Across-animal odor decoding by probabilistic manifold alignment

Anonymous Author(s)
Affiliation
Address
email

Abstract

Identifying the common structure of neural dynamics across subjects is key for extracting unifying principles of brain computation and for many brain machine interface applications. Here, we propose a novel probabilistic approach for aligning stimulus-evoked responses from multiple animals in a common low dimensional manifold and use hierarchical inference to identify which stimulus drives neural activity in any given trial. Our probabilistic decoder is robust to a range of features of the neural responses and significantly outperforms existing neural alignment procedures. When applied to recordings from the mouse olfactory bulb, our approach reveals low-dimensional population dynamics that are odor specific and have consistent structure across animals. Thus, our decoder can be used for increasing the robustness and scalability of neural-based chemical detection.

1 Introduction

Collective network dynamics are the foundation of neural computation, from early sensory encoding [1], to working memory [2], decision making [3], or motor control [4]. Neural population activity often has low-dimensional structure that is qualitatively preserved across sessions [3], behavioral states [5], and even animals [6, 7]. Aligning neural datasets from multiple recordings and animal subjects into a common latent space could provide a powerful tool for extracting unifying principles of brain computations. Nonetheless, progress is hampered by the lack of robust statistical tools for extracting shared neural population dynamics across datasets.

The alignment of low-dimensional neural manifolds is equally important in brain-computer interface (BCI) applications, where the assumption of a common low dimensional structure in neural responses is used to compensate for instability in neural recordings and ensure robust performance over time [8, 9]. This is not only restricted to motor control; in the sensory domain, a new generation of chemical detectors exploit the unsurpassed sensitivity of rodent olfactory receptors, but bypass the limitations of training animals to report odor identity by directly decoding it from neural responses [10–12]. Unfortunately, the applicability of this idea is limited by the need to learn the mapping between neural activity and chemical identity on an animal-by-animal basis, which involves costly data collection [13]. Statistical tools that make use of data from previous animals to quickly calibrate the decoder in a new animal could dramatically increase the practical use of such technologies.

The basic structure of the rodent olfactory bulb (OB) is well-understood [14]: a vast number of different olfactory receptors project to distinct OB subpopulations (glomeruli) whose spatial organization is largely preserved across animals [15, 16]. Odor identity is encoded in low-dimensional transients of their population dynamics [1, 17, 12], though how the glomerular topography translates into stimulus-dependent network dynamics is less clear. Nonetheless, given the anatomical stereotypy it is likely that the OB response to individual odors is itself somewhat stereotyped across animals. With appropriate alignment, one could construct ‘universal’ odor decoders, applicable across animals.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

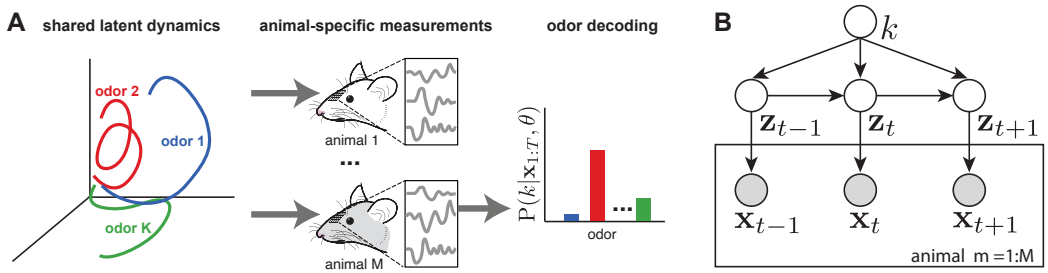

Figure 1: **Overview of the model. A.** Different odors elicit distinct trajectories in a shared neural latent dynamical space, with animal-specific measurements resulting in distinct observations. The goal is to decode the identity of a test odor, given parameters estimated from a dataset that combines measurements from all animals. **B.** amLDS graphical model with latent dynamics $\mathbf{z}_t$ specified by odor-specific k parameters, and animal-specific observation model for measurements $\mathbf{x}_t$.

Here we propose a novel probabilistic approach to aligning observations across neural recordings taken from different mice so as to extract commonalities in their circuit-level representations of different odors. Our estimation procedure, which uses our 'aligned mixture of latent dynamical systems' (amLDS) graphical model, learns independent mappings of different recordings into a common latent manifold. In this space, neural trajectories are similar across animals but distinct across odors which allows for robust odor identification. Moreover, our probabilistic decoder provides uncertainty estimates which can be used for optimal decision making. We used extended numerical simulations to evaluate the properties of the decoder in response to a range of biologically relevant manipulations, such as variations in measurement noise, dimensionality and across trial variability in the latent dynamics, (mis)alignment of electrodes across animals, etc. We then tested our procedure on recordings from the OB of mice in response to several monomolecular odors. When fit to this data, our model confirmed that encoding of odors relies on transient dynamics in a shared low-dimensional manifold and that these dynamics are largely preserved across animals. Taking advantage of this reproducible structure, we were able to obtain remarkable improvements in odor identification performance relative to alternative approaches.

2 Shared odor encoding dynamics across animals

We start by mathematically formalizing the idea of shared stimulus-specific neural dynamics with animal-specific measurements (Fig. 1A). The amLDS graphical model assumes d-dimensional latent variables $\mathbf{z}_t$, which evolve according to stochastic linear dynamics with odor-specific parameters (Fig. 1B):

$$\mathbf{z}_t = \mathbf{A}_k \mathbf{z}_{t-1} + \mathbf{b}_{k,t} + \mathbf{w}_t, \tag{1}$$

where $t = \{2 \dots T\}$ indexes time within trial; $\mathbf{A}_k$ and $\mathbf{b}_k$ denote parameters specifying the dynamics in response to odor k, with independent noise, $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$. The prior for the initial conditions x_1 is normal, with zero mean and covariance $\mathbf{Q}_\mathbf{0}$.

The shared odor-specific dynamics are mapped into measured neural responses via an animal specific linear Gaussian observation model, shared across all odor conditions

$$\mathbf{x}_t^{(m)} = \mathbf{C}_m \mathbf{z}_t + \mathbf{v}_t, \tag{2}$$

with i.i.d. noise $\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_m)$; $\mathbf{C}_m$ and $\mathbf{R}_m$ are observation model parameters specific to animal m. We will use N_m to denote the size of observations for animal m. Note that the animal-specific observation model means that it is not strictly required for the neural measurements to be tightly matched. For instance, different animals may have different numbers of measurement channels (e.g. due to some electrodes becoming damaged). Moreover, the model accounts for variability in the SNR of different electrodes (across channels and devices). Assuming that the measurement process is consistent in terms of the properties of the electrodes and their alignment on the olfactory bulb surface, one can additionally define a prior that encourages $\mathbf{C}$s to be similar across animals. To keep the procedure general, we only consider the unconstrained version here.

