

# MARKED INDUCING POINT CASCADED SDES FOR NEURAL MANIFOLD LEARNING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

The manifold hypothesis suggests that high-dimensional neural time series lie on a low-dimensional manifold shaped by simpler underlying dynamics. To uncover this structure, latent dynamical variable models such as state-space models, recurrent neural networks, neural ordinary differential equations, and Gaussian process latent variable models are widely used. We propose MIP-CSDE (Marked Inducing Point Cascaded Stochastic Differential Equations), a novel cascaded stochastic differential equation model that balances computational efficiency with interpretability and addresses key limitations of existing approaches. Our model assumes that a sparse set of trajectory samples suffices to reconstruct the underlying smooth manifold. The manifold dynamics are modeled using a set of Brownian bridge SDEs, with points – specified in both time and value – drawn from a multivariate marked point process. These Brownian bridges define the drift of a second set of SDEs, where their trajectories are mapped to the observed data. This cascade model structure enables continuous, differentiable latent processes that can model arbitrarily complex time series as the number of inducing points increases. For MIP-CSDE, we derive training and inference procedures, and demonstrate that computational complexity of its inference step scales as  $\mathcal{O}(P \cdot N)$ , exhibiting linear dependence on the observation data length  $N$ , where  $P$  is the number of particles. We then show its application in both synthetic data and neural recordings, where our proposed model accurately recovers the underlying manifold structure and scales effectively with data dimensionality.

## 1 INTRODUCTION

The manifold hypothesis proposes that high-dimensional neural time series lie on a low-dimensional manifold shaped by simpler latent dynamics (Whiteley et al., 2024). Evidence for such structure can be found in auditory cortex activity (Bondanelli et al., 2021) and in speech signals constrained by vocal tract mechanics (Gonzalez-Castillo et al., 2023). Methods for uncovering latent manifolds include state-space models (SSMs) (Särkkä and Svensson, 2023), dynamical autoencoders (Girin et al., 2020), switching SSMs (Ghahramani and Hinton, 2000), Gaussian and Dirichlet processes (Fox et al., 2008; Eleftheriadis et al., 2017; Wang et al., 2005), t-SNE and UMAP (Van der Maaten and Hinton, 2008; McInnes et al., 2018), and Latent Neural ODEs (Rubanova et al., 2019). In this paper, we focus on SSMs for high-dimensional neural time series. Classical models include Linear Gaussian SSMs (Kitagawa and Gersch, 1996) and Hidden Markov Models (Rabiner, 2002), while modern variants include Deep SSMs (Rangapuram et al., 2018), Deep Kalman Filters (Krishnan et al., 2015a), Gaussian Process Dynamical Models (GPDMs) (Wang et al., 2005), Gaussian Process SSMs (GPSSMs) (Eleftheriadis et al., 2017), and nonlinear latent-variable models such as Latent Factor Analysis via Dynamical Systems (LFADS) (Sussillo et al., 2016) and Gaussian Process Factor Analysis (GPFA) (Yu et al., 2008). A critical challenge is to develop robust and computationally efficient inference and training procedures, where recent approaches such as SDE Inference via Natural Gradients (SING) (Hu et al., 2025) improve inference for latent SDEs. Despite these advances, limitations remain: some models fail to capture oscillatory dynamics, others require structural constraints, and deep neural network (DNN) based approaches are data-hungry. Sequential models like recurrent neural networks (RNNs) and long short-term memory networks (LSTMs) (Chang et al., 2024) capture nonlinear dependencies but pose interpretability and training challenges (Glorot and Bengio, 2010). To address these limitations, we propose MIP-CSDE which balances interpretability and expressive power. Inspired by findings that neural manifolds evolve smoothly along low-dimensional trajectories (Cunningham and Yu, 2014a; Gosztolai et al., 2023), our proposed model assumes that a sparse set of trajectory samples suffices to reconstruct the mani-

054  
055  
056  
057  
058  
059  
060  
061  
062  
063  
064  
065  
066  
067  
068  
069  
070  
071  
072  
073  
074  
075  
076  
077  
078  
079  
080  
081  
082  
083  
084  
085  
086  
087  
088  
089  
090  
091  
092  
093  
094  
095  
096  
097  
098  
099  
100  
101  
102  
103  
104  
105  
106  
107

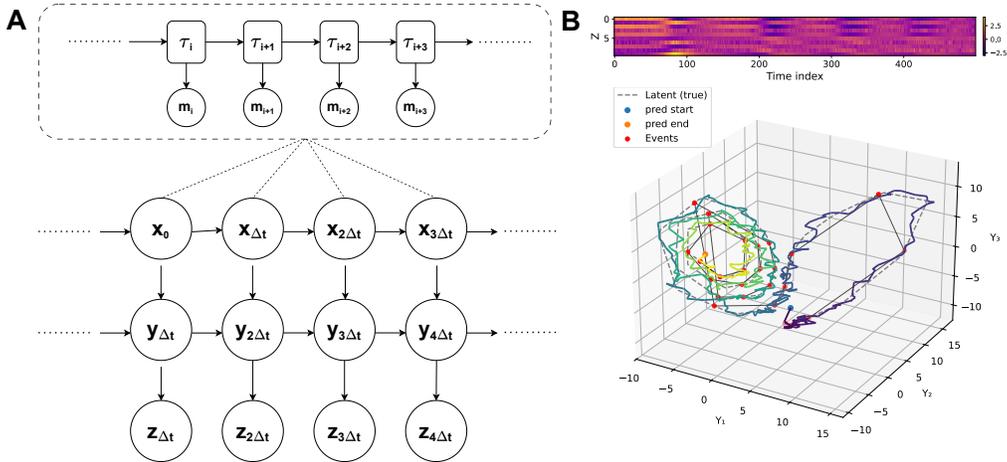


Figure 1: **MIP-CDSE Model Structure and its Application in Inferring Lorenz System Trajectory:** (A) Shows the graphical representation of our proposed model, where inducing points ( $\tau_i, m_i$ ) go through two layers of SDEs ( $X_t, Y_t$ ) followed by projection to the observation domain ( $Z_k$ ). (B) Top row shows the observed data used as input to our proposed model. This data is generated by projecting the Lorenz trajectory into a 10-dimensional observation space, where the mapping is defined by a linear projection with additive multivariate Gaussian noise. The bottom row shows the underlying Lorenz trajectory and its estimation by our model, along with the timing of the inducing points. MIP-CDSE simultaneously learns the mapping from the manifold to the observed space and infers the trajectory of the Lorenz attractor. The timing and values of the inducing points reflect their adaptive behavior in capturing both fast and slow transitions along the trajectory.

fold. The first layer generates trajectories with Brownian bridge SDEs, using inducing points from a multivariate marked point process (Daley and Vere-Jones, 2006; Oksendal, 2013). These trajectories define the drift for a second set of SDE layer, whose outputs map to observed data. This cascaded structure yields a continuous, differentiable latent process capable of modeling arbitrarily complex dynamics as inducing points increase. We derive efficient inference and training procedures, with their computational complexity scale linearly with data length. We then show that MIP-CDSE robustly and accurately recovers the latent manifolds in both synthetic and neural datasets. Figure 1 illustrates the MIP-CDSE graphical model and its application to synthetic data generated using the Lorenz dynamics.

The paper is organized as follows. Section 2 presents MIP-CDSE, its properties, and inference procedure. Section 3 shows model application in simulated and neural data for manifold inference and decoding tasks. Section 4 and 5 respectively cover the discussion and conclusion sections.

## 2 MATERIALS & METHODS

Neural population activity is widely observed to evolve along low-dimensional latent manifolds whose trajectories exhibit smooth temporal structure Cunningham and Yu (2014b). Motivated by this observation, we propose a cascade SDE framework combined with inducing points spanning both the time and value spaces to characterize neural population dynamics and their relationship to an underlying smooth manifold. In this section, we first define the components of our model, including the generation of inducing points and the SDEs that map these points to the observed time series data. We then establish the universal approximation properties of the proposed model and develop its training and inference procedures.

### 2.1 CASCADE SDE AND INDUCING POINT MOTIVATION

Our objective is to construct a generative model capable of producing trajectories that are continuous and smooth, i.e., with well-defined temporal derivatives. Standard diffusion processes or single-layer SDEs typically produce sample paths that are continuous but nowhere differentiable, making them ill-suited for modeling smooth neural trajectories.

To generate flexible continuous paths, we first employ a Brownian-bridgebased SDE in which inducing points define a piecewise-linear interpolation, enabling the construction of complex continuous functions. However, such trajectories are not smooth. To obtain differentiable paths, we introduce a second SDE whose drift is driven by the trajectory of the first SDE. By pushing the noise variance of this second SDE toward zero, the resulting trajectories converge to smooth, differentiable functions while preserving the flexibility provided by the inducing-point structure.

In this cascade system, the first SDE with inducing points offers a piecewise-linear approximation of any continuous function, while the second SDE ensures that the overall trajectory is smooth and differentiable, aligning with the smooth manifold structure underlying neural population activity.

## 2.2 CASCADE SDE FRAMEWORK

Figure 1A illustrates the model structure. The inducing points are a set of event-value pairs that sample the underlying manifold in both time and value space. Each event is characterized by a time  $t_i$  and an associated mark vector  $\vec{m}_i$ . The model considers an arbitrary finite sequence of these pairs, represented as the set  $\mathcal{I} = \{(\vec{m}_i, t_i); i = 1, 2, \dots\}$ . The joint probability distribution of event-value pairs over a period  $T$  for  $L$  events is given by:

$$p\left(\{(t_i, \vec{m}_i)\}_{i=1}^L, T\right) = \exp\left(-\int_0^T \lambda(t | \mathcal{H}_t) dt\right) \prod_{i=1}^L \lambda(t_i | \mathcal{H}_{t_i}) p(\vec{m}_i | t_i, \mathcal{H}_{t_i}) \quad (1)$$

where  $\lambda(t_i | \mathcal{H}_{t_i})$  is the event occurrence rate conditioned on the history of previous events  $\mathcal{H}_i$ , and  $p(\vec{m}_i | t_i, \mathcal{H}_{t_i})$  defines the mark distribution conditioned on the event time  $t_i$  and the history  $\mathcal{H}_i$  (Jacobsen, 2006). The sequence of events can also be described using waiting times, defined as  $\tau_i = t_i - t_{i-1}$ , which transforms the process into a renewal marked point process (Daley and Vere-Jones, 2006).

