Shallow Recurrent Decoders for Neural and Behavioral Dynamics

Amy S. Rude
Department of Applied Mathematics
University of Washington
Seattle, WA 98105
amysrude@uw.edu

J. Nathan Kutz
Department of Applied Mathematics &
Department of Electrical and Computer Engineering
Seattle, WA 98105
kutz@uw.edu

Abstract

Machine learning algorithms are affording new opportunities for building bio-inspired and data-driven models characterizing neural activity. Critical to understanding decision making and behavior is quantifying the relationship between the activity of neuronal population codes and individual neurons. We leverage a SHallow REcurrent Decoding (SHRED) architecture for mapping the dynamics of population codes to individual neurons and other proxy measures of neural activity and behavior. SHRED is constructed from a temporal sequence model, which encodes the temporal dynamics of limited sensor data, and a shallow decoder, which reconstructs the corresponding high-dimensional neuronal and/or behavioral states. It is a robust and flexible sensing strategy which allows for decoding the diversity of neural measurements with only a few sensor measurements. Thus estimates of whole brain activity, behavior, and individual neurons can be constructed with only a few neural time-series recordings. Several examples in this paper further highlight the potential of leveraging non-invasive or minimally invasive measurements to estimate large-scale brain dynamics.We empirically demonstrate the capabilities of the method on the neural data from *C. elegans* and mice.

1 Introduction

High-dimensional networked biological systems are ubiquitous and characterized by a large connectivity graph whose structure determines how the system operates as a whole [49, 35]. Typically the connectivity is so complex that the functionality, control and robustness of the network of interest is impossible to characterize using currently available methods. Moreover, with few exceptions, underlying nonlinearities impair our ability to construct analytically tractable solutions, forcing one to rely on experiments and/or modern high-performance computing to study a given system. However, advances in experimental and computational neuroscience over the past decades have revealed two critical, and seemingly ubiquitous, phenomena: (i) that the nervous system leverages dimensionality reduction when it encodes behavior and decision making [48, 38, 3, 37, 17], and (ii) that sparsity is fundamental in encoding and decoding information in the network [15, 34, 25]. Thus, instead of relying on individual neurons, one can consider the dynamic encoding space, or low-dimensional manifold [9], of the brain for representing complex decision making and cognitive capabilities.

A full understanding of the computational process encoded throughout a nervous system must reconcile these two key observations that transform sensory input into decision making and motor representations. These two observations are the basis of our proposed SHallow REcurrent Decoder (SHRED); a deep learning architecture that allows for an efficient and robust mathematical framework connecting the activity of individual (sparse) neurons to population codes and behavioral measurements. Originally developed as a sensing strategy, SHRED [51] is a promising deep learning architecture with impressive reconstruction and forecasting results in the low-data limit [27, 10, 40, 45, 16]. In

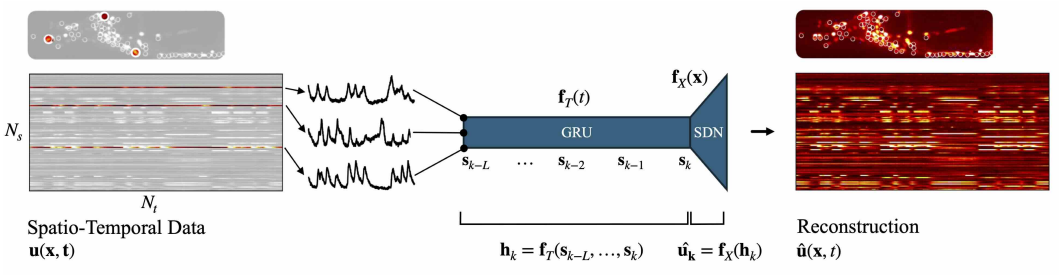

Figure 1: **Diagram of *SHallow REcurrent Decoder* (SHRED) architecture.** Time-series data from three neurons $\{s_i\}_{k-L}^{k}$ is fed into a *Gated Recurrent Unit* (GRU). The output from GRU $\mathbf{h_k}$ is the input for the *Shallow Decoder Network* (SDN), mapping back to the full state space. *C. elegan* neuron image (top) is adapted from [22].

this paper, we empirically demonstrate SHRED's ability to uncover population neural codes from a limited number of sensors. While SHRED has traditionally been applied to problems in physics and engineering [51, 27, 40, 45], the dynamic, high-dimensional, and coupled nature of neural population codes makes neuroscience an ideal domain for this architecture.

Importantly, SHRED's basic architecture connects sparse sampling of signals to population codes. Based on Takens' embedding theorem [44], SHRED leverages the fact that time-series measurements of a single state variable dynamically coupled to the system are sufficient to reconstruct a state space diffeomorphic to the true dynamics. In the context of neuroscience and supervised machine learning, this implies that recordings from a single neuron could, in principle, be sufficient to recover population activity. Theoretically rooted in the classic PDE method of separation of variables [51, 45], the decoding-only strategy of SHRED circumvents the computation of inverse pairs, i.e. an encoder and the corresponding decoder. By decoding only, SHRED avoids this problem and learns a single embedding without the corresponding inversion. Moreover, the SHRED architecture provides an approximation to an operator which maps the time-trajectories to full state estimates. Tomasetto et al [45] compares the SHRED architecture to leading operator learning methods showing the superior and robust performance advantage. Figure 1 details the basic architecture of the neural network, where the latent space of a temporal encoder is used to reconstruct full activity across neurons in the *C. elegans*.

SHRED is leveraged on a number of toy models in neuroscience, including the nematode Caenorhabditis elegans (*C. elegans*) and the mouse. Additional model systems (Sec. B), control comparisons (Sec. C), and sensitivity analyses (Sec. D) are included in the appendix. For each model system, SHRED is shown to be robust and accurate in approximating population codes using only a limited number of sensors. The architecture is thus promising for extracting critical information on decision making and motor commands with only limited, and potentially non-invasive, or minimally invasive, measurements.

2 Mathematical background for SHRED

SHRED consists of a combination of recurrence and decoding, with the former operation capturing the temporal behavior of limited sensor data and the latter operation performing a spatial upscaling to recover the corresponding high-dimensional state. A similar split in spatial and temporal components is taken into account by the well-known method of separation of variables for solving linear PDEs [12], which assumes that the solution has the form $u(\mathbf{x}, t) = T(t)X(\mathbf{x})$. Indeed, it has been shown that in the limit of linear dynamics with constant coefficients, the SHRED architecture has a closed form analytic solution [51, 45]. Specifically, for linear PDEs, the solution is given by a sum of modes with exponential time dynamics. Similarly, in the linear limit the analytic solution of the dynamic mode decomposition (DMD) [42, 26] is a sum of modes with exponential time dynamics. Of note is that the architecture can be used to produce estimates for uncertainty quantification [41] along with multiscale features [4].

Thus SHRED is a generalization of separation of variables with nonlinear (neural network) approximations. A compact representation of the SHRED algorithm, as shown in Fig. 1 is as two neural

networks jointly trained to the form

$$\mathbf{A} = \mathbf{X}(x)\mathbf{T}(t) \rightarrow \mathbf{A} = f_{\boldsymbol{\theta}_1}(x) \circ g_{\boldsymbol{\theta}_2}(t) \tag{1}$$

where $\mathbf{A}$ is a data matrix whose columns are time samples and whose rows are the full state space, $g_{\boldsymbol{\theta}_2}(t)$ is trained largely to encode sequential (time) data and $f_{\boldsymbol{\theta}_1}(x)$ is largely responsible for mapping from the latent space of the sequential encoding to the full state-space. Of course, the time and space encoding are not independent as they are jointly trained and the decoder is dependent on the time history of the sequence model.

To be more precise, as a deep learning architecture, SHRED aims to reconstruct high-dimensional neuronal data from N_s individual neural recordings. Instead of learning the *one-shot* reconstruction map $\mathbf{s}_k = \{u(\mathbf{x}_s, t_k)\}_{s=1}^{N_s} \rightarrow \mathbf{u}_k$ for $k = 1, \ldots, N_t$ – as considered by, e.g., [33, 31, 30, 53] – we take advantage of the temporal history of the individual neuron values. Specifically, a recurrent neural network $\mathbf{f}_T$ encodes the neuron measurements over a time window of length $L \leq N_t$ into a latent representation of dimension N_l, that is

$$\begin{aligned}
&\mathbf{f}_T : \underbrace{\mathbb{R}^{N_s} \times \ldots \times \mathbb{R}^{N_s}}_{L \text{ times}} \rightarrow \mathbb{R}^{N_l}, \\
&\mathbf{h}_k = \mathbf{f}_T(\mathbf{s}_{k-L}, \ldots, \mathbf{s}_k),
\end{aligned}$$

where the hyperparameter L identifies the number of lags considered by SHRED, and may be selected according to the problem-specific evolution rate or periodicity. The high-dimensional state is then approximated through a decoder $\mathbf{f}_X$, which performs a nonlinear projection of the latent representation onto the state space

$$\begin{aligned}
&\mathbf{f}_X : \mathbb{R}^{N_l} \rightarrow \mathbb{R}^{N_h}, \\
&\mathbf{u}_k \approx \hat{\mathbf{u}}_k = \mathbf{f}_X(\mathbf{h}_k) = \mathbf{f}_X(\mathbf{f}_T(\mathbf{s}_{k-L}, \ldots, \mathbf{s}_k)).
\end{aligned}$$

In this work, we consider a *gated recurrent unit* (GRU) network [7] to model the temporal dependency of sensor data, and a *shallow decoder network* (SDN) as the latent-to-state map. SHRED is a modular framework, allowing for various forms of temporal encoding and spatial decoding, including convolutional neural networks [28], LSTM [21], echo state networks [29], transformers [47] and variational autoencoders [23]. Importantly, the decoding only strategy circumvents learning an encoder-decoder pair [18] which leads to numerical stiffness [13, 20].

