
A general recurrent state space framework for modeling neural dynamics during decision-making

David M. Zoltowski¹ Jonathan W. Pillow^{1,2} Scott W. Linderman^{3,4}

Abstract

An open question in systems and computational neuroscience is how neural circuits accumulate evidence towards a decision. Fitting models of decision-making theory to neural activity helps answer this question, but current approaches limit the number of these models that we can fit to neural data. Here we propose a general framework for modeling neural activity during decision-making. The framework includes the canonical drift-diffusion model and enables extensions such as multi-dimensional accumulators, variable and collapsing boundaries, and discrete jumps. Our framework is based on constraining the parameters of recurrent state space models, for which we introduce a scalable variational Laplace EM inference algorithm. We applied the modeling approach to spiking responses recorded from monkey parietal cortex during two decision-making tasks. We found that a two-dimensional accumulator better captured the responses of a set of parietal neurons than a single accumulator model, and we identified a variable lower boundary in the responses of a parietal neuron during a random dot motion task. We expect this framework will be useful for modeling neural dynamics in a variety of decision-making settings.

1. Introduction

Evidence accumulation is central to many models of perceptual decision-making (Gold & Shadlen, 2007; Ratcliff & McKoon, 2008; Ratcliff et al., 2016). However, despite progress in identifying neural correlates of decision-making

it remains an open question how evidence accumulation is implemented in the brain (Brody & Hanks, 2016). One approach to address this question is to formulate models of decision-making behavior as generative models of neural activity (Ditterich, 2006; Bollimunta et al., 2012; Latimer et al., 2015; Zoltowski et al., 2019; DePasquale et al., 2019). Fitting these models to single-trial neural responses during decision-making tasks provides a direct test of how well the theorized model explains neural dynamics.

However, there are a number of decision-making models and features that are challenging to fit to neural activity using existing approaches. Most models of decision-making behavior rely on either 1) approximations of analytic solutions of the joint distribution of a binary choice and decision time (Wiecki et al., 2013, e.g.); or 2) numerical solutions to stochastic differential equations with boundary constraints (Brunton et al., 2013; Shinn et al., 2020, e.g.). Analytic solutions are not available for models of neural activity where there are observations throughout the decision-period, and numerical solutions are limited to low-dimensional accumulator models. Thus, there is a clear need for tractable methods for modeling neural dynamics during decision-making that can accommodate multiple choice options and/or inputs (Churchland et al., 2008; Brunton et al., 2013) and dynamics governed by multi-dimensional accumulators (Scott et al., 2015).

Here we propose a general framework for fitting decision-making models to neural activity.¹ Our key observation is that many decision-making models can be formulated as recurrent switching linear dynamical systems (rSLDS) models (Linderman et al., 2017) with appropriate constraints. This allows us to instantiate and fit a number of models of interest including the classic drift-diffusion model (Ratcliff & McKoon, 2008), one- and multi-dimensional accumulator models (Gold & Shadlen, 2007; Brunton et al., 2013), and the ramping and stepping models from (Latimer et al., 2015). The framework naturally includes collapsing boundaries and trial history effects (O’Connell et al., 2018; Lueckmann et al., 2018; Urai et al., 2019). It also enables new extensions to models with variable, probabilistic boundaries and with non-

¹Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA ²Department of Psychology, Princeton University, Princeton, NJ, USA ³Department of Statistics, Stanford University, Palo Alto, CA, USA ⁴Wu Tsai Neurosciences Institute, Stanford University, Palo Alto, CA, USA. Correspondence to: David M. Zoltowski <zoltowski@princeton.edu>.

¹Code is available at <https://github.com/davidzoltowski/ssmdm>.

constant boundary dynamics (Resulaj et al., 2009; Evans et al., 2018), which are not naturally described by traditional methods that assume constant, absorbing boundaries.

We also develop a variational Laplace EM algorithm for inference in rSLDS models. This method combines variational and Laplace approximations over the discrete and continuous latent variables, respectively. In synthetic data, we find that the algorithm is faster than particle EM and black box variational inference methods and provides comparable or greater accuracy. We show how it allows us to fit new decision-making models to Poisson spike train data, where we find that multi-dimensional accumulator models outperform existing methods for modeling neural responses during decision-making.

2. Background

We review decision-making models for behavioral and neural data and the recurrent switching linear dynamical systems we will use to represent, fit, and generalize them.

2.1. Perceptual decision-making

Perceptual decision-making is the process of using sensory information to make a categorical choice (Hanks & Summerfield, 2017). A classic example is determining the net motion direction of a cloud of randomly moving dots (Figure 1A). The accumulation of evidence over time in favor of a decision until a threshold or boundary is crossed (accumulation-to-bound) is a key process in models of such perceptual decisions (Gold & Shadlen, 2007; Ratcliff & McKoon, 2008).

A simple example model of perceptual decisions is the drift-diffusion model (DDM). The DDM is a 1-dimensional dynamics model of a binary choice (Figure 1A). It states that a scalar variable x evolves in time according to a biased random walk. In discrete time this is formalized as

$$x_t = x_{t-1} + \beta_c + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2). \quad (1)$$

where β_c , the drift term, corresponds to the strength of sensory evidence and noise ϵ_t corresponds to noise in the sensory input. Once x crosses an upper or lower threshold (boundary) the decision-process stops and the variable x is fixed to one of the two boundary values. The choice produced by the model depends on which of the upper or lower thresholds is hit; each boundary represents one of the two choices. The initial value may be set to a constant x_0 or drawn from an initial distribution $x_0 \sim \mathcal{N}(\mu_0, \sigma_0^2)$. Many variants of drift-diffusion models exist; see (Ratcliff et al., 2016) for a recent review.

2.2. Relating neural activity to decisions

Neural correlates of decision-making have been identified in numerous brain regions (Gold & Shadlen, 2007; Brody & Hanks, 2016; Hanks & Summerfield, 2017). While the neural correlates of decisions have traditionally been observed when averaging neural activity across many decisions, a prominent line of work has attempted to relate single-trial neural responses to decision-making dynamics (Ditterich, 2006; Churchland et al., 2011; Bollimunta et al., 2012; Latimer et al., 2015; Hanks et al., 2015; Zoltowski et al., 2019). A prominent example is the ramping model from (Latimer et al., 2015) where a modified DDM without a lower boundary was formulated as a statistical model of spike trains recorded in the monkey parietal cortex, following theoretical models (Mazurek et al., 2003).

2.3. Recurrent switching state space models

An rSLDS is an extension of a linear dynamical system model that approximates nonlinear dynamics with a discrete set of linear regimes (Linderman et al., 2017; Nassar et al., 2019). The generative process is as follows. At each time point, the dynamics are given by

$$x_t = A_{z_t} x_{t-1} + V_{z_t} u_t + b_{z_t} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, Q_{z_t}) \quad (2)$$

where x_t is the continuous state, z_t is one of K discrete states, and u_t is the input at time t . There are separate dynamics parameters A_k, V_k, b_k for each discrete state $k \in \{1, 2, \dots, K\}$. The key feature that makes the models “recurrent” is that transitions between discrete states depend on both the previous continuous and discrete latent variables. We parameterize the transition probabilities using a multi-class logistic regression

$$p(z_t | z_{t-1}, x_{t-1}) \propto \exp \{ \gamma (R_{z_t-1} + r x_{t-1}) \} \quad (3)$$

with parameters R_k for each discrete state and a vector $r \in \mathbb{R}^d$ describing the dependence on the previous continuous latent state. The hyperparameter γ sets the sharpness of the decision boundaries. As $\gamma \rightarrow 0$, the transitions become uniform, and as $\gamma \rightarrow \infty$ they become deterministic.

In our applications, we use a Poisson observation distribution

$$y_t \sim \text{Poisson}(f(Cx_t + d) \Delta_t) \quad (4)$$

where $f(x) = \log(1 + \exp(x))$ is the softplus function and Δ_t is the time bin size. However, it is possible to use other observation distributions for alternative types of data. The nonlinearity could also be changed to model accumulator tuning curves (Hanks et al., 2015; DePasquale et al., 2019) or saturating or accelerating relationships between the diffusion process and firing rates (Howard et al., 2018; Zoltowski et al., 2019). The graphical model of an rSLDS is shown in Figure 1C.

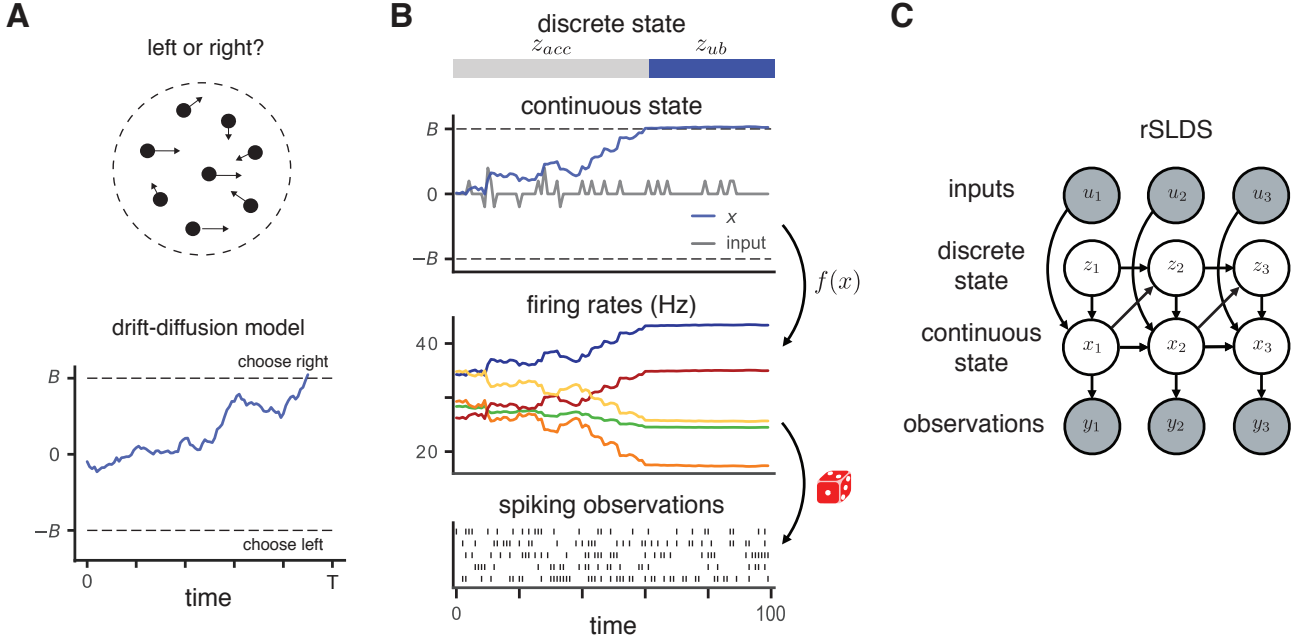


Figure 1. Constraining rSLDS to follow accumulation-to-bound dynamics. **A.** An example perceptual decision-making task is determining the direction of a cloud of randomly moving dots. The expected value of the motion direction is either to the left or right. The drift-diffusion model (DDM) is commonly used to explain decisions in these types of tasks. **B.** An rSLDS with constraints on the parameters can implement accumulation-to-bound. A simulated trial from this model is shown. The continuous latent variable accumulates input in the accumulation state before switching to the upper boundary state after it crosses the upper threshold. Firing rates are generated from an affine mapping of the continuous variable and Poisson spike counts are generated. **C.** The rSLDS is a switching linear dynamical system with additional dependencies between x_t and z_{t+1} , as shown in this graphical model.

