
Transient Neural Dynamics Reconstruction

Anima Kujur¹ Zahra Monfared^{1,2} Shervin Safavi^{3,4}

{anima.kujur, zahra.monfared}@iwr.uni-heidelberg.de, shervin.safavi@cmclab.org

¹Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University, Heidelberg, Germany

²Department of Mathematics and Computer Science, Heidelberg University, Heidelberg, Germany

³Computational Neuroscience, Department of Child and Adolescent Psychiatry,

Faculty of Medicine, Technische Universität Dresden, Dresden 01307, Germany

⁴Department of Computational Neuroscience, Max Planck Institute for Biological Cybernetics, Tübingen, Germany

Abstract

Sharp wave-ripples (SWRs) are high-frequency ($\sim 100\text{--}250$ Hz) oscillatory bursts often observed in hippocampal local field potentials (LFPs), and are involved in a wide range of cognitive functions (memory consolidation to off-line and online planning). Reconstructing LFPs in the SWR regime is challenging due to the complexity of signals and the transient nature of these bursts. While many algorithms provide reasonable short-term predictions, most fail to reproduce long-term dynamics while preserving fast transients. In this study, we combine principal component analysis (PCA) with dynamic mode decomposition (DMD) to approximate the Koopman operator in a reduced latent space, allowing for the efficient reconstruction of multi-channel hippocampal LFPs. The Koopman framework shifts the focus from state-space trajectories to observable evolution in an infinite-dimensional function space, enhancing interpretability and understanding of nonlinear systems. Using 200,000 samples, our PCA-DMD framework achieves superior reconstruction accuracy compared to state-of-the-art DMD variants. Our results highlight PCA-DMD’s robustness in capturing complex neural dynamics and offer a powerful tool for analysis of transient dynamics (e.g., SWR) with significant implications for neuroscience research and clinical applications.

1 Introduction

Local field potentials (LFPs) reflect the extracellular activity of neuronal ensembles [36, 9, 14, 35], providing a mesoscopic representation of circuit dynamics [9, 44, 12, 45]. In the hippocampus, LFPs are particularly informative during sharp wave-ripples (SWRs) [20, 19, 28], consisting of brief high-frequency oscillations ($\sim 100\text{--}250$ Hz) superimposed on sharp waves ($\sim 50\text{--}100$ ms). SWRs reflect synchronous neuronal firing with a characteristic multi-scale dynamics [7, 32, 23, 1, 43, 40], which are critical for a wide range of cognitive functions, for instance, episodic memory consolidation and planning [20, 27, 8]. Accurate reconstruction of LFPs with SWRs restores latent neural dynamics that noise or artifacts can mask. This is important for both basic neuroscience and clinical settings where pathological ripples disrupt cognition in epilepsy, schizophrenia, and Alzheimer’s disease [8]. Beyond recovery, reconstruction also enables the extraction of dynamical features such as Lyapunov exponents, Koopman modes, and eigenfunctions, which are valuable for mechanistic analysis and signal classification [50]. Prior work has largely emphasized *detection* and *prediction* of SWRs, often leveraging machine learning (ML) and deep learning (DL) techniques [31, 29]. Representative examples include convolutional neural networks (CNN) for automated SWR detection [37], recurrent architectures such as RippleNet [15], consensus guidelines distinguishing SWRs from high-frequency oscillations [30], and spatiotemporal models differentiating SWRs from epileptiform discharges [34]. Additional studies focus on feature extraction: topological analyses of SWR waveforms [47], ML

toolboxes for cross-species analysis [39], artifact removal via SVD-based filtering [11], and dynamical network models for EEG/MEG reconstruction [10]. Despite these advances, direct reconstruction of LFP time series in the SWR regime remains underexplored. Dynamic mode decomposition (DMD) [46] and Koopman operator theory [17] provide a principled data-driven framework to approximate nonlinear dynamics by decomposing spatiotemporal data into coherent modes with interpretable frequencies and growth rates. These techniques have uncovered hidden structures in diverse systems, such as magnetohydrodynamic plasma states [48], and are well suited for hippocampal dynamics, where SWRs originate from recurrent excitatory circuits in CA3 [45]. However, applications of DMD to hippocampal LFP reconstruction remain limited [19, 42]. A key challenge is the scale and complexity of LFP data: recordings span many samples across multiple channels, and dynamics are transient and fairly complex. Many DL models, such as recurrent neural networks (RNNs), struggle with LFP-scale data due to exploding/vanishing gradients in long sequences, excessive memory and compute requirements, and poor alignment with transient oscillations. Windowing can mitigate memory demands by segmenting recordings, but downsampling the signal causes information loss. For very long signals it fails to capture dependencies across distant windows, leading to loss of global temporal structure. In high-dimensional settings, architectures like Sampled RNNs [4] can change the frequency content of the ground-truth signal, creating distorted oscillations that reduce quality in the SWR band. Moreover, DL-based approaches generally provide *limited interpretability* compared to Koopman-based and DMD methods, which yield explicit spectral decompositions. Structured latent recurrent models, such as shPLRNN [16], also face challenges for long reconstructions. Their performance is highly sensitive to hyperparameter choices (e.g., latent dimension and teacher forcing parameter). Memory usage further scales poorly with sequence length, limiting applicability to large-scale LFP data. Similarly, existing DMD variants—including Classical DMD [24, 46], Sparse DMD (SpDMD) [21], multi-resolution DMD (MrDMD) [25], and higher-order DMD (HODMD) [26]—struggle to preserve high-frequency content, require large memory, or become numerically unstable on long windows. These limitations highlight the need for a *scalable* and *interpretable* reconstruction method for oscillatory LFP transients.

In this study, we develop a PCA-DMD framework (Fig. 1) for multichannel LFP reconstruction, combining principal component analysis (PCA) [33] with Koopman-based modeling. PCA projects high-dimensional LFP segments into a low-rank subspace, and DMD approximates the Koopman operator to capture linear evolution in this latent space. An overlap-add reconstruction ensures temporal continuity and retention of transient oscillations. Our pipeline is mathematically related to Hankel-/delay-based DMD [3, 49] and HAVOK [5], but differs in three key aspects: (i) we apply PCA as a separate pre-compression step on vectorized windows (rather than a single Hankel SVD), (ii) we fit the Koopman operator in this PCA latent space for improved stability and reduced regression size, and (iii) we reconstruct the full signal using overlap-add with tapering to preserve transient SWR oscillations; see also Table 1. To the best of our knowledge, this explicit PCA–latent Koopman–overlap-add workflow has not appeared verbatim in prior work. Its application to large-scale hippocampal LFPs makes it impactful for both scale and fidelity. Our contributions are: (1) Demonstrating reconstruction of hippocampal LFPs up to 200k samples, beyond the practical reach of typical RNN or standard DMD approaches. (2) Showing that ablation-guided tuning of window size, step size, and latent dimension allows PCA-DMD to preserve SWR oscillations (100–250 Hz). (3) Benchmarking against Classical DMD, SpDMD, MrDMD, and HODMD via pyDMD [13, 18], and leveraging advances in Koopman analysis [2, 22, 6]. PCA-DMD achieves superior reconstruction accuracy, enhancing the recovery of oscillatory transients, which is central to understand their role in replay and decision-making [20].

2 Data

We used hippocampal LFP recordings from rodents for SWR detection [35], sourced from [38]. The dataset includes eight channels of LFP signals, annotated with SWR events, recorded at 30 kHz to capture high-frequency oscillations (100–250 Hz). Each channel has 22,326,272 samples, totaling 744 seconds. For analysis, we utilized a subset of 200k samples, with a maximum of 800k for successful reconstruction. Data were preprocessed by segmenting the signal into overlapping windows, with a window size of 3,000 samples and a step size of 30 samples determined through an ablation study to optimize performance for the applied dynamical systems methods (see Appendix E).

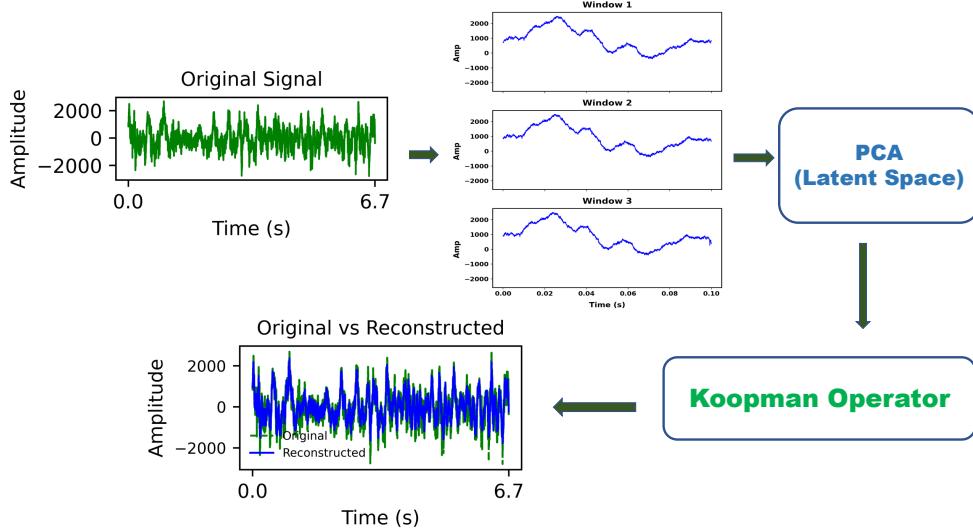


Figure 1: Overview of the PCA–Koopman reconstruction pipeline applied to hippocampal LFP signals. The raw LFP segment is first divided into short overlapping windows. These windowed observations are then projected into a lower-dimensional latent space using PCA, providing a compact representation of the system’s dominant spatial–temporal patterns. The Koopman operator is subsequently learned in this latent space, enabling linear evolution of the nonlinear neural dynamics. Finally, the reconstructed signal is obtained by mapping them back to the original signal space. This pipeline allows efficient long-term reconstruction of LFP dynamics while preserving fast transient components relevant for sharp wave–ripples.