Unless noted otherwise, the data used for training SHRED is accomplished with an 80:10:10 train/test/val split. Each data set is unique, and amount of data in the 80:10:10 split can differ significantly. If a different train/test split is used, it is explicitly outlined below. Table 1 gives a comprehensive summary of all mean-squared errors for Sections 3 A and B. For access to the code and data used to produce the results, see Sec. A in the Appendix. Furthermore, additional computational results (Sec. B), control comparisons (Sec. C), and sensitivity analyses (Sec. D) are also presented in the Appendix.

3 Model systems

The SHRED architecture maps limited, time-series sensor measurements to the full, high-dimensional neural or behavior data. In what follows, we show how measuring only a limited number of neurons or sensors can estimate the full state space in two different model systems.

3.1 *C. elegans* nematode

The nematode (*C. elegans*) is a perfect model organism to consider in the context of neuro-sensory integration as it is comprised of only 302 sensory, motor and inter-neurons whose electro-physical connections (i.e. its connectome) are known and stereotyped [50, 5]. In addition, it possesses only a small number of sensory neurons, often linked to specific stimuli [2]; neural activity can be imaged in real time, and its range of behavioral responses, during chemotaxis for instance, are varied yet limited, approximately confined to five observable motor states that are related to forward and backward crawling, omega turns, head sweeps and brief pause states. Thus it is reasonable to posit a complete model of its neuro-sensory integration capabilities.

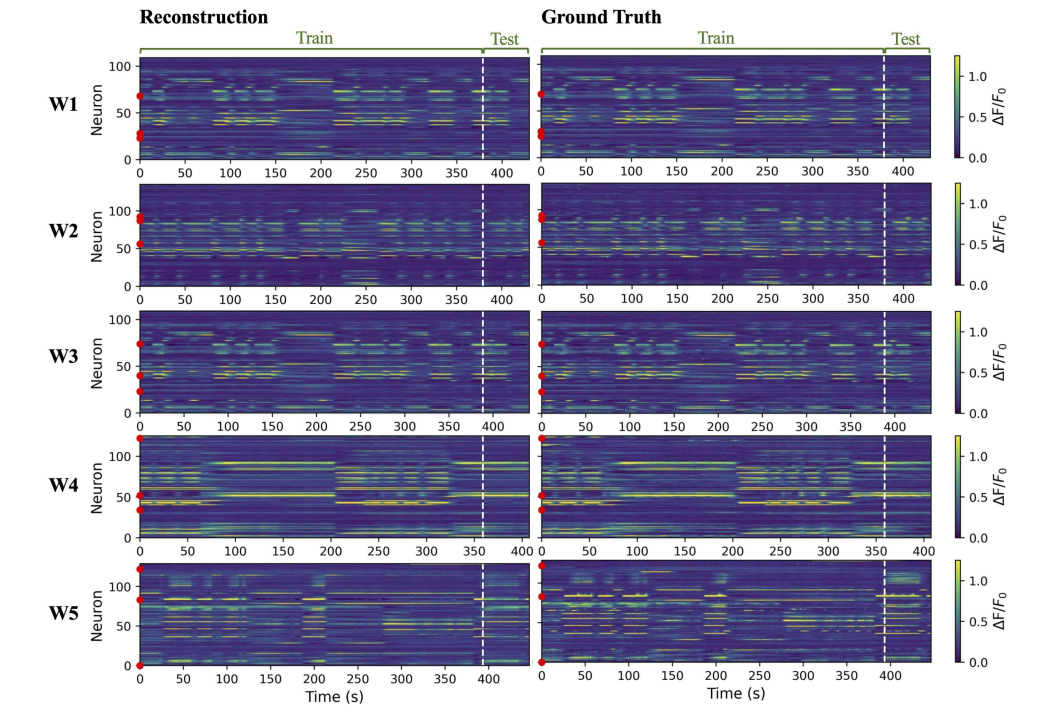

Figure 2: **SHRED can accurately reconstruct population-wide activity from only three neurons.** SHRED reconstruction (left) of neural activity (right) for all five worms. General activity patterns are preserved in the reconstruction of both training and testing datasets.

We obtained whole-brain calcium imaging recordings of immobilized nematodes (*C. elegans*) from a publicly available dataset provided in [22]. This dataset contains data for five worms with recordings from 107 - 131 identified neurons in an environmentally controlled condition. Recordings were made for a duration of roughly 18 minutes at an average rate of 2.85 Hz. Since bleaching causes damping of Ca^{2+} imaging signals over time, we truncated the data to only include the first half of each session [22]. In the sections that follow, we apply SHRED to this dataset for individual (Section 3.1.1) and cross-worm predictions (Section 3.1.2). The results presented in these sections demonstrate the capabilities of SHRED to decode population-wide dynamics from sparse neural measurements. Sensitivity analysis for the model is presented in Section D.

3.1.1 SHRED reconstruction of *C. elegan* neural activity

We reconstructed whole population activity for all five worms using SHRED. As shown in the carpet plots in Figure 2, some neurons remained inactive for the entire duration of the recording. To ensure that our model performance was not negatively impacted by the selection of these inactive neurons, we created subgroups of 34-41 *High Variance* (HV) neurons and 75 - 95 *Low Variance* (LV) neurons. Details of this split can be found in Section D. To perform the reconstruction, we randomly selected three neurons from the HV group and the first 1500 timepoints (approximately 8.5 minutes) from each worm was split 80:10:10 (train:test:validation). With a lag $L = 100$ the model was trained for 200 epochs. The corresponding normalized Mean Square Error (MSE) was calculated for both the training and testing datasets. Baseline control comparisons are conducted in Section C.

$$MSE = \frac{||\mathbf{f}_X(\mathbf{h}_k) - u_k||_2}{||u_k||_2} \tag{2}$$

SHRED does a remarkable job learning spatial and temporal patterns in data (Figure 2 & Table 1). To the naked eye, the reconstructions presented in Figure 2 are near identical to the original

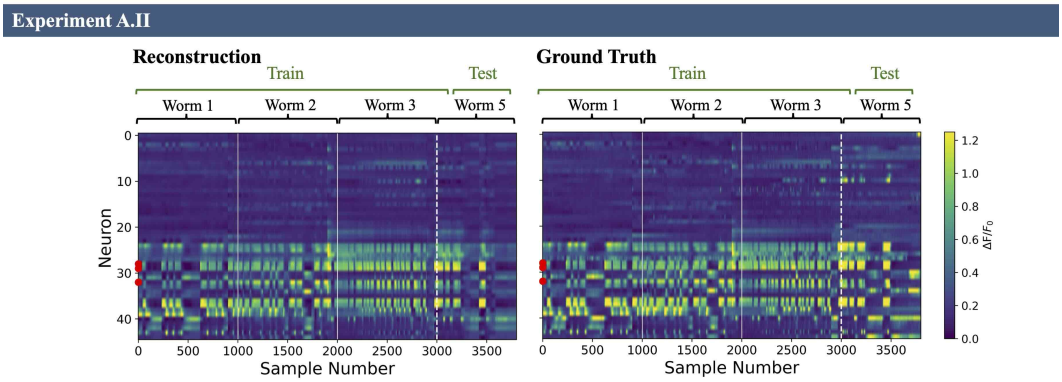

Figure 3: **SHRED can generalize across worms.** A SHRED model trained to reconstruct neural activity of Worms 1-3 from three high variance neurons can extrapolate to reconstruct activity for Worm 5. Red dots on the y-axis indicate sensors selected as input. Refer to Table 1 Exp A.II for error metrics.

Experiment A.I		
Worm	Train	Test
1	0.146	0.276
2	0.182	0.351
3	0.141	0.299
4	0.175	0.286
5	0.282	0.461

Experiment A.II		
Worm	Train	Test
1-5	0.240	0.571

Experiment B.I		
Probe	Train	Test
DG-A	0.355	0.391
DG-C	0.512	0.630

Experiment B.II		
Stim.	Train	Test
DG	0.174	0.173
FL	0.6663	0.713

Table 1: Mean Square Error (MSE) for all experiments presented in Sections 3.1 and 3.2. DG - drifting grating stimulus. FL - flash stimulus.

data. Activity traces are reconstructed in a temporally precise fashion —the peaks and troughs in the activity of each neuron are well-preserved. Even for the withheld test dataset, the model performs remarkably well with the exception of Worm 5 (W5). Several neurons exhibit heightened activity throughout the duration of the test data. This pattern in activity is sustained for a shorter duration in the reconstruction, resulting in a higher error. With requiring only three sensor recordings as input, SHRED has the potential to enable less invasive recording procedures and even recover information from low-quality sensors. In the following section, we increase the difficulty by reconstructing activity from a worm that was completely withheld during training. Such cross-individual predictions are particularly valuable, as they further reduce the need for repeated invasive experiments.