With the inducing points generated, we now introduce the remaining components of the model that map these points to the observed data. Let  $Z_k \in \mathbb{R}^M$  denote the observed data at discrete time points  $k = 1, \dots, K$ , which are modeled as functions of an underlying continuous latent process  $Y_t \in \mathbb{R}^D$ , where  $D \ll M$ .  $Y_t$  evolves according to another latent process  $X_t$ , which has the same dimension as  $Y_t$ . The latent process  $X_t = \{x_t^d\}_{d=1}^D$  is modeled using a set of SDEs, where the process is constrained to reach the mark values at times specified by the inducing time points - i.e, a Brownian bridge SDE (Pitman and Yor, 1999). Evolution of state process in each dimension  $d = 1, \dots, D$  is defined by:

$$dx_t^d = \mu_t^d dt + \sigma_t^d dw_t^d, \quad (2)$$

where  $w_t^d$  is a standard Wiener process,  $\mu_t^d$  is the drift term, and  $\sigma_t^d$  is the time-dependent diffusion coefficient. For  $t \in [t_i, t_{i+1})$ , which corresponds to the waiting period  $\tau_{i+1}$ , the drift and diffusion terms are defined as:

$$\mu_t^d = \frac{m_{i+1}^d - x_t^d}{t_{i+1} - t}, \quad \sigma_t^d = \sqrt{\frac{(t_{i+1} - t)(t - t_i)}{t_{i+1} - t_i}}, \quad (3)$$

where  $m_{i+1}^d$  is the  $d$ -th component of the mark vector  $\vec{m}_{i+1}$ , the value that the process must reach at the event time  $t_{i+1}$ . The latent process  $Y_t = \{y_t^d\}_{d=1}^D$  evolves according to:

$$dy_t^d = (-\zeta y_t^d + x_t^d) dt + \sigma_y^d d\nu_t^d, \zeta > 0 \quad (4)$$

where the drift term for  $y_t^d$  is defined by  $x_t^d$ ,  $\nu_t^d$  is a standard Wiener process, and  $\sigma_y^d$  is the diffusion coefficient for the  $d$ -th component. where we generate  $\sigma_y^d = 0$ , which guarantees first-order differentiability (no sharp jump or edge). We generally set  $\zeta$  toward 1, but it can be set to smaller values or trained. Finally, the observations  $Z_k$  are defined as:

$$Z_k = W \begin{pmatrix} y_{k\Delta t}^1 \\ \vdots \\ y_{k\Delta t}^D \end{pmatrix} + \varepsilon_k, \quad (5)$$

where  $\Delta t$  denotes the sampling interval at which the observations  $Z_k$  are collected,  $W \in \mathbb{R}^{M \times D}$  is a linear projection matrix, and  $\varepsilon_k \sim \mathcal{N}(0, R)$  represents Gaussian noise with covariance  $R \in \mathbb{R}^{M \times M}$ . Here, we assume a linear projection with additive noise; more generally, the framework can accommodate more complex and non-linear mappings. For instance, for the hippocampus data analyzed in Section 3.3, the mapping between  $Y_t$  and  $Z_k$  is a non-linear function characterizing the rate function for a point-process observation.

## 2.3 MODEL PROPERTIES

In this section, we discuss two key attributes of MIP-CSDE: its universal approximation capability and its computational cost. Other aspects of the model, including its nonparametric nature and strategies for managing the growth of inducing points, are discussed in the Appendix A.1.

### 2.3.1 UNIVERSAL APPROXIMATION PROPERTY

When the process is deterministic and the inducing points are equally spaced, we can rely on the sampling theorem which suggests that a signal can be completely reconstructed from its samples (Shannon, 1949). In simple terms, any continuous function can be reconstructed from properly sampled data points. Here, we extend a similar idea to the cascade SDEs using the inducing points.

**Theorem:** Let  $f \in C([0, T])$  be a continuous function and let  $\varepsilon > 0$  be arbitrary. Then there exists a choice of inducing points such that the expected one-dimensional component of the process,  $\mathbb{E}[s_t]$ , uniformly approximates the integral

$$g(t) = \int_0^t f(s) ds \quad (6)$$

within error  $\varepsilon$ , in the sense that

$$\sup_{t \in [0, T]} |\mathbb{E}[s_t] - g(t)| < \varepsilon. \quad (7)$$

Here,  $s_t$  denotes the  $d$ -th component  $y_t^d$  of the process  $Y_t = \{y_t^d\}_{d=1}^D$ .

**Proof:** In our model, each component is given by  $s_t = \int_0^t x_s^d ds + \sigma W_t^d$ , where  $x_t^d$  is a Brownian bridge,  $W_t^d$  is a standard Brownian motion, and  $\sigma$  is the noise variance.

Let  $\{t_i\}_{i=1}^N$  be uniformly spaced inducing points, each associated with a mark  $m_i$ . Define  $x_t^d$  by piecewise linear interpolation of these inducing points. Since piecewise linear functions are dense in  $C([0, T])$ , we can choose  $\{m_i\}$  such that  $x_t^d \rightarrow f(t)$  uniformly on  $[0, T]$ . Define  $s_t = \int_0^t x_s^d ds$ . Because integration preserves uniform convergence, it follows that

$$\mathbb{E}[s_t] \rightarrow g(t) = \int_0^t f(s) ds \quad \text{uniformly on } [0, T] \quad (8)$$

The noise term satisfies  $\mathbb{E}[\sigma W_t^d] = 0$  and  $\text{Var}(\sigma W_t^d) = \sigma^2 t$ . Applying Chebyshev's inequality, we have  $P(|s_t - \mathbb{E}[s_t]| \geq \eta) \leq \frac{\sigma^2 T}{\eta^2}$ , so for sufficiently small  $\sigma$ , the process  $s_t$  concentrates around its expectation. In Appendix A.2, we establish that as  $N \rightarrow \infty$ , the spacing of the inducing points converges to  $T/N$ . Thus, by selecting appropriate inducing points  $\{(t_i, m_i)\}$ , we ensure

$$\sup_{t \in [0, T]} |\mathbb{E}[s_t] - g(t)| < \varepsilon \quad (9)$$

**Corollary:** Let  $Y_t \in \mathbb{R}^D$  be a continuous vector-valued function on  $[0, T]$ . For any  $\varepsilon > 0$ , there exists a choice of inducing points and model parameters such that each component of the MIP-CSDE latent process  $Y_t^{(\text{model})}$  satisfies

$$\sup_{t \in [0, T]} \left\| \mathbb{E}[Y_t^{(\text{model})}] - Y_t \right\|_2 < \varepsilon \quad (10)$$

In particular, any continuous multidimensional time series that can be expressed as a continuous transformation of such a latent process (e.g., through Eq. (5)) can be approximated arbitrarily well by MIP-CSDE, for suitable latent dimension and inducing-point density.

**Proof:** This corollary extends the one-dimensional approximation result to continuous multidimensional trajectories. It does not specify a unique manifold representation, but shows that such a representation can be constructed given sufficient latent dimensionality and a suitably dense set of inducing points. We can assume that each dimension of  $Y_t$  is independent of the others; thus, we only need to adjust the corresponding inducing point values to capture each specific dynamic. Note that as the number of sample points increases, convergence to the bound is achieved using the same set of inducing point times across dimensions.

This proof does not specify the manifold representation but shows that any such representation can be constructed with suitable dimensions and a sparse set of inducing points. In practice, marks and times are adjusted from the observations, creating dependencies across dimensions and among the inducing points.

### 2.3.2 COMPUTATIONAL COST

A key advantage of the proposed model lies in its favorable computational complexity compared to other non-parametric models such as GPs. Standard GPs require inversion of an  $N \times N$  covariance matrix in their prediction step, which results in a computational complexity of  $\mathcal{O}(N^3)$  (where  $N$  is the number of samples or time points) (Seeger, 2004). This scaling substantially restricts the applicability of GP for long temporal sequences or high-frequency data. In contrast, as we show in the next section, inference for the discrete representation of our model can be performed using a sequential Monte Carlo (SMC) (Doucet et al., 2001) approach. When using particle-based methods such as Particle Marginal Metropolis-Hastings (PMMH) Andrieu et al. (2010), the computational cost of trajectory inference scales as  $\mathcal{O}(P \cdot N)$ , where  $P$  is the number of particles and  $N$  again denotes the number of time points. This linear scaling with respect to  $N$  enables our proposed model to handle long trajectories more efficiently, making it well suited for characterization of high-resolution neural data. Moreover, the computational cost of inference is independent of the number of inducing points, since neither their number nor their values affect the SMC procedure, which will be discussed next. While generating inducing points incurs additional computational cost, their sampling rate is adapted to the underlying dynamics and scales as  $O(L)$ , where  $L$  is the maximum number of inducing points.

## 2.4 MODEL TRAINING AND INFERENCE

The training objective is to maximize the marginal likelihood (or evidence) of the observed data  $\{Z_k\}_{k=0}^K$ . This requires updating multiple sets of parameters, which include the event-value posterior distributions, and inferring the trajectories of the latent processes  $X_t$  and  $Y_t$  over  $t \in [0, T]$ . For this purpose, we use an Expectation–Maximization (EM) algorithm, which naturally accommodates latent processes (Brown and Kass, 2018). A variational inference alternative (Blei et al., 2017) is discussed in Appendix A.3.

To develop the EM algorithm, we first derive a discrete-time representation of the model. This representation is obtained using the renewal waiting-time process introduced in Section 2.1. Under this assumption, the non-Markovian dependence of  $X_t$  is removed, which is a critical modeling step for deriving recursive inference methods such as SMC. Intuitively, at time  $t$ , the next inducing point and its mark are already known, which eliminates dependence of the current  $X_t$  and  $Y_t$  on the future trajectory of  $X_t$ .

Within the discrete representation, the times of the inducing points (events) are defined in continuous space; thus, the trajectories of  $X_t$  and  $Y_t$  can be reconstructed at any temporal resolution. Appendix A.4 provides additional details on discrete SDE approximations. With the discrete representation of MIP-CSDE and the renewal process for the inducing points, the full likelihood of the model is given by

$$P(Z_{1:K}, X_{0:K}, Y_{0:K}, \tau_{1:n_u}, \vec{m}_{1:n_u}; \Omega, \Psi) = p(X_0, Y_0) \prod_{k=1}^K \left[ p(Z_k | Y_k, W, R) p(Y_k | X_{k-1}, Y_{k-1}, \{\sigma_y^d\}_{d=1}^D) \right. \\ \left. \times p(X_k | X_{k-1}, \tau_{1:n_s}, \{\sigma_x^d\}_{d=1}^D, \vec{m}_{1:n_s}) \right] \times \prod_{n=1}^{n_u} \left[ p(\tau_n | H_n) p(\vec{m}_n | \tau_n, H_n) \right] p_e p(\Psi) \quad (10.a)$$

$$n_s = \min\{n : \tau_n > k\} \quad (10.b)$$

The parameter set is  $\Omega = \{W, R, \{\sigma_y^d\}_{d=1}^D, \{\sigma_x^d\}_{d=1}^D, \{\mu_{0,x}^d, \sigma_{0,x}^d\}_{d=1}^D, \{\mu_{0,y}^d, \sigma_{0,y}^d\}_{d=1}^D\}$ , representing the SDE and observation-model parameters. The set  $\Psi$  contains the parameters describing the time-value distribution. We assume the waiting times  $\tau_i$  are independent of previous events and follow a Gamma distribution. The marks  $\vec{m}_i$  are assumed independent of previous events and of the current waiting time, and follow a multivariate distribution with zero mean and diagonal covariance. Thus,  $\Psi = \{\{\mu_m^d, \sigma_m^d, \alpha^d, \beta^d\}_{d=1}^D\}$ , where  $\mu_m^d, \sigma_m^d$  are the mark-distribution parameters and  $\alpha^d, \beta^d$  are the Gamma-distribution parameters.

