
Generative Modeling of Neural Dynamics via Latent Stochastic Differential Equations

Ahmed ElGazzar¹ Marcel van Gerven¹

Abstract

We propose a probabilistic framework for developing computational models of biological neural systems. In this framework, physiological recordings are viewed as discrete-time partial observations of an underlying continuous-time stochastic dynamical system which implements computations through its state evolution. To model this dynamical system, we employ a system of coupled stochastic differential equations with differentiable drift and diffusion functions and use variational inference to infer its states and parameters. This formulation enables seamless integration of existing mathematical models in the literature, neural networks, or a hybrid of both to learn and compare different models. We demonstrate this in our framework by developing a generative model that combines coupled oscillators with neural networks to capture latent population dynamics from single-cell recordings. Evaluation across three neuroscience datasets spanning different species, brain regions, and behavioral tasks show that these hybrid models achieve competitive performance in predicting stimulus-evoked neural and behavioral responses compared to sophisticated black-box approaches while requiring an order of magnitude fewer parameters, providing uncertainty estimates, and offering a natural language for interpretation.

1. Introduction

Unlike many scientific fields, neuroscience lacks a unified theoretical framework that bridges micro-scale mechanisms with macro-scale observations. That is, while the biophysical processes underlying individual neuron spiking have been well-understood for decades (Hodgkin & Huxley, 1952), the computational principles governing collective

neural activity that give rise to cognition and behavior remain elusive.

Dynamical systems theory provides a promising framework to address this gap (Durstewitz et al., 2023). Within this formalism, the brain is viewed as a complex dynamical system where neural activity patterns evolve through a high-dimensional state space according to well-defined governing equations, providing a mathematical language for understanding how distributed neural circuits implement computation and generate behavior (Van Gelder, 1998; Izhikevich, 2007; Deco et al., 2008; Breakspear, 2017; Favela, 2021). This perspective has led to diverse mathematical models, ranging from mean-field approximations of biophysical models that preserve mechanistic interpretability (Wilson & Cowan, 1972; Jirsa & Haken, 1996; Robinson et al., 1997), to phenomenological models inspired by statistical physics that capture emergent collective behavior (Amari, 1977; Hopfield, 1982; Sompolinsky et al., 1988). However, while these models have provided key insights into various neural phenomena (Goldman-Rakic, 1995; Wang, 2002; Fries, 2005), fitting them to single-trial neural data has proven challenging (Urai et al., 2022).

In recent years, latent variable models have emerged as a powerful alternative. These models represent high-dimensional neural activity in terms of low-dimensional latent states evolving according to Markovian dynamics (Paninski et al., 2010; Hurwitz et al., 2021a). Neural computations are thus assumed to be implemented through dynamic motifs in a low-dimensional latent space (e.g., attractors, oscillations, bifurcations) (Vyas et al., 2020; Khona & Fiete, 2022; Durstewitz et al., 2023). The popularity of this approach stems from converging empirical evidence that neural population activity during simple tasks lives on a low-dimensional manifold (Gao & Ganguli, 2015). Coupled with ongoing developments in machine learning (Kingma & Welling, 2013; Fabius & Van Amersfoort, 2014; Rubanova et al., 2019), latent variable models are becoming a standard tool in systems and computational neuroscience (Chen et al., 2018; Glaser et al., 2020; Hurwitz et al., 2021b; Zhou & Wei, 2020; Kim et al., 2021).

However, selecting an appropriate model for representing latent neural dynamics presents a significant challenge. The

¹Donders Institute for Brain, Cognition and Behaviour Radboud University Nijmegen, The Netherlands. Correspondence to: Ahmed ElGazzar <ahmed.elgazzar@donders.ru.nl>.

model must be expressive enough to capture complex non-linear dynamics, supporting flexible neural computations, while remaining amenable to interpretation and analysis to be useful for hypothesis generation and testing. Furthermore, model predictions should account for various sources of uncertainty typically present in such modeling endeavors, such as measurement noise, process noise, and model uncertainty. Finally, to utilize the inferred model for online applications as in brain computer interfaces (BCIs), there should be an efficient tractable method for sampling from the model in real-time.

We posit that recent developments in scientific machine learning, where domain-specific models — often expressed as differential equations — are combined with neural networks and scalable optimization techniques provide an untapped potential to address these challenges (Raissi et al., 2019; Rackauckas et al., 2020; Lai et al., 2021; Karniadakis et al., 2021). Additionally, ongoing innovations in training (neural) stochastic differential equations (SDEs) (Tzen & Raginsky, 2019a; Li et al., 2020; Zeng et al., 2023; Course & Nair, 2024), provide an opportunity to develop scalable probabilistic models in neuroscience, a field wherein the systems in question are inherently stochastic (Laing & Lord, 2009; Rolls & Deco, 2010) and the existing measurement tools only provide a coarse proxy of the system under study (Stevenson & Kording, 2011; Urai et al., 2022).

Contributions We develop a general framework for modeling neural dynamics using latent SDEs with differentiable state- and input-dependent drift and diffusion functions. Within this framework, we introduce a specific generative model, combining coupled oscillators with neural networks, to capture neural population dynamics. We evaluate our approach across multiple experimental scenarios, demonstrating its effectiveness and efficiency in predicting stimulus-evoked neural and behavioral responses in both reaching and decision-making tasks. Through systematic comparison of latent ODE- and SDE-based models under varying noise levels and state dimensionality in simulated systems, we highlight the critical importance of modeling process noise; a factor often overlooked in current latent variable approaches. Finally, our results suggest that data-driven coupled oscillators could provide an expressive yet interpretable framework for understanding neural dynamics, while maintaining computational tractability.

2. Framework

Notations Let $x: [0, \tau] \rightarrow \mathbb{R}^{d_x}$ denote a continuous path of latent neural states, which we observe through neural recordings (e.g., spike counts, fMRI time-series) denoted by $\mathcal{Y} = \{\mathbf{y}_i\}_{i=1}^T$, where each $\mathbf{y}_i \in \mathbb{R}^{d_y}$ is observed at time $t_i \in [0, \tau]$. Additionally, we may observe behavioral

responses (e.g., movement trajectories, choice responses) represented as $\mathcal{B} = \{\mathbf{b}_i\}_{i=1}^T$, where each $\mathbf{b}_i \in \mathbb{R}^{d_b}$. In some settings, these observations are driven by external stimuli or task variables (e.g., visual stimuli, auditory cues) denoted by $\mathcal{V} = \{\mathbf{v}_i\}_{i=1}^T$, where each $\mathbf{v}_i \in \mathbb{R}^{d_v}$. A single trial is thus defined as $\mathcal{S} = (\mathcal{Y}, \mathcal{B}, \mathcal{V})$, where \mathcal{Y} , \mathcal{B} and/or \mathcal{V} could be missing. Given a dataset $\mathcal{D} = \{\mathcal{S}_k\}_{k=1}^K$ consisting of K trials, our objectives are to: (1) infer the underlying neural dynamics that generated the observations, represented by the posterior distribution $p(x | \mathcal{D})$, (2) learn a generative model capable of predicting stimulus-evoked responses, represented by the joint distribution $p(\mathcal{Y}, \mathcal{B} | \mathcal{V})$, and (3) extract interpretable dynamical features that provide insights into neural computations.

2.1. Generative model

We model latent neural dynamics x as a continuous-time stochastic process, specified as the solution to the following Itô SDE:

$$dx(t) = \mu_\theta(x(t), u(t))dt + \sigma_\theta(x(t), u(t))dw(t) \quad (1)$$

with $x(0) \sim \mathcal{P}_0$, where \mathcal{P}_0 is the probability distribution of the initial condition $x(0)$, and $u(t) = \sum_{i=1}^T \eta_\theta(\mathbf{v}_i)\psi_i(t)$ is an interpolated encoded representation of \mathcal{V} . Here, $\{\psi_i\}_{i=1}^T$ are basis functions chosen for the desired interpolation scheme and $\eta_\theta: \mathbb{R}^{d_v} \rightarrow \mathbb{R}^{d_u}$ is an input encoder whose form depends on the input modality. The functions $\mu_\theta: \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{d_x}$ and $\sigma_\theta: \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{d_x \times d_w}$ are Lipschitz-continuous differentiable functions defining the drift and diffusion coefficients of the SDE, respectively. Finally, $w: [0, \tau] \rightarrow \mathbb{R}^{d_w}$ denotes a standard d_w -dimensional Wiener process.

