
Learning High-Dimensional Transient Neural Dynamics for Zero-Shot Cross-Subject Reconstruction

Anonymous Authors¹

Abstract

Accurate and scalable reconstruction of high-dimensional neural time series remains a central challenge in dynamical systems modeling. In this work, we study long-horizon reconstruction of hippocampal local field potentials (LFPs) using a PCA-DMD framework, with Dynamic Mode Decomposition (DMD) as a data-driven approximation to Koopman spectral analysis. We extend prior analyses by scaling to sequences of up to 300,000 samples and introducing a cross-subject zero-shot evaluation protocol to assess the generalizability of learned latent representations across four subjects. PCA-DMD consistently achieves strong reconstruction quality under both long-horizon and zero-shot transfer settings, with correlation exceeding 0.95 and very low reconstruction errors. In contrast, standard DMD variants degrade substantially as signal length increases and show limited cross-subject robustness. These results demonstrate that PCA-DMD provides a scalable and transferable framework for reconstructing high-dimensional neural signals, highlighting its potential for modeling long-range neural dynamics in challenging multi-subject settings.

1. Introduction

High-dimensional time series arise in many scientific domains, including neuroscience, biology, and climate science, where the goal is often to learn compact representations that preserve the essential structure of the underlying dynamics. In neural systems, this challenge is particularly pronounced due to the large scale, high temporal resolution, and nonlinear structure of recorded signals. Accurate long-horizon reconstruction is therefore an important benchmark for evaluating whether a model captures meaningful dynamical structure over extended time intervals. LFPs provide a rich example of such data. These signals reflect coordinated population-level neural activity and exhibit complex temporal organization across multiple channels and timescales (Buzsáki et al., 2012; Buzsáki, 2015). Their high sampling rate and long recording duration make them a challeng-

ing testbed for scalable dynamical modeling. In practice, many existing signal models struggle with this regime due to memory limitations, instability over long horizons, and limited interpretability. This motivates the search for methods that are both computationally efficient and dynamically interpretable. Koopman operator theory offers a principled framework for analyzing nonlinear dynamical systems through linear evolution in an appropriate function space of observables (Koopman, 1931). DMD provides a practical data-driven approximation of this idea and has been widely used to extract coherent spatiotemporal patterns from complex systems (Schmid, 2010; Rowley et al., 2009; Kutz et al., 2016a). However, direct application of DMD to large neural recordings remains challenging due to the high dimensionality of windowed signals and the difficulty of maintaining stable reconstruction over long trajectories. To address these limitations, recent work introduced PCA-DMD as a scalable reconstruction framework that combines dimensionality reduction with Koopman-based latent evolution (Kujur et al., 2025). By projecting high-dimensional windows into a compact PCA subspace before fitting a linear operator (that approximates Koopman evolution), PCA-DMD substantially improves computational tractability and reconstruction stability. This approach supports long-horizon reconstruction up to 200k samples and outperforms several standard DMD variants in signal fidelity. In this work, we extend that framework in two important directions. First, we evaluate reconstruction at an even longer horizon, up to 300k samples, to test the scalability of the method. Second, we introduce a cross-subject zero-shot protocol across four subjects to assess whether the learned latent dynamics generalize beyond a single recording. This setting is particularly important for neural data, where robust transfer across subjects is both difficult and practically valuable. Beyond reconstruction metrics, we further analyze the learned Koopman representations through mode-level visualization across subjects and DMD variants. This analysis reveals that classical DMD-based methods tend to produce highly redundant and spatially smooth mode structures, indicating a concentration of energy in a limited set of dominant dynamics. In contrast, PCA-DMD yields a richer and more diverse set of Koopman modes, exhibiting both structured global patterns and higher-frequency variations across chan-

nels and time. This suggests that PCA-based dimensionality reduction not only improves computational scalability but also enhances the expressiveness and conditioning of the learned dynamical representation. Our results show that PCA-DMD consistently maintains high reconstruction quality in both long-horizon and zero-shot settings, achieving correlation above 0.95 together with very low Hellinger distance (HD) and Kullback-Leibler divergence (KLD) across source-target subject pairs. In contrast, classical DMD variants show clear degradation as signal length increases and fail to generalize reliably across subjects. These findings establish PCA-DMD as a scalable and transferable approach for reconstructing high-dimensional neural time series.

The contributions of this work are summarized as follows:

- We extend PCA-DMD to very long-horizon reconstruction of hippocampal LFP signals up to 300k samples.
- We introduce a cross-subject zero-shot evaluation framework for assessing generalization of learned dynamical representations across four subjects.
- We show that PCA-DMD consistently outperforms standard DMD variants in both reconstruction fidelity and transfer robustness.
- We provide Koopman spectral and mode-level analysis demonstrating improved structural richness and representation quality in PCA-DMD.

2. Related Work

Neural signals and hippocampal LFPs. LFPs provide a mesoscopic view of coordinated neural population activity and are widely used to study circuit-level dynamics (Buzsáki et al., 2012; Buzsáki, 2015). In the hippocampus, LFPs exhibit rich multi-scale temporal structure and spatial organization (Joo & Frank, 2018; Wilson & McNaughton, 1994; Diba & Buzsáki, 2007; Foster, 2017), making them both scientifically valuable and computationally challenging due to their high dimensionality and long duration.

Neural time-series reconstruction. Neural time-series modeling has been explored using latent variable models and recurrent architectures (Cunningham & Yu, 2014; Pandarinath et al., 2018; Durstewitz et al., 2023). While these approaches can capture complex nonlinear dependencies, they struggle with stability on very long horizons (e.g., LFPs) and memory, while offering limited interpretability (Brenner et al., 2022; 2024; Hess et al., 2023) unlike Koopman methods. Practical solutions such as truncation or downsampling can reduce computational cost but may compromise temporal consistency.

Koopman operator theory and DMD. Koopman operator theory enables linear representations of nonlinear dynamics in an infinite-dimensional lifted space (Koopman, 1931). DMD provides a data-driven approximation (in finite dimensions) and has been widely applied for extracting coherent spatiotemporal structures (Schmid, 2010; Rowley et al., 2009; Kutz et al., 2016a). Its interpretability through modes and spectra is a key advantage (Brunton & Kutz, 2022; Schmid, 2022). Extensions such as sparsity-promoting DMD (Jovanović et al., 2014), higher-order DMD (Le Clainche & Vega, 2017), multiresolution DMD (Kutz et al., 2016b), and delay embeddings (HAVOK) (Brunton et al., 2017) improve robustness, but scaling to long, high-dimensional neural recordings remains difficult.

Dimensionality reduction and scalable reconstruction in neural dynamics. Dimensionality reduction, particularly PCA, is widely used to compress neural recordings while preserving dominant variance (Cunningham & Yu, 2014). When combined with dynamical modeling, it reduces computational complexity and improves numerical stability. This makes it a natural complement to Koopman-based approaches for large-scale time-series analysis. (Kujur et al., 2025) introduced PCA-DMD as a scalable Koopman-based framework for long-horizon LFP reconstruction, demonstrating strong performance up to 200k samples. The present work extends this approach to longer horizons (300k samples) and introduces cross-subject zero-shot evaluation, enabling assessment of both scalability and transferability.

Position of this work. This work lies at the intersection of neural time-series reconstruction, Koopman-based modeling, and generalization across subjects. Unlike prior approaches focused on model complexity or event detection, we emphasize scalable reconstruction and the structure of learned dynamical representations. Our analysis shows that PCA-DMD produces more expressive and less redundant Koopman modes compared to standard DMD variants, while maintaining strong reconstruction and transfer performance.

