44. X. Grosmaître, A. Vassalli, P. Mombaerts, G. M. Shepherd, M. Ma, Odorant responses of olfactory sensory neurons expressing the odorant receptor mor23: A patch clamp analysis in gene-targeted mice. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 1970–1975 (2006).
45. P. Pfister *et al.*, Odorant receptor inhibition is fundamental to odor encoding. *Curr. Biol.* **30**, 2574–2587.e6 (2020).
46. L. Xu *et al.*, Widespread receptor-driven modulation in peripheral olfactory coding. *Science* **368**, eaaz5390 (2020).
47. C. Lienkaemper, M. A. Younger, G. K. Ocker, When noncanonical olfaction is optimal. *Proc. Natl. Acad. Sci. U.S.A.* **122**, e2508439122 (2025).
48. S. M. Potter *et al.*, Structure and emergence of specific olfactory glomeruli in the mouse. *J. Neurosci.* **21**, 9713–9723 (2001).
49. M. Herre *et al.*, Non-canonical odor coding in the mosquito. *Cell* **185**, 3104–3123.e28 (2022).
50. V. Bhandawat, S. R. Olsen, N. W. Gouwens, M. L. Schlieff, R. I. Wilson, Sensory processing in the *Drosophila* antennal lobe increases reliability and separability of ensemble odor representations. *Nat. Neurosci.* **10**, 1474–1482 (2007).
51. E. D. Adavi *et al.*, Olfactory receptor coexpression and co-option in the dengue mosquito. bioRxiv [Preprint] (2024). <https://doi.org/10.1101/2024.08.21.608847> (Accessed 3 September 2025).
52. T. Tsukahara *et al.*, A transcriptional rheostat couples past activity to future sensory responses. *Cell* **184**, 6326–6343.e32 (2021).
53. B. Bintu, Y. Isogai, X. Zhuang, C. Dulac, Molecular and spatial organization of the primary olfactory system and its responses to social odors. bioRxiv [Preprint] (2025). <https://doi.org/10.1101/2025.05.02.651832> (Accessed 3 September 2025).

54. D. H. Brann *et al.*, A spatial code governs olfactory receptor choice and aligns sensory maps in the nose and brain. *bioRxiv* [Preprint] (2025). <https://doi.org/10.1101/2025.05.02.651738> (Accessed 3 September 2025).
55. N. Y. Jun, G. Field, J. Pearson, "Efficient coding, channel capacity, and the emergence of retinal mosaics" in *Advances in Neural Information Processing Systems*, S. Koyejo *et al.*, Eds. (Curran Associates, Inc., 2022), vol. 35, pp. 32311–32324.
56. D. McAllester, K. Stratos, "Formal limitations on the measurement of mutual information" in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research*, S. Chiappa, R. Calandra, Eds. (PMLR, 2020), vol. 108, pp. 875–884.
57. R. D. Hjelm *et al.*, "Learning deep representations by mutual information estimation and maximization" in *International Conference on Learning Representations* (2019).
58. F. Huszár, How (not) to train your generative model: Scheduled sampling, likelihood, adversary? *arXiv* [Preprint] (2015). <https://doi.org/10.48550/arXiv.1511.05101> (Accessed 3 September 2025).
59. S. Nowozin, B. Cseke, R. Tomioka, "F-gan: Training generative neural samplers using variational divergence minimization" in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, R. Garnett, Eds. (Curran Associates, Inc., 2016), vol. 29.
60. D. M. Bear, J. M. Lassance, H. E. Hoekstra, S. R. Datta, The evolving neural and genetic architecture of vertebrate olfaction. *Curr. Biol.* **26**, R1039–R1049 (2016).
61. A. M. Hermundstad *et al.*, Variance predicts salience in central sensory processing. *eLife* **3**, e03722 (2014).
62. G. Tkačik, J. S. Prentice, V. Balasubramanian, E. Schneidman, Optimal population coding by noisy spiking neurons. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 14419–14424 (2010).