This SDE induces a path measure on the space of continuous paths which implicitly define the distribution $p(x | \mathcal{V})$. Samples from this distribution are obtained by solving the SDE with different realizations of $x(0)$ and w :

$$x(t) = x(0) + \int_0^t \mu_\theta(x(s), u(s))ds + \int_0^t \sigma_\theta(x(s), u(s))dw(s) \quad (2)$$

The continuous-time latent states $x_i = x(t_i)$ are mapped to discrete-time observations through:

$$\mathbf{y}_i \sim p(\cdot | \lambda_\theta(x_i)), \quad \mathbf{b}_i \sim p(\cdot | \rho_\theta(x_i)) \quad (3)$$

where $p(\cdot | \lambda_\theta(x))$ and $p(\cdot | \rho_\theta(x))$ denote the observation models (e.g., Poisson model for spike counts, Gaussian model for continuous measurements), with λ_θ and ρ_θ being differentiable functions mapping the latent state to the parameters of these distributions.

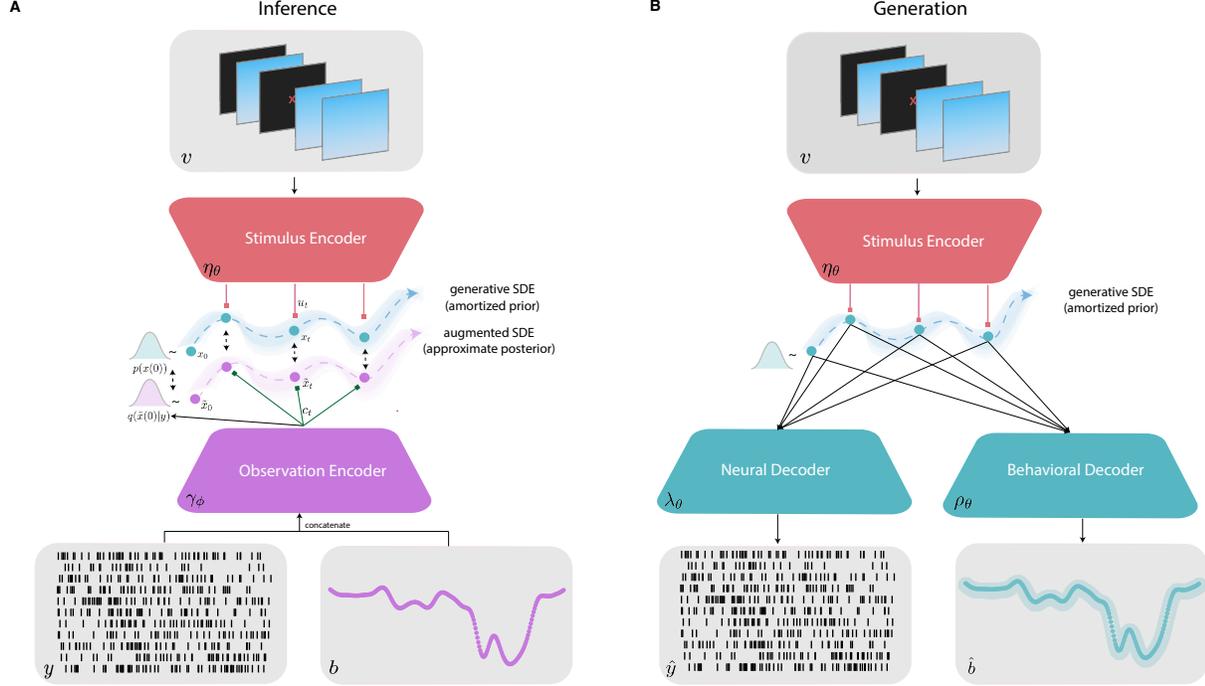


Figure 1. **Schematic illustration of the framework.** **A** Inference model for inferring the approximate posterior and learning the parameters of the model from neural and/or behavioral observations. **B** Generative model used for downstream applications after training.

2.2. Inference and learning

The true parameters θ of our generative model are unknown and need to be learned from the data \mathcal{D} . Moreover, the exact posterior distribution $p(x | \mathcal{D})$ is intractable to compute. Variational inference provides an efficient and scalable approach to tackle this problem. However unlike standard latent variable models, the latent state in our model are realizations of a continuous-time stochastic process. Thus we need to define a parametric distribution over paths rather than random-valued vectors.

Following (Tzen & Raginsky, 2019a,b; Li et al., 2020), we define the approximate posterior distribution $q(x | \mathcal{D})$ via an augmented neural SDE. Specifically, here we define the approximate posterior path $\tilde{x}: [0, \tau] \rightarrow \mathbb{R}^{d_x}$ as realizations of the following SDE (we omit time indices for ease of notation):

$$d\tilde{x} = \nu_\phi(\tilde{x}, u, c)dt + \sigma_\theta(\tilde{x}, u)dw \quad (4)$$

where $\tilde{x}(0) \sim \mathcal{Q}_0$ and $c(t) = \sum_{i=1}^T \xi_\phi(\mathbf{y}_i, \mathbf{b}_i)\psi_i(t)$ is a continuous path encoding the observations, with $\xi_\phi: \mathbb{R}^{d_y \times d_b} \rightarrow \mathbb{R}^{d_c}$ being a modality-specific observation encoder. We define the distribution of the initial condition as $\mathcal{Q}_0 = \mathcal{N}(\alpha_\phi(\mathbf{y}_{1:c}), \beta_\phi(\mathbf{y}_{1:c}))$, where α_ϕ and β_ϕ are neural networks, and $\mathbf{y}_{1:c} = \{\mathbf{y}_1, \dots, \mathbf{y}_c\}$. We further set $t_c \ll t_n$

to prevent over-parameterizing the initial conditions and rendering the dynamics redundant (Zabihi et al., 2019).

Here $\nu_\phi: \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \times \mathbb{R}^{d_c} \rightarrow \mathbb{R}^{d_x}$ is a neural network representing the drift of the augmented SDE. Note that while the drift term is different from that of the generative model, the diffusion term σ_θ is shared. As illustrated in (Tzen & Raginsky, 2019a), this specification enables leveraging Girsanov’s theorem (Girsanov, 1960) to obtain a tractable KL divergence between the the generative (prior) and augmented (approximate posterior) distributions:

$$D_{\text{KL}}(\mathcal{Q}_\tau \| \mathcal{P}_\tau) = \mathbb{E}_{\tilde{x}} \left[\int_0^\tau \frac{1}{2} \|\Delta(\tilde{x}, u, c)\|^2 dt \right] \quad (5)$$

where $\Delta(\tilde{x}, u, c) = \sigma_\theta(\tilde{x}, u)^{-1} (\nu_\phi(\tilde{x}, u, c) - \mu_\theta(\tilde{x}, u))$ with \mathcal{P}_τ and \mathcal{Q}_τ denoting the measures induced by the generative and augmented SDE, respectively.

Using this we can define an evidence lower bound (ELBO) on the conditional marginal likelihood of the observations:

$$\begin{aligned} \log p(\mathcal{Y}, \mathcal{B} | \mathcal{V}) \geq \\ \mathbb{E}_{\tilde{x}} \left[\sum_{i=1}^T \log p(\mathbf{y}_i | \tilde{x}_i) + \sum_{i=1}^T \log p(\mathbf{b}_i | \tilde{x}_i) \right] \\ - D_{\text{KL}}(\mathcal{Q}_0 \| \mathcal{P}_0) - D_{\text{KL}}(\mathcal{Q}_\tau \| \mathcal{P}_\tau) \quad (6) \end{aligned}$$

The ELBO consists of three main terms: (1) the expected log-likelihood of the observations under the approximate posterior, which includes both neural recordings and behavioral responses, (2) the KL divergence between the initial state distributions, and (3) the path-wise KL divergence between the approximate posterior and prior processes derived via Girsanov’s theorem (see Appendix B for detailed derivation of the loss).