Table 1: DMD variants grouped by SVD workflow and relation to our PCA-DMD.

Method	Type	Key Notes / Relation
Classical DMD [24]	Snapshot-SVD	Koopman regression on snapshot basis.
SpDMD [21]	Snapshot-SVD	Adds sparse mode selection; otherwise same as Classical.
MrDMD [25]	Window-SVD	Multi-resolution partitioning; trades global coherence for local detail.
HODMD [26]	Window-SVD	Delay embedding; captures high-frequency oscillations but unstable for long windows.
Hankel-DMD [3]	Window-SVD	Hankel matrix from delays; foundation of HAVOK.
HAVOK [5]	Window-SVD	Hankel embedding + forcing term; chaotic signals.
Our PCA-DMD	Windowing + PCA	PCA pre-compression + Koopman in latent space + overlap-add reconstruction. Improves stability, scalability, and transient SWR preservation.

3 PCA-DMD: A Koopman-based Multichannel LFP Reconstruction

We propose a PCA-DMD method for reconstructing multichannel LFP signals. The raw LFP recordings are first segmented into overlapping windows, then reduced in dimensionality using PCA to obtain compact latent representations. A linear Koopman operator (Appendix A) is subsequently learned in this latent space to capture the temporal evolution of the dynamics across channels. Finally, the latent states are mapped back to the original signal space through the inverse PCA projection, and the reconstructed windows are stitched together to approximate the full multichannel LFP signal while preserving SWR dynamics. Let the multichannel LFP signal be denoted by $s = \{s_t\}_{t=1}^T$, $s_t \in \mathbb{R}^m$, where m is the number of channels and T the total number of samples. We segment the signal into overlapping windows of length w with step size δ :

$$X_i = [s_i^\top \ s_{i+1}^\top \ \cdots \ s_{i+w-1}^\top]^\top \in \mathbb{R}^{w \times m}, \quad i = 1, 1 + \delta, \dots$$

Each window is vectorized into $x_i = \text{vec}(X_i) \in \mathbb{R}^{wm}$, and the N windows together form the dataset

$$X = [x_1^\top \ x_2^\top \ \cdots \ x_N^\top]^\top \in \mathbb{R}^{N \times (wm)}.$$

We compute a rank- d PCA projection $W \in \mathbb{R}^{(wm) \times d}$ with $d \ll wm$. The latent representations are

$$Z = XW = [z_1^\top \ z_2^\top \ \cdots \ z_N^\top]^\top \in \mathbb{R}^{N \times d}, \quad z_i \in \mathbb{R}^d.$$

Define the shifted latent pairs as $Z_{\text{past}} = [z_1 \ z_2 \ \cdots \ z_{N-1}]$ and $Z_{\text{next}} = [z_2 \ z_3 \ \cdots \ z_N]$. We model the temporal evolution in the latent space as $z_{i+1} \approx \mathbf{K}z_i$, $\mathbf{K} \in \mathbb{R}^{d \times d}$ where the optimal \mathbf{K} is obtained via least-squares: $\mathbf{K} = \arg \min_{A \in \mathbb{R}^{d \times d}} \|Z_{\text{next}} - AZ_{\text{past}}\|_F = Z_{\text{next}}Z_{\text{past}}^\dagger$ with \dagger the Moore–Penrose pseudoinverse. Then, the predicted latent states evolve as $\hat{z}_{i+k} = \mathbf{K}^k z_i$. These are mapped back to the window space via the inverse PCA projection $\hat{x}_i = \hat{z}_i W^\top \in \mathbb{R}^{wm}$ and $\hat{X}_i = \text{unvec}(\hat{x}_i) \in \mathbb{R}^{w \times m}$. Finally, the reconstructed multichannel signal $\hat{s}_t \in \mathbb{R}^m$ is obtained by overlap-add of the windows \hat{X}_i , with a tapering function (e.g., Hann) to reduce boundary artifacts. See Table 1 for a comparison of our method with other state-of-the-art DMD variants.

4 Results

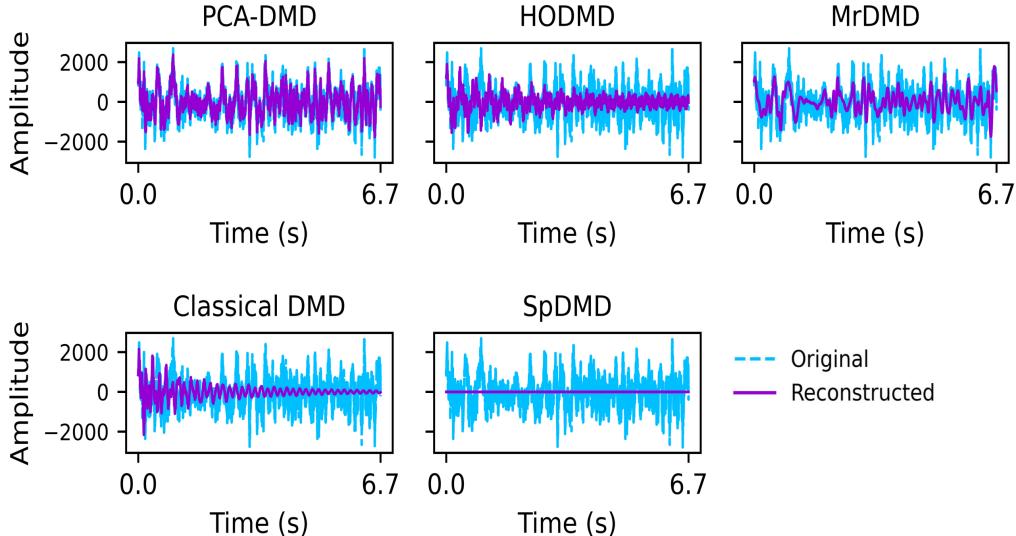


Figure 2: Original and reconstructed hippocampal LFP signals using five DMD methodologies. PCA-DMD achieves the closest match to the ground-truth signal, preserving the underlying dynamics of sharp wave-ripples. In contrast, HODMD, Classical DMD, MrDMD, and SpDMD show limitations in capturing the dynamics required for accurate long-term LFP reconstruction.

The comparative evaluation in Table 2 shows that PCA-DMD significantly outperforms Classical DMD, SpDMD, MrDMD, and HODMD in LFP signal reconstruction. Performance was quantified using Kullback–Leibler divergence (KLD) and Hellinger distance (HD); see Appendix D. Visual analysis of overlay line plots in Fig. 2 and the heatmaps in Fig. 3 further supports these findings: the PCA-DMD heatmap closely mirrors the amplitude patterns of the original signal in all eight channels and time samples, accurately capturing spatial (channel-wise) and temporal structures, while the other methods exhibit distortions and loss of high-frequency structure. These results underscore the efficacy of combining PCA-based dimensionality reduction with linear Koopman operator modeling, positioning our approach as a potential method for multi-channel LFP signal reconstruction. Ablation experiments further confirm these findings. Reducing the SVD rank from 8 to 5 caused significant degradation in all methods, as indicated by higher KLD and HD values. When scaling to 800,000 samples, only PCA-DMD preserved high fidelity, closely replicating the original signal’s spatiotemporal patterns, while other methods faced computational challenges. For more details, see Appendix F. Moreover, Koopman modes are illustrated in Appendix G. These modes offer complementary insights into the LFP signal by capturing both its spectral content (oscillatory frequencies) and its temporal dynamics.

