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], **Joseph D. Zak** [5], **Venkatesh N. Murthy** [1,3,‡,✉], **Cengiz Pehlevan** [1,4,6,‡,✉]

[1]Center for Brain Science, Harvard University; [2]Department of Physics, Harvard University; [3]Department of Molecular and Cellular Biology, Harvard University; [4]John A. Paulson School of Engineering and Applied Sciences, Harvard University; [5]Department of Biological Sciences, University of Illinois at Chicago; [6]Kempner Institute for the Study of Natural and Artificial Intelligence, Harvard University

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.

✉ **For correspondence:**
jzavatoneveth@g.harvard.edu,
paul_masset@fas.harvard.edu,
vnmurthy@fas.harvard.edu,
cpehlevan@seas.harvard.edu

[†]JAZ-V and PM contributed equally to this work.

[‡]VNM and CP jointly supervised this work.

Funding: This work was supported by NSF grant DMS-2134157 to CP, NIH grants 5R01DC017311 and 5R01DC016289 to VNM, and NTT Research award A47994 to VNM. PM was partially supported by a grant from the Harvard Mind Brain Behavior Interfaculty Initiative. JDZ was partially supported by NIH grant K99DC017754.

Competing interests: The authors declare no competing interests.

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 (Uexküll, 1909; Yong, 2022). Throughout evolution, organisms have developed specialized sensory mechanisms to extract specific information about the 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 (Hubel and Wiesel, 1959; Hubel and Wiesel, 1962; Kandler, Clause, and Noh, 2009). 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 (Attneave, 1954; Barlow, 2012; Srinivasan et al., 1982; Linsker, 1988; Hateren, 1992; Atick and Redlich, 1992; Olshausen and Field, 1996; Bell and Sejnowski, 1997; Brunel and Nadal, 1998; Rao and Ballard, 1999; Butts and Goldman, 2006).

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 (Murthy, 2011; Lee et al., 2022; Qian et al., 2023). In the absence of geometric intuitions, the principles of compressed sensing (CS) have emerged as an alternative paradigm for understanding olfactory coding (Koulakov and Rinberg, 2011; Grabska-Barwińska, Beck, et al., 2013; Grabska-Barwińska, Barthelmé, et al., 2017; Tootoonian and Lengyel, 2014; Kepple, Cazakoff, et al., 2018; Kepple, Giaffar, et al., 2019; Qin et al., 2019; Krishnamurthy et al., 2022). 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 (Baraniuk, 2007; Donoho, 2006; Candès and Tao, 2006; Candès, Romberg, and Tao, 2006; Raginsky, Willett, et al., 2010; Raginsky, Jafarpour, et al., 2011). 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* (Khan, Parthasarathy, and Bhalla, 2010; Ma, Qiu, et al., 2012; Isaacson, 2010; Grabska-Barwińska, Barthelmé, et al., 2017). 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 (Ma, Chen, and Fox, 2015; Masset, Zavatone-Veth, et al., 2022; Amari, 1998; Girolami and Calderhead, 2011; Chung and Abbott, 2021).

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 (§A neural circuit architecture for Poisson compressed sensing). Importantly, this mapping goes beyond basic counting of cell types; it includes detailed biological features like symmetric coupling and state-dependent inhibition (§Biological interpretation and predictions of the circuit model).
- 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 (§Geometry, speed, and capacity). This fast inference is enabled by considering the geometry of the olfactory receptor code. This consideration leads to a distributed odor coding, resolving a major problem with previous CS circuit models.
- We extend our circuit model to allow Bayesian inference of uncertainty in odor concentrations (§Fast sampling for uncertainty estimation).

In total, our results show how normative modeling can help us understand how distributed neural codes can enable efficient computation in biological circuits.

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; Lledo, Gheusi, and Vincent, 2005; Ache and Young, 2005; Su, Menuz, and Carlson, 2009; Uchida, Poo, and Haddad, 2014), and this distinctive circuit structure is thought to play a key role in the computational function of the olfactory bulb (Nagayama, Homma, and Imamura, 2014; Burton, 2017; Wu, Yu, and Komiyama, 2020). 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 (Ache and Young, 2005). Importantly, most ORs have broad affinity profiles, and the OSN code for odor identity is combinatorial (Malnic et al., 1999). In contrast to the immune system's highly adaptable chemical recognition capabilities arising from somatic recombination (Murphy, Weaver, and Janeway, 2017), ORs are hard coded into the genome as single genes (Buck and Axel, 1991), and therefore can only change over evolutionary timescales (Niimura and Nei, 2007; Niimura, Matsui, and Touhara, 2014). Some adaptation of expression levels across the receptors is possible (Teşileanu et al., 2019), 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 (Burton, 2017). 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 (Egger and Urban, 2006).

In our model, we will focus on a particular task: odor component identification within a single sniff (Kepecs, Uchida, and Mainen, 2005; Ackels et al., 2021). This computational problem differs from many experimental tasks focusing on discrimination between two odorants (Uchida and Mainen, 2003), which underlie most scientific work on rodent olfactory decision making (Uchida and Mainen, 2003; Spors et al., 2012; Reinert and Fukunaga, 2022). Here, the goal is to identify the components of a complex olfactory scene (Rokni et al., 2014; Lebovich et al., 2021; Li et al., 2023; Berners-Lee et al., 2023). 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 (Igarashi et al., 2012; Chae et al., 2022). 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 (Gupta, Albeanu, and Bhalla, 2015).

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) (Grabska-Barwińska, Beck, et al., 2013; Grabska-Barwińska, Barthelmé, et al., 2017; Tootoonian and Lengyel, 2014; Zhang and Sharpee, 2016; Zwicker, Murugan, and Brenner, 2016; Kepple, Cazakoff, et al., 2018; Kepple, Giaffar, et al., 2019; Qin et al., 2019; Singh, Tchernookov, and Balasubramanian, 2021; Krishnamurthy et al., 2022). Classical CS theory shows that sparse high-dimensional signals can be recovered from a small number of random projections (Baraniuk, 2007; Donoho, 2006; Candès and Tao, 2006; Candès, Romberg, and Tao, 2006; Raginsky, Willett, et al., 2010; Raginsky, Jafarpour, et al., 2011). Inspired by these results, previous works have used the principles of CS to build circuit models for olfactory coding (Grabska-Barwińska, Beck, et al., 2013; Grabska-Barwińska, Barthelmé, et al., 2017; Tootoonian and Lengyel, 2014; Zhang and Sharpee, 2016; Zwicker, Murugan, and Brenner, 2016; Kepple, Cazakoff, et al., 2018; Kepple, Giaffar, et al., 2019; Qin et al., 2019; Singh, Tchernookov, and Balasubramanian, 2021; Krishnamurthy et al., 2022). 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 Grabska-Barwińska, Beck, et al., 2013; Grabska-Barwińska, Barthelmé, et al., 2017). Instead, OSN activity is better captured by a Poisson noise model (Duchamp-Viret et al., 2005; Kostal, Lansky, and Rospars, 2007). Some theoretical guarantees for Poisson CS are known, but the situation is less well-understood than in the Gaussian case (Raginsky, Willett, et al., 2010; Raginsky, Jafarpour, et al., 2011; Dytso, Fauß, and Poor, 2020; Jin et al., 2012).

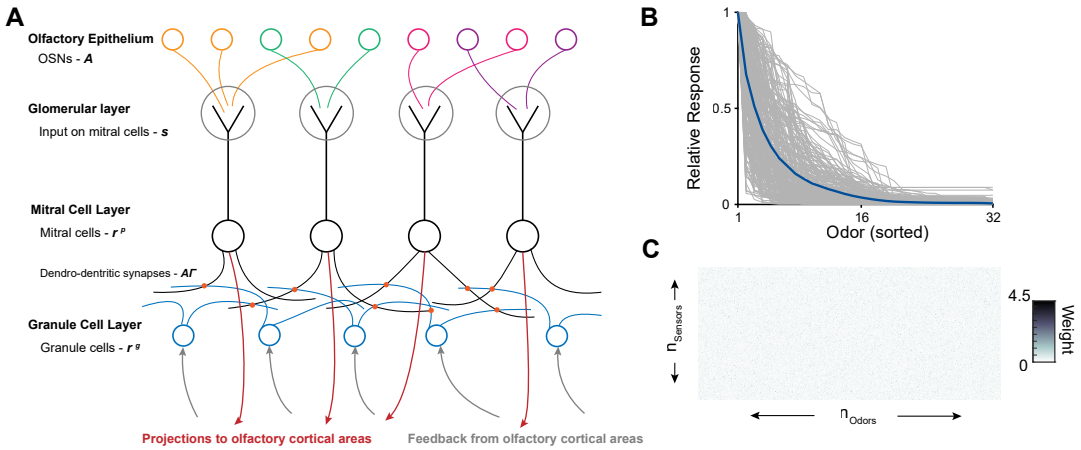

Figure 1. Circuit architecture of the mammalian olfactory bulb. **A.** Outline of the anatomy of OB circuits. **B.** Fit of the responses of 228 OSN glomerular responses to a panel of 32 odorants. (Gray line, response of single glomeruli, Blue line, Fitted average response.) See Appendix 6 for details of how these data were collected. **C.** Affinity matrix generated from the data-driven model ensemble.

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 $\boldsymbol{\lambda} \in \mathbb{R}_+^{n_{\text{odor}}}$:[1]

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

Given this likelihood and prior, we construct a neural circuit to compute the maximum *a posteriori* (MAP) estimate of the concentration $\mathbf{c}$. Here, we sketch the derivation, deferring some details to Appendix 2. Our starting point is the gradient ascent equation

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

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, Beck, et al. (2013).

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 (§Related work and review of the olfactory sensing problem). To distribute the code, we instead use a projected setup where the firing rates $\mathbf{g} \in \mathbb{R}^{n_{\text{g}}}$ of the neurons are mapped to the concentration estimate $\mathbf{c}$ through a matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{n_{\text{odor}} \times n_{\text{g}}}$: $\mathbf{c}(t) = \boldsymbol{\Gamma}\mathbf{g}(t)$. Even if $\boldsymbol{\Gamma}$ is non-square—in particular, if $n_{\text{g}} > n_{\text{odor}}$—so long as $\boldsymbol{\Gamma}\boldsymbol{\Gamma}^\top$ is positive-definite and the rates follow the dynamics $\tau_{\text{g}}\dot{\mathbf{g}}(t) = (\mathbf{A}\boldsymbol{\Gamma})^\top[\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g}) - \mathbf{1}] + \boldsymbol{\Gamma}^\top[(\boldsymbol{\alpha} - \mathbf{1}) \oslash (\boldsymbol{\Gamma}\mathbf{g}) - \boldsymbol{\lambda}]$ for some time constant τ_{g}, the concentration estimates will still converge to the MAP. This corresponds to preconditioned gradient ascent (Amari, 1998; Martens, 2020). 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 $\mathbf{g} \in \mathbb{R}_+^{n_{\text{g}}}$ and choose $\boldsymbol{\Gamma}$ to be positivity-preserving.

[1] See Appendix 1 for a detailed description of our notational conventions.

These dynamics include two divisive non-linearities, which can be challenging to implement in biophysical models of single neurons (Holt and Koch, 1997). Using the approach proposed by Chalk et al. (2017), 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

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

for cell-type-specific time constants τ_{g}, τ_{p}, and τ_{z}, where $\odot$ denotes elementwise multiplication. In the limit $\tau_{\text{p}}, \tau_{\text{z}} \ll \tau_{\text{g}}$, 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 (Chalk et al., 2017) (see Appendix 3 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 $\mathbb{R}_+^{n_{\text{p}}}$, 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.

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 (Marr and Poggio, 1976; Zavatone-Veth, Badwan, and Clark, 2020).

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.

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

The first salient feature of the dynamics (3) is that the cell types $\mathbf{g}$ and $\mathbf{p}$ 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 $\mathbf{A}\boldsymbol{\Gamma}$ of the connections from $\mathbf{g}$ to $\mathbf{p}$ neurons mirror exactly the weights $(\mathbf{A}\boldsymbol{\Gamma})^\top$ of the connections from $\mathbf{p}$ to $\mathbf{g}$ 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 (Crick, 1989; Guerguiev, Lillicrap, and Richards, 2017; Lillicrap et al., 2020)? 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; Urban and Arevian, 2009; Burton, 2017; Egger and Urban, 2006).

