SynapsNet: Enhancing Neuronal Population Dynamics Modeling via Learning Functional Connectivity

Parsa Delavari
Neuroscience
University of British Columbia
parsadlr@student.ubc.ca

Ipek Oruc
Ophthalmology and Visual Sciences
University of British Columbia
ipor@mail.ubc.ca

Timothy H. Murphy
Psychiatry
University of British Columbia
thmurphy@mail.ubc.ca

Abstract

The availability of large-scale neuronal population datasets necessitates new methods to model population dynamics and extract interpretable, scientifically translatable insights. While deep learning offers potential solutions, existing models often overlook the biological mechanisms underlying population activity and thus exhibit suboptimal performance with neural data and provide little to no interpretable information about neurons and their interactions. In response, we introduce SYNAPSNET, a novel deep-learning framework that effectively models population dynamics and the interactions between neurons. Within this biologically realistic framework, each neuron, characterized by a latent embedding, sends and receives currents through directed connections. A shared decoder uses the input current, previous neuronal activity, neuron embedding, and behavioral data to predict the population activity in the next time step. Our experiments demonstrate that SYNAPSNET consistently outperforms existing models in forecasting population activity. Additionally, our experiments on synthetic data showed that SYNAPSNET accurately recovers ground-truth connections between neurons.

1 Introduction

Recent advancements in brain recording techniques have enabled simultaneous in-vivo recordings from hundreds of neurons. This availability of neuronal population activity has motivated numerous studies on population dynamics, which can address neuroscientific questions such as explaining brain function and behavior [1, 2], as well as brain decoding applications like brain-computer interfaces (BCIs) [3, 4]. It is worth noting that by "population activity," we specifically refer to the activity of a population of neurons recorded at single-cell resolution, represented as a vector of individual neuron activities. Various approaches for studying neural dynamics have been proposed in the literature, including neural manifolds [5] and latent variable models (LVMs) [6, 7], which aim to capture low-dimensional patterns in neural activity by reproducing this activity. However, the current models often face significant limitations, particularly in neuroscience applications. Firstly, they typically offer minimal interpretability, providing little insight into the brain's underlying mechanisms. In contrast, explainable deep learning models have demonstrated the potential to combine computational power with interpretability, enabling the analysis of large datasets while providing insights into the model's computation, which can even lead to scientific discovery [8, 9]. Secondly, many deep learning approaches neglect the underlying biological mechanisms [10], limiting their performance by

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

applying generic architectures adapted from other domains to neural data. These limitations highlight the need for more biologically informed and interpretable models of neuronal population dynamics.

In this work, we introduce SYNAPSNET, a biologically inspired deep learning framework designed to model neuronal populations while uncovering functional connectivity. It integrates past neural activity, input currents from connected neurons, intrinsic neuron properties, and behavioral states to predict neural dynamics. Building on previous work [11], which models individual neuron dynamics, SYNAPSNET goes further by also learning functional connectivity to capture directed interactions between neurons. We validate SYNAPSNET through experiments on both synthetic and real datasets of mouse cortical activity recorded via the most common population recording methods, calcium imaging and Neuropixels [12]. Our experiments demonstrate that SYNAPSNET outperforms conventional time-series models in predicting neural activity showcasing its high capability of capturing population dynamics. Also, SYNAPSNET provides insights into the underlying interactions between neurons by accurately inferring functional connectivity.

2 SYNAPSNET

2.1 Overview

The aim of SYNAPSNET is to develop an interpretable model of neuronal populations capable of predicting future dynamics based on the current and past states of the brain and behavior. To this end, our model employs a biologically inspired framework in which a dynamical model predicts the future activity of each neuron by taking: ❶ its past activity, ❷ the input currents coming from the connected neurons, ❸ the intrinsic properties of the target neuron, and ❹ the animal's behavioral variables (e.g., running speed). We use these components as previous research confirms their role in neuronal responses. The Hodgkin-Huxley equations [13], a foundational work in computational neuroscience, describe a neuron's membrane potential dynamics as a function of its current voltage, input currents from presynaptic neurons, and a set of neuron-specific physiological properties. Additionally, behavior has been shown to significantly impact neural responses, with the first principal components of single-trial population activity corresponding to behavioral variables such as running speed and pupil size in various brain regions, including the visual cortex [14].