3. Decision-making dynamics as constrained rSLDS

Our key observation is that the dynamics of accumulation-to-bound models can be instantiated as constrained recurrent switching linear dynamical systems. We first illustrate this by writing a one-dimensional accumulation-to-bound as a constrained rSLDS. We then describe how this approach generalizes to multi-dimensional accumulators, ramping and stepping models, variable boundaries, and other features.

3.1. Illustrative example: one-dimensional accumulator model

Consider a one-dimensional accumulation-to-bound model with upper and lower decision boundaries. This model has a continuous latent variable $x_t \in \mathbb{R}$ that starts near zero and accumulates sensory input until it reaches one of two decision boundaries at $\pm B$. The model has $K = 3$ discrete states: an accumulation state ($z_t = acc$), an upper boundary state ($z_t = ub$), and a lower boundary state ($z_t = lb$).

In the accumulation state, the sensory input u_t at each time point is accumulated via the following dynamics

$$x_t = x_{t-1} + V_{acc}u_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_{acc}^2) \quad (5)$$

where V_{acc} is a weight on the sensory input. In the boundary states there is no dependence on the input ($V_{ub} = V_{lb} = 0$) such that

$$x_t = x_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_{ub/lb}^2) \quad (6)$$

and the diffusion variances σ_{ub}^2 and σ_{lb}^2 are set to small values relative to the scale of the dynamics. In all three states, $A = 1$ and $b = 0$.

To implement the decision boundaries, we parameterize the transitions such that the discrete state switches from the accumulation to the boundary states once x_{t-1} has crossed one of the boundaries. The transition parameters in equation 3 depend on a boundary location B parameter such that

$$R_{acc} = \begin{bmatrix} 0 \\ -B \\ -B \end{bmatrix}, \quad R_{ub} = \begin{bmatrix} -\infty \\ 0 \\ -\infty \end{bmatrix}, \quad R_{lb} = \begin{bmatrix} -\infty \\ -\infty \\ 0 \end{bmatrix}, \quad r = \begin{bmatrix} 0 \\ +1 \\ -1 \end{bmatrix}.$$

Here we set $B = 1$ and $\gamma = 500$ to enforce sharp transition boundaries at ± 1 . Figure 1B shows a simulation from this model with Poisson spike count observations, as described in equation 4.

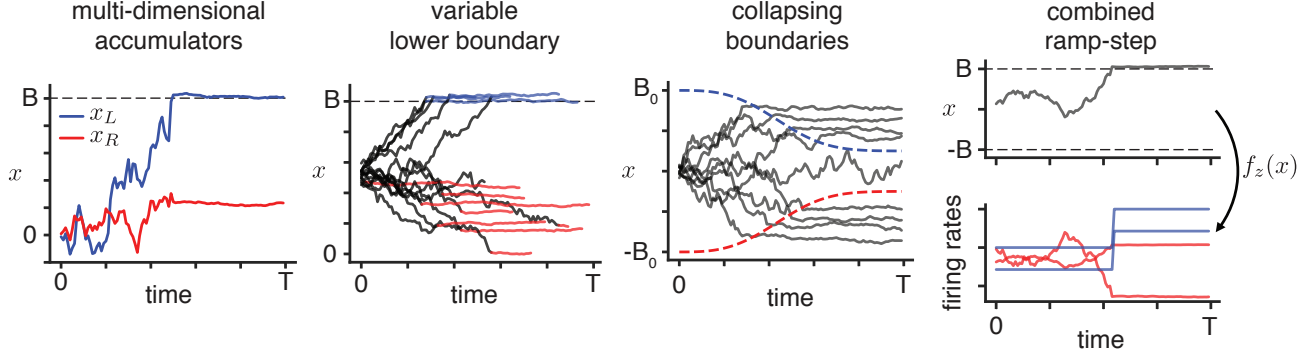


Figure 2. Examples of models in this framework. Constrained rSLDS models include multi-dimensional accumulators such as a 2D race accumulator model, accumulators with variable boundaries and collapsing boundaries, and a model with both discrete steps and ramps as special cases.

3.1.1. SOFT DECISION BOUNDARIES

We can implement soft decision boundaries where switches from the boundary state back to the accumulation state are allowed if x_t falls below the boundary. Here the transition probabilities depend only on the previous value of the continuous latent variable with

$$R_{\text{acc}} = R_{\text{ub}} = R_{\text{lb}} = \begin{bmatrix} 0 \\ -B \\ -B \end{bmatrix}, \quad r = \begin{bmatrix} 0 \\ +1 \\ -1 \end{bmatrix}. \quad (7)$$

3.2. Multi-dimensional accumulator models

Models with multiple accumulating dimensions are natural for settings with multiple input streams and choice options. Additional dimensions may also be useful for incorporating other task features such as context (Shvartsman et al., 2015). In this framework, it is straightforward to generalize the one-dimensional accumulator model to multi-dimensional accumulators by adding dimensions to the continuous and discrete states. For a D -dimensional accumulator, the continuous latent variable $x \in \mathbb{R}^D$ and the input $u \in \mathbb{R}^D$ are vectors. We set the dynamics, input, and covariance matrices $\{A_{\text{acc}}, V_{\text{acc}}, Q_{\text{acc}}\}$ to be diagonal so that each dimension independently accumulates one stream of input. However, we can relax this assumption to have interactions in the latent space (non-diagonal A_{acc}) or correlated noise (non-diagonal Q_{acc}).

We set the discrete state transitions to follow “race” accumulator dynamics, where the different accumulator dimensions race to reach an upper boundary. The discrete state switches when one of the dimensions of x crosses the boundary. As in the 1D model, the boundary states have zero dependence on the input and small dynamics variance. In this setup, there is one accumulation state and D boundary states such that $K = D + 1$. A set of simulated latent trajectories from this model is shown in Figure 2.

3.3. Ramping and stepping models

The ramping model is a one-dimensional accumulator model with an upper boundary (Latimer et al., 2015). It has a constant drift for each stimulus category such that u_t is a one-hot vector denoting the stimulus on each trial and V_{acc} is a vector with a drift for each category. The model has no lower boundary and the initial continuous state has mean $x_0 \in (-\infty, 1)$. This model can be written as a constrained recurrent state space model with $K = 2$ discrete states. Alternatively, the stepping model from Latimer et al. (2015) posits that single-trial firing rates start in an initial state and may either step up or down at some point during the trial.

The recurrent state space framework admits a generalization of the ramping and stepping models that can be either a ramp or a step depending on free parameters. In this model, the underlying latent follows 1D accumulation-to-bound dynamics with upper and lower boundaries and the emission mean parameter depends on the discrete state

$$y_t \sim \text{Poisson}(f(Cx_t + d_{z_t}) \Delta t). \quad (8)$$

The model is a ramp if the mean $d_{z_t}^{(n)}$ for a neuron n is the same for each discrete state and $C^{(n)}$ is non-zero. If $C^{(n)} = 0$ and $d_{z_t}^{(n)}$ is different for each state, then the model steps to different firing rates when the discrete state switches. An illustration of an underlying ramp generating firing rates for two ramping and two stepping neurons is shown in Figure 2 (far right).

3.4. Variable and learned boundaries

We may desire to relax the assumption of sharp boundaries governing the discrete transitions. As the scale parameter γ is increased, the distribution over transitions flattens and the transitions may occur at values below or above the boundary (Figure 2). It is also possible to learn the parameters governing the transitions. In our experiments, we present

Table 1. Decision-making models and features that fit in the framework.

feature	component	details
leaky or unstable dynamics	continuous dynamics	learn A_{acc}
input-dependent noise	continuous dynamics	$\sigma_{\text{acc}}^2 = \sigma_a^2 + u_t^2 \sigma_s^2$
history bias: drift	continuous dynamics	augment input $u_t = [u_t, c_{\text{prev}}]$
relaxed boundary dynamics	continuous dynamics	learn $A_{\text{lb}}, A_{\text{ub}}$ and/or $\sigma_{\text{lb}}^2, \sigma_{\text{ub}}^2$
variable drift	continuous dynamics	hierarchical model of V_{acc} across trials
variable boundaries	discrete transitions	decrease γ and/or learn boundary params
collapsing boundaries	discrete transitions	switch to ub or lb when $ x_{t-1} > B - f(t)$
urgency signal	discrete transitions and emissions	add urgency signal g_t to transitions and emissions
history bias: start	initial state	$x_0 \sim \mathcal{N}(\mu_{z_0} + V_c c_{\text{prev}}, \sigma_{z_0}^2)$
multi-dimensional	all	increase dimensions
non-decision time	all	start in additional state $z_0 = \text{nd}$ before acc

a model with a sharp upper boundary and a learned lower boundary.

3.5. Other features

The modeling framework includes numerous other features in decision-making models, such as sensory-dependent vs. accumulation noise (Brunton et al., 2013), collapsing boundaries and non-decision time (Ratcliff et al., 2016), and trial-history effects that for example depend on the previous choice c_{prev} (Urai et al., 2019). We list them in Table 1 and describe input-dependent noise, collapsing boundaries, and relaxed boundary dynamics below. While we do not fit models with all of the features in this paper, we present them for completeness and we fit models with collapsing boundaries and trial-history effects to simulated data in Appendix A.

Input-dependent dynamics variance In decisions with non-constant input, dynamics variability can be separated into input-driven and dynamics components (Brunton et al., 2013). This corresponds to adding an input-dependent term to the accumulation state variance

$$x_t = A_{\text{acc}} x_{t-1} + V_{\text{acc}} u_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_a^2 + u_t^2 \sigma_s^2). \quad (9)$$

Collapsing boundaries Collapsing decision boundaries can account for decision accuracy as a function of decision-time (Ratcliff et al., 2016; O’Connell et al., 2018). Linearly collapsing boundaries (Bowman et al., 2012) can be implemented by adding a term to the transition probabilities that depends on time t such that

$$p(z_t | z_{t-1}, x_{t-1}) \propto \exp \{ \gamma (R_{z_{t-1}} + r x_{t-1} + W t) \} \quad (10)$$

where $W = [0, \beta, \beta]^\top$ (Figure 2). The parameter β determines the slope of the linear boundary decrease. This formulation can be directly extended for nonlinear collapsing boundaries (Drugowitsch et al., 2012; Hawkins et al.,

2015). See Appendix A for more details and a simulated experiment.

Relaxed boundary dynamics The models introduced above make strong assumptions that the dynamics in the boundary state are constant with nearly zero noise. We can relax these assumptions by learning the dynamics parameters of the boundary states, which could correspond to working memory or changes of mind (Resulaj et al., 2009) and have been used to describe mouse escape behavior (Evans et al., 2018).

3.6. Differences with decision-making models

Our recurrent state space formulation of accumulation to a boundary departs from some formulations of decision-making models in two ways. First, the continuous state x_t can cross the boundary and fluctuate, rather than reaching a constant level after hitting the boundary. We consider this a feature because it enables modeling neural dynamics after the boundary is reached. However, using a small boundary variance and identity dynamics matrix mitigates this difference, as described above. Additionally, to enforce the continuous state to remain at the boundary we can threshold x_t when passing it to the firing rate. The second difference is that the transitions are probabilistic. We also consider this a feature because it allows for generalizations such as variable boundary locations. However, we can set the transitions to be effectively deterministic using sharp transition boundaries, as described previously.