This interpretation accounts for the interactions between the cell types $\mathbf{g}$ and $\mathbf{p}$ in the dynamics (3), but what about the cell type $\mathbf{z}$? These cells give direct excitatory input to the granule cells $\mathbf{g}$ with weights $\boldsymbol{\Gamma}^\top$, 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 (Davis and Macrides, 1981; Boyd et al., 2012; Markopoulos

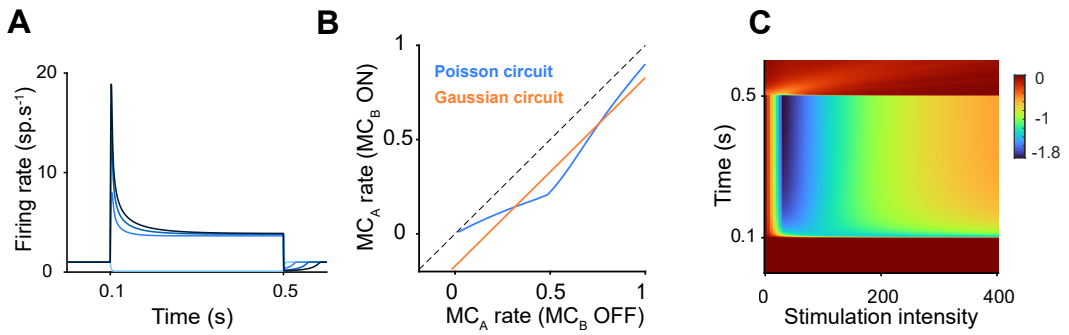

Figure 2. State-dependent inhibition of mitral cells. **A.** Upon stimulation, mitral cells exhibit a transient burst of activity followed by relaxation to a plateau. Darker color indicates stronger stimulation. **B.** In the circuit for Poisson CS we propose (blue line), the inhibition due to the activation 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, Kapoor, and Urban (2008) 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. Arevian, Kapoor, and Urban (2008) Fig. 3B.

et al., 2012). 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 $\boldsymbol{\alpha} = \mathbf{1}$, in which case it reduces to an exponential and feedback is not explicitly required.

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 (Baraniuk, 2007; Donoho, 2006; Candès and Tao, 2006; Candès, Romberg, and Tao, 2006; Raginsky, Willett, et al., 2010; Raginsky, Jafarpour, et al., 2011; Krishnamurthy et al., 2022). 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 (del Mármol, Yedlin, and Ruta, 2021; Buck and Axel, 1991; Niimura and Nei, 2007). 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 6 for details; these data were previously published in Zak et al., 2020). 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 6). We then define our ensemble of sensing matrices $\mathbf{A}$ by drawing their elements as independently and identically distributed $\mathrm{Gamma}(0.37, 0.36)$ random variables. An example matrix drawn from this ensemble is shown in Fig. 1C.

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 (Kepple, Giaffar, et al., 2019; Koulakov and Rinberg, 2011; Kepple, Cazakoff, et al., 2018). 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 (Chalk et al., 2017). 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 4; Rao and Ballard, 1999). In Fig. 2A, we

show example timecourses of model mitral cell activity following the onset of a stimulus. Consistent with experimentally-measured responses (Cury and Uchida, 2010; Shusterman et al., 2011), a sharp transient response at the onset of stimulation is followed by decay to a low level of tonic activity (Fig. 2A).

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 (Arevian, Kapoor, and Urban, 2008). 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 (Arevian, Kapoor, and Urban, 2008).

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, Kapoor, and Urban (2008)'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 4). 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 Arevian, Kapoor, and Urban, 2008 Fig. 3b).

Geometry, speed, and capacity

We have argued that the model introduced in §A neural circuit architecture for Poisson compressed sensing 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 (Donoho and Tanner, 2009; Maleki and Donoho, 2010). 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).

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 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 (Ma, Chen, and Fox, 2015; Masset, Zavatone-Veth, et al., 2022; Amari, 1998; Girolami and Calderhead,

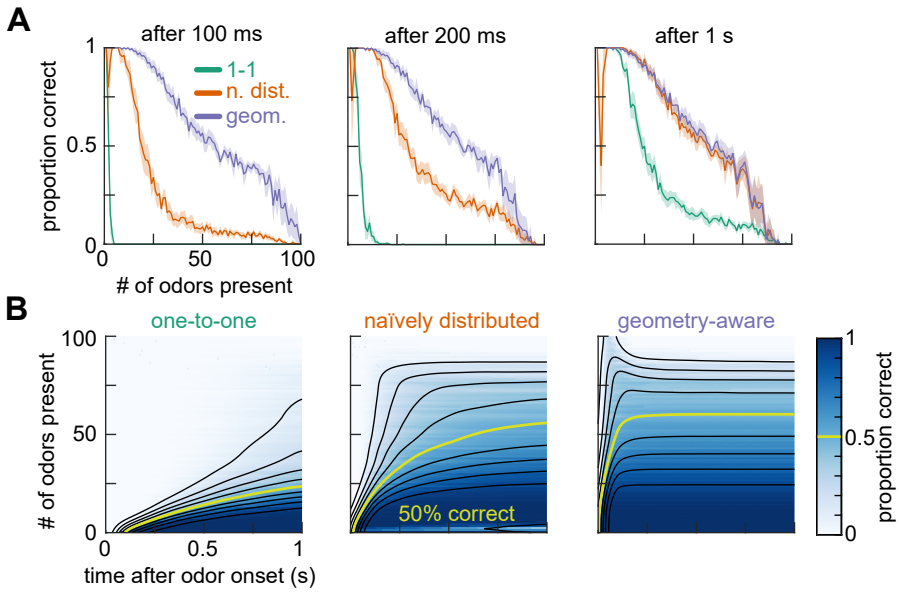

Figure 3. Fast detection of many odorants. **A**. Fraction of odorants correctly detected within 100 ms (*left*), 200 ms (*center*), and 1 s (*right*) after odorant onset as a function of the number of odorants present, for models with one-to-one, naïvely distributed, and geometry-aware codes. **B**. Heatmap with overlaid smoothed contours of correct detection fraction as a function of the number of present odorants and time window for models with one-to-one (*left*), naïvely distributed (*center*), and geometry-aware (*right*) codes. See Appendix 7 for details of our numerical methods.

2011) (see Appendix 7 for details). 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.

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 (Knill and Pouget, 2004; Körding and Wolpert, 2004; Fiser et al., 2010; Ott, Masset, and Kepecs, 2018; Masset, Zavatone-Veth, et al., 2022; Reddy, Murthy, and Vergassola, 2022). 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 2, 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 §A neural circuit architecture

for Poisson compressed sensing up to the addition of Gaussian noise to the granule cell dynamics:

$$\mathbf{c}(t) = \mathbf{\Gamma}\mathbf{g}(t) \qquad \tau_{\mathrm{g}}\dot{\mathbf{g}}(t) = (\mathbf{A}\mathbf{\Gamma})^{\top}(\mathbf{p} - \mathbf{1}) + \mathbf{\Gamma}^{\top}(\mathbf{z} - \boldsymbol{\lambda}) + \boldsymbol{\xi}(t)$$
$$\tau_{\mathrm{p}}\dot{\mathbf{p}}(t) = \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A}\mathbf{\Gamma}\mathbf{g}) \qquad \tau_{\mathrm{z}}\dot{\mathbf{z}}(t) = \boldsymbol{\alpha} - \mathbf{1} - \mathbf{z} \odot \mathbf{c}. \tag{4}$$

Here, $\boldsymbol{\xi}(t)$ is a vector of n_{g} independent zero-mean Gaussian noise processes with covariance

$$\mathbb{E}[\xi_j(t)\xi_{j'}(t')] = 2\tau_{\mathrm{g}}\delta_{jj'}\delta(t - t'), \tag{5}$$

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 (Ma, Chen, and Fox, 2015; Girolami and Calderhead, 2011). 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 $n_{\mathrm{g}} > n_{\mathrm{odor}}$—and is important for biological realism if we interpret the sampling noise as resulting from fluctuations in membrane potential due to synaptic noise (Calvin and Stevens, 1967).

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 (Rokni et al., 2014; Lebovich et al., 2021). 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 (Neal, 1993), 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.

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 (Marr and Poggio, 1976; Zavatone-Veth, Badwan, and Clark, 2020). 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 7). 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 (Nagayama, Homma, and Imamura, 2014; Burton, 2017). 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 (Zhu and Rozell, 2015). As a first step, in Supp. Fig. 3 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.

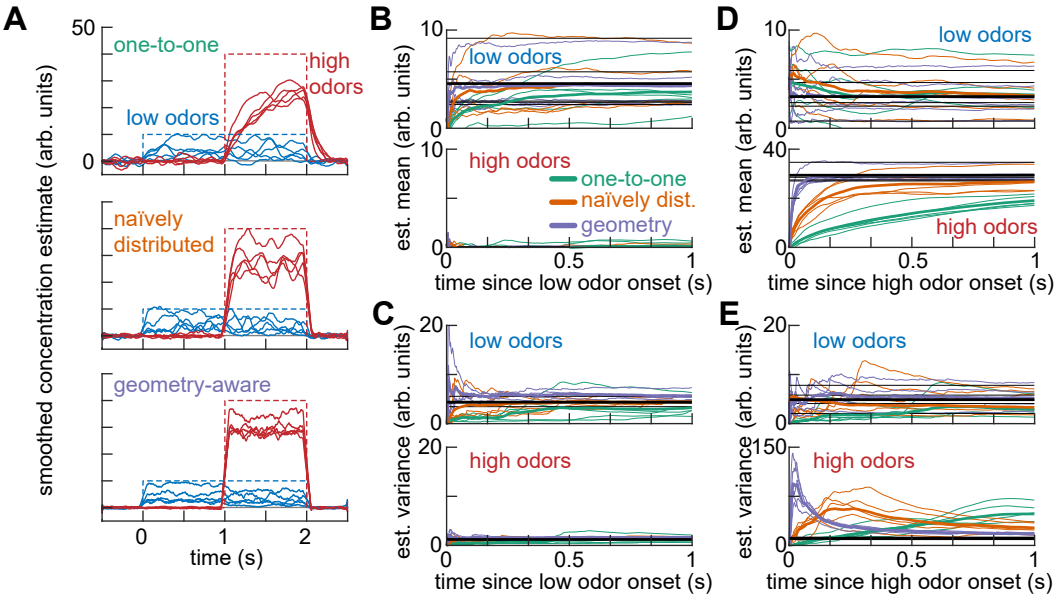

Figure 4. Fast uncertainty estimation using Langevin sampling of the posterior. Here, we use a simple concentration estimation task in which 5 randomly-selected 'low' odorants out of a panel of 1000 appear at concentration 10 at time 0 s, and then a further 5 randomly-selected 'high' odorants appear 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*), naïvely distributed (*middle*), and geometry-aware (*bottom*) codes. Background odorant estimates are shown as mean ± 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, naïvely distributed, and geometry-aware codes. Black lines indicate baseline estimates of the posterior mean. Thick lines indicate means over odorants, and thin lines 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 7 for details of our numerical methods.

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 (Reddy, Zak, et al., 2018; Zak et al., 2020; Gorur-Shandilya et al., 2017), which are known to affect the performance of Gaussian CS models for olfaction (Qin et al., 2019; Kadakia and Emonet, 2019). 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 (Savin and Denève, 2014; Masset, Zavatone-Veth, et al., 2022). 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 (Masset, Zavatone-Veth, et al., 2022). 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 5, 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 (Fang et al., 2022).

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 (Teşileanu et al., 2019; Tsukahara et al., 2021). Some past works have sought to incorporate plasticity of the mitral-granule cell synapses into decoding models (Hiratani and Latham, 2020; Kepple, Cazakoff, et al., 2018), tying into a larger body of research on how plasticity can enable flexible feature extraction in olfaction (Chapochnikov, Pehlevan, and Chklovskii, 2021; Pehlevan, Genkin, and Chklovskii, 2017). 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. Characterizing and carefully modeling these changes in responses across timescales will be an interesting avenue for future work (Masset, Qin, and Zavatone-Veth, 2022).

Though the circuit model derived in §A neural circuit architecture for Poisson compressed sensing 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-occurence statistics. Adaptive priors as encoded by cortical feedback would allow circuits to leverage this structure, enabling faster and more accurate inference (Davis and Macrides, 1981; Boyd et al., 2012; Markopoulos et al., 2012; Otazu et al., 2015; Wu, Yu, Chen, et al., 2020).

Our capacity 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 (Gerkin and Castro, 2015; Meister, 2015), 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 (Mayhew et al., 2022). Thus, an important objective for future work will be to scale this model 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} (Baraniuk, 2007; Donoho, 2006; Candès and Tao, 2006; Candès, Romberg, and Tao, 2006; Raginsky, Willett, et al., 2010; Raginsky, Jafarpour, et al., 2011). As a first step, in Supp. Fig. 1 and 2, we reproduce Fig. 3A-B for varying numbers of possible odorants from 500 to 2000, showing that performance does indeed drop off slowly with increasing odor space dimension. 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 (Zhang and Sharpee, 2016; Qin et al., 2019; Kepple, Cazakoff, et al., 2018; Grabska-Barwińska, Beck, et al., 2013; Grabska-Barwińska, Barthelmé, et al., 2017).

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 (Whittington et al., 2023; Plumbley, 2002; Pehlevan, Mohan, and Chklovskii, 2017; Lipshutz, Pehlevan, and Chklovskii, 2022; Higgins et al., 2021; Chang and Tsao, 2017; Bao et al., 2020; Rigotti et al., 2013; Hirokawa et al., 2019; Dubreuil et al., 2022). 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 (Bronstein et al., 2021). 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 (Chung and Abbott, 2021; Bronstein et al., 2021; Zavatone-Veth, Yang, et al., 2023).