Parameter learning

The dataset $\mathcal{D}$ combines measurements from M animals, jointly covering all K stimulus conditions. For each animal-stimulus pair, (m, k), we have $I_{m,k}$ trials.[1] The goal of the parameter estimation procedure is to derive maximum likelihood estimates of all model parameters, jointly denoted by $\theta = \{\mathbf{A}_{1:K}, \mathbf{b}_{1:K}, \mathbf{Q}_{0:K}, \mathbf{C}_{1:M}, \mathbf{R}_{1:M}\}$. We use expectation maximization (EM) to optimize the parameters, by introducing the variational distribution $q(\mathbf{z})$ and using it to lower-bound the marginal likelihood [18]:[2]

$$\begin{aligned}
\log \mathrm{P}(\mathcal{D}|\theta) &= \sum_{i,m,k} \int_{\mathbf{z}} \log \mathrm{P}(\mathbf{x}|\theta) q(\mathbf{z}) \mathbf{dz} \\
&= \sum_{i,m,k} \int_{\mathbf{z}} q(\mathbf{z}) \left(\log \mathrm{P}(\mathbf{z}, \mathbf{x}|\theta) - \log \mathrm{P}(\mathbf{z}|\mathbf{x}, \theta) \right) \mathbf{dz} \\
&= \sum_{i,m,k} \left(\int_{\mathbf{z}} q(\mathbf{z}) \log \frac{\mathrm{P}(\mathbf{x}, \mathbf{z}|\theta)}{q(\mathbf{z})} \mathbf{dz} - \sum_{i,m,k} \int_{\mathbf{z}} q(\mathbf{z}) \log \frac{\mathrm{P}(\mathbf{z}|\mathbf{x}, \theta)}{q(\mathbf{z})} \mathbf{dz} \right) \\
&= \mathcal{L}(q, \theta) + \sum_{i,m,k} \mathrm{KL}(q(\mathbf{z}) || \mathrm{P}(\mathbf{z}|\mathbf{x}, \theta)).
\end{aligned}$$

Due to the Markovian structure of the dynamics, the joint probability of the latent and observations has a relatively simple joint Gaussian structure. The E-step optimizes $\mathcal{L}$ w.r.t. q, by setting $q(\mathbf{z}) = \mathrm{P}(\mathbf{z}|\mathbf{x}, \theta)$; this is computed by traditional Kalman smoothing using the appropriate trial-specific parameters (Suppl. S1).

The M-step optimizes $\mathcal{L}$ w.r.t. to θ. The cost function decomposes into the sum of terms that only depend on the subject-specific observation parameters $(\mathbf{C}_m, \mathbf{R}_m)$ and the condition-specific latent dynamic parameters $(\mathbf{A}_k, \mathbf{Q}_k, \mathbf{Q}_0, \mathbf{b}_{k,t})$, respectively, and can be optimized separately. Taking the derivative of the loss and setting it to zero results in parameter updates:

$$\mathbf{b}_{k,1}^{new} = \frac{1}{I_k} \sum_i \mathbb{E}[\mathbf{z}_1] \tag{3}$$

$$\mathbf{Q}_0^{new} = \frac{1}{I_k} \sum_i \left(\mathbb{E}[\mathbf{z}_1 {\mathbf{z}_1}^\top] - \mathbb{E}[\mathbf{z}_1] \mathbb{E}[\mathbf{z}_1]^\top \right), \tag{4}$$

where the indicator i covers all trials of the same stimulus type, across all animals $I_k = \sum_m I_{m,k}$; expectations are taken under posterior for trial i (explicit indexing omitted for brevity).

The same set of trials are used for updating the latent dynamics, and the input drive:

$$\mathbf{A}_k^{new} = \left(\sum_{i=1}^{I_k} \sum_{t=2}^{T} \left(\mathbb{E}[\mathbf{z}_t {\mathbf{z}_{t-1}}^\top] - \mathbf{b}_{k,t} \mathbb{E}[{\mathbf{z}_{t-1}}^\top] \right) \right) \left(\sum_{k=1}^{I_k} \sum_{t=2}^{T} \mathbb{E}[\mathbf{z}_{t-1} {\mathbf{z}_{t-1}}^\top] \right)^{-1} \tag{5}$$

$$\mathbf{b}_{k,t}^{new} = \frac{1}{I_k} \sum_{i=1}^{I_k} \left(\mathbb{E}[\mathbf{z}_t] - \mathbf{A}_k^{new} \mathbb{E}[\mathbf{z}_{t-1}] \right). \tag{6}$$

Finally, the updated noise covariance is computed as

$$\begin{aligned}
\mathbf{Q}_k^{new} = \frac{1}{I_k(T-1)} \sum_{i=1}^{I_k} \sum_{t=2}^{T} \big(&\mathbb{E}[\mathbf{z}_t {\mathbf{z}_t}^\top] - \mathbf{A}_k^{new} \mathbb{E}[\mathbf{z}_{t-1} {\mathbf{z}_t}^\top] - \mathbb{E}[\mathbf{z}_t {\mathbf{z}_{t-1}}^\top] \mathbf{A}_k^{new\top} \\
&+ \mathbf{A}_k^{new} \mathbb{E}[\mathbf{z}_{t-1} {\mathbf{z}_{t-1}}^\top] \mathbf{A}_k^{new\top} - \mathbb{E}[\mathbf{z}_t] \mathbf{b}_{k,t}^{new\top} - \mathbf{b}_{k,t}^{new} \mathbb{E}[\mathbf{z}_t]^\top + \mathbf{A}_k^{new} \mathbb{E}[\mathbf{z}_{t-1}] \mathbf{b}_{k,t}^{new\top} \\
&+ \mathbf{b}_{k,t}^{new} \mathbb{E}[\mathbf{z}_{t-1}]^\top \mathbf{A}_k^{new\top} + \mathbf{b}_{k,t}^{new} \mathbf{b}_{k,t}^{new\top} \big).
\end{aligned} \tag{7}$$

[1] Note that it is not necessary that all stimuli are presented to all animals, nor that the number of stimulus presentations is identical across animals and conditions.

[2] To keep notation simple, we drop the explicit indexing by time, trial, stimulus type, and animal of $\mathbf{x}$ and $\mathbf{z}$.