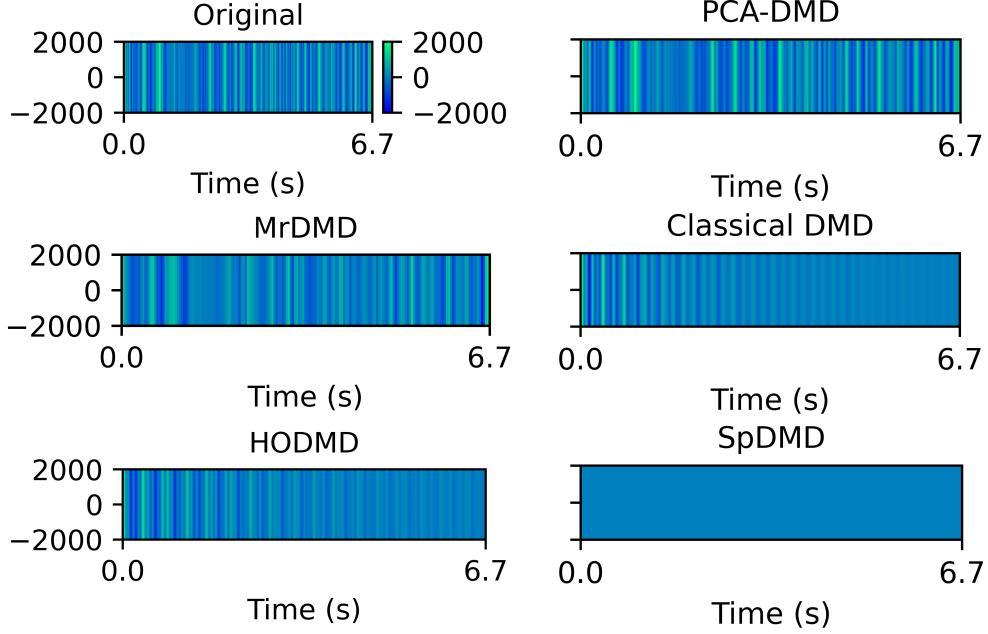


Figure 3: Comparison of the original signal with reconstructed methods for 200,000 samples.

Table 2: Comparison with State of the Art Methods

Methods	KLD	HD
Classical DMD	0.8710	0.4067
MrDMD	0.4852	0.2238
HODMD	0.6135	0.3050
SpDMD	14.4061	0.9074
PCA-DMD	0.0761	0.0847

5 Conclusion

We presented a PCA-DMD framework for multichannel LFP reconstruction in the hippocampal SWR regime. On 200k-sample recordings, PCA-DMD consistently outperformed state-of-the-art DMD methods in both statistical (KLD) and temporal (HD) measures, demonstrating its capacity to recover oscillatory transients and spatiotemporal structure more faithfully. The findings suggest that embedding Koopman operator learning into a reduced latent space is a powerful framework for uncovering the intrinsic low-dimensional dynamics underlying neural field potentials. A natural extension of this work is scaling the framework to reconstruct longer temporal sequences, up to the full dataset (~ 22 M samples), which will be crucial for evaluating stability, generalization, and robustness in real-world large-scale recordings.

Limitations: PCA-DMD may overlooks some complexities of LFPs [14], which can be further enhanced by incorporating recent developments in neural manifold learning [41]. The approach also requires careful hyperparameter tuning (window size, step size, latent dimension), and large-scale PCA computations are computationally demanding. Despite these constraints, PCA-DMD offers a scalable and interpretable alternative to DL approaches for reconstructions of long neural signals with fine-grained transient dynamics.

Acknowledgements

Z.M. is grateful to the Bundesministerium für Forschung, Technologie und Raumfahrt (BMFTR, Federal Ministry for Research, Technology and Space) for funding through project OIDLITDSM, No. 01IS24061.

References

- [1] J Abadchi, Mojtaba Nazari-Ahangarkolae, Sandra Gattas, Edgar Bermudez-Contreras, Artur Luczak, Bruce L McNaughton, Majid H Mohajerani, et al. Spatiotemporal patterns of neocortical activity around hippocampal sharp-wave ripples (*biorxiv*). 2020.
- [2] Daniel J Alford-Lago, Christopher W Curtis, Alexander T Ihler, and Opal Issan. Deep learning enhanced dynamic mode decomposition. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 32(3), 2022.
- [3] Hassan Arbabi and Igor Mezić. Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the koopman operator. *SIAM Journal on Applied Dynamical Systems*, 16(4):2096–2126, 2017.
- [4] Erik Lien Bolager, Ana Cukarska, Iryna Burak, Zahra Monfared, and Felix Dietrich. Gradient-free training of recurrent neural networks. *arXiv preprint arXiv:2410.23467*, 2024.
- [5] Steven L Brunton, Bingni W Brunton, Joshua L Proctor, Eurika Kaiser, and J Nathan Kutz. Chaos as an intermittently forced linear system. *Nature communications*, 8(1):19, 2017.
- [6] Steven L Brunton, Marko Budišić, Eurika Kaiser, and J Nathan Kutz. Modern koopman theory for dynamical systems. *arXiv preprint arXiv:2102.12086*, 2021.
- [7] György Buzsáki. Neural syntax: cell assemblies, synapsembles, and readers. *Neuron*, 68(3):362–385, 2010.
- [8] György Buzsáki. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10):1073–1188, 2015.
- [9] György Buzsáki, Costas A Anastassiou, and Christof Koch. The origin of extracellular fields and currents—eeg, ecog, lfp and spikes. *Nature reviews neuroscience*, 13(6):407–420, 2012.
- [10] Miao Cao, Daniel Galvis, Simon J Vogrin, William P Woods, Sara Vogrin, Fan Wang, Wessel Woldman, John R Terry, Andre Peterson, Chris Plummer, et al. Virtual intracranial eeg signals reconstructed from meg with potential for epilepsy surgery. *Nature communications*, 13(1):994, 2022.
- [11] Long Chen, Zhebing Ren, and Jing Wang. An svd-based method for dbs artifact removal: High-fidelity restoration of local field potential. *Biomedical Signal Processing and Control*, 108:107908, 2025.
- [12] Jan Cimbálník, JS Garcia-Salinas, N Hamedí, S Prathapagiri, LF Sarmiento Rivera, M Galanina, J Dolezal, L Jurkovicova, Pavel Daniel, Martin Kojan, et al. Human brain local field potential recordings during a battery of multilingual cognitive and eye-tracking tasks. *Scientific Data*, 12(1):899, 2025.
- [13] Nicola Demo, Marco Tezzele, and Gianluigi Rozza. PyDMD: Python dynamic mode decomposition. *Journal of Open Source Software*, 3(22):530, 2018. doi: 10.21105/joss.00530. URL <https://doi.org/10.21105/joss.00530>.
- [14] Gaute T Einevoll, Christoph Kayser, Nikos K Logothetis, and Stefano Panzeri. Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nature Reviews Neuroscience*, 14(11):770–785, 2013.
- [15] Espen Hagen, Anna R Chambers, Gaute T Einevoll, Klas H Pettersen, Rune Enger, and Alexander J Stasik. Ripplenet: a recurrent neural network for sharp wave ripple (spw-r) detection. *Neuroinformatics*, 19(3):493–514, 2021.
- [16] Florian Hess, Zahra Monfared, Manuel Brenner, and Daniel Durstewitz. Generalized Teacher Forcing for Learning Chaotic Dynamics. In *Proceedings of the 40th International Conference on Machine Learning*. PMLR, July 2023. ISSN: 2640-3498.
- [17] Zhenxing Hu, Manoj Aravind, Xu Lei, J Nathan Kutz, and Jean-Julien Aucouturier. Learning the bistable cortical dynamics of the sleep-onset period. *bioRxiv*, pages 2025–07, 2025.

[18] Sara M. Ichinaga, Francesco Andreuzzi, Nicola Demo, Marco Tezzele, Karl Lapo, Gianluigi Rozza, Steven L. Brunton, and J. Nathan Kutz. Pydmd: A python package for robust dynamic mode decomposition. *Journal of Machine Learning Research*, 25(417):1–9, 2024. URL <http://jmlr.org/papers/v25/24-0739.html>.

[19] Seung-Woo Jin, Hee-Seung Ha, and Inah Lee. Selective reactivation of value-and place-dependent information during sharp-wave ripples in the intermediate and dorsal hippocampus. *Science Advances*, 10(32):eadn0416, 2024.

[20] Hannah R Joo and Loren M Frank. The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. *Nature Reviews Neuroscience*, 19(12):744–757, 2018.

[21] Mihailo R Jovanovic, Peter J Schmid, and JW Nichols. Low-rank and sparse dynamic mode decomposition. *Center for Turbulence Research Annual Research Briefs*, 2012:139–152, 2012.

[22] Eurika Kaiser, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of koopman eigenfunctions for control. *Machine Learning: Science and Technology*, 2(3):035023, 2021.

[23] Raphael Kaplan, Mohit H Adhikari, Rikkert Hindriks, Dante Mantini, Yusuke Murayama, Nikos K Logothetis, and Gustavo Deco. Hippocampal sharp-wave ripples influence selective activation of the default mode network. *Current Biology*, 26(5):686–691, 2016.

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

[25] J. Nathan Kutz, Xing Fu, and Steven L. Brunton. Multiresolution dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 15(2):713–735, 2016. doi: 10.1137/15M1023543. URL <https://doi.org/10.1137/15M1023543>.

[26] Soledad Le Clainche and José M Vega. Higher order dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 16(2):882–925, 2017.