Acknowledgements

We thank Naoki Hiratani, Shanshan Qin, and Vikrant Kapoor for useful discussions. This manuscript was typeset using the LaPreprint LaTeX template (https://github.com/roaldarbol/LaPreprint).

References

Ache, Barry W. and Janet M. Young (2005). "Olfaction: Diverse Species, Conserved Principles". In: *Neuron* 48.3, pp. 417–430. ISSN: 0896-6273. DOI: https://doi.org/10.1016/j.neuron.2005.10.022.

Ackels, Tobias, Andrew Erskine, Debanjan Dasgupta, Alina Cristina Marin, Tom P. A. Warner, Sina Tootoonian, Izumi Fukunaga, Julia J. Harris, and Andreas T. Schaefer (2021). "Fast odour dynamics are encoded in the olfactory system and guide behaviour". In: *Nature* 593.7860, pp. 558–563. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03514-2.

Amari, Shun-ichi (1998). "Natural Gradient Works Efficiently in Learning". In: *Neural Computation* 10.2, pp. 251–276. ISSN: 0899-7667. DOI: 10.1162/089976698300017746.

Arevian, Armen C., Vikrant Kapoor, and Nathaniel N. Urban (2008). "Activity-dependent gating of lateral inhibition in the mouse olfactory bulb". In: *Nature Neuroscience* 11.1, pp. 80–87. ISSN: 1546-1726. DOI: 10.1038/nn2030.

Atick, Joseph J. and A. Norman Redlich (1992). "What Does the Retina Know about Natural Scenes?" In: *Neural Computation* 4.2, pp. 196–210. ISSN: 0899-7667. DOI: 10.1162/neco.1992.4.2.196.

Attneave, Fred (1954). "Some informational aspects of visual perception". In: *Psychological Review* 61.3, p. 183. DOI: https://doi.org/10.1037/h0054663.

Bao, Pinglei, Liang She, Mason McGill, and Doris Y. Tsao (2020). "A map of object space in primate inferotemporal cortex". In: *Nature* 583.7814, pp. 103–108. ISSN: 1476-4687. DOI: 10.1038/s41586-020-2350-5.

Baraniuk, Richard G. (2007). "Compressive Sensing [Lecture Notes]". In: *IEEE Signal Processing Magazine* 24.4, pp. 118–121. DOI: 10.1109/MSP.2007.4286571.

Barlow, Horace B. (2012). "Possible Principles Underlying the Transformations of Sensory Messages". In: *Sensory Communication*. The MIT Press. ISBN: 9780262518420. DOI: 10.7551/mitpress/9780262518420.003.0013.

Bell, Anthony J. and Terrence J. Sejnowski (1997). "The "independent components" of natural scenes are edge filters". In: *Vision Research* 37.23, pp. 3327–3338. ISSN: 0042-6989. DOI: https://doi.org/10.1016/S0042-6989(97)00121-1.

Berners-Lee, Alice, Elizabeth Shtrahman, Julien Grimaud, and Venkatesh N. Murthy (2023). "Experience-dependent evolution of odor mixture representations in piriform cortex". In: *PLOS Biology* 21.4, pp. 1–29. URL: https://doi.org/10.1371/journal.pbio.3002086.

Boyd, Alison M., James F. Sturgill, Cindy Poo, and Jeffry S. Isaacson (2012). "Cortical Feedback Control of Olfactory Bulb Circuits". In: *Neuron* 76.6, pp. 1161–1174. ISSN: 0896-6273. DOI: https://doi.org/10.1016/j.neuron.2012.10.020.

Bronstein, Michael M, Joan Bruna, Taco Cohen, and Petar Veličković (2021). "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges". In: *arXiv preprint arXiv:2104.13478*. URL: https://geometricdeeplearning.com/.

Brunel, Nicolas and Jean-Pierre Nadal (1998). "Mutual Information, Fisher Information, and Population Coding". In: *Neural Computation* 10.7, pp. 1731–1757. ISSN: 0899-7667. DOI: 10.1162/089976698300017115.

Buck, Linda and Richard Axel (1991). "A novel multigene family may encode odorant receptors: a molecular basis for odor recognition". In: *Cell* 65.1, pp. 175–187. DOI: https://doi.org/10.1016/0092-8674(91)90418-X.

Burton, Shawn D. (2017). "Inhibitory circuits of the mammalian main olfactory bulb". In: *Journal of Neurophysiology* 118.4, pp. 2034–2051. DOI: 10.1152/jn.00109.2017.

Butts, Daniel A and Mark S Goldman (2006). "Tuning Curves, Neuronal Variability, and Sensory Coding". In: *PLoS Biology* 4.4. DOI: 10.1371/journal.pbio.0040092.

Calvin, William H. and Charles F. Stevens (1967). "Synaptic Noise as a Source of Variability in the Interval between Action Potentials". In: *Science* 155.3764, pp. 842–844. DOI: 10.1126/science.155.3764.842.

Candès, Emmanuel J., Justin K. Romberg, and Terence Tao (2006). "Stable signal recovery from incomplete and inaccurate measurements". In: *Communications on Pure and Applied Mathematics* 59.8, pp. 1207–1223. DOI: https://doi.org/10.1002/cpa.20124.

Candès, Emmanuel J. and Terence Tao (2006). "Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?" In: *IEEE Transactions on Information Theory* 52.12, pp. 5406–5425. DOI: 10.1109/TIT.2006.885507.

Chae, Honggoo, Arkarup Banerjee, Marie Dussauze, and Dinu F. Albeanu (2022). "Long-range functional loops in the mouse olfactory system and their roles in computing odor identity". In: *Neuron* 110.23, 3970–3985.e7. ISSN: 0896-6273. DOI: https://doi.org/10.1016/j.neuron.2022.09.005.

Chalk, Matthew, Paul Masset, Sophie Deneve, and Boris Gutkin (2017). "Sensory noise predicts divisive reshaping of receptive fields". In: *PLOS Computational Biology* 13.6, pp. 1–26. DOI: 10.1371/journal.pcbi.1005582.

Chang, Le and Doris Y. Tsao (2017). "The Code for Facial Identity in the Primate Brain". In: *Cell* 169.6, 1013–1028.e14. ISSN: 0092-8674. DOI: https://doi.org/10.1016/j.cell.2017.05.011.

Chapochnikov, Nikolai M., Cengiz Pehlevan, and Dmitri B. Chklovskii (2021). "Normative and mechanistic model of an adaptive circuit for efficient encoding and feature extraction". In: *bioRxiv*. DOI: 10.1101/2021.09.24.461723.

Chung, SueYeon and L.F. Abbott (2021). "Neural population geometry: An approach for understanding biological and artificial neural networks". In: *Current Opinion in Neurobiology* 70. Computational Neuroscience, pp. 137–144. ISSN: 0959-4388. DOI: https://doi.org/10.1016/j.conb.2021.10.010.

Crick, Francis (1989). "The recent excitement about neural networks". In: *Nature* 337.6203, pp. 129–132. ISSN: 1476-4687. DOI: 10.1038/337129a0.

Cury, Kevin M. and Naoshige Uchida (2010). "Robust Odor Coding via Inhalation-Coupled Transient Activity in the Mammalian Olfactory Bulb". In: *Neuron* 68.3, pp. 570–585. ISSN: 0896-6273. DOI: https://doi.org/10.1016/j.neuron.2010.09.040.

Davis, Barry J. and Foteos Macrides (1981). "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". In: *Journal of Comparative Neurology* 203.3, pp. 475–493. DOI: https://doi.org/10.1002/cne.902030310.

del Mármol, Josefina, Mackenzie A. Yedlin, and Vanessa Ruta (2021). "The structural basis of odorant recognition in insect olfactory receptors". In: *Nature* 597.7874, pp. 126–131. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03794-8.

Donoho, D.L. (2006). "Compressed sensing". In: *IEEE Transactions on Information Theory* 52.4, pp. 1289–1306. DOI: 10.1109/TIT.2006.871582.

Donoho, David and Jared Tanner (2009). "Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367.1906, pp. 4273–4293. DOI: 10.1098/rsta.2009.0152.

Dubreuil, Alexis, Adrian Valente, Manuel Beiran, Francesca Mastrogiuseppe, and Srdjan Ostojic (2022). "The role of population structure in computations through neural dynamics". In: *Nature Neuroscience* 25.6, pp. 783–794. DOI: https://doi.org/10.1038/s41593-022-01088-4.

Duchamp-Viret, Patricia, Lubomir Kostal, Michel Chaput, Petr Lánsky, and Jean-Pierre Rospars (2005). "Patterns of spontaneous activity in single rat olfactory receptor neurons are different in normally breathing and tracheotomized animals". In: *Journal of Neurobiology* 65.2, pp. 97–114. DOI: https://doi.org/10.1002/neu.20177.

Dytso, Alex, Michael Fauß, and H. Vincent Poor (2020). "The Vector Poisson Channel: On the Linearity of the Conditional Mean Estimator". In: *IEEE Transactions on Signal Processing* 68, pp. 5894–5903. DOI: 10.1109/TSP.2020.3025525.

Egger, Veronica and Nathaniel N. Urban (2006). "Dynamic connectivity in the mitral cell-granule cell microcircuit". In: *Seminars in Cell & Developmental Biology* 17.4. Olfaction Animal Stem Cell Types, pp. 424–432. ISSN: 1084-9521. DOI: https://doi.org/10.1016/j.semcdb.2006.04.006.

Fang, Michael Y.-S., Mayur Mudigonda, Ryan Zarcone, Amir Khosrowshahi, and Bruno A. Olshausen (2022). "Learning and Inference in Sparse Coding Models With Langevin Dynamics". In: *Neural Computation* 34.8, pp. 1676–1700. ISSN: 0899-7667. DOI: 10.1162/neco_a_01505.

Fiser, József, Pietro Berkes, Gergő Orbán, and Máté Lengyel (2010). "Statistically optimal perception and learning: from behavior to neural representations". In: *Trends in Cognitive Sciences* 14.3, pp. 119–130. DOI: 10.1016/j.tics.2010.01.003.

Gerkin, Richard C and Jason B Castro (2015). "The number of olfactory stimuli that humans can discriminate is still unknown". In: *eLife* 4. Ed. by Alexander Borst, e08127. ISSN: 2050-084X. DOI: 10.7554/eLife.08127.

Girolami, Mark and Ben Calderhead (2011). "Riemann manifold Langevin and Hamiltonian Monte Carlo methods". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73.2, pp. 123–214. DOI: 10.1111/j.1467-9868.2010.00765.x.

Gorur-Shandilya, Srinivas, Mahmut Demir, Junjiajia Long, Damon A Clark, and Thierry Emonet (2017). "Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli". In: *eLife* 6. Ed. by Fred Rieke, e27670. ISSN: 2050-084X. DOI: 10.7554/eLife.27670.

Grabska-Barwińska, Agnieszka, Simon Barthelmé, Jeff Beck, Zachary F. Mainen, Alexandre Pouget, and Peter E. Latham (2017). "A probabilistic approach to demixing odors". In: *Nature Neuroscience* 20.1, pp. 98–106. ISSN: 1546-1726. DOI: 10.1038/nn.4444.

Grabska-Barwińska, Agnieszka, Jeff Beck, Alexandre Pouget, and Peter Latham (2013). "Demixing odors - fast inference in olfaction". In: *Advances in Neural Information Processing Systems*. Ed. by C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger. Vol. 26. Curran Asso-

ciates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2013/file/2bcab9d935d219641434683dd9d18a
Paper.pdf.

Guerguiev, Jordan, Timothy P Lillicrap, and Blake A Richards (2017). "Towards deep learning with segregated dendrites". In: *eLife* 6. Ed. by Peter Latham, e22901. ISSN: 2050-084X. DOI: 10.7554/eLife.22901.

Gupta, Priyanka, Dinu F Albeanu, and Upinder S Bhalla (2015). "Olfactory bulb coding of odors, mixtures and sniffs is a linear sum of odor time profiles". In: *Nature Neuroscience* 18.2, pp. 272–281. DOI: https://doi.org/10.1038/nn.3913.

Hateren, Johannes H. van (1992). "A theory of maximizing sensory information". In: *Biological Cybernetics* 68.1, pp. 23–29. ISSN: 1432-0770. DOI: 10.1007/BF00203134.

Higgins, Irina, Le Chang, Victoria Langston, Demis Hassabis, Christopher Summerfield, Doris Tsao, and Matthew Botvinick (2021). "Unsupervised deep learning identifies semantic disentanglement in single inferotemporal face patch neurons". In: *Nature Communications* 12.1, p. 6456. ISSN: 2041-1723. DOI: 10.1038/s41467-021-26751-5.

Hiratani, Naoki and Peter E. Latham (2020). "Rapid Bayesian learning in the mammalian olfactory system". In: *Nature Communications* 11.1, p. 3845. ISSN: 2041-1723. DOI: 10.1038/s41467-020-17490-0.