The updates for the observation model are the standard Kalman ones, with the distinction that the index i iterates across all trials and stimulus conditions for a particular animal (total number $I_m = \sum_k I_{m,k}$):

$$\mathbf{C}_m^{new} = \left(\sum_{i=1}^{I_m} \sum_{t=1}^{T} \mathbf{x}_t \mathbb{E}[\mathbf{z}_t] \right) \left(\sum_{i=1}^{I_m} \sum_{t=1}^{T} \mathbb{E}[\mathbf{z}_t \mathbf{z}_t{}^\top] \right)^{-1} \tag{8}$$

$$\mathbf{R}_m^{new} = \frac{1}{I_m T} \sum_{i,t} \left(\mathbf{x}_t \mathbf{x}_t{}^\top - \mathbf{C}_m^{new} \mathbb{E}[\mathbf{z}_t] \mathbf{x}_t{}^\top - \mathbf{x}_t \mathbb{E}[\mathbf{z}_t]^\top \mathbf{C}_m^{new\top} + \mathbf{C}_m^{new} \mathbb{E}[\mathbf{z}_t \mathbf{z}_t{}^\top] \mathbf{C}_m^{new\top} \right) \tag{9}$$

All the required posterior moments are obtained during the E step via Kalman smoothing as $\mathbb{E}[\mathbf{z}_t] = \mu_{t|T}$, $\mathbb{E}[\mathbf{z}_t \mathbf{z}_t{}^\top] = \mathbf{\Sigma}_{t|T} - \boldsymbol{\mu}_{t|T} \boldsymbol{\mu}_{t|T}{}^\top$ (see Suppl. S1 for further details).

As an important practical side note, the quality of the estimated parameter depends critically on a good initialization of θ. We achieve this by fitting a simpler factor analysis model (FA), which ignores all temporal dependencies. This directly yields initial conditions for $\mathbf{C}_m$ and $\mathbf{R}_m$. The inputs $\mathbf{b}_{k,t}$ are initialized to the average of the corresponding inferred latent variable $\mathbf{z}_t$, using all trials from stimulus condition k, while parameters $\mathbf{A}_k$ and $\mathbf{Q}_{0:K}$ are initialized heuristically by linear regression using the FA-inferred posterior mean estimates of the latents.

We used Bayesian model comparison to determine the dimensionality of the latent space from data (evaluated on on a separate validation set). As a more intuitive measure of goodness-of-fit, we also estimated the reconstruction error when predicting unobserved measurement dimensions (also known as leave-neuron-out error), commonly used for assessing the quality of models of neural population dynamics [19, 20]. More specifically, we use trials from the validation dataset, infer the posterior over the latent trajectory given data measurements in all electrodes except the j-th by Kalman smoothing $\mathrm{P}\left(\mathbf{z}_{1:T} | \mathbf{x}_{\neg j}, \theta\right)$, which is used to predict the values of the response of the jth neuron, $\hat{\mathbf{x}}_t^j = \mathbf{C}_m^j \mu_{t|T}^{(\neg j)}$, where $\mathbf{C}_m^j$ is the jth row of $\mathbf{C}_m$. We report reconstruction error as the squared distance between the reconstructed and measured activity, $\langle \| \mathbf{x}_t^j - \hat{\mathbf{x}}_t^j \|_2^2 \rangle_{j,i}$, where $\| \cdot \|_2$ is the L2 norm and $\langle \cdot \rangle_{j,i}$ is the mean over trials and choices for j.

Odor decoding as hierarchical inference

At test time, one needs to infer the identity of the stimulus given the observations $\mathbf{x}_{1:T}$ in animal M, and the estimated parameters θ, obtained by Bayes rule, $\mathrm{P}(k|\mathbf{x}_{1:T}, \theta) \propto \mathrm{P}(\mathbf{x}_{1:T}|k, \theta)\mathrm{P}(k)$. This requires estimating the marginal likelihood for each possible stimulus condition, which can be computed iteratively during the Kalman filtering procedure. More specifically, using the chain rule and then the Markov structure of the stimulus-specific LDS, we have

$$\mathrm{P}(\mathbf{x}_{1:T}|k, \theta) = \prod_{t=1}^{T} \mathrm{P}(\mathbf{x}_t|\mathbf{x}_{1:t-1}, k, \theta) = \prod_{t=1}^{T} \int \mathrm{P}(\mathbf{x}_t, \mathbf{z}_t|\mathbf{x}_{1:t-1}, k, \theta) \mathbf{dz_t}.$$

Taking the logarithm, rearranging the terms and taking advantage of the fact that the posterior marginals are normal with parameters obtained by Kalman filtering (see Suppl. Info. S1), $\mathrm{P}(\mathbf{z}_t|x_{1:t-1}, k, \theta) = \mathcal{N}(\mu_{t|t-1}, \Sigma_{t|t-1})$, yields:

$$\begin{aligned} \log \mathrm{P}(\mathbf{x}_{1:T}|k, \theta) &= \sum_{t=1}^{T} \log \mathrm{P}(\mathbf{x}_t|\mathbf{z}_t, \mathbf{x}_{1:t-1}, \theta_k) \mathrm{P}(\mathbf{z}_t|x_{1:t-1}, k, \theta) \\ &= \sum_{t=1}^{T} \log \mathcal{N}\left(\mathbf{x}_t; \mathbf{C}_m \mu_{t|t-1}, \mathbf{C}_m \Sigma_{t|t-1} \mathbf{C}_m{}^\top + \mathbf{R}_m\right) \end{aligned}$$

Lastly, here we assume a uniform prior over possible stimuli, to match the statistics of our OB dataset. Nonetheless, for real world applications the prior should reflect natural stimulus statistics.

3 Numerical experiments on simulated data

To validate the parameter estimation and decoding procedure, we constructed artificial datasets with the same statistical regularities assumed in the graphical model. More specifically, we defined a shared