3.1.2 Cross-worm predictions using SHRED

So far, SHRED has shown remarkable ability for reconstructing neural activity from a select number of neurons for a single individual. In this task, we leveraged the conserved patterns in activity of *C. elegans* to reconstruct neural activity across all five worms. For each worm, we selected the first 1000 time points. Only 45 neurons were recorded from across all 5 worms. Therefore, only these neurons were selected to be included in this trial. The training data was defined to be the first 3000 time points which corresponds to the neural activity of the first three worms. The validation data was defined to be the next 1000 time points corresponding to the neural data of worm four. The test data was defined to be the last 1000 time points which corresponds to the neural data of worm five. Reconstruction of the test dataset only includes the first 900 time points to account for the lag (L) of 100. The data was preprocessed to include 18 neurons in the HV group and 27 neurons in the LV group. We randomly selected three neurons from the *High Variance* group for the reconstruction. Results are presented in Figure 3.

The training error and test error are reported in Table 1(A.II). As expected, since the test set is extrapolating to a new worm not included in the training, the test error is substantially higher than the training MSE. Nevertheless, the reconstructed activity for the fifth worm still preserves the overall

trends (Figure 3). Importantly, this dataset provides reliable neural measurements for five worms from the same neurons, enabling cross-individual predictions. This example illustrates the potential of SHRED to aid in minimally invasive or non-invasive recording techniques. Thus once training is accomplished on the full data set, it can be deployed to new subjects with access only to non-invasive sensor sets. Moreover, by successfully extrapolating beyond individuals used for training, SHRED captures distinct spatio-temporal patterns without overfitting to the training data.

Applications of SHRED on the nematode *C. elegans* dataset demonstrated the flexibility and adaptability of this neural-network architecture. We next increased the challenge by leveraging SHRED for a more complex dataset: mouse unit and LFP recordings.

3.2 Mouse

The mouse is a popular model organism to study behavior and human diseases [36]. In recent years, we have seen a rapid rise in techniques for large-scale neural recordings during animal behavior, opening the doors to understanding network-level neural codes in the mouse brain [46, 6]. The next dataset we utilized is from the The Visual Coding – Neuropixels project developed by the Allen Brain Observatory. The main objective of the Neuropixel project was to study neuronal response to visual stimuli. Mice were presented with various visual cues such as gabors, flashing lights, and drifting gratings. The neural activity, both individual neurons and Local Field Potentials (LFP), in response to these stimuli were recorded from 6 neuropixel probes inserted throughout the visual cortex.

Each probe contained 374 or 383 channels with each channel containing two bands. The "spike band" measured action potentials of nearby neurons and was recorded at 30 kHz. The "LFP band" measured voltage fluctuations and was recorded at 2.5 kHz. LFPs capture lower frequency voltage fluctuations resulting from transmembrane currents from multiple neurons. Further processing and filtering were conducted to produce the publicly available Neurodata Without Borders (NWB) files. Pupil diameter measurements were also provided for the duration of the recordings. For convention, individual neurons were referred to as "units" in the spiking dataset since it can not be guaranteed that each recording originates from a single cell. Therefore, we also refer to neurons as "units."

To simplify analyses, we only selected data from a single mouse. A 96-day-old male mouse (ID 756029989) was selected due to the quality of its recordings. Both unit and LFP data were used in the following sections to demonstrate the robustness of SHRED (Sections 3.2.1, 3.2.2, B.1, B.2).

3.2.1 SHRED reconstruction of neural response to visual stimuli

Our first objective was to reconstruct Local Field Potential (LFP) responses to visual stimuli. In this section, we obtained data for the first three presentations of a 2-second 180° drifting grating stimulus. We selected LFP recordings from both probe A and probe C. Probe A was inserted into the anteromedial (AM) region of the visual cortex and contained a total of 74 channel recordings from the anteromedial area (VISam), cornu ammonis 1 (CA1), dentate gyrus (DG), and the anterior pretectal nucleus (APN). Probe C was inserted into the primary visual cortex (V1) and contained a total of 76 channel recordings from the primary visual cortex (VISp), subiculum (SUB), postsubiculum (POST), superior colliculus (SCig), midbrain (MB), anterior pretectal nucleus (APN), and the posterior limiting nucleus of the thalamus (POL).

For both experiments, we downsampled the data so that each response period contained 1000 time points. Consequently, the 2-s DG response was sampled at 0.5 kHz. The training data was split to include only the first two stimulus presentations. The first half of the third stimulus presentation was set aside as the test dataset, and the latter half of the third stimulus presentation was assigned to the validation set. For reconstructions of probe A, we randomly selected three channels from the cornu ammonis 1 (CA1) region of the hippocampus and trained with a lag of 100 for 200 epochs. For probe C, we randomly selected three channels from the subiculum (SUB) and trained with a lag of 100 for 200 epochs. No variance pre-processing was conducted for either stimuli.

SHRED more accurately reconstructed LFP data from Probe A compared to Probe C as shown in Table 1. By looking at the original dataset, it is clear that there is a simple oscillatory pattern in the voltage recorded from probe A which is a fairly straightforward task for many sparse-sensing algorithms. However,for probe C, some sensors in the primary visual cortex (VISp) and subiculum (SUB) exhibited a heightened level of activity which SHRED was unable to reconstruct. Nevertheless, general patterns in activity are conserved in the reconstruction of both probes. The similarity between

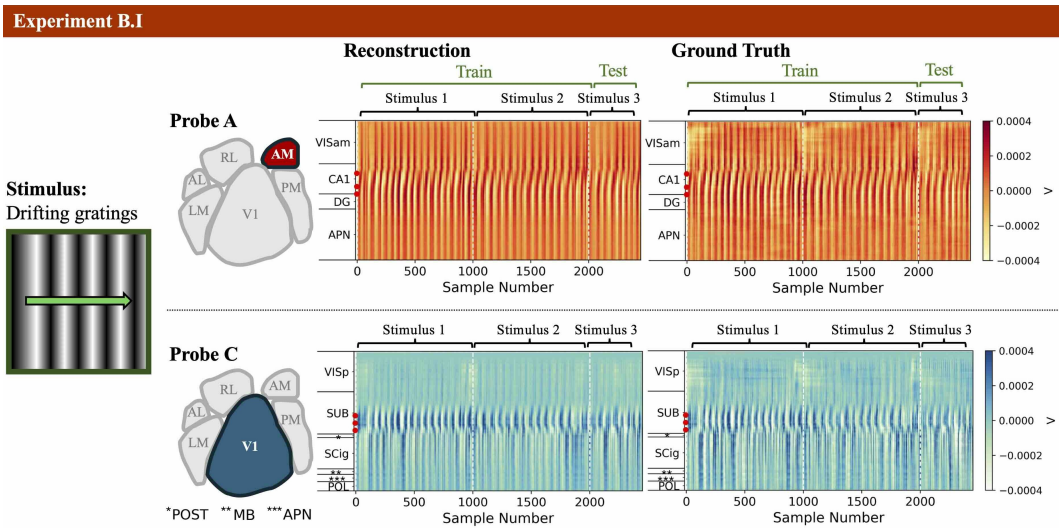

Figure 4: **SHRED can accurately reconstruct local field potential (LFP) in response to a 2s drifting grating. (Stimulus)** First three 2-second presentations of a 180°, 0.8-contrast drifting grating stimulus. **(Probe A)** LFPs were recorded from the anteromedial (AM) region of the mouse visual cortex. Red dots on the y-axis indicate the three neurons selected for training. **(Probe C)** LFPs were recorded from the primary visual cortex (V1) of the mouse. Red dots on the y-axis indicate the three neurons selected for training. **(Reconstruction)** For both probes, SHRED accurately reconstructed visually evoked LFP responses.

train and test MSE further highlights SHRED's generalizability (Table 1B.I). This algorithm can learn larger scale spatio-temporal patterns in population codes. By leveraging these patterns, SHRED can reconstruct whole-population activity from only three sensors.

3.2.2 SHRED reconstruction of average unit responses to visual stimuli

We next aimed to reconstruct averaged unit spike trains in response to all 2-second drifting grating and 25-millisecond flash stimuli. Spike train data is inherently stochastic and consists of discrete, binary events, presenting a significant challenge for many machine learning methods. We averaged spike train data across all presentations of each stimulus type to obtain a more continuous estimate of each unit's response. For each 2-second presentation of the drifting grating (dg) stimulus and 25-ms flash stimulus, we obtained 1000 time points of spiking data for 243 units. We then averaged this data across presentations to calculate the average unit responses. The data was pre-processed and split into 102 HV units and 141 LV units. Three units with high variability in activity were randomly selected from the HV group and the model was trained for 200 epochs.