Hirokawa, Junya, Alexander Vaughan, Paul Masset, Torben Ott, and Adam Kepecs (2019). "Frontal cortex neuron types categorically encode single decision variables". In: *Nature* 576.7787, pp. 446–451. DOI: https://doi.org/10.1038/s41586-019-1816-9.

Holt, Gary R. and Christof Koch (1997). "Shunting Inhibition Does Not Have a Divisive Effect on Firing Rates". In: *Neural Computation* 9.5, pp. 1001–1013. ISSN: 0899-7667. DOI: 10.1162/neco.1997.9.5.1001.

Hubel, D. H. and T. N. Wiesel (1959). "Receptive fields of single neurones in the cat's striate cortex". In: *The Journal of Physiology* 148.3, pp. 574–591. DOI: https://doi.org/10.1113/jphysiol.1959.sp006308.

— (1962). "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex". In: *The Journal of Physiology* 160.1, pp. 106–154. DOI: https://doi.org/10.1113/jphysiol.1962.sp006837.

Igarashi, Kei M., Nao Ieki, Myungho An, Yukie Yamaguchi, Shin Nagayama, Ko Kobayakawa, Reiko Kobayakawa, Manabu Tanifuji, Hitoshi Sakano, Wei R. Chen, and Kensaku Mori (2012). "Parallel Mitral and Tufted Cell Pathways Route Distinct Odor Information to Different Targets in the Olfactory Cortex". In: *Journal of Neuroscience* 32.23, pp. 7970–7985. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.0154-12.2012.

Isaacson, Jeffry S (2010). "Odor representations in mammalian cortical circuits". In: *Current Opinion in Neurobiology* 20.3. Sensory systems, pp. 328–331. ISSN: 0959-4388. DOI: https://doi.org/10.1016/j.conb.2010.02.004.

Jin, An, Birsen Yazici, Angelique Ale, and Vasilis Ntziachristos (2012). "Preconditioning of the fluorescence diffuse optical tomography sensing matrix based on compressive sensing". In: *Optics Letters* 37.20, pp. 4326–4328. DOI: 10.1364/OL.37.004326.

Kadakia, Nirag and Thierry Emonet (2019). "Front-end Weber-Fechner gain control enhances the fidelity of combinatorial odor coding". In: *eLife* 8. Ed. by Fred Rieke, Catherine Dulac, and Katherine I. Nagel, e45293. ISSN: 2050-084X. DOI: 10.7554/eLife.45293.

Kandler, Karl, Amanda Clause, and Jihyun Noh (2009). "Tonotopic reorganization of developing auditory brainstem circuits". In: *Nature neuroscience* 12.6, pp. 711–717. DOI: https://doi.org/10.1038/nn.2332.

Kepecs, Adam, Naoshige Uchida, and Zachary F. Mainen (2005). "The Sniff as a Unit of Olfactory Processing". In: *Chemical Senses* 31.2, pp. 167–179. ISSN: 0379-864X. DOI: 10.1093/chemse/bjj016.

Kepple, Daniel, Brittany N. Cazakoff, Heike S. Demmer, Dennis Eckmeier, Stephen D. Shea, and Alexei A. Koulakov (2018). "Computational algorithms and neural circuitry for compressed sensing in the mammalian main olfactory bulb". In: *bioRxiv*. DOI: 10.1101/339689.

Kepple, Daniel R, Hamza Giaffar, Dmitry Rinberg, and Alexei A Koulakov (2019). "Deconstructing odorant identity via primacy in dual networks". In: *Neural Computation* 31.4, pp. 710–737. DOI: https://doi.org/10.1162/neco_a_01175.

Khan, Adil Ghani, K. Parthasarathy, and Upinder Singh Bhalla (2010). "Odor representations in the mammalian olfactory bulb". In: *Wiley Interdisciplinary Reviews: Systems Biology and Medicine* 2.5, pp. 603–611. DOI: https://doi.org/10.1002/wsbm.85.

Knill, David C and Alexandre Pouget (2004). "The Bayesian brain: the role of uncertainty in neural coding and computation". In: *Trends in Neurosciences* 27.12, pp. 712–719. DOI: 10.1016/j.tins.2004.10.007.

Körding, Konrad P and Daniel M Wolpert (2004). "Bayesian integration in sensorimotor learning". In: *Nature* 427.6971, pp. 244–247. DOI: 10.1038/nature02169.

Kostal, Lubomir, Petr Lansky, and Jean-Pierre Rospars (2007). "Neuronal coding and spiking randomness". In: *European Journal of Neuroscience* 26.10, pp. 2693–2701. DOI: https://doi.org/10.1111/j.1460-9568.2007.05880.x.

Koulakov, Alexei A. and Dmitry Rinberg (2011). "Sparse Incomplete Representations: A Potential Role of Olfactory Granule Cells". In: *Neuron* 72.1, pp. 124–136. ISSN: 0896-6273. DOI: https://doi.org/10.1016/j.neuron.2011.07.031.

Krishnamurthy, Kamesh, Ann M Hermundstad, Thierry Mora, Aleksandra M Walczak, and Vijay Balasubramanian (2022). "Disorder and the neural representation of complex odors". In: *Frontiers in Computational Neuroscience* 16. DOI: 10.3389/fncom.2022.917786.

Lebovich, Lior, Michael Yunerman, Viviana Scaiewicz, Yonatan Loewenstein, and Dan Rokni (2021). "Paradoxical relationship between speed and accuracy in olfactory figure-background segregation". In: *PLOS Computational Biology* 17, pp. 1–16. DOI: 10.1371/journal.pcbi.1009674.

Lee, Brian K., 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 (2022). "A Principal Odor Map Unifies Diverse Tasks in Human Olfactory Perception". In: *bioRxiv*. DOI: 10.1101/2022.09.01.504602.

Li, Yan, 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 (2023). "Robust odor identification in novel olfactory environments in mice". In: *Nature Communications* 14.1, p. 673. ISSN: 2041-1723. DOI: 10.1038/s41467-023-36346-x.

Lillicrap, Timothy P., Adam Santoro, Luke Marris, Colin J. Akerman, and Geoffrey Hinton (2020). "Backpropagation and the brain". In: *Nature Reviews Neuroscience* 21.6, pp. 335–346. ISSN: 1471-0048. DOI: 10.1038/s41583-020-0277-3.

Linsker, R. (1988). "Self-organization in a perceptual network". In: *Computer* 21.3, pp. 105–117. DOI: 10.1109/2.36.

Lipshutz, David, Cengiz Pehlevan, and Dmitri B. Chklovskii (2022). "Biologically plausible single-layer networks for nonnegative independent component analysis". In: *Biological Cybernetics* 116.5, pp. 557–568. ISSN: 1432-0770. DOI: 10.1007/s00422-022-00943-8.

Lledo, Pierre-Marie, Gilles Gheusi, and Jean-Didier Vincent (2005). "Information Processing in the Mammalian Olfactory System". In: *Physiological Reviews* 85.1, pp. 281–317. DOI: 10.1152/physrev.00008.2004.

Ma, Yi-An, Tianqi Chen, and Emily Fox (2015). "A Complete Recipe for Stochastic Gradient MCMC". In: *Advances in Neural Information Processing Systems*. Ed. by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett. Vol. 28. Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2015/file/9a4400501febb2a95e79248486a5f6d3-Paper.pdf.

Ma, Limei, Qiang Qiu, Stephen Gradwohl, Aaron Scott, Elden Q. Yu, Richard Alexander, Winfried Wiegraebe, and C. Ron Yu (2012). "Distributed representation of chemical features and tunotopic organization of glomeruli in the mouse olfactory bulb". In: *Proceedings of the National Academy of Sciences* 109.14, pp. 5481–5486. DOI: 10.1073/pnas.1117491109.

Maleki, Arian and David L. Donoho (2010). "Optimally Tuned Iterative Reconstruction Algorithms for Compressed Sensing". In: *IEEE Journal of Selected Topics in Signal Processing* 4.2, pp. 330–341. DOI: 10.1109/JSTSP.2009.2039176.

Malnic, Bettina, Junzo Hirono, Takaaki Sato, and Linda B Buck (1999). "Combinatorial Receptor Codes for Odors". In: *Cell* 96.5, pp. 713–723. ISSN: 0092-8674. DOI: https://doi.org/10.1016/S0092-8674(00)80581-4.

Markopoulos, Foivos, Dan Rokni, David H. Gire, and Venkatesh N. Murthy (2012). "Functional Properties of Cortical Feedback Projections to the Olfactory Bulb". In: *Neuron* 76.6, pp. 1175–1188. ISSN: 0896-6273. DOI: https://doi.org/10.1016/j.neuron.2012.10.028.

Marr, David and Tomaso Poggio (1976). *From understanding computation to understanding neural circuitry*. Tech. rep. Massachussetts Institute of Technology AI Laboratory. URL: https://dspace.mit.edu/handle/1721.1/5782.

Martens, James (2020). "New Insights and Perspectives on the Natural Gradient Method". In: *Journal of Machine Learning Research* 21.146, pp. 1–76. URL: http://jmlr.org/papers/v21/17-678.html.

Masset, Paul, Shanshan Qin, and Jacob A Zavatone-Veth (2022). "Drifting neuronal representations: Bug or feature?" In: *Biological Cybernetics*, pp. 1–14. DOI: doi.org/10.1007/s00422-021-00916-3.

Masset, Paul, Jacob Zavatone-Veth, J. Patrick Connor, Venkatesh Murthy, and Cengiz Pehlevan (2022). "Natural gradient enables fast sampling in spiking neural networks". In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., pp. 22018–22034. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/8a0fd48510590071e3c129a79b8b8527-Paper-Conference.pdf.

Mathis, Alexander, Dan Rokni, Vikrant Kapoor, Matthias Bethge, and Venkatesh N. Murthy (2016). "Reading Out Olfactory Receptors: Feedforward Circuits Detect Odors in Mixtures without Demixing". In: *Neuron* 91.5, pp. 1110–1123. ISSN: 0896-6273. DOI: https://doi.org/10.1016/j.neuron.2016.08.007.

Mayhew, Emily J., 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 (2022). "Transport features predict if a molecule is odorous". In: *Proceedings of the National Academy of Sciences* 119.15, e2116576119. DOI: 10.1073/pnas.2116576119.

Meister, Markus (2015). "On the dimensionality of odor space". In: *eLife* 4. Ed. by Alexander Borst, e07865. ISSN: 2050-084X. DOI: 10.7554/eLife.07865.

Murphy, Kenneth, Casey Weaver, and Charles Janeway (2017). *Janeway's Immunobiology*. 9th edition. New York: Garland Science. ISBN: 9780815345053.

Murthy, Venkatesh N. (2011). "Olfactory Maps in the Brain". In: *Annual Review of Neuroscience* 34.1, pp. 233–258. DOI: 10.1146/annurev-neuro-061010-113738.

Nagayama, Shin, Ryota Homma, and Fumiaki Imamura (2014). "Neuronal organization of olfactory bulb circuits". In: *Frontiers in Neural Circuits* 8, p. 98. DOI: 10.3389/fncir.2014.00098.

Neal, Radford M (1993). *Probabilistic inference using Markov chain Monte Carlo methods*. Department of Computer Science, University of Toronto Toronto, ON, Canada. URL: https://www.cs.toronto.edu/~radford/review.abstract.html.

Niimura, Yoshihito, Atsushi Matsui, and Kazushige Touhara (2014). "Extreme expansion of the olfactory receptor gene repertoire in African elephants and evolutionary dynamics of orthologous gene groups in 13 placental mammals". In: *Genome Research* 24.9, pp. 1485–1496. DOI: 10.1101/gr.169532.113.

Niimura, Yoshihito and Masatoshi Nei (2007). "Extensive Gains and Losses of Olfactory Receptor Genes in Mammalian Evolution". In: *PLOS ONE* 2.8, pp. 1–8. DOI: 10.1371/journal.pone.0000708.

Olshausen, Bruno A. and David J. Field (1996). "Emergence of simple-cell receptive field properties by learning a sparse code for natural images". In: *Nature* 381.6583, pp. 607–609. ISSN: 1476-4687. DOI: 10.1038/381607a0.

Otazu, Gonzalo H, Honggoo Chae, Martin B Davis, and Dinu F Albeanu (2015). "Cortical feedback decorrelates olfactory bulb output in awake mice". In: *Neuron* 86.6, pp. 1461–1477. DOI: https://doi.org/10.1016/j.neuron.2015.05.023.

Ott, Torben, Paul Masset, and Adam Kepecs (2018). "The neurobiology of confidence: From beliefs to neurons". In: *Cold Spring Harbor Symposia on Quantitative Biology*. Vol. 83. Cold Spring Harbor Laboratory Press, pp. 9–16. DOI: 10.1101/sqb.2018.83.038794.

Pehlevan, Cengiz, Alexander Genkin, and Dmitri B. Chklovskii (2017). "A clustering neural network model of insect olfaction". In: *51st Asilomar Conference on Signals, Systems, and Computers*, pp. 593–600. DOI: 10.1109/ACSSC.2017.8335410.