3. Methodology

Fig. 1 illustrates the overall pipeline of the proposed framework. Multichannel hippocampal LFP recordings are treated as long-horizon time series with multiple channels and up to 300k samples. The signal is segmented into overlapping temporal windows, which are vectorized into high-dimensional snapshot representations. These snapshots are projected into a compact PCA latent space, where a Koopman operator is learned to model linear evolution of the latent states. The predicted latent trajectories are then mapped back to the original signal space and aggregated via overlap-add averaging to produce a continuous reconstructed signal. Finally, the model is evaluated under a cross-subject zero-shot pro-

Algorithm 1 PCA-DMD for Long-Horizon Reconstruction

Input: signal $s(t)$, window size w , step δ , latent dimension r
 Extract overlapping windows $W_i = s[i\delta : i\delta + w]$
 Vectorize windows: $x_i = \text{vec}(W_i)$
 Form consecutive window pairs (X, Y)
 Fit PCA and compute latent coordinates $z_i = P^\top x_i$
 Estimate Koopman operator: $K^\top = Z^\dagger Z_+$
for each window x_i **do**
 Predict latent state: $\hat{z}_i = K z_i$
 Reconstruct window: $\hat{x}_i = P \hat{z}_i + \mu$
end for
 Apply overlap-add averaging to obtain $\hat{s}(t)$
Output: reconstructed signal $\hat{s}(t)$

tocol: parameters trained on one subject are applied directly to unseen subjects without fine-tuning.

3.1. Problem Setup and Data

We consider the problem of very long-horizon reconstruction of multichannel hippocampal LFPs recorded from rodents (Meier et al., 2020; Navas-Olive et al., 2024). The dataset consists of four subjects, denoted by $q \in \{S1, S2, S3, S4\}$, with each recording containing $C = 8$ channels sampled at $f_s = 30,000$ Hz. Each channel contains approximately 22 million samples (~ 744 seconds), providing a challenging setting for scalable dynamical modeling. For computational tractability, we retain up to $T = 300,000$ samples per subject in the long-horizon experiments. Let $s^{(q)}(t) \in \mathbb{R}^C$ denote the multichannel signal from subject q at time index t . We segment each recording into overlapping windows of length $w = 3000$ and stride $\delta = 30$. The i -th window is denoted by

$$W_i = s^{(q)}[i\delta : i\delta + w - 1] \in \mathbb{R}^{w \times C},$$

which is vectorized as $x_i = \text{vec}(W_i) \in \mathbb{R}^{d_x}$, where $d_x = wC$. Consecutive windows define paired samples

$$X = \begin{bmatrix} x_1^\top \\ \vdots \\ x_{N-1}^\top \end{bmatrix}, \quad Y = \begin{bmatrix} x_2^\top \\ \vdots \\ x_N^\top \end{bmatrix},$$

where N is the number of windows. This yields a one-step window-to-window prediction over long temporal contexts.

3.2. PCA-DMD Reconstruction Framework

Our reconstruction model, PCA-DMD, combines dimensionality reduction with Koopman-inspired linear evolution in a compact latent space. Each high-dimensional window x_i is projected onto a PCA basis $P \in \mathbb{R}^{d_x \times r}$ (with

Algorithm 2 Cross-Subject Zero-Shot Evaluation

Input: subjects $\{S_1, \dots, S_n\}$, reconstruction model f_θ
for each source subject S_a **do**
 Train model parameters θ_a using only S_a
for each target subject $S_b, b \neq a$ **do**
 Reconstruct target signal: $\hat{s}^{(b)} = f_{\theta_a}(s^{(b)})$
 Compute reconstruction metrics between $s^{(b)}$ and $\hat{s}^{(b)}$
end for
end for
Output: cross-subject reconstruction results

$r = 100$), yielding latent coordinates $z_i = P^\top x_i$ and $\tilde{z}_i = P^\top y_i$. Latent dynamics are modeled via a linear operator $K \in \mathbb{R}^{r \times r}$ such that $\tilde{z}_i \approx K z_i$, which is estimated by least squares as $K = \arg \min_A \|Z_+ - Z A^\top\|_F^2$, where $Z = [z_1, \dots, z_{N-1}]^\top$ and $Z_+ = [\tilde{z}_1, \dots, \tilde{z}_{N-1}]^\top$, with closed-form solution $K^\top = Z^\dagger Z_+$. Prediction is performed in latent space as $\hat{z}_i = K z_i$ and $\hat{x}_i = P \hat{z}_i + \mu$, after which \hat{x}_i is reshaped into window form $\hat{W}_i \in \mathbb{R}^{w \times C}$ and combined via overlap-add averaging $\hat{s}(t) = \frac{1}{n_t} \sum_{i: t \in [i\delta, i\delta + w]} \hat{W}_i(t - i\delta)$ to produce a continuous full-length reconstruction. This formulation enables stable long-horizon prediction while significantly reducing the effective state dimension and improving numerical conditioning compared to DMD variants.

3.3. Cross-Subject Zero-Shot Evaluation

We evaluate generalization using a cross-subject zero-shot protocol. For each ordered pair (a, b) with $a \neq b$, the model is trained on subject a and directly applied to subject b , i.e., $\hat{s}^{(b)} = f_{\theta_a}(s^{(b)})$. This setting tests whether the learned latent dynamics capture subject-invariant structure. Building on prior results showing strong performance up to 200k samples, we extend evaluation to 300k samples and systematically assess transfer across all subject pairs. This allows us to jointly study scalability and generalization under a unified framework. Reconstruction quality is evaluated using global similarity measures between the original and reconstructed signals. Let $x = \text{vec}(s)$ and $\hat{x} = \text{vec}(\hat{s})$. We compute Pearson correlation $\text{Corr}(x, \hat{x}) = \frac{\text{Cov}(x, \hat{x})}{\sigma_x \sigma_{\hat{x}}}$, together with HD(p, \hat{p}) = $\frac{1}{\sqrt{2}} \|\sqrt{p} - \sqrt{\hat{p}}\|_2$ and $\text{KLD}(p \parallel \hat{p}) = \sum_j p_j \log \frac{p_j}{\hat{p}_j}$, where p and \hat{p} are histogram estimates of signal amplitudes. These complementary metrics capture both temporal agreement and distributional similarity, enabling consistent comparison across subjects and transfer settings. To further characterize the learned dynamics, we analyze the Koopman spectrum of the latent operator. Let $K v_k = \lambda_k v_k$ denote its eigenpairs, yielding the expansion $z_t \approx \sum_{k=1}^r c_k \lambda_k^t v_k$, where dominant behavior is governed by eigenvalues of largest magnitude. In practice, eigenvalues cluster near the unit circle, indicat-

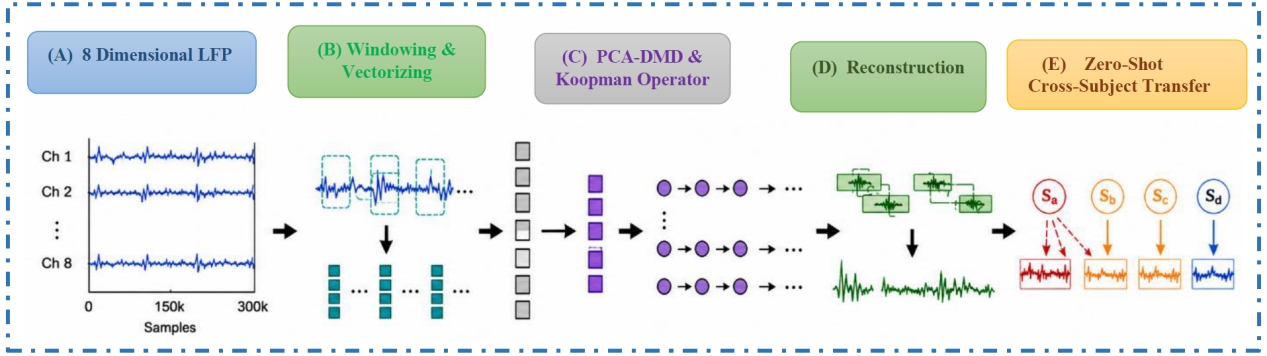


Figure 1. **Method overview.** (A) Multichannel hippocampal LFP recordings are first represented as long-horizon neural time series with multiple channels and up to 300k samples. (B) The signal is segmented into overlapping temporal windows and vectorized into high-dimensional snapshot representations. (C) These snapshots are projected into a compact PCA latent space, where a Koopman operator is learned to propagate latent states over time. (D) The predicted latent dynamics are mapped back to signal space and combined through overlap-add averaging to obtain a continuous reconstructed LFP signal. Finally, (E) the learned model is evaluated in a cross-subject zero-shot setting, where parameters learned from one source subject are directly applied to unseen target subjects without fine-tuning.

ing stable and persistent latent evolution. We additionally examine the corresponding Koopman modes, which capture the spatial-temporal structure of the learned dynamics. The low-rank structure induced by PCA further supports scalability. Given $X = U\Sigma V^T$, the rank- r approximation $X_r = U_r\Sigma_r V_r^T$ retains dominant variance while substantially reducing dimensionality. This compression enables efficient Koopman regression and stable long-horizon reconstruction for large multichannel neural recordings.

