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

Anima Kujur ANIMA.KUJUR@IWR.UNI-HEIDELBERG.DE
Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University, Heidelberg, Germany

Zahra Monfared ZAHRA.MONFARED@IWR.UNI-HEIDELBERG.DE
Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University, Heidelberg, Germany
Department of Mathematics and Computer Science, Heidelberg University, Heidelberg, Germany

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.

Keywords: Neural Dynamics, High-Dimensional Time Series, Signal Reconstruction, Cross-Subject Generalization, Zero-Shot Transfer, Dynamical Systems

1. Introduction

High-dimensional time series arise in neuroscience, biology, and climate science, where the goal is to learn compact representations that preserve the essential structure of the underlying dynamics. This challenge is especially pronounced in neural recordings because of their large scale, high temporal resolution, and nonlinear structure. Long-horizon reconstruction is therefore an important benchmark for assessing whether a model captures meaningful dynamical organization over extended time intervals. LFPs provide a rich example, reflecting coordinated population-level neural activity across multiple channels and timescales [7, 8]. Their high sampling rate and long duration make them a challenging testbed for scalable dynamical modeling, since many existing signal models suffer from memory limitations, long-horizon instability, and limited interpretability. Koopman operator theory offers a principled framework for analyzing nonlinear dynamics through linear evolution in an observable space [16], while DMD provides a practical data-driven approximation for extracting coherent spatiotemporal patterns [18, 24, 25]. However, applying DMD directly to large neural recordings remains difficult because of the high dimensionality of windowed signals and unstable long-horizon reconstruction. Recent work introduced PCA-DMD as a scalable framework that combines dimensionality reduction with Koopman-based latent evolution [17]. By projecting

© A. Kujur & Z. Monfared.

high-dimensional windows into a compact PCA subspace before fitting a linear Koopman operator, PCA-DMD improves computational tractability and reconstruction stability, supporting reconstruction up to 200k samples while outperforming standard DMD variants.

In this work, we extend PCA-DMD to longer horizons of up to 300k samples and introduce a cross-subject zero-shot protocol across four subjects to evaluate whether learned latent dynamics generalize beyond a single recording. This setting is important for neural data, where robust transfer across subjects is difficult but practically valuable. We further analyze learned Koopman representations through mode-level visualization, showing that PCA-DMD yields richer and less redundant modes than classical DMD variants, with structured global patterns and higher-frequency channel-time variations. Our results show that PCA-DMD maintains high reconstruction quality in long-horizon and zero-shot settings, achieving correlation above 0.95 with very low Hellinger distance (HD) and Kullback-Leibler divergence (KLD) across source-target pairs. In contrast, classical DMD variants degrade as signal length increases and fail to transfer reliably across subjects. Overall, this work demonstrates that PCA-DMD provides scalable, transferable, and interpretable reconstruction of high-dimensional hippocampal LFP dynamics.

2. Related Work

Neural signals and reconstruction. LFPs provide a mesoscopic view of coordinated neural population activity and are widely used to study circuit-level dynamics [7, 8]. In the hippocampus, they exhibit multi-scale temporal structure and spatial organization [10, 12, 14, 27], making them valuable but challenging due to high dimensionality and long duration. Neural time-series reconstruction has been studied with latent variable models and recurrent architectures [9, 11, 23], but these methods struggle with long-horizon stability, memory constraints, and interpretability [3, 4, 13].

Koopman modeling and scalable reconstruction. Koopman operator theory represents nonlinear dynamics through linear evolution in a lifted space [16], and DMD provides a finite-dimensional approximation for extracting coherent spatiotemporal structures [18, 24, 25]. Its modes and spectra provide useful interpretability [5, 26]. Several extensions, including sparsity-promoting DMD [15], higher-order DMD [20], multiresolution DMD [19], and delay-embedding-based HAVOK [6], improve robustness but remain difficult to scale to long, high-dimensional neural recordings.

Position of this work. Dimensionality reduction, particularly PCA, is widely used to compress neural recordings while preserving dominant variance [9]. Building on this idea, Kujur et al. [17] introduced PCA-DMD for scalable long-horizon LFP reconstruction up to 200k samples. We extend this framework to 300k samples and cross-subject zero-shot evaluation, emphasizing scalable reconstruction, transferability, and the structure of learned Koopman representations.