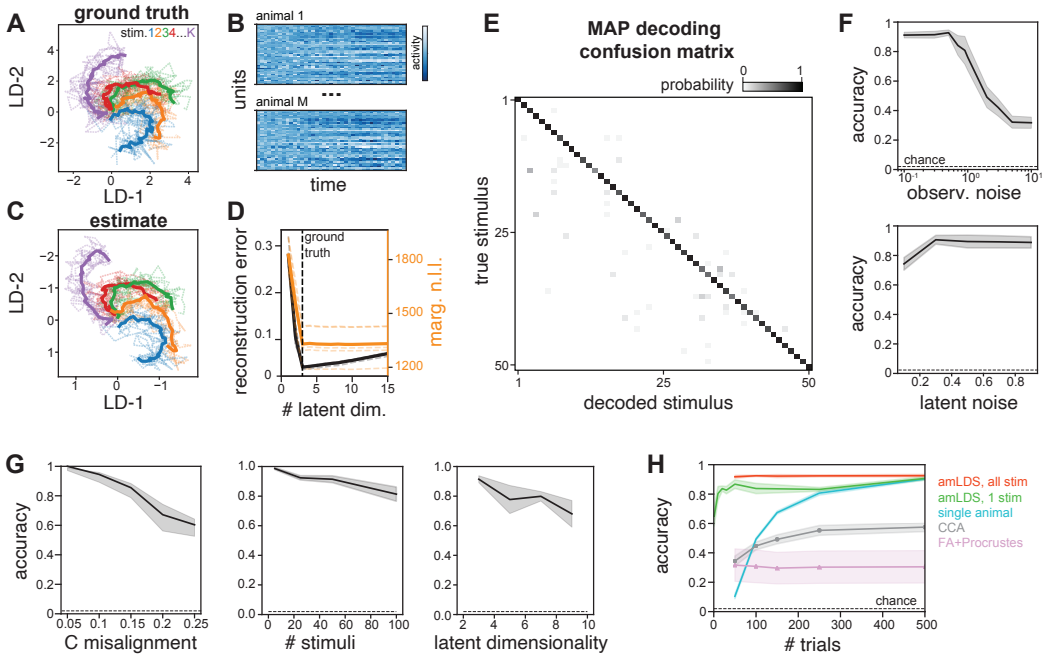

Figure 2: **Parameter estimation and decoding on simulated data**. **A.** Example latent dynamics for different odor stimuli; solid lines indicate the average latent trajectory and dashed lines are five random trials. **B.** Example measurements of an example trial in two animals. **C.** Latents recovered with amLDS: dashed lines shows posterior mean of individual trials, solid lines show across-trials average. **D.** Reconstruction error as a function of assumed latent dimensionality, d (ground truth $d = 3$). **E.** amLDS decoding performance measured using MAP estimates of the stimulus; $K = 50$. **F.** Decoding accuracy as a function of the magnitude of noise in the observations (top), and the latent dynamics (bottom). **G.** Decoding accuracy for parametric changes in model hyperparameters: alignment of axes across animals, $\mathbf{C}$ (left), the number of stimulus classes, K (middle) and the dimensionality of the latent space, d (right). **H.** Decoding accuracy as a function of the amount of data available for a target animal ($M = 5$) for amLDS when the target data covers all stimuli (red) or only one of the $K = 50$ stimuli (green), compared to using the same probabilistic model in only the target animal (blue), alignment by canonical correlation analysis (gray) and Procrustes alignment combined with factor analysis dimensionality reduction (pink). Solid lines show mean accuracy, shaded areas s.e.m. estimated across 5 independent experiments (random seeds).

low dimensional manifold ($d = 3$) and embedded $K = 50$ latent trajectories evolving over $T = 41$ time steps (Fig.2A). The stimulus dependent inputs $\mathbf{b}_k$ were constructed using a common template with stimulus specific amplitude (individual dimensions scaled by value drawn from $\mathcal{N}(1; 0.0004)$) and rotation (evenly spaced over 170 degrees). The other latent dynamics parameters were set randomly: matrices $\mathbf{A}_k$ have diagonal entries drawn from $\mathcal{N}(0.4; 0.01)$ and off-diagonal drawn from $\mathcal{N}(0; 0.04)$ and the noise covariances $\mathbf{Q}_k$ and $\mathbf{Q}_0$ are diagonal with variances drawn from $\mathcal{N}(0.55; 0.0025)$. To model across-animal variability in the observation model we started from a prototype projection matrix $\mathbf{C}^*$, which we then morph into animal specific parameters $\mathbf{C}_m$ by corrupting the individual latent axes with independent additive noise $\mathcal{N}(0; \alpha^2)$, and rescaling their norm to a noisy value around 1.0 (gaussian noise with variance 0.03). Parameter α controls the degree of alignment of measurements across animals, with default value 0.1. Finally, the observation noise covariance is set to be diagonal with variances set to the absolute value of independent draws from $\mathcal{N}(0; 0.25)$ (Fig.2B). This setup results in latent dynamics qualitatively similar to odor responses, and has the advantage of explicitly parametrizing the similarity across animals and stimuli.

As a first test of the estimator we simulated 5 animals, with $I_{m,k} = 50$ trials for each stimulus condition, which is on the same order of magnitude to the amount of data one may be able to collect in experiments. amLDS has low computational and memory demands; on a 2.9GHz CPU it takes 58 minutes for parameter learning (with 12500 trials; average over 5 runs) and 4 seconds to infer

stimulus class in a given trial (averaged over 2500 trials). In this setup, model comparison based on reconstruction error was able to correctly discover the dimensionality of the latent dynamics (Fig.2D) and reconstruct the underlying trajectories (Fig.2C), up to the expected degeneracies in the latent space scale and alignment [18]. Moreover, the decoding of stimulus identity of unseen test trials (20 per condition) revealed that hierarchical inference is able to correctly identify the true stimulus out of the 50 possible classes (Fig.2E), where we measure decoding accuracy as the fraction of test trials for which the MAP estimate of the stimulus class recovers the ground truth stimulus identity.

Qualitatively similar results can be obtained for a wide range of model choices. Performance is consistent across a wide range of noise levels in the observations (Fig.2F, top), which account for differences in recording quality. The same robustness can be measured for noise in the latent dynamics, due to potential changes in encoding owed to masking odors, variability in sniffing, etc ($\mathbf{Q} \sim \mathcal{N}(\mu, 0.0025)$, $\mu \in [0.1, 0.9]$; Fig.2F, bottom). The estimation is also relatively robust to different degrees of variability of the projections $\mathbf{C}_m$ across animals (Fig.2G, left), which may arise for example due to variability in the alignment of the electrodes on the OB surface. We also varied the dimensionality of the latent dynamics (Fig.2G, middle) and the number of stimulus classes (Fig.2G, right) and found that our decoding procedure shows robust performance for a wide range of settings. Given the relatively modest data requirements (50 trials per subject, and condition), these observations suggest that our approach is robust enough to be applicable to experimentally relevant data regimes.

To better quantify decoding performance as a function of the amount of data that needs to be collected in a new animal, we fixed the amount of data collected from 4 'source' animals ($K = 50$, $I_{m,k} = 50$) and varied the amount of new data collected in a fifth 'target' animal, for which we evaluated decoding accuracy. We found that in this setup amLDS yields high performance with a minimal amount of additional data (Fig.2F, red; one measurement per stimulus condition). Reaching the same decoding performance with data collected in the target animal alone would require many more trials, because the dynamics of all K stimuli need to be learned de novo (Fig.2F, blue). In the case of amLDS however, (effectively) one only needs to align the new measurements to the manifolds estimated based on data from previous animals.[3] In fact, observing a few trials for a single stimulus is enough to achieve reasonable performance in the target animal (Fig.2F, green).