SHRED more accurately reconstructed the average unit response to the drifting gradient stimulus compared to the flashes. This is most likely due to lower fluctuations in individual spiking activity across time for the DG stimulus compared to the flash (Figure 5). We observed lower fluctuations most likely due to two main reasons: 1) the drifting grating unit response is sampled with a lower sampling rate and 2) responses to drifting gratings of different orientations, speed, and contrasts are averaged together, smoothing out the unit's response. For example, consider a unit with a receptive field sensitive to light located in the bottom left-hand corner of the screen. When gratings are presented at different orientation and at different speeds, it is plausible that the receptive field will always be lit in at least one of these presentations. Therefore, averaging the data may result in a smoother activity trace, setting up an easier reconstruction task for SHRED.

The reconstruction errors are higher for the flash stimulus (Table 1(B.II)), but the example traces of unit 150 in Figure 5 C1-C2 show that general activity patterns are preserved. While individual spikes are not reconstructed, SHRED learns the time course of when the unit is firing more frequently. From this result, we hypothesized that the model is learning the overall firing rates rather than the timing of the individual spikes. Therefore in Section B.1, we present reconstruction of LFP from firing rate.

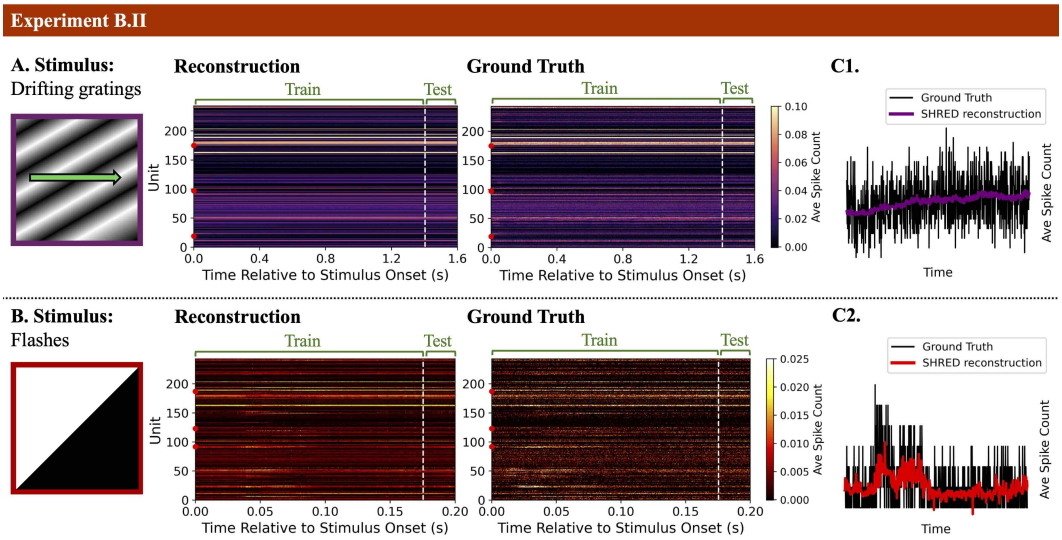

Figure 5: **SHRED accurately reconstructs and filters high frequency recordings of averaged unit activity.** **(A)** SHRED reconstruction of average unit activity for all 2-s trials with drifting grating stimuli. **(B)** SHRED reconstruction of average unit activity for all 25-ms trials with light flashing stimuli. **(C1-C2)** Average activity (black) and SHRED reconstruction (purple/red) of a randomly selected neuron. For both stimuli, averaged unit recordings contain high frequency fluctuations which are filtered out in the SHRED reconstruction. Overall trends in activity are well-preserved. Red dots on the y-axis indicate sensors selected as input.

4 Conclusions and Discussions

Modern neuroscience and behavioral sciences are now supplemented with exceptional data sets which measure a diversity of quantities, from individual neurons to whole brain activity to behavioral outcomes. Such emerging data sets, which are now routinely in the open source domain, offer significant scientific opportunities for learning and discovery in the biological sciences. Specifically, there are modern data-driven mathematical architectures and algorithms capable of extracting interpretable generative models from the diversity of data collected. Moreover, it is ideal if cheap and non-invasive sensing can be used instead of expensive and invasive sensors. The SHRED architecture introduced in this work provides a viable pathway towards using a highly limited number of sensors for reconstructing the full high-dimensional state-space of either neural activity or behavioral responses. And importantly, because of the globally convex landscape and ability to train with compressed data [27, 45], training SHRED with neural data can be accomplished on laptop-level computing, even for exceptionally large data sets.

SHRED has demonstrated in this work to have the capability of mapping limited neural or behavioral recordings accurately to full state-space measurements. The success of the method allows for deployment in practice to new patients or model organisms. For the *C. elegans* example, the ability of SHRED to be deployed in a parametric fashion affords many opportunities in clinical settings where non-invasive measurements can be used to approximate difficult or invasive measurements. In future work, learning generative models of the dynamics will be explored with the goal of extracting even more meaningful insights into biological activity [16]. Thus SHRED overall is an emerging deep learning algorithm which is well suited for diverse biological data.

Acknowledgments and Disclosure of Funding

We wish to acknowledge the support of the National Science Foundation AI Institute in Dynamic Systems grant 2112085 and support from the Air Force Office of Scientific Research (FA9550-24-1-0141).

References

[1] Misha B Ahrens, Michael B Orger, Drew N Robson, Jennifer M Li, and Philipp J Keller. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature methods*, 10(5):413, 2013.

[2] Z.F. Altun, L.A. Herndon, C. Crocker, and D. H. Hall. Wormatlas. *http://www.wormatlas.org*, 2002-2012.

[3] Bruno B Averbeck, Peter E Latham, and Alexandre Pouget. Neural correlations, population coding and computation. *Nature reviews neuroscience*, 7(5):358–366, 2006.

[4] Steven L Brunton and J Nathan Kutz. Methods for data-driven multiscale model discovery for materials. *Journal of Physics: Materials*, 2(4):044002, 2019.

[5] B. L. Chen, D. H. Hall, and D. B. Chlovskii. Wiring optimization can relate neuronal structure and function. *Proc. Natl. Acad. Sci.*, 103:4723–4728, 2006.

[6] Ritchie Chen, Andres Canales, and Polina Anikeeva. Neural recording and modulation technologies. *Nature reviews materials*, 2(2):1–16, 2017.

[7] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

[8] Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, 2012.

[9] John P Cunningham and Byron M Yu. Dimensionality reduction for large-scale neural recordings. *Nature neuroscience*, 17(11):1500–1509, 2014.

[10] Megan R. Ebers, Jan P. Williams, Katherine M. Steele, and J. Nathan Kutz. Leveraging arbitrary mobile sensor trajectories with shallow recurrent decoder networks for full-state reconstruction. *IEEE Access*, 12:97428–97439, 2024.

[11] Claudia E Feierstein, Ruben Portugues, and Michael B Orger. Seeing the whole picture: a comprehensive imaging approach to functional mapping of circuits in behaving zebrafish. *Neuroscience*, 296:26–38, 2015.

[12] Gerald B Folland. *Introduction to partial differential equations*, volume 102. Princeton university press, 1995.

[13] George E Forsythe, Michal A Malcolm, and Cleve B Moler. Computer methods for mathematical computation prentice-hall. *Englewood Cliffs, New Jersey*, 1977.

[14] Rainer W Friedrich, Christopher J Habermann, and Gilles Laurent. Multiplexing using synchrony in the zebrafish olfactory bulb. *Nature neuroscience*, 7(8):862–871, 2004.

[15] S. Ganguli and H. Sompolinsky. Compressed sensing, sparsity, and dimensionality in neuronal information processing. *Ann. Rev. Neuro.*, 35:485, 2012.

[16] Mars Liyao Gao, Jan P. Williams, and J. Nathan Kutz. Sparse identification of nonlinear dynamics and koopman operators with shallow recurrent decoder networks. *arXiv preprint arXiv:2501.13329*, 2025.

[17] Peiran Gao, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya Ganguli. A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*, 2017.

[18] Craig Gin, Bethany Lusch, Steven L Brunton, and J Nathan Kutz. Deep learning models for global coordinate transformations that linearise PDEs. *European Journal of Applied Mathematics*, pages 1–25, 2020.

[19] Georgia G Gregoriou, Stephen J Gotts, and Robert Desimone. Cell-type-specific synchronization of neural activity in fef with v4 during attention. *Neuron*, 73(3):581–594, 2012.

[20] Nicholas J Higham. *Accuracy and stability of numerical algorithms*. SIAM, 2002.

[21] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, November 1997.

[22] **Saul Kato, Harris S Kaplan, Tina Schrödel, Susanne Skora, Theodore H Lindsay, Eviatar Yemini, Shawn Lockery, and Manuel Zimmer. Global brain dynamics embed the motor command sequence of caenorhabditis elegans. *Cell*, 163(3):656–669, 2015.