[27] Timothy K Leonard, Jonathan M Mikkila, Emad N Eskandar, Jason L Gerrard, Daniel Kaping, Shaun R Patel, Thilo Womelsdorf, and Kari L Hoffman. Sharp wave ripples during visual exploration in the primate hippocampus. *Journal of Neuroscience*, 35(44):14771–14782, 2015.

[28] Zhenrui Liao, Kevin C Gonzalez, Deborah M Li, Catalina M Yang, Donald Holder, Natalie E McClain, Guofeng Zhang, Stephen W Evans, Mariya Chavarha, Jane Simko, et al. Functional architecture of intracellular oscillations in hippocampal dendrites. *Nature Communications*, 15(1):6295, 2024.

[29] Christos Panagiotis Lisgaras, Richard J Staba, and Liset M de la Prida. How can animal models advance research into high frequency oscillations: Guidelines for recording, detection and analysis. *Epilepsy Currents*, page 15357597251336792, 2025.

[30] Anli A Liu, Simon Henin, Saman Abbaspoor, Anatol Bragin, Elizabeth A Buffalo, Jordan S Farrell, David J Foster, Loren M Frank, Tamara Gedankien, Jean Gotman, et al. A consensus statement on detection of hippocampal sharp wave ripples and differentiation from other fast oscillations. *Nature communications*, 13(1):6000, 2022.

[31] Anli A Liu, Simon Henin, Saman Abbaspoor, Anatol Bragin, Elizabeth A Buffalo, Jordan S Farrell, David J Foster, Loren M Frank, Tamara Gedankien, Jean Gotman, et al. A consensus statement on detection of hippocampal sharp wave ripples and differentiation from other fast oscillations. *Nature communications*, 13(1):6000, 2022.

[32] Nikos K Logothetis, O Eschenko, Y Murayama, M Augath, T Steudel, HC Evrard, M Besserve, and A Oeltermann. Hippocampal–cortical interaction during periods of subcortical silence. *Nature*, 491(7425):547–553, 2012.

[33] Andrzej Maćkiewicz and Waldemar Ratajczak. Principal components analysis (pca). *Computers & Geosciences*, 19(3):303–342, 1993.

[34] Anna Maslarova, Jiyun Shin, Andrea Navas-Olive, Mihály Vöröslakos, Hajo Hamer, Arnd Doerfler, Simon Henin, György Buzsáki, and Anli Liu. Spatiotemporal patterns differentiate hippocampal sharp-wave ripples from interictal epileptiform discharges in mice and humans. *bioRxiv*, pages 2025–02, 2025.

[35] Kolja Meier, Andrea Merseburg, Dirk Isbrandt, Stephan Lawrence Marguet, and Fabio Morellini. Dentate gyrus sharp waves, a local field potential correlate of learning in the dentate gyrus of mice. *Journal of Neuroscience*, 40(37):7105–7118, 2020.

[36] Matthias Mölle, Oxana Yeshenko, Lisa Marshall, Susan J Sara, and Jan Born. Hippocampal sharp wave-ripples linked to slow oscillations in rat slow-wave sleep. *Journal of neurophysiology*, 2006.

[37] Andrea Navas-Olive, Rodrigo Amaducci, Maria-Teresa Jurado-Parras, Enrique R Sebastian, and Liset M de la Prida. Deep learning-based feature extraction for prediction and interpretation of sharp-wave ripples in the rodent hippocampus. *Elife*, 11:e77772, 2022.

[38] Andrea Navas-Olive, Adrian Rubio, Saman Abbaspour, Kari L Hoffman, and Liset M de la Prida. A machine learning toolbox for the analysis of sharp-wave ripples reveals common waveform features across species. *Communications Biology*, 7(1):211, 2024.

[39] Andrea Navas-Olive, Adrian Rubio, Saman Abbaspour, Kari L Hoffman, and Liset M de la Prida. A machine learning toolbox for the analysis of sharp-wave ripples reveals common waveform features across species. *Communications Biology*, 7(1):211, 2024.

[40] Noam Nitzan, Rachel Swanson, Dietmar Schmitz, and György Buzsáki. Brain-wide interactions during hippocampal sharp wave ripples. *Proceedings of the National Academy of Sciences*, 119(20):e2200931119, 2022.

[41] Matthew G Perich, Devika Narain, and Juan A Gallego. A neural manifold view of the brain. *Nature Neuroscience*, pages 1–16, 2025.

[42] Alain Rios, Minori Usui, and Yoshikazu Isomura. Modulation of hippocampal sharp-wave ripples by behavioral states and body movements in head-fixed rodents. *eneuro*, 12(7), 2025.

[43] Shervin Safavi, Nikos K Logothetis, and Michel Besserve. From univariate to multivariate coupling between continuous signals and point processes: a mathematical framework. *Neural computation*, 33(7):1751–1817, 2021.

[44] Shervin Safavi, Theofanis I Panagiotaropoulos, Vishal Kapoor, Juan F Ramirez-Villegas, Nikos K Logothetis, and Michel Besserve. Uncovering the organization of neural circuits with generalized phase locking analysis. *PLOS Computational Biology*, 19(4):e1010983, 2023.

[45] Natalie Schieferstein, Ana Del Toro, Roberta Evangelista, Barbara Imbroisci, Aarti Swaminathan, Dietmar Schmitz, Nikolaus Maier, and Richard Kempter. Propagation of sharp wave-ripple activity in the mouse hippocampal ca3 subfield in vitro. *The journal of physiology*, 602(19):5039–5059, 2024.

[46] Peter J Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28, 2010.

[47] Enrique R Sebastian, Juan P Quintanilla, Alberto Sánchez-Aguilera, Julio Esparza, Elena Cid, and Liset M de la Prida. Topological analysis of sharp-wave ripple waveforms reveals input mechanisms behind feature variations. *Nature neuroscience*, 26(12):2171–2181, 2023.

[48] Roy Taylor, J Nathan Kutz, Kyle Morgan, and Brian A Nelson. Dynamic mode decomposition for plasma diagnostics and validation. *Review of Scientific Instruments*, 89(5), 2018.

[49] Jonathan H. Tu, Clarence W. Rowley, Dirk M. Luchtenburg, Steven L. Brunton, and J. Nathan Kutz. On dynamic mode decomposition: Theory and applications, 2014. ISSN 2158-2491.

[50] Hassaan Wasim. Classification of noisy eeg signals in bci using rnns: A dynamical systems perspective. Master’s thesis, TUM School of Computation, Information and Technology, Munich, 2025.

Technical Appendices and Supplementary Material

The supplementary materials provided in this appendix offer extended information to complement the main text, including additional plots of the dataset utilized for LFP reconstruction experiments, in-depth overviews of implemented state-of-the-art methods, as well as additional figures, ablation studies, and implementation details to enhance reproducibility and provide deeper insights into the PCA-DMD framework.

The **PCA-DMD** approach for LFP signal reconstruction was evaluated against five state-of-the-art DMD methods: SpDMD, Classical DMD, HODMD, and MrDMD. Our method employs overlapping windowing of multi-channel LFP data, followed by dimensionality reduction to 8 components using PCA and learning a linear Koopman operator via regression to predict subsequent states, with full signal reconstruction achieved through weighted averaging of overlapping windows. In contrast, the DMD-based methods utilize PyDMD implementations, incorporating data scaling, eigenvalue-based Koopman operators, and various decomposition strategies (e.g., sparsity in SpDMD, multi-scale analysis in MrDMD), followed by similar windowed reconstruction. Comparative analysis demonstrates that our PCA-DMD method achieves superior reconstruction accuracy, as evidenced by lower average KLD and HD across channels. Visualizations, including multi-channel signal comparison plots and heatmaps of amplitude across time and channels, further confirm that our approach more effectively captures the underlying dynamics of LFP signals compared to these established DMD techniques.

A Koopman Operator

The Koopman operator provides a linear perspective on nonlinear dynamical systems by analyzing the evolution of *observables* (functions of the state) rather than the states themselves. Consider a discrete-time system

$$_{t+1} = f(t), \quad t \in \mathcal{H} \subseteq \mathbb{R}^{d_h}. \quad (1)$$

For an observable $\phi : \mathcal{H} \rightarrow \mathbb{C}$, the Koopman operator \mathcal{K} is defined as

$$[\mathcal{K}\phi](\cdot) = \phi(f(\cdot)).$$

Although f may be nonlinear, \mathcal{K} is always linear (but typically infinite-dimensional), allowing spectral methods to be applied to nonlinear dynamics. Koopman eigenfunctions φ_k and eigenvalues λ_k satisfy

$$\varphi_k(t+1) = \lambda_k \varphi_k(t),$$

and under suitable assumptions the state can be expanded in terms of these eigenfunctions. This leads to the *Koopman mode decomposition (KMD)*:

$$_t = \sum_k \lambda_k^t \phi_{\lambda_k}(0) \varphi_k, \quad (2)$$

where φ_k are the Koopman modes associated with the observable Φ . Thus, KMD expresses nonlinear dynamics as a superposition of modes evolving linearly in time, forming the theoretical foundation for data-driven methods such as dynamic mode decomposition (DMD).