We compared the decoding performance of our method to two state-of-the-art methods: 1) across condition alignment by canonical correlation analysis (CCA) [21, 22] and 2) Procrustes alignment with FA-based dimensionality reduction (FA+Procrustes) [8, 6]. CCA takes pairs of measures from two animals and projects them into a shared low dimensional manifold so as to maximize the correlation between them. Since there is no meaningful way of pairing single trials, CCA requires averaging responses across trials within condition and pairing these average responses by stimulus type. Moreover, unlike our approach, both procedures treat measurements as independent over time, ignoring temporal structure in the neural responses. Since there is no straightforward way to expand the procedure to more than two animals, we increase the number of trials in the single source animal to match the other methods.

The alternative FA+Procrustes alignment is similar to CCA in that it requires across-trial averages and only considers similarity between pairs of measures; however, it can be applied multiple times to align animals $2 - M$ to the prototype axes given by animal 1 (the target animal data was used for FA and Procrustes parameter estimates in this case). More formally, we use orthogonal Procrustes [23, 8] where given a set of prototype $\mathbf{z}^{k*}_{1:T}$ and misaligned $\mathbf{O}\mathbf{z}^{k}_{1:T}$ trajectory pairs finds a rotation between the two, by minimizing the average misalignment after applying a rotation O, $\mathrm{argmin}_{\mathbf{O}} \sum_k \|\mathbf{z}^{k*}_{1:T} - \mathbf{O}\mathbf{z}^{k}_{1:T}\|^2_F$, under the constraint of orthogonality, $\mathbf{O}\mathbf{O}^\top = \mathbf{I}$; where $\| \cdot \|_F$ is the Frobenius norm. The stimulus-specific latent trajectories $\mathbf{z}^{k}_{1:T}$ are obtained via a dimensionality reduction step, performed independently in each animal by factor analysis. As there is no standard way for establishing the shared latent dimensionality based on data in either procedure, we set it to the ground truth d for both CCA and FA+Procrustes. Having established a linear map between observations and a common latent space by either procedure, individual trials across conditions were all projected in this space. We used a linear SVM with L2 regularization to decode stimulus identity from the latent trajectories [12, 24]. All controls were constructed using the library scikit-learn [25]. We found that although

[3] While the estimation procedure still uses EM based on the full data, the fractional contribution of the new measurements to the latent parameter updates is comparatively small, and becomes negligible as the number of sources increases.

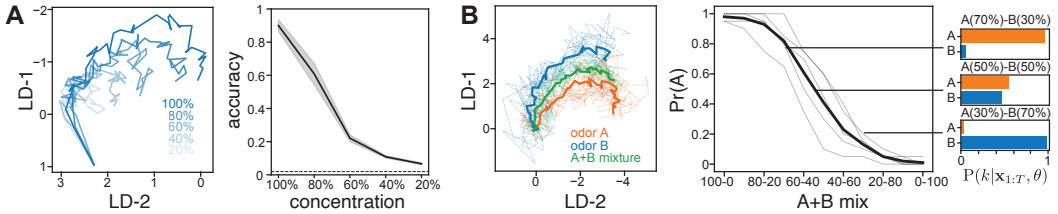

Figure 3: **Mixtures and concentration**. **A.** Estimation of odor identity across concentrations (right); example trajectories where decreasing concentration induces a compression of the stimulus trajectories (left). Shaded area shows s.e.m. across simulations. **B.** Example trajectories for two odors (A-orange and B-blue) and a 50-50% mixture that linearly interpolates between them (left). Frequency of the MAP estimate detecting odor A, as a function of the mixture ratio (middle) and three example posteriors (right); dashed lines show individual simulations, with average in solid line.

both alternatives performed above chance, neither came close to matching the across animal decoding performance of amLDS.

Robustness to fluctuations in concentration, odor mixtures

Odor decoding in the wild needs to overcome further challenges that are not explicitly captured in the amLDS graphical model, such as fluctuations in concentration that rescale the amplitude and speed of odor specific responses [26]. Moreover, odorants are rarely present in isolation and have to be detected when presented in mixtures [27]. We used further numerical simulations to test how the estimator would respond in such scenarios.

We model across-trial variability in odor concentration as changes in the amplitude of the input drive $\mathbf{b}_{t,k}$ [28], with all other parameters the same as before. The parameter learning used data from the highest concentration and the decoder was applied to test trials with lower concentrations (Fig.3A). We found that the decoding accuracy degrades gracefully as the amplitude of responses decreases, remaining above chance for a wide range of concentrations.

It remains unclear how mixtures are encoded in OB responses [29], but a reasonable first approximation is to use a linear interpolation of sample trajectories for the individual odors present [30, 31]. For simplicity, we used $K = 2$ and further tied the parameters $\mathbf{A}$ and $\mathbf{Q}$ of the latent dynamics. We constructed two component mixtures which interpolate between stimuli labeled A and B (Fig.3B, left). Using a model trained on individual odors, we estimated decoding performance on mixture test trials while varying the ratio between the two components (Fig.3B, middle). The posterior over the two odors appropriately reflects the inherent ambiguity in the observations (Fig.3B, right). Overall, these results suggest that the decoding procedure is robust to naturalistic odor variations.

4 Across-animal decoding in the rodent olfactory bulb

We tested our model on neural recordings from a 64-site grid-electrode stereotaxically implanted over the dorsal part of the olfactory bulb in five mice. We simultaneously recorded a pressure readout of the animal's sniff cycle. The animals were presented with five different monomolecular odorants (methyl valerate, ethyl tiglate, benzaldehyde, hexanal and background air), delivered in a temporally precise window (20-100 trials each). We extracted glomerular responses in a time window of 210ms from inhalation onset. The neural responses were preprocessed by removing the high frequency component (100Hz low-pass filter) and electrical noise (60Hz notch filter). Single electrode signals with peaks exceeding 2mV in a 5 second period were considered damaged and removed from the study. The electrode-specific signals were extracted by subtracting the instantaneous mean across all viable electrode sites.