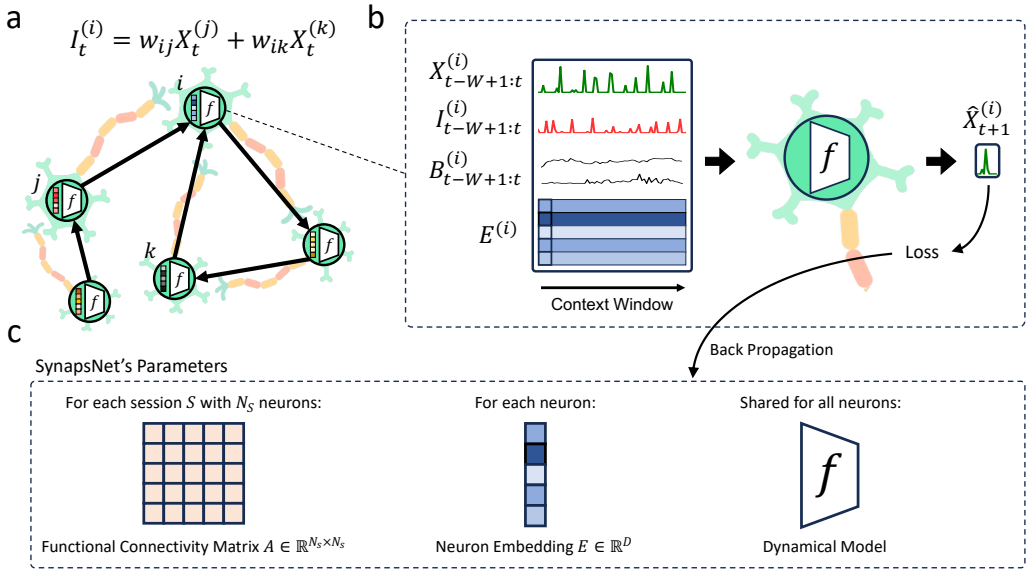

Figure 1: SYNAPSNET overview. (a) How input current is inferred based on functional connectivity and population activity (b) An example input frame to the dynamical model which includes past activity over the context window, past input current, past behavioral data, and the unique embedding of the target neuron. (c) The three sets of parameters in SYNAPSNET.

Although the Hodgkin-Huxley point-neuron model can accurately predict neuron dynamics, it requires detailed information on the activity of all presynaptic neurons, their synaptic strengths, and an exhaustive physiological description of the neurons. This level of detail is nearly impossible to obtain with current population recording techniques, which typically capture only a limited subset of neurons and provide minimal to no information on the population's physiological and anatomical properties. In this context, functional connectivity has been used as an alternative to anatomical connectivity for modeling interactions between neurons [15]. These functional interactions, which can be inferred from population activity, have been shown to predict neural responses even when only a small number of neurons are recorded [16]. Therefore, we estimate each neuron's input current based on the activity of other neurons and the population's functional connectivity. To address the lack of detailed physiological features of neurons, we use NeuPRINT [11] to learn time-invariant embeddings for each neuron based solely on population activity. The innovative approach proposed in NeuPRINT removes the dependency on neuron-specific characteristics from the dynamical model, placing it in the input so that the dynamical model becomes independent of intrinsic properties and can be shared across all neurons.

2.2 Implementation

We represent population datasets as a set of recording sessions $\mathbb{S}$, where each session $S_j \in \mathbb{S}$ includes a set of recorded neurons $\mathcal{N}_j$, population activity over time $\mathbf{X}_j \in \mathbb{R}^{|\mathcal{N}_j| \times T}$, and available behavioral data over time $\mathbf{B}_j \in \mathbb{R}^{N_B \times T}$. Let $M = |\bigcup_{j=1}^{|\mathbb{S}|} \mathcal{N}_j|$ be the total number of unique neurons across all recording sessions, acknowledging that some neurons may appear in multiple sessions. We define three sets of learnable parameters in SYNAPSNET (Figure 1c): (1) neuron embeddings $\mathbf{E} \in \mathbb{R}^{M \times D}$, where D is the dimensionality of the embedding vectors, (2) functional connectivity matrices $\{\mathbf{A}_j\}_{j=1}^{|\mathbb{S}|}$, where $\mathbf{A}_j \in \mathbb{R}^{|\mathcal{N}_j| \times |\mathcal{N}_j|}$, and (3) a dynamical model $f(.)$ shared among all neurons across all sessions. For simplicity, we drop the session number subscript j in the following descriptions and provide expressions for a single session.

Inferred input current calculation. Using population activity and functional connectivity, we calculate the input current to neuron i at time t as $I_t^{(i)} = \sum_{k=1}^{|\mathcal{N}|} A^{(i,k)} X_t^{(k)}$ (Figure 1a).

Dynamical model. We express the neuronal dynamics as

$$\hat{X}_{t+1}^{(i)} = f(X_{t-W+1:t}^{(i)}, I_{t-W+1:t}^{(i)}, B_{t-W+1:t}, E^{(i)})$$

where W is the context window that determines the history considered by the dynamical model (Figure 1b). The dynamical model f can be any sequential model designed for processing time-series data. The results presented in this paper are obtained using a GRU [17] (for results with other sequential models, see the supplementary materials).

Training. We formulate an optimization problem using the loss $\mathcal{L}$, calculated based on the ground truth and predicted activity, to learn the three sets of parameters in SYNAPSNET:

$$\min_{f, \{A_j\}, E} \left(\mathbb{E}_X \left[\mathcal{L}(\hat{X}_{t+1}^{(i)}, X_{t+1}^{(i)}) \right] + \lambda \sum_j ||A_j||_2 \right)$$

where $\mathcal{L}(.)$ is mean squared error (MSE), λ is the regularization weight for preventing overfitting through $\{A_j\}_j$ and $||\ .\ ||_2$ denotes L2 norm. For more details about the training procedure and the hyperparameters used, please refer to the supplementary materials.

3 Experiments

3.1 Population Activity Forecasting

Datasets and Preprocessing. We evaluated SYNAPSNET's performance in forecasting neuronal population activity using two distinct public datasets: [18] and [19]. The dataset from [18] consists

Table 1: Performance on neural data forecasting. Mean correlation and loss (%) ± standard error of the mean[†].