Pehlevan, Cengiz, Sreyas Mohan, and Dmitri B. Chklovskii (2017). "Blind Nonnegative Source Separation Using Biological Neural Networks". In: *Neural Computation* 29.11, pp. 2925–2954. ISSN: 0899-7667. DOI: 10.1162/neco_a_01007.

Plumbley, Mark (2002). "Conditions for nonnegative independent component analysis". In: *IEEE Signal Processing Letters* 9.6, pp. 177–180. DOI: 10.1109/LSP.2002.800502.

Pnevmatikakis, Eftychios A. and Andrea Giovannucci (2017). "NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data". In: *Journal of Neuroscience Methods* 291, pp. 83–94. ISSN: 0165-0270. DOI: https://doi.org/10.1016/j.jneumeth.2017.07.031.

Qian, Wesley W, Jennifer N Wei, Benjamin Sanchez-Lengeling, Brian K Lee, Yunan Luo, Marnix Vlot, Koen Dechering, Jian Peng, Richard C Gerkin, and Alexander B Wiltschko (2023). "Metabolic activity organizes olfactory representations". In: *eLife* 12. Ed. by Upinder Singh Bhalla and Andrew J King, e82502. ISSN: 2050-084X. DOI: 10.7554/eLife.82502.

Qin, Shanshan, Qianyi Li, Chao Tang, and Yuhai Tu (2019). "Optimal compressed sensing strategies for an array of nonlinear olfactory receptor neurons with and without spontaneous activity". In: *Proceedings of the National Academy of Sciences* 116.41, pp. 20286–20295. DOI: 10.1073/pnas.1906571116.

Raginsky, Maxim, Sina Jafarpour, Zachary T. Harmany, Roummel F. Marcia, Rebecca M. Willett, and Robert Calderbank (2011). "Performance Bounds for Expander-Based Compressed Sensing in Poisson Noise". In: *IEEE Transactions on Signal Processing* 59.9, pp. 4139–4153. DOI: 10.1109/TSP.2011.2157913.

Raginsky, Maxim, Rebecca M. Willett, Zachary T. Harmany, and Roummel F. Marcia (2010). "Compressed Sensing Performance Bounds Under Poisson Noise". In: *IEEE Transactions on Signal Processing* 58.8, pp. 3990–4002. DOI: 10.1109/TSP.2010.2049997.

Rao, Rajesh P. N. and Dana H. Ballard (1999). "Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects". In: *Nature Neuroscience* 2.1, pp. 79–87. ISSN: 1546-1726. DOI: 10.1038/4580.

Reddy, Gautam, Venkatesh N. Murthy, and Massimo Vergassola (2022). "Olfactory Sensing and Navigation in Turbulent Environments". In: *Annual Review of Condensed Matter Physics* 13.1, pp. 191–213. DOI: 10.1146/annurev-conmatphys-031720-032754.

Reddy, Gautam, Joseph D Zak, Massimo Vergassola, and Venkatesh N Murthy (2018). "Antagonism in olfactory receptor neurons and its implications for the perception of odor mixtures". In: *eLife* 7. Ed. by Fred Rieke, e34958. ISSN: 2050-084X. DOI: 10.7554/eLife.34958.

Reinert, Janine K. and Izumi Fukunaga (2022). "The facets of olfactory learning". In: *Current Opinion in Neurobiology* 76, p. 102623. ISSN: 0959-4388. DOI: https://doi.org/10.1016/j.conb.2022.102623.

Rigotti, Mattia, Omri Barak, Melissa R. Warden, Xiao-Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi (2013). "The importance of mixed selectivity in complex cognitive tasks". In: *Nature* 497.7451, pp. 585–590. ISSN: 1476-4687. DOI: 10.1038/nature12160.

Rokni, Dan, Vivian Hemmelder, Vikrant Kapoor, and Venkatesh N. Murthy (2014). "An olfactory cocktail party: figure-ground segregation of odorants in rodents". In: *Nature Neuroscience* 17.9, pp. 1225–1232. ISSN: 1546-1726. DOI: 10.1038/nn.3775.

Savin, Cristina and Sophie Denève (2014). "Spatio-temporal representations of uncertainty in spiking neural networks". In: *Advances in Neural Information Processing Systems* 27. URL: https://proceedings.neurips.cc/paper/2014/hash/4e2545f819e67f0615003dd7e04a6087-Abstract.html.

Shusterman, Roman, Matthew C. Smear, Alexei A. Koulakov, and Dmitry Rinberg (2011). "Precise olfactory responses tile the sniff cycle". In: *Nature Neuroscience* 14.8, pp. 1039–1044. ISSN: 1546-1726. DOI: 10.1038/nn.2877.

Singh, Vijay, Martin Tchernookov, and Vijay Balasubramanian (2021). "What the odor is not: Estimation by elimination". In: *Phys. Rev. E* 104 (2), p. 024415. DOI: 10.1103/PhysRevE.104.024415.

Spors, Hartwig, Dinu Florin Albeanu, Venkatesh N. Murthy, Dmitry Rinberg, Naoshige Uchida, Matt Wachowiak, and Rainer W. Friedrich (2012). "Illuminating Vertebrate Olfactory Processing". In: *Journal of Neuroscience* 32.41, 14102–14108a. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.3328-12.2012.

Srinivasan, Mandyam Veerambudi, Simon Barry Laughlin, A. Dubs, and George Adrian Horridge (1982). "Predictive coding: a fresh view of inhibition in the retina". In: *Proceedings of the Royal Society of London. Series B. Biological Sciences* 216.1205, pp. 427–459. DOI: 10.1098/rspb.1982.0085.

Su, Chih-Ying, Karen Menuz, and John R. Carlson (2009). "Olfactory Perception: Receptors, Cells, and Circuits". In: *Cell* 139.1, pp. 45–59. ISSN: 0092-8674. DOI: https://doi.org/10.1016/j.cell.2009.09.015.

Teşileanu, Tiberiu, Simona Cocco, Rémi Monasson, and Vijay Balasubramanian (2019). "Adaptation of olfactory receptor abundances for efficient coding". In: *eLife* 8. Ed. by Upinder Singh Bhalla, Catherine Dulac, and David Zwicker, e39279. ISSN: 2050-084X. DOI: 10.7554/eLife.39279.

Tootoonian, Sina and Mate Lengyel (2014). "A Dual Algorithm for Olfactory Computation in the Locust Brain". In: *Advances in Neural Information Processing Systems*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger. Vol. 27. Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2014/file/c8ba76c279269b1c6bc8a07e38e78fa4-Paper.pdf.

Tsukahara, Tatsuya, David H. Brann, Stan L. Pashkovski, Grigori Guitchounts, Thomas Bozza, and Sandeep Robert Datta (2021). "A transcriptional rheostat couples past activity to future sensory responses". In: *Cell* 184.26, 6326–6343.e32. ISSN: 0092-8674. DOI: https://doi.org/10.1016/j.cell.2021.11.022.

Uchida, Naoshige and Zachary F. Mainen (2003). "Speed and accuracy of olfactory discrimination in the rat". In: *Nature Neuroscience* 6.11, pp. 1224–1229. ISSN: 1546-1726. DOI: 10.1038/nn1142.

Uchida, Naoshige, Cindy Poo, and Rafi Haddad (2014). "Coding and Transformations in the Olfactory System". In: *Annual Review of Neuroscience* 37.1, pp. 363–385. DOI: 10.1146/annurev-neuro-071013-013941.

Uexküll, Jakob Johann von (1909). *Umwelt und Innenwelt der Tiere*. Springer. DOI: 10.1007/978-3-662-24819-5.

Urban, Nathaniel N. and Armen C. Arevian (2009). "Computing with Dendrodendritic Synapses in the Olfactory Bulb". In: *Annals of the New York Academy of Sciences* 1170.1, pp. 264–269. DOI: https://doi.org/10.1111/j.1749-6632.2009.03899.x.

Whittington, James C. R., Will Dorrell, Surya Ganguli, and Timothy Behrens (2023). "Disentanglement with Biological Constraints: A Theory of Functional Cell Types". In: *The Eleventh International Conference on Learning Representations*. URL: https://openreview.net/forum?id=9Z_GfhZnGH.

Wu, An, Bin Yu, Qiyu Chen, Gillian A. Matthews, Chen Lu, Evan Campbell, Kay M. Tye, and Takaki Komiyama (2020). "Context-dependent plasticity of adult-born neurons regulated by cortical feedback". In: *Science Advances* 6.42, eabc8319. DOI: 10.1126/sciadv.abc8319.

Wu, An, Bin Yu, and Takaki Komiyama (2020). "Plasticity in olfactory bulb circuits". In: *Current Opinion in Neurobiology* 64. Systems Neuroscience, pp. 17–23. ISSN: 0959-4388. DOI: https://doi.org/10.1016/j.conb.2020.01.007.

Yong, Ed (2022). *An Immense World: How Animal Senses Reveal the Hidden Realms Around Us*. First. New York: Random House, p. 464. ISBN: 9780593133231.

Zak, Joseph D., Gautam Reddy, Massimo Vergassola, and Venkatesh N. Murthy (2020). "Antagonistic odor interactions in olfactory sensory neurons are widespread in freely breathing mice". In: *Nature Communications* 11.1, p. 3350. ISSN: 2041-1723. DOI: 10.1038/s41467-020-17124-5.

Zavatone-Veth, Jacob A., Bara A. Badwan, and Damon A. Clark (2020). "A minimal synaptic model for direction selective neurons in *Drosophila*". In: *Journal of Vision* 20.2, pp. 2–2. ISSN: 1534-7362. DOI: 10.1167/jov.20.2.2.

Zavatone-Veth, Jacob A., Sheng Yang, Julian A. Rubinfien, and Cengiz Pehlevan (2023). "Neural networks learn to magnify areas near decision boundaries". In: *arXiv*. DOI: https://doi.org/10.48550/arXiv.2301.11375.

Zhang, Yilun and Tatyana O. Sharpee (2016). "A Robust Feedforward Model of the Olfactory System". In: *PLOS Computational Biology* 12.4, pp. 1–15. DOI: 10.1371/journal.pcbi.1004850.

Zhu, Mengchen and Christopher J. Rozell (2015). "Modeling Inhibitory Interneurons in Efficient Sensory Coding Models". In: *PLoS Computational Biology* 11.7, pp. 1–22. DOI: 10.1371/journal.pcbi.1004353.

Zwicker, David, Arvind Murugan, and Michael P. Brenner (2016). "Receptor arrays optimized for natural odor statistics". In: *Proceedings of the National Academy of Sciences* 113.20, pp. 5570–5575. DOI: 10.1073/pnas.1600357113.

Appendix 1

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 $\lambda \in \mathbb{R}_+^n$, we write

$$\mathbf{x} \sim \text{Poisson}(\boldsymbol{\lambda}) \tag{6}$$

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

$$p(\mathbf{x}) = \prod_{j=1}^{n} \frac{\lambda_j^{x_j}}{x_j!} e^{-\lambda_j}. \tag{7}$$

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

$$\mathbf{x} \sim \text{Gamma}(\alpha, \lambda) \tag{8}$$

if its probability density function is

$$p(\mathbf{x}) = \prod_{j=1}^{n} \frac{x_j^{\alpha-1}}{\Gamma(\alpha)} e^{-\lambda x_j}. \tag{9}$$

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

$$\mathbf{x} \sim \text{Gamma}(\boldsymbol{\alpha}, \boldsymbol{\lambda}) \tag{10}$$

if

$$p(\mathbf{x}) = \prod_{j=1}^{n} \frac{x_j^{\alpha_j-1}}{\Gamma(\alpha_j)} e^{-\lambda_j x_j}. \tag{11}$$

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

$$(\mathbf{x} \odot \mathbf{y})_i = x_i y_i \tag{12}$$

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

$$(\mathbf{x} \oslash \mathbf{y})_i = \frac{x_i}{y_i}. \tag{13}$$

Appendix 2

Detailed derivation of circuit model

In this Appendix, we give a detailed derivation of the circuit models introduced in Sections A neural circuit architecture for Poisson compressed sensing and Fast sampling for uncertainty estimation of the main text. Here, we focus on the setting of the sampling circuit; the circuit algorithm to compute the MAP introduced in Section A neural circuit architecture for Poisson compressed sensing 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 $p(\mathbf{c} \mid \mathbf{s})$ over concentrations $\mathbf{c}$ given OSN activity $\mathbf{s}$, for a Poisson likelihood

$$\mathbf{s} \mid \mathbf{c} \sim \mathrm{Poisson}(\mathbf{r}_0 + \mathbf{A}\mathbf{c}) \tag{14}$$

and a Gamma prior

$$\mathbf{c} \sim \mathrm{Gamma}(\boldsymbol{\alpha}, \boldsymbol{\lambda}). \tag{15}$$