We split the animals into 4 sources, and one target. All data from the source animals and 10% of the trials from the target animal were used for parameter learning and hyperparameter estimation. The remaining data from the target animal was used to assess odor decoding performance. When fitting the model to the neural recordings we found that the inferred latent trajectories were separable across stimuli and similar across animals (Fig. 4A). The model comparison revealed a relatively small

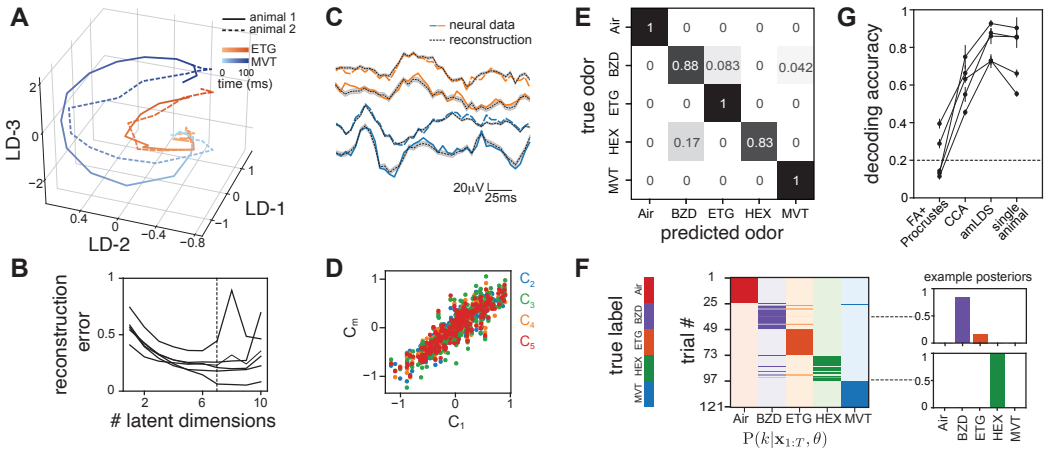

Figure 4: **Decoding performance on neural data from mouse olfactory bulb**. **A.** Inferred trajectories for two monomolecular odors (colors) in two animals (solid and dashed lines), averaged across 24 trials. **B.** Model selection based on reconstruction error performance; dashed lines show individual targets, solid line their average. **C.** Example odor-evoked responses and amLDS reconstructions for 2 stimuli and 2 animals. Shaded areas show 95% confidence intervals based on inferred posterior. **D.** Similarity of observation model across animals. **E.** Confusion matrix for amLDS decoding. **F.** Corresponding single trial posteriors; colors mark different inferred odors, trials are grouped by presented odor. **G.** Decoding performance for amLDS and alternative approaches based on CCA, Procrustes alignment and a probabilistic decoder trained using data from the target animal alone (50% training, 50% test). Lines correspond to different choices for the target animal. Error bars show s.e.m. across 10 folds of the target data, for CCA s.e.m. across sources.

latent dimensionality in the data (Fig. 4B; we used $d = 7$ for all subsequent analyses). The model accurately captured variability in observed electrodes in the target animals, despite the very limited target data (2-6 trials per stimulus; Fig. 4C). The inferred map from the neural measurements into the latent space was similar across animals (Fig. 4D; see also Suppl. Fig. S1), likely reflecting the stereotypy of OB encoding and a precise alignment of the electrodes on the bulb surface.

We evaluated the decoding performance of amLDS on a target animal which revealed high accuracy MAP inference based decoding (Fig. 4E), with generally sharp single trial posteriors (Fig. 4F). We measured decoding accuracy for all possible choices of target animal and compared it to alternative estimators (Fig. 4G). The CCA and FA+Procrustes estimation procedures were similar to those described for the simulated data, with the same latent dimensionality as inferred for amLDS. We used across trial averages within each stimulus condition from the source animals and the 10% fraction of the target animal data for dimensionality reduction and across animal alignment. We used the same data to train a linear SVM classifier based on the latent trajectories, which we then used to predict odor identity for the withheld target data. For CCA [21] we used single source animals and the training target data for the alignment and SVM training, averaging decoding performance on test data across all possible sources. Our probabilistic decoder significantly outperformed both alternative alignment methods. For all possible choices of target, amLDS also systematically outperformed probabilistic decoding based on data from the target animal alone, confirming that pooling together data across animals improves data efficiency and overall decoding performance. While the single animal performance is close to that of amLDS given enough data, $82\% \pm 8\%$ vs. $76\% \pm 3\%$ when training with 50% of the target animal's dataset (Fig. 4G), the difference between the two becomes dramatic in the low data regime used for amLDS (10% of target data used for training), with single animal performance going down to $60\% \pm 8\%$. This reinforces the idea that across-animal decoding is key to limiting the amount of data that needs to be collected for calibrating any new animal at deployment time.

5 Discussion

Animals are unparalleled in their ability to sense odors, which makes them an invaluable resource for a wide range of applications from security to medicine [32–34]. BCI approaches to odor detection attempt to bypass costly behavioral training by reading out odor information directly from the olfactory system, but they remain data limited [13]. Here we proposed a novel probabilistic procedure for mapping neural responses to odors into a shared manifold. This representation allows for robust odor decoding in a new animal with minimal amounts of further data collection. Calibrating the decoder to a new animal can be achieved with a small subset of stimuli, limiting the demand of rare/expensive chemicals. For instance, in our numerical experiments the decoder could extract odor identity in a new animal with as few as 2 trials for each stimulus, or 50 trials of a single stimulus. Moreover, we also found the estimation procedure to be robust to modeled fluctuations in odor concentration and the presence of background odors. These benefits were confirmed in data from mouse OB, where we found that decoding quality improved dramatically compared to previously-proposed alignment procedures, particularly when using a limited amount of data.

Our across-animals probabilistic decoding procedure brings neural chemical detection closer to practical applicability. Not only does it improve scalability and robustness, but it also significantly reduces data collection requirements. Furthermore, most odor detection applications have asymmetric costs and need to negotiate complex trade-offs between sensitivity and accuracy. For instance, when detecting *Clostridium difficile* [35], the medical costs of missing a positive diagnosis needs to be balanced against the financial costs of unnecessarily quarantining someone due to a false alarm. Our probabilistic decoder provides an explicit report of uncertainty associated with possible outcomes, which is critical for optimal decision making in these contexts.

Alternative procedures for neural datasets alignment are usually not explicitly probabilistic.[4] Instead, they rely on a sequence of independent steps: first a (usually off-the-shelf) dimensionality reduction procedure. Dimensionality reduction is followed by across-condition data alignment within the latent manifold, using approaches such as CCA (and its multivariate generalizations [37, 38]), or Procrustes alignment [6, 39, 8]. Lastly, a common decoder is trained on the aligned across-condition data [21, 9, 8]. These procedures are popular as a replacement for costly and inconvenient BCI re-calibration in prosthetics [40], but their application to across-animal alignment remains rare [7, 6]. The multi-step procedure, although formally suboptimal, allows for some degree of flexibility in the individual component steps. For instance, the map into the latent manifold could be nonlinear, e.g. using Isomap [7, 41]. There also exist several nonlinear alignment procedures, such as kernel CCA or Distance Covariance Analysis , although such extensions have yet to translate into BCI practice [42]. In contrast, the hierarchical model proposed here is jointly estimated and linear in both the map and the latent dynamics. The linearity assumption here is not just a simplifying mathematical assumption, but it reflects domain-specific knowledge about the underlying neurophysiology of the system.

