
Neural Circuits for Fast Poisson Compressed Sensing in the Olfactory Bulb

Jacob A. Zavatone-Veth^{1,2,*}, Paul Masset^{1,3,*}, William L. Tong^{1,4,5}, Joseph D. Zak⁶,
Venkatesh N. Murthy^{1,3,†}, Cengiz Pehlevan^{1,4,5,†}

¹Center for Brain Science, ²Department of Physics,

³Department of Molecular and Cellular Biology,

⁴John A. Paulson School of Engineering and Applied Sciences,

⁵Kempner Institute for the Study of Natural and Artificial Intelligence,

Harvard University

Cambridge, MA 02138

⁶Department of Biological Sciences,

University of Illinois at Chicago

Chicago, IL 60607

jzavatoneveth@g.harvard.edu, paul_masset@fas.harvard.edu,

vnmurthy@fas.harvard.edu, cpehlevan@seas.harvard.edu

Abstract

Within a single sniff, the mammalian olfactory system can decode the identity and concentration of odorants wafted on turbulent plumes of air. Yet, it must do so given access only to the noisy, dimensionally-reduced representation of the odor world provided by olfactory receptor neurons. As a result, the olfactory system must solve a compressed sensing problem, relying on the fact that only a handful of the millions of possible odorants are present in a given scene. Inspired by this principle, past works have proposed normative compressed sensing models for olfactory decoding. However, these models have not captured the unique anatomy and physiology of the olfactory bulb, nor have they shown that sensing can be achieved within the 100-millisecond timescale of a single sniff. Here, we propose a rate-based Poisson compressed sensing circuit model for the olfactory bulb. This model maps onto the neuron classes of the olfactory bulb, and recapitulates salient features of their connectivity and physiology. For circuit sizes comparable to the human olfactory bulb, we show that this model can accurately detect tens of odors within the timescale of a single sniff. We also show that this model can perform Bayesian posterior sampling for accurate uncertainty estimation. Fast inference is possible only if the geometry of the neural code is chosen to match receptor properties, yielding a distributed neural code that is not axis-aligned to individual odor identities. Our results illustrate how normative modeling can help us map function onto specific neural circuits to generate new hypotheses.

1 Introduction

Sensory systems allow organisms to detect physical signals in their environments, enabling them to maximize fitness by acting adaptively. This experience of the physical environment, also known as the *Umwelt*, depends on the sensors and sensory organs of each organism [1, 2]. Throughout evolution, organisms have developed specialized sensory mechanisms to extract specific information about the

*JAZ-V and PM contributed equally to this work.

†VNM and CP jointly supervised this work.

physical world. In vision and audition—the most studied sensory modalities in neuroscience—stimuli are characterized by intuitive metrics such as orientation or frequency, which have been shown to map onto neural representations from the earliest stages of the sensory systems [3–5]. One can continuously vary the orientation of an object or the pitch of a tone and quantify resulting changes in perception and neural representations. From a computational point of view, this structure in the representations can be viewed as optimizing the information transfer in the network [6–16].

In contrast, the geometric structure of the olfactory world is far less clear: How can one ‘rotate’ a smell? Despite significant effort, attempts to find such structure in olfactory stimuli and link that geometry to maps in olfactory areas have succeeded only in identifying coarse principles for high-level organization, far from the precision of orientation columns or tonotopy in visual and auditory cortices [17–19]. In the absence of geometric intuitions, the principles of compressed sensing (CS) have emerged as an alternative paradigm for understanding olfactory coding [20–27]. This framework provides a partial answer to the question of how an organism could identify which of millions of possible odorants are present given the activity of only a few hundred receptor types [28–33]. However, existing CS circuit models do not admit convincing biologically-plausible implementations that can perform fast inference at scale. Indeed, many proposals assume that the presence of each odorant is represented by a single, specialized neuron, which is inconsistent with the distributed odor coding observed *in vivo* [22, 34–36]. This axis-aligned coding does not leverage the geometric structure of the sensory space, which can be defined even in the absence of interpretable dimensions [37–41].

In this paper, we propose a Poisson CS model for the mammalian olfactory bulb. Our primary contributions are as follows:

- We derive a normative CS circuit model which can be mapped onto the circuits of the bulb (§3). Importantly, this mapping goes beyond basic counting of cell types; it includes detailed biological features like symmetric coupling and state-dependent inhibition (§4).
- We show that this model enables fast, accurate inference of odor identity in a biologically reasonable regime where tens of odorants are present in a given scene (§5). This fast inference is enabled by considering the geometry of the olfactory receptor code. This consideration leads to distributed odor coding, resolving a major tension between previous CS circuit models and neural data.
- We extend our circuit model to allow Bayesian inference of uncertainty in odor concentrations (§6).

In total, our results demonstrate the importance of considering representational geometry when trying to understand neural coding in the olfactory bulb (OB). Importantly, we show that it is the geometry in the space defined by the receptor affinity (or OSN activation) that controls the speed of inference. This view is distinct from previous geometric theories of olfaction, which have focused on the space of odorants [42, 43]. We propose that thinking in terms of the geometry of OSN coding will allow for deeper understanding of early olfactory processing.

2 Related work and review of the olfactory sensing problem

We begin with a review of the principles of olfactory coding, and of previous CS models for olfactory circuits. The structural logic of early olfactory processing is broadly conserved across the animal kingdom (Fig. 1A) [44–47], and this distinctive circuit structure is thought to play a key role in the computational function of the olfactory bulb [48–50]. In mammals, volatile odorants are first detected by olfactory receptor (ORs) proteins expressed on the surface of olfactory sensory neurons (OSNs). Each OSN expresses only a single OR type; in humans there are around 300 distinct ORs, in mice around 1000 [45]. Importantly, most ORs have broad affinity profiles, and the OSN code for odor identity is combinatorial [51]. In contrast to the immune system’s highly adaptable chemical recognition capabilities arising from somatic recombination [52], ORs are hard coded into the genome as single genes [53], and therefore can only change over evolutionary timescales [54, 55]. Some adaptation of expression levels across the receptors is possible [56], but the chemical affinity of the receptor array is fixed. OSNs expressing the same OR then converge onto the same glomerulus, synapsing onto the principal projection cells of the olfactory bulb (OB), the mitral and tufted cells. These in turn send signals to olfactory cortical areas. The OB contains several types of inhibitory interneurons, whose computational role remains to be clarified [49]. Importantly, the excitatory mitral and tufted cells are not reciprocally connected across glomeruli. Instead, they are connected through a network of inhibitory granule cells, the most numerous cell type in the OB [57].




Figure 1: Circuit architecture of the mammalian olfactory bulb. **A.** Outline of the anatomy of OB circuits. **B.** Fit of the responses of 228 OSN glomerular responses to a panel of 32 odorants. (Gray line, response of single glomeruli, Blue line, Fitted average response.) See Appendix F for details of how these data were collected. **C.** Affinity matrix generated from the data-driven model ensemble. For illustrative purposes, we use only 30 receptors and 100 odorants; in simulations we use 300 receptors to roughly match humans [45], and 1000 or more odorants.

In our model, we will focus on a particular task: odor component identification within a single sniff [58, 59]. This computational problem differs from many experimental tasks focusing on discrimination between two odorants [60], which underlie most scientific work on rodent olfactory decision making [60–62]. Here, the goal is to identify the components of a complex olfactory scene [63–66]. Importantly, the limits of human performance in this setting remain to our knowledge unknown [67, 68]. To render the problem more tractable, we will make a number of simplifications of the anatomy and physiology of the OB. We will not distinguish between mitral and tufted cells, the two classes of projection neurons. Recent works have shed some light on the distinct computational roles of these cell types; these distinctions are likely to become important when considering richer environmental dynamics than we do here [69, 70]. We will also ignore the fact that odorant concentrations can vary over many orders of magnitude, and OSN responses to large changes in concentration are strongly nonlinear. Here, we will focus on concentration changes of one or two orders of magnitude. Within such ranges, the responses of OB neurons are well characterized by linear models [71].

On the theoretical side, it is widely recognized that the tremendous compression of dimensionality inherent in the transformation from odorant molecules to receptor activity means that the olfactory decoding problem is analogous to the one faced in compressed sensing (CS) [21–27, 72–74]. Classical CS theory shows that sparse high-dimensional signals can be recovered from a small number of random projections [28–33]. Inspired by these results, previous works have used the principles of CS to build circuit models for olfactory coding [21–27, 72–74]. However, these model circuits do not map cleanly onto the neural circuits of the OB and the olfactory cortical areas, particularly because they often assume each granule cell encodes exactly one odorant. Moreover, these works usually assume a Gaussian noise model for OSNs, which is biologically unrealistic (but see [21, 22]). Instead, OSN activity is better captured by a Poisson noise model [75, 76]. Some theoretical guarantees for Poisson CS are known, but the situation is less well-understood than in the Gaussian case [32, 33, 77, 78].

3 A neural circuit architecture for Poisson compressed sensing

We now derive a normative, rate-based neural circuit model that performs Poisson CS, which in subsequent sections we will map onto the circuitry of the OB. The design of this model will follow general biological principles, without initially drawing on specific knowledge of the OB. With the goal of sensing within a single sniff in mind, the circuit’s objective is to rapidly infer the odorant concentrations $\mathbf{c} \in \mathbb{R}_+^{n_{\text{odor}}}$ underlying a single, static sample of OSN activity $\mathbf{s} \in \mathbb{R}_+^{n_{\text{OSN}}}$. For simplicity, we model the mean activity of OSNs as a linear function of the concentration, with a receptor affinity matrix $\mathbf{A} \in \mathbb{R}_+^{n_{\text{OSN}} \times n_{\text{odor}}}$ and a baseline rate $\mathbf{r}_0 \in \mathbb{R}_+^{n_{\text{OSN}}}$. As motivated above,

we use a Poisson noise model for OSN activity given the underlying concentration signal $\mathbf{c}$, and correspondingly a Gamma prior over concentrations with shape $\boldsymbol{\alpha} \in \mathbb{R}_+^{n_{\text{odor}}}$ and scale $\boldsymbol{\lambda} \in \mathbb{R}_+^{n_{\text{odor}}}$.³

$$\mathbf{s} \mid \mathbf{c} \sim \text{Poisson}(\mathbf{r}_0 + \mathbf{A}\mathbf{c}), \quad \mathbf{c} \sim \text{Gamma}(\boldsymbol{\alpha}, \boldsymbol{\lambda}). \quad (1)$$

Given this likelihood and prior, we construct a neural circuit to compute the maximum *a posteriori* (MAP) estimate of the concentration $\mathbf{c}$ using gradient ascent on the log-posterior probability. Here, we sketch the derivation, deferring some details to Appendix B. Our starting point is the gradient ascent equation

$$\dot{\mathbf{c}}(t) = \nabla_{\mathbf{c}} \log p(\mathbf{c} \mid \mathbf{s}) = \mathbf{A}^\top [\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\mathbf{c}) - \mathbf{1}] + (\boldsymbol{\alpha} - \mathbf{1}) \oslash \mathbf{c} - \boldsymbol{\lambda}, \quad (2)$$

where $\oslash$ denotes elementwise division and $\dot{\mathbf{c}}(t) = d\mathbf{c}/dt$. Here, the estimate $\mathbf{c}$ is formally constrained to $\mathbb{R}_+^{n_{\text{odor}}}$; in numerical simulations we will sometimes ignore this constraint. Circuit algorithms of this form were studied in previous work by Grabska-Barwińska et al. [21].

However, in this most basic setup there is a one-to-one mapping between neurons and odorants, which is at variance with our knowledge of biological olfaction (§2). To distribute the code, we instead use a projected setup where the firing rates $\mathbf{g} \in \mathbb{R}^{n_g}$ of the neurons are mapped to the concentration estimate $\mathbf{c}$ through a matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{n_{\text{odor}} \times n_g}$: $\mathbf{c}(t) = \boldsymbol{\Gamma}\mathbf{g}(t)$. Even if $\boldsymbol{\Gamma}$ is non-square—in particular, if $n_g > n_{\text{odor}}$ —so long as $\boldsymbol{\Gamma}\boldsymbol{\Gamma}^\top$ is positive-definite and the rates follow the dynamics $\tau_g \dot{\mathbf{g}}(t) = (\mathbf{A}\boldsymbol{\Gamma})^\top [\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g}) - \mathbf{1}] + \boldsymbol{\Gamma}^\top [(\boldsymbol{\alpha} - \mathbf{1}) \oslash (\boldsymbol{\Gamma}\mathbf{g}) - \boldsymbol{\lambda}]$ for some time constant τ_g , the concentration estimates will still converge to the MAP. This corresponds to preconditioned gradient ascent [39, 79]. Again, we formally require the constraint that $\mathbf{c} \in \mathbb{R}_+^{n_{\text{odor}}}$, which translates into a constraint on $\mathbf{g}$. Most simply, we may take $\mathbf{g} \in \mathbb{R}_+^{n_g}$ and choose $\boldsymbol{\Gamma}$ to be positivity-preserving.

These dynamics include two divisive non-linearities, which can be challenging to implement in biophysical models of single neurons [80]. Using the approach proposed by Chalk et al. [81], we can linearize the inference by introducing two additional cell types that have as their fixed points the elementwise divisions $\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g})$ and $(\boldsymbol{\alpha} - \mathbf{1}) \oslash (\boldsymbol{\Gamma}\mathbf{g})$. Concretely, we introduce cell types with rates $\mathbf{p} \in \mathbb{R}^{n_{\text{OSN}}}$ and $\mathbf{z} \in \mathbb{R}^{n_{\text{odor}}}$ such that their fixed-point rates for fixed $\mathbf{g}$ are $\mathbf{p}^* = \mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g})$ and $\mathbf{z}^* = (\boldsymbol{\alpha} - \mathbf{1}) \oslash (\boldsymbol{\Gamma}\mathbf{g})$, respectively. This yields the coupled circuit dynamics

$$\begin{aligned} \mathbf{c}(t) &= \boldsymbol{\Gamma}\mathbf{g}(t), & \tau_g \dot{\mathbf{g}}(t) &= (\mathbf{A}\boldsymbol{\Gamma})^\top (\mathbf{p} - \mathbf{1}) + \boldsymbol{\Gamma}^\top (\mathbf{z} - \boldsymbol{\lambda}), \\ \tau_p \dot{\mathbf{p}}(t) &= \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g}), & \tau_z \dot{\mathbf{z}}(t) &= \boldsymbol{\alpha} - \mathbf{1} - \mathbf{z} \odot \mathbf{c}, \end{aligned} \quad (3)$$

for cell-type-specific time constants τ_g , τ_p , and τ_z , where $\odot$ denotes elementwise multiplication. In the limit $\tau_p, \tau_z \ll \tau_g$, this circuit will recover the MAP gradient ascent. If τ_p and τ_z are not infinitely fast relative to τ_g , we expect these dynamics to approximate the desired dynamics [81] (see Appendix C for a preliminary analysis of the linear stability of the MAP fixed-point). We will test the accuracy of this approximation for biologically-reasonable time constants using numerical experiments. Moreover, $\mathbf{p}$ should formally be constrained to $\mathbb{R}_+^{n_p}$, such that the non-negativity of the target ratio is respected. Finally, we note that in the special case $\boldsymbol{\alpha} = \mathbf{1}$ in which the Gamma prior reduces to an exponential prior, the introduction of the cell type $\mathbf{z}$ is no longer required.

4 Biological interpretation and predictions of the circuit model

We now argue that the normative model derived in the preceding section can be mapped onto the circuitry of the OB. In particular, though the model was derived based only on general biological principles, its specific features are biologically implementable based on the detailed anatomy and physiology of the OB. In terms of the levels of understanding of neural circuits proposed by David Marr, this is an example of how normative modeling can bridge the gap between algorithmic and mechanistic understanding [82, 83].

As foreshadowed by our notation, we interpret the cell type $\mathbf{g}$ as the granule cells of the OB, and the cell type $\mathbf{p}$ as the mitral cells, which are projection neurons. Provided that the elements of the matrix $\mathbf{A}\boldsymbol{\Gamma}$ are non-negative, this interpretation is justified at the coarsest level by the signs with which the two cell types appear in the dynamics: the $\mathbf{p}$ neurons excite the $\mathbf{g}$ neurons, which in turn inhibit the $\mathbf{p}$ neurons. Finally, we interpret the cell type $\mathbf{z}$ as representing a form of cortical feedback. In the remainder of this section, we will justify this mapping in detail.

³See Appendix A for a detailed description of our notational conventions.




Figure 2: State-dependent inhibition of mitral cells. **A.** Upon stimulation, mitral cells exhibit a transient burst of activity followed by relaxation to a plateau. Darker color indicates stronger stimulation. **B.** In the circuit for Poisson CS we propose (blue line), the inhibition due to the activation of a second mitral cell (MC_B) is gated by the activity of the cell we are recording from (MC_A), as observed experimentally by Arevian et al. [95] in their Fig. 2D. This state-dependent gating does not occur in a circuit derived from a Gaussian noise model (orange line). Rates are normalized to the maximal stimulation of the principal cell. Dashed line indicates the unity line. **C.** Strength of inhibition as a function of time and of stimulation intensity; c.f. [95] Fig. 3B.