Here, we use the conventions of Appendix 1 for Poisson and Gamma random vectors. Discarding normalization constants that do not depend on $\mathbf{c}$, the density of the posterior is then

$$p(\mathbf{c} \mid \mathbf{s}) \propto p(\mathbf{s} \mid \mathbf{c}) p(\mathbf{c}) \tag{16}$$

$$\propto \left[\prod_{k=1}^{n_{\mathrm{OSN}}} (\mathbf{r}_0 + \mathbf{A}\mathbf{c})_k^{s_k} e^{-(\mathbf{r}_0 + \mathbf{A}\mathbf{c})_k}\right] \left[\prod_{j=1}^{n_{\mathrm{odor}}} c_j^{\alpha_j - 1} e^{-\lambda_j c_j}\right]. \tag{17}$$

We will sample from this distribution using Langevin dynamics, without explicitly constraining the concentration estimate to be non-negative (Grabska-Barwińska, Beck, et al., 2013; Grabska-Barwińska, Barthelmé, et al., 2017). This corresponds to the stochastic dynamics

$$\dot{\mathbf{c}}(t) = \boldsymbol{\nabla}_{\mathbf{c}} \log p(\mathbf{c} \mid \mathbf{s}) + \boldsymbol{\eta}(t), \tag{18}$$

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

$$\frac{\partial \log p(\mathbf{c} \mid \mathbf{s})}{\partial c_j} = \sum_{k=1}^{n_{\mathrm{OSN}}} \left(\frac{s_k}{(\mathbf{r}_0 + \mathbf{A}\mathbf{c})_k} - 1\right) A_{kj} + \frac{\alpha_j - 1}{c_j} - \lambda_j, \tag{19}$$

or, in vector notation,

$$\boldsymbol{\nabla}_{\mathbf{c}} \log p(\mathbf{c} \mid \mathbf{s}) = \mathbf{A}^{\top}[\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\mathbf{c}) - \mathbf{1}] + (\boldsymbol{\alpha} - \mathbf{1}) \oslash \mathbf{c} - \boldsymbol{\lambda}. \tag{20}$$

The Langevin dynamics then become

$$\dot{\mathbf{c}}(t) = \mathbf{A}^{\top}[\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\mathbf{c}) - \mathbf{1}] + (\boldsymbol{\alpha} - \mathbf{1}) \oslash \mathbf{c} - \boldsymbol{\lambda} + \boldsymbol{\eta}(t). \tag{21}$$

This is closely related to the Langevin sampling algorithm introduced in equation (3.9) of Grabska-Barwińska, Beck, et al. (2013), but here we are not separately inferring odor concentration and odor presence. We note that, for $\alpha_j > 1$, the term $(\boldsymbol{\alpha} - \mathbf{1}) \oslash \mathbf{c}$ 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,

Zavatone-Veth, et al. (2022) in using the "complete recipe" for stochastic gradient MCMC (Ma, Chen, and Fox, 2015). From the "complete recipe", we know that the dynamics

$$\tau_{\mathrm{g}}\dot{\mathbf{c}}(t) = \boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\boldsymbol{\nabla}_{\mathbf{c}} \log p(\mathbf{c} \,|\, \mathbf{s}) + \boldsymbol{\Gamma}\boldsymbol{\xi}(t), \tag{22}$$

for a (expectantly named) time constant $\tau_{\mathrm{g}} > 0$ and any (potentially non-square) matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{n_{\mathrm{odor}} \times n_{\mathrm{g}}}$ 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

$$\mathbb{E}[\xi_j(t)\xi_{j'}(t')] = 2\tau_{\mathrm{g}}\delta_{jj'}\delta(t-t'). \tag{23}$$

Then, we can write the concentration estimate as

$$\mathbf{c}(t) = \boldsymbol{\Gamma}\mathbf{g}(t) \tag{24}$$

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

$$\tau_{\mathrm{g}}\dot{\mathbf{g}}(t) = \boldsymbol{\Gamma}^{\top}\boldsymbol{\nabla}_{\mathbf{c}} \log p(\mathbf{g} \,|\, \mathbf{s})\Big|_{\mathbf{c}=\boldsymbol{\Gamma}\mathbf{g}} + \boldsymbol{\xi}(t). \tag{25}$$

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

$$\tau_{\mathrm{g}}\dot{\mathbf{g}}(t) = (\mathbf{A}\boldsymbol{\Gamma})^{\top}[\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g}) - \mathbf{1}] + \boldsymbol{\Gamma}^{\top}[(\boldsymbol{\alpha} - \mathbf{1}) \oslash (\boldsymbol{\Gamma}\mathbf{g})] - \boldsymbol{\Gamma}^{\top}\lambda + \boldsymbol{\xi}(t). \tag{26}$$

These dynamics include two divisive non-linearities, which can be complex to implement in biophysical models of single neurons (Holt and Koch, 1997). Using the approach proposed in Chalk et al. (2017), 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 $(\boldsymbol{\alpha} - \mathbf{1}) \oslash (\boldsymbol{\Gamma}\mathbf{g})$. We then introduce a population $\mathbf{p}$ of n_{OSN} neurons, with dynamics

$$\tau_{\mathrm{p}}\dot{\mathbf{p}}(t) = \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g}), \tag{27}$$

such that the fixed point is

$$\mathbf{p}^* = \mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g}). \tag{28}$$

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

$$\tau_{\mathrm{z}}\dot{\mathbf{z}}(t) = \boldsymbol{\alpha} - \mathbf{1} - \mathbf{z} \odot (\boldsymbol{\Gamma}\mathbf{g}), \tag{29}$$

such that the fixed point is

$$\mathbf{z}^* = (\boldsymbol{\alpha} - \mathbf{1}) \oslash (\boldsymbol{\Gamma}\mathbf{g}). \tag{30}$$

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

Putting everything together, this gives the circuit dynamics

$$\mathbf{c}(t) = \boldsymbol{\Gamma}\mathbf{g}(t) \tag{31}$$

$$\tau_{\mathrm{g}}\dot{\mathbf{g}}(t) = (\mathbf{A}\boldsymbol{\Gamma})^{\top}(\mathbf{p} - \mathbf{1}) + \boldsymbol{\Gamma}^{\top}(\mathbf{z} - \lambda) + \boldsymbol{\xi}(t) \tag{32}$$

$$\tau_{\mathrm{p}}\dot{\mathbf{p}}(t) = \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g}) \tag{33}$$

$$\tau_{\mathrm{z}}\dot{\mathbf{z}}(t) = \boldsymbol{\alpha} - \mathbf{1} - \mathbf{z} \odot \mathbf{c}, \tag{34}$$

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

$$\mathbb{E}[\xi_j(t)\xi_{j'}(t')] = 2\tau_{\mathrm{g}}\delta_{jj'}\delta(t-t'). \tag{35}$$

We note that the two constant terms in the dynamics of $\mathbf{g}$ can be grouped into an overall leak term $-\boldsymbol{\Gamma}^{\top}(\mathbf{A}^{\top}\mathbf{1}+\lambda)$. 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.

Appendix 3

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

$$\begin{aligned} \mathbf{c}(t) &= \boldsymbol{\Gamma}\mathbf{g}(t) \\ \tau_{\mathrm{g}}\dot{\mathbf{g}}(t) &= (\mathbf{A}\boldsymbol{\Gamma})^{\top}(\mathbf{p}-\mathbf{1}) - \boldsymbol{\Gamma}^{\top}\boldsymbol{\lambda} \\ \tau_{\mathrm{p}}\dot{\mathbf{p}}(t) &= \mathbf{s} - \mathbf{p}\odot(\mathbf{r}_0 + \mathbf{A}\boldsymbol{\Gamma}\mathbf{g}). \end{aligned} \tag{36}$$

Analysis of a two-cell circuit

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

$$\begin{aligned} c(t) &= \gamma g(t) \\ \tau_{\mathrm{g}}\frac{dg}{dt} &= a\gamma(p-1) - \gamma\lambda \\ \tau_{\mathrm{p}}\frac{dp}{dt} &= s - p(r_0 + a\gamma g). \end{aligned} \tag{37}$$

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

$$\tilde{t} = \frac{a\gamma}{\tau_{\mathrm{g}}}t, \tag{38}$$

and define the dimensionless parameters and input

$$\tilde{\tau} = \frac{\tau_{\mathrm{p}}}{\tau_{\mathrm{g}}}, \tag{39}$$

$$\tilde{r}_0 = \frac{r_0}{\alpha\gamma}, \tag{40}$$

$$\tilde{\lambda} = \frac{\lambda}{a}, \tag{41}$$

$$\tilde{s} = \frac{s}{\alpha\gamma}. \tag{42}$$

Then, the dynamics are

$$\frac{dg}{d\tilde{t}} = p - 1 - \tilde{\lambda} \tag{43}$$

$$\tilde{\tau}\frac{dp}{d\tilde{t}} = \tilde{s} - p(\tilde{r}_0 + g). \tag{44}$$

These dynamics have a single fixed point at

$$g_* = \frac{\tilde{s}}{1+\tilde{\lambda}} - \tilde{r}_0 \tag{45}$$

$$p_* = 1 + \tilde{\lambda}. \tag{46}$$

Linearizing about this fixed point, we have

$$\frac{d}{d\tilde{t}}\begin{pmatrix}\delta g \\ \delta g\end{pmatrix} = \mathbf{M}\begin{pmatrix}\delta g \\ \delta g\end{pmatrix} \tag{47}$$

for

$$\mathbf{M} = \frac{1}{\tilde{\tau}} \begin{pmatrix} 0 & \tilde{\tau} \\ -(1+\tilde{\lambda}) & -(1+\tilde{\lambda})^{-1}\tilde{s} \end{pmatrix}. \tag{48}$$

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

$$\Lambda_{\pm}(\mathbf{M}) = \frac{-\tilde{s} \pm \sqrt{\tilde{s}^2 - 4\tilde{\tau}(1+\tilde{\lambda})^3}}{2\tilde{\tau}(1+\tilde{\lambda})}. \tag{49}$$

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

$$\operatorname{Re}\Lambda_{\pm}(\mathbf{M}) = \frac{-\tilde{s} \pm \operatorname{Re}\sqrt{\tilde{s}^2 - 4\tilde{\tau}(1+\tilde{\lambda})^3}}{2\tilde{\tau}(1+\tilde{\lambda})} \tag{50}$$

$$= \begin{cases} \dfrac{-\tilde{s} \pm \sqrt{\tilde{s}^2 - 4\tilde{\tau}(1+\tilde{\lambda})^3}}{2\tilde{\tau}(1+\tilde{\lambda})} & \tilde{s}^2 - 4\tilde{\tau}(1+\tilde{\lambda})^3 > 0 \\ \dfrac{-\tilde{s}}{2\tilde{\tau}(1+\tilde{\lambda})} & \text{otherwise,} \end{cases} \tag{51}$$

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

$$\Lambda_{\pm}(\mathbf{M})\Big|_{\tilde{s}=0} = \pm 2i\sqrt{\frac{1+\tilde{\lambda}}{\tilde{\tau}}}, \tag{52}$$

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.

Analysis of the full circuit

We now consider the full MAP circuit

$$\begin{aligned} \mathbf{c}(t) &= \mathbf{\Gamma}\mathbf{g}(t) \\ \tau_{\mathrm{g}}\dot{\mathbf{g}}(t) &= (\mathbf{A}\mathbf{\Gamma})^{\top}(\mathbf{p} - \mathbf{1}) - \mathbf{\Gamma}^{\top}\lambda \\ \tau_{\mathrm{p}}\dot{\mathbf{p}}(t) &= \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{A}\mathbf{\Gamma}\mathbf{g}). \end{aligned} \tag{53}$$

We assume that $n_{\mathrm{OSN}} < n_{\mathrm{odor}} \leq n_{\mathrm{g}}$ and that $\mathbf{A}\mathbf{\Gamma}$ is of full row rank, i.e., it has rank n_{OSN}. We recall that $\mathbf{g} \in \mathbb{R}_+^{n_{\mathrm{g}}}$ and $\mathbf{p} \in \mathbb{R}_+^{n_{\mathrm{g}}}$. 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

$$\mathbf{q} \equiv \mathbf{A}\mathbf{\Gamma}\mathbf{g}. \tag{54}$$

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

$$\begin{aligned} \tau_{\mathrm{g}}\dot{\mathbf{q}} &= (\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^{\top}\mathbf{A}^{\top})(\mathbf{p} - \mathbf{1}) - \mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^{\top}\lambda \\ \tau_{\mathrm{p}}\dot{\mathbf{p}} &= \mathbf{s} - \mathbf{p} \odot (\mathbf{r}_0 + \mathbf{q}) \end{aligned} \tag{55}$$

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

$$(\mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^{\top}\mathbf{A}^{\top})(\mathbf{p}^* - \mathbf{1}) = \mathbf{A}\mathbf{\Gamma}\mathbf{\Gamma}^{\top}\lambda \tag{56}$$

$$\mathbf{p}^* \odot (\mathbf{r}_0 + \mathbf{q}^*) = \mathbf{s} \tag{57}$$

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

$$\mathbf{p}^* = \mathbf{1} + (\mathbf{A\Gamma\Gamma}^\top\mathbf{A}^\top)^{-1}\mathbf{A\Gamma\Gamma}^\top\boldsymbol{\lambda}. \tag{58}$$