63. T. Sakurai *et al.*, Targeted disruption of a single sex pheromone receptor gene completely abolishes in vivo pheromone response in the silkworm. *Sci. Rep.* **5**, 11001 (2015).
64. J. del Marmol, M. A. Yedlin, V. Ruta, The structural basis of odorant recognition in insect olfactory receptors. *Nature* **597**, 126–131 (2021).
65. W. Meyerhof *et al.*, The molecular receptive ranges of human tas2r bitter taste receptors. *Chem. Senses* **35**, 157–170 (2009).
66. P. Pregitzer *et al.*, Plant odorants interfere with detection of sex pheromone signals by male *Heliothis virescens*. *Front. Cell. Neurosci.* **6**, 42 (2012).
67. M. Renou, V. Party, A. Rouyar, S. Anton, Olfactory signal coding in an odor background. *Biosystems* **136**, 35–45 (2015).
68. S. Namiki, S. Iwabuchi, R. Kanzaki, Representation of a mixture of pheromone and host plant odor by antennal lobe projection neurons of the silkworm *bombyx mori*. *J. Comp. Physiol. A* **194**, 501–515 (2008).
69. K. Mika, R. Benton, Olfactory receptor gene regulation in insects: Multiple mechanisms for singular expression. *Front. Neurosci.* **15**, 738088 (2021).
70. G. L. Glotzer, P. D. H. Pastor, D. J. C. Kronauer, Transcriptional interference gates monogenic odorant receptor expression in ants. *Curr. Biol.* **35**, 5033–5047.e5 (2025).
71. D. Task *et al.*, Chemoreceptor co-expression in *Drosophila melanogaster* olfactory neurons. *eLife* **11**, e2599 (2022).
72. A. L. Goldman, W. Van der Goes, D. van Naters, CG Warr. Lessing, J. R. Carlson, Coexpression of two functional odor receptors in one neuron. *Neuron* **45**, 661–666 (2005).
73. S. Lebreton *et al.*, A *Drosophila* female pheromone elicits species-specific long-range attraction via an olfactory channel with dual specificity for sex and food. *BMC Biol.* **15**, 88 (2017).
74. A. Sánchez-Gracia, F. G. Vieira, J. Rozas, Molecular evolution of the major chemosensory gene families in insects. *Heredity* **103**, 208–216 (2009).
75. K. A. Adipietro, J. D. Mainland, H. Matsunami, Functional evolution of mammalian odorant receptors. *PLoS Genet.* **8**, e1002821 (2012).
76. B. Malnic, P. A. Godfrey, L. B. Buck, The human olfactory receptor gene family. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 2584–2589 (2004).
77. K. Mori, H. Nagao, Y. Yoshihara, The olfactory bulb: Coding and processing of odor molecule information. *Science* **286**, 711–715 (1999).
78. P. Mombaerts, Molecular biology of odorant receptors in vertebrates. *Annu. Rev. Neurosci.* **22**, 487–509 (1999).
79. S. Laughlin, A simple coding procedure enhances a neuron's information capacity. *Z. Naturforsch. C* **36**, 910–912 (1981).
80. H. Chang *et al.*, Odorant detection in a locust exhibits unusually low redundancy. *Curr. Biol.* **33**, 5427–5438.e5 (2023).
81. R. Baraniuk, Compressive sensing [lecture notes]. *IEEE Signal Process. Mag.* **24**, 118–121 (2007).
82. T. Blumensath, Compressed sensing with nonlinear observations and related nonlinear optimization problems. *IEEE Trans. Inf. Theory* **59**, 3466–3474 (2013).
83. J. Zavatone-Veth *et al.*, "Neural circuits for fast Poisson compressed sensing in the olfactory bulb" in *Advances in Neural Information Processing Systems*, A. Oh *et al.*, Eds. (Curran Associates, Inc., 2023), vol. 36, pp. 64793–64828.