3. Methodology

Fig. 1 summarizes the PCA-DMD pipeline. Multichannel hippocampal LFPs are segmented into overlapping windows, vectorized, projected into a compact PCA latent space, evolved using a learned Koopman/DMD operator, and reconstructed by inverse projection and overlap-add averaging. The model is evaluated in a cross-subject zero-shot setting, where parameters learned from one subject are applied directly to unseen subjects without fine-tuning.

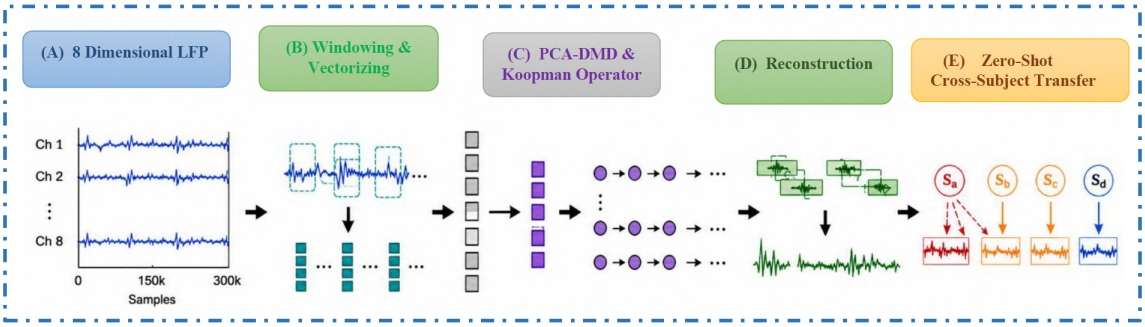

Figure 1: **Method overview.** LFP windows are projected into a PCA latent space, evolved by a Koopman/DMD operator, reconstructed by overlap-add averaging, and evaluated under cross-subject zero-shot transfer.

3.1. Problem Setup

We consider very long-horizon reconstruction of rodent hippocampal LFPs [21, 22]. The dataset contains four subjects, $q \in \{S1, S2, S3, S4\}$, each with $C = 8$ channels sampled at $f_s = 30{,}000$ Hz and approximately 22 million samples per channel. For scalability experiments, we retain up to $T = 300{,}000$ samples per subject. Let $s^{(q)}(t) \in \mathbb{R}^C$ denote the multichannel signal. Each recording is divided into overlapping windows of length $w = 3000$ and stride $\delta = 30$, where $W_i = s^{(q)}[i\delta : i\delta + w - 1] \in \mathbb{R}^{w \times C}$ is vectorized as $x_i = \mathrm{vec}(W_i) \in \mathbb{R}^{d_x}$ with $d_x = wC$. Consecutive windows form paired samples $X = [x_1^\top, \ldots, x_{N-1}^\top]^\top$ and $Y = [x_2^\top, \ldots, x_N^\top]^\top$ for one-step window prediction.

3.2. PCA-DMD Reconstruction

PCA-DMD projects each high-dimensional window x_i onto a PCA basis $P \in \mathbb{R}^{d_x \times r}$ with $r = 100$, giving $z_i = P^\top x_i$ and $\tilde{z}_i = P^\top y_i$. A Koopman operator $K \in \mathbb{R}^{r \times r}$ is learned such that $\tilde{z}_i \approx K z_i$, using least squares $K = \arg\min_A \|Z_+ - ZA^\top\|_F^2$ with $K^\top = Z^\dagger Z_+$. Prediction is performed as $\hat{z}_i = K z_i$ and $\hat{x}_i = P\hat{z}_i + \mu$. The reconstructed windows are reshaped and combined by overlap-add averaging, $\hat{s}(t) = \frac{1}{n_t} \sum_{i: t \in [i\delta, i\delta + w)} \hat{W}_i(t - i\delta)$, yielding a continuous long-horizon reconstruction.

3.3. Zero-Shot Evaluation and Metrics

For each ordered subject pair (a, b) with $a \neq b$, the model is trained on subject a and applied directly to subject b, i.e., $\hat{s}^{(b)} = f_{\theta_a}(s^{(b)})$. Reconstruction quality is measured using Pearson correlation, Hellinger distance, and Kullback–Leibler divergence: $\mathrm{Corr}(x, \hat{x}) = \frac{\mathrm{Cov}(x, \hat{x})}{\sigma_x \sigma_{\hat{x}}}$, $\mathrm{HD}(p, \hat{p}) = \frac{1}{\sqrt{2}} \|\sqrt{p} - \sqrt{\hat{p}}\|_2$, and $\mathrm{KLD}(p \| \hat{p}) = \sum_j p_j \log \frac{p_j}{\hat{p}_j}$, where $x = \mathrm{vec}(s)$, $\hat{x} = \mathrm{vec}(\hat{s})$, and $p, \hat{p}$ are histogram-based amplitude distributions. Additional details on PCA compression, Koopman spectra, and mode analysis are provided in Appendix A.