4. Results

We evaluate reconstruction performance under long-horizon and cross-subject zero-shot settings using hippocampal LFP recordings from four subjects. Across all experiments, PCA-DMD consistently achieves stronger reconstruction fidelity and transfer robustness than standard DMD variants.

Long-Horizon Cross-Subject Reconstruction. Fig. 3 presents representative zero-shot reconstructions produced by PCA-DMD across subjects. The reconstructed signals closely track the original waveforms over long horizons, preserving global temporal structure and amplitude variation. In contrast, baseline DMD variants exhibit clear degradation, including amplitude attenuation, waveform distortion, and reduced temporal alignment, which becomes more pronounced as the reconstruction length increases. Quantitative results are summarized in Table 1. PCA-DMD consistently achieves high correlation (0.9504-0.9800) together with very low HD (0.0010-0.0072) and KLD (0.0005-0.0022) across all source-target pairs. By comparison, DMD, HODMD, and HankelDMD show near-failure regimes with low correlation and large distributional mismatch, while SpDMD provides only limited improvement. These results demonstrate that PCA-DMD uniquely maintains stable long-horizon reconstruction while general-




Figure 2. PCA-DMD Koopman mode structure for subject S1: real and imaginary components of representative dominant modes across window time and channels. See Appendix B for all baseline DMD methods.

izing reliably under cross-subject transfer.

Comparison Across DMD Variants. All methods are evaluated under identical preprocessing and windowing settings (Appendix Table 3), with a fixed window size and step. PCA-DMD additionally introduces a low-dimensional latent representation ($r = 100$), which substantially reduces the effective state dimension of the windowed multichannel signal and improves numerical conditioning of the Koopman regression. Quantitative comparisons are summarized in Appendix Fig. 5 and Table 1. PCA-DMD consistently achieves the highest correlation and the lowest HD and KLD across all source-target pairs. In contrast, classical DMD variants (DMD, HODMD, HankelDMD) exhibit severe degradation in the zero-shot setting, with near-random correlation and large distributional mismatch, while SpDMD shows only marginal improvement. These results indicate that standard DMD algorithms struggle to maintain stable long-horizon reconstructions under distribution shift, whereas PCA-DMD provides both accuracy and robustness.

Koopman Spectral Analysis. We analyze the learned dynamics via the eigenvalue spectrum of the PCA-DMD Koop-

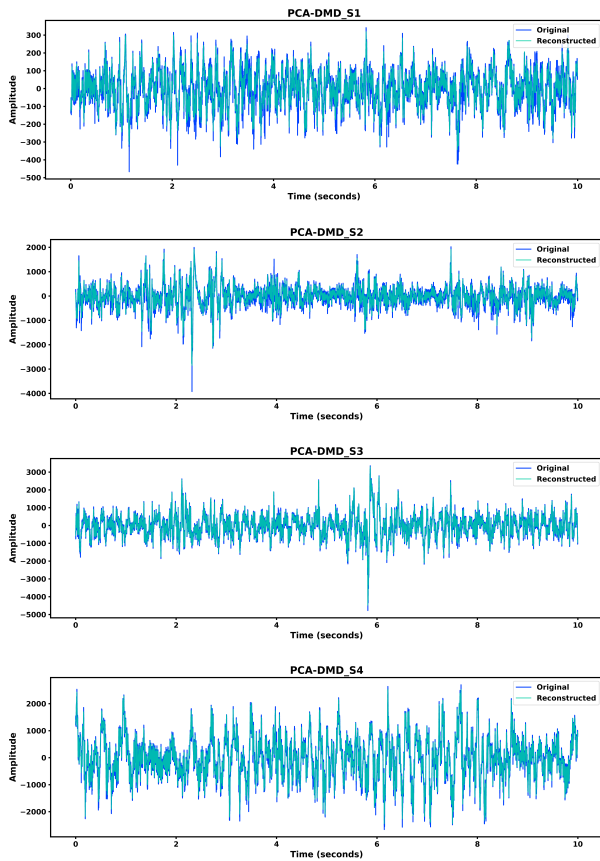


Figure 3. Zero-shot PCA-DMD reconstructions across four subjects closely follow original LFPs, preserving global temporal structure and amplitude over long durations.

man operator (Fig. 4). Dominant eigenvalues cluster near the unit circle across all zero-shot subjects, indicating stable, persistent latent evolution. Fig. 2 shows the Koopman mode structure. The spectral structure is consistent across subjects, indicating the learned operator captures shared dynamical organization rather than overfitting to subject-specific signal characteristics. This cross-subject consistency aligns with the strong transfer performance in Table 1, confirming that PCA-DMD learns a compact, well-conditioned latent representation of the underlying dynamics.

5. Conclusion

We presented a systematic study of long-horizon neural signal reconstruction using PCA-DMD under a cross-subject zero-shot setting. Extending prior work, we demonstrated that PCA-DMD scales to signals up to 300k samples while maintaining strong reconstruction fidelity across unseen subjects. Across all source-target pairs, PCA-DMD achieves high correlation and minimal distributional mismatch (low HD and KLD), significantly outperforming DMD, HODMD,

Table 1. Cross-subject zero-shot reconstruction performance. PCA-DMD consistently achieves high correlation and minimal distributional mismatch across all source - target pairs, substantially outperforming standard DMD variants. Full results are provided in Appendix D.

Method	Corr \uparrow	HD \downarrow	KLD \downarrow
PCA-DMD	0.9504-0.9800	0.0010-0.0072	0.0005-0.0022
DMD	~ 0.10	$\sim 0.10-0.26$	$\sim 3-5$
HODMD	~ 0.10	$\sim 0.10-0.26$	$\sim 3-5$
HankelDMD	~ 0.10	$\sim 0.10-0.26$	$\sim 3-5$
SpDMD	$\sim 0.15-0.27$	$\sim 0.07-0.25$	$\sim 1.7-3.2$

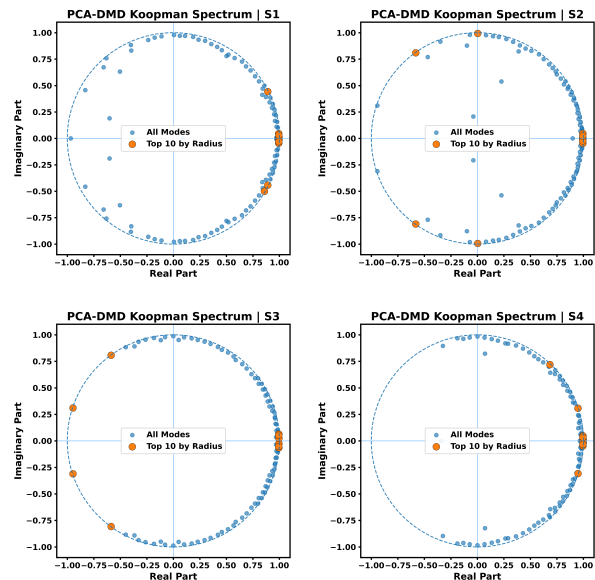


Figure 4. Koopman eigenvalue spectra of PCA-DMD under cross-subject zero-shot reconstruction. Across subjects, dominant eigenvalues cluster near the unit circle, indicating stable latent dynamics and consistent temporal structure. Additional spectra for baseline methods are provided in the appendix C.