Data Modality	Model	Natural Scenes		Drifting Gratings		Spontaneous	
		$Corr(\%)\uparrow$	$Loss\downarrow$	$Corr(\%)\uparrow$	$Loss\downarrow$	$Corr(\%)\uparrow$	$Loss\downarrow$
Ca Imaging	RNN	$24.09_{\pm 0.88}$	$0.920_{\pm 0.040}$	$19.95_{\pm 0.81}$	$1.017_{\pm 0.050}$	$13.15_{\pm 0.75}$	$0.924_{\pm 0.017}$
	GRU	$24.21_{\pm 0.87}$	$0.915_{\pm 0.040}$	$20.08_{\pm 0.89}$	$1.009_{\pm 0.049}$	$13.31_{\pm 0.85}$	$0.920_{\pm 0.017}$
	LSTM	$24.60_{\pm 0.87}$	$0.911_{\pm 0.040}$	$20.47_{\pm 0.85}$	$1.006_{\pm 0.049}$	$13.97_{\pm 0.90}$	$0.917_{\pm 0.017}$
	LFADS	$25.70_{\pm 0.73}$	$1.056_{\pm 0.040}$	$22.96_{\pm 0.62}$	$1.138_{\pm 0.049}$	$12.33_{\pm 0.79}$	$0.941_{\pm 0.017}$
	GWNET	$33.04_{\pm 0.87}$	$0.829_{\pm 0.040}$	$29.89_{\pm 0.85}$	$0.910_{\pm 0.049}$	$28.20_{\pm 0.90}$	$0.874_{\pm 0.017}$
	NEUPRINT	$33.82_{\pm 0.85}$	$0.884_{\pm 0.039}$	$31.17_{\pm 1.00}$	$0.973_{\pm 0.052}$	$28.82_{\pm 1.25}$	$0.867_{\pm 0.015}$
	SYNAPSNET	$\mathbf{37.43_{\pm 1.05}}$	$\mathbf{0.846_{\pm 0.037}}$	$\mathbf{36.14_{\pm 1.05}}$	$\mathbf{0.927_{\pm 0.48}}$	$\mathbf{30.60_{\pm 1.29}}$	$\mathbf{0.855_{\pm 0.016}}$
	NEUPRINT (multi-session)	$34.16_{\pm 0.91}$	$0.865_{\pm 0.038}$	$31.64_{\pm 1.38}$	$1.045_{\pm 0.014}$	$29.25_{\pm 1.43}$	$0.903_{\pm 0.014}$
	SYNAPSNET (multi-session)	$\mathbf{37.94_{\pm 1.15}}$	$\mathbf{0.825_{\pm 0.035}}$	$\mathbf{36.79_{\pm 1.62}}$	$\mathbf{0.992_{\pm 0.015}}$	$\mathbf{31.60_{\pm 1.58}}$	$\mathbf{0.886_{\pm 0.015}}$
NeuroPixels	RNN	$18.50_{\pm 0.30}$	$0.963_{\pm 0.012}$	$20.20_{\pm 0.35}$	$0.992_{\pm 0.032}$	$16.67_{\pm 0.44}$	$0.960_{\pm 0.015}$
	GRU	$18.11_{\pm 0.32}$	$0.963_{\pm 0.012}$	$20.33_{\pm 0.37}$	$0.989_{\pm 0.032}$	$16.68_{\pm 0.47}$	$0.989_{\pm 0.015}$
	LSTM	$18.02_{\pm 0.33}$	$0.963_{\pm 0.012}$	$20.39_{\pm 0.39}$	$0.987_{\pm 0.031}$	$16.74_{\pm 0.47}$	$0.987_{\pm 0.015}$
	LFADS	$18.26_{\pm 0.73}$	$0.987_{\pm 0.040}$	$20.65_{\pm 0.62}$	$0.953_{\pm 0.049}$	$17.38_{\pm 0.79}$	$0.927_{\pm 0.017}$
	GWNET	$21.10_{\pm 0.87}$	$0.945_{\pm 0.040}$	$22.54_{\pm 0.85}$	$0.935_{\pm 0.049}$	$20.21_{\pm 0.90}$	$0.941_{\pm 0.017}$
	NEUPRINT	$21.68_{\pm 0.21}$	$0.942_{\pm 0.011}$	$22.93_{\pm 0.25}$	$0.967_{\pm 0.029}$	$20.79_{\pm 0.48}$	$0.967_{\pm 0.014}$
	SYNAPSNET	$\mathbf{24..38_{\pm 0.28}}$	$\mathbf{0.926_{\pm 0.011}}$	$\mathbf{25.55_{\pm 0.31}}$	$\mathbf{0.953_{\pm 0.029}}$	$\mathbf{22.51_{\pm 0.53}}$	$\mathbf{0.953_{\pm 0.014}}$
	NEUPRINT (multi-session)	$22.10_{\pm 0.19}$	$0.951_{\pm 0.012}$	$23.49_{\pm 0.31}$	$0.941_{\pm 0.011}$	$21.31_{\pm 0.42}$	$0.919_{\pm 0.014}$
	SYNAPSNET (multi-session)	$\mathbf{25.48_{\pm 0.23}}$	$\mathbf{0.932_{\pm 0.012}}$	$\mathbf{26.80_{\pm 0.031}}$	$\mathbf{0.921_{\pm 0.011}}$	$\mathbf{23.83_{\pm 0.44}}$	$\mathbf{0.905_{\pm 0.014}}$

[†] The bold values correspond to the best performance metric separately for multi-/single-session training and each data modality.
Blue cells indicate significant difference with the next best-performing model (p-value ≤ 0.05 achieved by paired-t-tests).