[23] Diederik P. Kingma and Max Welling. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019.

[24] Sung-Kyun Ko, Xiaoqiang Chen, Juyoung Yoon, and Injae Shin. Zebrafish as a good vertebrate model for molecular imaging using fluorescent probes. *Chemical Society Reviews*, 40(5):2120–2130, 2011.

[25] J Nathan Kutz. Neurosensory network functionality and data-driven control. *Current Opinion in Systems Biology*, 13:31–36, 2019.

[26] J Nathan Kutz, Steven L Brunton, Bingni W Brunton, and Joshua L Proctor. *Dynamic mode decomposition: data-driven modeling of complex systems*. SIAM, 2016.

[27] J. Nathan Kutz, Maryam Reza, Farbod Faraji, and Aaron Knoll. Shallow recurrent decoder for reduced order modeling of plasma dynamics. *arXiv preprint arXiv:2405.119955*, 2024.

[28] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[29] Mantas Lukoševičius. *A Practical Guide to Applying Echo State Networks*, pages 659–686. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

[30] Zhaohui Luo, Longyan Wang, Jian Xu, Meng Chen, Jianping Yuan, and Andy C. C. Tan. Flow reconstruction from sparse sensors based on reduced-order autoencoder state estimation. *Physics of Fluids*, 35(7):075127, 2023.

[31] Romit Maulik, Kai Fukami, Nesar Ramachandra, Koji Fukagata, and Kunihiko Taira. Probabilistic neural networks for fluid flow surrogate modeling and data recovery. *Physical Review Fluids*, 5:104401, 2020.

[32] Jason R Meyers. Zebrafish: development of a vertebrate model organism. *Current Protocols Essential Laboratory Techniques*, 16(1):e19, 2018.

[33] Nirmal J. Nair and Andres Goza. Leveraging reduced-order models for state estimation using deep learning. *Journal of Fluid Mechanics*, 897:R1, 2020.

[34] B. Olshausen and D. Field. Sparse coding of sensory inputs. *Current Opinion in Neurobiology*, 14:481–487, 2004.

[35] H.-J. Park and K. Friston. Structural and functional brain networks: From connections to cognition. *Science*, 342:1238411–1–1238411–8, 2013.

[36] Robert L Perlman. Mouse models of human disease: an evolutionary perspective. *Evolution, medicine, and public health*, 2016(1):170–176, 2016.

[37] M. Rabinovich, R. Huerta, P. Varona, and V.S. Afraimovich. Transient cognitive dynamics, metastability, and decision making. *PLoS Comp. Bio.*, 4:e1000072, 2008.

[38] M. Rabinovich, A. Volkovskii, P. Lecanda, R. Huerta, H.D.I. Abarbanel, and G. Laurent. Dynamical encoding by networks of competing neuron groups: Winnerless competition. *Phys. Rev. Lettl.*, 87:068102, 2001.

[39] Ryan V Raut, Zachary P Rosenthal, Xiaodan Wang, Hanyang Miao, Zhanqi Zhang, Jin-Moo Lee, Marcus E Raichle, Adam Q Bauer, Steven L Brunton, Bingni W Brunton, et al. Arousal as a universal embedding for spatiotemporal brain dynamics. *BioRxiv*, pages 2023–11, 2025.

[40] Stefano Riva, Carolina Introini, Antonio Cammi, and J. Nathan Kutz. Robust state estimation from partial out-core measurements with shallow recurrent decoder for nuclear reactors. *arXiv preprint arXiv:2409.12550*, 2024.

[41] Diya Sashidhar and J Nathan Kutz. Bagging, optimized dynamic mode decomposition (bop-dmd) for robust, stable forecasting with spatial and temporal uncertainty-quantification. *arXiv preprint arXiv:2107.10878*, 2021.

[42] Peter J. Schmid and Joern Sesterhenn. Dynamic mode decomposition of numerical and experimental data. In *61st Annual Meeting of the APS Division of Fluid Dynamics*. American Physical Society, 2008. http://meetings.aps.org/link/BAPS.2008.DFD.MR.7.

[43] Elad Schneidman, Michael J Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.

[44] Floris Takens. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980, Lecture Notes in Mathematics*, volume 898, pages 366–381. Springer Berlin Heidelberg, 1981. https://doi.org/10.1007/BFb0091924.

[45] Matteo Tomasetto, Jan P Williams, Francesco Braghin, Andrea Manzoni, and J Nathan Kutz. Reduced order modeling with shallow recurrent decoder networks. *arXiv preprint arXiv:2502.10930*, 2025.

[46] Anne E Urai, Brent Doiron, Andrew M Leifer, and Anne K Churchland. Large-scale neural recordings call for new insights to link brain and behavior. *Nature neuroscience*, 25(1):11–19, 2022.

[47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2023.

[48] Amit Vinograd, Aditya Nair, Joseph H Kim, Scott W Linderman, and David J Anderson. Causal evidence of a line attractor encoding an affective state. *Nature*, 634(8035):910–918, 2024.

[49] D J Watts and S H Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.

[50] J. G. White, E. Southgate, J.N. Thomson, and S. Brenner. The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Phil. Trans. Roy. Soc. Lond. B*, 314:1–340, 1986.

[51] Jan P. Williams, Olivia Zahn, and J. Nathan Kutz. Sensing with shallow recurrent decoder networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 480(2298):20240054, 2024.

[52] Jan P Williams, Olivia Zahn, and J Nathan Kutz. Sensing with shallow recurrent decoder networks. *Proceedings of the Royal Society A*, 480(2298):20240054, 2024.

[53] Qian Zhang, Dmitry Krotov, and George Em Karniadakis. Operator learning for reconstructing flow fields from sparse measurements: an energy transformer approach. *arXiv preprint arXiv:2501.08339*, 2025.

A Code and data

All code and data are publicly available. Figures and experiments can be reproduced from the following GitHub: `https://github.com/amysrude/neuralSHRED`.

The data sets used in the analyses are from the following:

***C. elegans*:**
`https://osf.io/2395t/`

Mouse:
Download data:
`https://allensdk.readthedocs.io/en/latest/_static/examples/nb/ecephys_data_access.html`
Documentation:
`https://allensdk.readthedocs.io/en/latest/visual_coding_neuropixels.html`

Zebrafish:
`https://static-content.springer.com/esm/art%3A10.1038%2Fnmeth.2434/MediaObjects/41592_2013_BFnmeth2434_MOESM209_ESM.avi` and [1]

B Additional Results

Experiment B. III		
Data	Train	Test
FR*	0.060	0.061
LFP	0.658	1.070

Experiment B. IV		
Data	Train	Test
Pupil*	0.043	0.017
LFP	0.861	1.032

Experiment C.I		
Data	Train	Test
Video	0.290	0.396

Table 2: Mean Square Error (MSE) for all experiments presented in the Appendix.
DG - drifting grating stimulus. FL - flash stimulus. FR - firing rate. LFP - local field potential.
* - data for input into the model. Rows separated by horizontal lines represent separate models trained for the corresponding dataset.

B.1 SHRED Reconstruction of Local Field Potential from Individual Neuron Firing Rate

In Section 3.2.2, we observed that SHRED learns overall firing rates of individual units when reconstructing averaged spike data. Therefore, we then hypothesized that SHRED could reconstruct local field potential (LFP) from solely the firing rate. Although LFP recordings contain substantially higher frequency fluctuations than the firing rate, our focus was to first reconstruct the lower frequency signals. We obtained a spike train from a single unit from probe A in response to the first 2-second, $180°$, 0.8 contrast drifting grating stimulus. This unit was selected because of its high firing rate ($fr > 10$). We then obtained the LFP recording from the same channel, which reflects voltage fluctuations in the extracellular space near the selected unit.

The firing rate of the selected unit was calculated with a gaussian filter. Firing rate (v) of a neuron is defined as

$$v = \frac{n_s}{t}$$

where n_s is the number of spikes in time t. To obtain an accurate measurement of the spike rate, we used a one-dimensional gaussian filter. The sampling rate (sr) for the unit data was down sampled to around 1.25 kHz. The standard deviation for the gaussian was set to

$$\sigma = \frac{0.02}{sr}$$

thus the spike train was smoothed over a 20 millisecond window. After applying the filter, we multiplied the filtered trace by the sampling rate to obtain the firing rate in spikes/s. The smoothed

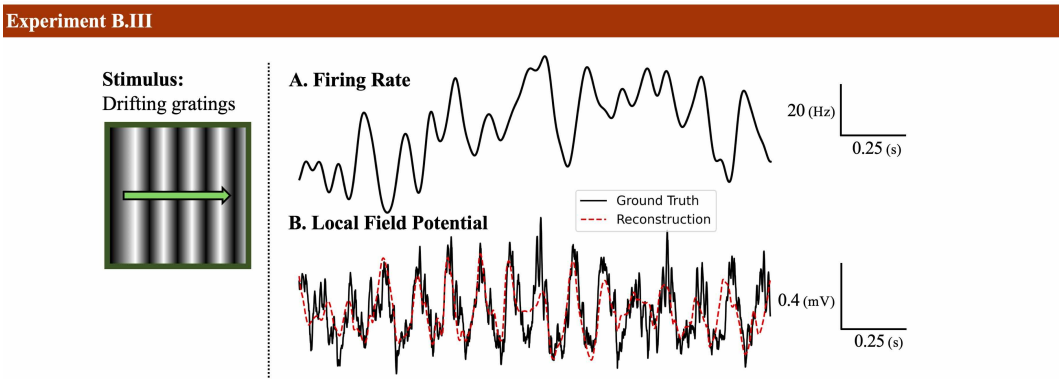

Figure 6: **SHRED can reconstruct Local Field Potential (LFP) from firing rate of nearby neuron. (A)** Firing rate of a neuron in the visual cortex located near the probe inserted in the anteromedial (AM) region served as the input to the SHRED model. **(B)** The SHRED model was trained to reconstruct LFP recording closest to the neuron selected as input. The model reconstruction (red) closely follows the general patterns of the LFP recoding (black). Red dots on the y-axis indicate sensors selected as input. Refer to Table 2 Exp B.III for error metrics

firing rate is shown in Figure 6 A. We included 2002 time points in the training, 250 in the validation, and 250 in the testing dataset. SHRED was trained to predict the LFP and reconstruct the firing rate.