B SOTA Methods and their Results on LFP Reconstruction

B.1 Classical DMD

The Classical DMD implementation mirrors the windowing approach, utilizing PyDMD's DMD with an SVD rank of 8 on the transposed windows. Reconstruction involves extracting the real part of the DMD-reconstructed data for predictions, deriving a diagonal Koopman matrix from eigenvalues, and reconstructing the full signal through overlapping window averaging tailored to the signal length. The Classical DMD method for LFP signal reconstruction begins by windowing the signal $x(t) \in \mathbb{R}^N$ into overlapping snapshots $\mathbf{X}_{\text{full}} = \{x(t_i : t_i + w)\}_{i=1}^M$, with window size $w = 3000$ and step $\delta = 30$. The transposed snapshot matrix $\mathbf{X}_{\text{full}}^T \in \mathbb{R}^{w \times M}$ is decomposed using DMD with SVD rank $r = 8$, solving:

$$\mathbf{X}_{\text{next}}^T = \mathbf{A} \mathbf{X}^T,$$

where $\mathbf{X} = \mathbf{X}_{\text{full}}[:, -1]$, $\mathbf{X}_{\text{next}} = \mathbf{X}_{\text{full}}[1 :, :]$, and \mathbf{A} is approximated via DMD modes Φ and eigenvalues Λ :

$$\mathbf{X}_{\text{full}}^T \approx \Phi \Lambda^t \mathbf{b},$$

with \mathbf{b} as the initial amplitude. The predicted snapshots are $\mathbf{X}_{\text{pred}} = \text{Re}(\Phi \Lambda \mathbf{b})$. The full reconstructed signal is obtained as:

$$x_{\text{full}}(t) = \frac{1}{c(t)} \sum_{i=1}^{M-1} \mathbf{X}_{\text{pred},i}(t - i\delta), \quad t \in [i\delta, i\delta + w],$$

where $c(t)$ counts overlapping windows, and the signal is trimmed to $t \in [\delta, \min(N, T - \delta)]$.

B.2 SpDMD

For SpDMD, the LFP reconstruction starts with windowing the signal similarly into overlapping segments, then applying PyDMD's SpDMD with an SVD rank of 8 and a sparsity parameter rho of 1e-6 directly on the transposed window matrix. The reconstructed data from SpDMD is used to predict next snapshots, taking the real part, and a diagonal Koopman matrix is formed from the eigenvalues. Full signal reconstruction employs a custom averaging function over overlapping windows, adjusted to match the expected signal length.

SpDMD method for LFP signal reconstruction processes a single-channel signal $x(t) \in \mathbb{R}^N$ by forming overlapping windows $\mathbf{X}_{\text{full}} = \{x(t_i : t_i + w)\}_{i=1}^M$, with window size $w = 3000$ and step $\delta = 30$. The transposed snapshot matrix $\mathbf{X}_{\text{full}}^T \in \mathbb{R}^{w \times M}$ is decomposed using SpDMD with SVD rank $r = 8$ and sparsity parameter $\rho = 10^{-6}$, computing modes and eigenvalues via:

$$\mathbf{X}_{\text{full}}^T \approx \Phi \Lambda \Phi^{-1},$$

where Φ contains SpDMD modes, and Λ is a diagonal matrix of eigenvalues. The predicted snapshots are obtained as $\mathbf{X}_{\text{pred}} = \text{Re}(\Phi \Lambda \mathbf{b})$, with \mathbf{b} derived from the initial snapshot projection. Using $\mathbf{X} = \mathbf{X}_{\text{full}}[:, -1]$ and $\mathbf{X}_{\text{next}} = \mathbf{X}_{\text{full}}[1 :, :]$, the full reconstructed signal is computed as:

$$x_{\text{full}}(t) = \frac{1}{c(t)} \sum_{i=1}^{M-1} \mathbf{X}_{\text{pred},i}(t - i\delta), \quad t \in [i\delta, i\delta + w],$$

where $c(t)$ is the count of overlapping windows at time t , and the signal is trimmed to $t \in [\delta, \min(N, T - \delta)]$.

B.3 HODMD

HODMD for LFP reconstruction incorporates data scaling with StandardScaler before fitting PyDMD's HODMD with SVD rank 8 and delay d=2 on scaled transposed windows. The real part of reconstructed data is inverse-scaled, and predictions are made on shifted windows. Signal rebuilding uses the same overlapping averaging method. HODMD method for LFP reconstruction begins by windowing the signal $x(t) \in \mathbb{R}^N$ into overlapping snapshots $\mathbf{X}_{\text{full}} = \{x(t_i : t_i + w)\}_{i=1}^M$, with window size $w = 3000$ and step $\delta = 30$, scaled to $\tilde{\mathbf{X}}_{\text{full}} = \text{StandardScaler}(\mathbf{X}_{\text{full}})$. HODMD, with SVD rank $r = 8$ and delay $d = 2$, is applied to the transposed snapshot matrix $\tilde{\mathbf{X}}_{\text{full}}^T \in \mathbb{R}^{w \times M}$, constructing an augmented Hankel matrix and decomposing it as:

$$\tilde{\mathbf{X}}_{\text{full}}^T \approx \Phi \Lambda^t \mathbf{b},$$

where Φ and Λ are the DMD modes and eigenvalues, and \mathbf{b} is the amplitude vector. The reconstructed snapshots $\tilde{\mathbf{X}}_{\text{rec}}$ are inverse-scaled to \mathbf{X}_{rec} . Predictions use $\mathbf{X} = \mathbf{X}_{\text{full}}[:, -1]$, with $\mathbf{X}_{\text{pred}} = \text{Re}(\text{HODMD}(\text{StandardScaler}(\mathbf{X}))^T)$. The full signal is reconstructed via:

$$x_{\text{full}}(t) = \frac{1}{c(t)} \sum_{i=1}^{M-1} \mathbf{X}_{\text{pred},i}(t - i\delta), \quad t \in [i\delta, i\delta + w],$$

where $c(t)$ counts overlapping windows, and the signal is trimmed to $t \in [\delta, \min(N, T - \delta)]$.

B.4 MrDMD

MrDMD implementation scales the windowed data and employs PyDMD’s MrDMD with a base DMD of SVD rank 8, max level=2, and max cycles=1, fitting on scaled transposed matrices for multi-scale decomposition. Reconstructions and predictions use the real part inverse-scaled, with full signal assembly via overlapping averages. MrDMD method processes the LFP signal $x(t) \in \mathbb{R}^N$ by forming overlapping windows $\mathbf{X}_{\text{full}} = \{x(t_i : t_i + w)\}_{i=1}^M$, with window size $w = 3000$ and step $\delta = 30$, scaled to $\tilde{\mathbf{X}}_{\text{full}} = \text{StandardScaler}(\mathbf{X}_{\text{full}})$. The MrDMD model, with a base DMD of SVD rank $r = 8$, maximum decomposition level $L = 2$, and maximum cycles $C = 1$, decomposes the transposed snapshot matrix $\tilde{\mathbf{X}}_{\text{full}}^T \in \mathbb{R}^{w \times M}$ into multi-scale modes:

$$\tilde{\mathbf{X}}_{\text{full}}^T \approx \sum_{l=1}^L \sum_{c=1}^C \Phi_{l,c} \Lambda_{l,c}^t \mathbf{b}_{l,c},$$

where $\Phi_{l,c}$, $\Lambda_{l,c}$, and $\mathbf{b}_{l,c}$ are the modes, eigenvalues, and amplitudes at level l and cycle c . The reconstructed snapshots $\tilde{\mathbf{X}}_{\text{rec}}$ are inverse-scaled to \mathbf{X}_{rec} , and predictions use $\mathbf{X} = \mathbf{X}_{\text{full}}[:, -1]$ and $\mathbf{X}_{\text{pred}} = \text{Re}(\text{MrDMD}(\text{StandardScaler}(\mathbf{X}))^T)$. The full signal is reconstructed as:

$$x_{\text{full}}(t) = \frac{1}{c(t)} \sum_{i=1}^{M-1} \mathbf{X}_{\text{pred},i}(t - i\delta), \quad t \in [i\delta, i\delta + w],$$

with $c(t)$ as the overlap count, trimmed to $t \in [\delta, \min(N, T - \delta)]$.