of two-photon calcium imaging data recorded from the mouse primary visual cortex (V1). Mice were exposed to three types of visual stimuli: natural scenes, drifting gratings, and blank screens (spontaneous activity). This dataset includes calcium traces from four mice, recorded across 17 sessions of approximately 20 minutes each. Each session captured between 178 and 868 neurons, totaling 9728 neurons. Neurons were recorded at a 4.3 Hz sampling rate on 7 imaging planes at different cortical depths, resulting in volumetric recordings. The dataset also provides neuron positions, cell types, and behavioral data such as running speed and pupil size. We normalized each neuron's calcium trace and behavioral variables. The dataset from [19] includes electrophysiology spiking data from mouse cortical neurons, recorded using six Neuropixels probes [12]. This dataset primarily covers V1 and higher visual areas, as well as deep subcortical areas like the LGN and hippocampus CA1. It features two session types: "brain observatory" sessions for natural scenes and drifting gratings, and "functional connectivity" sessions for spontaneous activity (blank screen). We included sessions where neuron locations were available, resulting in 46 sessions (46 animals) and 31,408 recorded neurons. The dataset also provides running speed data. We binned the spikes with a 33.3 ms bin width and normalized each neuron's response and behavioral data.

Benchmark Models. We compare SYNAPSNET with the following existing models: standard sequential models (RNN [20], GRU [17], LSTM [21]); NeuPRINT [11], a self-supervised method of learning neuronal representations based on population dynamics; LFADS[6], an RNN based variational auto-encoder designed specifically for neuronal population dynamics; GWNet [22], a graph neural network designed to model both spatial interactions and temporal patterns in data.

Results. We assessed SYNAPSNET and the benchmark models on two different data modalities across three tasks. We measured test MSE loss and Pearson's correlation between true and predicted population activity, averaging over all times for each neuron and across all neurons in each recording session. The means and standard errors were calculated across sessions (Ca imaging: $n = 17$; Neuropixels natural scenes and drifting gratings: $n = 22$; Neuropixels spontaneous: $n = 24$). Performance measures are summarized in Table 1. SYNAPSNET consistently surpassed all benchmark models across both data modalities and all three tasks, demonstrating its superior capability in modeling neuronal population dynamics. Both NeuPRINT and SYNAPSNET demonstrated improved performance when trained on multiple sessions, indicating their scalability. SYNAPSNET achieved approximately 15% higher correlation scores compared to the general models and 5% higher compared to NeuPRINT on calcium imaging data, translating to around 80% and 15% relative improvements, respectively. For the Neuropixels modality, these values were about 8% and 3%, corresponding to approximately 30% and 15% relative improvements.

3.2 Evaluation of Learned Functional Connectivity on Synthetic Data

Simulation. To further validate SYNAPSNET's ability to recover functional interactions, we simulated the activity of a neuron population with known ground-truth connectivity (Figure 2a). We simulated 600 interconnected neurons with connection weights randomly initialized to create a sparse connectivity matrix A[1]. We generated four random variables, two representing behavioral data (B) and two representing task variables (V). Each neuron's firing rate was calculated as: $R_t^{(i)} = \sigma\left(\bar{R}^{(i)} + w_1 \sum_k A^{(i,k)} X_{t-1}^{(k)} + w_2 B_{t-1} + w_3 V_{t-1} - w_4 X_{t-1}^{(i)} + n_t\right)$ where $\sigma(.)$ is the non-linearity ($tanh$ here), $\bar{R}$ is the neuron's mean firing rate, n_t is noise. Parameters w_1 to w_4 control the effect of input current, behavior, task, and self-inhibition respectively, and are randomly initialized for each neuron. Then, the activity of each neuron is determined by a Poisson process: $X_t^{(i)} \sim Poisson(R_t^{(i)})$. To mimic practical recording conditions, a subset of 200 neurons is randomly chosen as "recorded" neurons. This approach reflects the real-world scenario where only a fraction of the neurons in a population can be recorded, leaving the activity of many neurons inaccessible. Consequently, input currents to recorded neurons often originate from both recorded and unrecorded neurons, better approximating the incomplete knowledge of neuronal activity typically encountered in experimental settings. Note that the model has access to the behavioral data but not the task variables. Similar to real data, population activity and behavioral variables are normalized.

Benchmark Methods for Functional Connectivity. We compared SYNAPSNET with two benchmark methods for inferring functional connectivity: pair-wise correlation and CURBD [23]. Pair-wise correlation is a popular statistical method for calculating functional connectivity in neuroscience across various recording modalities. We used variants of this method with delays ($D = 0, 1, 2$) to capture causal interactions and achieve directed FCs. CURBD is an RNN-based model designed to infer inter-region currents by learning functional connectivity in a network of interconnected neurons capable of reproducing real data.

Performance on Recovering Ground-Truth Functional Connectivity. Figure 2b shows the ground-truth FC alongside the inferred FC by each method in a single simulation run, while Figure 2c compares the mean accuracy over 50 runs. SYNAPSNET reconstructed a sparse and asymmetric FC with over 80% accuracy. CURBD, despite being trained to nearly perfectly reproduce the data (explained variance = 0.93), failed to estimate the ground-truth FC. This likely results from CURBD's RNN overfitting to the data without generalizing the learned FC to a held-out test set. Pair-wise correlation with $D = 1$ partially recovered the true FC, achieving a mean accuracy of 30%. Notably, SYNAPSNET successfully inferred the ground-truth functional connectivity from synthetic data while its performance in predicting population activity was not perfect, with a correlation score of 36% and a test loss of 0.83, matching its performance on real data. This suggests that SYNAPSNET's strength in inferring FC is not constrained by the challenges of predicting highly stochastic spiking activity.