The quantity  $n_u$  denotes the number of events in the interval  $[0, T]$ . The term  $p_e$  is the likelihood contribution for the non-event period after the final event  $n_u$  until time  $T$ , given by

$$p_e = \prod_{d=1}^D \exp\left(-\int_0^{T-t_{n_u}} f(s; \alpha^d, \beta^d) ds\right) \quad (11)$$

**Algorithm 1** SMC Algorithm for Inferring Inducing Points and State Estimation

---

```

270 Algorithm 1 SMC Algorithm for Inferring Inducing Points and State Estimation
271
272 1: Set Algorithm Hyperparameters:
273 2: Set number of particles  $U$ 
274 3: Define initial distributions  $p(x_0)$  and  $p(y_0)$ 
275 4: Define proposal density  $\pi_k(\cdot)$ 
276 5: Set hyperparameters  $\alpha_0, \beta_0$  for  $\tau$  distribution
277 6: Set hyperparameters  $\mu_0, \xi_0$  for  $m$  distribution
278 7: Initialization:
279 8: for  $u = 1$  to  $U$  do
280 9:   Sample  $x_0^u \sim p(x_0), y_0^u \sim p(y_0)$  ▷ Draw initial latent variables
281 10:   Set  $m_0^u = \vec{0}, \tau_0^u = 0, n_u = 0$  ▷ Initialize inducing-point memory
282 11:   Set initial weight  $w_k^u = \frac{1}{U}$  ▷ Uniform particle weights
283 12:   Initialize particle  $D_0^u = \{x_0^u, y_0^u, \tau_0^u, m_0^u, n_u\}$  ▷ Store complete particle state
284 13: end for
285 14: Inference:
286 15: for  $k = 1$  to  $K$  do
287 16:   1. Time & Mark Sampling:
288 17:   for  $u = 1$  to  $U$  do
289 18:     if  $k \cdot \Delta t > \tau_{\max}^{n_u}$  then
290 19:       Sample  $\tau_{\text{new}}^u \sim \Gamma(\tau; \alpha_0, \beta_0)$  ▷ Sample a new inducing-point time
291 20:       Sample  $m_{\text{new}}^u \sim \mathcal{N}(m; \mu_0, \xi_0)$  ▷ Sample corresponding inducing-point mark
292 21:       Update  $D_k^u = \{\dots\}$  ▷ Append inducing point to particle history
293 22:     end if
294 23:   end for
295 24:   2. Sampling:
296 25:   for  $u = 1$  to  $U$  do
297 26:     Sample  $(x_k^u, y_k^u) \sim \pi_k(\cdot)$  ▷ Draw new latent states from proposal
298 27:     Compute importance weight:
299
300 
$$w_k^u = w_{k-1}^u \cdot \frac{p(z_k | y_k^u) p(y_k^u | x_{k-1}^u) p(x_k^u | \tau_{0:n_u}^u, m_{0:n_u}^u)}{\pi_k(x_k^u, y_k^u | x_{0:k-1}^u, y_{0:k-1}^u, z_k, \tau_{0:n_u}^u, m_{0:n_u}^u)}$$

301 ▷ Correct for discrepancy between proposal and true posterior
302
303 28:   end for
304 29:   3. Normalization:
305 30:   for  $u = 1$  to  $U$  do
306
307 
$$\hat{w}_k^u = \frac{w_k^u}{\sum_{v=1}^U w_k^v}$$

308 ▷ Normalize weights to obtain posterior over particles
309
310 31:   end for
311 32:   4. Resampling:
312 33:   Resample  $U$  particles  $D_k^u$  with probabilities  $\hat{w}_k^u$  ▷ Mitigate particle degeneracy
313 34:   for  $u = 1$  to  $U$  do
314 35:     Reset weight:  $w_k^u = \frac{1}{U}$  ▷ Equalize weights after resampling
315 36:   end for
316 37: end for

```

---

where  $f(\cdot; \alpha^d, \beta^d)$  is the Gamma density function. Finally,  $p(\Psi)$  denotes the prior distribution over the model parameters.

With the likelihood function, the  $Q$ -function of the EM algorithm is defined by

$$\begin{aligned}
Q &= \mathbb{E}_{p(X_{0:K}, Y_{0:K}, \tau_{1:n_u}, \vec{m}_{1:n_u} | Z_{1:K}; \Omega^{(r)}, \Psi^{(r)})} [\log P(Z_{1:K}, X_{0:K}, Y_{0:K}, \tau_{1:n_u}, \vec{m}_{1:n_u}; \Omega, \Psi)] \\
&\approx \frac{1}{P} \sum_{p=1}^P \log P(Z_{1:K}, X_{0:K}^p, Y_{0:K}^p, \tau_{1:n_u}^p, \vec{m}_{1:n_u}^p; \Omega, \Psi) \quad (12)
\end{aligned}$$

Where  $p(X_{0:K}, Y_{0:K}, \tau_{1:n_u}, \vec{m}_{1:n_u} | Z_{1:K}; \Omega^{(r)}, \Psi^{(r)})$  is the posterior distribution of the latent variables given the current parameter estimates  $\Omega^{(r)}$  and  $\Psi^{(r)}$  at the  $r$ -th EM iteration.

The final line of Eq.12 provides an approximation of the  $Q$ -function using  $P$  sample trajectories of the latent process. Algorithm 1 describes the steps used to construct these trajectories. Each trajectory may terminate with a different number of inducing points; the model explicitly accounts for this. At each time index, the procedure determines whether an inducing point should be removed from the trajectory or whether a new inducing point should be added. Consequently, at every time index the proposed SMC method produces a projected trajectory for future time points. Under the renewal formulation, the process retains the Markov property required for the recursive update of the inducing points.

### 3 RESULTS

In this section, we evaluate our framework on both simulated and real datasets. We begin with a one-dimensional chirp time series and then reconstruct a Lorenz system trajectory embedded in a higher-dimensional observation space, comparing MIP-CSDE performance against recent and established methods for time series analysis and manifold inference. Finally, we apply the model to neural recordings: decoding rat hippocampal place cell activity during navigation on a W-shaped maze Joo and Frank (2018) and inferring low-dimensional manifolds from monkey M1 and PMd activity during a center out reach task (Pei et al., 2021; Churchland et al., 2012a). Together, these analyses illustrate the key characteristics and advantages of our method.

#### 3.1 CHIRP SIGNAL

We use MIP-CSDE to reconstruct a chirp signal whose frequency slowly decreases over time. Chirp signals provide a simple test of whether the model can place inducing points in ways that capture changing dynamics. We generated 500 noisy samples at 20 Hz, with frequency decreasing linearly from 0.2 to 0.1 Hz over 25 s. Both latent and observed processes were one-dimensional, and inducing point times were drawn from a Gamma prior that yields roughly one event every two seconds about 1.25 times the Nyquist rate for the highest chirp frequency. Inducing point values were drawn from a standard normal prior. We fit the model using our SMC-EM procedure with 1,000 particles over 12 iterations; smaller particle counts were also tested, and 1,000 provided stable fits across runs.

Figure 2 summarizes the results. The model accurately reconstructs the underlying trajectory and can generate new samples resembling chirp signals. Early in the sequence, the model increases inducing point magnitudes to track stronger oscillations; later, as the frequency slows, it uses more inducing points with smaller magnitudes. Unlike classical sampling approaches which typically reduce sample rates when frequency content decreases MIP-CSDE adapts through coordinated changes in both event timing and point values. These behaviors highlight the flexibility of the model and how its priors shape the resulting representation.

#### 3.2 LORENZ DYNAMIC RECONSTRUCTION

We applied our framework to simulated data generated by the Lorenz attractor system (Lorenz, 1963), a standard test for models that aim to recover complex, nonlinear dynamics from noisy measurements. The true three-dimensional Lorenz trajectory was projected into a ten-dimensional observation space and sampled at 10 Hz for 500 time points, with added correlated Gaussian noise. Both latent processes,  $X_t$  and  $Y_t$ , were set to three dimensions. We chose noise levels so that  $X_t$  captured the rough dynamics while  $Y_t$  remained smooth. During training, the model jointly learned the projection matrix  $W$ , the latent trajectory, and the timing and values of the inducing points.

Figure 1B shows the observed data and the inferred manifold. We used 20,000 particles and 25 EM iterations, which was sufficient for the estimates and likelihood to stabilize. The recovered trajectory closely follows the true Lorenz attractor, capturing both its overall shape and the finer curvature within each lobe. The inducing points adapt to these transitions: near the entrances and exits of each lobe, events become more frequent and their mark values change direction to anticipate the turn, whereas within a lobe the events become sparser and the marks remain smaller. This adaptive behavior is a core strength of MIP-CSDE it does not rely on uniform sampling or fixed basis functions, but instead adjusts where and how inducing events are placed, allowing it to reconstruct a highly nonlinear attractor from noisy, mixed observations.

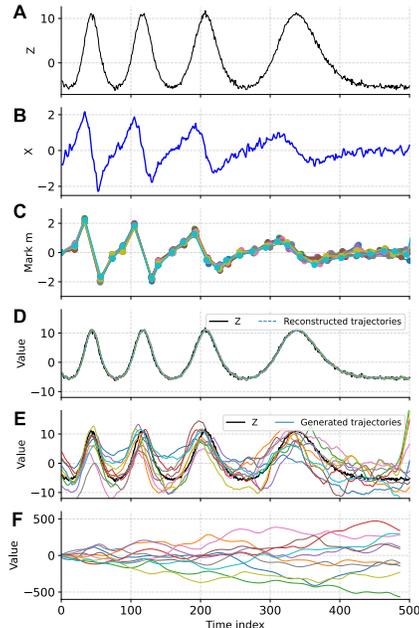


Figure 2: **Chirp Signal Reconstruction:** (A) Simulated noisy chirp signal. (B) Posterior mean of the latent process. (C) Inferred inducing events. (D) Reconstructed trajectory closely follows the observed data. (E-F) Samples generated from the trained and initial models. MIP-CSDE adapts event timing and values to track time varying frequency structure.

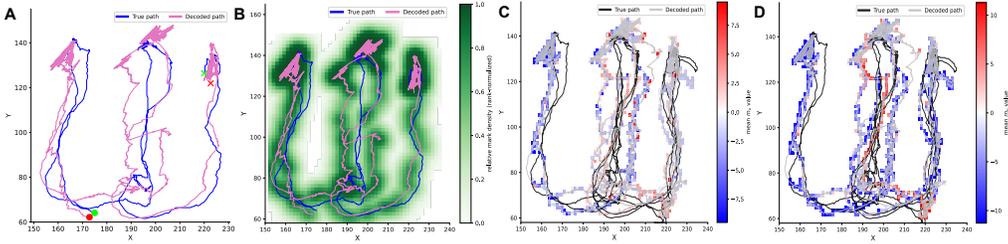


Figure 3: MIP-CSDE place cell decoder on rat hippocampus data. (A) True and decoded position for a left to right maze traversal. (B-D) Inducing events and their spatial marks concentrate near turns and rapid transitions, remaining sparse along straight segments. MIP-CSDE provides accurate and interpretable trajectory decoding from place cell spiking.

Table 1: Performance of MIP-CSDE compared to existing state of the art decoders. Root mean squared error (RMSE) quantifies the average distance between the decoded and true positions, and the 95% highest posterior density (HPD) coverage reports the fraction of time the true position lies inside the model’s 95% credible region.

Method	MIP-CSDE	GMM	Gaussian Approx	4D GMM	Exact Point Process
RMSE, HPD%	6.1, 92	17.2, 86	27.3, 62	15.7, 88	14.0, 91

To test whether the model recovers the correct dimensionality, we varied the latent dimension  $d \in \{1, 2, 3, 4\}$  (Appendix Figure 7). Models with one or two dimensions collapsed the dynamics, and four dimensions introduced unnecessary freedom without improving accuracy. The three-dimensional model matching the true system gave the best reconstruction and prediction, confirming that MIP-CSDE identifies an appropriate latent dimensionality for this dataset.