SpDMD, and HankelDMD. These results show that combining PCA-based dimensionality reduction with Koopman-inspired latent evolution yields a scalable and robust framework for high-dimensional time-series reconstruction. Koopman spectral analysis further reveals consistent eigenvalue structure across subjects, supporting the stability and transferability of the learned latent dynamics. Overall, PCA-DMD provides an interpretable and computationally efficient approach for modeling long-range neural dynamics in large multichannel recordings.

Future Work. Future directions include extending the framework toward multi-scale and adaptive representations, incorporating localized spectral modeling, and exploring hybrid dynamical systems formulations that preserve the scalability and interpretability of PCA-DMD while capturing richer temporal structure.

References

- Brenner, M., Hess, F., Mikhaeil, J. M., Bereska, L. F., Monfared, Z., Kuo, P.-C., and Durstewitz, D. Tractable dendritic RNNs for reconstructing nonlinear dynamical systems. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S. (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 2292–2320. PMLR, 17–23 Jul 2022.
- Brenner, M., Hess, F., Mikhaeil, J. M., Bereska, L., Monfared, Z., Kuo, P.-C., and Durstewitz, D. A scalable generative model for dynamical system reconstruction from neuroimaging data, 2024. URL <https://arxiv.org/abs/2411.02949>.
- Brunton, S. L. and Kutz, J. N. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2022.
- Brunton, S. L., Brunton, B. W., Proctor, J. L., Kaiser, E., and Kutz, J. N. Chaos as an intermittently forced linear system. *Nature Communications*, 8:19, 2017. doi: 10.1038/s41467-017-00030-8.
- Buzsáki, G. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10):1073–1188, 2015. doi: 10.1002/hipo.22488.
- Buzsáki, G., Anastassiou, C. A., and Koch, C. The origin of extracellular fields and currents—eeg, ecog, lfp and spikes. *Nature Reviews Neuroscience*, 13(6):407–420, 2012. doi: 10.1038/nrn3241.
- Cunningham, J. P. and Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, 2014. doi: 10.1038/nn.3776.
- Diba, K. and Buzsáki, G. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10(10):1241–1242, 2007. doi: 10.1038/nn1961.
- Durstewitz, D., Koppe, G., and Thurm, M. I. Reconstructing computational system dynamics from neural data with recurrent neural networks. *Nature Reviews Neuroscience*, 24(11):693–710, 2023. doi: 10.1038/s41583-023-00740-7.
- Foster, D. J. Replay comes of age. *Annual Review of Neuroscience*, 40:581–602, 2017. doi: 10.1146/annurev-neuro-072116-031538.
- Hess, F., Monfared, Z., Brenner, M., and Durstewitz, D. Generalized Teacher Forcing for Learning Chaotic Dynamics. In *Proceedings of the 40th International Conference on Machine Learning*. PMLR, July 2023. ISSN: 2640-3498.
- Joo, H. R. and Frank, L. M. The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation. *Nature Reviews Neuroscience*, 19(12):744–757, 2018. doi: 10.1038/s41583-018-0077-1.
- Jovanović, M. R., Schmid, P. J., and Nichols, J. W. Sparsity-promoting dynamic mode decomposition. *Physics of Fluids*, 26(2):024103, 2014. doi: 10.1063/1.4863670.
- Koopman, B. O. Hamiltonian systems and transformation in hilbert space. *Proceedings of the National Academy of Sciences*, 17(5):315–318, 1931. doi: 10.1073/pnas.17.5.315.
- Kujur, A., Monfared, Z., and Safavi, S. Transient neural dynamics reconstruction. In *NeurIPS 2025 Workshop on Learning from Time Series for Health*, 2025.
- Kutz, J. N., Brunton, S. L., Brunton, B. W., and Proctor, J. L. *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016a. doi: 10.1137/1.9781611974508.
- Kutz, J. N., Fu, X., and Brunton, S. L. Multiresolution dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 15(2):713–735, 2016b. doi: 10.1137/15M1023543.
- Le Clainche, S. and Vega, J. M. Higher order dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 16(2):882–925, 2017. doi: 10.1137/15M1054924.
- Meier, K., Merseburg, A., Isbrandt, D., Marguet, S. L., and Morellini, F. 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.
- Navas-Olive, A., Rubio, A., Abbaspoor, S., Hoffman, K. L., and de la Prida, L. M. A machine learning toolbox for the analysis of sharp-wave ripples reveals common waveform features across species. *Communications Biology*, 7(1): 211, 2024.
- Pandarínath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., others, and Sussillo, D. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10):805–815, 2018. doi: 10.1038/s41592-018-0109-9.
- Rowley, C. W., Mezić, I., Bagheri, S., Schlatter, P., and Henningson, D. S. Spectral analysis of nonlinear flows. *Journal of Fluid Mechanics*, 641:115–127, 2009. doi: 10.1017/S0022112009992059.
- Schmid, P. J. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656: 5–28, 2010. doi: 10.1017/S0022112010001217.

330 Schmid, P. J. Dynamic mode decomposition and its variants.
331 *Annual Review of Fluid Mechanics*, 54:225–254, 2022.
332 doi: 10.1146/annurev-fluid-030121-015835.

333
334 Wilson, M. A. and McNaughton, B. L. Reactivation of hip-
335 pocampal ensemble memories during sleep. *Science*, 265
336 (5172):676–679, 1994. doi: 10.1126/science.8036517.

337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384

Technical Appendices and Supplementary Material

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

Fig. 5 presents boxplot comparisons of reconstruction metrics across all methods under cross-subject zero-shot evaluation. PCA-DMD consistently achieves the highest correlation and the lowest HD and KLD, with substantially smaller variability than the baseline DMD variants. These distributions further support the robustness of PCA-DMD for long-horizon neural signal reconstruction.

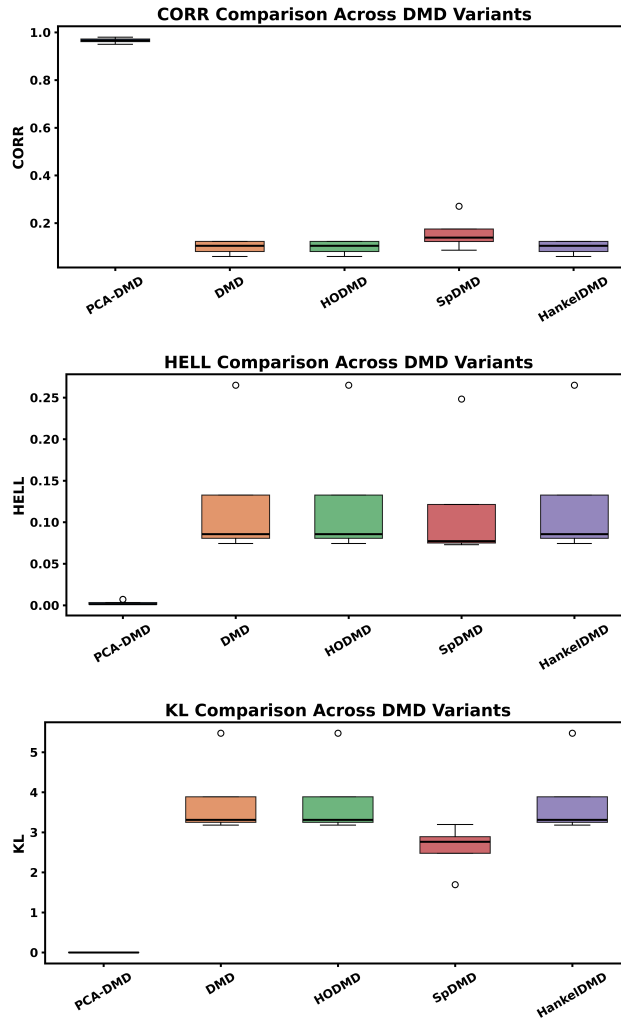


Figure 5. Comparison of reconstruction metrics across all methods in the cross-subject zero-shot setting: correlation (top), HD (middle), and KLD (bottom). PCA-DMD consistently outperforms the baseline DMD variants in both fidelity and robustness.