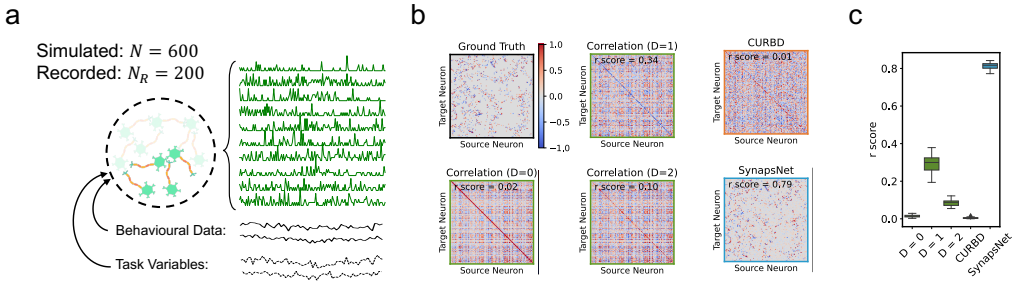

Figure 2: Synthetic data experiment. (a) Simulation process. (b) Functional connectivity matrices inferred using SYNAPSNET and other baselines compared with the ground-truth in a single simulation run. (c) Functional connectivity reconstruction accuracies achieved by each method over 50 runs of simulation with random initializations.

[1]For better visual comparison, the locations of non-zero elements in the connectivity matrix were based on an image from the Mandelbrot set fractal to include visual patterns in the ground-truth adjacency matrix.

4 Limitations

SYNAPSNET is specifically designed to model neuronal populations, which may limit its applicability to other types of neural data. Additional research is necessary to assess its generalizability to modalities where each channel does not correspond to an individual neuron, such as electrocorticography (ECoG). Also, since SYNAPSNET encodes the activity of each neuron separately, architectural modifications are required for tasks other than population activity forecasting, such as neural decoding. While our experiments demonstrated promising results in analyzing the functional connectivity and input currents inferred by SYNAPSNET on both real and synthetic data, fully validating the learned connections remains challenging due to the lack of ground-truth anatomical connections in population activity datasets. Therefore, the anatomical translation of SYNAPSNET's functional connectivity requires future validation.

5 Conclusions

We introduced SYNAPSNET, a biologically inspired deep learning framework that advances neuronal population modeling by combining high predictive performance with interpretability. Our approach addresses the limitations of existing deep learning methods, which often fail to account for the biological mechanisms underlying neuronal activity and offer limited interpretability. We showed that SYNAPSNET consistently surpassed conventional time-series models, neural data-specific models, and graph neural networks in capturing population dynamics. Additionally, experiments on synthetic data revealed SYNAPSNET's ability to accurately identify underlying connectivity, offering biologically meaningful insights into neuron interactions and inferring the currents flowing between them.

References

[1] Saurabh Vyas, Matthew D Golub, David Sussillo, and Krishna V Shenoy. Computation through neural population dynamics. *Annual review of neuroscience*, 43:249–275, 2020.

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

[3] Francis R Willett, Erin M Kunz, Chaofei Fan, Donald T Avansino, Guy H Wilson, Eun Young Choi, Foram Kamdar, Matthew F Glasser, Leigh R Hochberg, Shaul Druckmann, et al. A high-performance speech neuroprosthesis. *Nature*, 620(7976):1031–1036, 2023.

[4] Christopher Heelan, Jihun Lee, Ronan O'Shea, Laurie Lynch, David M Brandman, Wilson Truccolo, and Arto V Nurmikko. Decoding speech from spike-based neural population recordings in secondary auditory cortex of non-human primates. *Communications biology*, 2(1):466, 2019.

[5] Lea Duncker and Maneesh Sahani. Dynamics on the manifold: Identifying computational dynamical activity from neural population recordings. *Current opinion in neurobiology*, 70:163–170, 2021.

[6] Chethan Pandarinath, Daniel J O'Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D Stavisky, Jonathan C Kao, Eric M Trautmann, Matthew T Kaufman, Stephen I Ryu, Leigh R Hochberg, et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature methods*, 15(10):805–815, 2018.

[7] Michael Nolan, Bijan Pesaran, Eli Shlizerman, and Amy Orsborn. Multi-block rnn autoencoders enable broadband ecog signal reconstruction. *bioRxiv*, pages 2022–09, 2022.

[8] Pablo Lemos, Niall Jeffrey, Miles Cranmer, Shirley Ho, and Peter Battaglia. Rediscovering orbital mechanics with machine learning. *Machine Learning: Science and Technology*, 4(4):045002, 2023.

[9] Parsa Delavari, Gulcenur Ozturan, Lei Yuan, Özgür Yilmaz, and Ipek Oruc. Artificial intelligence, explainability, and the scientific method: A proof-of-concept study on novel retinal biomarker discovery. *PNAS nexus*, 2(9):pgad290, 2023.

[10] Tim C Kietzmann, Patrick McClure, and Nikolaus Kriegeskorte. Deep neural networks in computational neuroscience. *BioRxiv*, page 133504, 2017.

[11] Lu Mi, Trung Le, Tianxing He, Eli Shlizerman, and Uygar Sümbül. Learning time-invariant representations for individual neurons from population dynamics. *Advances in Neural Information Processing Systems*, 36, 2024.

[12] James J Jun, Nicholas A Steinmetz, Joshua H Siegle, Daniel J Denman, Marius Bauza, Brian Barbarits, Albert K Lee, Costas A Anastassiou, Alexandru Andrei, Çağatay Aydın, et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236, 2017.

[13] Allan L Hodgkin and Andrew F Huxley. Currents carried by sodium and potassium ions through the membrane of the giant axon of loligo. *The Journal of physiology*, 116(4):449, 1952.