4.1 Circuit anatomy: weight transport, dendro-dendritic coupling, and cortical feedback

The first salient feature of the dynamics (3) is that the cell types g and p do not make direct lateral connections amongst themselves. Rather, they connect only indirectly through the neurons of the opposing cell type, matching the connectivity structure of mitral and granule cells (Fig. 1A). Moreover, the synaptic weights $\mathbf{A}\mathbf{\Gamma}$ of the connections from g to p neurons mirror exactly the weights $(\mathbf{A}\mathbf{\Gamma})^\top$ of the connections from p to g neurons. Naïvely, this creates a weight transport problem of the form that renders backpropagation biologically implausible: how should the exact transpose of a matrix be copied to another synapse [84–86]? However, in the unique case of the OB this does not pose a substantial obstacle, as the mitral and granule cells are coupled by dendro-dendritic synapses, meaning that bi-directional connectivity occurs at a single physical locus (Fig. 1) [49, 57, 87].

This interpretation accounts for the interactions between the cell types g and p in the dynamics (3), but what about the cell type z ? These cells give direct excitatory input to the granule cells g with weights $\mathbf{\Gamma}^\top$, and represent the concentration-dependent contribution to the log-prior gradient. We can therefore interpret these cells as representing feedback from olfactory cortical areas to the OB, which arrives at the granule cells [88–90]. Though our model can thus flexibly incorporate cortical feedback, for our subsequent simulations we will focus on the simple case in which the prior is static and has $\alpha = 1$, in which case it reduces to an exponential and feedback is not explicitly required.

Given this mapping of the cell types of our model to the cell types of the bulb, we will henceforth choose the membrane time constants to match experiment, taking $\tau_p = 20$ ms [91] and $\tau_g = 30$ ms [92]. This matching of timescales is required for our comparison of model inference timescales to the timescale of a single sniff to be meaningful.

4.2 Choosing the affinity matrix

In theories of Poisson CS, the optimal sensing matrix $\mathbf{A}$ is one that has columns that are in some precise sense as orthogonal as possible, so that it acts as an approximate isometry [27–33]. However, biologically, the olfactory system is not free to choose optimal sensors. Rather, the affinity profiles of each receptor are dictated by biophysics and by evolutionary history [53, 54, 93]. To build a more realistic model for OSN sensing, we therefore turn to biological data. Using two-photon calcium imaging, we recorded the responses of 228 mouse OSNs glomeruli to 32 odorants (see Appendix F for details; these data were previously published in [94]). In Fig. 1B, we show that the distribution of responses is well-fit by a Gamma distribution with shape 0.37 and scale 0.36 (Appendix F). We then define our ensemble of sensing matrices $\mathbf{A}$ by drawing their elements as independently and identically distributed Gamma(0.37, 0.36) random variables. An example matrix drawn from this ensemble is shown in Fig. 1C.

4.3 Divisive predictive coding by mitral cells

In our model (3), odorant concentration estimates are decoded from granule cell activity. This feature is shared with previous Gaussian CS models of olfactory coding [20, 24, 25]. Our model, however, matches biology more closely than these previous works because it allows for a distributed code rather than assuming that each granule cell codes for a single odorant.

What, then, is the functional role of the mitral cells in our model? We can interpret their dynamics as implementing a form of predictive coding, in which they are trying to cancel their input by the current prediction. Because the mitral cell activity converges to the ratio of their input to the prediction, this is a divisive form of predictive coding [81]. In contrast, a Gaussian noise model gives a subtractive form of predictive coding in which the activity converges to the difference between input and prediction (Appendix D) [15]. In Fig. 2A, we show example timecourses of model mitral cell activity following the onset of a stimulus. Consistent with experimentally-measured responses [96, 97], a sharp transient response at the onset of stimulation is followed by decay to a low level of tonic activity (Fig. 2A).

4.4 State-dependent inhibition of mitral cells

A salient feature of our circuit model is that the inhibition from the granule cells onto a mitral cell is gated by the activity of the mitral cell itself (3). This state-dependent inhibition is reminiscent of *in vitro* experiments showing that granule cell mediated inhibition is activity dependent [95]. In these experiments, Arevian et al. measured the activity of a primary mitral cell (MC_A) while increasing its level of stimulation under two conditions. In the first condition, no other cells in the circuit are being stimulated. In the second condition, they also activate another mitral cell (MC_B). The activation of the second mitral cell leads to the activation of granule cells connected to both mitral cells and a reduction in the firing evoked by stimulation of the primary cell alone. Strikingly, these authors showed that this inhibition is dependent on the activity of the primary cell [95].

Here, we show that our proposed circuit reproduces these observations of state-dependent inhibition, and that they do not arise in a similarly-constructed circuit for a Gaussian noise model. To model Arevian et al. [95]’s *in vitro* experiments, we simulated a reduced circuit with 2 mitral cells and 10 granule cells. We stimulated MC_A with an input $s_A \in [1 : 400]$ while toggling on or off the stimulation $s_B = 80$ of the second mitral cell MC_B . As observed experimentally, when MC_B is activated, the inhibition of the primary mitral cell MC_A is state dependent (Fig. 2B). To show that this effect arises from our Poisson circuit, we build a similar inference circuit with a Gaussian noise model (Appendix D). Under similar conditions, the inhibition in the Gaussian circuit is independent of the activity of the primary mitral cell MC_A (Fig. 2B). Furthermore, the dynamics of the relative inhibition qualitatively recapitulate those observed experimentally, with sustained inhibition throughout the stimulation period at the stimulation level of MC_A leading to maximal inhibition and inhibition followed by relaxation for stronger stimulation levels of MC_A (Fig. 2C and see [95] Fig. 3b).

5 Geometry, speed, and capacity

We have argued that the model introduced in §3 could be implemented biologically, but can it perform at scale? Concretely, can a circuit of this architecture with neuron counts comparable to the human OB correctly identify which among a large set of odorants are present in a given scene? This is precisely the question of the capacity of the CS algorithm [98, 99]. In Fig. 3A, we present our algorithm with scenes composed of varying numbers of randomly-selected odorants out of a panel of 1000, which for simplicity we take to be at the same concentration. We first ask how many odorants can be reliably detected within a single sniff, i.e., 200 ms. To convert MAP concentration estimates into presence estimates, we simply binarize the estimated concentrations based on whether they are larger than half of the true odorant concentration. The ability of the one-to-one code to successfully detect odorants falls off rapidly, with the detection fraction falling below one-half even if only a handful of odorants are present (Fig. 3A).⁴

The limited capacity of the one-to-one code can be overcome by distributing the code in a way that takes into account the geometry of the sensing problem. Here, the information geometry of the

⁴We remark that, with a one-to-one code, our model is identical to that proposed by Grabska-Barwińska et al. [21] except for the introduction of the granule cells.




Figure 3: Fast detection of many odorants. **A**. Fraction of odorants correctly detected within 100 ms (*left*), 200 ms (*center*), and 1 s (*right*) after odorant onset as a function of the number of odorants present, for models with **one-to-one**, **naïvely distributed**, and **geometry-aware** codes. Here, we consider a repertoire of 1000 possible odorants. **B**. Heatmap with overlaid smoothed contours of correct detection fraction as a function of number of present odorants and time window for models with **one-to-one** (*left*), **naïvely distributed** (*center*), and **geometry-aware** (*right*) codes. As in **A**, a repertoire of 1000 possible odorants is used. **C**. Threshold number of odorants for which half can be correctly detected as a function of total repertoire size within 100 ms (*left*), 200 ms (*center*), and 1 s (*right*) after odorant onset. See also Supp. Figs. **G.1** and **G.2** for versions of panels **A** and **B** with varying repertoire sizes, from which the capacities shown here are derived. See Appendix **G** for details of our numerical methods. Shaded patches show ± 1.96 SEM over realizations throughout.

problem is governed by the sensing matrix $\mathbf{A}$, which introduces correlations in the input signals to mitral cells because the off-diagonal components of $\mathbf{A}^\top \mathbf{A}$ are non-negligible. To counteract these detrimental correlations, we can choose a readout matrix $\mathbf{\Gamma}$ such that $\mathbf{\Gamma} \mathbf{\Gamma}^\top \approx (\mathbf{A}^\top \mathbf{A})^+$ up to constants of proportionality, where we must take a pseudoinverse because $\mathbf{A}^\top \mathbf{A}$ is highly rank-deficient (see Appendix **G** for details). Geometrically, this corresponds an approximation of natural gradient descent, in which we use the Gauss-Newton matrix instead of the Fisher information matrix because the latter is state-dependent for Poisson likelihoods [37–40]. As a control, we also consider a naïvely distributed code with $\mathbf{\Gamma} \mathbf{\Gamma}^\top \approx \mathbf{I}_{n_{\text{odor}}}$. Distributing the code, even without accounting for the geometry of inference, markedly improves the single-sniff capacity to around 10-20, and taking into account the geometry produces a further improvement (Fig. 3A).

To gain a more granular view of how tuned geometry enables faster inference, in Fig. 3B we test the three models’ detection capabilities at sub-sniff resolution. We can see that at long times—around a second after odorant onset—the naïvely distributed and geometry-aware codes achieve similar capacities of around 50-60 odorants. However, the geometry-aware code reaches this detection capacity within a single sniff, whereas the naïvely distributed code requires the full second of

processing time. The detection capacity of the one-to-one code reaches only around 20 odorants after 1 second, and even then does not appear to have reached its asymptote. These results illustrate two important conceptual points: First, when the strength of individual synapses is bounded, distributed coding can speed up dynamics by increasing the effective total input to a given neuron. Second, given a sensing matrix that induces strong correlations, geometry-aware distributed coding can accelerate inference by counteracting that detrimental coupling.

These tests show that our model can reliably detect tens of odorants from a panel of thousands of possible odorants, but do not probe how the model’s capacity scales to larger odor spaces. While the true dimensionality of odor space remains unknown [100, 101], there may be orders of magnitude more than thousands of possible odorants. As an upper bound, there are on the order of millions of known volatile compounds that are plausibly odorous [102]. Thus, it is important to determine how our model scales to more realistically-sized odor spaces. From the literature on compressed sensing performance bounds, we expect the threshold sparsity to decay slowly—roughly logarithmically—with increasing n_{odor} [28–33]. As a first step, in Supp. Figs. G.1 and G.2, we reproduce Fig. 3A–B for between 500 and 8000 possible odorants, showing that performance does indeed drop off slowly with increasing odor space dimension. To get a more precise estimate of how performance scales with repertoire size, in Fig. 3C we plot the threshold number of odorants for which half can be reliably detected as a function of the repertoire size, showing that for the geometry-aware code it decays only a bit faster than logarithmically. Our ability to simulate larger systems was limited by computational resources. This limitation is present in previous works, and the repertoires tested here are comparable to—or substantially larger than—those used in past studies [21, 22, 24, 26, 72].

6 Fast sampling for uncertainty estimation

Thus far, we have focused on a circuit that performs MAP estimation of odorant concentrations. However, to successfully navigate a dynamic, noisy world, animals must estimate sensory uncertainty at the timescale of perception [38, 103–107]. Fortunately, our circuit model can be easily extended to perform Langevin sampling of the full posterior distribution, allowing for uncertainty estimation while maintaining its attractive structural features. In Appendix B, we provide a detailed derivation of a model that implements Langevin sampling through the granule cells for a Poisson likelihood and Gamma prior. This yields a circuit that is identical to the MAP estimation circuit introduced in §3 up to the addition of Gaussian noise to the granule cell dynamics:

$$\begin{aligned} \mathbf{c}(t) &= \mathbf{\Gamma} \mathbf{g}(t) & \tau_g \dot{\mathbf{g}}(t) &= (\mathbf{A} \mathbf{\Gamma})^\top (\mathbf{p} - \mathbf{1}) + \mathbf{\Gamma}^\top (\mathbf{z} - \boldsymbol{\lambda}) + \boldsymbol{\xi}(t) \\ \tau_p \dot{\mathbf{p}}(t) &= \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A} \mathbf{\Gamma} \mathbf{g}) & \tau_z \dot{\mathbf{z}}(t) &= \boldsymbol{\alpha} - \mathbf{1} - \mathbf{z} \odot \mathbf{c}. \end{aligned} \quad (4)$$

Here, $\boldsymbol{\xi}(t)$ is a vector of n_g independent zero-mean Gaussian noise processes with covariance $\mathbb{E}[\xi_j(t) \xi_{j'}(t')] = 2\tau_g \delta_{jj'} \delta(t - t')$, and once again the rates $\mathbf{g}$ and $\mathbf{p}$ should in principle be constrained to be non-negative. In this case, the readout matrix $\mathbf{\Gamma}$ both preconditions the effective gradient force and shapes the structure of the effective noise $\mathbf{\Gamma} \boldsymbol{\xi}$, allowing us to mold the geometry of the sampling manifold [37, 40]. By using a projected readout, we maintain the independence of the noise processes for different neurons. This both allows us to have many independent samplers—if $n_g > n_{\text{odor}}$ —and is important for biological realism if we interpret the sampling noise as resulting from fluctuations in membrane potential due to synaptic noise [108].

As a test of how this sampling circuit performs, we consider a simple setup in which one set of odorants appears at a low concentration, and then a second set of odorants appear at a higher concentration while the low odorants are still present. This setup tests both the circuit’s ability to converge rapidly enough to give accurate posterior samples within a single sniff, and its ability to correctly infer odorant concentrations even in the presence of distractors [63, 66]. In Fig. 4, we show that circuits with one-to-one or naïvely distributed codes do not give accurate estimates of the concentration mean within 200 ms, while tuning the geometry to match the receptor affinities enables fast convergence. All three circuits overestimate the posterior variance at short times, consistent with what one would expect for unadjusted Langevin samplers [109], but the geometry-aware model’s estimate decays most rapidly towards the target. Therefore, when the synaptic weights are tuned, our circuit model can enable fast, robust estimation of concentration statistics, even in the presence of distractors.




Figure 4: Fast uncertainty estimation using Langevin sampling of the posterior. Here, we use a simple concentration estimation task in which 5 randomly-selected ‘low’ odorants out of a panel of 1000 appear at concentration 10 at time 0 s, and then a further 5 randomly-selected ‘high’ odorants appear at concentration 40 at time 1 s. **A**. Smoothed timeseries of instantaneous concentration estimates for **low**, **high**, and **background** odorants, for models with one-to-one (*top*), naively distributed (*middle*), and geometry-aware (*bottom*) codes. Background odorant estimates are shown as mean $\pm$ standard deviation over odorants. Dashed lines show true concentrations over time. **B**. Cumulative estimates of concentration mean for low (*top*) and high (*bottom*) odorants after the onset of the low odorants for **one-to-one**, **naively distributed**, and **geometry-aware** codes. Black lines indicate baseline estimates of the posterior mean. Thick colored lines indicate means over odorants. **C**. As in **B**, but for the estimated variance. **D**. As in **B**, but after the onset of high odorants. **E**. As in **C**, but after the onset of high odorants. See Appendix G for details of our numerical methods, and for individual-odor traces.

7 Discussion

In this paper, we have derived a novel, minimal rate model for fast Bayesian inference in early olfaction. Unlike previously-proposed algorithms for CS in olfaction, this model has a clear mapping onto the circuits of the mammalian OB. We showed that this model successfully performs odorant identification across a biologically-relevant range of scene sparsities and circuit sizes. This model therefore exemplifies how normative approaches can blur the lines between algorithmic and mechanistic understanding of neural circuits [82, 83]. We now conclude by discussing possible avenues for future inquiry, as well as some of the limitations of our work.

One limitation of our simulations is that we have chosen to distribute the neural code randomly, and have allowed for negative entries in the mitral-granule synaptic weight matrix $\mathbf{A}\mathbf{\Gamma}$ (see methods in Appendix G). These features are not entirely biologically satisfactory. However, our model only captures the mitral cells and granule cells, overlooking a number of other inhibitory cell types that could contribute to solving this problem, e.g., through feedforward inhibition onto mitral cells or lateral inhibition across granule cells [48, 49]. Fundamentally, negative values in the geometry-aware decoding matrix $\mathbf{\Gamma}$ arise as although the affinity matrix is positive, its inverse will contain negative elements. A biologically-plausible realization of the geometry-aware code through the introduction of additional cell types could be achieved by decomposing the inverse into several components, yielding sparse, consistently-signed connectivity [110]. As a first step, in Supp. Fig. G.4 we show that similar performance to Fig. 4 can be achieved using a sparse non-negative randomly distributed code. As a result, one objective for future work will be to develop better models for the decoding matrix $\mathbf{\Gamma}$ that result in more realistic connectivity.

Though our model captures two of the interesting features of the anatomy and physiology of the OB—symmetric dendrodendritic coupling and state-dependent inhibition of mitral cells—there are

many biological details which we have not addressed. First, our linear model for OSN mean firing neglects receptor antagonism, gain control, and other nonlinear effects [94, 111, 112], which are known to affect the performance of Gaussian CS models for olfaction [26, 113]. Second, our models are rate-based, while neurons in the OB spike. In spiking implementations of sampling networks, the noise is not uncorrelated across neurons, complicating their biological interpretation [38, 114]. Constructing models that capture these richly nonlinear effects will be an important objective for future work. A first step towards such a nonlinear model would be to build a spiking network that approximates the rate-based models considered here, which could be accomplished using the efficient balanced network formalism for distributed spiking networks [38]. Another step would be to add a Hill function nonlinearity to the OSN model to approximate competitive binding, as studied for Gaussian compressed sensing by Qin et al. [26]. One challenge in constructing models that incorporate additional nonlinearity is that the simple linear strategy for distributing the code used here may no longer be directly applicable. In Appendix E, we illustrate this obstacle for the relatively simple case of a model with the same linear OSNs and Poisson likelihood but an L_0 instead of Gamma prior, building on recent work on circuits for Gaussian CS with L_0 priors [115].