4. Inference

Inference in rSLDS models via direct maximum likelihood scales exponentially in the length of the time-series and is therefore computationally intractable (Barber, 2012). A number of alternative methods have been developed for inference in rSLDS models including expectation-

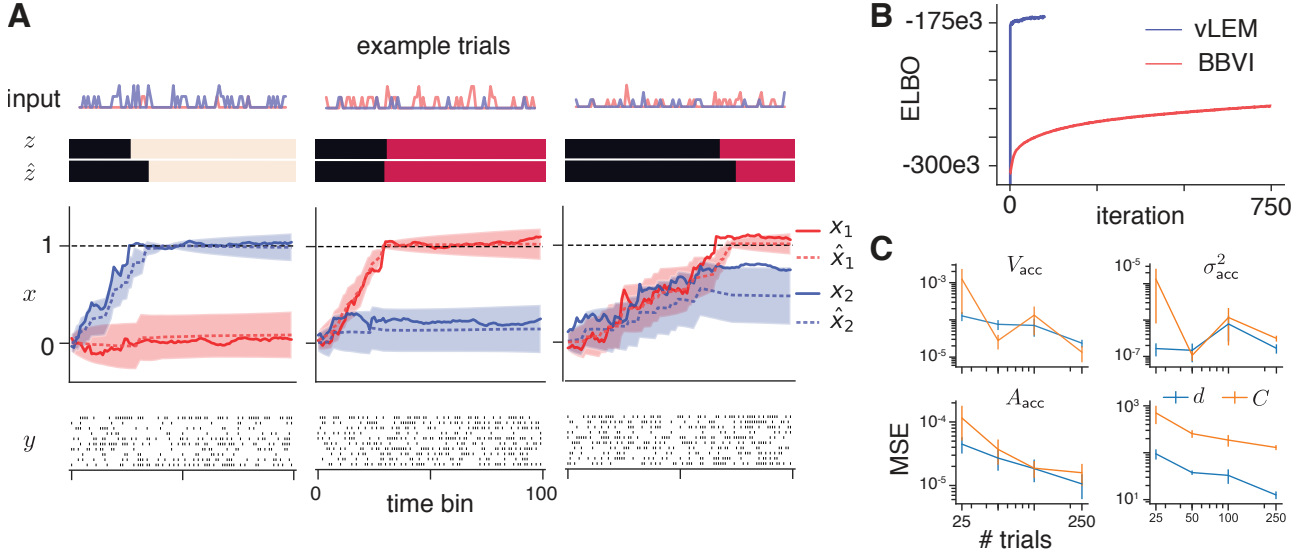


Figure 3. Simulated 2D race accumulator model with Poisson observations. **A.** The true (z, x) and inferred (\hat{z}, \hat{x}) discrete and continuous latent variables are similar. The colored bars (*top*) indicate the true and inferred most likely discrete states at each time point. Shaded areas indicate two standard deviations of the posterior $q(x)$. **B.** The ELBO as a function of inference iteration for vLEM and BBVI. **C.** Average mean squared error \pm standard error between the true and inferred parameters across 10 simulated experiments for each number of trials. The different colored lines on the V_{acc} , σ_{acc}^2 , and A_{acc} panels denote the parameters for the two different dimensions.

propagation (Barber, 2006) and augmented Gibbs sampling algorithms (Linderman et al., 2017; Nassar et al., 2019). However, our use of multinomial logistic regression for the transition distributions and Poisson observations render the model non-conjugate. Further, we found that inference using black-box variational inference (BBVI) approaches (Archer et al., 2015; Gao et al., 2016; Linderman et al., 2019) performed relatively poorly in this setting with strong boundary parameter constraints (Figure 3, Appendix A).

We therefore developed an approximate inference algorithm that exploits information about the parameter constraints yet also has favorable scaling. The method combines a variational approximation of the posterior over the discrete states and a Laplace approximation of the posterior over the continuous states. Here we present an overview of the method, variational Laplace EM (vLEM), with additional details in Appendix B.²

We approximate the true posterior over the latent variables with a structured, factorized approximate posterior $p(z, x | \theta, y) \approx q(z)q(x)$, where $q(z)$ is a variational approximation and $q(x)$ is computed via a Laplace approximation. This admits a lower-bound on the marginal likelihood

²A Python implementation of the vLEM algorithm for fitting (r)SLDS models is available at <https://github.com/slinderman/ssm>.

(ELBO)

$$\begin{aligned} \log p(y | \theta) &\geq \mathcal{L}_q(\theta) \\ &= \mathbb{E}_{q(z)q(x)}[\log p(z, x, y | \theta)] \\ &\quad - \mathbb{E}_{q(z)}[\log q(z)] - \mathbb{E}_{q(x)}[\log q(x)]. \end{aligned} \quad (11)$$

We alternate between updating the two approximate posteriors and the model parameters in three steps.

First, the discrete state approximate posterior is updated using the optimal coordinate ascent variational inference update (Bishop, 2006; Blei et al., 2017)

$$q^*(z) \propto \exp(\mathbb{E}_{q(x)}[\log p(z, x, y | \theta)]). \quad (12)$$

This step locally maximizes the ELBO given a fixed $q(x)$ and model parameters θ . We compute the expectation in the update using Monte Carlo samples from $q(x)$. Conditioned on these samples, the posterior $q^*(z)$ has the same factor graph as an HMM and we use the forward-backward algorithm to compute the posterior distributions over the discrete states and the marginal likelihood.

The second step is to update $q(x)$ using a Laplace approximation around the most likely latent path x^* (Paninski et al., 2010; Macke et al., 2011; 2015). That is, we set $q^*(x)$ to be

$$q^*(x) = \mathcal{N}(x^*, -H^{-1}) \quad (13)$$

where

$$x^* = \arg \max_x \mathbb{E}_{q(z)} [\log p(x, z, y | \theta)] \quad (14)$$

$$H = \nabla_x^2 \mathbb{E}_{q(z)} [\log p(x, z, y | \theta)] \Big|_{x=x^*}. \quad (15)$$

Importantly, the Hessian is block-tridiagonal such that linear solves with the Hessian and sampling from $q(x)$ scale linearly in the length of the time-series. We note that this step is not guaranteed to increase the ELBO.

Finally, the model parameters are updated by optimizing the ELBO with respect to the parameters. This corresponds to setting θ to

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{q(z)q(x)} [\log p(x, z, y | \theta)]. \quad (16)$$

We approximate this update using a single sample from $q(x)$ and marginalizing the discrete states. That is, we find $\theta^* = \arg \max_{\theta} \mathbb{E}_{q(z)} [\log p(\hat{x}, z, y | \theta)]$ where $\hat{x} \sim q(x)$. We then set the parameters to a convex combination of the previous and new parameters

$$\theta^{(i+1)} = \alpha \theta^{(i)} + (1 - \alpha) \theta^*. \quad (17)$$

We use either $\alpha = 0.0$ or $\alpha = 0.5$ in our experiments.

This method is a generalization of the Laplace EM method for inference in single-state linear dynamical systems with non-conjugate observations (Paninski et al., 2010; Macke et al., 2011; 2015). When there is only one discrete state the presented method is equivalent to Laplace EM. Methods that combine variational and Laplace approximations were proposed by (Wang & Blei, 2013), and they are particularly well suited to this problem because of the block-tridiagonal structure of the Hessian.

5. Experiments

We demonstrate the constrained recurrent state space modeling framework and vLEM inference with one simulated example and two applications to recordings of neural activity during decision-making tasks.

5.1. Simulated 2D accumulator

We first tested the vLEM inference algorithm on data simulated from a 2D race accumulator model with 2D pulse inputs and Poisson spike count observations. We simulated 250 time series (i.e. trials) of spike counts from $N = 10$ neurons with $T = 100$ time points in each trial, which is similar to the amount of data collected in experiments (Yates et al., 2017). To generate realistic spike counts, we scaled the rates of the Poisson distribution to have a mean of about 0.4 spike counts per 10ms time bin. We fit a 2D race accumulator model to the simulated data using 100 iterations

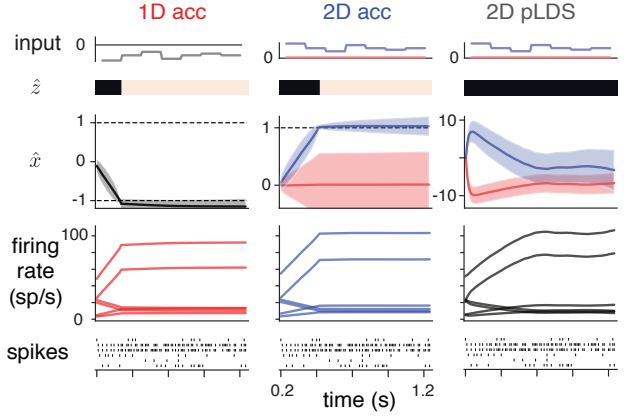


Figure 4. Inferred latent variables on a single trial from 1D and 2D accumulator models and a 2D Poisson LDS fit to LIP responses during a discrete-pulse motion direction discrimination task (Yates et al., 2017). Given the same input and observed spikes, the three models infer similar smoothed firing rates with different underlying latent structure. The most likely discrete state \hat{z} is indicated by the colored bar, and the shaded areas indicate the mean plus two standard deviations of the continuous state posterior $q(x)$.

of the vLEM algorithm (Figure 3A-B). The inferred discrete and continuous states closely matched the true states (Figure 3A). To test parameter recovery with vLEM, we fit the model to 10 different simulated datasets for each of 25, 50, 100, and 250 trials and 10 neurons. We used the same parameters for each simulation. As expected, the MSE between the true and inferred parameters decreased as the number of simulated trials increased (Figure 3C).

We next compared vLEM with BBVI and particle EM. For BBVI, we used a jointly Gaussian posterior over the continuous latent variables with block-tridiagonal structure in the precision matrix and we marginalized the discrete states (Archer et al., 2015; Gao et al., 2016; Linderman et al., 2019). For particle EM, we used a Rao-Blackwellized particle filter to sample from the marginal posterior over the continuous latent variables and we optimized the expected log joint probability (Appendix C). We found that vLEM achieved a higher ELBO than BBVI (Figure 3B) and provided more accurate latent variable estimates and parameters (Appendix A). Notably, vLEM provided comparable accuracy to particle EM at a much lower computational cost (Appendix A).