### 3.3 RAT HIPPOCAMPUS: DECODING SPATIAL TRAJECTORIES FROM CA1 SPIKING

We applied MIP-CSDE to hippocampal CA1 place cell recordings from a rat running on a W-shaped maze Joo and Frank (2018). Decoding position from place cell activity is a standard benchmark, with prior work using Gaussian mixture models and deep neural network based decoders to improve robustness and accuracy Yousefi et al. (2019); Karlsson and Frank (2008); Brown et al. (1998). The dataset contains spiking activity from 62 cells sampled in 33 ms bins. We modeled the latents  $X_k$  and  $Y_k$  as two-dimensional, where  $X_k$  captures movement related dynamics and  $Y_k$  represents position. Each cells activity in bin  $k$  is  $Z_k = \{z_k^i\}_{i=1}^{62}$ , where  $z_k^i \in \{0, 1\}$ . Firing rates  $\lambda_i$  were estimated using standard kernel methods Yousefi et al. (2019), and the likelihood for each cell followed a Poisson point process model:

$$p(Y_k | z_k^i) \propto p(z_k^i | \lambda_{i,k}) = (\Delta k \lambda_{i,k})^{z_k^i} \exp(-\lambda_{i,k} \Delta k) \tag{13}$$

with  $\Delta k = 33$  ms. The full likelihood is the product across cells. We used the first 80% of the trajectory to fit rate models and decoded position on the remaining 20%, following common practice Yousefi et al. (2019). Decoding used a particle filter with 10,000 particles, which is consistent with earlier hippocampal decoders Yousefi et al. (2019); Brown et al. (1998) and provided stable posterior trajectories. For the inducing point process, we used a Gamma(8, 2) prior for waiting times (mean 4 s) and a Gaussian prior  $\mathcal{N}([0, 0]^T, 0.5I_{2 \times 2})$  for mark values. These choices mildly favor several second gaps between events and prevent unrealistically large mark jumps, while remaining flexible; moderate changes to the hyperparameters did not affect the results.

Figure 3 shows the decoded trajectory and inducing points. The posterior mean path closely follows the rat’s movement along the maze arms (panel A). Inducing events cluster near corners and the central corridor (panel B), where the trajectory changes direction more rapidly, and remain sparse along straight segments. Additional diagnostics in Appendix Figure 8 show stable latent dynamics across the full session. Table 1 compares MIP-CSDE with existing decoders: it achieves lower RMSE and comparable or higher 95% HPD coverage. With a runtime of about 2 ms per 33 ms bin, MIP-CSDE provides an accurate, temporally aligned, and interpretable decoder for hippocampal place cell activity.

### 3.4 MANIFOLD DIFFERENTIATION DURING MONKEY REACHING TASK

We applied MIP-CSDE to the Neural Latents Benchmark (NLB) MC\_Maze dataset (Pei et al., 2021), which contains high resolution recordings from macaque PMd and M1 during a center out reaching task (Churchland et al., 2012a). The dataset includes spiking activity from 182 neurons sampled at 1 ms, along with hand position, velocity, and cursor information. Prior work shows that population activity at movement onset strongly predicts the upcoming reach (Churchland et al., 2012a; Pei et al., 2021).

Our goal was to learn a low-dimensional representation of neural activity during preparation and movement, and to characterize how these phases relate. Building on studies suggesting that preparatory activity sets the initial conditions for movement on a shared manifold (Churchland et al., 2012a; Kaufman et al., 2014; Elsayed et al., 2016), we tested whether preparatory and reach trajectories are linked by a simple linear transformation. For interpretability, we restricted MIP-CSDE to a two-dimensional latent space (though the model can accommodate higher dimensions). Each neurons spiking was modeled with a Bernoulli GLM, where the log odds of spiking at time  $k$  is a linear function of the latent state, allowing the latents to be interpreted as low-dimensional predictors of firing probability.

Figure 4 shows the inferred 2D manifolds for two example trials. Within each trial, preparatory and reach activity occupy overlapping regions of the same latent space and form smooth, connected trajectories, while across trials the paths diverge according to the reach target. To quantify the link between preparation and movement, we aligned the two trajectories and fit a single affine map  $Y(t) \approx AX(t) + b$ . This  $2 \times 2$  transform explained about 90% of the variance in the reach phase latents ( $R^2 \approx 0.90$ ), showing that movement activity is well approximated as a linear transformation of preparatory activity within the same 2D manifold. This finding aligns with prior MC\_Maze results and supports the view that preparatory dynamics organize the neural state, which is then smoothly transformed into the reach trajectory.

### 3.5 COMPARATIVE ANALYSIS ON SIMULATED DATASETS

To evaluate the performance of MIP-CSDE, we benchmark it against several continuous time baselines in terms of both predictive accuracy and computational efficiency. The baselines include models of increasing flexibility: a Linear SDE, a Gaussian Process SDE (GP-SDE) (Duncker et al., 2019), and the Gaussian Process Switching Linear Dynamical System (GP-SLDS) (Hu et al., 2024). To ensure a robust comparison, all baseline models were fitted using the highly efficient SING inference framework (Hu et al., 2025).

We report performance using the MSE in the observation space. This choice of metric is necessitated by a key property of the baseline models, where their latent spaces are only identified up to an arbitrary affine transformation. This implies that a direct comparison between their inferred latent trajectories and the ground truth is not meaningful without a post hoc alignment procedure. The reconstructed latent along with realigned version is presented in appendix A.6. The quantitative results

Table 2: Performance comparison with baseline continuous time models. For simulated data (Chirp, Lorenz), we report MSE. (Results are reported as mean  $\pm$  standard error across 5 trials.)

Dataset	MIP-CSDE	Linear SDE	GP-SDE	GP-SLDS
Chirp Signal	0.30 $\pm$ 0.02	0.48 $\pm$ 0.05	0.39 $\pm$ 0.06	0.36 $\pm$ 0.02
Lorenz System	0.18 $\pm$ 0.01	0.28 $\pm$ 0.04	0.25 $\pm$ 0.03	0.21 $\pm$ 0.03

in Table 2 indicate that MIP-CSDE attains the lowest error on the Lorenz system (MSE  $0.18 \pm 0.01$ ), while performing comparably to GP-SDE on the chirp signal (MSE  $0.30 \pm 0.02$  vs.  $0.29 \pm 0.06$ ). On the Lorenz benchmark, our implementation required approximately 290 s, compared with 12 s for Linear SDE, 38 s for GP-SDE, and 252 s for GP-SLDS. We also observed higher memory usage for GP-SLDS on longer sequences, which is consistent with the scaling behavior of GP-based kernels. For MIP-CSDE, runtime scales approximately linearly with the number of particles increasing particles improves accuracy at additional computational cost, reflecting an explicit accuracy compute trade off. Although implementation details and hardware choices affect absolute timings, the ob-

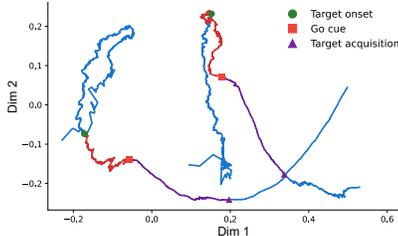


Figure 4: **Inferred neural manifolds during preparation and reach.** Two-dimensional latent trajectories from two MC\_Maze trials. For each trial, preparatory and reach activity lie in the same 2D space and form smooth, connected paths, while trajectories differ across reach targets. A single affine map fit between preparation and movement explains about 90% of the variance in the reach phase, showing that movement activity largely reflects a linear transformation of preparatory activity within this shared manifold.

486 served trends are consistent with the computational analysis presented in Section 2.3.2. As expected,  
487 measured runtimes scale approximately linearly with both particle count and sequence length (Ap-  
488 pendix A.7). Overall, these results support MIP-CSDE as a more accurate and computationally  
489 competitive approach for recovering latent dynamics and manifolds in continuous time.

## 490 4 DISCUSSION

491  
492 We developed a training and inference pipeline for a discretized version of our framework and eval-  
493 uated it on both simulated and neural datasets. Across these examples, MIP-CSDE consistently  
494 adapted its inducing events to capture the structure of the underlying dynamics. In the chirp signal,  
495 the model allocated events in response to changes in frequency; in the Lorenz system, it recovered  
496 the nonlinear geometry of a chaotic attractor and scaled effectively as dimensionality increased. In  
497 the rat hippocampus task, the model served as a robust decoder of spatial position, and in the mon-  
498 key reaching task it revealed a shared manifold linking preparatory and movement activity through a  
499 simple linear transformation. Together, these findings show that the model is flexible enough to han-  
500 dle diverse forms of temporal structure and strong nonlinearities while remaining computationally  
501 efficient. We also compared MIP-CSDE to continuous time baselines including Linear SDE, GP-  
502 SDE, and GP-SLDS. On simulated datasets, MIP-CSDE matched or exceeded their reconstruction  
503 accuracy while retaining linear scaling in data length. These comparisons highlight the advantages  
504 of combining adaptive inducing events with a cascaded SDE formulation: the model avoids the cu-  
505 bic scaling of GP based approaches while still achieving competitive performance, even on systems  
506 with complex latent geometry. The simulated results further support the universal approximation  
507 properties of the framework, as the model accurately reconstructed both smooth and chaotic dynam-  
508 ics.

509 While models such as DKF and GP based methods (Krishnan et al., 2015b; Casale et al., 2018) can  
510 also infer latent processes, they often depend on large datasets, specific kernel choices, or more com-  
511 plex inference procedures. MIP-CSDE requires no pre-defined kernel and offers two complementary  
512 modes of interpretation: one can examine the latent trajectories directly or study the inducing point  
513 timings and values that drive them. This dual interpretability is valuable in neuroscience, where link-  
514 ing latent dynamics to behavior or external variables is often a key goal (Churchland et al., 2012b).

515 Although we used a discretized version of the model, a continuous formulation similar to neural  
516 ODEs (Chen et al., 2018) could further improve inference stability and reduce numerical error when  
517 generating trajectories. Several parameters, such as the noise variance in  $X_t$  and the decay term  
518 in  $Y_t$ , influence the trade off between smoothness and flexibility, and learning them automatically  
519 rather than setting them manually is a natural next step. Our inference method is based on a custom  
520 SMC algorithm that handles both Gaussian and point process observations. In settings where the  
521 latent to observation mapping is linear and Gaussian, the model reduces to a standard linear state  
522 space system, suggesting a hybrid approach in which SMC infers inducing events while the Kalman  
523 filter updates states, offering additional computational gains.

524 Finally, we assumed independence between waiting times and marks, which simplifies inference  
525 but may limit adaptability across temporal scales. More expressive choices such as time dependent  
526 or history dependent distributions, or joint modeling of marks and waiting times may allow the  
527 model to better handle rapid transitions. The model also requires choosing a latent dimensionality.  
528 Although we examined several choices in the Lorenz example, a principled selection method such  
529 as automatic relevance determination (ARD) Wipf and Nagarajan (2007) may enable the model to  
530 infer the effective latent dimension directly from data.

## 531 5 CONCLUSION

532  
533 Here, we introduce a cascade SDE framework to infer the underlying latent structure and manifold  
534 present in high-dimensional time series using a sparse set of inducing points, adaptively placed in  
535 both time and value space. The model achieves a high level of expressive power, while its computa-  
536 tional cost grows only linearly in both data dimensions and time. The comparative analyses indicate  
537 that it achieves performance accuracy on par with or superior to state-of-the-art models. These re-  
538 sults suggest that the model holds promise as an unsupervised dimensionality reduction tool and can  
539 be robustly applied as a dynamical neural decoder or adaptive feature extractor across a range of  
neuroscience applications and time-series analysis.