A closely related point is that we model the weights of the synapses between mitral and granule cells as fixed, and do not consider synaptic plasticity. In particular, we assume that they are tuned to the statistics of the receptor affinities without specifying a mechanism by which this tuning could take place. In biology, receptor abundances and other OSN properties display activity-dependent adaptation over long timescales, meaning that the optimal tuning is unlikely to be static [56, 116]. Some past works have sought to incorporate plasticity of the mitral-granule cell synapses into decoding models [24, 117], tying into a larger body of research on how plasticity can enable flexible feature extraction in olfaction [118, 119]. This learning should lead to measurable changes in the population response to odorant panels, which our model predicts should be linked in a precise way to account for receptor-induced correlations in the responses to the most frequently present odorants. Experimentally characterizing and carefully modeling these changes in responses across timescales will be an interesting avenue for future work [120]. Experimental techniques to probe these ideas at the neural and behavioral level have recently been proposed [121, 122] which allow more precise control of stimulus and subjective geometry.

Though the circuit model derived in §3 incorporates a general Gamma prior represented by cortical feedback, our simulations focus on the special case in which the prior reduces to an exponential, in which the feedback neurons are not needed. Future work will therefore be required to carefully probe the effect of incorporating a Gamma prior with non-unit shape and to dissect the structure of the resulting modeled cortical feedback. More generally, it will be interesting to extend our framework to incorporate data-adaptive priors. Importantly, the stimuli used in this paper constitute an extremely impoverished model for the richness of the true odor world; we do not account for its rich dynamical structure and co-occurrence statistics. Adaptive priors as encoded by cortical feedback would allow circuits to leverage this structure, enabling faster and more accurate inference [88–90, 123, 124].

We conclude by noting that our work provides an example of how distributed coding can lead to faster inference than axis-aligned disentangled coding. In recent years, the question of when axis-aligned coding is optimal has attracted significant attention in machine learning and neuroscience [125–135]. Much of this work focuses on the question of when axis-aligned codes are optimal for energy efficiency or for generalization, whereas here we focus on the question of which code yields the fastest inference dynamics. These ideas are one example of the broader question of how agents and algorithms should leverage the rich geometry of the natural world to enable fast, robust learning and inference [136]. We believe that investigating how task demands and biological constraints affect the optimal representational geometry for that task is a promising avenue for illuminating neural information processing in brains and machines [41, 136, 137].

Acknowledgments and Disclosure of Funding

We thank Naoki Hiratani, Shanshan Qin, and Vikrant Kapoor for useful discussions. This work was supported by NSF grants DMS-2134157 and CAREER IIS-2239780 to CP, NIH grants R01DC017311 and R01DC016289 to VNM, and NTT Research award A47994 to VNM. CP received additional support from a Sloan Research Fellowship. PM was partially supported by a grant from the Harvard Mind Brain Behavior Interfaculty Initiative. JDZ was partially supported by NIH grant K99DC017754. 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. A subset of the computations in this paper were run on the FASRC cluster supported by the FAS Division of Science Research Computing Group at Harvard University.

References

- [1] Jakob Johann von Uexküll. *Umwelt und Innenwelt der Tiere*. Springer, 1909. doi:[10.1007/978-3-662-24819-5](https://doi.org/10.1007/978-3-662-24819-5).
- [2] Ed Yong. *An Immense World: How Animal Senses Reveal the Hidden Realms Around Us*. Random House, New York, first edition, 2022. ISBN 9780593133231.
- [3] D. H. Hubel and T. N. Wiesel. Receptive fields of single neurones in the cat’s striate cortex. *The Journal of Physiology*, 148(3):574–591, 1959. doi:<https://doi.org/10.1113/jphysiol.1959.sp006308>. URL <https://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1959.sp006308>.
- [4] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of Physiology*, 160(1):106–154, 1962. doi:<https://doi.org/10.1113/jphysiol.1962.sp006837>. URL <https://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1962.sp006837>.
- [5] Karl Kandler, Amanda Clause, and Jihyun Noh. Tonotopic reorganization of developing auditory brainstem circuits. *Nature neuroscience*, 12(6):711–717, 2009. doi:<https://doi.org/10.1038/nn.2332>.
- [6] Fred Attneave. Some informational aspects of visual perception. *Psychological Review*, 61(3):183, 1954. doi:<https://doi.org/10.1037/h0054663>.
- [7] Horace B. Barlow. Possible Principles Underlying the Transformations of Sensory Messages. In *Sensory Communication*. The MIT Press, 09 2012. ISBN 9780262518420. doi:[10.7551/mitpress/9780262518420.003.0013](https://doi.org/10.7551/mitpress/9780262518420.003.0013). URL <https://doi.org/10.7551/mitpress/9780262518420.003.0013>.
- [8] Mandyam Veerambudi Srinivasan, Simon Barry Laughlin, A. Dubs, and George Adrian Horridge. Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 216(1205):427–459, 1982. doi:[10.1098/rspb.1982.0085](https://doi.org/10.1098/rspb.1982.0085).
- [9] R. Linsker. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988. doi:[10.1109/2.36](https://doi.org/10.1109/2.36).
- [10] Johannes H. van Hateren. A theory of maximizing sensory information. *Biological Cybernetics*, 68(1):23–29, Nov 1992. ISSN 1432-0770. doi:[10.1007/BF00203134](https://doi.org/10.1007/BF00203134). URL <https://doi.org/10.1007/BF00203134>.
- [11] Joseph J. Atick and A. Norman Redlich. What Does the Retina Know about Natural Scenes? *Neural Computation*, 4(2):196–210, 03 1992. ISSN 0899-7667. doi:[10.1162/neco.1992.4.2.196](https://doi.org/10.1162/neco.1992.4.2.196). URL <https://doi.org/10.1162/neco.1992.4.2.196>.
- [12] Bruno A. Olshausen and David J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, Jun 1996. ISSN 1476-4687. doi:[10.1038/381607a0](https://doi.org/10.1038/381607a0). URL <https://doi.org/10.1038/381607a0>.

- [13] Anthony J. Bell and Terrence J. Sejnowski. The “independent components” of natural scenes are edge filters. *Vision Research*, 37(23):3327–3338, 1997. ISSN 0042-6989. doi:[https://doi.org/10.1016/S0042-6989\(97\)00121-1](https://doi.org/10.1016/S0042-6989(97)00121-1). URL <https://www.sciencedirect.com/science/article/pii/S0042698997001211>.
- [14] Nicolas Brunel and Jean-Pierre Nadal. Mutual Information, Fisher Information, and Population Coding. *Neural Computation*, 10(7):1731–1757, 10 1998. ISSN 0899-7667. doi:[10.1162/089976698300017115](https://doi.org/10.1162/089976698300017115).
- [15] Rajesh P. N. Rao and Dana H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, Jan 1999. ISSN 1546-1726. doi:[10.1038/4580](https://doi.org/10.1038/4580). URL <https://doi.org/10.1038/4580>.
- [16] Daniel A Butts and Mark S Goldman. Tuning curves, neuronal variability, and sensory coding. *PLoS Biology*, 4(4), 03 2006. doi:[10.1371/journal.pbio.0040092](https://doi.org/10.1371/journal.pbio.0040092). URL <https://doi.org/10.1371/journal.pbio.0040092>.
- [17] Venkatesh N. Murthy. Olfactory maps in the brain. *Annual Review of Neuroscience*, 34(1): 233–258, 2011. doi:[10.1146/annurev-neuro-061010-113738](https://doi.org/10.1146/annurev-neuro-061010-113738). URL <https://doi.org/10.1146/annurev-neuro-061010-113738>.
- [18] Brian K. Lee, Emily J. Mayhew, Benjamin Sanchez-Lengeling, Jennifer N. Wei, Wesley W. Qian, Kelsie Little, Matthew Andres, Britney B. Nguyen, Theresa Moloy, Jane K. Parker, Richard C. Gerkin, Joel D. Mainland, and Alexander B. Wiltschko. A principal odor map unifies diverse tasks in human olfactory perception. *bioRxiv*, 2022. doi:[10.1101/2022.09.01.504602](https://doi.org/10.1101/2022.09.01.504602). URL <https://www.biorxiv.org/content/early/2022/12/13/2022.09.01.504602>.
- [19] Wesley W Qian, Jennifer N Wei, Benjamin Sanchez-Lengeling, Brian K Lee, Yunan Luo, Marnix Vlot, Koen Dechering, Jian Peng, Richard C Gerkin, and Alexander B Wiltschko. Metabolic activity organizes olfactory representations. *eLife*, 12:e82502, may 2023. ISSN 2050-084X. doi:[10.7554/eLife.82502](https://doi.org/10.7554/eLife.82502). URL <https://doi.org/10.7554/eLife.82502>.
- [20] Alexei A. Koulakov and Dmitry Rinberg. Sparse incomplete representations: A potential role of olfactory granule cells. *Neuron*, 72(1):124–136, 2011. ISSN 0896-6273. doi:<https://doi.org/10.1016/j.neuron.2011.07.031>. URL <https://www.sciencedirect.com/science/article/pii/S0896627311006891>.
- [21] Agnieszka Grabska-Barwińska, Jeff Beck, Alexandre Pouget, and Peter Latham. Demixing odors - fast inference in olfaction. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/2bcab9d935d219641434683dd9d18a03-Paper.pdf.
- [22] Agnieszka Grabska-Barwińska, Simon Barthelmé, Jeff Beck, Zachary F. Mainen, Alexandre Pouget, and Peter E. Latham. A probabilistic approach to demixing odors. *Nature Neuroscience*, 20(1):98–106, Jan 2017. ISSN 1546-1726. doi:[10.1038/nn.4444](https://doi.org/10.1038/nn.4444). URL <https://doi.org/10.1038/nn.4444>.
- [23] Sina Tootoonian and Mate Lengyel. A dual algorithm for olfactory computation in the locust brain. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL https://proceedings.neurips.cc/paper_files/paper/2014/file/c8ba76c279269b1c6bc8a07e38e78fa4-Paper.pdf.
- [24] Daniel Kepple, Brittany N. Cazakoff, Heike S. Demmer, Dennis Eckmeier, Stephen D. Shea, and Alexei A. Koulakov. Computational algorithms and neural circuitry for compressed sensing in the mammalian main olfactory bulb. *bioRxiv*, 2018. doi:[10.1101/339689](https://doi.org/10.1101/339689). URL <https://www.biorxiv.org/content/early/2018/06/05/339689>.
- [25] Daniel R Kepple, Hamza Giaffar, Dmitry Rinberg, and Alexei A Koulakov. Deconstructing odorant identity via primacy in dual networks. *Neural Computation*, 31(4):710–737, 2019. doi:https://doi.org/10.1162/neco_a_01175.

- [26] Shanshan Qin, Qianyi Li, Chao Tang, and Yuhai Tu. Optimal compressed sensing strategies for an array of nonlinear olfactory receptor neurons with and without spontaneous activity. *Proceedings of the National Academy of Sciences*, 116(41):20286–20295, 2019. doi:[10.1073/pnas.1906571116](https://doi.org/10.1073/pnas.1906571116). URL <https://www.pnas.org/doi/abs/10.1073/pnas.1906571116>.
- [27] Kamesh Krishnamurthy, Ann M Hermundstad, Thierry Mora, Aleksandra M Walczak, and Vijay Balasubramanian. Disorder and the neural representation of complex odors. *Frontiers in Computational Neuroscience*, 16, 2022. doi:[10.3389/fncom.2022.917786](https://doi.org/10.3389/fncom.2022.917786).
- [28] Richard G. Baraniuk. Compressive sensing [lecture notes]. *IEEE Signal Processing Magazine*, 24(4):118–121, 2007. doi:[10.1109/MSP.2007.4286571](https://doi.org/10.1109/MSP.2007.4286571).
- [29] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006. doi:[10.1109/TIT.2006.871582](https://doi.org/10.1109/TIT.2006.871582).
- [30] Emmanuel J. Candès and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, 2006. doi:[10.1109/TIT.2006.885507](https://doi.org/10.1109/TIT.2006.885507).
- [31] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006. doi:<https://doi.org/10.1002/cpa.20124>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.20124>.
- [32] Maxim Raginsky, Rebecca M. Willett, Zachary T. Harmany, and Roummel F. Marcia. Compressed sensing performance bounds under Poisson noise. *IEEE Transactions on Signal Processing*, 58(8):3990–4002, 2010. doi:[10.1109/TSP.2010.2049997](https://doi.org/10.1109/TSP.2010.2049997).
- [33] Maxim Raginsky, Sina Jafarpour, Zachary T. Harmany, Roummel F. Marcia, Rebecca M. Willett, and Robert Calderbank. Performance bounds for expander-based compressed sensing in Poisson noise. *IEEE Transactions on Signal Processing*, 59(9):4139–4153, 2011. doi:[10.1109/TSP.2011.2157913](https://doi.org/10.1109/TSP.2011.2157913).
- [34] Adil Ghani Khan, K. Parthasarathy, and Upinder Singh Bhalla. Odor representations in the mammalian olfactory bulb. *Wiley Interdisciplinary Reviews: Systems Biology and Medicine*, 2(5):603–611, 2010. doi:<https://doi.org/10.1002/wsbm.85>.
- [35] Limei Ma, Qiang Qiu, Stephen Gradwohl, Aaron Scott, Elden Q. Yu, Richard Alexander, Winfried Wiegand, and C. Ron Yu. Distributed representation of chemical features and tunotopic organization of glomeruli in the mouse olfactory bulb. *Proceedings of the National Academy of Sciences*, 109(14):5481–5486, 2012. doi:[10.1073/pnas.1117491109](https://doi.org/10.1073/pnas.1117491109).
- [36] Jeffrey S Isaacson. Odor representations in mammalian cortical circuits. *Current Opinion in Neurobiology*, 20(3):328–331, 2010. ISSN 0959-4388. doi:<https://doi.org/10.1016/j.conb.2010.02.004>. URL <https://www.sciencedirect.com/science/article/pii/S0959438810000267>. Sensory systems.
- [37] Yi-An Ma, Tianqi Chen, and Emily Fox. A complete recipe for stochastic gradient MCMC. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL https://proceedings.neurips.cc/paper_files/paper/2015/file/9a4400501febb2a95e79248486a5f6d3-Paper.pdf.
- [38] Paul Masset, Jacob Zavatone-Veth, J. Patrick Connor, Venkatesh Murthy, and Cengiz Pehlevan. Natural gradient enables fast sampling in spiking neural networks. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 22018–22034. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/8a0fd48510590071e3c129a79b8b8527-Paper-Conference.pdf.
- [39] Shun-ichi Amari. Natural Gradient Works Efficiently in Learning. *Neural Computation*, 10(2):251–276, 02 1998. ISSN 0899-7667. doi:[10.1162/089976698300017746](https://doi.org/10.1162/089976698300017746). URL <https://doi.org/10.1162/089976698300017746>.