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

$$\mathbf{q}^* = \mathbf{s} \oslash \mathbf{p}^* - \mathbf{r}_0, \tag{59}$$

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

$$\tilde{t} = \frac{t}{\tau_{\mathrm{g}}}, \tag{60}$$

and

$$\tilde{\tau} = \frac{\tau_{\mathrm{p}}}{\tau_{\mathrm{g}}}, \tag{61}$$

which yields the linearized dynamics

$$\frac{d}{d\tilde{t}}\begin{pmatrix}\delta\mathbf{q}\\ \delta\mathbf{p}\end{pmatrix} = \mathbf{M}\begin{pmatrix}\delta\mathbf{q}\\ \delta\mathbf{p}\end{pmatrix} \tag{62}$$

for

$$\mathbf{M} = \begin{pmatrix}\mathbf{0} & (\mathbf{A\Gamma\Gamma}^\top\mathbf{A}^\top)\\ -\tilde{\tau}^{-1}\operatorname{diag}(\mathbf{p}^*) & -\tilde{\tau}^{-1}\operatorname{diag}(\mathbf{r}_0+\mathbf{q}^*)\end{pmatrix} \tag{63}$$

$$= \begin{pmatrix}\mathbf{0} & (\mathbf{A\Gamma\Gamma}^\top\mathbf{A}^\top)\\ -\tilde{\tau}^{-1}\operatorname{diag}(\mathbf{p}^*) & -\tilde{\tau}^{-1}\operatorname{diag}(\mathbf{s}\oslash\mathbf{p}^*)\end{pmatrix}. \tag{64}$$

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

$$\det(\mathbf{M}-\mu\mathbf{I}_{2n_{\mathrm{OSN}}}) = \det\begin{pmatrix}\mu\mathbf{I}_{n_{\mathrm{OSN}}} & -(\mathbf{A\Gamma\Gamma}^\top\mathbf{A}^\top)\\ \tilde{\tau}^{-1}\operatorname{diag}(\mathbf{p}^*) & \tilde{\tau}^{-1}\operatorname{diag}(\mathbf{s}\oslash\mathbf{p}^*)+\mu\mathbf{I}_{n_{\mathrm{OSN}}}\end{pmatrix} \tag{65}$$

$$= \det[\mu^2\mathbf{I}_{n_{\mathrm{OSN}}} + \mu\tilde{\tau}^{-1}\operatorname{diag}(\mathbf{s}\oslash\mathbf{p}^*) + \tilde{\tau}^{-1}(\mathbf{A\Gamma\Gamma}^\top\mathbf{A}^\top)\operatorname{diag}(\mathbf{p}^*)]. \tag{66}$$

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

$$\det(\mathbf{M}-\mu\mathbf{I}_{2n_{\mathrm{OSN}}}) = \prod_{j=1}^{n_{\mathrm{OSN}}}\left[\mu^2 + \frac{s_j}{\tilde{\tau}p_j^*}\mu + \frac{a_j p_j^*}{\tilde{\tau}}\right], \tag{67}$$

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

$$\mu_{j,\pm} = \frac{1}{2}\left[-\frac{s_j}{\tilde{\tau}p_j^*} \pm \sqrt{\left(\frac{s_j}{\tilde{\tau}p_j^*}\right)^2 - 4\frac{a_j p_j^*}{\tilde{\tau}}}\right], \tag{68}$$

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

$$\mathbf{m} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} \tag{69}$$

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

$$\mathbf{M}\mathbf{m} = \mu\mathbf{m} \tag{70}$$

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

$$\frac{1}{\tau_{\mathrm{g}}}(\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})\mathbf{v} = \mu\mathbf{u} \tag{71}$$

$$-\frac{1}{\tau_{\mathrm{p}}}\operatorname{diag}(\mathbf{p}^{*})\mathbf{u} - \frac{1}{\tau_{\mathrm{p}}}\operatorname{diag}(\mathbf{s} \oslash \mathbf{p}^{*})\mathbf{v} = \mu\mathbf{v}. \tag{72}$$

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 $\mathbf{v}$:

$$\mu^{2}\mathbf{v} + \mu\frac{1}{\tau_{\mathrm{p}}}\operatorname{diag}(\mathbf{s} \oslash \mathbf{p}^{*})\mathbf{v} + \frac{1}{\tau_{\mathrm{p}}\tau_{\mathrm{g}}}\operatorname{diag}(\mathbf{p}^{*})(\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})\mathbf{v} = 0. \tag{73}$$

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

$$\mu^{2} + \frac{a}{\tau_{\mathrm{p}}}\mu + \frac{b}{\tau_{\mathrm{p}}\tau_{\mathrm{g}}} = 0. \tag{74}$$

where we define the coefficients

$$a \equiv \mathbf{v}^{\dagger}\operatorname{diag}(\mathbf{s} \oslash \mathbf{p}^{*})\mathbf{v} \tag{75}$$

and

$$b \equiv \mathbf{v}^{\dagger}\operatorname{diag}(\mathbf{p}^{*})(\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})\mathbf{v}. \tag{76}$$

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

$$\tilde{\lambda} \equiv (\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})^{-1}\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\boldsymbol{\lambda} \tag{77}$$

such that

$$\mathbf{p}^{*} = \mathbf{1} + \tilde{\lambda}. \tag{78}$$

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

$$\operatorname{Re}(b) = \mathbf{v}^{\dagger}\frac{\operatorname{diag}(\mathbf{p}^{*})(\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top}) + (\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})\operatorname{diag}(\mathbf{p}^{*})}{2}\mathbf{v} \tag{79}$$

$$= \mathbf{v}^{\dagger}(\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})\mathbf{v} + \mathbf{v}^{\dagger}\frac{\operatorname{diag}(\tilde{\lambda})(\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top}) + (\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})\operatorname{diag}(\tilde{\lambda})}{2}\mathbf{v}, \tag{80}$$

while

$$\operatorname{Im}(b) = \mathbf{v}^{\dagger}\frac{\operatorname{diag}(\mathbf{p}^{*})(\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top}) - (\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})\operatorname{diag}(\mathbf{p}^{*})}{2i}\mathbf{v} \tag{81}$$

$$= \mathbf{v}^{\dagger}\frac{\operatorname{diag}(\tilde{\lambda})(\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top}) - (\mathbf{A}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^{\top}\mathbf{A}^{\top})\operatorname{diag}(\tilde{\lambda})}{2i}\mathbf{v}. \tag{82}$$

Then, solving the quadratic equation for μ, we have

$$\mu = \frac{1}{\tau_{\mathrm{p}}} \frac{-a \pm \sqrt{a^2 - 4\tilde{\tau} b}}{2}, \tag{83}$$

for

$$\tilde{\tau} \equiv \frac{\tau_{\mathrm{p}}}{\tau_{\mathrm{g}}}. \tag{84}$$

This gives

$$\mathrm{Re}(\mu) = -\frac{a}{2\tau_{\mathrm{p}}} \pm \frac{1}{2\tau_{\mathrm{p}}} \mathrm{Re}\sqrt{a^2 - 4\tilde{\tau} b} \tag{85}$$

$$= -\frac{a}{2\tau_{\mathrm{p}}} \pm \frac{1}{2\tau_{\mathrm{p}}} \sqrt{\frac{\sqrt{(a^2 - 4\tilde{\tau}\,\mathrm{Re}(b))^2 + 16\tilde{\tau}^2\,\mathrm{Im}(b)^2} + a^2 - 4\tilde{\tau}\,\mathrm{Re}(b)}{2}}. \tag{86}$$

If $\mathrm{Im}(b) = 0$, 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

$$a > \sqrt{\frac{\sqrt{(a^2 - 4\tilde{\tau}\,\mathrm{Re}(b))^2 + 16\tilde{\tau}^2\,\mathrm{Im}(b)^2} + a^2 - 4\tilde{\tau}\,\mathrm{Re}(b)}{2}}. \tag{87}$$

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

$$\mathrm{Re}(b) > 0 \tag{88}$$

and

$$\sqrt{\tilde{\tau}}\,|\,\mathrm{Im}(b)| < a\sqrt{\mathrm{Re}(b)}. \tag{89}$$

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 $\mathrm{span}(\mathbf{\Gamma}^\top) \supset \mathrm{span}[(\mathbf{A}\mathbf{\Gamma})^\top]$. We can first exclude components outside $\mathrm{span}(\mathbf{\Gamma}^\top)$, as they will be unchanged by the dynamics and do not affect the concentration estimate. Recalling the non-negativity constraint, components in $\mathrm{span}(\mathbf{\Gamma}^\top) \setminus \mathrm{span}[(\mathbf{A}\mathbf{\Gamma})^\top]$ will decay linearly to zero. Then, the components in $\mathrm{span}[(\mathbf{A}\mathbf{\Gamma})^\top]$ 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.

Appendix 4

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

$$\mathbf{s} \mid \mathbf{c} \sim \mathcal{N}(\mathbf{r}_0 + \mathbf{A}\mathbf{c}, \sigma^2 \mathbf{I}_{n_{\mathrm{OSN}}}) \tag{90}$$

variance σ^2, and an exponential prior

$$\mathbf{c} \sim \mathrm{Exp}(\lambda). \tag{91}$$

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

$$\dot{\mathbf{c}}(t) = \frac{1}{\sigma^2}\mathbf{A}^\top[\mathbf{s} - (\mathbf{r}_0 + \mathbf{A}\mathbf{c})] - \lambda, \tag{92}$$

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

$$\begin{aligned} \mathbf{c}(t) &= \mathbf{\Gamma}\mathbf{g}(t) \\ \tau_{\mathrm{g}}\dot{\mathbf{g}}(t) &= \frac{1}{\sigma^2}(\mathbf{A}\mathbf{\Gamma})^\top \mathbf{p} - \mathbf{\Gamma}^\top \lambda \\ \tau_{\mathrm{p}}\dot{\mathbf{p}}(t) &= -\mathbf{p} + \mathbf{s} - (\mathbf{r}_0 + \mathbf{A}\mathbf{\Gamma}\mathbf{g}). \end{aligned} \tag{93}$$

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.

Appendix 5

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

$$p(c_i) = \varpi e^{-\lambda c_i} + (1 - \varpi)\delta(c_i), \tag{94}$$

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. (2022), 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:

$$\mathbf{c} = f(\mathbf{u}), \tag{95}$$

where

$$f(u) = \begin{cases} 0 & u < u_0 \\ u - u_0 & u \geq u_0 \end{cases} \tag{96}$$

is the soft-thresholding function for threshold

$$u_0 = -\frac{1}{\lambda} \log \varpi. \tag{97}$$

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

$$\dot{\mathbf{u}}(t) = \left[\nabla_{\mathbf{c}} \log p(\mathbf{s} \mid \mathbf{c})\right]_{\mathbf{c}=f(|\mathbf{u}|)} \odot \Theta(|\mathbf{u}| - \mathbf{u}_0) - \lambda \operatorname{sign}(\mathbf{u}) + \boldsymbol{\eta}(t) \tag{98}$$

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. (2022) 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

$$\dot{\mathbf{u}}(t) = \{\mathbf{A}^\top[\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}f(|\mathbf{u}|)) - \mathbf{1}]\} \odot \Theta(|\mathbf{u}| - \mathbf{u}_0) - \lambda \operatorname{sign}(\mathbf{u}) + \boldsymbol{\eta}(t) \tag{99}$$

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 2 by writing

$$\mathbf{u} = \boldsymbol{\Gamma}\mathbf{g}, \tag{100}$$

we will have

$$\tau_{\mathrm{g}}\dot{\mathbf{g}}(t) = \boldsymbol{\Gamma}^\top\left[\{\mathbf{A}^\top[\mathbf{s} \oslash (\mathbf{r}_0 + \mathbf{A}f(|\boldsymbol{\Gamma}\mathbf{g}|)) - \mathbf{1}]\} \odot \Theta(|\boldsymbol{\Gamma}\mathbf{g}| - \mathbf{u}_0) - \lambda \operatorname{sign}(\boldsymbol{\Gamma}\mathbf{g})\right] + \boldsymbol{\xi}(t), \tag{101}$$

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

$$\mathbb{E}[\xi_j(t)\xi_{j'}(t')] = 2\tau_{\mathrm{g}}\delta_{jj'}\delta(t - t'). \tag{102}$$

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

$$\sum_k \Gamma_{ki} A_{jk} \Theta(|u_k| - u_0) \tag{103}$$

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

$$\mathbf{c}(t) = f(|\mathbf{\Gamma g}(t)|) \tag{104}$$

$$\tau_{\mathrm{g}} \dot{\mathbf{g}}(t) = \mathbf{\Gamma}^{\top} \left[[\mathbf{A}^{\top}(\mathbf{p} - \mathbf{1})] \odot \Theta(|\mathbf{\Gamma g}| - \mathbf{u}_0) - \lambda \operatorname{sign}(\mathbf{\Gamma g}) \right] dt + \boldsymbol{\xi}(t) \tag{105}$$

$$\tau_{\mathrm{p}} \dot{\mathbf{p}}(t) = \mathbf{s} - \mathbf{p} \odot [\mathbf{r}_0 + \mathbf{A} f(|\mathbf{\Gamma g}|)]. \tag{106}$$

This issue illustrates the limitations of the simple, linear approach to distributing the neural code used in Appendix 2. 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.