## REFERENCES

- 540  
541  
542 C. Andrieu, A. Doucet, and R. Holenstein. Particle markov chain monte carlo methods. *Journal of*  
543 *the Royal Statistical Society: Series B (Statistical Methodology)*, 72(269):269–342, 2010.
- 544  
545 D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians.  
546 *Journal of the American statistical Association*, 112(518):859–877, 2017.
- 547  
548 G. Bondanelli, T. Deneux, B. Bathellier, and S. Ostojic. Network dynamics underlying off responses  
549 in the auditory cortex. *Elife*, 10:e53151, 2021.
- 550  
551 E. N. Brown and R. E. Kass. Estimating a state-space model from point process observations.  
552 *Unpublished Manuscript*, 2018.
- 553  
554 E. N. Brown, L. M. Frank, D. Tang, M. C. Quirk, and M. A. Wilson. A statistical paradigm for  
555 neural spike train decoding applied to position prediction from ensemble firing patterns of rat  
556 hippocampal place cells. *Journal of Neuroscience*, 18(18):7411–7425, 1998.
- 557  
558 F. P. Casale, A. V. Dalca, L. Saglietti, J. Listgarten, and N. Fusi. Gaussian process prior variational  
559 autoencoders. 31, 2018.
- 560  
561 H. Chang, Q. Zhang, Y. Wang, Z. Qin, L. Zhao, and H. Wang. Unlocking the power of lstm for long  
562 term time series forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38  
563 (4):4292–4300, 2024.
- 564  
565 R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud. Neural ordinary differential equations.  
566 31, 2018.
- 567  
568 M. M. Churchland, J. P. Cunningham, M. T. Kaufman, J. D. Foster, P. Nuyujukian, S. I. Ryu, and  
569 K. V. Shenoy. Neural population dynamics during reaching. *Nature*, 2012a.
- 570  
571 M. M. Churchland, J. P. Cunningham, M. T. Kaufman, S. I. Ryu, and K. V. Shenoy. Neural popula-  
572 tion dynamics during reaching. *Nature*, 487(7405):51–56, 2012b.
- 573  
574 J. P. Cunningham and B. M. Yu. Dimensionality reduction for large-scale neural recordings. *Nature*  
575 *neuroscience*, 17(11):1500–1509, 2014a.
- 576  
577 J. P. Cunningham and B. M. Yu. Dimensionality reduction for large-scale neural recordings. *Nature*  
578 *neuroscience*, 17(11):1500–1509, 2014b.
- 579  
580 D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes: volume I: elemen-*  
581 *tary theory and methods*. Springer Science & Business Media, 2006.
- 582  
583 A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer,  
584 New York, 2001.
- 585  
586 L. Duncker, G. Bohnert, J. Boussard, and M. Sahani. Learning interpretable continuous-time models  
587 of latent stochastic dynamical systems. In *International conference on machine learning*, pages  
588 1726–1734. PMLR, 2019.
- 589  
590 S. Eleftheriadis, T. Nicholson, M. Deisenroth, and J. Hensman. Identification of gaussian process  
591 state space models. *Advances in neural information processing systems*, 30, 2017.
- 592  
593 G. F. Elsayed, A. H. Lara, M. T. Kaufman, M. M. Churchland, and J. P. Cunningham. Reorganization  
between preparatory and movement population responses in motor cortex. *Nature Communica-*  
*tions*, 7:13239, 2016. doi: 10.1038/ncomms13239.
- E. Fox, E. Sudderth, M. Jordan, and A. Willsky. Nonparametric bayesian learning of switching linear  
dynamical systems. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in*  
*Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008.
- Z. Ghahramani and G. E. Hinton. Variational learning for switching state-space models. *Neural*  
*Computation*, 12(4):831–864, 2000. doi: 10.1162/089976600300015619.

- 594 L. Girin, S. Leglaive, X. Bie, J. Diard, T. Hueber, and X. Alameda-Pineda. Dynamical variational  
595 autoencoders: A comprehensive review. *arXiv preprint arXiv:2008.12595*, 2020.  
596
- 597 X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks.  
598 In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*  
599 (*AISTATS*), pages 249–256. PMLR, 2010.
- 600 J. Gonzalez-Castillo, I. S. Fernandez, K. C. Lam, D. A. Handwerker, F. Pereira, and P. A. Ban-  
601 dettini. Manifold learning for fmri time-varying functional connectivity. *Frontiers in Human*  
602 *Neuroscience*, 17:1134012, 2023.  
603
- 604 A. Gosztolai, R. L. Peach, A. Arnaudon, M. Barahona, and P. Vandergheynst. Interpretable statistical  
605 representations of neural population dynamics and geometry. *arXiv preprint*, 2023.  
606
- 607 A. Hu, D. Zoltowski, A. Nair, D. Anderson, L. Duncker, and S. Linderman. Modeling latent neu-  
608 ral dynamics with gaussian process switching linear dynamical systems. *Advances in Neural*  
609 *Information Processing Systems*, 37:33805–33835, 2024.
- 610 A. Hu, H. Smith, and S. Linderman. Sing: Sde inference via natural gradients. *arXiv preprint*  
611 *arXiv:2506.17796*, 2025.  
612
- 613 M. Jacobsen. *Point Process Theory and Applications: Marked Point and Piecewise Deterministic*  
614 *Processes*. Birkhäuser, Boston, 2006.
- 615 H. R. Joo and L. M. Frank. The hippocampal sharp wave–ripple in memory retrieval for immediate  
616 use and consolidation. *Nature Reviews Neuroscience*, 19(12):744–757, 2018.  
617
- 618 M. P. Karlsson and L. M. Frank. Network dynamics underlying the formation of sparse, informative  
619 representations in the hippocampus. *Journal of Neuroscience*, 28(52):14271–14281, 2008. doi:  
620 10.1523/JNEUROSCI.4261-08.2008.
- 621 M. T. Kaufman, M. M. Churchland, S. I. Ryu, and K. V. Shenoy. Cortical activity in the null  
622 space: permitting preparation without movement. *Nature Neuroscience*, 17(3):440–448, 2014.  
623 doi: 10.1038/nn.3643.  
624
- 625 G. Kitagawa and W. Gersch. Linear gaussian state space modeling. In *Smoothness priors analysis*  
626 *of time series*, pages 55–65. Springer, 1996.  
627
- 628 R. G. Krishnan, U. Shalit, and D. Sontag. Deep kalman filters, 2015a.
- 629 R. G. Krishnan, U. Shalit, and D. Sontag. Deep kalman filters. *arXiv preprint arXiv:1511.05121*,  
630 2015b.  
631
- 632 E. N. Lorenz. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2):130–141,  
633 1963.  
634
- 635 L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for  
636 dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- 637 B. Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science  
638 & Business Media, 2013.  
639
- 640 F. Pei, J. Ye, D. M. Zoltowski, A. Wu, R. H. Chowdhury, H. Sohn, J. E. ODoherty, K. V. Shenoy,  
641 M. T. Kaufman, M. Churchland, M. Jazayeri, L. E. Miller, J. Pillow, I. M. Park, E. L. Dyer,  
642 and C. Pandarinath. Neural latents benchmark 21: Evaluating latent variable models of neural  
643 population activity. In *Advances in Neural Information Processing Systems.*, 2021.
- 644 J. Pitman and M. Yor. Brownian bridge and related stochastic processes. *Probability Surveys*, 1:  
645 1–61, 1999.  
646
- 647 L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition.  
*Proceedings of the IEEE*, 77(2):257–286, 2002.

- 648 S. S. Rangapuram, M. W. Seeger, J. Gasthaus, L. Stella, Y. Wang, and T. Januschowski. Deep state  
649 space models for time series forecasting. *Advances in neural information processing systems*, 31,  
650 2018.
- 651 Y. Rubanova, R. T. Chen, and D. K. Duvenaud. Latent ordinary differential equations for irregularly-  
652 sampled time series. *Advances in neural information processing systems*, 32, 2019.
- 653 S. Särkkä and L. Svensson. *Bayesian filtering and smoothing*, volume 17. Cambridge university  
654 press, 2023.
- 655 M. Seeger. Gaussian processes for machine learning. *International journal of neural systems*, 14  
656 (02):69–106, 2004.
- 657 C. E. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21,  
658 1949.
- 659 D. Sussillo, R. Jozefowicz, L. Abbott, and C. Pandarinath. Lfads-latent factor analysis via dynamical  
660 systems. *arXiv preprint arXiv:1608.06315*, 2016.
- 661 G. E. Uhlenbeck and L. S. Ornstein. On the theory of the brownian motion. *Physical Review*, 36(5):  
662 823–841, 1930.
- 663 L. Van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning  
664 research*, 9(11), 2008.
- 665 J. Wang, A. Hertzmann, and D. J. Fleet. Gaussian process dynamical models. *Advances in neural  
666 information processing systems*, 18, 2005.
- 667 N. Whiteley, A. Gray, and P. Rubin-Delanchy. Statistical exploration of the manifold hypothesis,  
668 2024.
- 669 D. Wipf and S. Nagarajan. A new view of automatic relevance determination. *Advances in Neural  
670 Information Processing Systems (NeurIPS)*, 20, 2007.
- 671 A. Yousefi, A. K. Gillespie, J. A. Guidera, M. Karlsson, L. M. Frank, and U. T. Eden. Efficient  
672 decoding of multi-dimensional signals from population spiking activity using a gaussian mixture  
673 particle filter. *IEEE Transactions on Biomedical Engineering*, 66(12):3486–3498, 2019.
- 674 B. M. Yu, J. P. Cunningham, G. Santhanam, S. Ryu, K. V. Shenoy, and M. Sahani. Gaussian-process  
675 factor analysis for low-dimensional single-trial analysis of neural population activity. *Advances  
676 in neural information processing systems*, 21, 2008.
- 677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701

## A APPENDIX

### A.1 NONPARAMETRIC AND NON-MARKOVIAN PROPERTIES

In our model, the  $X_t$  exhibits dynamics similar to that of samples from a GP with specific covariance structures, for example, an Ornstein-Uhlenbeck process corresponding to an exponential kernel (Uhlenbeck and Ornstein, 1930). Within the model, the number of inducing points is not fixed in advance and it adapts dynamically to the complexity of the observed dynamics. The nonparametric and GP-like nature of our model makes it an alternative choice for machine learning and probabilistic applications. In our proposed model, the trajectory of  $X_t$  process is dependent to both past and future event-value pairs; thus it does not satisfy the Markovian property. This might complicate both training and inference within our framework. In section 2.2, we reformulated the inducing point distribution using a renewal marked point process, which lets us to define  $X_t$  and  $Y_t$  process with a Markovian property. In section 2.4, we leverage this reformulation of SDEs for the inference and training of the model.

### A.2 CONVERGENCE OF SPACING BETWEEN ADJACENT SAMPLES

Let  $X_1, X_2, \dots, X_N$  be i.i.d. random variables uniformly distributed on  $[0, T]$ , and let  $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(N)}$  denote their order statistics. Define the adjacent spacings  $d_i = X_{(i+1)} - X_{(i)}$  for  $i = 1, \dots, N - 1$ .

