Langevin Flows for Modeling Neural Latent Dynamics

Yue Song[1], T. Anderson Keller[2], Yisong Yue[1], Pietro Perona[1], and Max Welling[3]

[1] Caltech, Vision Lab, Pasadena, CA
[2] The Kempner Institute for the Study of Natural and Artificial Intelligence, Harvard University, Cambridge, MA
[3] University of Amsterdam, Amsterdam, The Netherlands

Abstract

We propose LangevinFlow, a sequential variational model for neural population activity, where latent dynamics are governed by the underdamped Langevin equation. This framework captures both intrinsic neural dynamics and external unobserved inputs through physically grounded priors – incorporating inertia, damping, stochasticity, and a learned potential landscape. The potential is parameterized as a locally coupled oscillator network, biasing the model toward oscillatory and flow-like behaviors observed in real neural circuits. Our architecture combines a recurrent encoder, a one-layer Transformer decoder, and structured Langevin dynamics in the latent space. LangevinFlow achieves strong empirical results: it closely tracks ground-truth firing rates on synthetic data driven by a Lorenz attractor, and outperforms prior methods on the Neural Latents Benchmark across four datasets in terms of both bits-per-spike and forward prediction. It also matches or exceeds baselines in decoding behavioral variables such as hand velocity. This work introduces a compact, physics-inspired, interpretable, and high-performing model for neural population dynamics.

Keywords: neural population dynamics; variational autoencoders; latent variable models

Introduction

Neural population activity exhibits low-dimensional latent dynamics that govern spiking variability over time (Shenoy et al., 2013; Vyas et al., 2020). Modeling these structures is crucial for understanding internal computation and external unobserved inputs (Gallego et al., 2017). Prior work has explored dynamical systems with inferred control inputs (Pandarinath et al., 2018a), attractor landscapes (Genkin et al., 2023), and Transformer-based models (Ye & Pandarinath, 2021) to model both autonomous and externally driven neural processes.

We propose LangevinFlow, a sequential latent variable model governed by underdamped Langevin dynamics. This physically inspired formulation integrates key elements of neural computation: inertia, damping, stochasticity, and a potential function shaping attractor-like behavior. The potential is parameterized as a locally coupled oscillator network, biasing the model toward oscillatory and flow-like patterns consistent with cortical rhythms and traveling waves (Buzsaki, 2006; Churchland et al., 2012). The model uses a recurrent encoder and a one-layer Transformer decoder within a variational autoencoding framework. The encoder captures local temporal structure, while the Transformer integrates global latent context to refine spike rate predictions.

Empirically, LangevinFlow accurately reconstructs synthetic Lorenz-generated neural data and achieves state-of-the-art performance on the Neural Latents Benchmark (NLB) (Pei et al., 2021), outperforming baselines on neuron likelihoods and behavior decoding across four datasets. Latent trajectories reveal smooth spatiotemporal wave patterns, offering interpretable insights into neural computation. Our results demonstrate the utility of combining physics-inspired dynamics with modern sequence models for neural data analysis.

Methodology

This section first introduces the underdamped Langevin equation, then present the sequential VAE framework, followed by the model architecture and training algorithm.

Underdamped Langevin Equation

We seek to build a latent variable model which integrates the desired beneficial inductive biases. From the physics literature, a canonical abstract model is the Langevin equation:

$$\frac{\partial \boldsymbol{z}}{\partial t} = \boldsymbol{v},\ \ m\frac{\partial \boldsymbol{v}}{\partial t} = -\nabla_{\boldsymbol{z}} U(\boldsymbol{z}) - m\gamma\boldsymbol{v} + \sqrt{2m\gamma k_B \tau}\boldsymbol{\eta}(t) \tag{1}$$