A. Additional Results

We provide additional reconstruction examples for the baseline DMD variants across subjects. Compared with PCA-DMD, these methods show visibly weaker temporal alignment, larger amplitude distortions, and reduced stability over long horizons. These qualitative results are consistent with the quantitative trends reported in the main text.

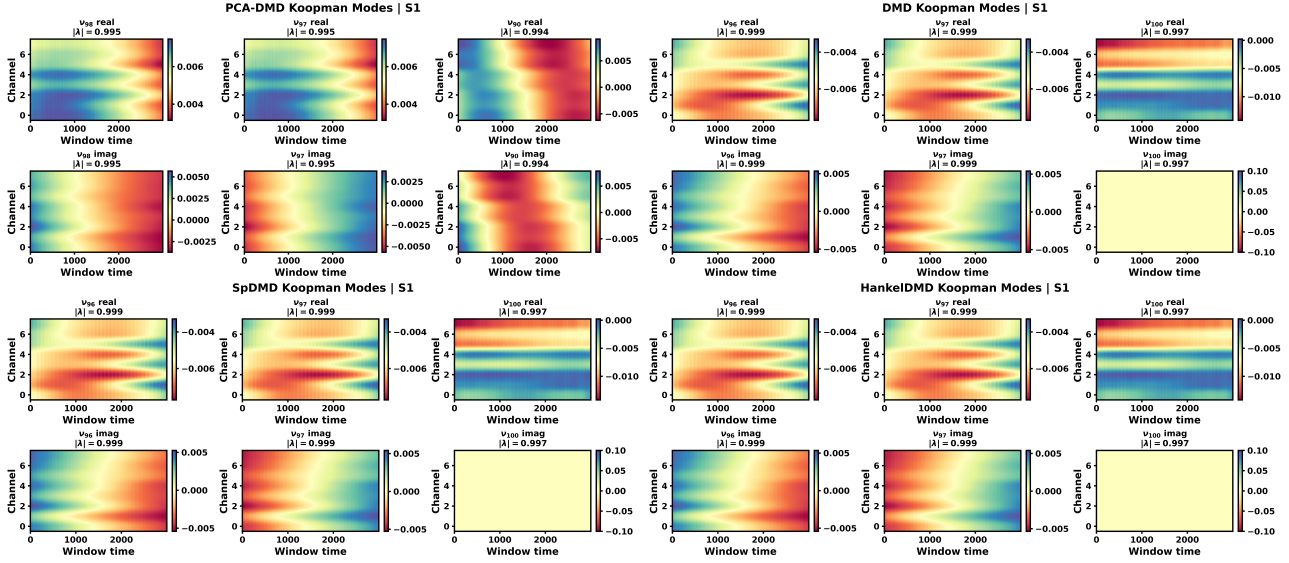


Figure 6. Koopman modes for DMD variants for Subject S1. The real and imaginary parts are shown separately.

B. Additional Koopman Mode

C. Additional Koopman Spectra

This section contains remaining koopman eigenvalue spectra for DMD variants across subjects. PCA-DMD yields spectra that are more consistently organized near the unit circle, indicating stable latent dynamics across recordings. In contrast, the baseline methods show less coherent spectral structure, reflecting their reduced reconstruction stability under long-horizon and cross-subject settings.

D. Full Cross-Subject Reconstruction

Table 2 reports the complete pairwise cross-subject zero-shot reconstruction results for PCA-DMD, together with aggregate ranges for the baseline DMD variants. PCA-DMD consistently maintains high correlation and very low HD and KLD across all ordered source-target pairs.

Table 2. Full cross-subject zero-shot reconstruction results. PCA-DMD shows consistently strong transfer across all source-target pairs, whereas the baseline DMD variants degrade substantially under long-horizon reconstruction.

Pair / Method	Corr \uparrow	HD \downarrow	KLD \downarrow
S1 \rightarrow S2	0.9504	0.0020	0.0019
S1 \rightarrow S3	0.9671	0.0011	0.0006
S1 \rightarrow S4	0.9778	0.0013	0.0007
S2 \rightarrow S1	0.9623	0.0072	0.0020
S2 \rightarrow S3	0.9693	0.0011	0.0007
S2 \rightarrow S4	0.9798	0.0013	0.0005
S3 \rightarrow S1	0.9624	0.0071	0.0022
S3 \rightarrow S2	0.9553	0.0017	0.0010
S3 \rightarrow S4	0.9800	0.0012	0.0005
S4 \rightarrow S1	0.9643	0.0071	0.0020
S4 \rightarrow S2	0.9589	0.0016	0.0013
S4 \rightarrow S3	0.9707	0.0010	0.0007
DMD/HODMD/HankelDMD (avg.)	~ 0.10	$\sim 0.10\text{--}0.26$	$\sim 3\text{--}5$
SpDMD (avg.)	$\sim 0.15\text{--}0.27$	$\sim 0.07\text{--}0.25$	$\sim 1.7\text{--}3.2$