Neural activity has complex temporal dynamics and structured across-trial variability [43]. Previous approaches to aligning population activity across conditions via CCA or Procrustes alignment fail to capture either of these key statistical regularities. First, they treat the measurements across time as independent. Second, they intrinsically rely on paired measures, and since the correspondence of individual moments in time across trials in different animals is essentially arbitrary, they are restricted to across-trial average responses that may obscure important features of the underlying computation [44, 45]. It is also less clear how to determine the latent dimensionality of the data using these procedures. In contrast, amLDS extracts common low-dimensional dynamics using single trials, exploiting known features of the underlying computation. Moreover, the animal-specific observation model provides a principled framework for determining data dimensionality and for handling different sources of experimental variability (e.g. due to grid misalignment or missing data from defective electrodes). The probabilistic formulation also lends itself to a variety of generalizations. One could easily convert from Gaussian to Poisson observation noise for modeling spiking data instead of the raw electric signals in our data [46, 47], or replace the linear dynamics with shared nonlinear equivalents [48, 49]. Hence, amLDS can be thought of as yet another building block in the expanding statistical toolkit for understanding neural computation through the lens of dynamics in low-dimensional manifolds.

Acknowledgements

[4] Some exceptions exist, such as probabilistic CCA [36], but they are not commonly used in practice.

References

[1] Ofer Mazor and Gilles Laurent. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*, 48(4):661–673, 2005.

[2] Christopher J. Cueva, Alex Saez, Encarni Marcos, Aldo Genovesio, Mehrdad Jazayeri, Ranulfo Romo, C. Daniel Salzman, Michael N. Shadlen, and Stefano Fusi. Low-dimensional dynamics for working memory and time encoding. *Proceedings of the National Academy of Sciences*, 117(37):23021–23032, 2020.

[3] Valerio Mante, David Sussillo, Krishna V Shenoy, and William T Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, 2013.

[4] Mark M. Churchland, John P. Cunningham, Matthew T. Kaufman, Stephen I. Ryu, and Krishna V. Shenoy. Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron*, 68(3):387–400, 2010.

[5] Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520, 2019.

[6] Hung-Tu Chen, Jeremy R. Manning, and Matthijs A. A. van der Meer. Between-subject prediction reveals a shared representational geometry in the rodent hippocampus. *bioRxiv*, 2020.

[7] Alon Rubin, Liron Sheintuch, Noa Brande-Eilat, Or Pinchasof, Yoav Rechavi, Nitzan Geva, and Yaniv Ziv. Revealing neural correlates of behavior without behavioral measurements. *Nature Communications*, 10(1):4745, 2019.

[8] Alan D. Degenhart, William E. Bishop, Emily R. Oby, Elizabeth C. Tyler-Kabara, Steven M. Chase, Aaron P. Batista, and Byron M. Yu. Stabilization of a brain–computer interface via the alignment of low-dimensional spaces of neural activity. *Nature Biomedical Engineering*, 4(7):672–685, 2020.

[9] He He and Dongrui Wu. Transfer learning for brain–computer interfaces: A euclidean space data alignment approach. *IEEE Transactions on Biomedical Engineering*, 67(2):399–410, 2020.

[10] Qi Dong, Liping Du, Liujing Zhuang, Rong Li, Qingjun Liu, and Ping Wang. A novel bioelectronic nose based on brain-machine interface using implanted electrode recording in vivo in olfactory bulb. *Biosensors and Bioelectronics*, 49:263—269, 2013.

[11] Debajit Saha, Darshit Mehta, Ege Altan, Rishabh Chandak, Mike Traner, Ray Lo, Prashant Gupta, Srikanth Singamaneni, Shantanu Chakrabartty, and Baranidharan Raman. Explosive sensing with insect-based biorobots. *Biosensors and Bioelectronics: X*, 6:100050, 2020.

[12] Erez Shor, Pedro Herrero-Vidal, Adam Dewan, Ilke Uguz, Vincenzo F. Curto, George G. Malliaras, Cristina Savin, Thomas Bozza, and Dmitry Rinberg. A mouse bio-electronic nose for sensitive and versatile chemical detection. *bioRxiv*, 2020.

[13] Vinay Jayaram, Morteza Alamgir, Yasemin Altun, Bernhard Scholkopf, and Moritz Grosse-Wentrup. Transfer learning in brain-computer interfaces. *IEEE Computational Intelligence Magazine*, 11(1):20–31, 2016.

[14] Rachel I. Wilson and Zachary F. Mainen. Early events in olfactory processing. *Annual Review of Neuroscience*, 29(1):163–201, 2006. PMID: 16776583.

[15] Markus Meister and Tobias Bonhoeffer. Tuning and topography in an odor map on the rat olfactory bulb. *Journal of Neuroscience*, 21(4):1351–1360, 2001.

[16] Edward R. Soucy, Dinu F. Albeanu, Antoniu L. Fantana, Venkatesh N. Murthy, and Markus Meister. Precision and diversity in an odor map on the olfactory bulb. *Nature Neuroscience*, 12(2):210–220, 2009.

[17] Brice Bathellier, Derek L. Buhl, Riccardo Accolla, and Alan Carleton. Dynamic ensemble odor coding in the mammalian olfactory bulb: Sensory information at different timescales. *Neuron*, 57(4):586–598, 2008.

[18] Sam Roweis and Zoubin Ghahramani. A unifying review of linear gaussian models. *Neural Computation*, 11(2):305–345, 1999.

[19] Byron M. Yu, John P. Cunningham, Gopal Santhanam, Stephen Ryu, Krishna V. Shenoy, and Maneesh Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis

of neural population activity. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2009.

[20] John P. Cunningham and Byron M. Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, 2014.

[21] Juan A. Gallego, Matthew G. Perich, Raeed H. Chowdhury, Sara A. Solla, and Lee E. Miller. Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience*, 23(2):260–270, 2020.

[22] T L Veuthey, K Derosier, S Kondapavulur, and K Ganguly. Single-trial cross-area neural population dynamics during long-term skill learning. *Nature Communications*, 11(1):4057, 2020.