The model parameters θ and variational parameters ϕ can be jointly optimized by maximizing this ELBO using stochastic gradient descent. Both the expectation over \tilde{x} and the path-wise integral in the KL divergence term are computed using numerical integration schemes suitable for SDEs (see Appendix C for implementation details).

This formulation enables us to simultaneously learn the parameters of the generative model while performing approximate posterior inference over the latent neural trajectories. The learned generative model can then be used for various online and offline downstream tasks including probabilistic prediction of stimulus-evoked responses and extraction of interpretable dynamical features.

3. Hybrid modeling

So far we have described a general framework for modeling the latent dynamics of a neural system from empirical observations. This framework allows us to leverage almost arbitrary differentiable functions to model the drift and diffusion terms of an SDE. A natural choice would be to use a neural SDE to act as the generative model because of their universal approximation properties. However, this comes at the cost of efficiency and interpretability. An alternative option, is to leverage prior models proposed in the literature and use neural networks either to approximate their parameters or to augment their dynamics. Such hybrid models can improve sample-efficiency, generalization performance and, are easier to interpret in comparison to models that rely only on neural networks (Rackauckas et al., 2020; El-Gazzar & van Gerven, 2024). In this section, we detail a specific instantiation of the framework which combines a well-understood dynamical system and neural networks to obtain an expressive, interpretable and parameter-efficient generative model of neural population dynamics.

3.1. Stochastic nonlinear coupled oscillators

A well-known class of models that are used to study several biological and physical systems are coupled oscillators. Coupled oscillators provide a rich yet simplified language for studying complex interactions in systems exhibiting emergent behavior such as synchronization, pattern formation and phase transitions (Winfree, 1980).

We focus on a specific instantiation of coupled oscillators

(Matthews et al., 1991) described by the following complex-valued ODE:

$$dz = ((\alpha + i\omega)z + |z|^2z + \kappa z)dt \quad (7)$$

where $z \in \mathbb{C}^{d_z}$ represents the position of d_z oscillators in the complex plane, $\alpha \in \mathbb{R}^{d_z}$ the bifurcation parameters, $\omega \in \mathbb{R}^{d_z}$ denotes the natural frequency, and $\kappa \in \mathbb{R}$ represents all-to-all coupling strength. This formulation captures the essential dynamics of coupled limit-cycle oscillators near a supercritical Hopf bifurcation, where in the weak coupling regime, the term $|z(t)|^2$ determines the limit cycle radius. Despite its simplicity, this system can exhibit diverse emergent dynamical regimes including frequency locking, amplitude death, and chaos through variations in its parameters and the initial conditions.

Here we propose to use this system as the drift of our generative SDE in our framework. Specifically, we use the position of the oscillator as the latent state $x(t)$ in the generative SDE. Additionally, we consider the coupling strength κ as a time-varying input dependent function of the external input $u(t)$.

Incorporating these updates and writing the model in terms of the real and imaginary parts of the oscillator $z(t) = a(t) + b(t)i$, we obtain the following system of SDEs:

$$\begin{aligned} da &= [\alpha a - \omega b - (a^2 + b^2)a + \kappa(u)a]dt + \Sigma_a(x, u)dw \\ db &= [\omega a + \alpha b - (a^2 + b^2)b + \kappa(u)b]dt + \Sigma_b(x, u)dw \end{aligned} \quad (8)$$

where κ , Σ_a , Σ_b are neural networks representing the input dependent connectivity and the diffusion terms for the real and imaginary states respectively. The trainable parameters θ of the generative SDE are thus the parameters of the neural networks as well as α and ω . We refer to this dynamics model as coupled oscillator SDE (CO-SDE).

3.2. Modeling multiple interacting populations

Often, we have access to neural activity recorded from neurons across multiple brain regions, simultaneously. Given access to recordings $\mathbf{y} = (\mathbf{y}^1, \dots, \mathbf{y}^J)$ from J brain regions, we can extend the generative model by factorizing the latent state $x = (x^1, \dots, x^J)$ where the number of oscillators N_j can vary depending on the population size and complexity. The observed activity of a population is then

$$\mathbf{y}_i^j \sim \text{Poisson}(\lambda) \quad (9)$$

where $\lambda = \exp(\zeta_\theta^j(x_i^j))$ is a vector of rate parameters which depends on a population-specific transformation $\zeta_\theta^j: \mathbb{R}^{N_j} \rightarrow \mathbb{R}^{N_j}$ that maps the latent state to the firing rates of the population.

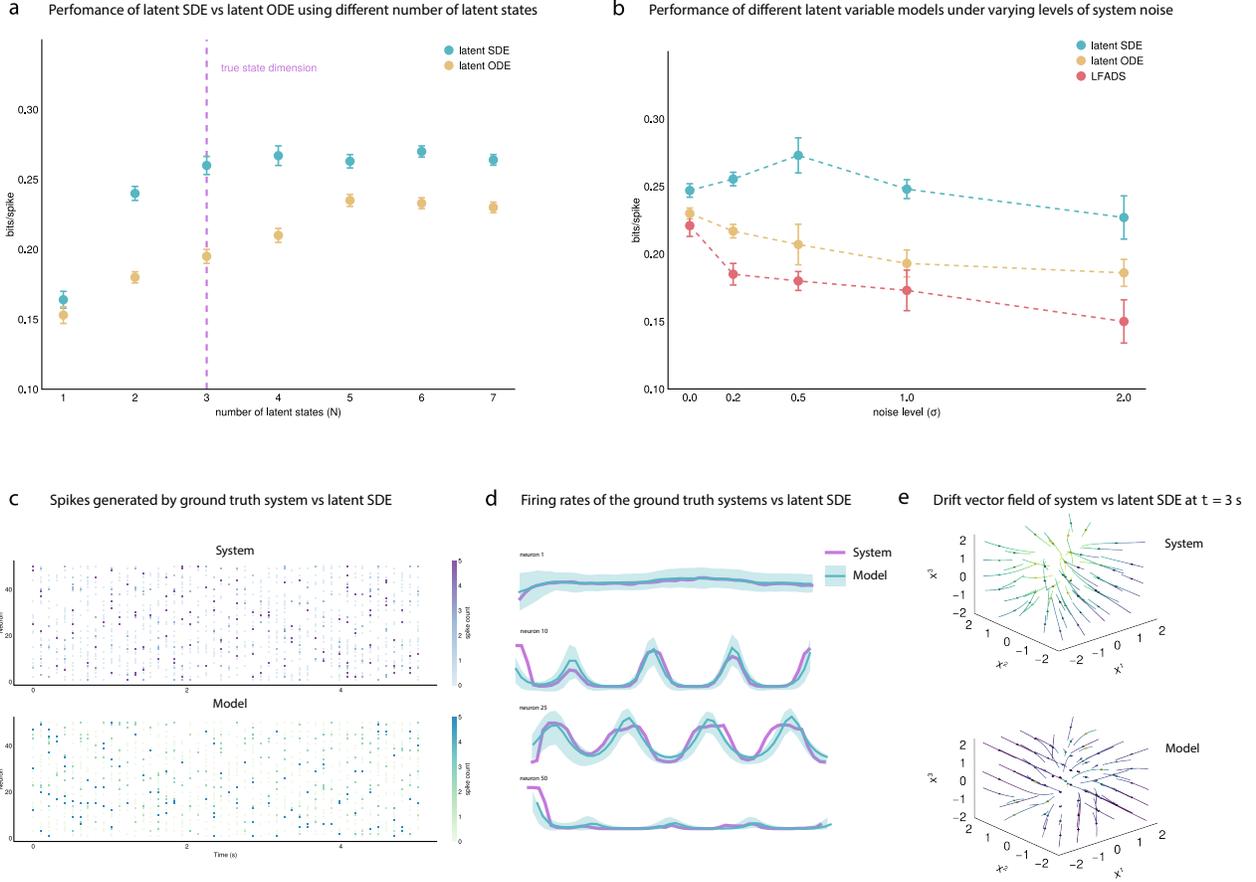


Figure 2. Results on simulated spiking neural system. **a** 5-fold cross validation results of comparing the performance of latent ode vs latent sde on fitting simulated data under different number of model latent states. **b** 5-fold cross validation results of different latent variable models under varying levels of process noise in the ground-truth system. **c** A sample of the spikes generated via the system vs spikes generated via the latent SDE model in response to the same input. **d** Samples of underlying firing rates generated via the system versus mean and sd of 30 samples from a trained latent SDE model. **e** Phase portrait of the ground truth system versus the drift vector field of a trained latent SDE model at one time point under a fixed input.