Appendix 6

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. (2020). 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.

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 $100mL/min$, 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* (Pnevmatikakis and Giovannucci, 2017). 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 (Mathis et al., 2016). 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.

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 $\mathrm{Gamma}(0.37, 0.36)$ distribution.

Appendix 7

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 less than 6000 CPU-hours of total compute time. Code will be made available on GitHub upon acceptance.

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, Kapoor, and Urban (2008) 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 where 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 93. For plotting purposes, we normalized the range of firing in Figure 2B to the maximum firing rate for each circuit model.

Capacity simulations

In our capacity simulations, we use $n_{\text{OSN}} = 300$ and $n_{\text{odor}} = 500$, 1000, or 2000. We take the odor stimulus to be a rectangular pulse, with varying numbers of randomly-selected odors appearing at concentration $c_j = 40$. 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 $\mathbf{\Gamma}$:

- For the one-to-one code, we let $n_{\text{g}} = n_{\text{odor}}$, and simply set $\mathbf{\Gamma} = \mathbf{I}_{n_{\text{odor}}} / \max(A_{ij})$.
- For the naïvely distributed code, we let $n_{\text{g}} = 5n_{\text{odor}}$, and choose $\mathbf{\Gamma}$ as follows: We sample a random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_{\text{g}}}$ with $\mathbf{Q}\mathbf{Q}^\top = \mathbf{I}_{n_{\text{odor}}}$ by drawing a Gaussian matrix and orthogonalizing its rows. Then, we define $\mathbf{\Gamma} = \mathbf{Q} / \max\{|(\mathbf{AQ})_{ij}|\}$.
- For the geometry-aware code, we let $n_{\text{g}} = 5n_{\text{odor}}$, and choose $\mathbf{\Gamma}$ as follows: We sample a random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_{\text{g}}}$ 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 $\mathbf{\Gamma} = \mathbf{BQ} / \max\{|(\mathbf{ABQ})_{ij}|\}$.

Our normalization convention for $\mathbf{\Gamma}$ 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.

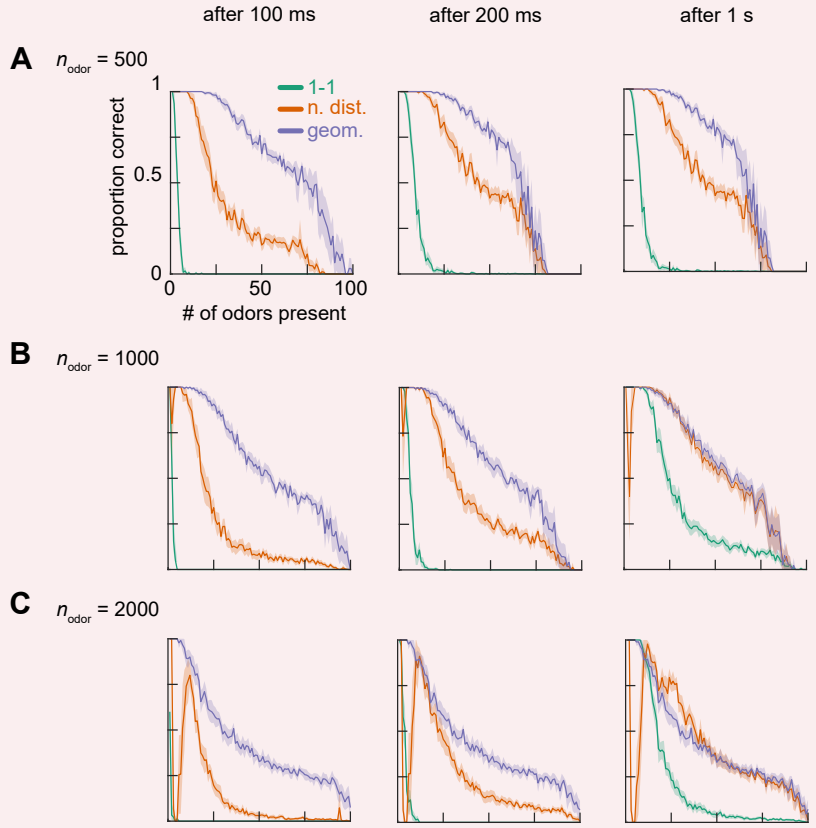

Appendix 7—figure 1. A. Fraction of odors correctly detected within 100 ms (*left*), 200 ms (*center*), and 1 s (*right*) after odor onset as a function of the number of odors present, with $n_{\text{odor}} = 500$ possible odors, for models with one-to-one, naïvely distributed, and geometry-aware codes. **B**. As in **A**, but for $n_{\text{odor}} = 1000$. This matches Fig. 3B. **C**. As in **A**, but for $n_{\text{odor}} = 2000$.

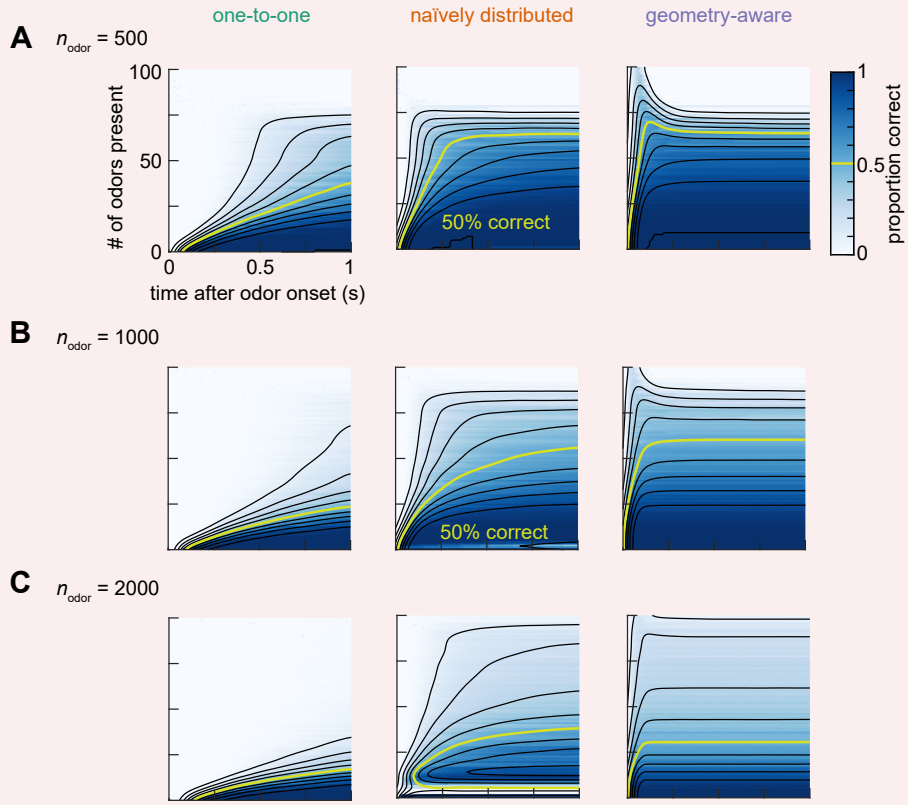

Appendix 7—figure 2. A. Heatmap with overlaid smoothed contours of the fraction of odors correctly detected as a function of number of present odors and time window out of a panel of $n_{\text{odor}} = 500$ possible odors for models with one-to-one (*left*), naïvely distributed (*center*), and geometry-aware (*right*) codes. **B**. As in **A**, but for $n_{\text{odor}} = 1000$. This matches Fig. 3B. **C**. As in **A**, but for $n_{\text{odor}} = 2000$.

Sampling simulations

In our sampling simulations, we use $n_{\text{OSN}} = 300$, $n_{\text{odor}} = 500$ or 1000, and $n_{\text{g}} = 5n_{\text{odor}}$. 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. 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 $\mathbf{\Gamma}$ as follows:

- For the one-to-one code, we let $n_{\text{g}} = n_{\text{odor}}$, and simply set $\mathbf{\Gamma} = \mathbf{I}_{n_{\text{odor}}} / \max(A_{ij})$.
- For the naïvely distributed code, we let $n_{\text{g}} = 5n_{\text{odor}}$, and choose $\mathbf{\Gamma}$ as follows: We sample a random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_{\text{g}}}$ with $\mathbf{Q}\mathbf{Q}^{\top} = \mathbf{I}_{n_{\text{odor}}}$ by drawing a Gaussian matrix and orthogonalizing its rows. Then, we define $\mathbf{\Gamma} = \mathbf{Q} / \max\{|(\mathbf{AQ})_{ij}|\}$.
- For the geometry-aware code, we let $n_{\text{g}} = 5n_{\text{odor}}$, and choose $\mathbf{\Gamma}$ as follows: We sample a random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_{\text{g}}}$ 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 $\mathbf{\Gamma} = \mathbf{BQ} / \max\{|(\mathbf{ABQ})_{ij}|\}$.

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

$$\mu_j(\tau) = \frac{\Delta t}{\tau} \sum_{t_{\text{low}} \le t \le t_{\text{low}}+\tau} \hat{c}_j(t) \tag{107}$$

and

$$\sigma_j^2(\tau) = \frac{\Delta t}{\tau} \sum_{t_{\text{low}} \le t \le t_{\text{low}}+\tau} \hat{c}_j(t)^2 - \mu_j(\tau)^2, \tag{108}$$

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. 3, we show a preliminary experiment with an alternative, mostly non-negative choice for $\mathbf{\Gamma}$. This gives the following three models:

- For the one-to-one code, we let $n_{\text{g}} = n_{\text{odor}}$, and simply set $\mathbf{\Gamma} = \mathbf{I}_{n_{\text{odor}}} / \max(A_{ij})$.
- For the naïvely distributed code, we let $n_{\text{g}} = 5n_{\text{odor}}$, and choose $\mathbf{\Gamma}$ as follows: We sample a sparse, non-negative random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_{\text{g}}}$ with entries that are non-zero with density 0.15, and the non-zero entries are drawn uniformly on $[0, 1]$. Then, we define $\mathbf{\Gamma} = \mathbf{Q} / \max\{|(\mathbf{AQ})_{ij}|\}$.
- For the geometry-aware code, we let $n_{\text{g}} = 5n_{\text{odor}}$, and choose $\mathbf{\Gamma}$ as follows:We sample a sparse, non-negative random matrix $\mathbf{Q} \in \mathbb{R}^{n_{\text{odor}} \times n_{\text{g}}}$ 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 $\mathbf{\Gamma} = \mathbf{BQ} / \max\{|(\mathbf{ABQ})_{ij}|\}$.

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 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 $\mathbf{\Gamma}$ that yields a fully non-negative, sparse synaptic weight matrix $\mathbf{A\Gamma}$.

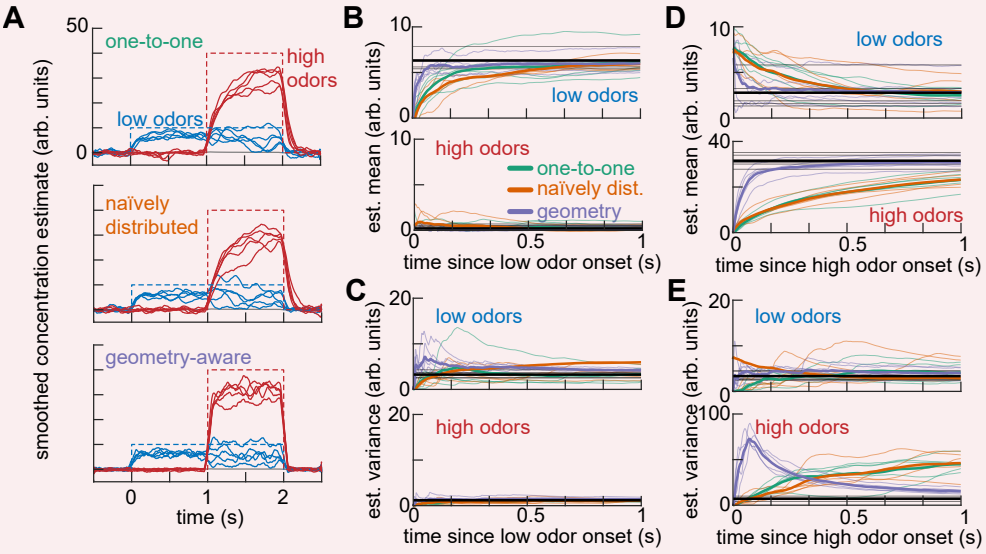

Appendix 7—figure 3. Langevin sampling of the posterior using sparse non-negative distributed codes. Here, we use a simple concentration estimation task in which 5 randomly-selected 'low' odors out of a panel of 500 appear at concentration 10 at time 0 s, and then a further 5 randomly-selected 'high' 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 7 for details of our numerical methods.