From properties of uniform order statistics, the expected value of the  $i$ -th order statistic is  $\mathbb{E}[X_{(i)}] = \frac{iT}{N+1}$ . It follows that

$$\mathbb{E}[d_i] = \mathbb{E}[X_{(i+1)} - X_{(i)}] = \frac{T}{N+1} \quad (14)$$

which satisfies  $\mathbb{E}[d_i] \rightarrow \frac{T}{N}$  as  $N \rightarrow \infty$ .

Moreover, the variance of  $d_i$  satisfies  $\text{Var}(d_i) = O(N^{-2})$ . By Chebyshevs inequality,

$$\mathbb{P}(|d_i - \mathbb{E}[d_i]| \geq \epsilon) \leq \frac{\text{Var}(d_i)}{\epsilon^2} = O\left(\frac{1}{N^2}\right) \quad (15)$$

which vanishes as  $N \rightarrow \infty$ . Therefore,  $d_i \rightarrow \frac{T}{N}$  in probability.

Furthermore, classical results on uniform spacings imply that the maximum spacing  $\max_i d_i$  satisfies

$$\max_{1 \leq i \leq N-1} d_i = \frac{T}{N} + O\left(N^{-1/2}\right) \quad (16)$$

with high probability, confirming that all gaps become uniformly close to  $\frac{T}{N}$  as  $N \rightarrow \infty$ .

### A.3 VARIATIONAL INFERENCE FOR MODEL TRAINING

To complement the EM approach, we also develop a variational inference (VI) algorithm for training our model. As in Section 2.2, we begin with the discrete-time representation of the renewal process, which breaks the non-Markovian dependence of  $X_t$  by ensuring that the next inducing point and its mark are known at any time  $t$ . This property makes recursive inference feasible, but instead of relying on exact latent sampling as in the E-step of EM, we approximate the intractable posterior using a variational family.

Specifically, we introduce the following structured mean-field approximation:

$$q_\phi(X_{0:K}, Y_{0:K}, \{\tau_i, m_i\}_{i=1}^L) = q_X(X_{0:K}; \phi_X) q_Y(Y_{0:K}; \phi_Y) q_{\tau, m}(\{\tau_i, m_i\}_{i=1}^L; \phi_{\tau, m}), \quad (17)$$

where  $\phi = \{\phi_X, \phi_Y, \phi_{\tau, m}\}$  are variational parameters. This factorization decouples states and inducing points while retaining the renewal structure. In practice, we amortize these distributions using neural networks that map observed data into variational parameters.

The training objective is the evidence lower bound (ELBO):

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{q_\phi} \left[ \log p_\theta(Z_{0:K}, X_{0:K}, Y_{0:K}, \{\tau_i, m_i\}_{i=1}^L) - \log q_\phi(X_{0:K}, Y_{0:K}, \{\tau_i, m_i\}_{i=1}^L) \right], \quad (18)$$

where  $\theta$  denotes the generative model parameters. Maximizing  $\mathcal{L}(\theta, \phi)$  yields both approximate posterior inference (via  $q_\phi$ ) and maximum likelihood estimation of  $\theta$ .

We employ stochastic gradient variational Bayes (SGVB) with the reparameterization trick to obtain low-variance gradient estimates. In this formulation, sampling of event–value pairs is embedded directly into the variational distribution  $q_{\tau, m}$ , which is parameterized by waiting-time and mark distributions. These distributions can be chosen flexibly, e.g., Gamma and Gaussian, or replaced with neural flows for greater expressivity.

---

#### Algorithm 2 Variational Inference for Inducing Point and State Estimation

---

- 1: **Initialize:** model parameters  $\theta$ , variational parameters  $\phi$
- 2: **for** each training iteration **do**
- 3:   Sample latent variables from  $q_\phi$ :

$$X_{0:K}, Y_{0:K}, \{\tau_i, m_i\}_{i=1}^L \sim q_\phi$$

- 4:   Compute stochastic ELBO estimate:

$$\hat{\mathcal{L}} = \log p_\theta(Z_{0:K}, X_{0:K}, Y_{0:K}, \{\tau_i, m_i\}_{i=1}^L) - \log q_\phi(X_{0:K}, Y_{0:K}, \{\tau_i, m_i\}_{i=1}^L)$$

- 5:   Update  $(\theta, \phi)$  via gradient ascent on  $\hat{\mathcal{L}}$
  - 6: **end for**
- 

Unlike EM, where process noise parameters are typically fixed, VI allows them to be included in the variational family and learned directly. In practice, however, we sometimes constrain these parameters to preserve stability of the underlying SDEs.

Finally, while the description above applies to a single observed trajectory, the variational framework naturally extends to multiple trials. Each trial maintains its own approximate posterior over inducing points and latent trajectories, while global parameters  $\theta$  are shared across trials. This amortized formulation enables scalable training across large experimental datasets.

#### A.4 DISCRETE REPRESENTATION OF HIERARCHICAL SDE

We focus on a discrete-time formulation of the model. To construct the discrete process, we assume that  $X_t$  and  $Y_t$  are sampled at regular intervals of  $\Delta t$ . The discrete representation of  $x_t^d$  is defined as:

$$x_{k+1}^d = x_k^d + \frac{m_{i+1}^d - x_k^d}{t_{i+1} - k\Delta t} \Delta t + \sqrt{\frac{(t_{i+1} - k\Delta t)(k\Delta t - t_i)}{t_{i+1} - t_i}} \Delta t \cdot w_k^d, \quad (19)$$

$$w_k^d \sim \mathcal{N}(0, \sigma_x^{d^2})$$

where  $w_k^d$  is a Gaussian noise term. Similarly, the discrete-time evolution of  $y_t^d$  is:

$$y_{k+1}^d = y_k^d + x_k^d \Delta t + \sqrt{\Delta t} \cdot \nu_k^d, \quad \nu_k^d \sim \mathcal{N}(0, \sigma_y^{d^2}) \quad (20)$$

where  $\nu_k^d$  is Gaussian noise. The discrete observation process is given by:

$$Z_k = WY_k + \xi_k, \quad \xi_k \sim \mathcal{N}(0, R) \quad (21)$$

810 where  $Y_k = \begin{pmatrix} y_k^1 \\ \vdots \\ y_k^D \end{pmatrix}$ ,  $W \in \mathbb{R}^{M \times D}$  is a projection matrix, and  $\xi_k$  is observation noise.  
811  
812  
813

814 To ensure accuracy,  $\Delta t$  must be much smaller than the minimum inter-event time, i.e.,  $\Delta t \ll$   
815  $\min(\tau_i)$ , so that no two inducing points fall within the same discrete time bin. This constraint can  
816 be satisfied by analyzing the posterior distribution of waiting times and adjusting  $\Delta t$  accordingly.  
817 In essence, we require an orderly event process—allowing at most one event per bin—which can be  
818 enforced by carefully selecting the bin size  $\Delta t$ .

819 When using the waiting time representation, at time  $k$ , we already know when the next inducing  
820 point (e.g., event 213) will occur and what its value will be. From a Markovian perspective, we  
821 can assume that the entire process is determined at time  $k$ . This assumption simplifies the inference  
822 procedure presented in Algorithm 1.  
823

#### 824 A.5 GAMMA DISTRIBUTION FOR WAITING TIMES AND PRIOR SELECTION FOR INDUCING 825 POINTS

826 To complete the Bayesian framework, we define priors for the model parameters. For the mark  
827 distribution parameters, we assume:  
828

$$829 \mu_i \mid \Sigma_i \sim N(\mu_0, \xi \Sigma_i), \quad \Sigma_i \sim \text{Inverse-Wishart}(\nu, \Psi) \quad (22)$$

831 where  $\mu_0$ ,  $\xi$ ,  $\nu$ , and  $\Psi$  are hyperparameters. For the Gamma distribution parameters  $\alpha$  and  $\xi$  govern-  
832 ing the waiting times  $\tau_i$ , we consider the following prior options based on domain-specific knowl-  
833 edge, though their specific forms remain to be fully specified in this study:  
834

- 835 • For  $\alpha$ :

$$836 - \alpha \sim \text{Gamma}(a_0, b_0) = \frac{b_0^{a_0}}{\Gamma(a_0)} \alpha^{a_0-1} e^{-b_0 \alpha},$$

$$837 - \alpha \sim \text{Exp}(\xi_0) = \xi_0 e^{-\xi_0 \alpha},$$

$$838 - \alpha \sim \text{Lognormal}(\mu_0, \sigma_0^2) = \frac{1}{\alpha \sqrt{2\pi\sigma_0^2}} \exp\left(-\frac{(\log \alpha - \mu_0)^2}{2\sigma_0^2}\right),$$

- 841 • For  $\xi$ :

$$842 - \xi \sim \text{Gamma}(c_0, d_0) = \frac{d_0^{c_0}}{\Gamma(c_0)} \xi^{c_0-1} e^{-d_0 \xi},$$

$$843 - \xi \sim \text{InvGamma}(\gamma_0, \delta_0) = \frac{\delta_0^{\gamma_0}}{\Gamma(\gamma_0)} \xi^{-(\gamma_0+1)} e^{-\delta_0/\xi},$$

845 where  $a_0, b_0, \xi_0, \mu_0, \sigma_0^2, c_0, d_0, \gamma_0, \delta_0$  are hyperparameters.  
846

847 In our model, the waiting times  $\tau_i$  and the marks  $\vec{m}_i$  associated with each event are generated  
848 according to specific probabilistic distributions:  
849

- 850 • **Waiting Times  $\tau_i$ :**

851 The waiting times between events are assumed to follow a Gamma distribution parameter-  
852 ized by a shape parameter  $\alpha$  and a rate parameter  $\lambda$ . The probability density function for  $\tau_i$   
853 is given by:

$$854 p(\tau_i) = \text{Gamma}(\tau_i; \alpha, \beta) \quad (23)$$

855 where the Gamma distribution is defined as:

$$856 \text{Gamma}(\tau_i; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \tau_i^{\alpha-1} e^{-\beta \tau_i}, \quad \tau_i > 0 \quad (24)$$

858 and  $\Gamma(\alpha)$  denotes the Gamma function evaluated at  $\alpha$ .  
859

#### 860 **Motivation for using the Gamma distribution:**

861 Consider  $N$  i.i.d. samples  $U_1, \dots, U_N \sim \text{Uniform}(0, T)$ , and denote their order statistics  
862 by  $U_{(1)} \leq \dots \leq U_{(N)}$ . Define the gaps between consecutive order statistics as  
863

$$\Delta_0 = U_{(1)}, \quad \Delta_i = U_{(i+1)} - U_{(i)} \text{ for } i = 1, \dots, N-1, \quad \Delta_N = T - U_{(N)} \quad (14)$$

As  $N \rightarrow \infty$ , it is well-known that each gap satisfies  $\Delta_i \xrightarrow{d} T/N$ , and the rescaled gaps  $N\Delta_i$  converge in distribution to an exponential random variable, that is,

$$N\Delta_i \xrightarrow{d} \text{Exp}(1) \quad (15)$$

Moreover, the normalized gaps  $(\Delta_0/T, \dots, \Delta_N/T)$  jointly follow a Dirichlet $(1, \dots, 1)$  distribution. Marginally, each normalized gap  $\Delta_i/T$  follows a Beta $(1, N)$  distribution. As  $N$  becomes large, the Beta $(1, N)$  distribution approximates a Gamma $(1, 1/N)$  distribution, because

$$N \cdot (\Delta_i/T) \xrightarrow{d} \text{Exp}(1) \quad (16)$$

which suggests that

$$\Delta_i \approx \text{Gamma}(1, T/N) \quad (17)$$

Thus, in the large-sample limit, the gaps between ordered uniform samples behave approximately like scaled exponential random variables.