[14] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):eaav7893, 2019.

[15] Daniele Poli, Vito P Pastore, and Paolo Massobrio. Functional connectivity in in vitro neuronal assemblies. *Frontiers in neural circuits*, 9:57, 2015.

[16] Ian H Stevenson, Brian M London, Emily R Oby, Nicholas A Sachs, Jacob Reimer, Bernhard Englitz, Stephen V David, Shihab A Shamma, Timothy J Blanche, Kenji Mizuseki, et al. Functional connectivity and tuning curves in populations of simultaneously recorded neurons. *PLoS computational biology*, 8(11):e1002775, 2012.

[17] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.

[18] Stephane Bugeon, Joshua Duffield, Mario Dipoppa, Anne Ritoux, Isabelle Prankerd, Dimitris Nicoloutsopoulos, David Orme, Maxwell Shinn, Han Peng, Hamish Forrest, et al. A transcriptomic axis predicts state modulation of cortical interneurons. *Nature*, 607(7918):330–338, 2022.

[19] Joshua H Siegle, Xiaoxuan Jia, Séverine Durand, Sam Gale, Corbett Bennett, Nile Graddis, Greggory Heller, Tamina K Ramirez, Hannah Choi, Jennifer A Luviano, et al. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, 592(7852):86–92, 2021.

[20] Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.

[21] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[22] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep spatial-temporal graph modeling. *arXiv preprint arXiv:1906.00121*, 2019.

[23] Matthew G Perich, Charlotte Arlt, Sofia Soares, Megan E Young, Clayton P Mosher, Juri Minxha, Eugene Carter, Ueli Rutishauser, Peter H Rudebeck, Christopher D Harvey, et al. Inferring brain-wide interactions using data-constrained recurrent neural network models. *BioRxiv*, pages 2020–12, 2020.

A Methods Details

A.1 Training Details and Hyperparameters

For each session of data, we split the population activity into three continuous partitions: approximately 80% for training, 10% for validation, and 10% for testing. The validation set was used to select the epoch with the best performance and to tune hyperparameters, while the test set was reserved for the final evaluation of the models and reporting performance metrics.

For SynapsNet, we used context window size $W = 5$, neuron embedding size of 32, and a single-layer GRU with hidden layer size of 40 as the dynamical model. During training, the training set is divided into identical (but overlapping) partitions of size $W + 1$, where the first W time points are the input and the last time point is the target. For batch sampling, we first select a random session from which all batch samples are chosen. This approach ensures that all samples in a batch have identical dimensionality, simplifying implementation, as different sessions contain different numbers of neurons.

For all datasets, we use a batch size of 32 and an initial learning rate of 10^{-3}. The models are trained for 100 epochs, with the learning rate halved every 10 epochs. Additionally, dropout layers with a rate of 0.1 are employed in all neural networks. SYNAPSNET is implemented using `PyTorch` in `Python`, and evaluations are run on a Linux machine with a GPU and 16GB of RAM. Training SYNAPSNET on each session of data (single-session training) takes approximately half an hour, depending on the number of recorded neurons and the session length.

For the sake of fair comparison, we use the same training, validation, and testing data for all benchmark models and ablated versions of SYNAPSNET.

A.2 Models Used for Comparison

The RNN, GRU, and LSTM models each consist of 2 layers with a hidden layer size of 100. The LFADS model we used has the dimensionality of the generator, encoder, and controller all set to 256, and the dimensionality of the factor and inferred inputs to the generator set to 128. The GWNet model is defined with 32 residual channels, 32 dilation channels, 128 skip channels, and 256 end channels, with a single layer, 4 blocks, and a kernel size of 2.

B Additional Results

B.1 Evaluation of SYNAPSNET's forecasting performance at the population level

We also assessed the performance of SYNAPSNET at the population level. Figure 3 compares the first three principal components (PCs) of the true and predicted population activity. Across all sessions, SYNAPSNET's predictions achieved mean correlation scores of approximately 85%, 75%, and 65% for the first three PCs, respectively. These correlations were significantly higher than those achieved by NeuPRINT.

B.2 Evaluating input currents inferred by SYNAPSNET

Figure 4a presents the functional connectivity (FC) matrices inferred by SYNAPSNET alongside the pair-wise correlations between neurons in two example sessions—one from calcium imaging and the other from the Neuropixels dataset. SYNAPSNET's inferred FC matrices appear sparser, more asymmetric, and more structured compared to the correlation matrices. Figure 4b offers a 3D visualization of neurons and the learned connections, revealing clear patterns of connection types and strengths based on neuron locations, even though neuron positions were not provided to the model during training.

We analyzed the input currents inferred by SYNAPSNET, which are latent variables used to predict future neural activity. As described in §2.2, we calculated input currents based on population activity and functional connectivity within a recording session using the equation $I = AX$ (Figure 5a). We then measured the cross-correlation (correlation as a function of relative delay) between the input current to each neuron and its activity. We employed three different methods to infer functional connectivity: SYNAPSNET, pair-wise correlation, and shuffled SYNAPSNET's FC. As illustrated