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,

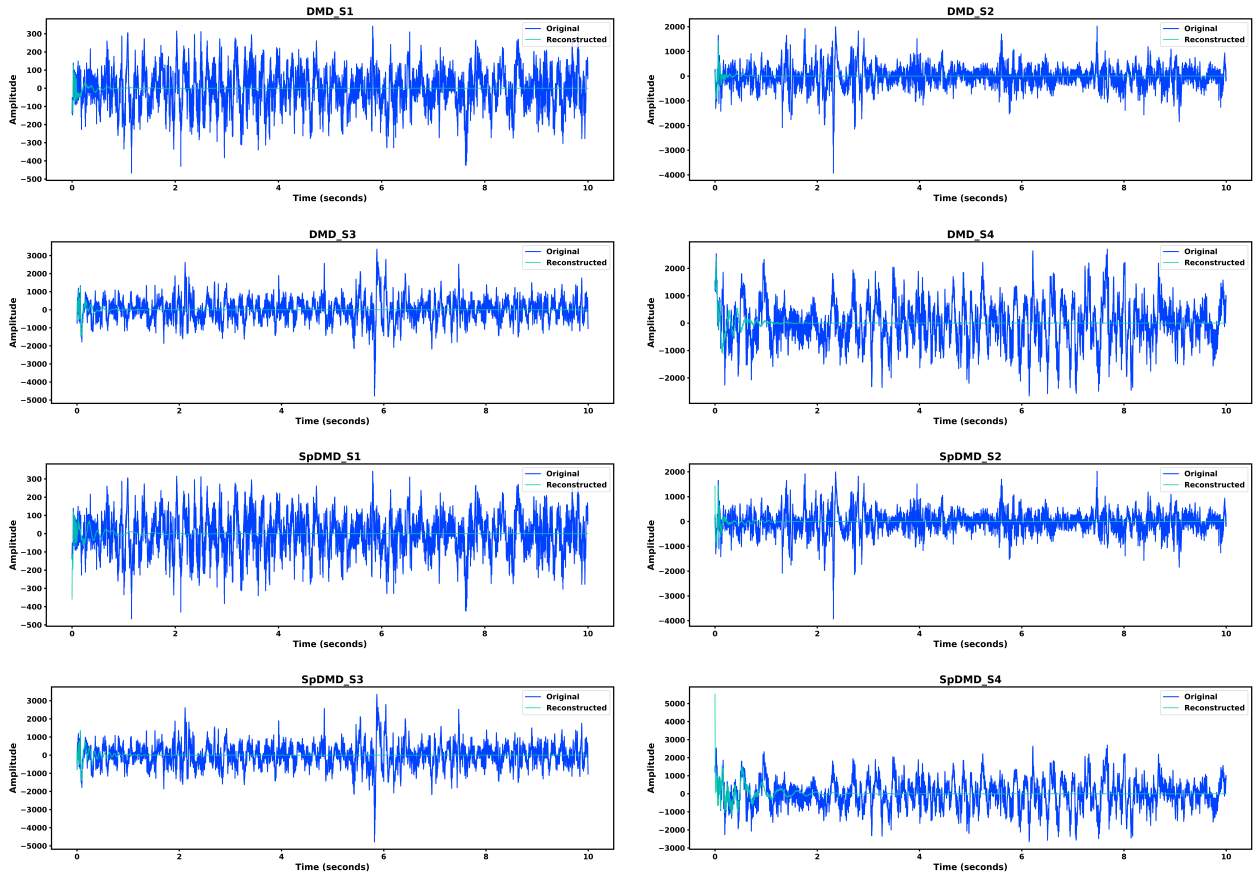


Figure 7. Additional cross-subject reconstruction examples for DMD variants.

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.

E. Additional Koopman Spectral Analysis

We provide additional analysis of the Koopman eigenvalue radius distributions across all subjects and methods. These histograms illustrate the distribution of $|\lambda|$ and provide insight into stability and spectral structure of the learned operators.

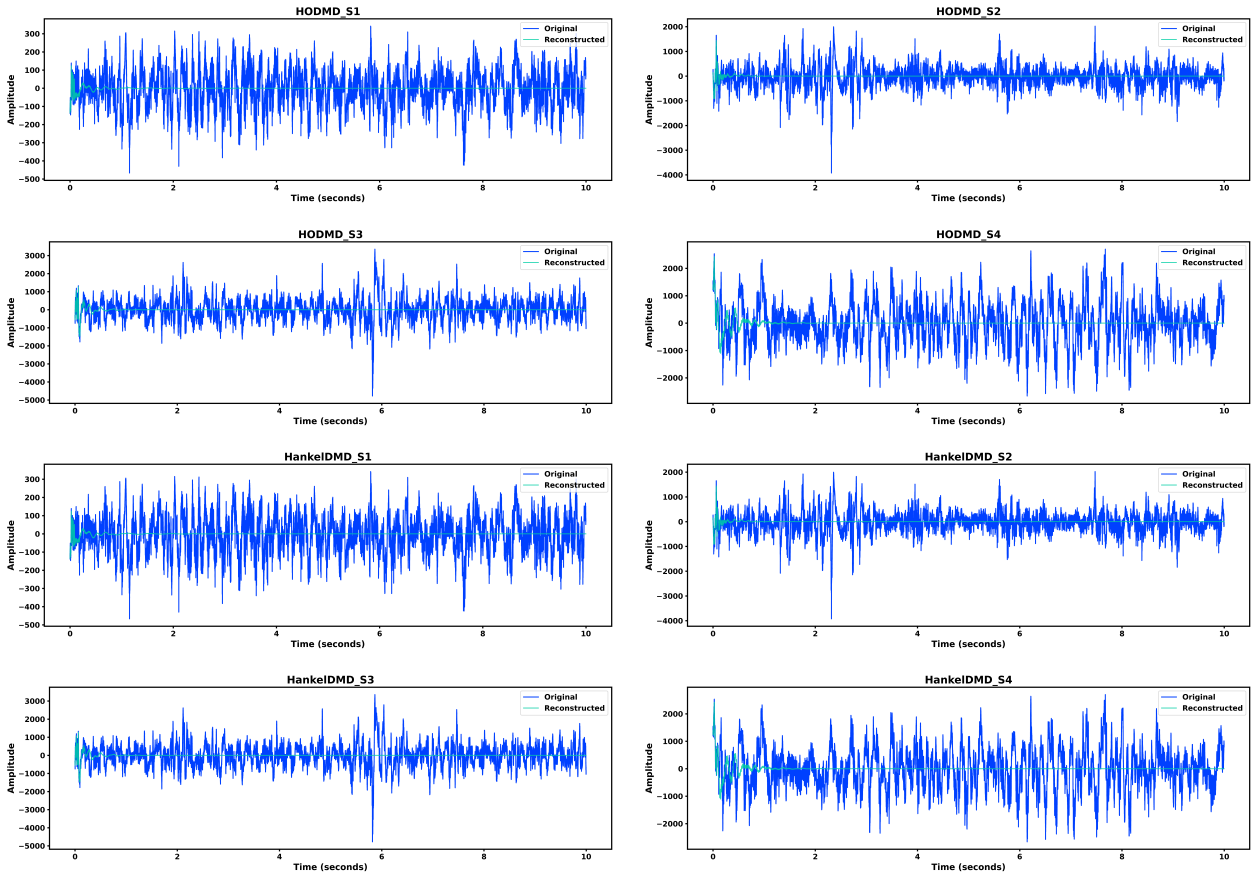


Figure 8. Additional cross-subject reconstruction examples for DMD variants. Each pair of panels corresponds to one method across subjects.

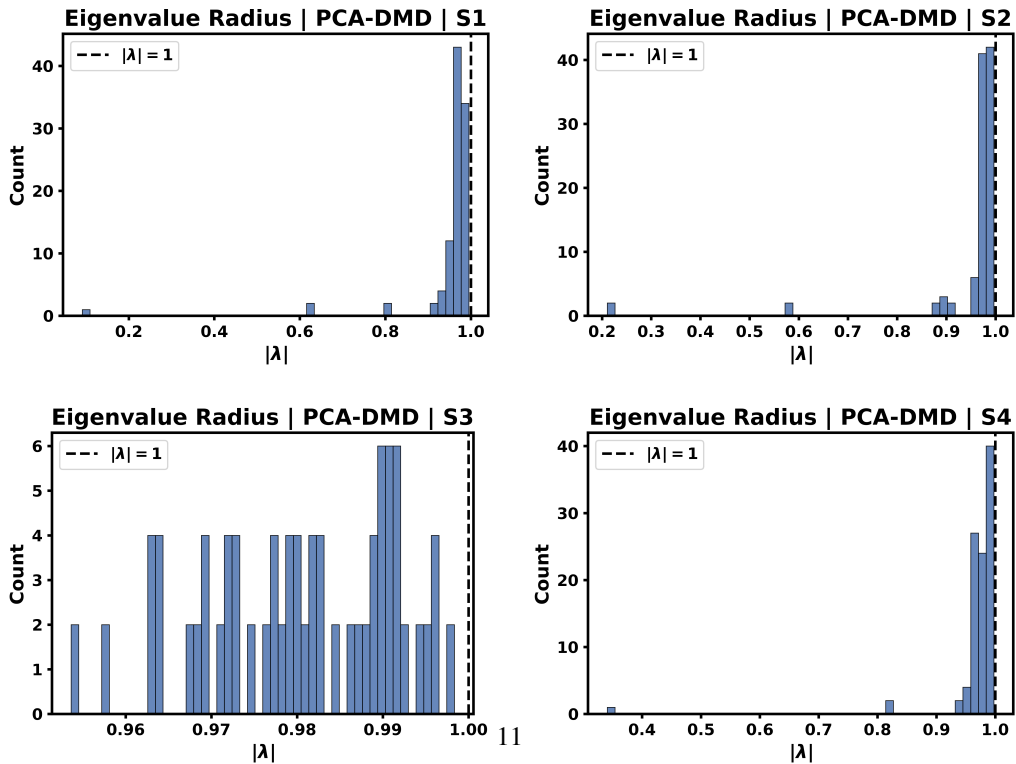


Figure 11. Eigenvalue radius distribution ($|\lambda|$) for PCA-DMD across all subjects.