C (Hyper)parameters

Table 3: Short Forms of Parameters

Parameter	Short Form
Sampling Frequency	fs
Maximum Samples	max_samples
Window Size	window_size
Step Size	step
Latent Dimension	latent_dim
Sparsity Parameter	rho
Delay Embeddings	d
Maximum Decomposition Levels	max_level
Maximum Cycles	max_cycles

Table 4: Parameters Used in PCA-DMD, Classical DMD, SpDMD, HODMD, and MrDMD

Parameter	PCA-DMD	Classical DMD	SpDMD	HODMD	MrDMD
fs	30,000	30,000	30,000	30,000	30,000
max_samples	200,000	200,000	200,000	200,000	200,000
window_size	3,000	3,000	3,000	3,000	3,000
step	30	30	30	30	30
latent_dim	8	8	8	8	8
rho	-	-	1e-8	-	-
d	-	-	-	50	-
max_level	-	-	-	-	5
max_cycles	-	-	-	-	2

D Evaluation Measures

D.1 Kullback-Leibler divergence (KLD)

To evaluate the geometrical agreement between the true and reconstructed LFP signals in PCA-DMD and DMD-based methods (Classical DMD, SpDMD, MrDMD, and HODMD), we employed the KLD

as a state space divergence metric. Specifically, for each of the eight LFP channels, we estimated probability distributions $p(x)$ and $q(x)$ from the trajectories of the true and reconstructed signals, respectively, in the observation space. The KLD was computed using histograms with 100 bins over the trimmed signal segments to approximate the distributions, with a small regularization term ($\epsilon = 10^{-10}$) added to avoid numerical issues. The KLD, averaged across channels, quantifies the discrepancy between the true and reconstructed attractor geometries, with lower values indicating better fidelity. Mathematically, the state space divergence is defined as:

$$KLD(p(x) \parallel q(x)) = \int_{\mathbb{R}^N} p(x) \log \frac{p(x)}{q(x)} dx,$$

where $p(x)$ and $q(x)$ represent the probability densities functions of the true and reconstructed trajectories, respectively.

D.2 Hellinger Distance

To assess the temporal agreement between the ground truth and reconstructed LFP signals, we utilized the (*HD*) as a temporal measure, bounded between 0 and 1, averaged across all eight dynamical variables (channels). For each channel, we computed the power spectra $f_i(\omega)$ and $g_i(\omega)$ for the true and reconstructed signals, respectively, using histogram-based approximations with 100 bins over the trimmed signal segments, normalized to satisfy $\int_{-\infty}^{\infty} f_i(\omega) d\omega = 1$ and $\int_{-\infty}^{\infty} g_i(\omega) d\omega = 1$. A regularization term ($\epsilon = 10^{-10}$) was applied to ensure numerical stability. The Hellinger distance, ranging from 0 (perfect agreement) to 1, was calculated per channel and averaged to produce HD. The Hellinger distance for the i -th channel is defined as:

$$HD(f_i(\omega), g_i(\omega)) = \sqrt{1 - \int_{-\infty}^{\infty} \sqrt{f_i(\omega)g_i(\omega)} d\omega},$$

where $f_i(\omega)$ and $g_i(\omega)$ are the normalized power spectra of the true and reconstructed signals for the i -th channel.

E Dataset

To provide a broader view of the dataset beyond the analyzed segments, this section includes visualizations of the full-length hippocampal LFP recordings, both with and without annotated SWR events, as well as the combined eight-channel signals. These plots highlight the large-scale structure of the recordings and illustrate the spatial and temporal distribution of SWRs across channels, complementing the windowed analysis used for reconstruction.

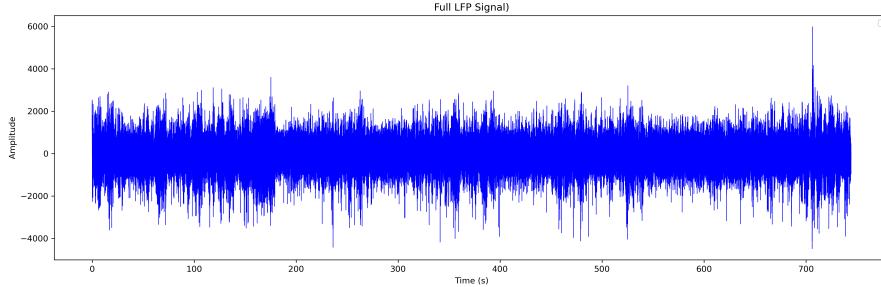


Figure 4: Full-length LFP signal.

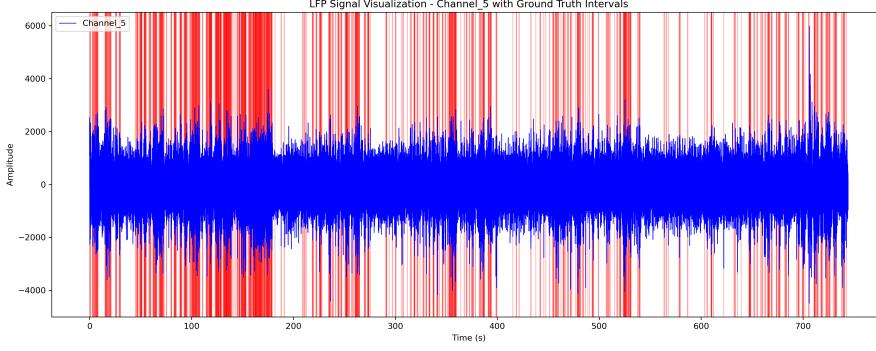


Figure 5: Full-length LFP signal with SWR.

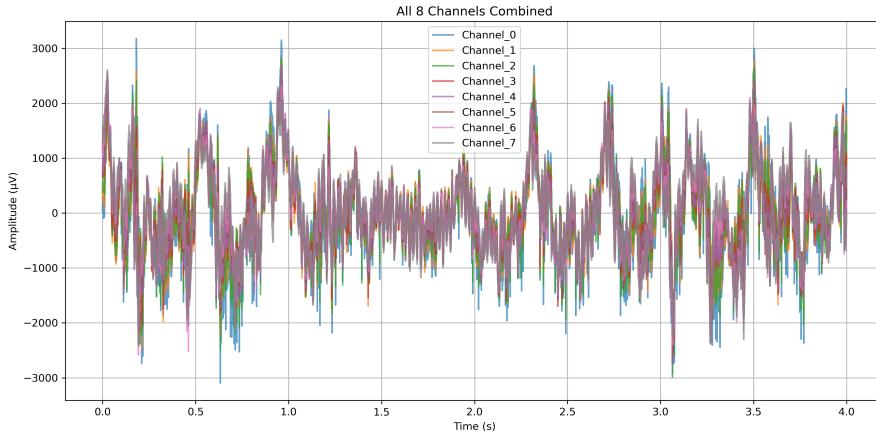


Figure 6: All Channels.

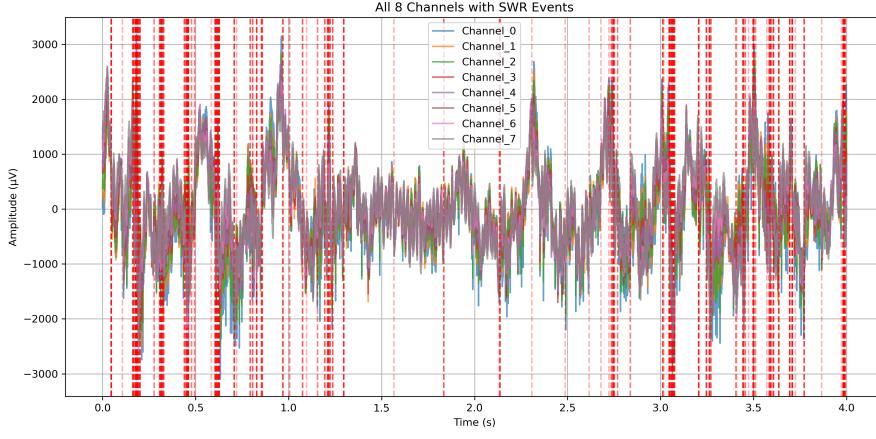


Figure 7: All Channels with SWR.

F Ablation Study and Results

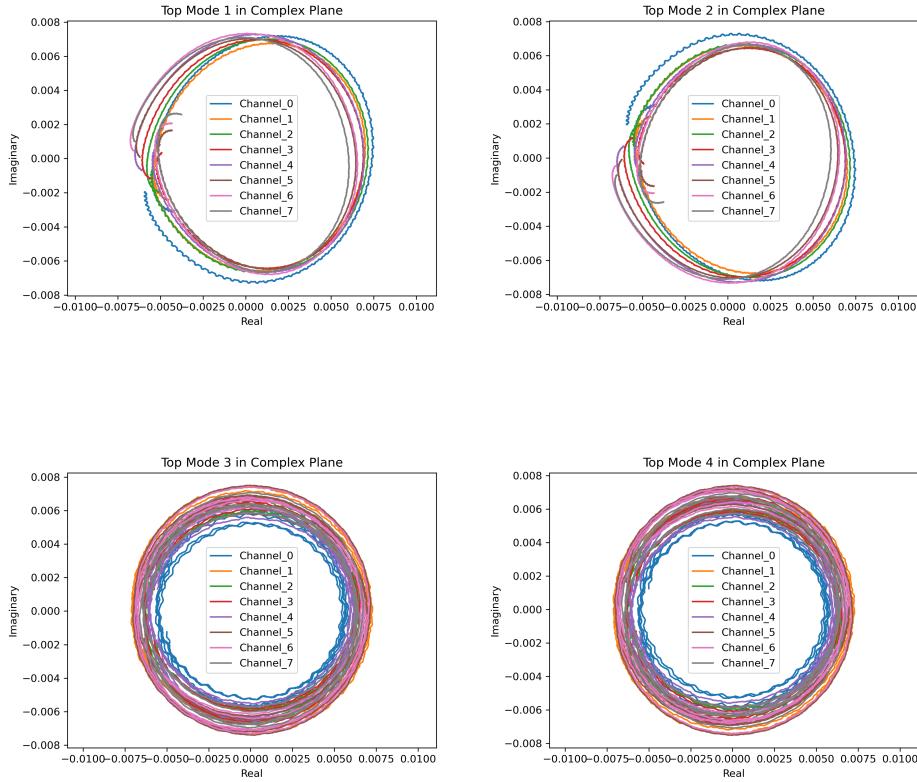
To evaluate the performance of the PCA-DMD, Classical DMD, SpDMD, HODMD, and MrDMD methods in analyzing the 8-dimensional system, an ablation study was conducted with consistent parameters. The study focuses on comparing the reconstructed signals, Koopman modes, and