4. Experiments

4.1. Simulated data

We first sought to evaluate the efficacy of our framework on modeling a partially observed neural system with different levels of process noise. To this end, we generated spiking data from the following system:

$$\begin{aligned}
 x(0) &\sim \mathcal{N}(0, I) \\
 dx &= \delta^{-1} (-x + \tanh(Ax + Bu)) dt + \sigma dw \\
 y &\sim \text{Poisson}(\exp(Rx))
 \end{aligned} \tag{10}$$

where $x(t) \in \mathbb{R}^3$ describes the population average firing rate of three neural populations driven by a sinusoidal input $u(t)$

with randomly sampled frequency and phase. $A \in \mathbb{R}^{3 \times 3}$ and $B \in \mathbb{R}^{3 \times 1}$ describe the recurrent connectivity and input interaction matrices, respectively. The time scale of the network dynamics is given by $\delta \in \mathbb{R}$ and $\sigma \in \mathbb{R}$ describes the magnitude of process noise in the system. The system output is given by $y(t) \in \mathbb{N}_0^{50}$ representing the observed spike counts of 50 neurons generated by sampling from an inhomogeneous Poisson process with rates computed by affine transformation of the population average activity via the output matrix $R \in \mathbb{R}^{50 \times 3}$. We fit our latent SDE model on the spiking data by maximizing the Poisson-log likelihood of the observations given predicted rate and evaluated performance using bits per spike (bps), a metric that quantifies how well the predicted firing rates explain the observed

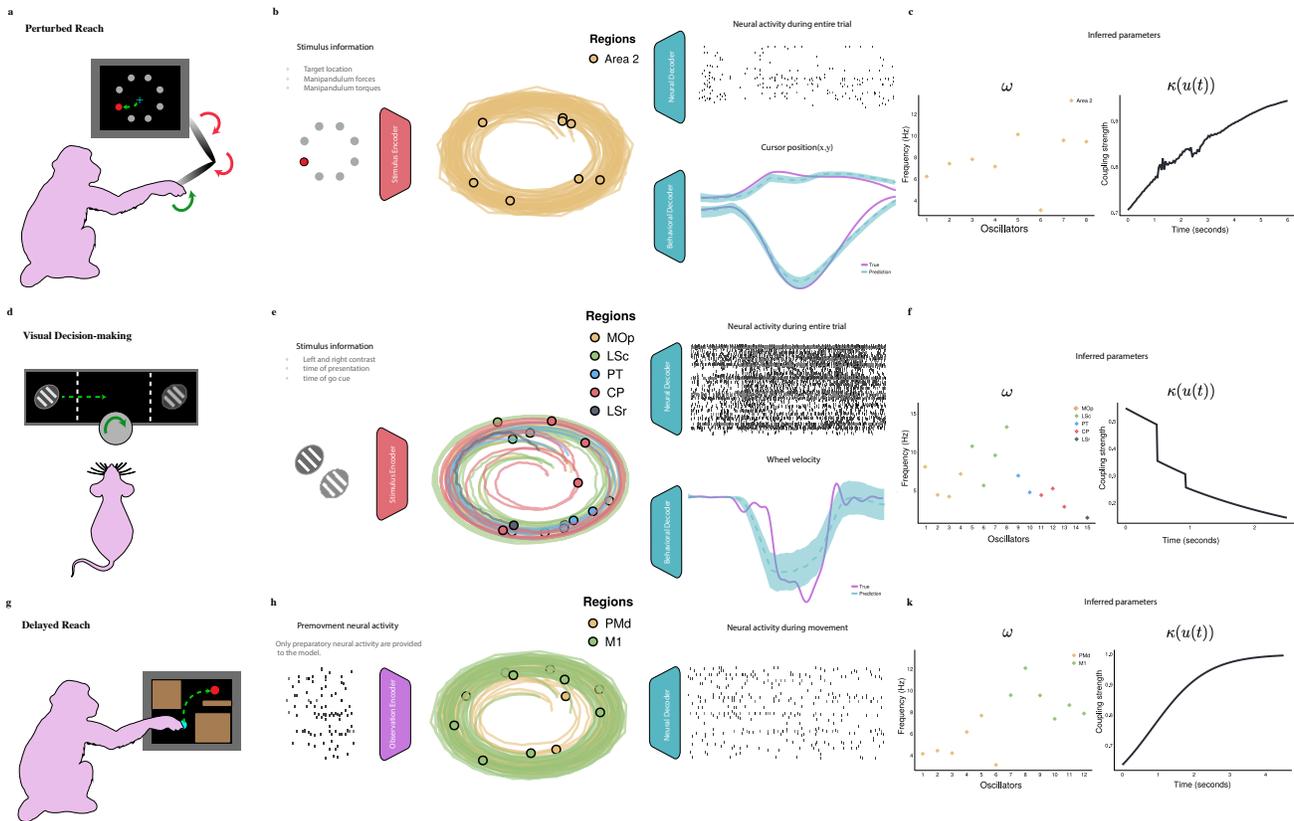


Figure 3. Generative modeling of neural and behavioral data via latent coupled oscillators across three different datasets. (a-c) Perturbed reach task: The generative model takes as input the target location and manipulandum forces/torques and is tasked with predicting neural responses of 64 Neurons in Area 2 and the behavioral response of the monkey as measured via the cursor position. The inferred frequency of the coupled oscillators as well as the coupling strength (for a sample trial) for one training run is show in **c**. **(d-f) Visual decision-making task:** Input consists of contrast levels as well as the timing of presentation and the go cue. The model is trained to predict both neural activity across multiple brain regions and wheel velocity. **(g-k) Delayed reach task:** Unlike the other tasks, this model receives only preparatory neural activity before movement onset and is trained to predict subsequent neural activity in the dorsal pre-motor (PMd) and primary motor cortex (M1) during movement execution.

spike patterns (Pei et al., 2021).

For our experiments, we varied the noise level $\sigma \in \{0.0, 0.5, 1.0, 2.0\}$ to test the robustness of the model. We used a latent neural SDE with varying sizes of latent dimension and compared against two baselines: a latent ODE (Rubanova et al., 2019) and LFADS (Pandarinath et al., 2018).

As demonstrated in Figure 2a, the latent SDE model can better explain the data variance when using fewer latent dimensions compared to a latent ODE model. Increasing the latent dimensions beyond the true state dimensions does not significantly improve the model fit unlike the ODE-based model which stagnates at a lower absolute value and higher latent dimension. Additionally, as shown in Figure 2b, the latent SDE model is more robust to process noise compared to latent ODE and LFADS models. These results suggest

that latent SDEs can be more reliable to model real-world data when the true state dimension is unknown and process noise can emerge from unobserved interactions. In this setting, a probabilistic initial state with deterministic dynamics, as in the case of latent ODEs and LFADS, may not be sufficient.

4.2. Empirical data

4.2.1. DATASETS & BASELINES

We evaluated our framework across three neuroscience datasets spanning different brain regions and behavioral tasks.