From just the neuron firing rate, SHRED can predict local field potential (Table 2(B.III)). The SHRED reconstruction of the entire dataset shown in Figure 6 closely follows the data. While there are some deviations between the predicted and the true LFP, overall activity trends are well-preserved. It is evident that the firing rate is dynamically coupled to the LFP, making this a suitable setting for applying SHRED. Based on the conclusions of this experiment, we hypothesized that any dataset that is dynamically coupled to neural recordings should be sufficient to reconstruct population-wide neural activity, even if it is simply a proxy of neural activity. Therefore, in the next experiment outlined in Section B.2, we reconstructed LFP data from just the pupil area recording.

B.2 SHRED Reconstruction of Local Field Potential from Pupil Area Response to a Visual Stimulus

Previous experiments demonstrated SHRED's ability to reconstruct neural activity from inputs directly correlated with neural activity. Here, we extend this approach by attempting to reconstruct LFP recordings from the time series of the pupil area. We hypothesized that pupil size reflects central nervous system activity, consistent with the proposed model in [39]. Because LFPs represent neural responses evoked by visual stimuli that are first encoded through the retina and gated by the light passing through the pupil, we proposed that the pupil area should be dynamically coupled to the evoked neural response. Therefore, in this experiment, we aimed to reconstruct LFP recordings from just the pupil area measurements. Probe A LFP and pupil area recordings were obtained from the first presentation of the 2-second, $180°$, 0.8 contrast drifting grating stimulus. To account for the different sampling rates, we applied a cubic spline to the pupil area measurements. The pupil area recordings were upsampled, to test whether finer-scale trends in the LFP data could be reconstructed. Thus pupil recordings were obtained for the same timepoints as the LFP recordings. We split the data into 75% train, 12.5% test, and 12.5% validation. With a lag of 50, pupil area recordings served as the input to the model which was trained to predict the LFP and reconstruct the pupil area recordings.

The reconstruction of population dynamics from pupil area introduces additional challenges compared to the tasks outlined in the previous section. Pupil area is not a direct correlate of neural activity. However, SHRED performed surprisingly well. The reconstruction shown in Figure 7 deviates from the ground truth, but the peaks and troughs are still maintained. Train and test MSE are presented in Table 2 (B.IV). While there is still more work needed to improve the prediction accuracy, this baseline result highlights the versatility and robustness of the SHRED architecture.

Applications of SHRED to the mouse dataset highlighted its diverse usage. Not only is SHRED capable of reconstructing population neural activity from a few sensors, it is also capable of reconstructing

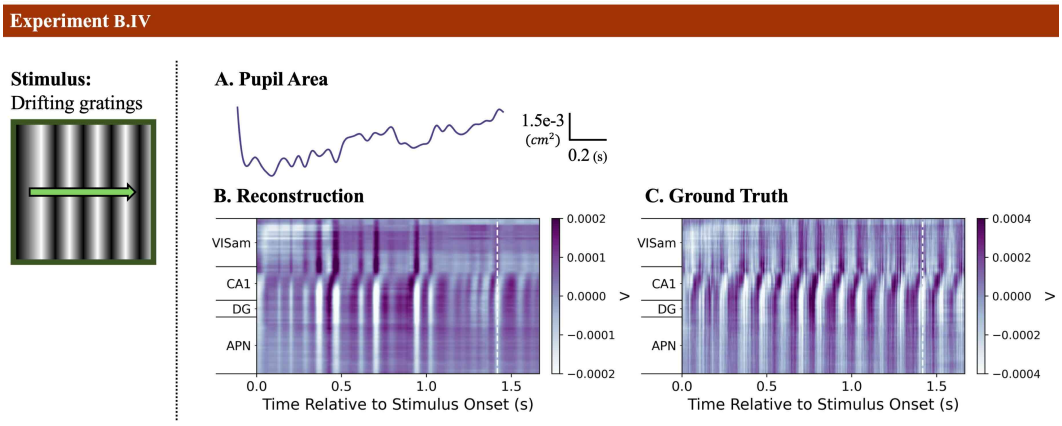

Figure 7: **SHRED has potential to reconstruct Local Field Potential (LFP) activity from pupil area. (A)** Pupil area in response to a 2-s drifting grating stimulus served as the input to the SHRED model. **(B)** The SHRED model was trained to reconstruct LFP activity from a probe inserted in the anteromedial (AM) region of the mouse visual cortex. The reconstruction successfully captures coarse fluctuations in spatial and temporal activity. Red dots on the y-axis indicate sensors selected as input. Refer to Table 2 Exp B.IV for error metrics

population dynamics from any coupled measurements. While cross-individual predictions were not feasible due to limitations of the dataset, individual reconstructions still offer significant value. They can help reduce the need for invasive recording procedures and even enable recovery of data from faulty or missing sensors. A model trained to predict full population activity from recordings in a single brain region, or even from proxy measures such as pupil area, which are easy to acquire, opens the door to a wide range of applications in neuroscience and clinical settings. By learning the patterns and connections in these datasets, SHRED was capable of predicting general brain activity from unit firing rate and pupil area measurements. To further explore the versatility of SHRED, we next reconstructed a video of the zebrafish forebrain activity.

B.3 Zebra Fish

The zebrafish, a small freshwater species in the minnow family, is a widely used model organism for studying vertebrate development and, more recently, systems neuroscience [24, 11]. Zebrafish offer several advantages in research, one of which is their transparency during early developmental stages, enabling clear and accessible imaging [24, 32]. Recent studies have focused on understanding neuronal ensembles code for sensorimotor pathways [3, 14, 43, 19, 8]. With improvements in imaging technologies, high-speed light-sheet microscopy can now be implemented to simultaneously record from almost all neurons in the zebrafish. This large scale recording allows for complex analysis of activity patterns in the whole brain [1].

For our third model system, we looked at a video of forebrain activity of a zebra fish obtained using high-speed light-sheet microscopy. This video was made publicly available by [1]. Neural projections were genetically encoded using a calcium indicator GCaMP5G. The 1041.859-second video contained 750 frames indicating a frame rate of approximately 0.72 Hz. We then converted the video to grayscale and cropped each frame to include 472×600 pixels thus a total of 283,200 pixels. The original video (forebrain.avi) was cropped to only include the left-hand side which shows the dorsal maximum-intensity projections of whole-brain functional activity and maximum-intensity projections of the reference anatomy (gray).

Variance pre-processing was applied to this dataset to ensure that the background of the forebrain slice was not included as the input. There were 92,634 pixels in the HV group and 190,566 pixels in the LV group. We aimed to reconstruct this video from the timeseries of just 250 pixels. More sensors were needed as input compared to the previous sections due to the greater number of total pixels. However, even with this increase, we still only included ~ 0.9 % of the total sensors as input. We split the data to include 700 time points in the training set, 25 time points in the testing set, and 25 time points in the validation set and trained the model for 200 epochs.

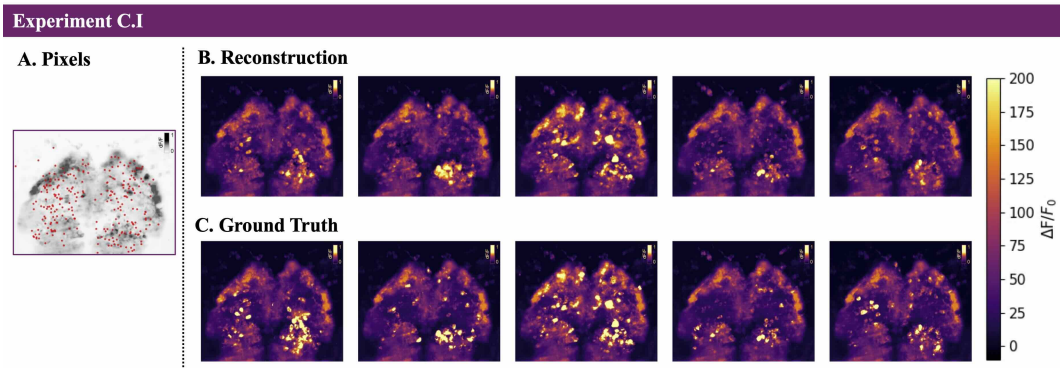

Figure 8: **SHRED accurately reconstructed a video of zebra fish forebrain activity. (A)** Pixels selected as input for the model. **(B/C)**SHRED reconstruction of 6 time points (top) follows general trends in activity of the original (bottom). Refer to Table 2 Exp C.I for error metrics.