- [40] Mark Girolami and Ben Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011. doi:[10.1111/j.1467-9868.2010.00765.x](https://doi.org/10.1111/j.1467-9868.2010.00765.x).
- [41] SueYeon Chung and L.F. Abbott. Neural population geometry: An approach for understanding biological and artificial neural networks. *Current Opinion in Neurobiology*, 70:137–144, 2021. ISSN 0959-4388. doi:<https://doi.org/10.1016/j.conb.2021.10.010>. URL <https://www.sciencedirect.com/science/article/pii/S0959438821001227>. Computational Neuroscience.
- [42] Yuansheng Zhou, Brian H. Smith, and Tatyana O. Sharpee. Hyperbolic geometry of the olfactory space. *Science Advances*, 4(8):eaq1458, 2018. doi:[10.1126/sciadv.aaq1458](https://doi.org/10.1126/sciadv.aaq1458). URL <https://www.science.org/doi/abs/10.1126/sciadv.aaq1458>.
- [43] Alexei Koulakov, Brian Kolterman, Armen Enikolopov, and Dmitry Rinberg. In search of the structure of human olfactory space. *Frontiers in Systems Neuroscience*, 5, 2011. ISSN 1662-5137. doi:[10.3389/fnsys.2011.00065](https://doi.org/10.3389/fnsys.2011.00065). URL <https://www.frontiersin.org/articles/10.3389/fnsys.2011.00065>.
- [44] Pierre-Marie Lledo, Gilles Gheusi, and Jean-Didier Vincent. Information processing in the mammalian olfactory system. *Physiological Reviews*, 85(1):281–317, 2005. doi:[10.1152/physrev.00008.2004](https://doi.org/10.1152/physrev.00008.2004).
- [45] Barry W. Ache and Janet M. Young. Olfaction: Diverse species, conserved principles. *Neuron*, 48(3):417–430, 2005. ISSN 0896-6273. doi:<https://doi.org/10.1016/j.neuron.2005.10.022>. URL <https://www.sciencedirect.com/science/article/pii/S0896627305008949>.
- [46] Chih-Ying Su, Karen Menuz, and John R. Carlson. Olfactory perception: Receptors, cells, and circuits. *Cell*, 139(1):45–59, 2009. ISSN 0092-8674. doi:<https://doi.org/10.1016/j.cell.2009.09.015>. URL <https://www.sciencedirect.com/science/article/pii/S0092867409011751>.
- [47] Naoshige Uchida, Cindy Poo, and Rafi Haddad. Coding and transformations in the olfactory system. *Annual Review of Neuroscience*, 37(1):363–385, 2014. doi:[10.1146/annurev-neuro-071013-013941](https://doi.org/10.1146/annurev-neuro-071013-013941).
- [48] Shin Nagayama, Ryota Homma, and Fumiaki Imamura. Neuronal organization of olfactory bulb circuits. *Frontiers in Neural Circuits*, 8:98, 2014. doi:[10.3389/fncir.2014.00098](https://doi.org/10.3389/fncir.2014.00098).
- [49] Shawn D. Burton. Inhibitory circuits of the mammalian main olfactory bulb. *Journal of Neurophysiology*, 118(4):2034–2051, 2017. doi:[10.1152/jn.00109.2017](https://doi.org/10.1152/jn.00109.2017).
- [50] An Wu, Bin Yu, and Takaki Komiyama. Plasticity in olfactory bulb circuits. *Current Opinion in Neurobiology*, 64:17–23, 2020. ISSN 0959-4388. doi:<https://doi.org/10.1016/j.conb.2020.01.007>. Systems Neuroscience.
- [51] Bettina Malnic, Junzo Hirono, Takaaki Sato, and Linda B Buck. Combinatorial receptor codes for odors. *Cell*, 96(5):713–723, 1999. ISSN 0092-8674. doi:[https://doi.org/10.1016/S0092-8674\(00\)80581-4](https://doi.org/10.1016/S0092-8674(00)80581-4). URL <https://www.sciencedirect.com/science/article/pii/S0092867400805814>.
- [52] Kenneth Murphy, Casey Weaver, and Charles Janeway. *Janeway’s Immunobiology*. Garland Science, New York, 9th edition edition, 2017. ISBN 9780815345053.
- [53] Linda Buck and Richard Axel. A novel multigene family may encode odorant receptors: a molecular basis for odor recognition. *Cell*, 65(1):175–187, 1991. doi:[https://doi.org/10.1016/0092-8674\(91\)90418-X](https://doi.org/10.1016/0092-8674(91)90418-X).
- [54] Yoshihito Niimura and Masatoshi Nei. Extensive gains and losses of olfactory receptor genes in mammalian evolution. *PLOS ONE*, 2(8):1–8, 08 2007. doi:[10.1371/journal.pone.0000708](https://doi.org/10.1371/journal.pone.0000708). URL <https://doi.org/10.1371/journal.pone.0000708>.

- [55] Yoshihito Niimura, Atsushi Matsui, and Kazushige Touhara. Extreme expansion of the olfactory receptor gene repertoire in african elephants and evolutionary dynamics of orthologous gene groups in 13 placental mammals. *Genome Research*, 24(9):1485–1496, 2014. doi:[10.1101/gr.169532.113](https://doi.org/10.1101/gr.169532.113).
- [56] Tiberiu Teşileanu, Simona Cocco, Rémi Monasson, and Vijay Balasubramanian. Adaptation of olfactory receptor abundances for efficient coding. *eLife*, 8:e39279, feb 2019. ISSN 2050-084X. doi:[10.7554/eLife.39279](https://doi.org/10.7554/eLife.39279). URL <https://doi.org/10.7554/eLife.39279>.
- [57] Veronica Egger and Nathaniel N. Urban. Dynamic connectivity in the mitral cell-granule cell microcircuit. *Seminars in Cell & Developmental Biology*, 17(4):424–432, 2006. ISSN 1084-9521. doi:<https://doi.org/10.1016/j.semcdb.2006.04.006>. URL <https://www.sciencedirect.com/science/article/pii/S1084952106000498>. Olfaction Animal Stem Cell Types.
- [58] Adam Kepecs, Naoshige Uchida, and Zachary F. Mainen. The Sniff as a Unit of Olfactory Processing. *Chemical Senses*, 31(2):167–179, 12 2005. ISSN 0379-864X. doi:[10.1093/chemse/bjj016](https://doi.org/10.1093/chemse/bjj016). URL <https://doi.org/10.1093/chemse/bjj016>.
- [59] Tobias Ackels, Andrew Erskine, Debanjan Dasgupta, Alina Cristina Marin, Tom P. A. Warner, Sina Tootoonian, Izumi Fukunaga, Julia J. Harris, and Andreas T. Schaefer. Fast odour dynamics are encoded in the olfactory system and guide behaviour. *Nature*, 593(7860): 558–563, May 2021. ISSN 1476-4687. doi:[10.1038/s41586-021-03514-2](https://doi.org/10.1038/s41586-021-03514-2). URL <https://doi.org/10.1038/s41586-021-03514-2>.
- [60] Naoshige Uchida and Zachary F. Mainen. Speed and accuracy of olfactory discrimination in the rat. *Nature Neuroscience*, 6(11):1224–1229, Nov 2003. ISSN 1546-1726. doi:[10.1038/nn1142](https://doi.org/10.1038/nn1142). URL <https://doi.org/10.1038/nn1142>.
- [61] Hartwig Spors, Dinu Florin Albeanu, Venkatesh N. Murthy, Dmitry Rinberg, Naoshige Uchida, Matt Wachowiak, and Rainer W. Friedrich. Illuminating vertebrate olfactory processing. *Journal of Neuroscience*, 32(41):14102–14108a, 2012. ISSN 0270-6474. doi:[10.1523/JNEUROSCI.3328-12.2012](https://doi.org/10.1523/JNEUROSCI.3328-12.2012).
- [62] Janine K. Reinert and Izumi Fukunaga. The facets of olfactory learning. *Current Opinion in Neurobiology*, 76:102623, 2022. ISSN 0959-4388. doi:<https://doi.org/10.1016/j.conb.2022.102623>. URL <https://www.sciencedirect.com/science/article/pii/S0959438822001179>.
- [63] Dan Rokni, Vivian Hemmelder, Vikrant Kapoor, and Venkatesh N. Murthy. An olfactory cocktail party: figure-ground segregation of odorants in rodents. *Nature Neuroscience*, 17(9): 1225–1232, Sep 2014. ISSN 1546-1726. doi:[10.1038/nn.3775](https://doi.org/10.1038/nn.3775). URL <https://doi.org/10.1038/nn.3775>.
- [64] Yan Li, Mitchell Swerdloff, Tianyu She, Asiyah Rahman, Naveen Sharma, Reema Shah, Michael Castellano, Daniel Mogel, Jason Wu, Asim Ahmed, James San Miguel, Jared Cohn, Nikesh Shah, Raddy L. Ramos, and Gonzalo H. Otazu. Robust odor identification in novel olfactory environments in mice. *Nature Communications*, 14(1):673, Feb 2023. ISSN 2041-1723. doi:[10.1038/s41467-023-36346-x](https://doi.org/10.1038/s41467-023-36346-x). URL <https://doi.org/10.1038/s41467-023-36346-x>.
- [65] Alice Berners-Lee, Elizabeth Shtrahman, Julien Grimaud, and Venkatesh N. Murthy. Experience-dependent evolution of odor mixture representations in piriform cortex. *PLOS Biology*, 21(4):1–29, 04 2023. URL <https://doi.org/10.1371/journal.pbio.3002086>.
- [66] Lior Lebovich, Michael Yunerman, Viviana Scaiewicz, Yonatan Loewenstein, and Dan Rokni. Paradoxical relationship between speed and accuracy in olfactory figure-background segregation. *PLOS Computational Biology*, 17:1–16, 12 2021. doi:[10.1371/journal.pcbi.1009674](https://doi.org/10.1371/journal.pcbi.1009674). URL <https://doi.org/10.1371/journal.pcbi.1009674>.
- [67] David G. Laing and G.W. Francis. The capacity of humans to identify odors in mixtures. *Physiology & Behavior*, 46(5):809–814, 1989. ISSN 0031-9384. doi:[https://doi.org/10.1016/0031-9384\(89\)90041-3](https://doi.org/10.1016/0031-9384(89)90041-3). URL <https://www.sciencedirect.com/science/article/pii/0031938489900413>.

- [68] A Jinks and David George Laing. The analysis of odor mixtures by humans: evidence for a configurational process. *Physiology & Behavior*, 72(1):51–63, 2001. ISSN 0031-9384. doi:[https://doi.org/10.1016/S0031-9384\(00\)00407-8](https://doi.org/10.1016/S0031-9384(00)00407-8). URL <https://www.sciencedirect.com/science/article/pii/S0031938400004078>.
- [69] Kei M. Igarashi, Nao Ieki, Myungho An, Yukie Yamaguchi, Shin Nagayama, Ko Kobayakawa, Reiko Kobayakawa, Manabu Tanifuji, Hitoshi Sakano, Wei R. Chen, and Kensaku Mori. Parallel mitral and tufted cell pathways route distinct odor information to different targets in the olfactory cortex. *Journal of Neuroscience*, 32(23):7970–7985, 2012. ISSN 0270-6474. doi:[10.1523/JNEUROSCI.0154-12.2012](https://doi.org/10.1523/JNEUROSCI.0154-12.2012).
- [70] Honggoo Chae, Arkarup Banerjee, Marie Dussauze, and Dinu F. Albeanu. Long-range functional loops in the mouse olfactory system and their roles in computing odor identity. *Neuron*, 110(23):3970–3985.e7, 2022. ISSN 0896-6273. doi:<https://doi.org/10.1016/j.neuron.2022.09.005>.
- [71] Priyanka Gupta, Dinu F Albeanu, and Upinder S Bhalla. Olfactory bulb coding of odors, mixtures and sniffs is a linear sum of odor time profiles. *Nature Neuroscience*, 18(2):272–281, 2015. doi:<https://doi.org/10.1038/nn.3913>.
- [72] Yilun Zhang and Tatyana O. Sharpee. A robust feedforward model of the olfactory system. *PLOS Computational Biology*, 12(4):1–15, 04 2016. doi:[10.1371/journal.pcbi.1004850](https://doi.org/10.1371/journal.pcbi.1004850). URL <https://doi.org/10.1371/journal.pcbi.1004850>.
- [73] David Zwicker, Arvind Murugan, and Michael P. Brenner. Receptor arrays optimized for natural odor statistics. *Proceedings of the National Academy of Sciences*, 113(20):5570–5575, 2016. doi:[10.1073/pnas.1600357113](https://doi.org/10.1073/pnas.1600357113). URL <https://www.pnas.org/doi/abs/10.1073/pnas.1600357113>.
- [74] Vijay Singh, Martin Tchernookov, and Vijay Balasubramanian. What the odor is not: Estimation by elimination. *Phys. Rev. E*, 104:024415, Aug 2021. doi:[10.1103/PhysRevE.104.024415](https://doi.org/10.1103/PhysRevE.104.024415). URL <https://link.aps.org/doi/10.1103/PhysRevE.104.024415>.
- [75] Patricia Duchamp-Viret, Lubomir Kostal, Michel Chaput, Petr Lánský, and Jean-Pierre Rospars. Patterns of spontaneous activity in single rat olfactory receptor neurons are different in normally breathing and tracheotomized animals. *Journal of Neurobiology*, 65(2):97–114, 2005. doi:<https://doi.org/10.1002/neu.20177>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/neu.20177>.
- [76] Lubomir Kostal, Petr Lansky, and Jean-Pierre Rospars. Neuronal coding and spiking randomness. *European Journal of Neuroscience*, 26(10):2693–2701, 2007. doi:<https://doi.org/10.1111/j.1460-9568.2007.05880.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9568.2007.05880.x>.
- [77] Alex Dytso, Michael Fauß, and H. Vincent Poor. The vector Poisson channel: On the linearity of the conditional mean estimator. *IEEE Transactions on Signal Processing*, 68:5894–5903, 2020. doi:[10.1109/TSP.2020.3025525](https://doi.org/10.1109/TSP.2020.3025525).
- [78] An Jin, Birsan Yazici, Angelique Ale, and Vasilis Ntziachristos. Preconditioning of the fluorescence diffuse optical tomography sensing matrix based on compressive sensing. *Optics Letters*, 37(20):4326–4328, Oct 2012. doi:[10.1364/OL.37.004326](https://doi.org/10.1364/OL.37.004326). URL <https://opg.optica.org/ol/abstract.cfm?URI=ol-37-20-4326>.
- [79] James Martens. New insights and perspectives on the natural gradient method. *Journal of Machine Learning Research*, 21(146):1–76, 2020. URL <http://jmlr.org/papers/v21/17-678.html>.
- [80] Gary R. Holt and Christof Koch. Shunting inhibition does not have a divisive effect on firing rates. *Neural Computation*, 9(5):1001–1013, 07 1997. ISSN 0899-7667. doi:[10.1162/neco.1997.9.5.1001](https://doi.org/10.1162/neco.1997.9.5.1001). URL <https://doi.org/10.1162/neco.1997.9.5.1001>.

- [81] Matthew Chalk, Paul Masset, Sophie Deneve, and Boris Gutkin. Sensory noise predicts divisive reshaping of receptive fields. *PLOS Computational Biology*, 13(6):1–26, 06 2017. doi:[10.1371/journal.pcbi.1005582](https://doi.org/10.1371/journal.pcbi.1005582). URL <https://doi.org/10.1371/journal.pcbi.1005582>.
- [82] David Marr and Tomaso Poggio. From understanding computation to understanding neural circuitry. Technical report, Massachusetts Institute of Technology AI Laboratory, 1976. URL <https://dspace.mit.edu/handle/1721.1/5782>.
- [83] Jacob A. Zavatone-Veth, Bara A. Badwan, and Damon A. Clark. A minimal synaptic model for direction selective neurons in *Drosophila*. *Journal of Vision*, 20(2):2–2, 02 2020. ISSN 1534-7362. doi:[10.1167/jov.20.2.2](https://doi.org/10.1167/jov.20.2.2). URL <https://doi.org/10.1167/jov.20.2.2>.
- [84] Francis Crick. The recent excitement about neural networks. *Nature*, 337(6203):129–132, Jan 1989. ISSN 1476-4687. doi:[10.1038/337129a0](https://doi.org/10.1038/337129a0). URL <https://doi.org/10.1038/337129a0>.
- [85] Jordan Guerguiev, Timothy P Lillicrap, and Blake A Richards. Towards deep learning with segregated dendrites. *eLife*, 6:e22901, dec 2017. ISSN 2050-084X. doi:[10.7554/eLife.22901](https://doi.org/10.7554/eLife.22901).
- [86] Timothy P. Lillicrap, Adam Santoro, Luke Marris, Colin J. Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, Jun 2020. ISSN 1471-0048. doi:[10.1038/s41583-020-0277-3](https://doi.org/10.1038/s41583-020-0277-3). URL <https://doi.org/10.1038/s41583-020-0277-3>.
- [87] Nathaniel N. Urban and Armen C. Arevian. Computing with dendrodendritic synapses in the olfactory bulb. *Annals of the New York Academy of Sciences*, 1170(1):264–269, 2009. doi:<https://doi.org/10.1111/j.1749-6632.2009.03899.x>. URL <https://nyaspubs.onlinelibrary.wiley.com/doi/abs/10.1111/j.1749-6632.2009.03899.x>.
- [88] Barry J. Davis and Foteos Macrides. The organization of centrifugal projections from the anterior olfactory nucleus, ventral hippocampal rudiment, and piriform cortex to the main olfactory bulb in the hamster: An autoradiographic study. *Journal of Comparative Neurology*, 203(3):475–493, 1981. doi:<https://doi.org/10.1002/cne.902030310>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cne.902030310>.
- [89] Alison M. Boyd, James F. Sturgill, Cindy Poo, and Jeffry S. Isaacson. Cortical feedback control of olfactory bulb circuits. *Neuron*, 76(6):1161–1174, 2012. ISSN 0896-6273. doi:<https://doi.org/10.1016/j.neuron.2012.10.020>. URL <https://www.sciencedirect.com/science/article/pii/S0896627312009403>.
- [90] Foivos Markopoulos, Dan Rokni, David H. Gire, and Venkatesh N. Murthy. Functional properties of cortical feedback projections to the olfactory bulb. *Neuron*, 76(6):1175–1188, 2012. ISSN 0896-6273. doi:<https://doi.org/10.1016/j.neuron.2012.10.028>. URL <https://www.sciencedirect.com/science/article/pii/S0896627312009488>.
- [91] Shawn D. Burton and Nathaniel N. Urban. Greater excitability and firing irregularity of tufted cells underlies distinct afferent-evoked activity of olfactory bulb mitral and tufted cells. *The Journal of Physiology*, 592(10):2097–2118, 2014. doi:<https://doi.org/10.1113/jphysiol.2013.269886>. URL <https://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.2013.269886>.
- [92] Shawn D. Burton and Nathaniel N. Urban. Rapid feedforward inhibition and asynchronous excitation regulate granule cell activity in the mammalian main olfactory bulb. *Journal of Neuroscience*, 35(42):14103–14122, 2015. ISSN 0270-6474. doi:[10.1523/JNEUROSCI.0746-15.2015](https://doi.org/10.1523/JNEUROSCI.0746-15.2015). URL <https://www.jneurosci.org/content/35/42/14103>.
- [93] Josefina del Marmol, Mackenzie A. Yedlin, and Vanessa Ruta. The structural basis of odorant recognition in insect olfactory receptors. *Nature*, 597(7874):126–131, Sep 2021. ISSN 1476-4687. doi:[10.1038/s41586-021-03794-8](https://doi.org/10.1038/s41586-021-03794-8). URL <https://doi.org/10.1038/s41586-021-03794-8>.