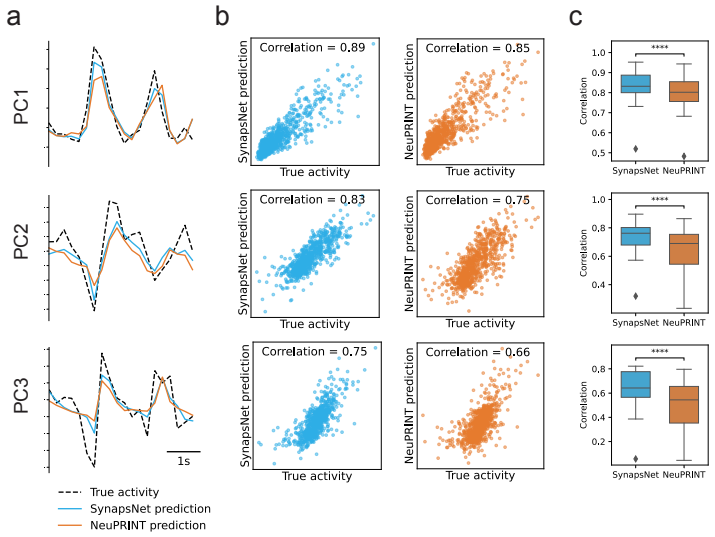

Figure 3: Performance on neural data forecasting. (a) The first three principal components (PCs) of the true and predicted population activity sampled from an example session. (b) Correlation between the first three PCs of the true and predicted activity for the all-time points in the test set of an example session. (c) Comparison between prediction correlations achieved by SynapsNet and NeurPRINT across the test portion of all sessions. **** indicates p-value $\leq .0001$ achieved by paired-t-test.

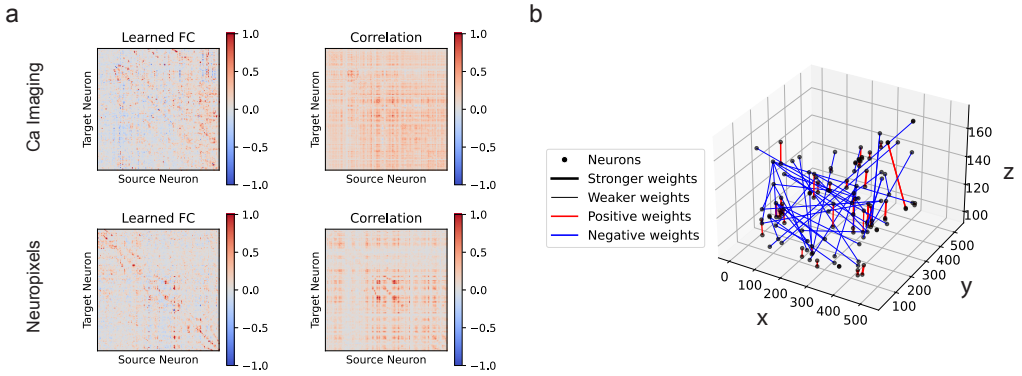

Figure 4: Illustration of the learned functional connectivity. (a) The connectivity matrix learned by SYNAPSNET compared with that of achieved by pair-wise Pearson's correlation. The top and bottom matrices correspond to a sample session from Ca imaging and Neuropixels modalities respectively. (b) 3D visualization of the learned functional connectivity by SYNAPSNET on a sample Ca imaging session. Dots represent neurons and lines show the type and strength of the connections. Coordinates are in μm.

in Figure 5b, the cross-correlation curve peaks at a positive delay, indicating that the input current is most correlated with future activity, demonstrating a predictive relationship. In contrast, the FC derived from pair-wise correlation peaks at a delay of zero. Notably, in the forecasting task, predicting activity at time t relies solely on information up to time $t-1$, making current and future time points inaccessible. Interestingly, the input currents derived by SYNAPSNET show a higher correlation with population activity at positive delays compared to the pair-wise correlation method, which are the only time points available in forecasting.

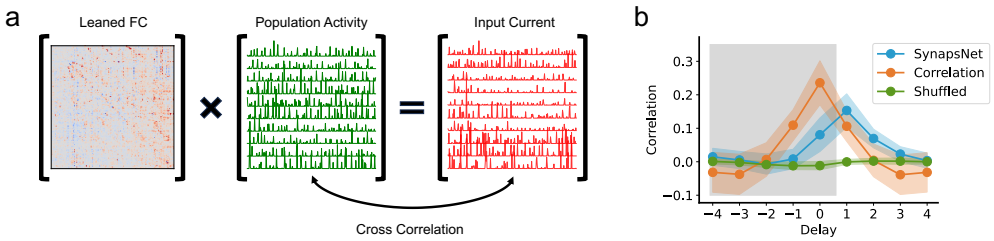

Figure 5: Input currents achieved by functional connectivity. (a) An example calculation of the input current to each neuron based on the population activity and the learned functional connectivity. (b) Average cross-correlations between each neuron's input currents and their activity. The input currents are calculated based on connectivity matrices achieved by three different methods: SYNAPSNET, pair-wise correlation, and shuffling the SYNAPSNET's connectivity matrix. The dots and color-shaded areas represent the mean and standard deviation of the correlations. The x-axis shows delays in time steps and the gray-shaded area marks the unavailable time points during the neuron activity forecasting task.

B.3 Ablation Study

We conducted ablation experiments on SYNAPSNET by removing each of its main components. The results, presented in Table 2, indicate that removing functional connectivity leads to the most significant increase in test loss across both data modalities. Furthermore, the model without functional connectivity exhibits the lowest correlation score on the Neuropixels dataset and nearly matches the lowest correlation score on the calcium imaging data. These results suggest the significant role of functional connectivity in capturing the population dynamics.

Table 2: Ablation experiments on neural data forecasting. Mean correlation and loss (%) $\pm$ standard error of the mean[†].