To simulate ordered points efficiently for a finite number  $M$  of samples, we propose sampling  $M$  independent gaps

$$\Delta_i \sim \text{Gamma}(1, T/M) \quad (18)$$

and constructing ordered points via the cumulative sums

$$U_{(i)} = \sum_{j=0}^{i-1} \Delta_j, \quad i = 1, \dots, M \quad (19)$$

This motivates our use of Gamma-distributed waiting times  $\tau_i$  in the model, capturing the natural variability in the timing of events.

- **Marks  $\vec{m}_i$ :**

The marks, representing additional information associated with each event, are modeled as drawn from a multivariate normal (Gaussian) distribution. Each mark vector  $\vec{m}_i$  has an associated mean vector  $\mu_i$  and covariance matrix  $\Sigma_i$ , with the distribution:

$$p(\vec{m}_i) = \mathcal{N}(\vec{m}_i; \mu_i, \Sigma_i) \quad (25)$$

explicitly given by:

$$\mathcal{N}(\vec{m}_i; \mu_i, \Sigma_i) = \frac{1}{(2\pi)^{d/2} |\Sigma_i|^{1/2}} \exp\left(-\frac{1}{2}(\vec{m}_i - \mu_i)^\top \Sigma_i^{-1}(\vec{m}_i - \mu_i)\right) \quad (26)$$

where  $d$  is the dimensionality of the mark vector.

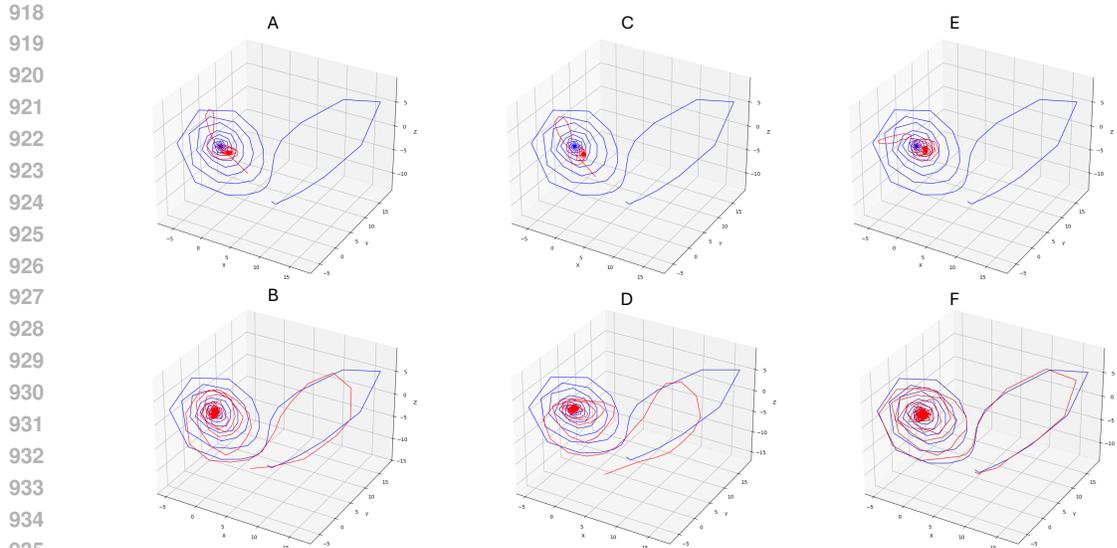
This modeling choice allows flexible and realistic characterization of the temporal dynamics  $\tau_i$  and the event-related features  $\vec{m}_i$  within the system under study.

## A.6 COMPARISON OF LATENT TRAJECTORIES

In this section, we provide additional analyses of the latent trajectories inferred by the compared models. As discussed in the main text, the latent spaces of baseline continuous-time models (e.g., Linear SDE, GP-SDE, GP-SLDS) are identifiable only up to an arbitrary affine transformation. To enable meaningful comparisons, we apply a Procrustes-based alignment procedure between the inferred latent trajectories and the ground-truth latent dynamics.

Representative examples are shown in Figure 5, where we compare raw latent trajectories (top panels) with their aligned counterparts (bottom panels) across baseline models. This visualization highlights the necessity of alignment for baseline approaches, as their raw latents are not directly comparable to the true dynamics.

The generative structure of MIP-CSDE naturally constrains its latent space, yielding trajectories that are more directly interpretable without alignment. Nonetheless, for fairness, all quantitative performance metrics reported in the main text are computed in the observation space.



936  
937  
938  
939

Figure 5: **Reconstruction of the Lorenz Trajectory Using Baseline Models.** Panels (A–B): Linear SDE; (C–D): GP-SDE; (E–F): GP-SLDS. For each model, the top panel shows the raw reconstructed latent trajectory, and the bottom panel shows the trajectory after Procrustes alignment to the ground truth.

940  
941

#### A.7 RUNTIME SCALABILITY EXPERIMENTS

942  
943  
944  
945  
946  
947  
948  
949

This appendix examines the empirical runtime scaling of the SMC–EM procedure used in our experiments. We vary two primary factors that drive computational cost: the number of particles  $P$  in the SMC layer and the sequence length  $K$  (number of time bins). For each setting, we run the Lorenz benchmark ten times with independent random seeds and report wall-clock time averaged over runs. Following common practice for GPU timing, we insert explicit CUDA synchronizations around the timed region and use a high-resolution host timer; we discard a short warm-up to avoid one-time kernel compilation and cache effects. All experiments use the same model configuration and batch size as in the main results to isolate the effect of  $P$  and  $K$ .

950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

Figure 6 summarizes the measurements. Panel (A) varies  $P$  at fixed  $K$ , plotting seconds per 1,000 time bins on a log scale. Panel (B) varies  $K$  at fixed  $P = 20,000$ , reporting seconds per 20,000 particles. In both regimes, ordinary least squares fits (orange) achieve  $R^2 \geq 0.995$  against the measured times (blue), consistent with the expected  $\mathcal{O}(P)$  and  $\mathcal{O}(K)$  complexity under our implementation. Absolute times depend on hardware, kernel fusion, and memory bandwidth, but the trends align with the cost analysis in Section 2.3.2. We note that memory usage grows linearly in  $P$  and modestly in  $K$  due to buffering of particle states; for large  $P$ , gradient checkpointing and mixed precision can reduce footprint without materially affecting the observed scaling.

972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

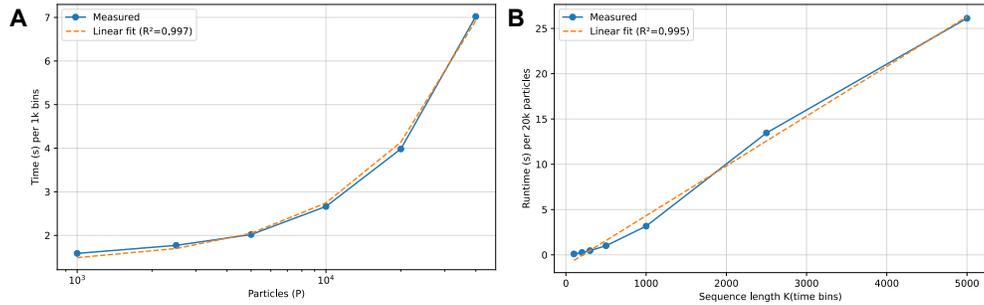


Figure 6: **Runtime of the SMC-EM Algorithm With Varying Numbers of Particles on the Lorenz Benchmark.** (A) Runtime vs. particle count  $P$ , shown on a log scale (seconds per 1k bins). (B) Runtime vs. sequence length  $K$  (seconds per 20k particles). Results are averaged over 10 runs. In both cases, measured runtime (blue) closely follow linear fits (orange,  $R^2 \geq 0.995$ ), confirming the expected  $O(P \cdot N)$  complexity discussed in Section 2.2.2. Experiments were run on an NVIDIA T4 GPU (16 GB) using PyTorch + CUDA.

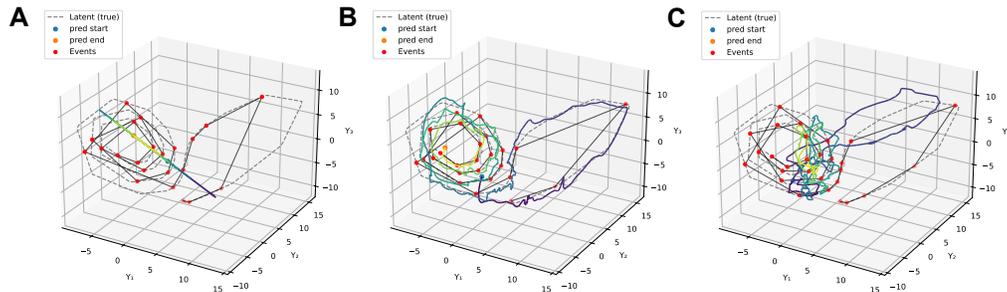
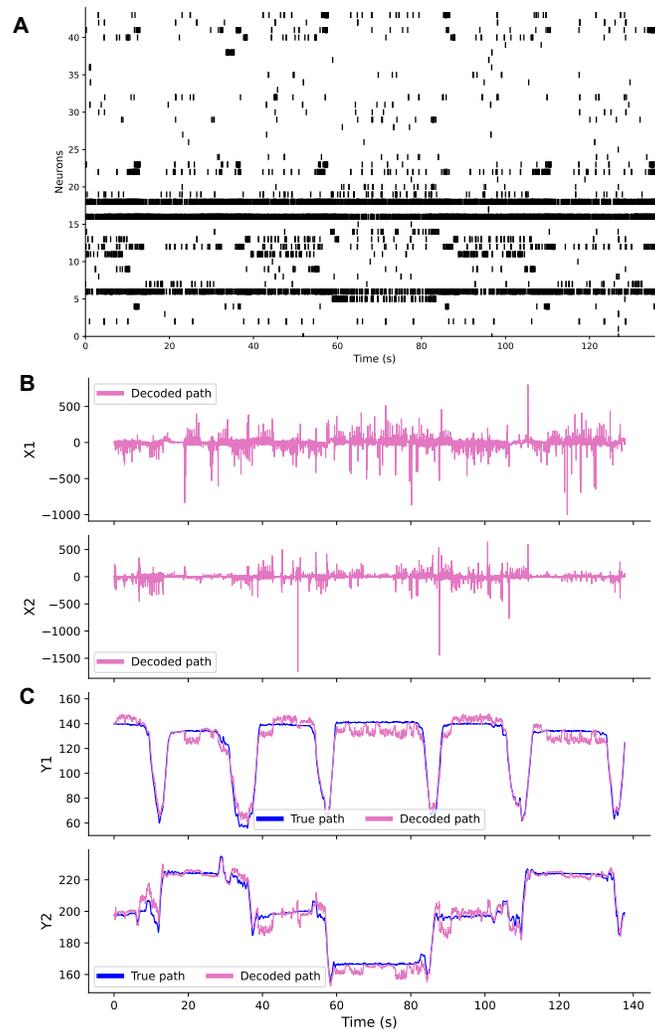


Figure 7: **Effect of latent dimensionality on MIP-CSDE decoding for the Lorenz system.** (A) One-dimensional latent model: decoded trajectory for a single observed coordinate of the Lorenz system, illustrating underfitting when the latent space is too restrictive. (B) Two-dimensional latent model: decoded trajectory with improved reconstruction of the underlying Lorenz dynamics compared to the 1D case. (C) Four-dimensional latent model: decoded trajectory capturing finer structure of the Lorenz dynamics, showing that increasing latent dimensionality allows MIP-CSDE to represent more of the systems nonlinear behavior.