Perturbed reach task In this experiment, a monkey is presented with a target location and tasked with moving the cursor from the center to the target location by controlling

Table 1. Performance comparison of latent sequential models across different tasks and metrics. Prediction of neural responses is evaluated via computing the bits per spike (bps), and behavioral responses evaluated via R^2 . The number of parameters of the generative dynamics for each model is shown in average across experiments. Encoders and decoder architectures are shared across all models. Arrows (\uparrow/\downarrow) indicate whether higher or lower values are better. Best performances are highlighted in **bold**. The reported results are the mean and std. of 5-fold cross validation.

Latent Dynamics	# Params \downarrow	Datasets				
		Perturbed Reach		Visual Decision		Delayed Reach
		Neural (bps) \uparrow	Behavior (R^2) \uparrow	Neural (bps) \uparrow	Behavior (R^2) \uparrow	Neural (bps) \uparrow
RNN	8320	0.175 \pm 0.008	0.84 \pm 0.02	0.14 \pm 0.015	0.59 \pm 0.03	0.291 \pm 0.012
GRU	24960	0.187 \pm 0.004	0.81 \pm 0.015	0.16 \pm 0.008	0.63 \pm 0.02	0.305 \pm 0.006
LSTM	33280	0.191 \pm 0.003	0.85 \pm 0.01	0.16 \pm 0.006	0.66 \pm 0.015	0.309 \pm 0.004
Neural ODE	4880	0.221 \pm 0.005	0.86 \pm 0.012	0.15 \pm 0.01	0.59 \pm 0.025	0.316 \pm 0.008
Coupled Oscillator ODE	465	0.221 \pm 0.006	0.85 \pm 0.015	0.21 \pm 0.009	0.61 \pm 0.02	0.309 \pm 0.007
Neural SDE	4944	0.243 \pm 0.007	0.86 \pm 0.015	0.18 \pm 0.012	0.70 \pm 0.02	0.317 \pm 0.009
Coupled Oscillator SDE	526	0.238 \pm 0.008	0.86 \pm 0.018	0.22 \pm 0.011	0.67 \pm 0.025	0.314 \pm 0.01

a manipulandum (Chowdhury et al., 2020) (Figure 3a). On one portion of the trials, the monkey’s arm was perturbed before the reach by applying resisting forces on the manipulandum. Neural activity was recorded from Brodmann’s area 2 of the somatosensory cortex. Here, we trained a model to predict stimulus-evoked neural responses in area 2 as well as the hand position of the monkey. The input to the model are the target location, as well as the perturbing force and torques on the manipulandum.

Visual decision making In this task, a mouse is presented with images of different contrasts and the mouse is tasked with turning a wheel to bring the image with higher contrast to the center of the screen when a go cue is randomly presented (Steinmetz et al., 2019) (Figure 3d). Neural activity is recorded in multiple regions across the brain using Neuropixels probes. Here, we trained the model to predict stimulus-invoked neural responses across several brain regions simultaneously as well as predict the wheel velocity. The input to the model are the contrast levels, the timing of their presentation, and the timing of the go cue.

Delayed reach task The task is a delayed center-out reach task with barriers, resulting in a variety of straight and curved trajectories (Mante et al., 2013) (Figure 3g). Neural activity was recorded from the dorsal premotor (PmD) and primary motor (M1) cortices, and cursor, monkey gaze position, and monkey hand position and velocity are also provided. Here we train a model to predict neural responses during movement in both areas given the preparatory activity before the go cue. No stimulus information is provided to the model during training or prediction.

We compare different latent variable models across the three datasets. To enable a fair comparison, we fix the encoder and decoder architectures across all models while varying the latent sequential model. That is, we compare the performance of RNN, GRU, LSTM, neural ODE, and coupled

oscillator ODE models with that of our proposed coupled oscillator SDE and a neural SDE model. All the baseline models had a probabilistic initial state and were trained via variational inference. Note that this formulation enables generalizing the baselines to several previous models proposed in literature (Pandarinath et al., 2018; Hurwitz et al., 2021b; Kim et al., 2021; Versteeg et al., 2023) (See Appendix D for details on the model architecture).

4.2.2. EMPIRICAL RESULTS

Table 1 presents 5-fold cross-validation results across the three datasets, evaluating both neural activity prediction (in bits per spike) and behavioral performance (via the coefficient of determination, R^2). First, we observe that our stochastic models (Neural SDE and CO-SDE) on average consistently outperform their deterministic counterparts. This suggests that incorporating stochasticity into the latent dynamics is crucial for capturing the variability inherent in neural and behavioral data and that a stochastic initial state might not be sufficient to capture uncertainty in the dynamics.

Second, we observe that model complexity, as measured by parameter count, does not correlate with performance. The coupled oscillator models achieve competitive or superior performance while using only 465-526 parameters, compared to 33,280 parameters for the LSTM model. Note that while this number does not account for the number of parameters required for modeling the approximate posterior SDE (3528), this remains a significant improvement in parameter efficiency for the generative model. This suggests that incorporating appropriate inductive biases through the coupled oscillator architecture enables more efficient modeling of the underlying dynamics. This also suggests that standard deep learning models may be too over-parameterized to learn the underlying dynamical system during simple

behavioral tasks.

Beyond quantitative results, oscillator-based models provide a direct way to interpret latent neural dynamics. We visualize the value of the inferred frequencies of the model after training as well as the value of the coupling strength in response to sample stimuli in Figure 3. In the perturbed reach task, where we have data from one somatosensory area, we observe that the connectivity strength increases over time (Figure 3c). This is also observed in the delayed reach task where we have recordings from two regions in the motor in the motor cortex (Figure 3k). In contrast, on the visual decision making task where we have recordings from multiple regions across the brain the coupling strength decreases over time (Figure 3f). This decrease is marked via two significant drops at two timepoints which mark the presentation of the stimulus and the go cue, respectively. Additionally, we observe that on average, the inferred natural frequency of the oscillators differ between brain regions. Notably, the average frequency of oscillators corresponding to M1 is higher than the average frequency of oscillators corresponding to PMd during the delayed reach task (Figure 3k), which may be linked to distinct functional roles related to action preparation and execution (Thura et al., 2022).

5. Discussion

This work introduces a probabilistic framework for modeling neural dynamics that combines the expressiveness of neural networks with the interpretability of mechanistic or phenomenological models. Our empirical results across simulated and real datasets demonstrate several key advantages of this approach.

First, explicitly modeling probabilistic dynamics through SDEs consistently improves model performance compared to deterministic alternatives, even when the latter incorporate uncertainty through probabilistic initial conditions. This suggests that neural variability stems not just from uncertain initial states but from inherently stochastic dynamics (Laing & Lord, 2009). The superior performance of SDE-based models in high-noise regimes as shown in the simulations results further supports this conclusion. Additional simulation results also show that using stochastic dynamics can improve model fit while relying on lower-dimensional latent states; an observation that is also found in stochastic low-rank RNNs (Pals et al., 2024).

Second, our hybrid coupled oscillator model achieves competitive or superior performance while requiring an order of magnitude fewer parameters than standard deep learning approaches. This parameter efficiency suggests that incorporating appropriate inductive biases through mechanistic components can effectively constrain the model space (Rack-

auckas et al., 2020; ElGazzar & van Gerven, 2024). The success of coupled oscillators in particular hints at their potential as a natural description language for neural population dynamics, consistent with existing theoretical work on neural oscillations (Buzsaki & Draguhn, 2004; Breakspear et al., 2010; Churchland et al., 2012; Bick et al., 2020). This also supports the potential of coupled oscillators as an expressive and efficient tool to model arbitrary time-series data, as seen in emerging machine learning research (Rusch & Mishra, 2020; Effenberger et al., 2022; Rusch & Rus, 2024).