Data Modality	Model	Natural Scenes	
		$Corr(\%) \uparrow$	$Loss \downarrow$
Ca Imaging	SYNAPSNET	$37.94_{\pm 1.15}$	$0.825_{\pm 0.035}$
	Without Neuronal Embeddings	$35.96_{\pm 1.04}$	$0.836_{\pm 0.035}$
	Without Previous Activity	$\mathbf{33.27}_{\pm \mathbf{1.22}}$	$0.851_{\pm 0.036}$
	Without Functional Connectivity	$33.94.09_{\pm 0.867}$	$\mathbf{0.920}_{\pm \mathbf{0.038}}$
NeuroPixels	SYNAPSNET	$25.48_{\pm 0.23}$	$0.932_{\pm 0.012}$
	Without Neuronal Embeddings	$24.68_{\pm 0.23}$	$0.933_{\pm 0.009}$
	Without Previous Activity	$24.59_{\pm 0.24}$	$0.933_{\pm 0.010}$
	Without Functional Connectivity	$\mathbf{21.85}_{\pm \mathbf{0.22}}$	$\mathbf{0.949}_{\pm \mathbf{0.009}}$

[†] The bold values correspond to the largest drop in the performance metric separately for multi-/single-session training and each data modality.

B.4 Sensitivity Analysis

Effect of Context Window Length on Forecasting Performance. We evaluated the effect of different context window lengths (W) on the forecasting performance of both SYNAPSNET and NeuPRINT (the second-best-performing model). The results for Ca imaging and Neuropixels datasets are presented separately in Figure 6. The context window length determines how far back in the past the model can access neural activity and behavioral data to predict the population's activity at the next time point. The results indicate that at least two past time steps are required to accurately forecast neural activity, as evidenced by the substantial performance difference between $W = 1$ and $W = 2$. This aligns with the observations made by the authors of NeuPRINT, who identified $W = 2$ as the optimal choice [11]. Performance plateaus after $W = 5$, with the negligible difference observed

among context window sizes of 5, 10, and 20. Therefore, all results reported in the main paper are based on $W = 5$ to maintain model simplicity without significantly sacrificing performance.

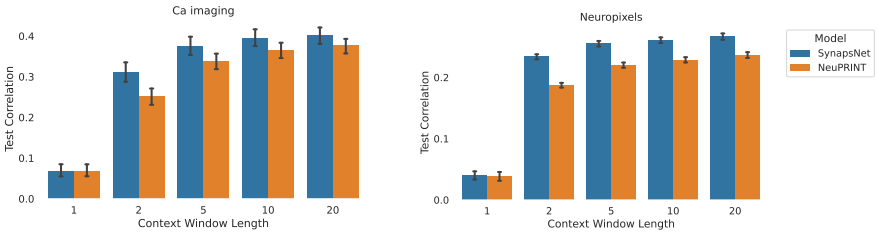

Figure 6: Sensitivity Analysis of Context Window Length (W). Test correlation scores for different context window lengths are plotted separately for the Ca imaging dataset (a) and the Neuropixels dataset (b).

Effect of Using Different Sequential Models on Forecasting Performance. We explored four different architectures for the SYNAPSNET's dynamical model—RNN, GRU, LSTM, and Transformer—to assess their impact on the model's performance. The results for each architecture are detailed in Table 3 (we only used natural scenes task for simplicity). Among these, the GRU architecture achieved the highest correlation score across both data modalities, leading us to base all the results presented in this paper on the GRU dynamical model.

Table 3: Sensitivity Analysis of the dynamical model's architecture. Test correlation scores for different neuronal embedding sizes are reported separately for the Ca imaging dataset and the Neuropixels dataset.

Data Modality	Dynamical Model	Natural Scenes	
		$Corr(\%) \uparrow$	$Loss \downarrow$
Ca Imaging	RNN	$36.99_{\pm 1.05}$	$0.787_{\pm 0.033}$
	GRU	$37.43_{\pm 1.05}$	$0.846_{\pm 0.037}$
	LSTM	$37.23_{\pm 1.09}$	$0.786_{\pm 0.033}$
	Transformer (one-layer)	$29.83_{\pm 1.86}$	$0.802_{\pm 0.036}$
	Transformer (two-layer)	$27.35_{\pm 2.48}$	$0.826_{\pm 0.039}$
Neuropixels	RNN	$24.04_{\pm 0.26}$	$0.927_{\pm 0.012}$
	GRU	$24.38_{\pm 0.28}$	$0.926_{\pm 0.011}$
	LSTM	$23.91_{\pm 0.29}$	$0.927_{\pm 0.012}$
	Transformer (one-layer)	$20.67_{\pm 1.32}$	$0.917_{\pm 0.023}$
	Transformer (two-layer)	$22.00_{\pm 1.18}$	$0.876_{\pm 0.012}$

† The bold values correspond to the largest drop in the performance metric separately for multi-/single-session training and each data modality.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims in the abstract and introduction are supported by the experiments presented in the paper.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Limitations are provided in the discussion.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [NA]

Justification: We do not have theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Implementation details are available in the main paper and supplementary materials.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We used publicly available datasets. We also share the codes used to produce these results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: Main text and supplementary materials describe the experiments in details.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Measure of variability and statistical significance are reported in the results.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Provided in the supplementary materials.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: This work does not violate the NeurIPS Code of Ethics.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

    Justification: Our work proposes a new model of neuronal populations and we do not see any social impacts associated with it.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: Our work proposes a new model of neuronal populations and we do not see any case of abuse associated with it.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: Papers and methods we used are cited in this paper.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: We do not introduce new assets.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: We do not have any human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: We do not have any human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.