- [94] Joseph D. Zak, Gautam Reddy, Massimo Vergassola, and Venkatesh N. Murthy. Antagonistic odor interactions in olfactory sensory neurons are widespread in freely breathing mice. *Nature Communications*, 11(1):3350, Jul 2020. ISSN 2041-1723. doi:[10.1038/s41467-020-17124-5](https://doi.org/10.1038/s41467-020-17124-5). URL <https://doi.org/10.1038/s41467-020-17124-5>.
- [95] Armen C. Arevian, Vikrant Kapoor, and Nathaniel N. Urban. Activity-dependent gating of lateral inhibition in the mouse olfactory bulb. *Nature Neuroscience*, 11(1):80–87, Jan 2008. ISSN 1546-1726. doi:[10.1038/nn2030](https://doi.org/10.1038/nn2030). URL <https://doi.org/10.1038/nn2030>.
- [96] Kevin M. Cury and Naoshige Uchida. Robust odor coding via inhalation-coupled transient activity in the mammalian olfactory bulb. *Neuron*, 68(3):570–585, 2010. ISSN 0896-6273. doi:<https://doi.org/10.1016/j.neuron.2010.09.040>. URL <https://www.sciencedirect.com/science/article/pii/S0896627310007828>.
- [97] Roman Shusterman, Matthew C. Smear, Alexei A. Koulakov, and Dmitry Rinberg. Precise olfactory responses tile the sniff cycle. *Nature Neuroscience*, 14(8):1039–1044, Aug 2011. ISSN 1546-1726. doi:[10.1038/nn.2877](https://doi.org/10.1038/nn.2877). URL <https://doi.org/10.1038/nn.2877>.
- [98] David Donoho and Jared Tanner. Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1906):4273–4293, 2009. doi:[10.1098/rsta.2009.0152](https://doi.org/10.1098/rsta.2009.0152). URL <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2009.0152>.
- [99] Arian Maleki and David L. Donoho. Optimally tuned iterative reconstruction algorithms for compressed sensing. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):330–341, 2010. doi:[10.1109/JSTSP.2009.2039176](https://doi.org/10.1109/JSTSP.2009.2039176).
- [100] Richard C Gerkin and Jason B Castro. The number of olfactory stimuli that humans can discriminate is still unknown. *eLife*, 4:e08127, jul 2015. ISSN 2050-084X. doi:[10.7554/eLife.08127](https://doi.org/10.7554/eLife.08127). URL <https://doi.org/10.7554/eLife.08127>.
- [101] Markus Meister. On the dimensionality of odor space. *eLife*, 4:e07865, jul 2015. ISSN 2050-084X. doi:[10.7554/eLife.07865](https://doi.org/10.7554/eLife.07865). URL <https://doi.org/10.7554/eLife.07865>.
- [102] Emily J. Mayhew, Charles J. Arayata, Richard C. Gerkin, Brian K. Lee, Jonathan M. Magill, Lindsey L. Snyder, Kelsie A. Little, Chung Wen Yu, and Joel D. Mainland. Transport features predict if a molecule is odorous. *Proceedings of the National Academy of Sciences*, 119(15):e2116576119, 2022. doi:[10.1073/pnas.2116576119](https://doi.org/10.1073/pnas.2116576119). URL <https://www.pnas.org/doi/abs/10.1073/pnas.2116576119>.
- [103] David C Knill and Alexandre Pouget. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719, 2004. doi:[10.1016/j.tins.2004.10.007](https://doi.org/10.1016/j.tins.2004.10.007).
- [104] Konrad P Körding and Daniel M Wolpert. Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247, 2004. doi:[10.1038/nature02169](https://doi.org/10.1038/nature02169).
- [105] József Fiser, Pietro Berkes, Gergő Orbán, and Máté Lengyel. Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3):119–130, 2010. doi:[10.1016/j.tics.2010.01.003](https://doi.org/10.1016/j.tics.2010.01.003).
- [106] Torben Ott, Paul Masset, and Adam Kepecs. The neurobiology of confidence: From beliefs to neurons. In *Cold Spring Harbor Symposia on Quantitative Biology*, volume 83, pages 9–16. Cold Spring Harbor Laboratory Press, 2018. doi:[10.1101/sqb.2018.83.038794](https://doi.org/10.1101/sqb.2018.83.038794).
- [107] Gautam Reddy, Venkatesh N. Murthy, and Massimo Vergassola. Olfactory sensing and navigation in turbulent environments. *Annual Review of Condensed Matter Physics*, 13(1):191–213, 2022. doi:[10.1146/annurev-conmatphys-031720-032754](https://doi.org/10.1146/annurev-conmatphys-031720-032754). URL <https://doi.org/10.1146/annurev-conmatphys-031720-032754>.

- [108] William H. Calvin and Charles F. Stevens. Synaptic noise as a source of variability in the interval between action potentials. *Science*, 155(3764):842–844, 1967. doi:10.1126/science.155.3764.842. URL <https://www.science.org/doi/abs/10.1126/science.155.3764.842>.
- [109] Radford M Neal. *Probabilistic inference using Markov chain Monte Carlo methods*. Department of Computer Science, University of Toronto Toronto, ON, Canada, 1993. URL <https://www.cs.toronto.edu/~radford/review.abstract.html>.
- [110] Mengchen Zhu and Christopher J. Rozell. Modeling inhibitory interneurons in efficient sensory coding models. *PLoS Computational Biology*, 11(7):1–22, 07 2015. doi:10.1371/journal.pcbi.1004353. URL <https://doi.org/10.1371/journal.pcbi.1004353>.
- [111] Gautam Reddy, Joseph D Zak, Massimo Vergassola, and Venkatesh N Murthy. Antagonism in olfactory receptor neurons and its implications for the perception of odor mixtures. *eLife*, 7: e34958, apr 2018. ISSN 2050-084X. doi:10.7554/eLife.34958. URL <https://doi.org/10.7554/eLife.34958>.
- [112] Srinivas Gorur-Shandilya, Mahmut Demir, Junjiajia Long, Damon A Clark, and Thierry Emonet. Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli. *eLife*, 6:e27670, jun 2017. ISSN 2050-084X. doi:10.7554/eLife.27670. URL <https://doi.org/10.7554/eLife.27670>.
- [113] Nirag Kadakia and Thierry Emonet. Front-end Weber-Fechner gain control enhances the fidelity of combinatorial odor coding. *eLife*, 8:e45293, jun 2019. ISSN 2050-084X. doi:10.7554/eLife.45293. URL <https://doi.org/10.7554/eLife.45293>.
- [114] Cristina Savin and Sophie Denève. Spatio-temporal representations of uncertainty in spiking neural networks. *Advances in Neural Information Processing Systems*, 27, 2014. URL <https://proceedings.neurips.cc/paper/2014/hash/4e2545f819e67f0615003dd7e04a6087-Abstract.html>.
- [115] Michael Y.-S. Fang, Mayur Mudigonda, Ryan Zarccone, Amir Khosrowshahi, and Bruno A. Olshausen. Learning and Inference in Sparse Coding Models With Langevin Dynamics. *Neural Computation*, 34(8):1676–1700, 07 2022. ISSN 0899-7667. doi:10.1162/neco_a_01505. URL https://doi.org/10.1162/neco_a_01505.
- [116] Tatsuya Tsukahara, David H. Brann, Stan L. Pashkovski, Grigori Guitchev, Thomas Bozza, and Sandeep Robert Datta. A transcriptional rheostat couples past activity to future sensory responses. *Cell*, 184(26):6326–6343.e32, 2021. ISSN 0092-8674. doi:<https://doi.org/10.1016/j.cell.2021.11.022>. URL <https://www.sciencedirect.com/science/article/pii/S0092867421013374>.
- [117] Naoki Hiratani and Peter E. Latham. Rapid Bayesian learning in the mammalian olfactory system. *Nature Communications*, 11(1):3845, Jul 2020. ISSN 2041-1723. doi:10.1038/s41467-020-17490-0. URL <https://doi.org/10.1038/s41467-020-17490-0>.
- [118] Nikolai M. Chapochnikov, Cengiz Pehlevan, and Dmitri B. Chklovskii. Normative and mechanistic model of an adaptive circuit for efficient encoding and feature extraction. *Proceedings of the National Academy of Sciences*, 120(29):e2117484120, 2023. doi:10.1073/pnas.2117484120. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2117484120>.
- [119] Cengiz Pehlevan, Alexander Genkin, and Dmitri B. Chklovskii. A clustering neural network model of insect olfaction. In *51st Asilomar Conference on Signals, Systems, and Computers*, pages 593–600, 2017. doi:10.1109/ACSSC.2017.8335410.
- [120] Paul Masset, Shanshan Qin, and Jacob A Zavatore-Veth. Drifting neuronal representations: Bug or feature? *Biological Cybernetics*, pages 1–14, 2022. doi:doi.org/10.1007/s00422-021-00916-3.

- [121] Edmund Chong, Monica Moroni, Christopher Wilson, Shy Shoham, Stefano Panzeri, and Dmitry Rinberg. Manipulating synthetic optogenetic odors reveals the coding logic of olfactory perception. *Science*, 368(6497):eaba2357, 2020. doi:[10.1126/science.aba2357](https://doi.org/10.1126/science.aba2357).
- [122] Hirofumi Nakayama, Richard C. Gerkin, and Dmitry Rinberg. A behavioral paradigm for measuring perceptual distances in mice. *Cell Reports Methods*, 2(6):100233, 2022. ISSN 2667-2375. doi:<https://doi.org/10.1016/j.crmeth.2022.100233>. URL <https://www.sciencedirect.com/science/article/pii/S2667237522001023>.
- [123] Gonzalo H Otazu, Honggoo Chae, Martin B Davis, and Dinu F Albeanu. Cortical feedback decorrelates olfactory bulb output in awake mice. *Neuron*, 86(6):1461–1477, 2015. doi:<https://doi.org/10.1016/j.neuron.2015.05.023>.
- [124] An Wu, Bin Yu, Qiyu Chen, Gillian A. Matthews, Chen Lu, Evan Campbell, Kay M. Tye, and Takaki Komiyama. Context-dependent plasticity of adult-born neurons regulated by cortical feedback. *Science Advances*, 6(42):eabc8319, 2020. doi:[10.1126/sciadv.abc8319](https://doi.org/10.1126/sciadv.abc8319). URL <https://www.science.org/doi/abs/10.1126/sciadv.abc8319>.
- [125] James C. R. Whittington, Will Dorrell, Surya Ganguli, and Timothy Behrens. Disentanglement with biological constraints: A theory of functional cell types. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=9Z_GfhZnGH.
- [126] Mark Plumbley. Conditions for nonnegative independent component analysis. *IEEE Signal Processing Letters*, 9(6):177–180, 2002. doi:[10.1109/LSP.2002.800502](https://doi.org/10.1109/LSP.2002.800502).
- [127] Cengiz Pehlevan and Dmitri B Chklovskii. Optimization theory of hebbian/anti-hebbian networks for pca and whitening. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1458–1465. IEEE, 2015.
- [128] Cengiz Pehlevan, Sreyas Mohan, and Dmitri B. Chklovskii. Blind Nonnegative Source Separation Using Biological Neural Networks. *Neural Computation*, 29(11):2925–2954, 11 2017. ISSN 0899-7667. doi:[10.1162/neco_a_01007](https://doi.org/10.1162/neco_a_01007). URL https://doi.org/10.1162/neco_a_01007.
- [129] David Lipshutz, Cengiz Pehlevan, and Dmitri B. Chklovskii. Biologically plausible single-layer networks for nonnegative independent component analysis. *Biological Cybernetics*, 116(5):557–568, Dec 2022. ISSN 1432-0770. doi:[10.1007/s00422-022-00943-8](https://doi.org/10.1007/s00422-022-00943-8). URL <https://doi.org/10.1007/s00422-022-00943-8>.
- [130] Irina Higgins, Le Chang, Victoria Langston, Demis Hassabis, Christopher Summerfield, Doris Tsao, and Matthew Botvinick. Unsupervised deep learning identifies semantic disentanglement in single inferotemporal face patch neurons. *Nature Communications*, 12(1):6456, Nov 2021. ISSN 2041-1723. doi:[10.1038/s41467-021-26751-5](https://doi.org/10.1038/s41467-021-26751-5). URL <https://doi.org/10.1038/s41467-021-26751-5>.
- [131] Le Chang and Doris Y. Tsao. The code for facial identity in the primate brain. *Cell*, 169(6):1013–1028.e14, 2017. ISSN 0092-8674. doi:<https://doi.org/10.1016/j.cell.2017.05.011>. URL <https://www.sciencedirect.com/science/article/pii/S009286741730538X>.
- [132] Pinglei Bao, Liang She, Mason McGill, and Doris Y. Tsao. A map of object space in primate inferotemporal cortex. *Nature*, 583(7814):103–108, Jul 2020. ISSN 1476-4687. doi:[10.1038/s41586-020-2350-5](https://doi.org/10.1038/s41586-020-2350-5). URL <https://doi.org/10.1038/s41586-020-2350-5>.
- [133] Mattia Rigotti, Omri Barak, Melissa R. Warden, Xiao-Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi. The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451):585–590, May 2013. ISSN 1476-4687. doi:[10.1038/nature12160](https://doi.org/10.1038/nature12160). URL <https://doi.org/10.1038/nature12160>.
- [134] Junya Hirokawa, Alexander Vaughan, Paul Masset, Torben Ott, and Adam Kepecs. Frontal cortex neuron types categorically encode single decision variables. *Nature*, 576(7787):446–451, 2019. doi:<https://doi.org/10.1038/s41586-019-1816-9>.

- [135] Alexis Dubreuil, Adrian Valente, Manuel Beiran, Francesca Mastrogiuseppe, and Srdjan Ostojic. The role of population structure in computations through neural dynamics. *Nature Neuroscience*, 25(6):783–794, 2022. doi:<https://doi.org/10.1038/s41593-022-01088-4>.
- [136] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021. URL <https://geometricdeeplearning.com/>.
- [137] Jacob A. Zavatone-Veth, Sheng Yang, Julian A. Rubinien, and Cengiz Pehlevan. Neural networks learn to magnify areas near decision boundaries. *arXiv*, 2023. doi:<https://doi.org/10.48550/arXiv.2301.11375>.
- [138] Eftychios A. Pnevmatikakis and Andrea Giovannucci. Normcorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *Journal of Neuroscience Methods*, 291:83–94, 2017. ISSN 0165-0270. doi:<https://doi.org/10.1016/j.jneumeth.2017.07.031>.
- [139] Alexander Mathis, Dan Rokni, Vikrant Kapoor, Matthias Bethge, and Venkatesh N. Murthy. Reading out olfactory receptors: Feedforward circuits detect odors in mixtures without demixing. *Neuron*, 91(5):1110–1123, 2016. ISSN 0896-6273. doi:<https://doi.org/10.1016/j.neuron.2016.08.007>. URL <https://www.sciencedirect.com/science/article/pii/S0896627316304998>.

A Notational conventions

In this Appendix, we define the notational conventions used throughout the paper. We write $\mathbb{R}_+^n = \{\mathbf{x} \in \mathbb{R}^n : x_1, \dots, x_n \geq 0\}$ for the non-negative reals, and $\mathbb{N} = \{0, 1, \dots\}$ for the natural numbers. For a vector $\boldsymbol{\lambda} \in \mathbb{R}_+^n$, we write

$$\mathbf{x} \sim \text{Poisson}(\boldsymbol{\lambda}) \quad (\text{A.1})$$

if the probability mass function of $\mathbf{x} \in \mathbb{N}^n$ is

$$p(\mathbf{x}) = \prod_{j=1}^n \frac{\lambda_j^{x_j}}{x_j!} e^{-\lambda_j}. \quad (\text{A.2})$$

For scalars $\alpha > 0$ and $\lambda > 0$, we write that a vector $\mathbf{x} \in \mathbb{R}_+^n$ is distributed as

$$\mathbf{x} \sim \text{Gamma}(\alpha, \lambda) \quad (\text{A.3})$$

if its probability density function is

$$p(\mathbf{x}) = \prod_{j=1}^n \frac{x_j^{\alpha-1}}{\Gamma(\alpha)} e^{-\lambda x_j}. \quad (\text{A.4})$$

Similarly, for vectors $\boldsymbol{\alpha}, \boldsymbol{\lambda} \in \mathbb{R}_+^n$, we write

$$\mathbf{x} \sim \text{Gamma}(\boldsymbol{\alpha}, \boldsymbol{\lambda}) \quad (\text{A.5})$$

if

$$p(\mathbf{x}) = \prod_{j=1}^n \frac{x_j^{\alpha_j-1}}{\Gamma(\alpha_j)} e^{-\lambda_j x_j}. \quad (\text{A.6})$$

For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}_+^n$, we write $\mathbf{x} \odot \mathbf{y}$ for their Hadamard (elementwise) product

$$(\mathbf{x} \odot \mathbf{y})_i = x_i y_i \quad (\text{A.7})$$

and $\mathbf{x} \oslash \mathbf{y}$ for their elementwise ratio

$$(\mathbf{x} \oslash \mathbf{y})_i = \frac{x_i}{y_i}. \quad (\text{A.8})$$

B Detailed derivation of circuit model

In this Appendix, we give a detailed derivation of the circuit models introduced in Sections 3 and 6 of the main text. Here, we focus on the setting of the sampling circuit; the circuit algorithm to compute the MAP introduced in Section 3 of the main text can be recovered at each step by dropping the additive noise terms from the dynamics.