5.2. 1D vs. 2D accumulator models

We used the modeling framework to compare accumulator models of neural responses in the monkey lateral intraparietal area (area LIP) during a discrete-pulse accumulation task (Yates et al., 2017). In this task, a series of seven motion pulses was presented to the animal over a period of

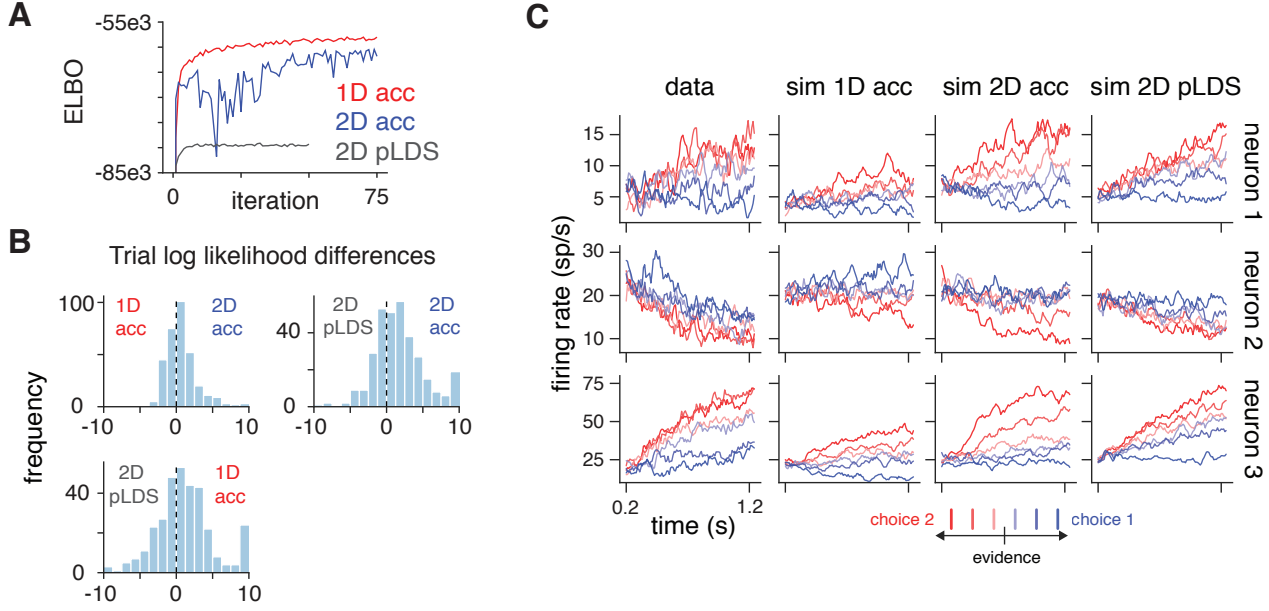


Figure 5. Models of LIP responses during a discrete-pulse motion direction discrimination task (Yates et al., 2017). **A.** ELBO as a function of vLEM iteration. **B.** The distribution of held-out log likelihoods $\Delta_{\mathcal{M}_1 - \mathcal{M}_2}^i = \log p(y_i | \mathcal{M}_1) - \log p(y_i | \mathcal{M}_2)$ across trials i for models \mathcal{M}_1 (right side of plot) and \mathcal{M}_2 (left side of plot). Positive values mean that the model on the right hand side had a higher held-out log likelihood. **C.** The experimental and simulated trial-averaged responses across stimulus conditions. The color indicates the net direction of motion and the darkness indicates the strength of motion.

1050ms. Each pulse had a variable strength of motion in one of two directions. The animal was trained to determine the net direction of motion across the seven pulses.

We fit 1D and 2D accumulation-to-bound models and a 2D Poisson LDS (pLDS) with inputs to the responses of 6 LIP neurons simultaneously recorded during 327 trials of this task (Figures 4-5). The pLDS is equivalent to an unconstrained 2D accumulation without boundaries, with interaction between the two accumulators. We fit the models to the recorded binned spike counts in 10ms bins from the period 200ms after motion onset until 200ms after motion offset. This is the time window in which the neurons putatively accumulate evidence. In the 1D model, the input was the net motion strength at each time point. In the 2D accumulator model, the input was separated into the two directions such that each dimension received only one direction of motion as input.

The fits from the models to an example trial are shown in Figure 4. All three models estimate similar smoothed firing rate trajectories on this trial, with different underlying structure. The accumulator models infer a switch between the accumulation and boundary states in the same direction and at about the same time across the two models. The pLDS model simply infers a smooth 2D latent trajectory underlying the responses. Interestingly, the firing rates in the pLDS model appear to plateau about halfway through the trial, which matches the 2D accumulator model even

Table 2. The mean held-out log likelihood across all trials on Yates et al. (2017) data for each model, with standard error bars computed across ten Monte Carlo sampling runs.

model	mean held-out log likelihood
1D acc	-349.47 ± 0.01
2D acc	-348.74 ± 0.01
2D pLDS	-351.29 ± 0.02

though the pLDS has no boundary.

The trial-averaged responses of three of the neurons are shown in Figure 5C, where the averages are computed over trials with similar net motion strengths. To qualitatively check how well the fit models correspond to the data, we simulated data from the inferred parameters of each of the models. For each stimulus in the training data we simulated spikes counts on three independent trials. We then computed the trial-averaged simulated responses in the same manner as the true responses (Figure 5C). The fit 1D model was unable to capture uniform increases or decreases in the trial-averaged responses across task conditions. However, the fit 2D models produced simulated responses that either decreased (neuron 2) or increased (neuron 3) across all conditions, as seen in the data. Qualitatively, the 2D models appear to better capture the trial-averaged activity.




Figure 6. Ramping model with variable lower boundary fit to the responses of an LIP neuron during a motion direction discrimination task (Roitman & Shadlen, 2002). **A.** The experimental and simulated trial-averaged responses for each task condition. **B.** The ELBO as a function of inference iteration. **C.** The inferred discrete and continuous latent states on three example trials.

We compared the models quantitatively with a five-fold cross validation. For each fold, we estimated the log likelihood of the held-out data with Monte Carlo sampling (Lattimer et al., 2015; Zoltowski et al., 2019). The 2D accumulator model had the highest held-out log likelihood, and both accumulator models had better predictive performance than the 2D pLDS (Figure 5, Table 2). Putting together the quantitative and qualitative comparisons, we find that the two dimensional accumulator model best explains the observed responses out of the considered models.

5.3. Ramping model with variable lower boundary

The transitions in our framework are probabilistic, which enables fitting a model with a variable lower boundary where switches from the accumulation state to the boundary state may occur at different values of x across trials. We fit a ramping model with a variable lower boundary to the spike count responses of an LIP neuron during a reaction-time version of the random dot motion task (Roitman & Shadlen, 2002). To do this, we introduced transition parameters B_{lb} and γ_{lb} that governed the location and sharpness of the lower boundary, and we learned these parameters during inference. We kept a sharp upper boundary fixed at $x = 1$. Example simulated trajectories from this model are shown in Figure 2. We fit the model to the binned spiking responses of an LIP neuron (10ms bins) during 225 trials of a reaction time version of the random dot motion task (Roitman & Shadlen, 2002). We included the period 200ms after motion onset until 50ms before the saccade in our analysis, and

the trials were separated into five categories of stimulus motion in terms of stimulus strength (+high to -high). We learned a separate ramping slope for each stimulus condition. True and simulated trial-averaged responses from the fit model are shown in Figure 6A. We found that a variable lower boundary stopped downward sloping trajectories, as we inferred switches to the lower boundary state at different values of x (Figure 6C).

6. Discussion

We introduced a unifying framework for decision-making models based on constrained recurrent switching linear dynamical systems. We also presented a variational Laplace EM algorithm for inference in rSLDS models. While we applied the algorithm to fit rSLDS models constrained by decision-making theory, we also expect the algorithm to be broadly useful for fitting unconstrained rSLDS models. Finally, we demonstrated our framework by fitting 1D and 2D accumulator models and a ramping model with a variable lower boundary to neural responses during decision-making.

Recent work has used HMM and LDS models to show that choice behavior is dynamic across trials of an experiment (Ashwood et al., 2020; Roy et al., 2020; Stone et al., 2020). Incorporating this observation is an interesting direction of future work, and we expect that using such models of choice behavior in conjunction with our framework will be useful for the study of decision making.

In contrast to the long line of work developing general purpose models of neural dynamics (Smith & Brown, 2003; Yu et al., 2009; Paninski et al., 2010; Macke et al., 2011; Linderman et al., 2017; Wu et al., 2017; Zhao & Park, 2017; Pandarinath et al., 2018; Duncker et al., 2019), here we leveraged long-standing theories of neural computation. The theory-driven approach is natural for this setting, as it offers interpretable latent states, statistical tests of how well specific decision-making dynamics describe neural activity, and a framework for adding complexity as warranted by the data (Linderman & Gershman, 2017). Importantly, the recurrent state space framework not only unifies many existing models, it also suggests generalizations (c.f. Table 1) and greatly expands our toolkit for modeling neural dynamics during decision-making.

Acknowledgements

We thank Jacob Yates and Alex Huk for sharing their LIP data and Jamie Roitman and Michael Shadlen for making their LIP data publicly available. We also thank Benjamin Antin for contributions to the vLEM implementation and Orren Karniol-Tambour for helpful discussions. D.M.Z. was supported by NIH grant T32MH065214. J.W.P. was supported by grants from the Simons Collaboration on

the Global Brain (SCGB AWD543027), the NIH BRAIN initiative (NS104899 and R01EB026946), and a U19 NIH-NINDS BRAIN Initiative Award (5U19NS104648). S.W.L. was supported by NIH grants U19NS113201 and R01NS113119.

References

- Archer, E., Park, I. M., Buesing, L., Cunningham, J., and Paninski, L. Black box variational inference for state space models. *arXiv preprint arXiv:1511.07367*, 2015.
- Ashwood, Z., Urai, A., Aguillon, V., Laboratory, I. B., Churchland, A., Pouget, A., and Pillow, J. Discrete latent states underlie sensory decision-making behavior in mice. In *Cosyne Abstracts*, 2020.
- Barber, D. Expectation correction for smoothed inference in switching linear dynamical systems. *Journal of Machine Learning Research*, 7(Nov):2515–2540, 2006.
- Barber, D. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- Bishop, C. M. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006. ISBN 0387310738.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Bollimunta, A., Totten, D., and Ditterich, J. Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex. *Journal of Neuroscience*, 32(37):12684–12701, 2012.
- Bowman, N. E., Kording, K. P., and Gottfried, J. A. Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. *Neuron*, 75(5):916–927, 2012.
- Brody, C. D. and Hanks, T. D. Neural underpinnings of the evidence accumulator. *Current Opinion in Neurobiology*, 37:149–157, 2016.
- Brunton, B. W., Botvinick, M. M., and Brody, C. D. Rats and humans can optimally accumulate evidence for decision-making. *Science*, 340(6128):95–98, 2013.
- Churchland, A. K., Kiani, R., and Shadlen, M. N. Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6):693, 2008.
- Churchland, A. K., Kiani, R., Chaudhuri, R., Wang, X.-J., Pouget, A., and Shadlen, M. N. Variance as a signature of neural computations during decision making. *Neuron*, 69(4):818–831, 2011.
- DePasquale, B., Pillow, J. W., and Brody, C. D. Accumulated evidence inferred from neural activity accurately predicts behavioral choice. In *Cosyne Abstracts*, 2019.
- Ditterich, J. Stochastic models of decisions about motion direction: behavior and physiology. *Neural Networks*, 19(8):981–1012, 2006.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., and Pouget, A. The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, 32(11):3612–3628, 2012.
- Duncker, L., Böhner, G., Boussard, J., and Sahani, M. Learning interpretable continuous-time models of latent stochastic dynamical systems. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1726–1734, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v97/duncker19a.html>.
- Evans, D. A., Stempel, A. V., Vale, R., Ruehle, S., Lefler, Y., and Branco, T. A synaptic threshold mechanism for computing escape decisions. *Nature*, 558(7711):590, 2018.
- Gao, Y., Archer, E. W., Paninski, L., and Cunningham, J. P. Linear dynamical neural population models through non-linear embeddings. In *Advances in Neural Information Processing Systems*, pp. 163–171, 2016.
- Gold, J. I. and Shadlen, M. N. The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–574, 2007.
- Hanks, T. D. and Summerfield, C. Perceptual decision making in rodents, monkeys, and humans. *Neuron*, 93(1):15–31, 2017.
- Hanks, T. D., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C., and Brody, C. D. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*, 520(7546):220–223, 2015.
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., and Brown, S. D. Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, 35(6):2476–2484, 2015.
- Howard, M. W., Luzardo, A., and Tiganj, Z. Evidence accumulation in a Laplace domain decision space. *Computational Brain & Behavior*, 1(3-4):237–251, 2018.
- Latimer, K. W., Yates, J. L., Meister, M. L., Huk, A. C., and Pillow, J. W. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244):184–187, 2015.