SHRED struggles to reconstruct individual changes in pixels, but captures flares of activity with great accuracy (Figure 8). Train and test MSE are presented in Table 2 (C.i). The overall shape of the forebrain was preserved and regions with increased activity was reflected in the reconstruction. In this section, we showed that SHRED can be used to reconstruct videos. SHRED is versatile and can be applied to a variety of complex datasets with minimal hyperparameter tuning.

C Control Comparison

We also performed linear regression as a baseline for each experiment, using the same neurons or sensors provided to SHRED as input. SHRED outperformed this baseline in nearly all tasks based on mean squared error (MSE). However, MSE alone was insufficient to capture the differences and strengths of each model. In this section, we therefore evaluate performance using the power spectral density (PSD), which highlights the reconstructed frequencies relevant to the dataset. It should be noted that neural networks generally act as bandpass filters due to the spectral bias in learning when using MSE in training. Regardless, SHRED generally does a better job of preserving the frequency components and the associated PSD than simple linear models.

As shown in Figure 9, SHRED outperformed linear regression in both training and test reconstructions of the first *C. elegans* worm from Experiment A.I (Sec. 3.1.1). These results were obtained using variance pre-processing with three neurons selected from the high-variance group. With the same neurons, linear regression performed surprisingly well (Fig. 9 B, C2, D). Nonetheless, further analysis indicates that SHRED (1) is more robust to neuron selection and (2) reconstructs lower-frequency patterns more accurately.

We compared SHRED and linear regression under two experimental conditions: training on three neurons from either the low-variance or high-variance subsets. With low-variance input, the linear model fails to reconstruct the full state space, reflecting its inability to extrapolate beyond the given signals. In contrast, SHRED accurately reconstructs population activity, including high-variance neurons (Fig. 9 A, C1). Power spectral density (PSD) analysis further highlights this difference: the linear model struggles to capture lower-frequency components, whereas SHRED closely follows the ground truth, even with limited input. Both the PSD and individual traces also show that SHRED acts as a filter, suppressing high-frequency fluctuations that persist in the linear model.

When high-variance neurons are used as input, the performance of both models improves (Fig. 9 B, C2). However, the linear model benefits more from this optimal selection, underscoring SHRED's robustness to neuron choice. The PSD in C2 indicates that both models capture low-frequency patterns more accurately, yet SHRED still outperforms linear regression in both train/test MSE and PSD MSE (Fig. 9D). This advantage arises because SHRED reconstructs lower-frequency dynamics more precisely and on the correct timescale (note the log scale in C1 and C2).

For this dataset, the primary goal was to recover the general activity patterns and timing of neuronal activation. High-frequency fluctuations are likely noise and therefore less critical to reconstruct. If

both low- and high-frequency components were desired, a potential future direction would be to combine the strengths of the two models. For example, SHRED could be used for low-frequency dynamics and linear regression for high-frequency detail. Nevertheless, for the purposes of these experiments, SHRED consistently outperforms the linear regression baseline, even without optimal sensor input.

When applying these models in a biological context, robustness to sensor selection is critical. If future applications aim to support minimally invasive or non-invasive sensing, the objective is to reconstruct activity from sensors that are easy to access and record from, not necessarily the sensors with the highest quality recordings.

A linear regression baseline was also applied to the experiments in Sections 3.1–3.2, with power spectral densities (PSDs) and errors shown in Figure 10. SHRED outperforms the linear model in all experiments except B.I DG-A, B.II DG, and C.I-250 pixels. Experiment B.I DG-A corresponds to reconstruction of mouse LFP recorded on probe A during the drifting grating (DG) stimulus (Fig. 4), which exhibits a strong oscillatory pattern readily captured by linear transformations from the input timeseries. Similarly, Experiment B.II DG explores the average unit reconstruction during the DG stimulus (Fig. 5), where units display a sustained level of activity that can be well approximated by a line. These cases therefore favor linear regression, as reflected in the error metrics. By contrast, when more challenging datasets are considered for the same tasks (e.g., Exp B.I DG-C and Flash-A/C; Exp B.II Flash), SHRED outperforms linear regression. For the video reconstructions in Experiment C.I, the error of the linear model increases substantially when fewer pixels are used as input. Overall, SHRED demonstrates greater versatility and is less constrained by input quality or quantity. Similar to what was observed in Figure 9, lower frequency patterns in data are better reconstructed by SHRED for all experiments presented in this manuscript thus far. More extensive comparisons of SHRED to other models have been conducted in [52, 45].

D Sensitivity analysis for *C. elegans* data

We also conducted some preliminary tests to study the sensitivity of the SHRED model to 1) the activity of the neuron selected, 2) the number of neurons selected and 3) the data split. For these series of experiments, a limited number of neurons were selected and SHRED was trained to predict population activity. Only the first half of the time-series data for the first *C. elegan* (W1) was used to simplify analysis.

Out of the 109 neurons recorded from W1, approximately half exhibited minimal fluctuations in activity over the entire recording period (Figure 11 A-B). This suggests that a substantial proportion of neurons remained inactive throughout the trial thus the selection of these neurons for population-wide predictions with SHRED may result in sub-optimal results. Standard deviations across time for each of the neurons were calculated and neurons with activity levels above the mean were classified in the *High Variance* (HV) group and those below were classified in the *Low Variance* (LV) group. To test the effect of neuronal variability on SHRED performance, three neurons were randomly selected from each of the following groups: HV, LV, and the full population (Rand). SHRED was trained to reconstruct the entire population activity. We repeated these trials for 25 independent runs and the corresponding train and test Mean Squared Error (MSE) were calculated.

Random selection from the HV group yielded the best overall performance. Across 25 independent runs, the averaged prediction for the HV condition (Train: 0.159 ± 0.004 Test: 0.504 ± 0.010) most closely matched the ground truth and exhibited lower variability in performance compared to the LV (Train: 0.241 ± 0.012 Test: 0.799 ± 0.049) and Rand conditions (Train: 0.202 ± 0.011 Test: 0.657 ± 0.032) as shown in Figure 11 C - D. While randomly selecting from the entire population resulted in lower overall MSE compared to the LV condition, there is still a substantial variability in accuracy across runs. Variance pre-processing is not a prerequisite for SHRED [27], but its utility depends on the dataset. Unlike other applications, primarily in physics, neural activity can often be represented using only a small subset of neurons or channels. Population coding allows complex and dynamic tasks to be captured by the activity of a fraction of neurons. Since SHRED relies on inputs that are dynamically coupled with other recordings, performance naturally declines when inactive neurons or channels are included. Thus, we recommend testing both variance-based selection and random sampling to assess potential performance differences.

Previous applications of SHRED —such as in sea surface temperature, turbulent fluid flow, and magnetoencephalography (MEG) recordings— have demonstrated that as few as three sensors are sufficient to reconstruct whole-system dynamics [52]. To support these results, we then tested the effect of the number of neurons selected on the reconstruction of this immobilized *C. elegan* neural dataset. A total of 15 independent runs were conducted for trials selecting from one, up to five neurons from the high variance group. Train and test MSE were calculated and averaged across trials.

As expected, three neurons are sufficient for accurate reconstruction of the entire population dynamics using SHRED. As shown in Figure 12, the prediction error plateaus after around 3 neurons. There is no substantial additional benefit when including more than three neurons. Analogous to how three sensors are sufficient to localize objects in two dimensions, reconstruction of spatio-temporal data also seems to improve with the inclusion of three sensors.

We also tested the effect of data splits on SHRED accuracy. For convention, an $x : y$ split in data represents a split with $x\%$ train, $y/2\%$ validation, and $y/2\%$ test. We conducted five independent runs each for the following splits {95:5, 90:10, 85:15, 80:20, 70:30, 60:40, 50:50,40:60,30:70} and averaged the train and test MSEs. Based on the results discussed above, we selected three high variance neurons for this analysis. The data split does not seem to have a significant impact on the accuracy of the reconstruction of the training or testing data. Since an $80 : 10 : 10$ split is customary, we selected this split for the majority of the experiments conducted for the nematode and the mice.