We recall that our goal is to sample the posterior $p(\mathbf{c} | \mathbf{s})$ over concentrations $\mathbf{c}$ given OSN activity $\mathbf{s}$, for a Poisson likelihood

$$\mathbf{s} | \mathbf{c} \sim \text{Poisson}(\mathbf{r}_0 + \mathbf{A}\mathbf{c}) \quad (\text{B.1})$$

and a Gamma prior

$$\mathbf{c} \sim \text{Gamma}(\boldsymbol{\alpha}, \boldsymbol{\lambda}). \quad (\text{B.2})$$

Here, we use the conventions of Appendix A for Poisson and Gamma random vectors. Discarding normalization constants that do not depend on $\mathbf{c}$, the density of the posterior is then

$$p(\mathbf{c} | \mathbf{s}) \propto p(\mathbf{s} | \mathbf{c}) p(\mathbf{c}) \quad (\text{B.3})$$

$$\propto \left[\prod_{k=1}^{n_{\text{OSN}}} (\mathbf{r}_0 + \mathbf{A}\mathbf{c})_k^{s_k} e^{-(\mathbf{r}_0 + \mathbf{A}\mathbf{c})_k} \right] \left[\prod_{j=1}^{n_{\text{odor}}} c_j^{\alpha_j-1} e^{-\lambda_j c_j} \right]. \quad (\text{B.4})$$

We will sample from this distribution using Langevin dynamics, without explicitly constraining the concentration estimate to be non-negative [21, 22]. This corresponds to the stochastic dynamics

$$\dot{\mathbf{c}}(t) = \nabla_{\mathbf{c}} \log p(\mathbf{c} | \mathbf{s}) + \boldsymbol{\eta}(t), \quad (\text{B.5})$$

where $\boldsymbol{\eta}(t)$ is an n_{odor} -dimensional zero-mean Gaussian noise process with $\mathbb{E}[\eta_j(t)\eta_{j'}(t')] = 2\delta_{jj'}\delta(t-t')$. The gradient of the log-posterior is

$$\frac{\partial \log p(\mathbf{c} | \mathbf{s})}{\partial c_j} = \sum_{k=1}^{n_{\text{OSN}}} \left(\frac{s_k}{(\mathbf{r}_0 + \mathbf{A}\mathbf{c})_k} - 1 \right) A_{kj} + \frac{\alpha_j - 1}{c_j} - \lambda_j, \quad (\text{B.6})$$

or, in vector notation,

$$\nabla_{\mathbf{c}} \log p(\mathbf{c} | \mathbf{s}) = \mathbf{A}^\top [\mathbf{s} \odot (\mathbf{r}_0 + \mathbf{A}\mathbf{c}) - \mathbf{1}] + (\boldsymbol{\alpha} - \mathbf{1}) \odot \mathbf{c} - \boldsymbol{\lambda}. \quad (\text{B.7})$$

The Langevin dynamics then become

$$\dot{\mathbf{c}}(t) = \mathbf{A}^\top [\mathbf{s} \odot (\mathbf{r}_0 + \mathbf{A}\mathbf{c}) - \mathbf{1}] + (\boldsymbol{\alpha} - \mathbf{1}) \odot \mathbf{c} - \boldsymbol{\lambda} + \boldsymbol{\eta}(t). \quad (\text{B.8})$$

This is closely related to the Langevin sampling algorithm introduced in equation (3.9) of Grabska-Barwińska et al. [21], but here we are not separately inferring odor concentration and odor presence. We note that, for $\alpha_j > 1$, the term $(\boldsymbol{\alpha} - \mathbf{1}) \odot \mathbf{c}$ gives a force that diverges for small c_j .

However, in this setup single neurons encode single odors. We instead want to allow for a distributed code, where the population of neurons responsible for sampling may not directly encode single odors. To distribute the code, we follow previous work by Masset et al. [38] in using the “complete recipe” for stochastic gradient MCMC [37]. From the “complete recipe”, we know that the dynamics

$$\tau_g \dot{\mathbf{c}}(t) = \boldsymbol{\Gamma} \boldsymbol{\Gamma}^\top \nabla_{\mathbf{c}} \log p(\mathbf{c} | \mathbf{s}) + \boldsymbol{\Gamma} \boldsymbol{\xi}(t), \quad (\text{B.9})$$

for a (expectantly named) time constant $\tau_g > 0$ and any (potentially non-square) matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{n_{\text{odor}} \times n_g}$ such that $\boldsymbol{\Gamma} \boldsymbol{\Gamma}^\top$ is positive-definite, will have as their stationary distribution the posterior $p(\mathbf{c} | \mathbf{s})$. here, we have replaced the n_{odor} -dimensional noise process $\boldsymbol{\eta}(t)$ with an n_g -dimensional noise process $\boldsymbol{\xi}(t)$ with zero mean and covariance

$$\mathbb{E}[\xi_j(t)\xi_{j'}(t')] = 2\tau_g \delta_{jj'} \delta(t-t'). \quad (\text{B.10})$$

Then, we can write the concentration estimate as

$$\mathbf{c}(t) = \boldsymbol{\Gamma} \mathbf{g}(t) \quad (\text{B.11})$$

where the activity of the neurons $\mathbf{g} \in \mathbb{R}^{n_g}$ follows the dynamics

$$\tau_g \dot{\mathbf{g}}(t) = \boldsymbol{\Gamma}^\top \nabla_{\mathbf{c}} \log p(\mathbf{g} | \mathbf{s}) \Big|_{\mathbf{c}=\boldsymbol{\Gamma} \mathbf{g}} + \boldsymbol{\xi}(t). \quad (\text{B.12})$$

Using the gradient of the log-posterior computed above, we then have

$$\tau_g \dot{\mathbf{g}}(t) = (\mathbf{A}\boldsymbol{\Gamma})^\top [\mathbf{s} \odot (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma} \mathbf{g}) - \mathbf{1}] + \boldsymbol{\Gamma}^\top [(\boldsymbol{\alpha} - \mathbf{1}) \odot (\boldsymbol{\Gamma} \mathbf{g})] - \boldsymbol{\Gamma}^\top \boldsymbol{\lambda} + \boldsymbol{\xi}(t). \quad (\text{B.13})$$

These dynamics include two divisive non-linearities, which can be complex to implement in biophysical models of single neurons [80]. Using the approach proposed in Chalk et al. [81], we can linearize the inference by introducing two additional cell types that have as their fixed points the elementwise divisions $\mathbf{s} \odot (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma} \mathbf{g})$ and $(\boldsymbol{\alpha} - \mathbf{1}) \odot (\boldsymbol{\Gamma} \mathbf{g})$. We then introduce a population $\mathbf{p}$ of n_{OSN} neurons, with dynamics

$$\tau_p \dot{\mathbf{p}}(t) = \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma} \mathbf{g}), \quad (\text{B.14})$$

such that the fixed point is

$$\mathbf{p}^* = \mathbf{s} \odot (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma} \mathbf{g}). \quad (\text{B.15})$$

We finally introduce a third population $\mathbf{z}$ of n_{odor} neurons, with dynamics

$$\tau_z \dot{\mathbf{z}}(t) = \boldsymbol{\alpha} - \mathbf{1} - \mathbf{z} \odot (\boldsymbol{\Gamma} \mathbf{g}), \quad (\text{B.16})$$

such that the fixed point is

$$\mathbf{z}^* = (\boldsymbol{\alpha} - \mathbf{1}) \odot (\boldsymbol{\Gamma} \mathbf{g}). \quad (\text{B.17})$$

To be more precise, the cell types $\mathbf{p}$ and $\mathbf{z}$ compute the desired elementwise divisions in the pseudo-steady-state regime $\tau_p, \tau_z \downarrow 0$.

Putting everything together, this gives the circuit dynamics

$$\mathbf{c}(t) = \mathbf{\Gamma} \mathbf{g}(t) \quad (\text{B.18})$$

$$\tau_g \dot{\mathbf{g}}(t) = (\mathbf{A}\mathbf{\Gamma})^\top (\mathbf{p} - \mathbf{1}) + \mathbf{\Gamma}^\top (\mathbf{z} - \boldsymbol{\lambda}) + \boldsymbol{\xi}(t) \quad (\text{B.19})$$

$$\tau_p \dot{\mathbf{p}}(t) = \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A}\mathbf{\Gamma} \mathbf{g}) \quad (\text{B.20})$$

$$\tau_z \dot{\mathbf{z}}(t) = \boldsymbol{\alpha} - \mathbf{1} - \mathbf{z} \odot \mathbf{c}, \quad (\text{B.21})$$

where we recall that the covariance of the zero-mean Gaussian noise process $\boldsymbol{\xi}(t)$ is

$$\mathbb{E}[\xi_j(t) \xi_{j'}(t')] = 2\tau_g \delta_{jj'} \delta(t - t'). \quad (\text{B.22})$$

We note that the two constant terms in the dynamics of $\mathbf{g}$ can be grouped into an overall leak term $-\mathbf{\Gamma}^\top (\mathbf{A}^\top \mathbf{1} + \boldsymbol{\lambda})$. As detailed in the main text, we interpret $\mathbf{p}$ as M/T cells, $\mathbf{g}$ as granule cells, and $\mathbf{z}$ as cortical feedback.

C Stability of the MAP fixed point

In this Appendix, we aim to get some understanding for how the introduction of the cell type $\mathbf{p}$ to represent the elementwise division affects the inference, particularly when τ_p is comparable to τ_g . For simplicity, we specialize to the case of an exponential prior (i.e., we set $\boldsymbol{\alpha} = \mathbf{1}$), in which case the general circuit (3) simplifies to

$$\begin{aligned} \mathbf{c}(t) &= \mathbf{\Gamma} \mathbf{g}(t) \\ \tau_g \dot{\mathbf{g}}(t) &= (\mathbf{A}\mathbf{\Gamma})^\top (\mathbf{p} - \mathbf{1}) - \mathbf{\Gamma}^\top \boldsymbol{\lambda} \\ \tau_p \dot{\mathbf{p}}(t) &= \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A}\mathbf{\Gamma} \mathbf{g}). \end{aligned} \quad (\text{C.1})$$

C.1 Analysis of a two-cell circuit

To build intuition, we first consider a circuit with $n_{\text{odor}} = n_{\text{OSN}} = n_g = 1$:

$$\begin{aligned} c(t) &= \gamma g(t) \\ \tau_g \frac{dg}{dt} &= a\gamma(p - 1) - \gamma\lambda \\ \tau_p \frac{dp}{dt} &= s - p(r_0 + a\gamma g). \end{aligned} \quad (\text{C.2})$$

In this case, it is useful to non-dimensionalize the system. Re-scale time as

$$\tilde{t} = \frac{a\gamma}{\tau_g} t, \quad (\text{C.3})$$

and define the dimensionless parameters and input

$$\tilde{\tau} = \frac{\tau_p}{\tau_g}, \quad (\text{C.4})$$

$$\tilde{r}_0 = \frac{r_0}{a\gamma}, \quad (\text{C.5})$$

$$\tilde{\lambda} = \frac{\lambda}{a}, \quad (\text{C.6})$$

$$\tilde{s} = \frac{s}{a\gamma}. \quad (\text{C.7})$$

Then, the dynamics are

$$\frac{dg}{d\tilde{t}} = p - 1 - \tilde{\lambda} \quad (\text{C.8})$$

$$\tilde{\tau} \frac{dp}{d\tilde{t}} = \tilde{s} - p(\tilde{r}_0 + g). \quad (\text{C.9})$$

These dynamics have a single fixed point at

$$g_* = \frac{\tilde{s}}{1 + \tilde{\lambda}} - \tilde{r}_0 \quad (\text{C.10})$$

$$p_* = 1 + \tilde{\lambda}. \quad (\text{C.11})$$

Linearizing about this fixed point, we have

$$\frac{d}{dt} \begin{pmatrix} \delta g \\ \delta p \end{pmatrix} = \mathbf{M} \begin{pmatrix} \delta g \\ \delta p \end{pmatrix} \quad (\text{C.12})$$

for

$$\mathbf{M} = \frac{1}{\tilde{\tau}} \begin{pmatrix} 0 & \tilde{\tau} \\ -(1 + \tilde{\lambda}) & -(1 + \tilde{\lambda})^{-1} \tilde{s} \end{pmatrix}. \quad (\text{C.13})$$

The eigenvalues of $\mathbf{M}$ are easily computed as

$$\Lambda_{\pm}(\mathbf{M}) = \frac{-\tilde{s} \pm \sqrt{\tilde{s}^2 - 4\tilde{\tau}(1 + \tilde{\lambda})^3}}{2\tilde{\tau}(1 + \tilde{\lambda})}. \quad (\text{C.14})$$

For any positive $\tilde{\tau}$ and non-negative $\tilde{\lambda}$, we have

$$\text{Re } \Lambda_{\pm}(\mathbf{M}) = \frac{-\tilde{s} \pm \text{Re} \sqrt{\tilde{s}^2 - 4\tilde{\tau}(1 + \tilde{\lambda})^3}}{2\tilde{\tau}(1 + \tilde{\lambda})} \quad (\text{C.15})$$

$$= \begin{cases} \frac{-\tilde{s} \pm \sqrt{\tilde{s}^2 - 4\tilde{\tau}(1 + \tilde{\lambda})^3}}{2\tilde{\tau}(1 + \tilde{\lambda})} & \tilde{s}^2 - 4\tilde{\tau}(1 + \tilde{\lambda})^3 > 0 \\ \frac{-\tilde{s}}{2\tilde{\tau}(1 + \tilde{\lambda})} & \text{otherwise,} \end{cases} \quad (\text{C.16})$$

hence it is easy to see that the real parts of both of these eigenvalues are strictly negative so long as $\tilde{s} > 0$, meaning that the system is stable. If $\tilde{s} = 0$ —which should be exceedingly rare if the OSNs have a baseline rate—then

$$\Lambda_{\pm}(\mathbf{M}) \Big|_{\tilde{s}=0} = \pm 2i \sqrt{\frac{1 + \tilde{\lambda}}{\tilde{\tau}}}, \quad (\text{C.17})$$

and there can be oscillations. Therefore, a linear stability analysis suggests that the MAP fixed point of this two-cell circuit should be stable even for large $\tilde{\tau}$ in the presence of a non-zero input, though we expect the relaxation timescales to grow as $\tilde{\tau}$ becomes larger.

C.2 Analysis of the full circuit

We now consider the full MAP circuit

$$\begin{aligned} \mathbf{c}(t) &= \mathbf{\Gamma} \mathbf{g}(t) \\ \tau_g \dot{\mathbf{g}}(t) &= (\mathbf{A} \mathbf{\Gamma})^\top (\mathbf{p} - \mathbf{1}) - \mathbf{\Gamma}^\top \boldsymbol{\lambda} \\ \tau_p \dot{\mathbf{p}}(t) &= \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A} \mathbf{\Gamma} \mathbf{g}). \end{aligned} \quad (\text{C.18})$$

We assume that $n_{\text{OSN}} < n_{\text{odor}} \leq n_g$ and that $\mathbf{A} \mathbf{\Gamma}$ is of full row rank, i.e., it has rank n_{OSN} . We recall that $\mathbf{g} \in \mathbb{R}_+^{n_g}$ and $\mathbf{p} \in \mathbb{R}_+^{n_g}$. Therefore, the dynamics of $\mathbf{g}$ span only an n_{odor} -dimensional subspace, and, in particular, the $\mathbf{p}$ -dependent term affects only an n_{OSN} -dimensional subspace.

We start by observing that the dynamics of $\mathbf{p}$ depend on $\mathbf{g}$ only through

$$\mathbf{q} \equiv \mathbf{A} \mathbf{\Gamma} \mathbf{g}. \quad (\text{C.19})$$

Importantly, if $\mathbf{A} \mathbf{\Gamma}$ is positivity-preserving, then $\mathbf{q} \in \mathbb{R}_+^{n_{\text{OSN}}}$. Then, we have the closed dynamics

$$\begin{aligned} \tau_g \dot{\mathbf{q}} &= (\mathbf{A} \mathbf{\Gamma} \mathbf{\Gamma}^\top \mathbf{A}^\top) (\mathbf{p} - \mathbf{1}) - \mathbf{A} \mathbf{\Gamma} \mathbf{\Gamma}^\top \boldsymbol{\lambda} \\ \tau_p \dot{\mathbf{p}} &= \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{q}) \end{aligned} \quad (\text{C.20})$$

for $(\mathbf{q}, \mathbf{p}) \in \mathbb{R}_+^{2n_{\text{OSN}}}$. The fixed point of this system is of course determined by the conditions $d(\mathbf{q}, \mathbf{p})/dt = \mathbf{0}$, which gives

$$(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \mathbf{A}^\top)(\mathbf{p}^* - \mathbf{1}) = \mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \boldsymbol{\lambda} \quad (\text{C.21})$$

$$\mathbf{p}^* \odot (\mathbf{r}_0 + \mathbf{q}^*) = \mathbf{s} \quad (\text{C.22})$$

subject to the non-negativity constraints. By our assumptions on the rank of $\mathbf{A}\mathbf{\Gamma}$, the symmetric matrix $\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \mathbf{A}^\top$ is positive-definite and thus can be inverted to solve the first condition for $\mathbf{p}^*$:

$$\mathbf{p}^* = \mathbf{1} + (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \mathbf{A}^\top)^{-1} \mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \boldsymbol{\lambda}. \quad (\text{C.23})$$

For the second condition to be satisfied, we can see that the elements of $\mathbf{p}^*$ must be strictly positive at the fixed point, which gives a self-consistency condition on $\mathbf{A}$, $\mathbf{\Gamma}$, and $\boldsymbol{\lambda}$. Assuming that this holds, we then have

$$\mathbf{q}^* = \mathbf{s} \oslash \mathbf{p}^* - \mathbf{r}_0, \quad (\text{C.24})$$

which again gives a self-consistency condition as non-negativity is required.