550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604




Figure 9. Additional Koopman eigenvalue spectra across subjects for DMD.

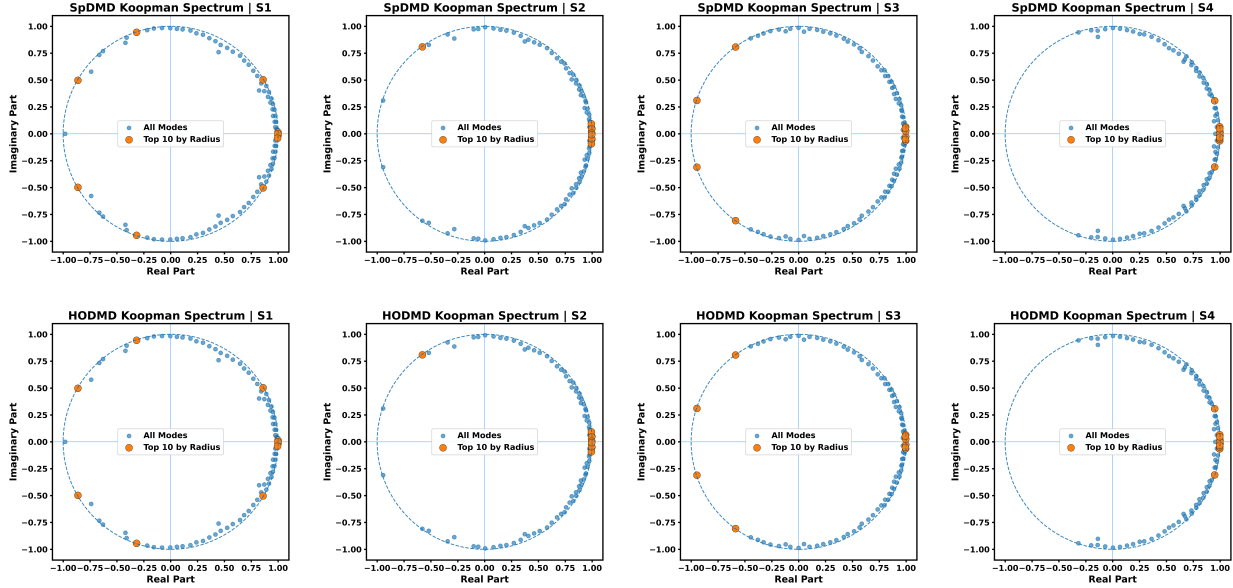


Figure 10. Additional Koopman eigenvalue spectra across for DMD variants.

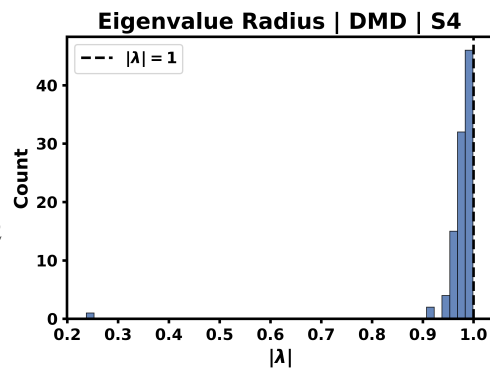
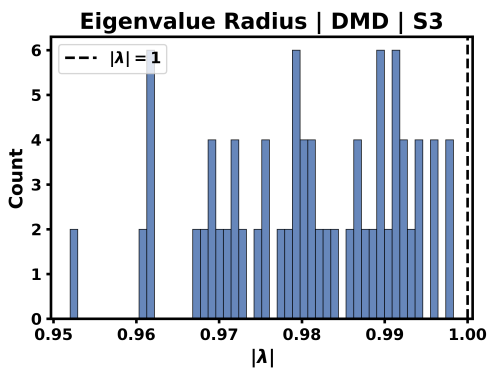
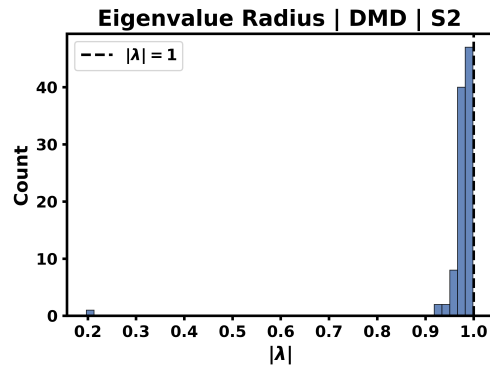
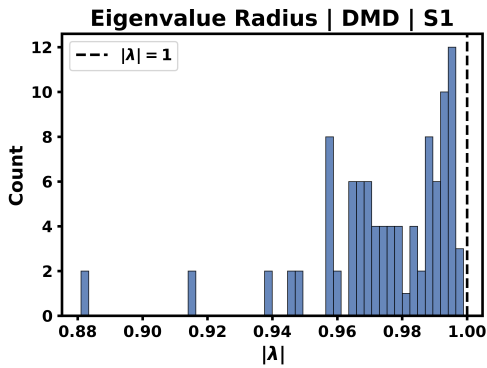


Table 3. Experimental settings for all reconstruction methods.

Method	Rank	Window	Step	Latent
PCA-DMD	100	3000	30	100
DMD	100	3000	30	–
HODMD	100	3000	30	–
SpDMD	100	3000	30	–
HankelDMD	100	3000	30	–

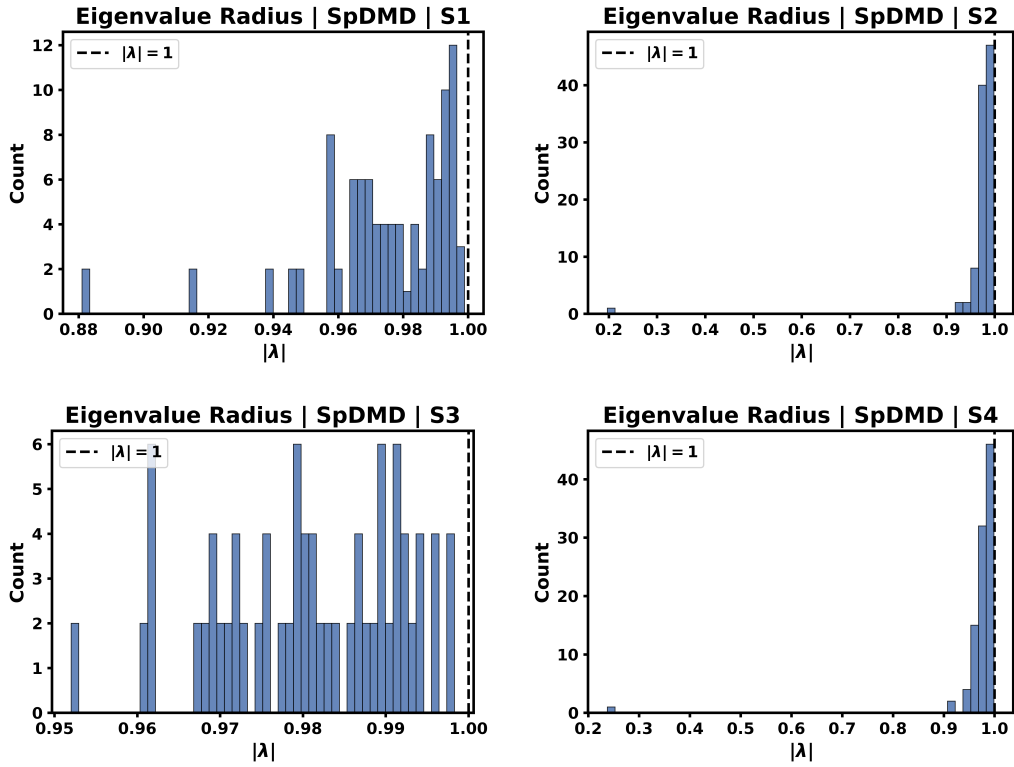


Figure 13. Eigenvalue radius distribution ($|\lambda|$) for SpDMD across all subjects.