- Linderman, S., Johnson, M., Miller, A., Adams, R., Blei, D., and Paninski, L. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 914–922, 2017.
- Linderman, S., Nichols, A., Blei, D., Zimmer, M., and Paninski, L. Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in *c. elegans*. *bioRxiv*, 2019. doi: 10.1101/621540. URL <https://www.biorxiv.org/content/early/2019/04/29/621540>.
- Linderman, S. W. and Gershman, S. J. Using computational theory to constrain statistical models of neural data. *Current opinion in neurobiology*, 46:14–24, 2017.
- Lueckmann, J.-M., Macke, J. H., and Nienborg, H. Can serial dependencies in choices and neural activity explain choice probabilities? *Journal of Neuroscience*, 38(14):3495–3506, 2018.
- Macke, J. H., Buesing, L., Cunningham, J. P., Byron, M. Y., Shenoy, K. V., and Sahani, M. Empirical models of spiking in neural populations. In *Advances in Neural Information Processing Systems*, pp. 1350–1358, 2011.
- Macke, J. H., Buesing, L., and Sahani, M. Estimating state and parameters in state space models of spike trains. In Chen, Z. (ed.), *Advanced State Space Methods for Neural and Clinical Data*, chapter 6, pp. 137–159. Cambridge University Press, 2015.
- Mazurek, M. E., Roitman, J. D., Ditterich, J., and Shadlen, M. N. A role for neural integrators in perceptual decision making. *Cerebral Cortex*, 13(11):1257–1269, 2003.
- Nassar, J., Linderman, S., Bugallo, M., and Park, I. M. Tree-structured recurrent switching linear dynamical systems for multi-scale modeling. In *International Conference on Learning Representations (ICLR)*, 2019. URL <https://openreview.net/forum?id=HkzRQhR9YX>.
- O’Connell, R. G., Shadlen, M. N., Wong-Lin, K., and Kelly, S. P. Bridging neural and computational viewpoints on perceptual decision-making. *Trends in Neurosciences*, 41(11):838–852, 2018.
- Pandarínath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., et al. Inferring single-trial neural population dynamics using sequential autoencoders. *Nature Methods*, pp. 1, 2018.
- Paninski, L., Ahmadian, Y., Ferreira, D. G., Koyama, S., Rad, K. R., Vidne, M., Vogelstein, J., and Wu, W. A new look at state-space models for neural data. *Journal of Computational Neuroscience*, 29(1-2):107–126, 2010.
- Ratcliff, R. and McKoon, G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922, 2008.
- Ratcliff, R., Smith, P. L., Brown, S. D., and McKoon, G. Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20(4):260–281, 2016.
- Resulaj, A., Kiani, R., Wolpert, D. M., and Shadlen, M. N. Changes of mind in decision-making. *Nature*, 461(7261):263, 2009.
- Roitman, J. D. and Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, 22(21):9475–9489, 2002.
- Roy, N. A., Bak, J. H., Akrami, A., Brody, C. D., Pillow, J. W., et al. Extracting the dynamics of behavior in decision-making experiments. *bioRxiv*, 2020.
- Scott, B. B., Constantinople, C. M., Erlich, J. C., Tank, D. W., and Brody, C. D. Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *eLife*, 4:e11308, 2015.
- Shinn, M., Lam, N. H., and Murray, J. D. A flexible framework for simulating and fitting generalized drift-diffusion models. *bioRxiv*, 2020.
- Shvartsman, M., Srivastava, V., and Cohen, J. D. A theory of decision making under dynamic context. In *Advances in Neural Information Processing Systems*, pp. 2485–2493, 2015.
- Smith, A. C. and Brown, E. N. Estimating a state-space model from point process observations. *Neural Computation*, 15(5):965–991, 2003.
- Stone, I., Bolkan, S., Witten, I., and Pillow, J. Latent-state models reveal a state-dependent contribution of the striatum to decision-making. In *Cosyne Abstracts*, 2020.
- Urai, A. E., De Gee, J. W., Tsetsos, K., and Donner, T. H. Choice history biases subsequent evidence accumulation. *eLife*, 8, 2019.
- Wang, C. and Blei, D. M. Variational inference in nonconjugate models. *Journal of Machine Learning Research*, 14(Apr):1005–1031, 2013.
- Wiecki, T. V., Sofer, I., and Frank, M. J. Hddm: Hierarchical bayesian estimation of the drift-diffusion model in python. *Frontiers in Neuroinformatics*, 7:14, 2013.
- Wu, A., Roy, N. A., Keeley, S., and Pillow, J. W. Gaussian process based nonlinear latent structure discovery in multivariate spike train data. In *Advances in Neural Information Processing Systems*, pp. 3496–3505, 2017.

- Yates, J. L., Park, I. M., Katz, L. N., Pillow, J. W., and Huk, A. C. Functional dissection of signal and noise in mt and lip during decision-making. *Nature Neuroscience*, 20(9): 1285, 2017.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., and Sahani, M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. In *Advances in Neural Information Processing Systems*, pp. 1881–1888, 2009.
- Zhao, Y. and Park, I. M. Variational latent gaussian process for recovering single-trial dynamics from population spike trains. *Neural Computation*, 29(5):1293–1316, 2017.
- Zoltowski, D. M., Latimer, K. W., Yates, J. L., Huk, A. C., and Pillow, J. W. Discrete stepping and nonlinear ramping dynamics underlie spiking responses of lip neurons during decision-making. *Neuron*, 2019. ISSN 0896-6273. doi: <https://doi.org/10.1016/j.neuron.2019.04.031>. URL <http://www.sciencedirect.com/science/article/pii/S0896627319303885>.

Supplementary Material: A general recurrent state space framework for modeling neural dynamics during decision-making

David M. Zoltowski¹, Jonathan W. Pillow^{1,2}, and Scott W. Linderman^{3,4}

¹Princeton Neuroscience Institute, Princeton University, Princeton, NJ

²Department of Psychology, Princeton University, Princeton, NJ

³Department of Statistics, Stanford University, Palo Alto, CA

⁴Wu Tsai Neurosciences Institute, Stanford University, Palo Alto, CA

A Additional Simulated Experiments

Here we demonstrate fitting models with collapsing boundaries and trial-history effects to simulated data. We also compare the variational Laplace EM inference algorithm with black box variational inference and particle EM approaches.

A.1 Nonlinear collapsing boundaries

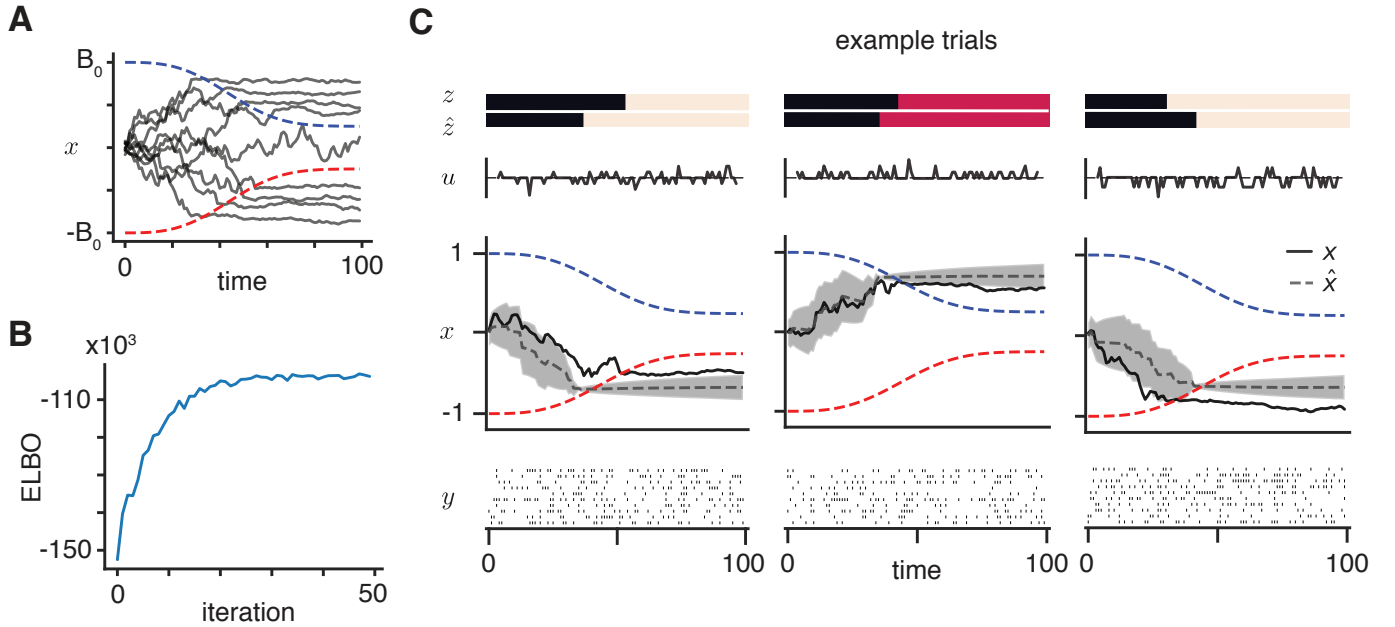


Figure A1. Simulated experiment with nonlinear collapsing boundaries. **A.** Simulated latent trajectories from a one-dimensional accumulation model with collapsing upper and lower boundaries. **B.** ELBO as a function of iteration. **C.** The simulated data and the true and inferred discrete and continuous states for three example trials.

We simulated spike counts from 10 Poisson neurons from a one-dimensional accumulation model with nonlinear collapsing boundaries (Figure A1A). The bin size was $\Delta_t = 0.01$, the trial length was $T = 100$, and the number of trials was 200. The inputs were the difference of two dimensional pulses. The accumulation state parameters were $V_{\text{acc}} = 0.005$ and $\sigma_{\text{acc}}^2 = 0.002$.

The boundary parameterization was

$$b_t = b_0 - (1 - e^{-(t/\lambda)^k})(b_0 - b_\infty) \quad (1)$$

where $b_0 = 1.0$ is the initial height of the boundary and $b_\infty = 0.25$ is the final height of the boundary. The upper and lower boundaries were symmetric across zero. We implemented this model by modifying the transition probabilities to depend on the time-varying boundary.

We fit the nonlinear collapsing boundaries model to the simulated data using 50 iterations of the vLEM algorithm (Figure A1B). The inferred continuous and discrete states from the algorithm were similar to the true latent states (Figure A1C).