Assuming these conditions hold, we can then linearize the dynamics about the fixed point. For convenience, we non-dimensionalize time through

$$\tilde{t} = \frac{t}{\tau_g}, \quad (\text{C.25})$$

and

$$\tilde{\tau} = \frac{\tau_p}{\tau_g}, \quad (\text{C.26})$$

which yields the linearized dynamics

$$\frac{d}{d\tilde{t}} \begin{pmatrix} \delta \mathbf{q} \\ \delta \mathbf{p} \end{pmatrix} = \mathbf{M} \begin{pmatrix} \delta \mathbf{q} \\ \delta \mathbf{p} \end{pmatrix} \quad (\text{C.27})$$

for

$$\mathbf{M} = \begin{pmatrix} \mathbf{0} & (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \mathbf{A}^\top) \\ -\tilde{\tau}^{-1} \text{diag}(\mathbf{p}^*) & -\tilde{\tau}^{-1} \text{diag}(\mathbf{r}_0 + \mathbf{q}^*) \end{pmatrix} \quad (\text{C.28})$$

$$= \begin{pmatrix} \mathbf{0} & (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \mathbf{A}^\top) \\ -\tilde{\tau}^{-1} \text{diag}(\mathbf{p}^*) & -\tilde{\tau}^{-1} \text{diag}(\mathbf{s} \oslash \mathbf{p}^*) \end{pmatrix}. \quad (\text{C.29})$$

Using the fact that the diagonal matrices $\text{diag}(\mathbf{p}^*)$ and $\text{diag}(\mathbf{s} \oslash \mathbf{p}^*)$ commute, the characteristic polynomial of $\mathbf{M}$ is

$$\det(\mathbf{M} - \mu \mathbf{I}_{2n_{\text{OSN}}}) = \det \begin{pmatrix} \mu \mathbf{I}_{n_{\text{OSN}}} & -(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \mathbf{A}^\top) \\ \tilde{\tau}^{-1} \text{diag}(\mathbf{p}^*) & \tilde{\tau}^{-1} \text{diag}(\mathbf{s} \oslash \mathbf{p}^*) + \mu \mathbf{I}_{n_{\text{OSN}}} \end{pmatrix} \quad (\text{C.30})$$

$$= \det[\mu^2 \mathbf{I}_{n_{\text{OSN}}} + \mu \tilde{\tau}^{-1} \text{diag}(\mathbf{s} \oslash \mathbf{p}^*) + \tilde{\tau}^{-1} (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \mathbf{A}^\top) \text{diag}(\mathbf{p}^*)]. \quad (\text{C.31})$$

One case that is particularly easy to solve is when the symmetric positive-definite matrix $\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top \mathbf{A}^\top$ is in fact diagonal, with positive diagonal entries a_j . Then, we have

$$\det(\mathbf{M} - \mu \mathbf{I}_{2n_{\text{OSN}}}) = \prod_{j=1}^{n_{\text{OSN}}} \left[\mu^2 + \frac{s_j}{\tilde{\tau} p_j^*} \mu + \frac{a_j p_j^*}{\tilde{\tau}} \right], \quad (\text{C.32})$$

meaning that the eigenvalues of $\mathbf{M}$ are

$$\mu_{j,\pm} = \frac{1}{2} \left[-\frac{s_j}{\tilde{\tau} p_j^*} \pm \sqrt{\left(\frac{s_j}{\tilde{\tau} p_j^*} \right)^2 - 4 \frac{a_j p_j^*}{\tilde{\tau}}} \right], \quad (\text{C.33})$$

which under the given assumptions always have strictly positive real part for non-negative inputs.

Now, more generally, assume that $\mathbf{M}$ is diagonalizable, and let $\mathbf{m}$ be a un-normalized eigenvector of $\mathbf{M}$ with eigenvalue μ . Then, writing

$$\mathbf{m} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} \quad (\text{C.34})$$

for $\mathbf{u}, \mathbf{v} \in \mathbb{C}^{n_{\text{OSN}}}$, the eigenvector condition

$$\mathbf{M}\mathbf{m} = \mu\mathbf{m} \quad (\text{C.35})$$

implies that $\mathbf{u}$ and $\mathbf{v}$ satisfy

$$\frac{1}{\tau_g}(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)\mathbf{v} = \mu\mathbf{u} \quad (\text{C.36})$$

$$-\frac{1}{\tau_p}\text{diag}(\mathbf{p}^*)\mathbf{u} - \frac{1}{\tau_p}\text{diag}(\mathbf{s} \oslash \mathbf{p}^*)\mathbf{v} = \mu\mathbf{v}. \quad (\text{C.37})$$

Assuming that μ is non-zero, we can solve the first equation for $\mathbf{u}$ and then substitute the result into the second to obtain a quadratic eigenproblem for μ and $\mathbf{v}$:

$$\mu^2\mathbf{v} + \mu\frac{1}{\tau_p}\text{diag}(\mathbf{s} \oslash \mathbf{p}^*)\mathbf{v} + \frac{1}{\tau_p\tau_g}\text{diag}(\mathbf{p}^*)(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)\mathbf{v} = 0. \quad (\text{C.38})$$

We will not attempt to solve this eigenproblem, but will instead attempt to extract information about possible values of μ for a fixed $\mathbf{v}$. Suppose (without loss of generality given the assumption that it is nonzero, as otherwise we may divide by its norm) that $\mathbf{v}$ is a unit vector. Then, acting with $\mathbf{v}^\dagger$ from the left, we have

$$\mu^2 + \frac{a}{\tau_p}\mu + \frac{b}{\tau_p\tau_g} = 0. \quad (\text{C.39})$$

where we define the coefficients

$$a \equiv \mathbf{v}^\dagger \text{diag}(\mathbf{s} \oslash \mathbf{p}^*)\mathbf{v} \quad (\text{C.40})$$

and

$$b \equiv \mathbf{v}^\dagger \text{diag}(\mathbf{p}^*)(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)\mathbf{v}. \quad (\text{C.41})$$

So long as $\mathbf{s}$ and $\mathbf{p}^*$ are positive, a is real and positive. Let

$$\tilde{\lambda} \equiv (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)^{-1}\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\lambda \quad (\text{C.42})$$

such that

$$\mathbf{p}^* = \mathbf{1} + \tilde{\lambda}. \quad (\text{C.43})$$

By assumption, the elements of $\tilde{\lambda}$ are strictly greater than -1 . Then,

$$\text{Re}(b) = \mathbf{v}^\dagger \frac{\text{diag}(\mathbf{p}^*)(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top) + (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)\text{diag}(\mathbf{p}^*)}{2}\mathbf{v} \quad (\text{C.44})$$

$$= \mathbf{v}^\dagger(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)\mathbf{v} + \mathbf{v}^\dagger \frac{\text{diag}(\tilde{\lambda})(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top) + (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)\text{diag}(\tilde{\lambda})}{2}\mathbf{v}, \quad (\text{C.45})$$

while

$$\text{Im}(b) = \mathbf{v}^\dagger \frac{\text{diag}(\mathbf{p}^*)(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top) - (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)\text{diag}(\mathbf{p}^*)}{2i}\mathbf{v} \quad (\text{C.46})$$

$$= \mathbf{v}^\dagger \frac{\text{diag}(\tilde{\lambda})(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top) - (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^\top\mathbf{A}^\top)\text{diag}(\tilde{\lambda})}{2i}\mathbf{v}. \quad (\text{C.47})$$

Then, solving the quadratic equation for μ , we have

$$\mu = \frac{1}{\tau_p} \frac{-a \pm \sqrt{a^2 - 4\tau_g b}}{2}, \quad (\text{C.48})$$

for

$$\tilde{\tau} \equiv \frac{\tau_p}{\tau_g}. \quad (\text{C.49})$$

This gives

$$\text{Re}(\mu) = -\frac{a}{2\tau_p} \pm \frac{1}{2\tau_p} \text{Re} \sqrt{a^2 - 4\tilde{\tau}b} \quad (\text{C.50})$$

$$= -\frac{a}{2\tau_p} \pm \frac{1}{2\tau_p} \sqrt{\frac{\sqrt{(a^2 - 4\tilde{\tau} \text{Re}(b))^2 + 16\tilde{\tau}^2 \text{Im}(b)^2} + a^2 - 4\tilde{\tau} \text{Re}(b)}{2}}. \quad (\text{C.51})$$

If $\text{Im}(b) = 0$, then we are in the same situation as we were for the two-neuron circuit, and the fixed point is thus stable.

More generally, for the system to be stable we want to have

$$a > \sqrt{\frac{\sqrt{(a^2 - 4\tilde{\tau} \text{Re}(b))^2 + 16\tilde{\tau}^2 \text{Im}(b)^2} + a^2 - 4\tilde{\tau} \text{Re}(b)}{2}}. \quad (\text{C.52})$$

For either sign of $a^2 - 4\tilde{\tau} \text{Re}(b)$, one can show that this holds provided that

$$\text{Re}(b) > 0 \quad (\text{C.53})$$

and

$$\sqrt{\tilde{\tau}} |\text{Im}(b)| < a \sqrt{\text{Re}(b)}. \quad (\text{C.54})$$

This amounts to a combined condition on the effective prior strength $\tilde{\lambda}$ and the relative time constants. Therefore, for sufficiently small $\tilde{\lambda}$, we have stability for a broad range of τ_p .

How can we go from stability of the $(\mathbf{q}, \mathbf{p})$ circuit to stability of the full circuit? Heuristically, this follows by decomposing the dynamics of $\mathbf{g}$ in terms of the different subspaces. We will not attempt to do this rigorously, as our goal is only to get a rough sense of how the system should behave. By the fact that their dimensions are strictly ordered, we know that $\text{span}(\mathbf{\Gamma}^\top) \supset \text{span}[(\mathbf{A}\mathbf{\Gamma})^\top]$. We can first exclude components outside $\text{span}(\mathbf{\Gamma}^\top)$, as they will be unchanged by the dynamics and do not affect the concentration estimate. Recalling the non-negativity constraint, components in $\text{span}(\mathbf{\Gamma}^\top) \setminus \text{span}[(\mathbf{A}\mathbf{\Gamma})^\top]$ will decay linearly to zero. Then, the components in $\text{span}[(\mathbf{A}\mathbf{\Gamma})^\top]$ should be controllable using the argument for the $(\mathbf{q}, \mathbf{p})$ space. Giving a fully rigorous analysis of the conditions under which the MAP fixed point is stable will be an interesting avenue for future investigation.

D Derivation for the Gaussian circuit model

To compare the result of the state dependent inhibition due to estimating the MAP for a Poisson noise model with the fixed inhibition resulting from Gaussian noise we build an alternative circuit model. This circuit still shares the separation into two cell types, but the granule cells $\mathbf{g}$ converge on the solution for a Gaussian noise model.

Our starting point is an isotropic Gaussian likelihood

$$\mathbf{s} | \mathbf{c} \sim \mathcal{N}(\mathbf{r}_0 + \mathbf{A}\mathbf{c}, \sigma^2 \mathbf{I}_{n_{\text{OSN}}}) \quad (\text{D.1})$$

variance σ^2 , and an exponential prior

$$\mathbf{c} \sim \text{Exp}(\boldsymbol{\lambda}). \quad (\text{D.2})$$

Gradient ascent on the resulting log-posterior over $\mathbf{c}$ leads to the dynamics:

$$\dot{\mathbf{c}}(t) = \frac{1}{\sigma^2} \mathbf{A}^\top [\mathbf{s} - (\mathbf{r}_0 + \mathbf{A}\mathbf{c})] - \boldsymbol{\lambda}, \quad (\text{D.3})$$

Distributing the code such that $\mathbf{c}(t) = \mathbf{\Gamma}\mathbf{g}$ and splitting the inference as we did in the Poisson case gives the circuit dynamics

$$\begin{aligned} \mathbf{c}(t) &= \mathbf{\Gamma}\mathbf{g}(t) \\ \tau_g \dot{\mathbf{g}}(t) &= \frac{1}{\sigma^2} (\mathbf{A}\mathbf{\Gamma})^\top \mathbf{p} - \mathbf{\Gamma}^\top \boldsymbol{\lambda} \\ \tau_p \dot{\mathbf{p}}(t) &= -\mathbf{p} + \mathbf{s} - (\mathbf{r}_0 + \mathbf{A}\mathbf{\Gamma}\mathbf{g}). \end{aligned} \quad (\text{D.4})$$

Unlike in the Poisson case, the inhibition onto the projection cells (mitral cells) is not gated by their activity, leading to the fixed offset shown in Fig. 2B.

E Extending the sampling circuit to incorporate an L_0 prior

In this Appendix, we consider the possibility of extending our circuit model to incorporate an L_0 spike-and-slab prior

$$p(c_i) = \varpi e^{-\lambda c_i} + (1 - \varpi)\delta(c_i), \quad (\text{E.1})$$

where ϖ is the probability that the odor is present and λ is the rate of the exponential prior on concentrations given that the odor is present. To sample from the resulting posterior using Langevin dynamics, we will follow the approach of Fang et al. [115], who developed a Langevin algorithm to perform sparse coding with this prior given a Gaussian likelihood. This approach only works in the sampling setting; it cannot be applied to MAP estimation because the Dirac mass at $c_i = 0$ means that the MAP estimate will always vanish.

In this approach, we define an auxiliary variable $\mathbf{u}$ that is mapped to concentration estimates $\mathbf{c}$ via element-wise soft thresholding:

$$\mathbf{c} = f(\mathbf{u}), \quad (\text{E.2})$$

where

$$f(u) = \begin{cases} 0 & u < u_0 \\ u - u_0 & u \geq u_0 \end{cases} \quad (\text{E.3})$$

is the soft-thresholding function for threshold

$$u_0 = -\frac{1}{\lambda} \log \varpi. \quad (\text{E.4})$$

We then posit the following Langevin dynamics for $\mathbf{u}$, given an observation $\mathbf{s}$:

$$\dot{\mathbf{u}}(t) = \left[\nabla_{\mathbf{c}} \log p(\mathbf{s} | \mathbf{c}) \right]_{\mathbf{c}=f(|\mathbf{u}|)} \odot \Theta(|\mathbf{u}| - \mathbf{u}_0) - \lambda \text{sign}(\mathbf{u}) + \boldsymbol{\eta}(t) \quad (\text{E.5})$$

with no constraint on the sign of $\mathbf{u}$, where $\boldsymbol{\eta}(t)$ is a white Gaussian noise process. Here, all nonlinearities are applied element-wise, and $\mathbf{u}_0 = u_0 \mathbf{1}$ is a vector with all elements equal to the threshold u_0 . Fang et al. [115] argue that the stationary distribution induced on $\mathbf{c} = f(\mathbf{u})$ by these dynamics should be the desired posterior with L_0 prior.

Using the Poisson likelihood gradient as computed before, we have

$$\dot{\mathbf{u}}(t) = \{\mathbf{A}^\top [\mathbf{s} \odot (\mathbf{r}_0 + \mathbf{A}f(|\mathbf{u}|)) - \mathbf{1}]\} \odot \Theta(|\mathbf{u}| - \mathbf{u}_0) - \lambda \text{sign}(\mathbf{u}) + \boldsymbol{\eta}(t) \quad (\text{E.6})$$

At this stage, we can see that the soft thresholding has picked out a preferred basis. If we apply the complete recipe in a way analogous to what we did in Appendix B by writing

$$\mathbf{u} = \mathbf{\Gamma} \mathbf{g}, \quad (\text{E.7})$$

we will have

$$\tau_g \dot{\mathbf{g}}(t) = \mathbf{\Gamma}^\top \left[\{\mathbf{A}^\top [\mathbf{s} \odot (\mathbf{r}_0 + \mathbf{A}f(|\mathbf{\Gamma} \mathbf{g}|)) - \mathbf{1}]\} \odot \Theta(|\mathbf{\Gamma} \mathbf{g}| - \mathbf{u}_0) - \lambda \text{sign}(\mathbf{\Gamma} \mathbf{g}) \right] + \boldsymbol{\xi}(t), \quad (\text{E.8})$$

where $\boldsymbol{\xi}(t)$ is a zero-mean Gaussian noise process with covariance

$$\mathbb{E}[\xi_j(t) \xi_{j'}(t')] = 2\tau_g \delta_{jj'} \delta(t - t'). \quad (\text{E.9})$$

Here, we cannot simply regroup terms; if we introduce an additional cell type $\mathbf{p}$ as before to compute the division we will have a sort of effective weight matrix

$$\sum_k \Gamma_{ki} A_{jk} \Theta(|u_k| - u_0) \quad (\text{E.10})$$

for the $\mathbf{p}$ -to- $\mathbf{g}$ connections, and an entirely different coupling $\mathbf{A}f(|\mathbf{\Gamma} \mathbf{g}|)$ for the $\mathbf{g}$ -to- $\mathbf{p}$ connections. Concretely, introducing $\mathbf{p}$ as before, we have the system

$$\mathbf{c}(t) = f(|\mathbf{\Gamma} \mathbf{g}(t)|) \quad (\text{E.11})$$

$$\tau_g \dot{\mathbf{g}}(t) = \mathbf{\Gamma}^\top \left[[\mathbf{A}^\top (\mathbf{p} - \mathbf{1})] \odot \Theta(|\mathbf{\Gamma} \mathbf{g}| - \mathbf{u}_0) - \lambda \text{sign}(\mathbf{\Gamma} \mathbf{g}) \right] dt + \boldsymbol{\xi}(t) \quad (\text{E.12})$$

$$\tau_p \dot{\mathbf{p}}(t) = \mathbf{s} - \mathbf{p} \odot [\mathbf{r}_0 + \mathbf{A}f(|\mathbf{\Gamma} \mathbf{g}|)]. \quad (\text{E.13})$$

This issue illustrates the limitations of the simple, linear approach to distributing the neural code used in Appendix B. In particular, the nonlinearity in some sense picks out a preferred basis, meaning that we can no longer perform a simple linear change of coordinates to distribute the code.