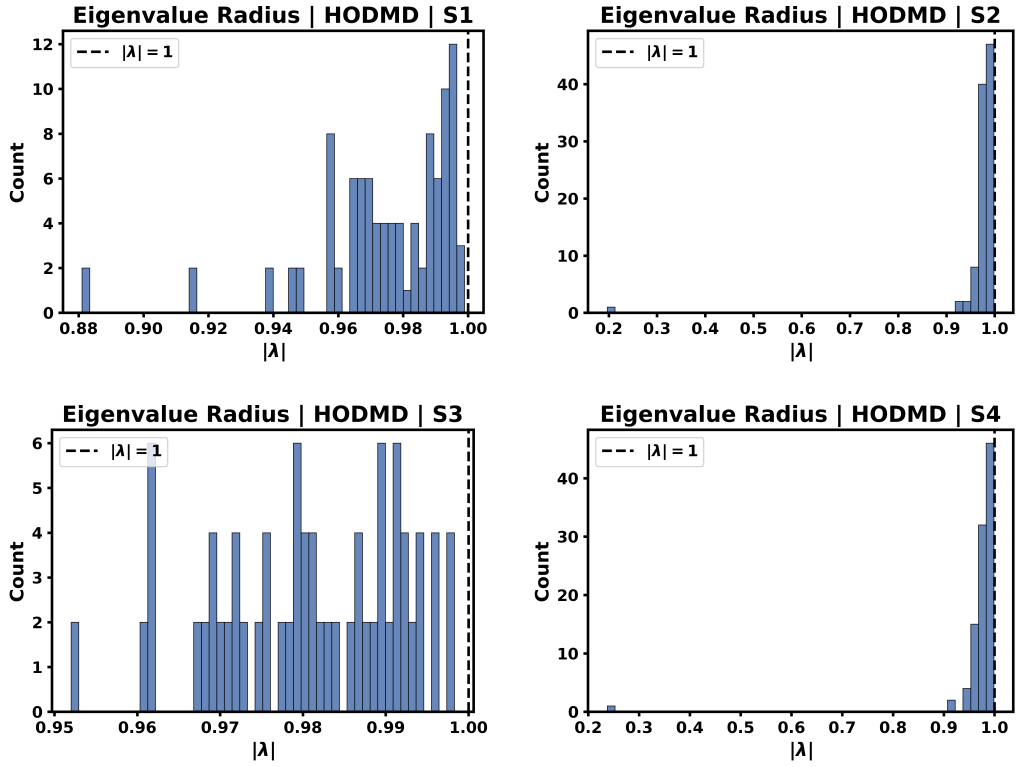


Figure 14. Eigenvalue radius distribution ($|\lambda|$) for HODMD across all subjects.




Figure 16. PCA-DMD reconstruction overlays for zero-shot transfer from subject S1 to S2, S3, and S4.

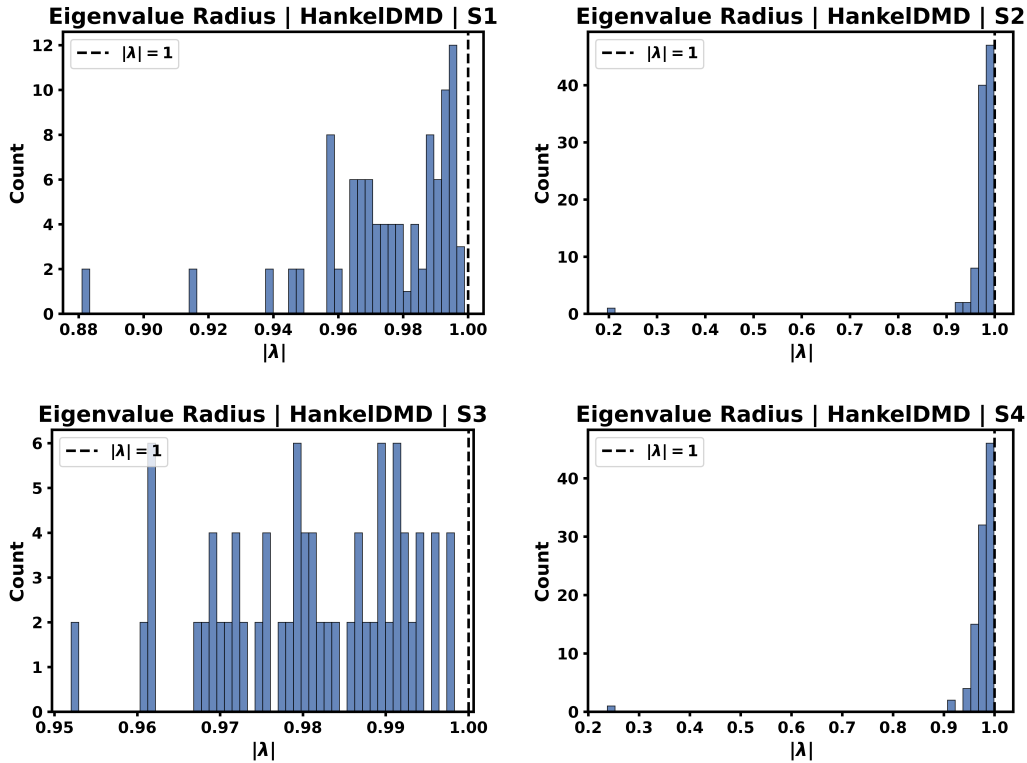


Figure 15. Eigenvalue radius distribution ($|\lambda|$) for HankelDMD across all subjects.

F. Additional Results for Cross-Subject Reconstruction (Transfer from Subject S1)

The reconstruction overlays show that all DMD-based models perform reasonably well for short signal windows of length 3000 samples. However, as the signal length increases to long-horizon sequences, reconstruction quality deteriorates significantly for standard DMD variants, including DMD, SpDMD, HODMD, and HankelDMD. In contrast, PCA-DMD remains stable and consistently achieves superior long-range reconstruction and cross-subject generalization performance.




Figure 17. DMD reconstruction overlays for zero-shot transfer from subject S1 to S2, S3, and S4.




Figure 18. SpDMD reconstruction overlays for zero-shot transfer from subject S1 to S2, S3, and S4.

G. Koopman Operator and DMD Variants

G.1. 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

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

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)*:

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

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).

G.2. 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,




Figure 19. HODMD reconstruction overlays for zero-shot transfer from subject S1 to S2, S3, and S4.




Figure 20. HankelDMD reconstruction overlays for zero-shot transfer from subject S1 to S2, S3, and S4.

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:

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

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 $t \in [\delta, \min(N, T - \delta)]$.

G.3. 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 $t \in [\delta, \min(N, T - \delta)]$.

G.4. 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 $t \in [\delta, \min(N, T - \delta)]$.

H. Evaluation Measures

H.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.

H.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,

990 ranging from 0 (perfect agreement) to 1, was calculated per channel and averaged to produce HD. The Hellinger distance for
991 the i -th channel is defined as:

$$992 \quad HD(f_i(\omega), g_i(\omega)) = \sqrt{1 - \int_{-\infty}^{\infty} \sqrt{f_i(\omega)g_i(\omega)} d\omega},$$

993
994 where $f_i(\omega)$ and $g_i(\omega)$ are the normalized power spectra of the true and reconstructed signals for the i -th channel.
995
996
997
998
999

1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044