84. S. Dasgupta, C. F. Stevens, S. Navlakha, A neural algorithm for a fundamental computing problem. *Science* **358**, 793–796 (2017).
85. M. Wachowiak, A. Dewan, T. Bozza, T. F. O'Connell, E. J. Hong, Recalibrating olfactory neuroscience to the range of naturally occurring odor concentrations. *J. Neurosci.* **45**, e1872242024 (2025).
86. J. Murlis, C. D. Jones, Fine-scale structure of odour plumes in relation to insect orientation to distant pheromone and other attractant sources. *Physiol. Entomol.* **6**, 71–86 (1981).
87. K. R. Mylne, P. J. Mason, Concentration fluctuation measurements in a dispersing plume at a range of up to 1000 m. *Q. J. R. Meteorol. Soc.* **117**, 177–206 (1991).
88. N. J. Vickers, Mechanisms of animal navigation in odor plumes. *Biol. Bull.* **198**, 203–212 (2000).
89. G. Reddy, J. D. Zak, M. Vergassola, V. N. Murthy, Antagonism in olfactory receptor neurons and its implications for the perception of odor mixtures. *eLife* **7**, e34958 (2018).
90. V. Singh, N. R. Murphy, V. Balasubramanian, J. D. Mainland, Competitive binding predicts nonlinear responses of olfactory receptors to complex mixtures. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 9598–9603 (2019).
91. M. Elad, Optimized projections for compressed sensing. *IEEE Trans. Signal Process.* **55**, 5695–5702 (2007).
92. P. Ramdya, R. Benton, Evolving olfactory systems on the fly. *Trends in Genet.* **26**, 307–316 (2010).
93. X. Ibarra-Soria *et al.*, Variation in olfactory neuron repertoires is genetically controlled and environmentally modulated. *eLife* **6**, e21476 (2017).
94. S. W. Santoro, C. Dulac, The activity-dependent histone variant h2be modulates the life span of olfactory neurons. *eLife* **1**, e00070 (2012).
95. G. J. Gutierrez, F. Rieke, E. T. Shea-Brown, Nonlinear convergence boosts information coding in circuits with parallel outputs. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e1921882118 (2021).
96. P. Y. Wang, Y. Sun, R. Axel, L. Abbott, G. R. Yang, Evolving the olfactory system with machine learning. *Neuron* **109**, 3879–3892.e5 (2021).
97. P. Cisek, B. Y. Hayden, Neuroscience needs evolution. *Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci.* **377**, 20200518 (2022).
98. N. Kadakia, T. Emonet, Front-end Weber-Fechner gain control enhances the fidelity of combinatorial odor coding. *eLife* **8**, e45293 (2019).
99. H. Giffar, S. Shuvaev, D. Rinberg, A. A. Koulakov, The primacy model and the structure of olfactory space. *PLoS Comput. Biol.* **20**, 1–23 (2024).
100. F. X. P. Bourassa, P. François, G. Reddy, M. Vergassola, Manifold learning for olfactory habituation to strongly fluctuating backgrounds. *bioRxiv* [Preprint] (2025). <https://doi.org/10.1101/2025.05.26.656161> (Accessed 3 September 2025).
101. N. M. Chapochnikov, C. Pehlevan, D. B. Chklovskii, Normative and mechanistic model of an adaptive circuit for efficient encoding and feature extraction. *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2117484120 (2023).
102. S. C. L. Durian, K. Bojanek, O. Marre, S. E. Palmer, Preserving predictive information under biologically plausible compression. *bioRxiv* [Preprint] (2025). <https://doi.org/10.1101/2025.03.12.642864> (Accessed 3 September 2025).
103. J. Bradbury *et al.*, JAX: Composable transformations of Python+NumPy programs (2018). <https://github.com/google/jax/tree/main?tab=readme-ov-file#citing-jax>.