A.2 Linear collapsing boundaries

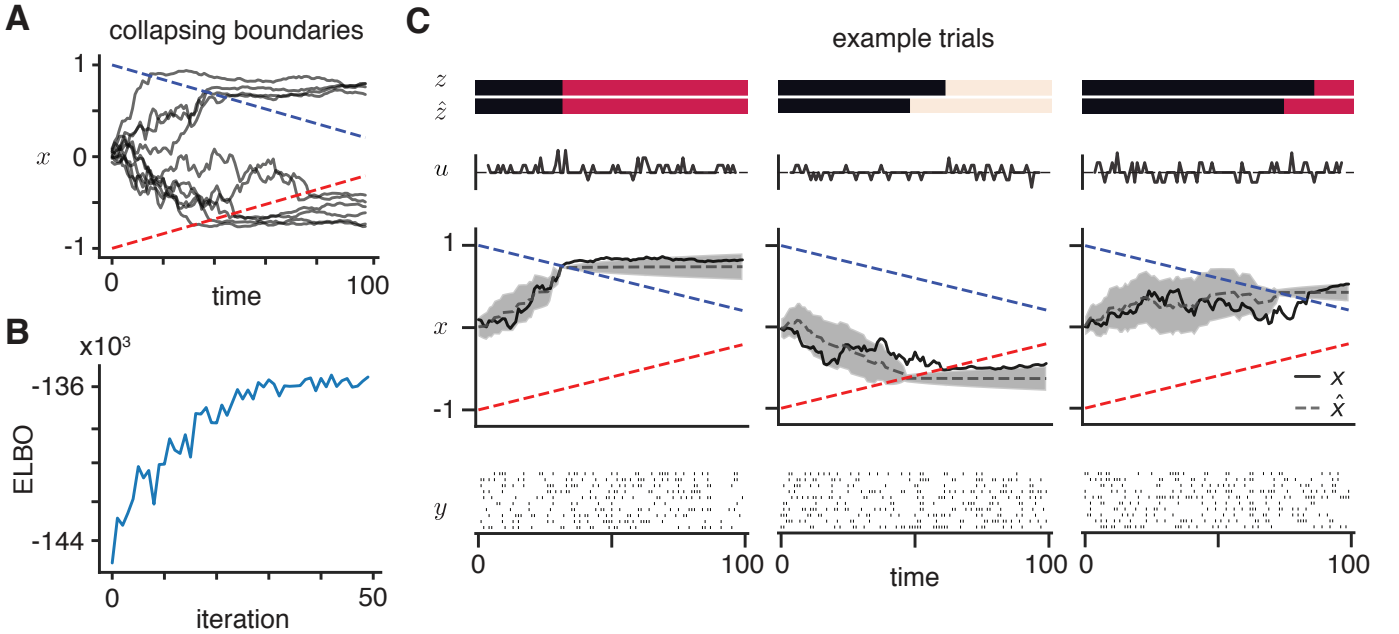


Figure A2. Simulated experiment with linear collapsing boundaries. **A.** Simulated latent trajectories **B.** ELBO as a function of iteration. **C.** The simulated data and the true and inferred discrete and continuous states for three example trials.

We replicated the nonlinear collapsing boundary simulation with a linear collapsing bound, with 10 Poisson neurons, $\Delta_t = 0.01$, $T = 100$, and 200 trials (Figure A2A-C). The inputs were the difference of two dimensional pulses. Here, the accumulation state parameters were $V_{\text{acc}} = 0.01$ and $\sigma_{\text{acc}}^2 = 0.001$.

The boundaries started at ± 1 and collapsed towards zero at a rate of 0.008 per time bin, which means that the final boundary values at $T = 100$ were ± 0.2 . We implemented this model with the following steps. First, we augmented the input vector with the current time of the trial such that $u_t = [s_t, t]^\top$ where s_t is the current stimulus input. Importantly, we set the second dimension of the input weight parameter $V_{\text{acc}}^{(2)}$ to zero so the time is not input to the continuous

dynamics x . We modified the transitions to depend on the input with the following parameterization

$$p(z_t | z_{t-1}, x_{t-1}) \propto \exp \left\{ \gamma (R_{z_{t-1}} + r x_{t-1} + W u_t) \right\}, \quad W = \begin{bmatrix} 0 & 0 \\ 0 & \beta \\ 0 & \beta \end{bmatrix} \quad (2)$$

where β is a scalar parameter that controls the slope of the boundary. We set the left column of W to zeros so the sensory input does not directly affect the transitions. We set $\beta = 0.008$, which corresponds to the rate of the collapsing boundaries as described above. We note that asymmetric collapsing boundaries can be implemented by having separate β parameters for each dimension. While we fix the slope parameter β , its value could be learned. The parameters $R_{z_{t-1}}$ and r have the same form as in the original 1D accumulator model.

A.3 Trial-history




Figure A3. Simulated experiment with trial-history effects. **A.** True and inferred averaged drift rates for positive going trials with and without the choice history bias. The previous choice biases the drift rate upwards (previous choice corresponds to upper boundary) and downwards (previous choice corresponds to lower boundary). **B.** The ELBO as a function of optimization iteration. **C.** True and inferred states for positive going trials with positive (*left*) and negative (*right*) biases.

The modeling framework allows for trial-history effects based on the previous choice, reward, or stimulus. Here we simulated data from a model where the previous trial choice affects the drift rate (Figure A3). We implemented this by including the previous trial choice $c_{\text{prev}} = \{-1, 1\}$ as an additional input covariate. The input at each time point on a given trial was $u_t = [s_t, c_{\text{prev}}]$. In this case, we learn each dimension of the input weights $V_{\text{acc}} \in \mathbb{R}^2$. The element in the second dimension corresponds to the bias in the drift rate. This parameterization enforces a symmetric drift bias, but it is again possible to relax the symmetry.

We simulated spike counts of 5 Poisson neurons from this model with a bin size $\Delta = 0.1$. Each trial had length $T = 100$ and we simulated $N = 200$ trials. In this simulation, the input on each trial was a constant drift of $s_t = 0.015$ for positive going trials and $s_t = -0.015$ for negative going trials. The drift bias was 0.005 and the variance was $\sigma_{\text{acc}}^2 = 0.001$. The average drift rate on positive going trials is shown in Figure A3. The bias increased the average drift when the previous choice was +1 (blue line) and decreased the average drift when the previous choice was -1.

We fit this model using 15 iterations of the vLEM algorithm (Figure A3B). The inferred drift rates in the fit model were similar to the true drift rates (Figure A3A). Next, the inferred latent states correctly followed the bias shown in the true latent states (Figure A3C).

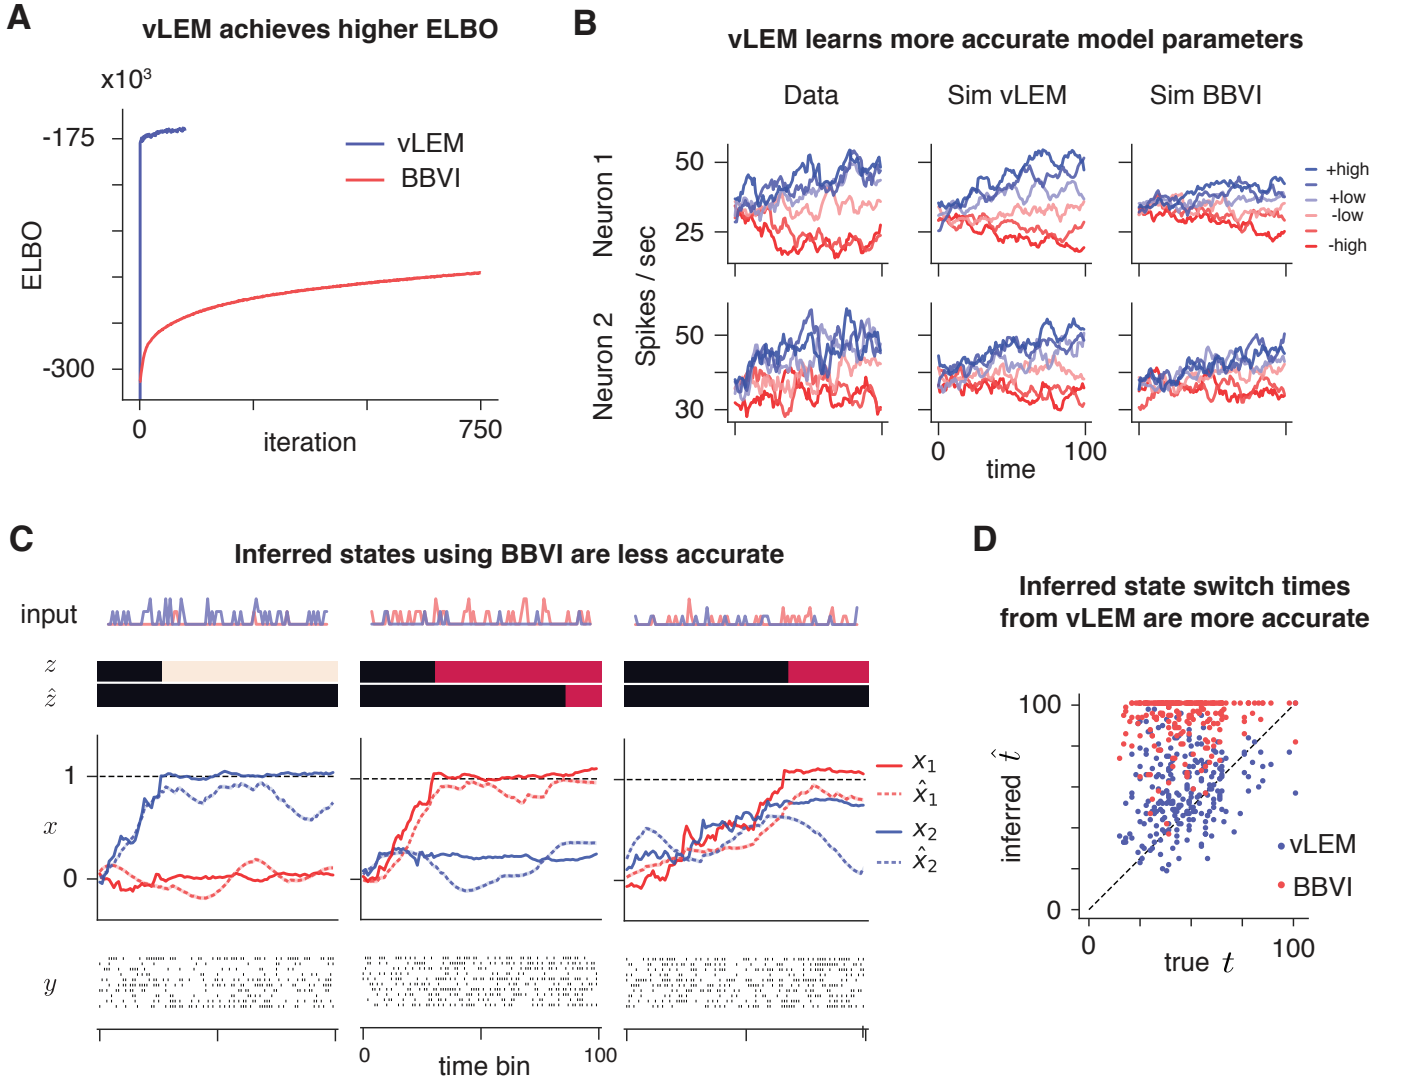


Figure A4. Comparison of vLEM and BBVI for fitting a 2D accumulator model with Poisson neurons. **A.** The ELBO as a function of algorithm iteration (these are the same values as presented in Figure 3). **B.** The true average neural responses across different evidence strengths (line colors) and simulated responses from the fit model using vLEM or BBVI. Here, “+high” is strong stimulus motion towards the “+” direction while “-high” is strong stimulus motion to the opposite “-” direction. **C.** The inferred (dashed lines, \hat{z} , and \hat{x}) and true (solid lines, z , and x) continuous and discrete states using BBVI for three typical example trials. **D.** The true t and inferred \hat{t} transition times from the accumulation state to the boundary state for vLEM and BBVI for all trials.

A.4 Comparison of vLEM and BBVI

Here we present additional results of the comparison between vLEM and BBVI from the simulated 2D accumulator experiment in Section 5.1. As stated previously, we simulated a 2D race accumulator model and the model using vLEM and BBVI. For BBVI, we used a jointly Gaussian posterior over the continuous latent variables with block-tridiagonal structure in the precision of the covariance matrix and we marginalized the discrete states (Archer et al., 2015; Gao et al., 2016; Linderman et al., 2019). We initialized the models with the same parameters and with the same posterior over the continuous latent variables.