F Experimental details and affinity matrix fitting

Here we briefly describe the experimental procedures used to collect the data used to fit the parameters of the affinity matrix $\mathbf{A}$ in Fig. 1. The full experimental methods are described in the paper that first presented the data, Zak et al. [94]. All the experiments were performed in accordance with the guidelines set by the National Institutes of Health and approved by the Institutional Animal Care and Use Committee at Harvard University.

F.1 *In vivo* recordings from mouse OB

Adult (> 8 weeks) OMP-GCaMP3 mice of both sexes were used in this study. A craniotomy was performed to provide optical access to olfactory sensory neuron axon terminals in both olfactory bulbs. A custom-built two-photon microscope was used for *in vivo* imaging. Images were acquired at 16-bit resolution and 4-8 frames/s. The pixel size was $0.6 \mu m$ and the fields of view were $720 \times 720 \mu m$. Monomolecular odorants (Allyl butyrate, Ethyl valerate, Methyl tiglate, and Isobutyl propionate) were used as stimuli and delivered by a 16-channel olfactometer controlled by custom-written software in LabView. For the odorant concentration series, the initial odorant concentration was between $0.08\% - 80\%(v/v)$ in mineral oil and further diluted 16 times with air. The relative odorant concentrations were measured by a photoionization detector, then normalized to the largest detected signal for each odorant. For all experiments, the airflow to the animal was held constant at $100 mL/min$, and odorants were injected into a carrier stream. Each odorant concentration was delivered 2–6 times in pseudorandom order.

Images were processed using both custom and available MATLAB scripts. Motion artifact compensation and denoising were done using *NoRMcorre* [138]. The $\Delta F/F$ signal was calculated by finding the peak signal following odorant onset and averaging with the two adjacent points. To account for changes in respiration and anesthesia depth, correlated variability was corrected for [139]. Thresholds for classifying responding ROIs were determined from a noise distribution of blank (no odorant) trials from which three standard deviations were used for responses.

F.2 Gamma distribution fitting

In order to fit the affinity matrix $\mathbf{A}$ to the experimentally recorded data, we normalized the response of each glomerulus to its maximum response across the panel of 32 odorants. We then vectorized the resulting matrix and fitted a Gamma distribution using the `gamfit` function in MATLAB. As mentioned in the main text, this resulted in a $\text{Gamma}(0.37, 0.36)$ distribution.

G Numerical methods and supplemental figures

All numerical simulations were performed using MATLAB 9.13 (R2022b, The MathWorks, Natick, MA, USA) either on desktop workstations (CPU: Intel i9-9900K or Xeon W-2145, 64GB RAM) or on the Harvard University FASRC Cannon HPC cluster (<https://www.rc.fas.harvard.edu>). Our simulations were not computationally intensive, and required around that 6000 CPU-hours of total compute time.

G.1 State-dependent inhibition simulations

For the simulations to highlight the signatures of state-dependent inhibition of the Poisson network and compare its behavior with the experimental observation by Arevian et al. [95] we use a reduced network with $n_{OSN} = 2$ and $n_g = 10$. We chose this reduced circuit both to reduce computational costs and to match the experimental parameters of the *in-vitro* experiment. The other parameters of the simulation were as follows: $\tau_p = 0.02$, $\tau_g = 0.03$, $r_0 = 10$, $\lambda = 2$, $dt = 1e^{-4}$. The simulation ran for $600 ms$, stimulation was applied starting at $t_{start} = 100 ms$ and ended at $t_{end} = 500 ms$. We stimulated the principal mitral cell MC_A with 80 equally spaced values in $s_A \in [1, 400]$. When the second mitral cell was active, its stimulation was set at $s_B = 80$. The entries of the Γ matrix were sampled from a normal distribution on which we applied a mask such that only 25% of its entries were non-zero. We sampled 32 'pairs' of cells by resampling the Γ matrix for each 'pair' of cells. The simulation for the Gaussian circuit used the same parameters except the dynamics followed those

from equation D.4. For plotting purposes, we normalized the range of firing in Figure 2B to the maximum firing rate for each circuit model.

G.2 Capacity simulations

In our capacity simulations, we use $n_{\text{OSN}} = 300$ and $n_{\text{odor}} = 500, 1000, \text{ or } 2000$. We take the odor stimulus to be a rectangular pulse, with varying numbers of randomly-selected odors appearing at concentration $c_j = 40$. As elsewhere, we take $\tau_p = 20$ ms [91] and $\tau_g = 30$ ms [92] to match experiment. We set the baseline rate to be $r_0 = 1$ and the prior mean to be $\lambda = 1$; based on some experimentation our results appear relatively insensitive to small variations in these choices. We integrate the MAP circuit dynamics (3) using the forward Euler method with timestep $\Delta t = 10^{-4}$ s. To determine which odors the model estimates as being present at a given timepoint, we simply check which concentration estimates exceed 20 at that time.

As discussed in the main text, we consider three variants of the MAP circuit, defined by different choices of the matrix Γ :

- For the **one-to-one** code, we let $n_g = n_{\text{odor}}$, and simply set $\Gamma = \mathbf{I}_{n_{\text{odor}}} / N(A_{ij})$.
- For the **naïvely distributed** code, we let $n_g = 5n_{\text{odor}}$, and choose Γ as follows: We sample a random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_g}$ with $\mathbf{Q}\mathbf{Q}^\top = \mathbf{I}_{n_{\text{odor}}}$ by drawing a Gaussian matrix and orthogonalizing its rows. Then, we define $\Gamma = \mathbf{Q} / N\{ |(\mathbf{A}\mathbf{Q})_{ij}| \}$.
- For the **geometry-aware** code, we let $n_g = 5n_{\text{odor}}$, and choose Γ as follows: We sample a random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_g}$ with $\mathbf{Q}\mathbf{Q}^\top = \mathbf{I}_{n_{\text{odor}}}$ by drawing a Gaussian matrix and orthogonalizing its rows. Then, we compute an approximate inverse square root of the low-rank matrix $\mathbf{A}^\top \mathbf{A}$ as $\mathbf{B} = (\mathbf{A}^\top \mathbf{A} + a\mathbf{I}_{n_{\text{odor}}})^{-1/2}$ for a small positive regularizing constant a . In our simulations, we set $a = 0.5$; we find empirically that our results are not substantially sensitive to small variations in a . Finally, we let $\Gamma = \mathbf{B}\mathbf{Q} / N\{ |(\mathbf{A}\mathbf{B}\mathbf{Q})_{ij}| \}$.

where N is a normalization function defined as

$$N(A) = \max(A) \cdot \frac{\sqrt{n_g}}{C}$$

with $C = 50$ identified as a reasonable choice. Our normalization convention for Γ is motivated by the idea that, in biology, the strength of individual synapses should be bounded. Moreover, changing the overall scale of the synaptic weights in our model corresponds to changing the effective time constant of the dynamics, which can produce a trivial speedup.




Figure G.1: **A.** Fraction of odors correctly detected within 100 ms (*left*), 200 ms (*center*), and 1 s (*right*) after odor onset as a function of the number of odors present, with $n_{\text{odor}} = 500$ possible odors, for models with **one-to-one**, **naïvely distributed**, and **geometry-aware** codes. **B.** As in **A**, but for $n_{\text{odor}} = 1000$. This matches Fig. 3B. **C, D, E.** As in **A**, but for $n_{\text{odor}} = 2000$, 4000, or 8000 possible odorants, respectively. Shaded patches show ± 1.96 SEM over realizations throughout.




Figure G.2: **A.** Heatmap with overlaid smoothed contours of the fraction of odors correctly detected as a function of number of present odors and time window out of a panel of $n_{\text{odor}} = 500$ possible odors for models with **one-to-one** (*left*), **naïvely distributed** (*center*), and **geometry-aware** (*right*) codes. **B.** As in **A**, but for $n_{\text{odor}} = 1000$. This matches Fig. 3B. **C, D, E.** As in **A**, but for $n_{\text{odor}} = 2000$, 4000, or 8000 possible odorants, respectively.

G.3 Sampling simulations

In our sampling simulations, we use $n_{\text{OSN}} = 300$, $n_{\text{odor}} = 500$ or 1000 , and $n_g = 5n_{\text{odor}}$. Here, the odor stimulus was composed of two rectangular pulses: At time $t_{\text{low}} = 0$ s, five “low” odors appear at concentration 10, and remain present at that concentration until time $t = 2$ s. Then, at time $t_{\text{high}} = 1$ s, five “high” odors appear at concentration 40, and remain present until time $t_{\text{off}} = 2$ s. As in our capacity simulations, we take $\tau_p = 20$ ms [91] and $\tau_g = 30$ ms [92] to match experiment, and we set the baseline rate to be $r_0 = 1$ and the prior mean to be $\lambda = 1$. As for our capacity results, our results appear relatively insensitive to small variations in these choices. We integrate the sampling circuit dynamics (4) using the forward Euler-Maruyama method with timestep $\Delta t = 10^{-5}$ s.

As in our capacity simulations, in Fig. 4 we choose Γ as follows:

- For the **one-to-one** code, we let $n_g = n_{\text{odor}}$, and simply set $\Gamma = \mathbf{I}_{n_{\text{odor}}} / \max(A_{ij})$.
- For the **naïvely distributed** code, we let $n_g = 5n_{\text{odor}}$, and choose Γ as follows: We sample a random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_g}$ with $\mathbf{Q}\mathbf{Q}^\top = \mathbf{I}_{n_{\text{odor}}}$ by drawing a Gaussian matrix and orthogonalizing its rows. Then, we define $\Gamma = \mathbf{Q} / \max\{ |(\mathbf{A}\mathbf{Q})_{ij}| \}$.
- For the **geometry-aware** code, we let $n_g = 5n_{\text{odor}}$, and choose Γ as follows: We sample a random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_g}$ with $\mathbf{Q}\mathbf{Q}^\top = \mathbf{I}_{n_{\text{odor}}}$ by drawing a Gaussian matrix and orthogonalizing its rows. Then, we compute an approximate inverse square root of the low-rank matrix $\mathbf{A}^\top \mathbf{A}$ as $\mathbf{B} = (\mathbf{A}^\top \mathbf{A} + a\mathbf{I}_{n_{\text{odor}}})^{-1/2}$ for a small positive regularizing constant a . In our simulations, we set $a = 0.5$; we find empirically that our results are not substantially sensitive to small variations in a . Finally, we let $\Gamma = \mathbf{B}\mathbf{Q} / \max\{ |(\mathbf{A}\mathbf{B}\mathbf{Q})_{ij}| \}$.

In Fig. 4A, we smooth the concentration estimate timeseries using a 100 ms moving average. In Fig. 4B-C, we show cumulative estimates of the mean and variance. Concretely, given a concentration timeseries $\hat{c}_j(t)$, we estimate the mean and variance as

$$\mu_j(\tau) = \frac{\Delta t}{\tau} \sum_{t_{\text{low}} \leq t \leq t_{\text{low}} + \tau} \hat{c}_j(t) \quad (\text{G.1})$$

and

$$\sigma_j^2(\tau) = \frac{\Delta t}{\tau} \sum_{t_{\text{low}} \leq t \leq t_{\text{low}} + \tau} \hat{c}_j(t)^2 - \mu_j(\tau)^2, \quad (\text{G.2})$$

respectively, where we assume τ and t are integer multiples of Δt . In Fig. 4D-E, we do the same except for times after t_{high} . In both cases, baselines were obtained by running the naïve sampling algorithm for 10^8 steps with a burn-in period of 10^7 steps. In Supp. Fig. G.3, we reproduce Fig. 4 showing estimates for individual odorants as well as the mean across odorants.

In Supp. Fig. G.4, we show a preliminary experiment with an alternative, mostly non-negative choice for Γ . This gives the following three models:

- For the **one-to-one** code, we let $n_g = n_{\text{odor}}$, and simply set $\Gamma = \mathbf{I}_{n_{\text{odor}}} / \max(A_{ij})$.
- For the **naïvely distributed** code, we let $n_g = 5n_{\text{odor}}$, and choose Γ as follows: We sample a sparse, non-negative random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_g}$ with entries that are non-zero with density 0.15, and the non-zero entries are drawn uniformly on $[0, 1]$. Then, we define $\Gamma = \mathbf{Q} / \max\{ |(\mathbf{A}\mathbf{Q})_{ij}| \}$.
- For the **geometry-aware** code, we let $n_g = 5n_{\text{odor}}$, and choose Γ as follows: We sample a sparse, non-negative random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_g}$ with entries that are non-zero with density 0.15, and the non-zero entries are drawn uniformly on $[0, 1]$. Then, we compute an approximate inverse square root of the low-rank matrix $\mathbf{A}^\top \mathbf{A}$ as $\mathbf{B} = (\mathbf{A}^\top \mathbf{A} + a\mathbf{I}_{n_{\text{odor}}})^{-1/2}$ for a small positive regularizing constant a . In our simulations, we set $a = 0.5$; we find empirically that our results are not substantially sensitive to small variations in a . Finally, we let $\Gamma = \mathbf{B}\mathbf{Q} / \max\{ |(\mathbf{A}\mathbf{B}\mathbf{Q})_{ij}| \}$.

This yields a **naïvely distributed** synaptic weight matrix $\mathbf{A}\Gamma$ that is entirely non-negative, while the **geometry-aware** weight matrix has a small number of negative elements due to the fact that the inverse of a non-negative matrix need not be non-negative. We see that the behavior of this circuit is




Figure G.3: Fast uncertainty estimation using Langevin sampling of the posterior, showing estimates for individual odorants. Here, we use a simple concentration estimation task in which 5 randomly-selected ‘low’ odorants out of a panel of 1000 appear at concentration 10 at time 0 s, and then a further 5 randomly-selected ‘high’ odorants appear at concentration 40 at time 1 s. **A.** Smoothed timeseries of instantaneous concentration estimates for **low**, **high**, and **background** odorants, for models with one-to-one (*top*), naively distributed (*middle*), and geometry-aware (*bottom*) codes. Background odorant estimates are shown as mean $\pm$ standard deviation over odorants. Dashed lines show true concentrations over time. **B.** Cumulative estimates of concentration mean for low (*top*) and high (*bottom*) odorants after the onset of the low odorants for **one-to-one**, **naively distributed**, and **geometry-aware** codes. Black lines indicate baseline estimates of the posterior mean. Thick colored lines indicate means over odorants, while thin lines show traces for individual odorants. **C.** As in **B**, but for the estimated variance. **D.** As in **B**, but after the onset of high odorants. **E.** As in **C**, but after the onset of high odorants. See Appendix G for details of our numerical methods.

similar to that observed in Fig. 4. As mentioned in the Discussion, an important topic for future work will be to devise a method to choose Γ that yields a fully non-negative, sparse synaptic weight matrix $A\Gamma$.

H Code and data availability

All code and data required to reproduce the figures presented is available under an MIT License at <https://github.com/Pehlevan-Group/olfaction-geometry/>.




Figure G.4: Langevin sampling of the posterior using sparse non-negative distributed codes. Here, we use a simple concentration estimation task in which 5 randomly-selected ‘low’ odors out of a panel of 500 appear at concentration 10 at time 0 s, and then a further 5 randomly-selected ‘high’ odors appear at concentration 40 at time 1 s. This figure replicates Fig. 4, except that the code is distributed using a sparse non-negative random matrix rather than an orthogonal matrix. Smoothed timeseries of instantaneous concentration estimates for low, high, and background odors, for models with one-to-one (*top*), naïvely distributed (*middle*), and geometry-aware (*bottom*) codes. Background odor estimates are shown as mean $\pm$ standard deviation over odors. Dashed lines show true odor concentrations over time. **B.** Cumulative estimates of odor concentration mean for low (*top*) and high (*bottom*) odors after the onset of the low odors for one-to-one, naïvely distributed, and geometry-aware codes. Black lines indicate baseline estimates of the posterior mean. Thick lines indicate means over odors, and thin lines individual odors. **C.** As in **B**, but for the estimated variance. **D.** As in **B**, but after the onset of high odors. **E.** As in **C**, but after the onset of high odors. See Appendix G for details of our numerical methods.