associated metrics such as KL divergence and Hellinger distance. Visualizations, including original vs. reconstructed signal plots, heatmaps of original and reconstructed signals, Koopman spectrum, top 5 Koopman modes in both time and complex domains, Koopman operator heatmaps, and the first 5 principal components, are provided for each method to assess their effectiveness in capturing the dynamics of the 8D system. Detailed results, including mode structures and reconstruction accuracy, are presented in the following figures and tables. Further experiments reducing the SVD rank from 8 to 5 resulted in significant degradation of reconstruction quality across all methods, with increased KLD and HD reflecting loss of critical dynamic modes. When scaling the reconstruction to 800,000 samples, only PCA-DMD maintained high fidelity, as visualized in heatmaps that nearly replicated the original signal's spatiotemporal patterns, while other methods struggled with computational complexity. Some of the additional results for the implemented methods are presented in the following sections.

G Koopman modes

G.1 Classical DMD Additional Results

This section presents the top four *eigenmodes* in the complex plane. In the time domain, the plots show the temporal evolution of the top five modes per channel, with examples provided for the first two channels.



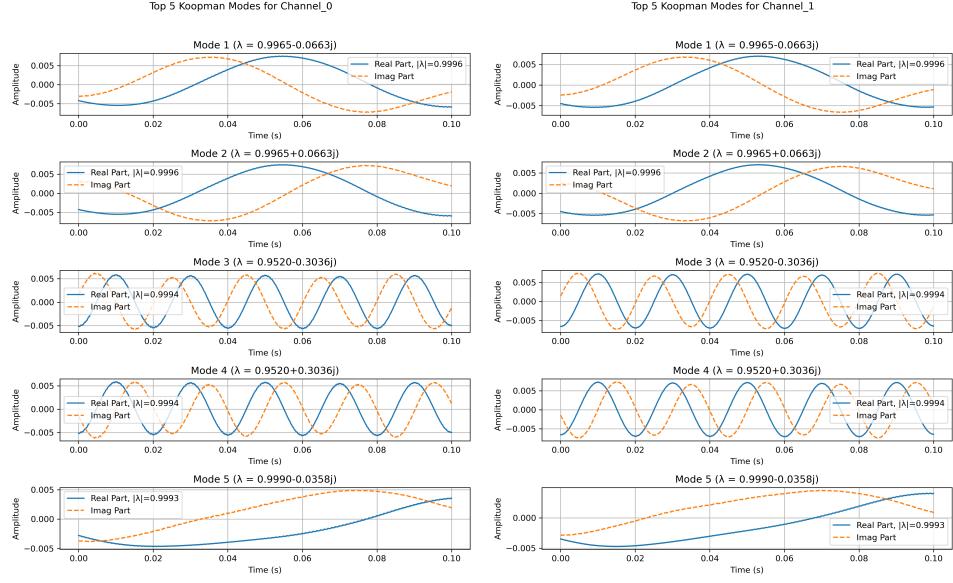
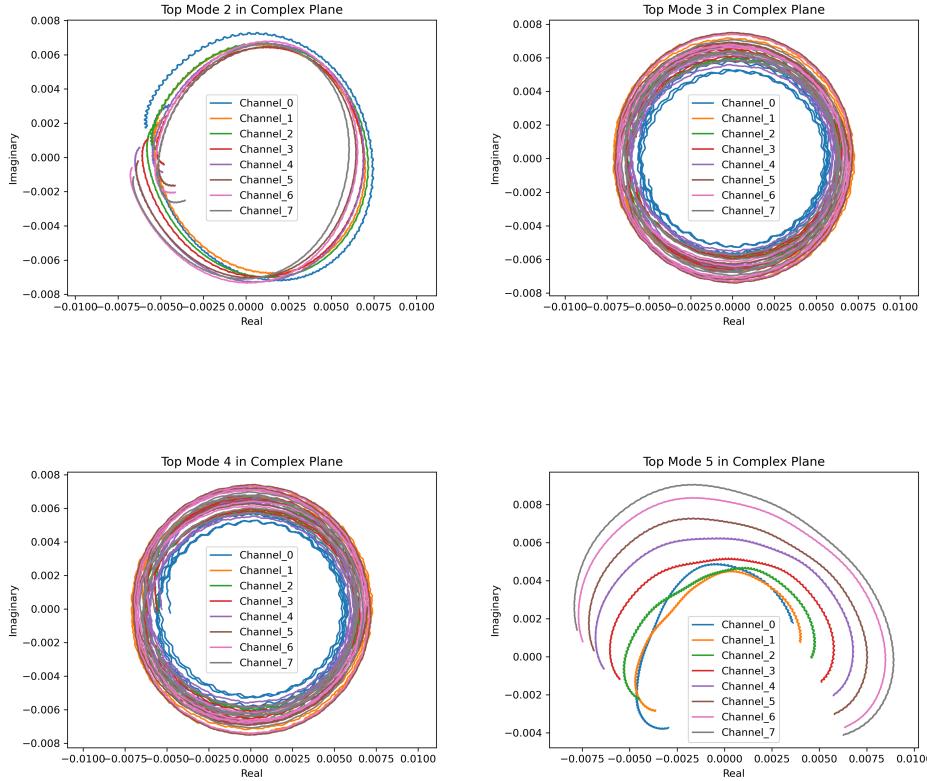


Figure 10: Four Modes in Complex Plane and Five Modes for first two channels in Time vs Amplitude (Classical DMD).

G.2 SpDMD Additional Results

The sparse formulation eliminates unnecessary modes, resulting in fewer; see Fig. 13.



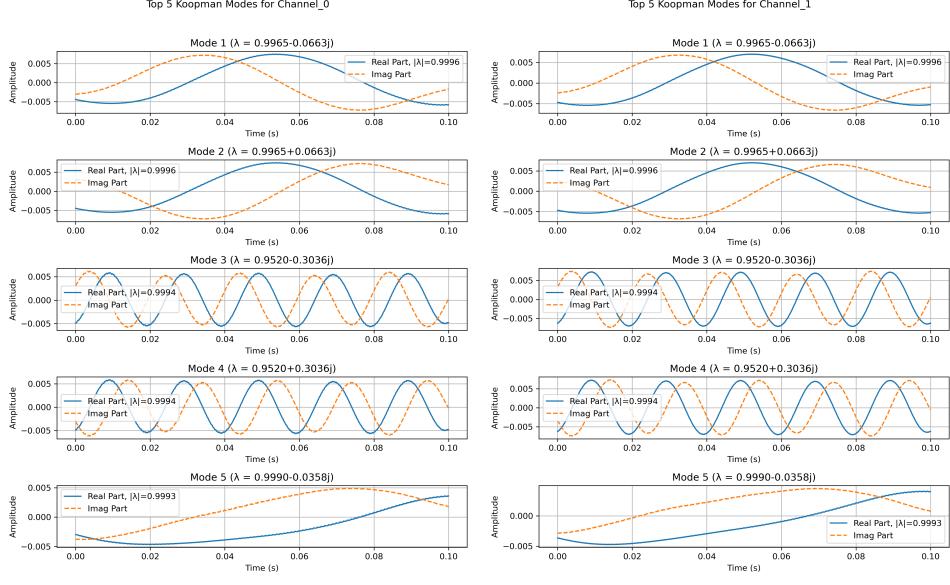
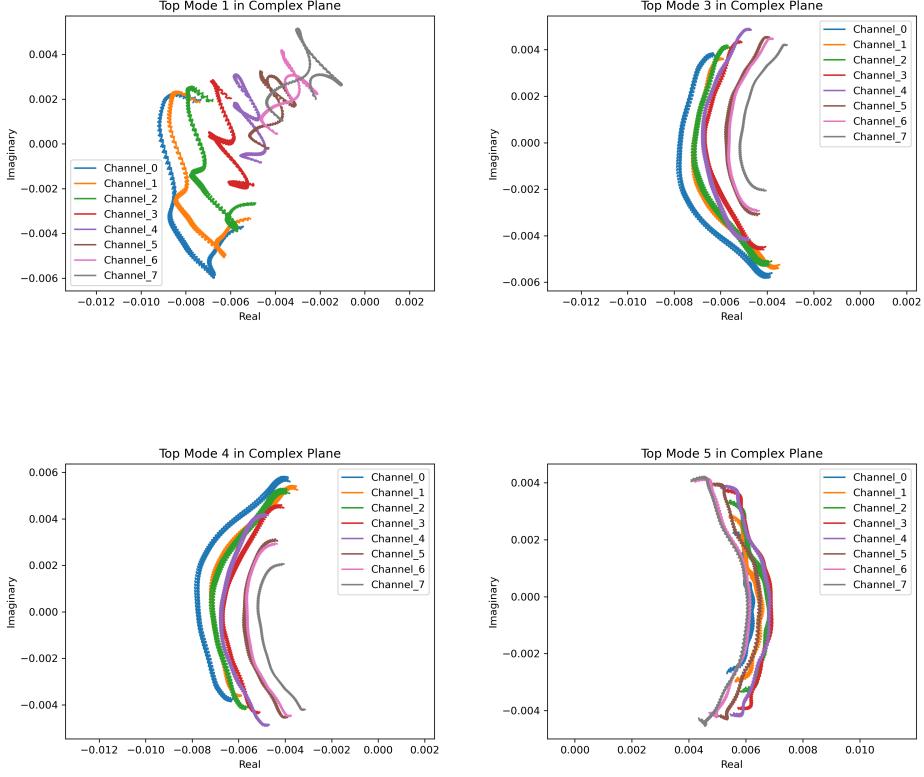