1026  
 1027  
 1028  
 1029  
 1030  
 1031  
 1032  
 1033  
 1034  
 1035  
 1036  
 1037  
 1038  
 1039  
 1040  
 1041  
 1042  
 1043  
 1044  
 1045  
 1046  
 1047  
 1048  
 1049  
 1050  
 1051  
 1052  
 1053  
 1054  
 1055  
 1056  
 1057  
 1058  
 1059  
 1060  
 1061  
 1062  
 1063  
 1064  
 1065  
 1066  
 1067  
 1068  
 1069  
 1070  
 1071  
 1072  
 1073  
 1074  
 1075  
 1076  
 1077  
 1078  
 1079



**Figure 8: MIP-CSDE decoding of 2D position from hippocampal population activity.** (A) Raster plot of spiking activity from 62 simultaneously recorded hippocampal neurons during a representative segment of the rats movement. Each row corresponds to one neuron and each tick marks a spike. (B) Inferred mean velocity along the horizontal and vertical axes over time, obtained from the MIP-CSDE latent state and mapped into velocity coordinates. The model captures the main changes in movement speed and direction. (C) Decoded mean 2D position compared with the rats true trajectory for the same segment. The close overlap between decoded and true paths illustrates that MIP-CSDE can accurately reconstruct position from the observed spike trains.

1080  
 1081  
 1082  
 1083  
 1084  
 1085  
 1086  
 1087  
 1088  
 1089  
 1090  
 1091  
 1092  
 1093  
 1094  
 1095  
 1096  
 1097  
 1098  
 1099  
 1100  
 1101  
 1102  
 1103  
 1104  
 1105  
 1106  
 1107  
 1108  
 1109  
 1110  
 1111  
 1112  
 1113  
 1114  
 1115  
 1116  
 1117  
 1118  
 1119  
 1120  
 1121  
 1122  
 1123  
 1124  
 1125  
 1126  
 1127  
 1128  
 1129  
 1130  
 1131  
 1132  
 1133

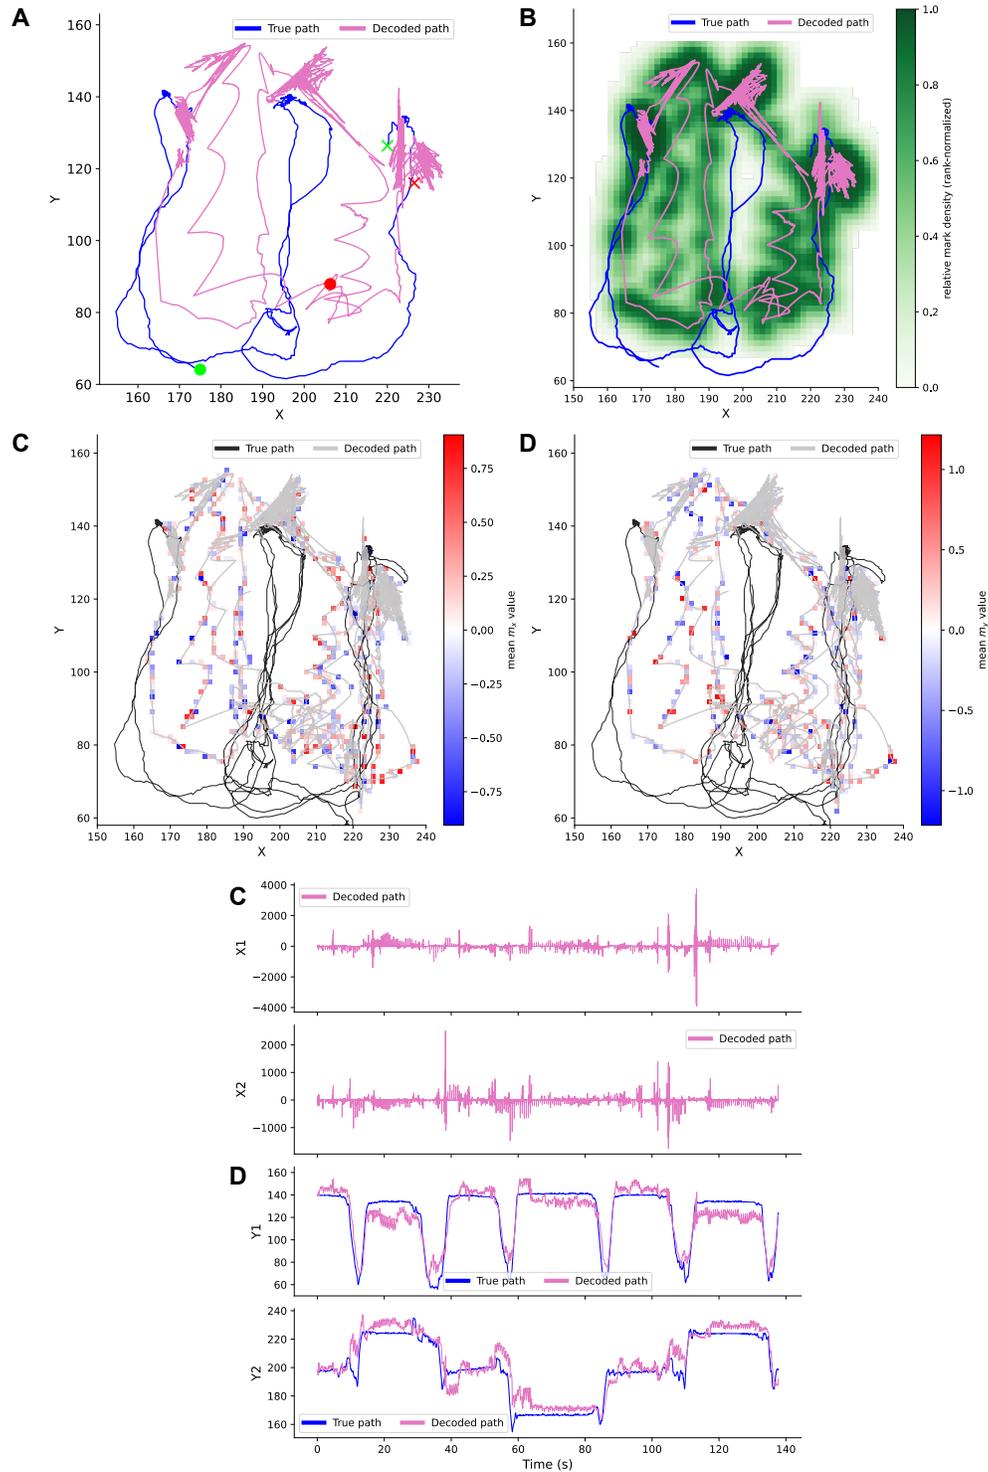


Figure 9: MIP-CSDE rat decoding with 1,000 particles. (A) One example traversal of the W-maze, showing the rat moving from the left arm to the right arm. The model is trained and decoded on the full recording session, but for visual clarity we display a single left-to-right trajectory. (B) Inducing-point locations in latent vs. observed position for this traversal. (C-D) Heatmaps of inferred  $x$ - and  $y$ -coordinate marks along the same latent trajectory. (E) Inferred mean horizontal and vertical velocity over time. (F) Decoded mean 2D position vs. true trajectory, showing accurate reconstruction with 1,000 particles.

1134  
 1135  
 1136  
 1137  
 1138  
 1139  
 1140  
 1141  
 1142  
 1143  
 1144  
 1145  
 1146  
 1147  
 1148  
 1149  
 1150  
 1151  
 1152  
 1153  
 1154  
 1155  
 1156  
 1157  
 1158  
 1159  
 1160  
 1161  
 1162  
 1163  
 1164  
 1165  
 1166  
 1167  
 1168  
 1169  
 1170  
 1171  
 1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187

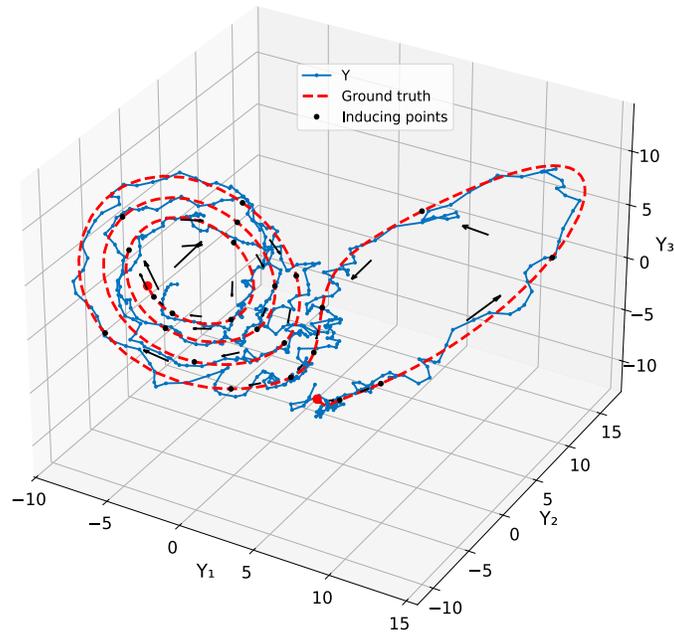


Figure 10: **Decoding of Lorenz Trajectory Using MIP-CSDE.** Here, the projection matrix from the latent process to observations and the additive noise covariance in the simulated data are known. The plot shows one inferred trajectory ( $Y$ ), and dots indicate a sample set of inducing points. The decoded trajectory is aligned with the generated trajectory.

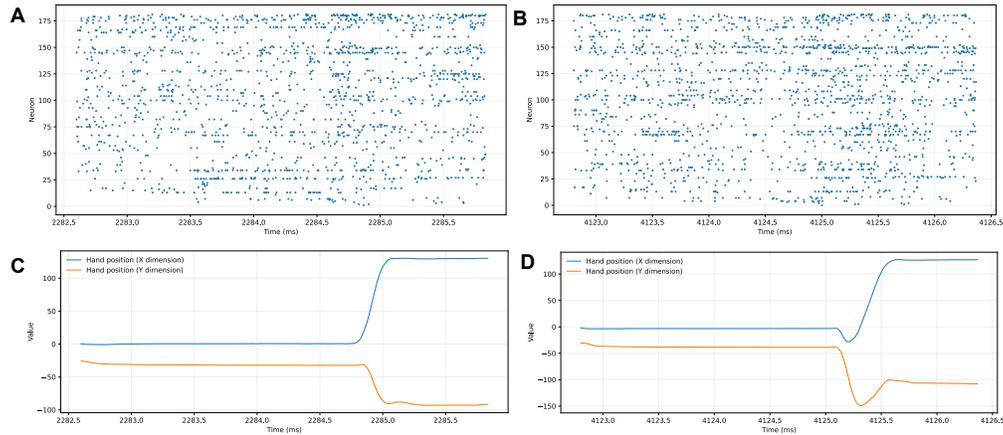


Figure 11: **Raster Plots of Monkey PDM and M1 Neurons Along with Hand Position During the Reach Task.** (A, B) Raster plots of 182 neurons from M1 and PMd, showing activity before target onset, the go cue, and target acquisition across two task trials. (C, D) Corresponding monkey hand positions during the same trials.