where $\boldsymbol{z}(t)$ is the latent position, $\boldsymbol{v}$ denotes the velocity, m is a diagonal matrix of masses, γ is the damping (or friction) coefficient, k_B is the Boltzmann constant, τ is the temperature, $\boldsymbol{\eta}(t)$ represents Gaussian white noise, and the potential function is parameterized as a locally coupled oscillator network: $U(\boldsymbol{z}) = \boldsymbol{z}^T \frac{\boldsymbol{W}_z}{||\boldsymbol{W}_z||_2}\boldsymbol{z}$ where $\boldsymbol{W}_z \in \mathbb{R}^{d\times d}$ is the symmetric matrix of coupling coefficients. This choice biases the latent space toward smooth spatiotemporal oscillatory dynamics (Diamant & Bortoff, 1969; Ermentrout & Kopell, 1984).

Sequential Variational Auto-Encoder

The observed spikes $\bar{\boldsymbol{x}}$ are modeled as conditionally Poisson-distributed given latent firing rates $\bar{\boldsymbol{r}}$. The joint generative distribution factorizes as:

$$p(\bar{\boldsymbol{x}},\bar{\boldsymbol{z}},\bar{\boldsymbol{v}}) = p(\boldsymbol{v}_0)p(\boldsymbol{z}_0)\prod_{t=1}^{T} p(\boldsymbol{v}_t)\delta(\boldsymbol{z}_t - f_{\boldsymbol{z}}(\boldsymbol{z}_{t-1},\boldsymbol{v}_{t-1}))\prod_{t=0}^{T} p(\boldsymbol{x}_t|\boldsymbol{z}_t,\boldsymbol{v}_t) \tag{2}$$

where $\boldsymbol{z}_t$ is deterministically updated via Hamiltonian flow, and $\boldsymbol{v}_t$ is sampled using a stochastic transition. We optimize the evidence lower bound (ELBO) using a sequential VAE framework. The approximate posterior follows:

$$q_\theta(\bar{\boldsymbol{z}},\bar{\boldsymbol{v}}|\bar{\boldsymbol{x}}) = q(\boldsymbol{z}_0|\boldsymbol{x}_0)q(\boldsymbol{v}_0|\boldsymbol{x}_0)\prod_{t=1}^{T}\delta(\boldsymbol{z}_t - f_{\boldsymbol{z}}(\boldsymbol{z}_{t-1},\boldsymbol{v}_{t-1}))q(\boldsymbol{v}_t|\boldsymbol{z}_{t-1},\boldsymbol{v}_{t-1}) \tag{3}$$

Table 1: Results on `MC_Maze` and `MC_RTT` with the sampling frequency of 20 ms.

Methods	MC-Maze				MC-RTT		
	co-bps (↑)	vel R2 (↑)	psth R2 (↑)	fp-bps (↑)	co-bps (↑)	vel R2 (↑)	fp-bps (↑)
Smoothing (Yu et al., 2008)	0.2076	0.6111	-0.0005	–	0.1454	0.3875	–
GPFA (Yu et al., 2008)	0.2463	0.6613	0.5574	–	0.1769	0.5263	–
SLDS (Linderman et al., 2017)	0.2117	0.7944	0.4709	-0.1513	0.1662	0.5365	-0.0509
NDT (Ye & Pandarinath, 2021)	0.3597	0.8897	0.6172	0.2442	0.1643	0.6100	0.1200
AutoLFADS (Pandarinath et al., 2018b)	0.3554	0.8906	0.6002	0.2454	0.1976	0.6105	0.1241
MINT (Perkins et al., 2023)	0.3295	**0.9005**	**0.7474**	0.2076	0.2008	0.6547	0.1099
LangevinFlow	**0.3641**	0.8940	0.6801	**0.2573**	**0.2010**	**0.6652**	**0.1389**

Table 2: Results on `Area2_Bump` and `DMFC_RSG` with the sampling frequency of 20 ms.