With that said, there are a number of limitations to the proposed approach which merits consideration. First, the framework relies on variational inference for training the model parameters. This could be suboptimal as variational inference is known to provide overconfident uncertainty estimates and is prone to several optimization challenges such as posterior collapse, especially when coupled with powerful decoders (Blei et al., 2017). Additionally, training the model parameters currently requires back-propagation through an SDE solver, which can be computationally expensive in terms of both time and memory requirements. Potential solutions could lie in emerging simulation-free approaches for training (neural) stochastic differential equations (Course & Nair, 2024; Zhang et al., 2024). Future work could also explore comparing different dynamical systems beyond coupled oscillators and focus on extending the current approach to model neural data from several subjects as well as other recording modalities.

Acknowledgements

This work is supported via the Dutch Brain Interface Initiative (DBI²) with project number 024.005.022 of the research programme Gravitation which is (partly) financed by the Dutch Research Council (NWO).

References

- Abrevaya, G., Dumas, G., Aravkin, A. Y., Zheng, P., Gagnon-Audet, J. C., Kozloski, J., Polosecki, P., Lajoie, G., Cox, D., Dawson, S. P., et al. Learning brain dynamics with coupled low-dimensional nonlinear oscillators and deep recurrent networks. *Neural Computation*, 33(8): 2087–2127, 2021.
- Abrevaya, G., Ramezani-Panahi, M., Gagnon-Audet, J.-C., Polosecki, P., Rish, I., Dawson, S. P., Cecchi, G., and Dumas, G. Effective latent differential equation models via attention and multiple shooting. In *The Symbiosis of Deep Learning and Differential Equations III*, 2023.
- Amari, S. Dynamics of pattern formation in lateral-

- inhibition type neural fields. *Biological cybernetics*, 27 (2):77–87, 1977.
- Bandyopadhyay, A., Ghosh, S., Biswas, D., Chakravarthy, V. S., and S. Bapi, R. A phenomenological model of whole brain dynamics using a network of neural oscillators with power-coupling. *Scientific Reports*, 13(1): 16935, 2023.
- Bick, C., Goodfellow, M., Laing, C. R., and Martens, E. A. Understanding the dynamics of biological and neural oscillator networks through exact mean-field reductions: a review. *The Journal of Mathematical Neuroscience*, 10 (1):9, 2020.
- 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.
- Breakspear, M. Dynamic Models of large-scale brain activity. *Nature Neuroscience*, 20(3):340–352, 2017.
- Breakspear, M., Heitmann, S., and Daffertshofer, A. Generative models of cortical oscillations: neurobiological implications of the Kuramoto model. *Frontiers in Human Neuroscience*, 4:190, 2010.
- Buzsaki, G. and Draguhn, A. Neuronal oscillations in cortical networks. *Science*, 304(5679):1926–1929, 2004.
- Chen, R. T., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 31, 2018.
- Chowdhury, R. H., Glaser, J. I., and Miller, L. E. Area 2 of primary somatosensory cortex encodes kinematics of the whole arm. *Elife*, 9:e48198, 2020.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., and Shenoy, K. V. Neural population dynamics during reaching. *Nature*, 487 (7405):51–56, 2012.
- Course, K. and Nair, P. Amortized reparametrization: efficient and scalable variational inference for latent SDEs. *Advances in Neural Information Processing Systems*, 36, 2024.
- Deco, G., Jirsa, V. K., Robinson, P. A., Breakspear, M., and Friston, K. The dynamic brain: From spiking neurons to neural masses and cortical fields. *PLoS Computational Biology*, 4(8):e1000092, 2008.
- Durstewitz, D., Koppe, G., and Thurm, M. I. Reconstructing computational system dynamics from neural data with recurrent neural networks. *Nature Reviews Neuroscience*, pp. 1–18, 2023.
- Effenberger, F., Carvalho, P., Dubinin, I., and Singer, W. A biology-inspired recurrent oscillator network for computations in high-dimensional state space, 2022.
- ElGazzar, A. and van Gerven, M. Universal differential equations as a common modeling language for neuroscience. *arXiv preprint arXiv:2403.14510*, 2024.
- Ermentrout, B. and Terman, D. H. *Mathematical Foundations of Neuroscience*, volume 35. Springer, 2010.
- Fabius, O. and Van Amersfoort, J. R. Variational recurrent auto-encoders. *arXiv preprint arXiv:1412.6581*, 2014.
- Favela, L. H. The dynamical renaissance in neuroscience. *Synthese*, 199(1-2):2103–2127, 2021.
- Fries, P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10):474–480, 2005.
- Fu, H., Li, C., Liu, X., Gao, J., Celikyilmaz, A., and Carin, L. Cyclical annealing schedule: A simple approach to mitigating KL vanishing. *arXiv preprint arXiv:1903.10145*, 2019.
- Gao, P. and Ganguli, S. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology*, 32:148–155, 2015.
- Girsanov, I. V. On transforming a certain class of stochastic processes by absolutely continuous substitution of measures. *Theory of Probability & Its Applications*, 5(3): 285–301, 1960.
- Glaser, J., Whiteway, M., Cunningham, J. P., Paninski, L., and Linderman, S. Recurrent switching dynamical systems models for multiple interacting neural populations. *Advances in Neural Information Processing Systems*, 33: 14867–14878, 2020.
- Goldman-Rakic, P. S. Cellular basis of working memory. *Neuron*, 14(3):477–485, 1995.
- Hodgkin, A. L. and Huxley, A. F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117 (4):500, 1952.
- Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
- Hu, A., Zoltowski, D., Nair, A., Anderson, D., Duncker, L., and Linderman, S. Modeling latent neural dynamics with gaussian process switching linear dynamical systems. *arXiv preprint arXiv:2408.03330*, 2024.

- Hurwitz, C., Kudryashova, N., Onken, A., and Hennig, M. H. Building population models for large-scale neural recordings: Opportunities and pitfalls. *Current Opinion in Neurobiology*, 70:64–73, 2021a.
- Hurwitz, C., Srivastava, A., Xu, K., Jude, J., Perich, M., Miller, L., and Hennig, M. Targeted neural dynamical modeling. *Advances in Neural Information Processing Systems*, 34:29379–29392, 2021b.
- Izhikevich, E. M. *Dynamical Systems in Neuroscience*. MIT press, 2007.
- Jirsa, V. K. and Haken, H. Field theory of electromagnetic brain activity. *Physical Review Letters*, 77(5):960, 1996.
- Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., and Yang, L. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- Khona, M. and Fiete, I. R. Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, 23(12):744–766, 2022.
- Kidger, P. On neural differential equations. *arXiv preprint arXiv:2202.02435*, 2022.
- Kim, T. D., Luo, T. Z., Pillow, J. W., and Brody, C. D. Inferring latent dynamics underlying neural population activity via neural differential equations. In *International Conference on Machine Learning*, pp. 5551–5561. PMLR, 2021.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kuramoto, Y. Self-entrainment of a population of coupled non-linear oscillators. In *International Symposium on Mathematical Problems in Theoretical Physics: January 23–29, 1975, Kyoto University, Kyoto/Japan*, pp. 420–422. Springer, 1975.
- Lai, Z., Mylonas, C., Nagarajaiah, S., and Chatzi, E. Structural identification with physics-informed neural ordinary differential equations. *Journal of Sound and Vibration*, 508:116196, 2021.
- Laing, C. and Lord, G. J. *Stochastic Methods in Neuroscience*. OUP Oxford, 2009.
- Li, X., Wong, T.-K. L., Chen, R. T., and Duvenaud, D. Scalable gradients for stochastic differential equations. In *International Conference on Artificial Intelligence and Statistics*, pp. 3870–3882. PMLR, 2020.
- Macke, J. H., Buesing, L., Cunningham, J. P., Yu, B. M., Shenoy, K. V., and Sahani, M. Empirical models of spiking in neural populations. *Advances in neural information processing systems*, 24, 2011.
- Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, 2013.
- Matthews, P. C., Mirollo, R. E., and Strogatz, S. H. Dynamics of a large system of coupled nonlinear oscillators. *Physica D: Nonlinear Phenomena*, 52(2-3):293–331, 1991.
- Oh, Y., Lim, D., and Kim, S. Stable neural stochastic differential equations in analyzing irregular time series data. *arXiv preprint arXiv:2402.14989*, 2024.
- Pals, M., Sağtekin, A. E., Pei, F., Gloeckler, M., and Macke, J. H. Inferring stochastic low-rank recurrent neural networks from neural data. *arXiv preprint arXiv:2406.16749*, 2024.
- Pandarinath, 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 auto-encoders. *Nature Methods*, 15(10):805–815, 2018.
- Paninski, L., Ahmadian, Y., Ferreira, D. G., Koyama, S., Rahnama Rad, K., Vidne, M., Vogelstein, J., and Wu, W. A new look at state-space models for neural data. *Journal of Computational Neuroscience*, 29:107–126, 2010.
- Pei, F., Ye, J., Zoltowski, D., Wu, A., Chowdhury, R. H., Sohn, H., O’Doherty, J. E., Shenoy, K. V., Kaufman, M. T., Churchland, M., et al. Neural latents benchmark’21: evaluating latent variable models of neural population activity. *arXiv preprint arXiv:2109.04463*, 2021.
- Rackauckas, C., Ma, Y., Martensen, J., Warner, C., Zubov, K., Supekar, R., Skinner, D., Ramadhan, A., and Edelman, A. Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*, 2020.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- Robinson, P. A., Rennie, C. J., and Wright, J. J. Propagation and stability of waves of electrical activity in the cerebral cortex. *Physical Review E*, 56(1):826, 1997.
- Rolls, E. T. and Deco, G. *The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function*. Oxford academic, 2010.
- Rubanov, Y., Chen, R. T., and Duvenaud, D. K. Latent ordinary differential equations for irregularly-sampled time series. *Advances in Neural Information Processing Systems*, 32, 2019.