[23] Peter H. Schönemann. A generalized solution of the orthogonal Procrustes problem. *Psychometrika*, 31(1):1–10, 1966.

[24] Kevin A. Bolding and Kevin M. Franks. Recurrent cortical circuits implement concentration-invariant odor coding. *Science*, 361(6407), 2018.

[25] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[26] Jeffrey A. Riffell. Processing a dynamic odor in a noisy chemical environment. *Flavour*, 3(1):O28, 2014.

[27] Jordi Fonollosa, Irene Rodríguez-Luján, Marco Trincavelli, Alexander Vergara, and Ramón Huerta. Chemical discrimination in turbulent gas mixtures with mox sensors validated by gas chromatography-mass spectrometry. *Sensors*, 14(10):19336–19353, 2014.

[28] Mark Stopfer, Vivek Jayaraman, and Gilles Laurent. Intensity versus identity coding in an olfactory system. *Neuron*, 39(6):991–1004, 2003.

[29] Thierry Thomas-Danguin, Charlotte Sinding, Sébastien Romagny, Fouzia El Mountassir, Boriana Atanasova, Elodie Le Berre, Anne-Marie Le Bon, and Gérard Coureaud. The perception of odor objects in everyday life: a review on the processing of odor mixtures. *Frontiers in Psychology*, 5:504, 2014.

[30] Max L. Fletcher. Analytical processing of binary mixture information by olfactory bulb glomeruli. *PLOS ONE*, 6:1–6, 2011.

[31] Kai Shen, Sina Tootoonian, and Gilles Laurent. Encoding of mixtures in a simple olfactory system. *Neuron*, 80(5):1246–1262, 2013.

[32] Matteo E. Bonfanti. From Sniffer Dogs to Emerging Sniffer Devices for Airport Security: An Opportunity to Rethink Privacy Implications? *Science and Engineering Ethics*, 20(3):791–807, 2014.

[33] In-Seok Seo, Hwan-Gon Lee, Bonkon Koo, Chin Su Koh, Hae-Yong Park, Changkyun Im, and Hyung-Cheul Shin. Cross detection for odor of metabolic waste between breast and colorectal cancer using canine olfaction. *PLOS ONE*, 13(2):1–9, 2018.

[34] Georgies F. Mgode, Christophe L. Cox, Stephen Mwimanzi, and Christiaan Mulder. Pediatric tuberculosis detection using trained African giant pouched rats. *Pediatric Research*, 84(1):99–103, 2018.

[35] Marije K. Bomers, Michiel A. van Agtmael, Hotsche Luik, Merk C. van Veen, Christina M. J. E. Vandenbroucke-Grauls, and Yvo M. Smulders. Using a dog's superior olfactory sensitivity to identify clostridium difficile in stools and patients: proof of principle study. *BMJ*, 345, 2012.

[36] Francis R. Bach and Michael I. Jordan. *A probabilistic interpretation of canonical correlation*. Technical Report 688, Department of Statistics, UC Berkeley, 2005.

[37] Chong Wang. Variational bayesian approach to canonical correlation analysis. *IEEE Transactions on Neural Networks*, 18(3):905–910, 2007.

[38] Xiaowei Zhuang, Zhengshi Yang, Karthik R. Sreenivasan, Virendra R. Mishra, Tim Curran, Rajesh Nandy, and Dietmar Cordes. Multivariate group-level analysis for task fmri data with canonical correlation analysis. *NeuroImage*, 194:25–41, 2019.

[39] Pedro Rodrigues, Christian Jutten, and Marco Congedo. Riemannian procrustes analysis: Transfer learning for brain-computer interfaces. *IEEE Transactions on Biomedical Engineering*, PP:1–1, 12 2018.

[40] Chethan Pandarinath, K. Cora Ames, Abigail A. Russo, Ali Farshchian, Lee E. Miller, Eva L. Dyer, and Jonathan C. Kao. Latent factors and dynamics in motor cortex and their application to brain-machine interfaces. *Journal of Neuroscience*, 38(44):9390–9401, 2018.

[41] Eva L. Dyer, Mohammad Gheshlaghi, Azar, Matthew G. Perich, Hugo L. Fernandes, Stephanie Naufel, Lee E. Miller, and Konrad P. Körding. A cryptography-based approach for movement decoding. *Nature Biomedical Engineering*, 1(12):967–976, 2017.

[42] Max Dabagia, Konrad P. Kording, and Eva L. Dyer. Comparing high-dimensional neural recordings by aligning their low-dimensional latent representations. *Nature Biomedical Engineering (to appear)*, 2020.

[43] Jonathan C. Kao, Paul Nuyujukian, Stephen I. Ryu, Mark M. Churchland, John P. Cunningham, and Krishna V. Shenoy. Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nature Communications*, 6(1):7759, 2015.

[44] Matthew R. Whiteway and Daniel A. Butts. The quest for interpretable models of neural population activity. *Current Opinion in Neurobiology*, 58:86–93, 2019. Computational Neuroscience.

[45] Alex H. Williams and Scott W. Linderman. Statistical neuroscience in the single trial limit. *arXiv*, 2021.

[46] Jakob H. Macke, Lars Buesing, and Maneesh Sahani. *Estimating state and parameters in state space models of spike trains*, page 137–159. Cambridge University Press, 2015.

[47] Yuanjun Gao, Lars Busing, Krishna V. Shenoy, and John P. Cunningham. High-dimensional neural spike train analysis with generalized count linear dynamical systems. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.

[48] Yuan Zhao and Il Memming Park. Interpretable nonlinear dynamic modeling of neural trajectories. *arXiv*, 2016.

[49] Lea Duncker, Gergo Bohner, Julien Boussard, and Maneesh Sahani. Learning interpretable continuous-time models of latent stochastic dynamical systems. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1726–1734. PMLR, 09–15 Jun 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Yes, it's accurate
   (b) Did you describe the limitations of your work? [Yes] See Section 2 and Discussion
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] Yes, it conforms
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 2
   (b) Did you include complete proofs of all theoretical results? [Yes] See Section 2 and Supplementary 1
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] A full description of the model and simulations is provided on Sections 2-4. Supplementary code is provided to replicate all the results. A complete description of the neural data is provided on Section 4 and the data will be made available upon reasonable request.

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Sections 2-4

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See figure captions and Sections 3 and 4

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 3. amLDS memory and computational demands are really low, both parameter learning and inference can run on a laptop computer.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] We use code from scikit-learn, which we credit in reference [25].

(b) Did you mention the license of the assets? [N/A]

(c) Did you include any new assets either in the supplemental material or as a URL? [Yes] The model and a user tutorial is provided as supplementary material.

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]