Methods	Area2-Bump				DMFC-RSG			
	co-bps (↑)	vel R2 (↑)	psth R2 (↑)	fp-bps (↑)	co-bps (↑)	tp corr (↓)	psth R2 (↑)	fp-bps (↑)
Smoothing (Yu et al., 2008)	0.1529	0.5319	-0.1840	–	0.1183	-0.5115	0.2830	–
GPFA (Yu et al., 2008)	0.1791	0.6094	0.5998	–	0.1378	-0.5506	0.3180	–
SLDS (Linderman et al., 2017)	0.1816	0.6967	0.5200	0.0132	0.1575	-0.5997	0.5470	0.0374
NDT (Ye & Pandarinath, 2021)	0.2624	0.8623	0.6078	0.1459	0.1757	-0.6928	0.5477	0.1649
AutoLFADS (Pandarinath et al., 2018b)	0.2542	0.8565	0.6552	0.1423	0.1871	**-0.7819**	0.5903	0.1791
MINT (Perkins et al., 2023)	0.2718	0.8803	**0.9049**	0.1489	0.1824	-0.6995	**0.7014**	0.1647
LangevinFlow	**0.2881**	**0.8810**	0.7641	**0.1647**	**0.1904**	-0.5981	0.6079	**0.1945**

The velocity evolves via a learned Ornstein–Uhlenbeck transition. The ELBO objective includes a Poisson likelihood term and KL regularization on latent variables and their transitions:

$$\begin{aligned}\log p(\bar{\boldsymbol{x}}) \geq \sum_{t=0}^{T} \mathbb{E}_{q_\theta}\left[\log p(\boldsymbol{x}_t|\boldsymbol{z}_t,\boldsymbol{v}_t)\right] \\ -\mathbb{E}_{q_\theta}\left[\mathrm{D}_{\mathrm{KL}}\left[q_\theta(\boldsymbol{z}_0|\boldsymbol{x}_0)||p(\boldsymbol{z}_0)\right]\right] - \mathbb{E}_{q_\theta}\left[\mathrm{D}_{\mathrm{KL}}\left[q_\theta(\boldsymbol{v}_0|\boldsymbol{x}_0)||p(\boldsymbol{v}_0)\right]\right] \\ -\sum_{t=1}^{T} \mathbb{E}_{q_\theta}\left[\mathrm{D}_{\mathrm{KL}}\left[q_\theta(\boldsymbol{v}_t|\boldsymbol{z}_{t-1},\boldsymbol{v}_{t-1})||p(\boldsymbol{v}_t)\right]\right]\end{aligned} \tag{4}$$

Model Architecture and Training Algorithm

The encoder is a GRU that processes input spikes and initializes latent variables. Latents then evolve via Langevin dynamics. A Transformer-based decoder maps the latent trajectory to firing rate predictions, capturing both local and global temporal dependencies. Training proceeds by alternating between Langevin updates and optimization of the ELBO. The full algorithm is summarized in Alg. 1.

Table 3: R_2 of the firing rates on Lorenz Attractor.

	AutoLFADS	NDT	LangevinFlow
$R_2(\uparrow)$	0.921±0.005	0.934±0.004	**0.944±0.003**

Results

Table 3 shows that LangevinFlow achieves the highest R^2 correlation with ground-truth firing rates, outperforming all baselines. This indicates its superior ability to capture the underlying neural dynamics. In Tables 1 and 2, LangevinFlow achieves state-of-the-art performance on held-out neuron likelihood (co-smoothing bits per spike) and forward prediction. It also performs competitively on behavioral metrics such as hand-velocity decoding. While it lags slightly on PSTH R^2, this aligns with the known trade-off between co-bps and trial-averaged correlation. Overall, the model demonstrates strong and robust performance across diverse neural decoding tasks.

Algorithm 1 Training algorithm of our Langevin flow.