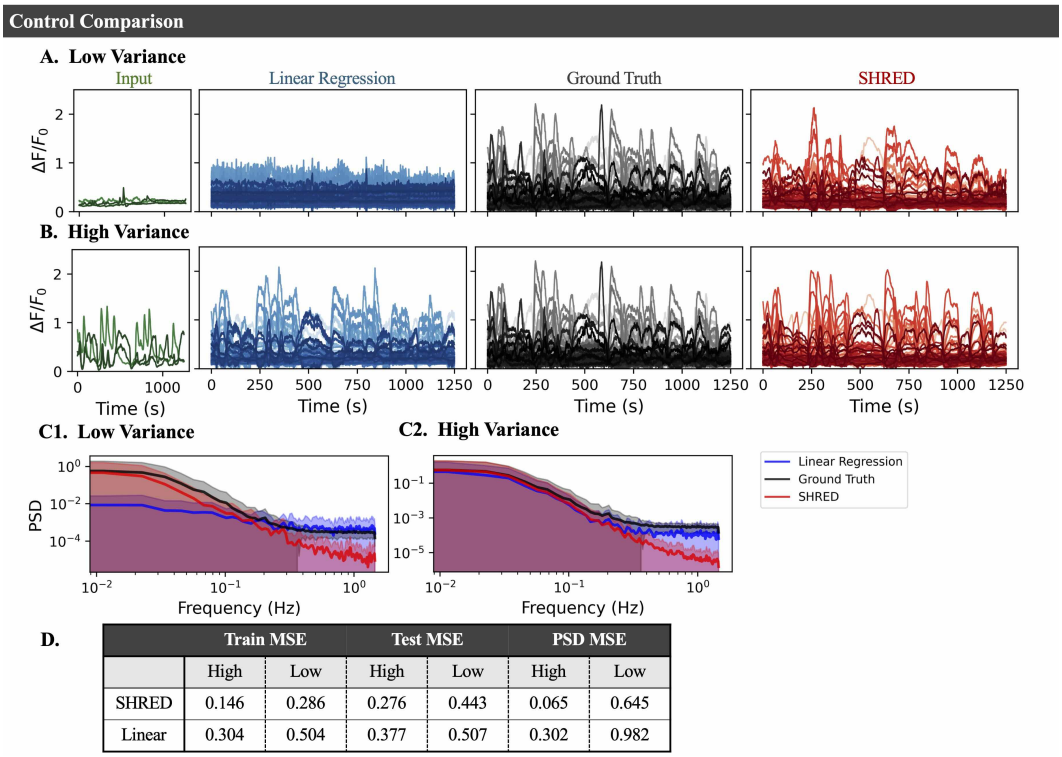

D.

	Train MSE		Test MSE		PSD MSE	
	High	Low	High	Low	High	Low
SHRED	0.146	0.286	0.276	0.443	0.065	0.645
Linear	0.304	0.504	0.377	0.507	0.302	0.982

Figure 9: **SHRED outperforms a linear regression baseline. (A–B)** Reconstructions from three neurons selected from either the high-variance (A) or low-variance (B) groups, showing ground truth activity alongside linear regression and SHRED outputs. **(C)** Power spectral densities of SHRED (red) and linear regression (blue) reconstructions compared to the ground truth (grey). **(D)** Summary of error metrics

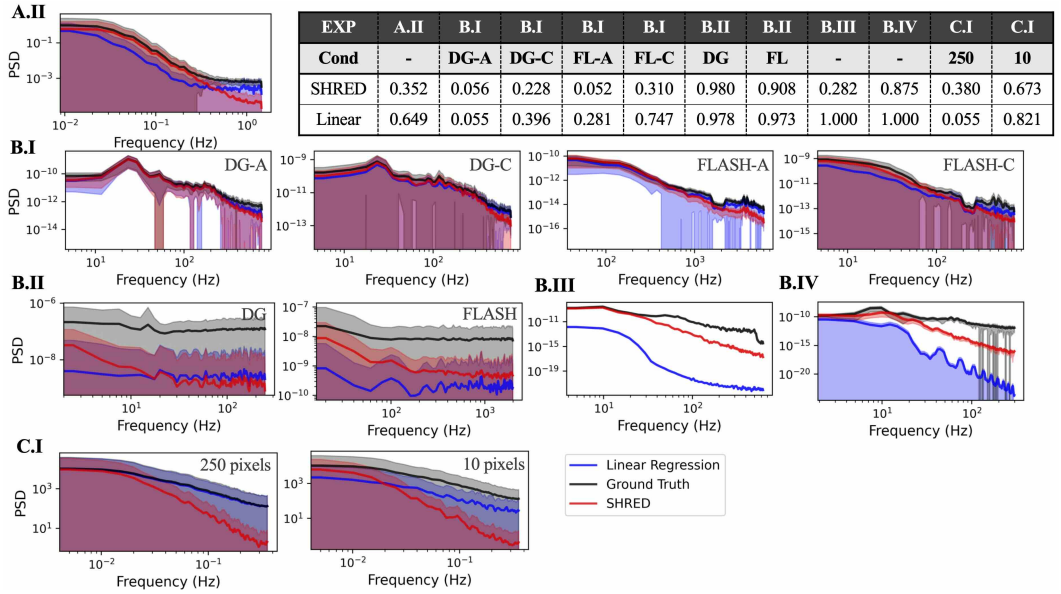

EXP	A.II	B.I	B.I	B.I	B.I	B.II	B.II	B.III	B.IV	C.I	C.I
Cond	-	DG-A	DG-C	FL-A	FL-C	DG	FL	-	-	250	10
SHRED	0.352	0.056	0.228	0.052	0.310	0.980	0.908	0.282	0.875	0.380	0.673
Linear	0.649	0.055	0.396	0.281	0.747	0.978	0.973	1.000	1.000	0.055	0.821

Figure 10: **Baseline comparisons.** Power spectral densities of SHRED (red) and linear regression (blue) reconstructions compared to the ground truth (grey) for each experiment. Error metrics are presented in the table at the top. DG = drifting grating.

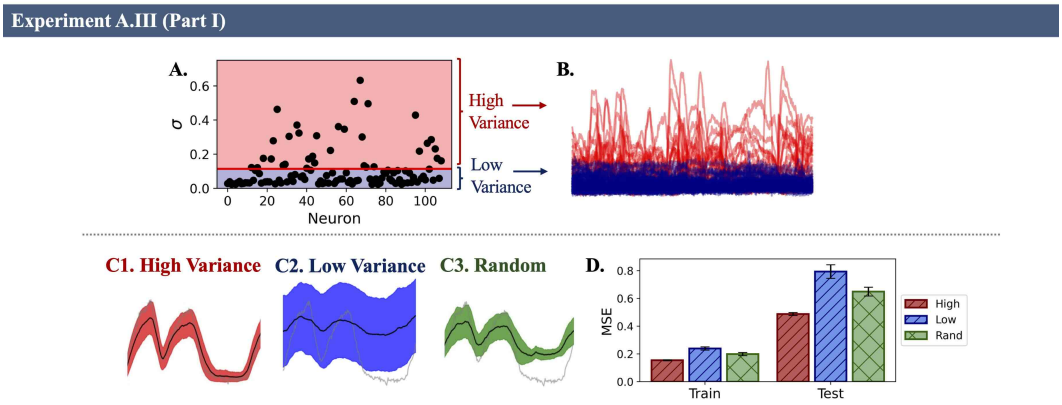

Figure 11: **Selecting neurons with highly variable activity improves SHRED performance. (A)** Scatterplot of 109 neurons (black dots) showing variability in activity, measured as the standard deviation across 1,568 time points. The red line indicates the mean standard deviation. Neurons with values above the average were classified as the *High Variance* (HV) group and those below were classified as the *Low Variance* (LV) group. **(B)** Activity traces from HV neurons (red) and LV neurons (blue) illustrating differences in signal variability. **(C)** A randomly selected neuron trace (grey) and average SHRED reconstruction (black) from 25 trials. Shaded region indicates ± 1 SD. **(D)** Average train and test Mean Squared Error (MSE) for 25 trials using 3 neurons sampled from the HV, LV, or full population (Random).

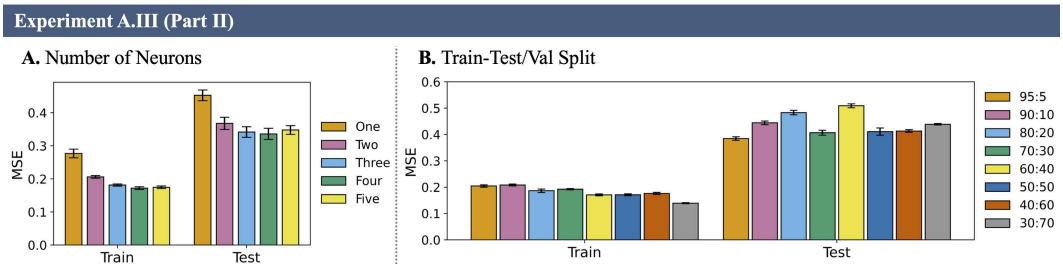

Figure 12: **Number of neurons and data split influence SHRED performance. (A)** Train and test MSE decrease with neuron count but plateau after three neurons. **(B)** The data split primarily affects reconstruction accuracy on the test dataset, with minimal impact on the training. A training and test+validation split of 80:10:10 resulted in the best performance. All trials conducted with variance pre-processing.