- Rusch, T. K. and Mishra, S. Coupled oscillatory recurrent neural network (coRNN): An accurate and (gradient) stable architecture for learning long time dependencies. *arXiv preprint arXiv:2010.00951*, 2020.
- Rusch, T. K. and Rus, D. Oscillatory state-space models. *arXiv preprint arXiv:2410.03943*, 2024.
- Sahani, M. *Latent Variable Models for Neural Data Analysis*. California Institute of Technology, 1999.
- Schimel, M., Kao, T.-C., Jensen, K. T., and Hennequin, G. iLQR-VAE: control-based learning of input-driven dynamics with applications to neural data. *BioRxiv*, pp. 2021–10, 2021.
- Sompolinsky, H., Crisanti, A., and Sommers, H.-J. Chaos in random neural networks. *Physical Review Letters*, 61(3):259, 1988.
- Steinmetz, N. A., Zatka-Haas, P., Carandini, M., and Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature*, 576(7786):266–273, 2019.
- Stevenson, I. H. and Kording, K. P. How advances in neural recording affect data analysis. *Nature Neuroscience*, 14(2):139–142, 2011.
- Thura, D., Cabana, J. F., Feghaly, A., and Cisek, P. Integrated neural dynamics of sensorimotor decisions and actions. *PLoS Biology*, 20, 12 2022.
- Tzen, B. and Raginsky, M. Neural stochastic differential equations: Deep latent Gaussian models in the diffusion limit. *arXiv preprint arXiv:1905.09883*, 2019a.
- Tzen, B. and Raginsky, M. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In *Conference on Learning Theory*, pp. 3084–3114. PMLR, 2019b.
- Urai, A. E., Doiron, B., Leifer, A. M., and Churchland, A. K. Large-scale neural recordings call for new insights to link brain and behavior. *Nature Neuroscience*, 25(1):11–19, 2022.
- Van Gelder, T. The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21(5):615–628, 1998.
- Versteeg, C., Sedler, A. R., McCart, J. D., and Pandarinath, C. Expressive dynamics models with nonlinear injective readouts enable reliable recovery of latent features from neural activity. *arXiv preprint arXiv:2309.06402*, 2023.
- Vyas, S., Golub, M. D., Sussillo, D., and Shenoy, K. V. Computation through neural population dynamics. *Annual Review of Neuroscience*, 43:249–275, 2020.
- Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–968, 2002.
- Wilson, H. R. and Cowan, J. D. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12(1):1–24, 1972.
- Winfree, A. T. *The Geometry of Biological Time*, volume 2. Springer, 1980.
- Wu, A., Roy, N. A., Keeley, S., and Pillow, J. W. Gaussian process based nonlinear latent structure discovery in multivariate spike train data. *Advances in Neural Information Processing Systems*, 30, 2017.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S., Shenoy, K. V., and Sahani, M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Advances in Neural Information Processing Systems*, 21, 2008.
- Zabihi, M., Oldehinkel, M., Wolfers, T., Frouin, V., Goyard, D., Loth, E., Charman, T., Tillmann, J., Banaschewski, T., Dumas, G., et al. Dissecting the heterogeneous cortical anatomy of autism spectrum disorder using normative models. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 4(6):567–578, 2019.
- Zeng, S., Graf, F., and Kwitt, R. Latent SDEs on homogeneous spaces. *arXiv preprint arXiv:2306.16248*, 2023.
- Zhang, X., Pu, Y., Kawamura, Y., Loza, A., Bengio, Y., Shung, D. L., and Tong, A. Trajectory flow matching with applications to clinical time series modeling. *arXiv preprint arXiv:2410.21154*, 2024.
- Zhou, D. and Wei, X.-X. Learning identifiable and interpretable latent models of high-dimensional neural activity using Pi-VAE. *Advances in Neural Information Processing Systems*, 33:7234–7247, 2020.

Appendix

A. Related Work

Latent variable models in neuroscience Our work builds on previous latent variable models which have been proposed to infer low-dimensional latent dynamics from single-trial data (Sahani, 1999; Yu et al., 2008; Macke et al., 2011; Wu et al., 2017; Pandarinath et al., 2018; Hurwitz et al., 2021b; Schimel et al., 2021; Kim et al., 2021). Most related to our work, is recent work by Hu et al. (2024) which leverage SDEs to represent latent neural dynamics. In their work, the authors developed a novel Gaussian process kernel function that defines a smooth, locally linear prior on dynamics, enabling it to act as an easily interpretable recurrent linear switching dynamical systems while providing posterior uncertainty estimates. A key advantage of our employed approach is the scalability and flexibility which is limited by the use of Gaussian processes. Relevant recent work also by Pals et al. (2024) extend low-rank RNNs to a stochastic setting with additive process noise. The authors use a specific form of the RNN which enables finding fixed-points effectively after training.

Scalable inference of latent SDEs Inferring the state and learning the parameters of a latent SDE from (high-dimensional) noisy observations is a challenging problem which appears through out science and engineering. With the advent of neural differential equations (Chen et al., 2018; Kidger, 2022), there has been a growing interest in developing scalable methods for fitting latent neural SDEs to data. Seminal work by Tzen & Raginsky (2019a;b) established the theoretical foundation for training generative Neural SDEs via variational inference through the lens of stochastic control. Subsequent research has focused on addressing key computational challenges, including improved memory efficiency and extension to the time domain (Li et al., 2020), enhanced numerical stability (Zeng et al., 2023; Oh et al., 2024), and computational efficiency (Course & Nair, 2024). Our work builds upon the formulation by Li et al. (2020), extending it to controlled and multi-modal settings in neuroscience. This specific formulation enables preserving the expressivity of the generative SDE through state- and input-dependent diffusion on arbitrary spaces.

Modeling neural dynamics with coupled oscillators Coupled oscillator models have a rich history in theoretical neuroscience, dating back to seminal work by Kuramoto (1975) and Wilson & Cowan (1972). These models have proven particularly effective in capturing rhythmic neural activity and investigating synchronization phenomena in neural circuits (Breakspear et al., 2010; Ermentrout & Terman, 2010). Of particular relevance to our work are modern approaches that combine oscillator models with machine learning techniques (Abrevaya et al., 2021; 2023; Bandyopadhyay et al., 2023) to learn dynamical systems that can reproduce complex neural trajectories while maintaining interpretable structure. Our approach builds upon these foundations by introducing a flexible SDE framework that can naturally capture oscillatory dynamics while accounting for process noise and external inputs. This bridges the gap between mechanistic oscillator models and data-driven approaches, allowing for both accurate trajectory prediction and meaningful scientific interpretation.