Figure 13: Four Modes in Complex Plane and Five Modes for first two channels in Time vs Amplitude (SpDMD).

G.3 MrDMD Additional Results

Fig. 16 illustrates the top four *eigenmodes* in the complex plane and the temporal evolution of the top five modes for the first two channels using the MrDMD method.



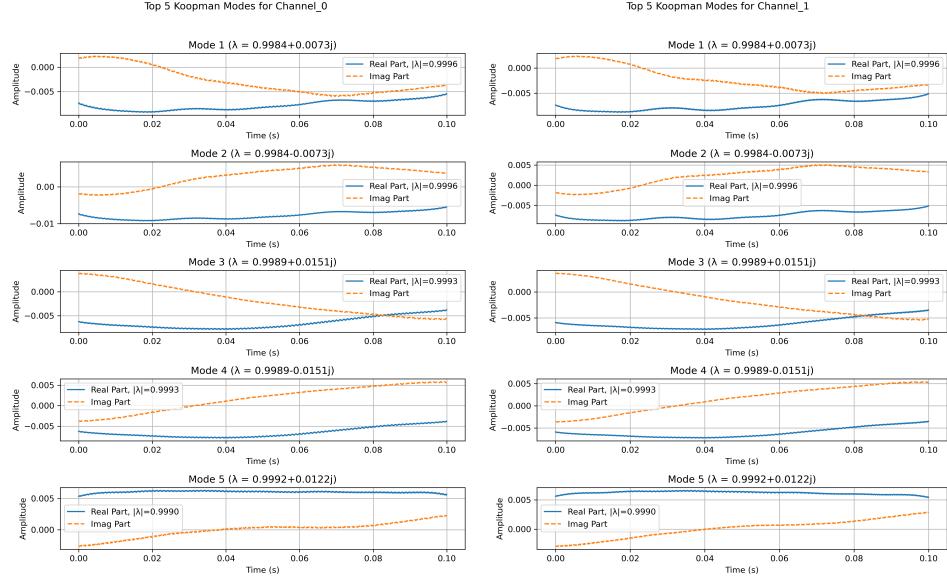
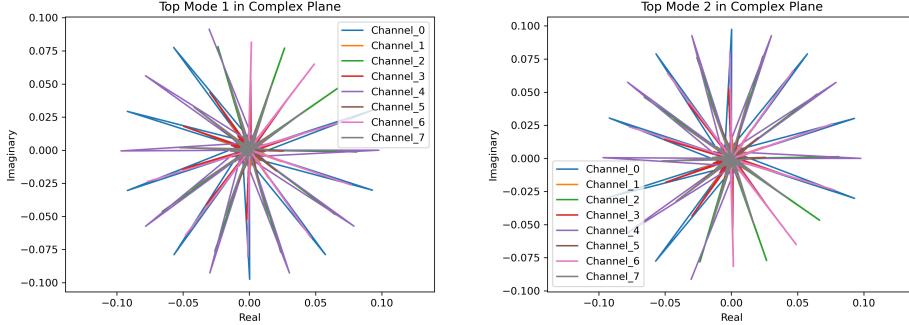


Figure 16: Four Modes in Complex Plane and Five Modes for first two channels in Time vs Amplitude (MrDMD).

G.4 HODMD Additional Results

The plot features a star-like pattern with multiple straight lines radiating from a central point, each line corresponding to one of the eight channels. HODMD extends the standard DMD by incorporating higher-order dynamics, often using techniques like delayed embedding to enhance the representation of complex systems where spectral complexity exceeds spatial complexity.



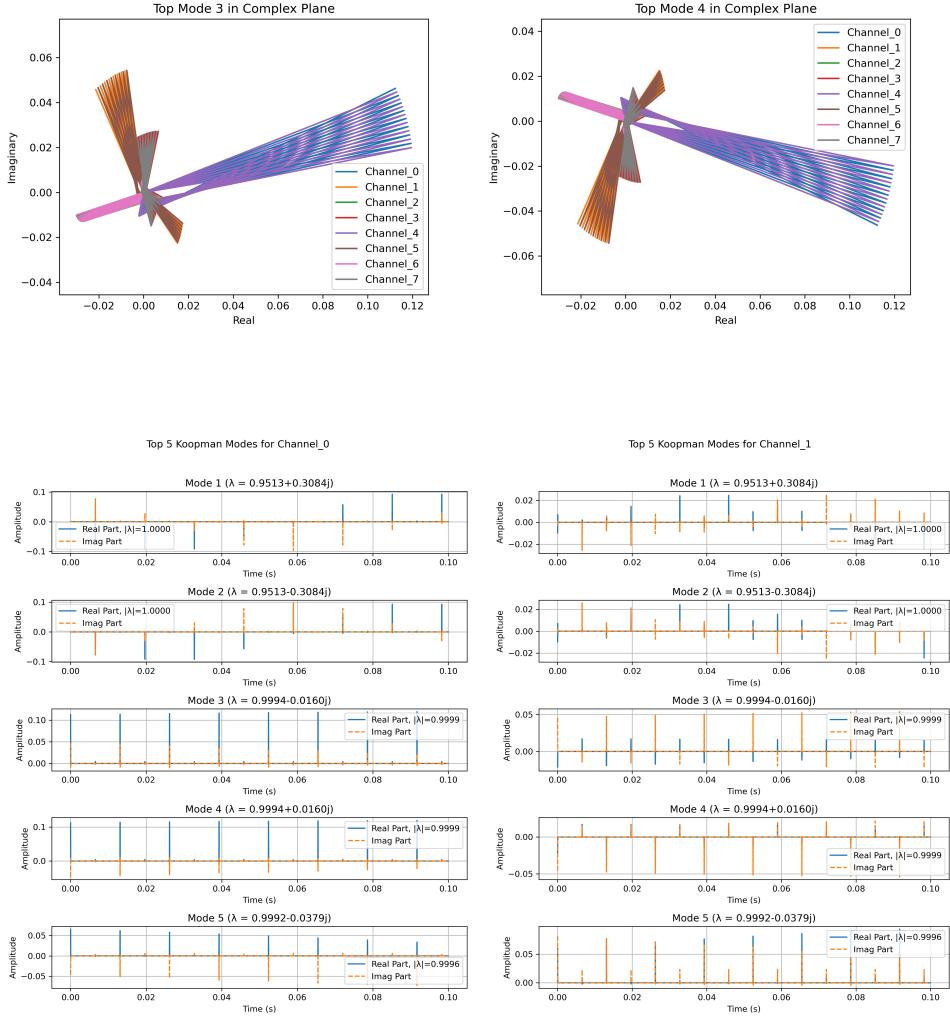


Figure 19: Four Modes in Complex Plane and Five Modes for first two channels in Time vs Amplitude (HODMD).

G.5 PCA-DMD Additional Results

Principal Component Heatmaps (Spatial Structure): Since PCA-DMD first projects the data into a reduced PCA subspace, the analysis is presented in terms of spatial principal components rather than eigenvalues in the complex plane. The top five heatmaps show the spatial distribution of the first five principal components across channels and time, highlighting the dominant variance patterns preserved by PCA.

Time Domain (Top 5 Modes): The temporal modes reconstructed in the reduced space show smoother dynamics, with each mode aligned to a specific principal component. Mode 1 typically captures the global trend, while higher modes encode localized oscillations in certain channels.