Results of fitting the model with BBVI are shown in Figure A4. First, we found that vLEM achieved substantially higher ELBO values (Figure A4A). Next, the learned model parameters from vLEM generated data that are more similar to the true simulated data than BBVI (Figure A4B). This is shown by the similarity in the PSTHs in the first two columns.

Crucially, vLEM provided more accurate inferences about the latent states (Figure A4C). BBVI had difficulty learning transitions from accumulation to boundary and had qualitatively poorer uncertainty estimates (Figure A4C-D). On many trials BBVI did not infer a switch from accumulation to boundary.

A.5 Comparison of vLEM and particle EM

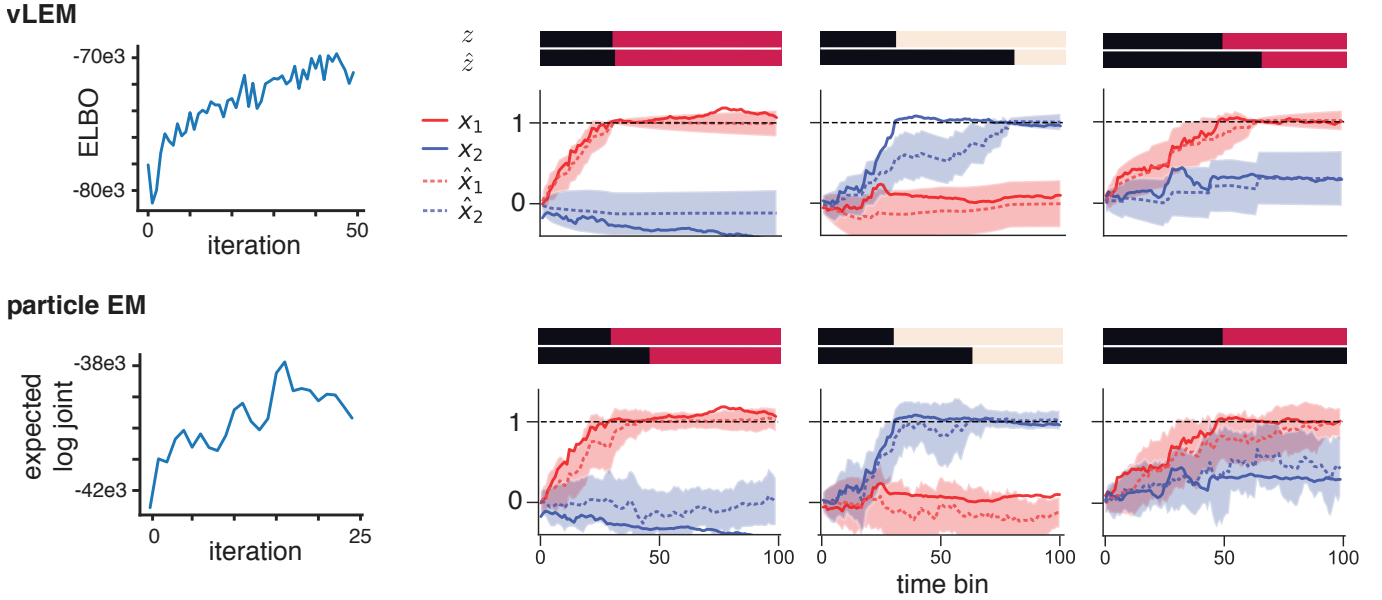


Figure A5. (*top*) ELBO as a function of iteration (*left*) and true and inferred latents on example trials (*right*) when using vLEM to fit the model. (*bottom*) Same as above, except for using particle EM to fit the model and with the expected log joint probability instead of the ELBO.

To test the accuracy of vLEM, we compared vLEM with a particle EM algorithm that used a Rao-Blackwellized particle filter to sample from the marginal posterior over the continuous latent states (see Appendix C). We simulated 100 trials from a 2D accumulator model with 10 neurons and used both vLEM and particle EM to fit the simulated data. (Figure A5). We used $S = 50$ particles in particle EM. This balanced variety in the particles with computational cost, as running the particle EM for 25 iterations took about four hours with our implementation (as opposed to ≈ 10 minutes for 50 iterations of vLEM). The mean squared error between the true and inferred (posterior mean) latent continuous trajectories was smaller for vLEM (0.047 ± 0.0008 for vLEM, 0.064 ± 0.0010 for particle EM). The relatively strong performance of vLEM in this limited data comparison is encouraging. We note that we could improve the particle EM method by optimizing the speed of our implementation, which would allow us to increase the number of particles without incurring a large computational cost. Nonetheless, we consider our current implementation to be a reasonable baseline and are encouraged that vLEM achieves comparable accuracy with lower computational cost.

B Variational Laplace-EM Inference

Here we describe in more detail the variational Laplace-EM inference method. As noted in the main text, we introduce a factorized approximate posterior $q(z)q(x) \approx p(z, x \mid y, \theta)$ over the discrete and continuous latent variables. With those distributions we lower-bound the marginal likelihood with

$$\begin{aligned}\mathcal{L}_q(\theta) &= \mathbb{E}_{q(z)q(x)}[\log p(x, z, y \mid \theta) - \log q(z)q(x)] \\ &= \mathbb{E}_{q(z)q(x)}[\log p(x, z, y \mid \theta)] - \mathbb{E}_{q(z)}[\log q(z)] - \mathbb{E}_{q(x)}[\log q(x)].\end{aligned}$$

To optimize this objective, we alternate between updating 1) $q(z)$, 2) $q(x)$ and 3) θ . The updates to $q(z)$ and θ follow from optimizing the lower bound $\mathcal{L}_q(\theta)$. The update to $q(x)$ is an approximate update and is therefore not guaranteed to increase the value of the lower bound.

B.1 Update discrete state posterior

We update $q(z)$ via the optimal coordinate ascent variational inference update

$$q^*(z) \propto \exp(\mathbb{E}_{q(x)}[\log p(x, z, y \mid \theta)]). \quad (3)$$

To compute this, we expand the expected log joint probability

$$\begin{aligned}\mathbb{E}_{q(x)}[\log p(x, z, y \mid \theta)] &= \mathbb{E}_{q(x)} \left[\log p(z_1, x_1 \mid \theta) + \sum_{t=2}^T \log p(x_t \mid x_{t-1}, z_t, \theta) \right. \\ &\quad \left. + \sum_{t=1}^{T-1} \log p(z_{t+1} \mid z_t, x_t, \theta) + \sum_{t=1}^T \log p(y_t \mid x_t, z_t, \theta) \right] \\ &= \phi(z_1, x_1) + \sum_{t=2}^T \phi(z_t, x_t, x_{t-1}) + \sum_{t=1}^{T-1} \phi(z_t, z_{t+1}, x_t) + \sum_{t=1}^T \phi(z_t, x_t, y_t)\end{aligned}$$

where we have introduced the potentials

$$\begin{aligned}\phi(z_1, x_1) &= \mathbb{E}_{q(x)}[\log p(z_1, x_1 \mid \theta)] \\ \phi(z_t, x_t, x_{t-1}) &= \mathbb{E}_{q(x)}[\log p(x_t \mid x_{t-1}, z_t, \theta)] \\ \phi(z_t, z_{t+1}, x_t) &= \mathbb{E}_{q(x)}[\log p(z_{t+1} \mid z_t, x_t, \theta)] \\ \phi(z_t, x_t, y_t) &= \mathbb{E}_{q(x)}[\log p(y_t \mid x_t, z_t, \theta)].\end{aligned}$$

We used samples from $q(x)$ to estimate the expectations in these potentials. We used a default of a single sample in our simulations and applications to data. We note that if the observations are independent of the discrete states when conditioned on the continuous states (i.e. $\log p(y_t \mid x_t, z_t, \theta) = \log p(y_t \mid x_t, \theta)$) then the emission potential $\phi(z_t, x_t, y_t)$ can be disregarded for updating $q(z)$.

We introduce the normalizing constant $Z(\theta)$ of the distribution such that

$$q(z) = \frac{1}{Z(\theta)} \exp \left(\phi(z_1, x_1) + \sum_{t=2}^T \phi(z_t, x_t, x_{t-1}) + \sum_{t=1}^{T-1} \phi(z_t, z_{t+1}, x_t) + \sum_{t=1}^T \phi(z_t, x_t, y_t) \right). \quad (4)$$

Conditioned on the estimates of the potentials, we have a factor graph equivalent to the factor graph of an HMM. Therefore we compute the unary and pairwise marginals over z and the normalizing constant using the forward-backwards

algorithm. We evaluate the entropy term in the ELBO using the potentials, the unary and pairwise marginals, and the normalizing constant as

$$\begin{aligned}
\mathbb{E}_{q(z)}[\log q(z)] &= \mathbb{E}_{q(z)} \left[-\log Z(\theta) + \phi(z_1, x_1) + \sum_{t=2}^T \phi(z_t, x_t, x_{t-1}) + \sum_{t=1}^{T-1} \phi(z_t, z_{t+1}, x_t) + \sum_{t=1}^T \phi(z_t, x_t, y_t) \right] \\
&= -\log Z(\theta) + \mathbb{E}_{q(z)}[\phi(z_1, x_1)] + \sum_{t=2}^T \mathbb{E}_{q(z)}[\phi(z_t, x_t, x_{t-1})] + \sum_{t=1}^{T-1} \mathbb{E}_{q(z)}[\phi(z_t, z_{t+1}, x_t)] \\
&\quad + \sum_{t=1}^T \mathbb{E}_{q(z)}[\phi(z_t, x_t, y_t)].
\end{aligned}$$

B.2 Update continuous state posterior

We update $q(x)$ with a Laplace approximation around the mode of $\mathbb{E}_{q(z)}[\log p(x, z, y | \theta)]$ such that

$$\begin{aligned}
q^*(x) &= \mathcal{N}(x^*, -H^{-1}) \\
x^* &= \arg \max_x \mathbb{E}_{q(z)}[\log p(x, z, y | \theta)] \\
H &= \nabla_x^2 \mathbb{E}_{q(z)}[\log p(x, z, y | \theta)] \Big|_{x=x^*}.
\end{aligned}$$

To compute the Hessian we expand the terms in the objective

$$\begin{aligned}
\mathcal{L}(x) &= \mathbb{E}_{q(z)}[\log p(x, z, y | \theta)] \\
&= \mathbb{E}_{q(z)} \left[\log p(z_1 | \theta) + \log p(x_1 | z_1, \theta) + \sum_{t=2}^T \log p(x_t | x_{t-1}, z_t, \theta) \right. \\
&\quad \left. + \sum_{t=1}^{T-1} \log p(z_{t+1} | z_t, x_t, \theta) + \sum_{t=1}^T \log p(y_t | x_t, z_t, \theta) \right] \\
&= \phi(x_1, z_1) + \sum_{t=2}^T \phi(x_t, x_{t-1}, z_t) + \sum_{t=1}^{T-1} \phi(x_t, z_t, z_{t+1}) + \sum_{t=1}^T \phi(x_t, y_t, z_t) + \text{const}
\end{aligned}$$

where

$$\begin{aligned}
\phi(x_1, z_1) &= \mathbb{E}_{q(z)}[\log p(x_1 | z_1, \theta)] = \sum_k q(z_1 = k) \log p(x_1 | z_1 = k, \theta) \\
\phi(x_t, x_{t-1}, z_t) &= \mathbb{E}_{q(z)}[\log p(x_t | x_{t-1}, z_t, \theta)] = \sum_k q(z_t = k) \log p(x_t | x_{t-1}, z_t = k, \theta) \\
\phi(x_t, z_t, z_{t+1}) &= \mathbb{E}_{q(z)}[\log p(z_{t+1} | z_t, x_t, \theta)] = \sum_k \sum_j q(z_t = k, z_{t+1} = j) \log p(z_{t+1} = j | z_t = k, x_t, \theta) \\
\phi(x_t, y_t, z_t) &= \mathbb{E}_{q(z)}[\log p(y_t | x_t, z_t, \theta)] = \sum_k q(z_t = k) \log p(y_t | x_t, z_t = k, \theta).
\end{aligned}$$

The above derivation was written in full generality. If the emission potential does not depend on the discrete state then the emission potential simplifies to $\phi(x_t, y_t, z_t) = \log p(y_t | x_t, \theta)$. Also, if there are no recurrent dependencies (as in a standard SLDS) then the transition term $\log p(z_{t+1} | z_t, x_t, \theta)$ is equal to $\log p(z_{t+1} | z_t, \theta)$ and therefore the transition potential $\phi(x_t, z_t, z_{t+1})$ no longer depends on x_t .