B. Learning the parameters of the generative model and inferring latent paths

Recall that our goal is to learn a generative model $p_\theta(y, b | u)$ via the following continuous-time state-space model:

$$dx(t) = \mu_\theta(x(t), u(t))dt + \sigma_\theta(x(t), u(t))dw(t), \quad t \in [0, \tau], \quad x_0 \sim \mathcal{N}(0, I) \quad (11)$$

$$y(t) \sim p(y(t) | \lambda_\theta(x(t))), \quad b(t) \sim p(b(t) | \rho_\theta(x(t))). \quad (12)$$

We leverage variational inference to train the model parameters θ in a tractable and scalable manner. The SDE induces a path measure \mathcal{P}_τ , which together with the initial distribution implicitly define our prior distribution $p_\theta(x|u)$ over continuous paths. To define the posterior distribution, we consider another continuous-time stochastic process defined via the following SDE:

$$dx(t) = \nu_\phi(x(t), y(t), b(t), u(t))dt + \sigma_\theta(x(t), u(t))dw(t), \quad t \in [0, \tau], \quad x_0 \sim \mathcal{N}(\alpha_\phi(y, b), \beta_\phi(y, b)) \quad (13)$$

This SDE induces a measure on the path space \mathcal{Q}_τ which together with the initial distribution implicitly define our posterior distribution $q_\phi(x|u, y, b)$ over continuous paths. The choice of using a similar diffusion function allows us to derive a tractable relationship between the two path measures induced via the two SDEs using Girsanov’s theorem. According to Girsanov’s theorem, the Radon-Nikodym derivative between these measures is given by:

$$\frac{d\mathcal{Q}_\tau}{d\mathcal{P}_\tau} = \exp\left(-\frac{1}{2} \int_0^\tau |\Delta(t)|^2 dt + \int_0^\tau \Delta(t)^T dw(t)\right) \quad (14)$$

where $\Delta(t) = \sigma_\theta(x(t), u(t))^{-1}(\nu_\phi(x(t), y(t), b(t), u(t)) - \mu_\theta(x(t), u(t)))$. We can thus derive a finite Kullback–Leibler divergence between the two path measures as:

$$D_{\text{KL}}(\mathcal{Q}_\tau || \mathcal{P}_\tau) = \frac{1}{2} \mathbb{E}_{\mathcal{Q}_\tau} \left[\int_0^\tau |\Delta(t)|^2 dt \right] \quad (15)$$

The evidence lower bound (ELBO) can then be written as:

$$\text{ELBO}(\theta, \phi) = \mathbb{E}_{q_\phi(x|u,y,b)} [\log p_\theta(y, b | x, u)] - D_{\text{KL}}(\mathcal{Q}_\tau || \mathcal{P}_\tau) - D_{\text{KL}}(\mathcal{N}(\alpha_\phi(y, b), \beta_\phi(y, b)) || \mathcal{N}(0, I)) \quad (16)$$

$$= \mathbb{E}_{q_\phi(x|u,y,b)} \left[\int_0^\tau \log p_\theta(y(t), b(t) | x(t), u(t)) dt \right] - \frac{1}{2} \mathbb{E}_{\mathcal{Q}_\tau} \left[\int_0^\tau |\Delta(t)|^2 dt \right] \quad (17)$$

$$- \frac{1}{2} (\text{tr}(\beta_\phi(y, b)) + |\alpha_\phi(y, b)|^2 - \log \det(\beta_\phi(y, b)) - d_x) \quad (18)$$

This optimization can be performed using stochastic gradient descent, where the path integrals are approximated using numerical integration schemes such as the Euler-Maruyama method, and the expectations are estimated using Monte Carlo sampling.

C. Implementation details

In the following, we summarize the general implementation details of our latent SDE. The models were trained to minimize ELBO on the training set by standard stochastic gradient backpropagation through the computational graph. This includes the numerical solver. That is, the intermediate steps are saved in memory during the forward pass and used to update the parameters during the backward pass. While this scales to $\mathcal{O}(n)$ in terms of memory requirement where n is the number of steps in the numerical solver compared to $\mathcal{O}(1)$ of the adjoint method, we found that this leads to more stable training in our experiments. We used an Euler-Maruyama solver for all our experiments. We found that setting a step size dt at 2 times the bin width of the spikes in our empirical experiments did not significantly influence the performance and helped improved the training speed and memory requirements.

Notably, we found that cyclic annealing the KL-divergence term in our ELBO significantly improved training stability. We follow the approach by Fu et al. (2019) and use a linear with 4 cycles with a linearly increasing value from 0 to 10 for half the cycle and then stays constant for the second half, until the cycle resets. In experiments which include reconstructing both neural and behavioral data, we use a constant weighting for each during training depending on the dataset. Our experiments were carried out on a NVIDIA A100-PCIE-40GB. Table 2 lists the main components of our latent SDE and their hyperparameters used across all empirical experiments.

Table 2. General hyperparameters and configuration settings used during training.

Component	Hyperparameter	Value
General	Batch size	128
	Optimizer	Adam
	Learning rate	1e-3
	Max n_epochs	2000
	Patience	20
	MC samples (training)	1
	MC samples (prediction)	30
Differential equation solver	KL annealing scheduler	Cyclic
	Solver	Euler-Maruyama
	Step size	$2(t_{i+1} - t_i)$
	Backpropagation	Discrete-adjoint (discretize-then-optimize)

D. Architecture

We summarize below the objective and parameters of each module in our proposed framework, detailing the configurations used in our empirical experiments.

Observation Encoder The observation encoder serves two purposes: (1) determining the parameters of the posterior distributions for the initial state, and (2) generating a lower-dimensional time-series representation (context) to condition the augmented SDE on the data. This is accomplished through two parallel components: the initialization network and the context network. The initialization network processes the first 30 observations through a 2-layer LSTM, producing a time-series of length 64. The context network processes the complete time-series of observations through a 2-layer LSTM, producing a time-series of length 64. This context vector undergoes linear interpolation to create a continuous-time representation for the augmented SDE.

Stimulus Encoder This module creates a lower-dimensional continuous representation of the stimulus. For low-dimensional stimuli (where $d_v \leq d_x$), we employ an identity transformation. We include time as an additional input channel. In scenarios without stimulus information (such as the maze reach experiment), we set $u(t) = t$. After encoding, we use linear interpolation to obtain a continuous-time representation of the stimulus.

Generative SDE The generative SDE defines our dynamics by mapping the current state and encoded stimulus to the change in the next state. The drift and diffusion structure is flexible and should incorporate prior knowledge when available. In this paper, we present a model combining coupled oscillators with neural networks to represent the SDE’s drift vector field. We also evaluate a neural SDE variant using a state size of 16 and a 2-layer MLP with hidden layer size 64 and tanh activation functions. The diffusion uses a sparse connection layer with non-zero diagonal parameters only, which is shared with the augmented SDE.

Augmented SDE The augmented SDE induces a path measure on continuous latent paths, implicitly defining our approximate posterior distribution. This neural SDE processes a concatenated vector of the state, encoded stimulus, and context at time t to determine the change in the next state. Our implementation uses a 2-layer MLP with hidden size 64 and tanh activation functions for the drift vector field. The diffusion function is identical to that of the generative SDE.

Decoder The decoder transforms latent states into observations. For systems with both neural and behavioral observations, we implement parallel neural and behavioral decoders. The neural decoder consists of an MLP with one hidden layer of size 128 and soft-plus activations. For K neural populations, we employ K similar MLPs, each with output size matching the neuron count in its corresponding population. These outputs’ exponents define the rate parameter of a time-inhomogeneous Poisson distribution. The behavioral decoder uses two parallel linear transformations on the latent states to generate the mean and squared variance of a Gaussian distribution at each time step, defining the predicted behavior distribution.