Require: Recurrent encoder `GRU`, Transformer-based sequence decoder `Transformer`, linear mapping for latent variables n, input spike sequence $\bar{\boldsymbol{x}}$, and posterior q_θ.

repeat
- Initial hidden states: $\boldsymbol{h}_0 = \mathtt{GRU}(\boldsymbol{x}_0)$
- Initial latent variables: $[\boldsymbol{z}_0, \boldsymbol{v}_0] = n(\boldsymbol{h}_0)$
- Time step counter: $i = 0$
- **while** $i \leq T-1$ **do**
  - Update position (deterministic step): $\boldsymbol{z}_{i+1} = \boldsymbol{z}_i + \boldsymbol{v}_i$
  - Update velocity (deterministic step): $\boldsymbol{v}_{i+\frac{1}{2}} = \boldsymbol{v}_i - {}^{\nabla_z U(\boldsymbol{z}_i)}\!/_{m}$
  - Update velocity (probabilistic step): $\boldsymbol{v}_{i+1} = (1-\gamma)\boldsymbol{v}_{i+\frac{1}{2}} + \sqrt{2m\gamma k_B \tau}\, \boldsymbol{\eta}(i)$
  - Update hidden states: $\boldsymbol{h}_{i+1} = \mathtt{GRU}(\boldsymbol{x}_{i+1}, \boldsymbol{h}_i)$
  - Concatenate variable sequences: $\bar{\boldsymbol{z}} = [\boldsymbol{z}_{0:i}, \boldsymbol{z}_{i+1}],\ \bar{\boldsymbol{v}} = [\boldsymbol{v}_{0:i}, \boldsymbol{v}_{i+1}],\ \bar{\boldsymbol{h}} = [\boldsymbol{h}_{0:i}, \boldsymbol{h}_{i+1}]$
  - Update time step counter: $i = i+1$
- **end while**
- Predict firing rates: $\bar{\boldsymbol{r}} = \mathtt{Transformer}(\bar{\boldsymbol{z}}, \bar{\boldsymbol{v}}, \bar{\boldsymbol{h}})$
- Optimize the ELBO.

until converged

References

Buzsaki, G. (2006). *Rhythms of the brain*. Oxford university press.

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, *487*(7405), 51–56.

Diamant, N., & Bortoff, A. (1969, February). Nature of the intestinal low-wave frequency gradient. *American Journal of Physiology-Legacy Content*, *216*(2), 301–307. doi: 10.1152/ajplegacy.1969.216.2.301

Ermentrout, G. B., & Kopell, N. (1984). Frequency plateaus in a chain of weakly coupled oscillators, i. *SIAM journal on Mathematical Analysis*, *15*(2), 215–237.

Gallego, J. A., Perich, M. G., Miller, L. E., & Solla, S. A. (2017). Neural manifolds for the control of movement. *Neuron*, *94*(5), 978-984. doi: https://doi.org/10.1016/j.neuron.2017.05.025

Genkin, M., Shenoy, K. V., Chandrasekaran, C., & Engel, T. A. (2023). The dynamics and geometry of choice in premotor cortex. *bioRxiv*. doi: 10.1101/2023.07.22.550183

Linderman, S., Johnson, M., Miller, A., Adams, R., Blei, D., & Paninski, L. (2017). Bayesian learning and inference in recurrent switching linear dynamical systems. In *Aistats.*

Pandarinath, C., O'Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., ... others (2018a). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature methods*.

Pandarinath, C., O'Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., ... others (2018b). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature methods*.

Pei, F., Ye, J., Zoltowski, D., Wu, A., Chowdhury, R. H., Sohn, H., ... others (2021). Neural latents benchmark'21: evaluating latent variable models of neural population activity. *NeurIPS*.

Perkins, S. M., Cunningham, J. P., Wang, Q., & Churchland, M. M. (2023). Simple decoding of behavior from a complicated neural manifold. *BioRxiv*, 2023–04.

Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annual review of neuroscience*, *36*(1), 337–359.

Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation through neural population dynamics. *Annual review of neuroscience*, *43*(1), 249–275.

Ye, J., & Pandarinath, C. (2021). Representation learning for neural population activity with neural data transformers. *Neurons, Behavior, Data analysis, and Theory*.

Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S., Shenoy, K. V., & Sahani, M. (2008). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *NeurIPS*.