We require the Hessian matrix for the Laplace approximation. This matrix is given by

$$\begin{aligned}\nabla_x^2 \mathcal{L}(x) &= \nabla_x^2 \mathbb{E}_{q(z)}[\log p(x, z, y \mid \theta)] \\ &= \nabla_x^2 \phi(x_1, z_1) + \sum_{t=2}^T \nabla_x^2 \phi(x_t, x_{t-1}, z_t) + \sum_{t=1}^{T-1} \nabla_x^2 \phi(x_t, z_t, z_{t+1}) + \sum_{t=1}^T \nabla_x^2 \phi(x_t, y_t, z_t)\end{aligned}$$

where

$$\begin{aligned}\nabla_x^2 \phi(x_1, z_1) &= \sum_k q(z_1 = k) \nabla_x^2 \log p(x_1 \mid z_1 = k, \theta) \\ \nabla_x^2 \phi(x_t, x_{t-1}, z_t) &= \sum_k q(z_t = k) \nabla_x^2 \log p(x_t \mid x_{t-1}, z_t = k, \theta) \\ \nabla_x^2 \phi(x_t, z_t, z_{t+1}) &= \sum_k \sum_j q(z_t = k, z_{t+1} = j) \nabla_x^2 \log p(z_{t+1} = j \mid z_t = k, x_t, \theta) \\ \nabla_x^2 \phi(x_t, y_t, z_t) &= \sum_k q(z_t = k) \nabla_x^2 \log p(y_t \mid x_t, z_t = k, \theta).\end{aligned}$$

Therefore, we can compute the Hessian by computing the contributions to the Hessian of the dynamics, emission, and transition potentials.

The Hessian has size $TD \times TD$ for a time series of length T with latent dimensionality D but has a sparse, block tridiagonal structure with blocks of size $D \times D$. The terms in the Hessian from the initial state, transition, and emission potentials only contribute terms to the primary block diagonal. The dynamics potentials contribute terms to both the primary and first off-diagonal blocks. Throughout, we only represent and store the main and lower diagonal blocks of the Hessian. This reduces storage from the full $(TD)^2$ to $(2T - 1)D^2$ such that it is linear in T . For linear solves and matrix inversions of the Hessian, we also use algorithms that exploit the block tridiagonal structure.

To find the most likely latent path x^* , we use Newton's method with a backtracking line search. However, we can also use optimization routines that require only gradient information (IBFGS) or require only gradient information and Hessian-vector products (Newton-CG or trust-region Newton-CG).

B.3 Update parameters

We update the model parameters by approximately optimizing the ELBO with respect to the parameters

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{q(z)q(x)}[\log p(x, z, y \mid \theta) - \log q(z)q(x)]. \quad (5)$$

Instead of optimizing the expectation under the full distribution of $q(x)$ we optimize

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{q(z)}[\log p(\hat{x}, z, y \mid \theta)] \quad (6)$$

where \hat{x} is a sample from $q(x)$ and we have dropped terms that do not depend on θ . Conditioned on \hat{x} , the update consists of M-steps on the transition, dynamics, and emission parameters. We use either exact updates (where applicable) or IBFGS to implement the M-steps. Finally, we set the parameters at iteration i via a convex combination of the new parameters θ^* and the parameters at the previous iteration

$$\theta_i = (1 - \alpha) \theta^* + \alpha \theta_{i-1}. \quad (7)$$

We note that we can also update the parameters using stochastic gradient ascent with samples from $q(x)$.

B.4 Initialization

We can exploit the known structure of the 1D and 2D accumulation-to-bound models to initialize some of the parameters. For the 1D and 2D accumulation-to-bound models we set the emission parameter d to the mean spike counts across trials in the first three time bins. In the 1D model, we set the emission parameter C using the firing rate at the end of trials with strong input to the upper (λ_{UB}) and lower (λ_{LB}) boundaries. Given those values for each neuron and the fact that the boundaries are at ± 1 for this model, we set $C = \frac{1}{2}(\lambda_{UB} - \lambda_{LB})$. In the 2D model, for each neuron we initialized the elements of C as the difference between the firing rate at the end of trials with strong net input and the mean rate d . We did this for each dimension of the input and corresponding element in C . For the models and data in this paper, we did not identify procedures to reliably estimate the initial underlying latent dynamics parameters. Therefore we randomly initialized the input weights and dynamics variance and set the initial dynamics matrix to $A_{acc} = I$.

C Particle EM

We compared vLEM with a particle EM inference method. The first step of this method is to use a Rao-Blackwellized particle filter to obtain S Monte Carlo samples from the marginal posterior over the latent continuous states

$$x_{1:T}^s \sim p(x_{1:T} \mid y_{1:T}, \theta) \quad (8)$$

for $s = 1, \dots, S$ and for each time series. The second step is to use the samples from the posterior to estimate the expected log joint probability

$$\mathbb{E}_{p(x,z|y,\theta)}[\log p(x, z, y \mid \theta)] \approx \frac{1}{S} \sum_{s=1}^S \mathbb{E}_{p(z|x^s,y,\theta)}[\log p(x^s, z, y \mid \theta)], \quad x^s \sim p(x_{1:T} \mid y_{1:T}, \theta) \quad (9)$$

$$= \hat{\mathbb{E}}_{p(x,z|y,\theta)}[\log p(x, z, y \mid \theta)] \quad (10)$$

where we have dropped the subscripts denoting the entire time series. Finally, we update the parameters by maximizing the sample expectation of the log joint

$$\theta^* = \arg \max_{\theta} \hat{\mathbb{E}}_{p(x,z|y,\theta)}[\log p(x, z, y \mid \theta)]. \quad (11)$$

We used the L-BFGS optimizer to maximize this objective.

C.1 Particle filter

Here we describe the Rao-Blackwellized particle filter used in the particle EM algorithm. The posterior over the time series x is

$$p(x_{1:t} \mid y_{1:t}) \propto p(x_{1:t}, y_{1:t}) \quad (12)$$

$$= p(y_t \mid x_{1:t}, y_{1:t-1}) p(x_t \mid x_{1:t-1}, y_{1:t-1}) p(x_{1:t-1} \mid y_{1:t-1}) \quad (13)$$

$$= p(y_t \mid x_t) p(x_t \mid x_{1:t-1}) p(x_{1:t-1} \mid y_{1:t-1}). \quad (14)$$

Let $q(x_{1:t} \mid y_{1:t})$ be the importance sampling density used in the particle filter. If we propose particles according to the prior distribution, $q(x_t \mid x_{1:t-1}, y_{1:t-1}) = p(x_t \mid x_{1:t-1})$, then the particle filter weights are

$$w_t^s \propto \frac{p(x_{1:t}^s \mid y_{1:t})}{q(x_{1:t}^s \mid y_{1:t})} \quad (15)$$

$$= \frac{p(y_t | x_t^s) p(x_t^s | x_{1:t-1}^s) p(x_{1:t-1}^s | y_{1:t-1})}{q(x_t^s | x_{1:t-1}^s, y_{1:t-1}) q(x_{1:t-1}^s | y_{1:t-1})} \quad (16)$$

$$= p(y_t | x_t^s) w_{t-1}^s. \quad (17)$$

The proposal can be written via a marginalization over z_t

$$p(x_t | x_{1:t-1}) = \sum_{z_t} p(x_t | z_t, x_{t-1}) p(z_t | x_{1:t-1}). \quad (18)$$

Therefore we can sample from $p(x_t | x_{1:t-1})$ with

$$z_t \sim p(z_t | x_{1:t-1}) \quad (19)$$

$$x_t \sim p(x_t | z_t, x_{t-1}) \quad (20)$$

and ignoring z_t . This procedure requires the ‘look-ahead’ posterior $p(z_t | x_{1:t-1})$ which can be computed using the filtered posterior $p(z_t | x_{1:t-1})$ and transition probabilities

$$p(z_t | x_{1:t-1}) = \sum_{z_{t-1}} p(z_t | z_{t-1}, x_{t-1}) p(z_{t-1} | x_{1:t-1}). \quad (21)$$

We compute and store the filtered posterior $p(z_t | x_{1:t})$ for each time point t .

The final Rao-Blackwellized particle filter is given in Algorithm 1.

Algorithm 1 Rao-Blackwellized particle filter for rSLDS

- 1: Input: observations y , inputs u , and number of particles S
- 2: Initialize particles x_1^s , weights w_1^s , and initial $p(z_1 | x_1^s)$ for $s = \{1, \dots, S\}$.
- 3: **for** $t = 2$ to T **do**
- 4: **for** $s = 1$ to S **do**
- 5: Compute look-ahead posterior using transition matrix

$$p(z_t | x_{1:t-1}^s) = \sum_{z_{t-1}} p(z_t | z_{t-1}, x_{t-1}^s) p(z_{t-1} | x_{1:t-1}^s)$$

- 6: Sample from look-ahead posterior and dynamics model

$$z_t \sim p(z_t | x_{1:t-1}^s)$$

$$x_t^s \sim p(x_t | z_t, x_{t-1}^s)$$

- 7: Compute posterior

$$p(z_t | x_{1:t}^s) \propto p(x_t^s | z_t) \sum_{z_{t-1}} p(z_t | z_{t-1}, x_{t-1}^s) p(z_{t-1} | x_{1:t-1}^s)$$

- 8: Compute likelihood $p(y_t | x_t^s)$
 - 9: Multiply weight $w_t^s = w_{t-1}^s p(y_t | x_t^s)$
 - 10: **end for**
 - 11: Normalize weights $w_t^s = w_t^s / \sum_t w_t^s$
 - 12: Resample particle indices $(I_1, \dots, I_S) \sim \text{Mu}(S, (w_1^1, \dots, w_t^S))$
 - 13: For each index assign $x_{1:t}^{s_{\text{new}}} = x_{1:t}^{I_s}$ and $p(z_t | x_{1:t}^{s_{\text{new}}}) = p(z_t | x_{1:t}^{I_s})$
 - 14: **end for**
 - 15: Return $x_{1:T}^s, w_{1:t}^s$ for all s
-

References

- Archer, E., Park, I. M., Buesing, L., Cunningham, J., and Paninski, L. (2015). Black box variational inference for state space models. *arXiv preprint arXiv:1511.07367*.
- Gao, Y., Archer, E. W., Paninski, L., and Cunningham, J. P. (2016). Linear dynamical neural population models through nonlinear embeddings. In *Advances in Neural Information Processing Systems*, pages 163–171.
- Linderman, S., Nichols, A., Blei, D., Zimmer, M., and Paninski, L. (2019). Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in c. elegans. *bioRxiv*.