4. Results

We evaluate PCA-DMD under long-horizon reconstruction, cross-subject zero-shot transfer, Koopman mode analysis, and external validation. Across all experiments, PCA-DMD achieves stronger

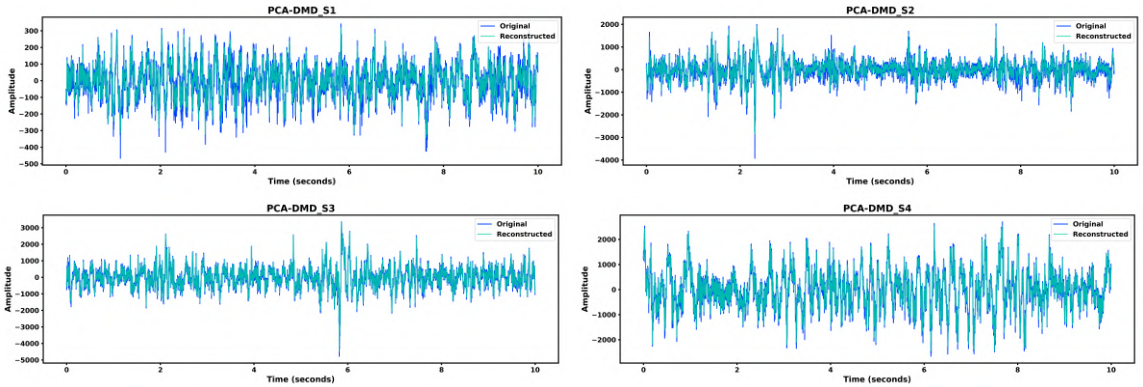

Figure 2: Zero-shot PCA-DMD reconstructions across four subjects.

reconstruction fidelity and transfer robustness than standard DMD variants. The final hyperparameters were selected from an ablation over window sizes $\{150, 200, 300, 500, 1000, 3000\}$, step sizes $\{10, 20, 30, 50, 60\}$, and latent dimensions $\{8, 20, 50, 100\}$. The setting $w = 3000$, $\delta = 30$, and $r = 100$ gave the best balance between long temporal context, stable overlapping-window transitions, and compact latent representation; therefore, it was used for all subsequent experiments.

Long-horizon cross-subject reconstruction. Fig. 2 shows representative zero-shot PCA-DMD reconstructions across all four subjects. The reconstructed signals closely follow the original LFP waveforms over long horizons, preserving global temporal structure and amplitude variation. In contrast, standard DMD variants show clear degradation, including amplitude attenuation, waveform distortion, and reduced temporal alignment, especially as the reconstruction length increases.

Comparison across DMD variants. All methods are evaluated under identical preprocessing and windowing settings, with a fixed window size and step. PCA-DMD additionally uses a low-dimensional latent representation ($r = 100$), which reduces the effective state dimension of the windowed multichannel signal and improves the numerical conditioning of Koopman regression. Quantitative results are summarized in Fig. 3B. PCA-DMD achieves high correlation (0.9504-0.9800) with very low HD (0.0010-0.0072) and KLD (0.0005-0.0022) across source–target pairs. In contrast, DMD, HODMD, and HankelDMD show near-failure regimes with low correlation and large distributional mismatch, while SpDMD provides only limited improvement. Full source–target results and comparison plots are provided in Appendix C.

Koopman mode analysis. Fig. 3A shows the PCA-DMD Koopman mode structure for subject S1, including real and imaginary components of representative dominant modes across window time and channels. The modes exhibit structured spatiotemporal organization rather than redundant smooth patterns, suggesting that PCA-DMD learns a richer and better-conditioned dynamical representation than standard DMD variants. Koopman eigenvalue spectra and their interpretation are moved to Appendix E to keep the main results compact.

4.1. External Validation on Allen Neuropixels LFP

To further assess scalability, we validate PCA-DMD on Neuropixels LFP recordings from the Allen Brain Observatory Visual Coding dataset [1, 2]. This experiment uses one Neuropixels probe with 93 channels and 300,000 samples sampled at approximately 1250 Hz, corresponding to 240 sec-

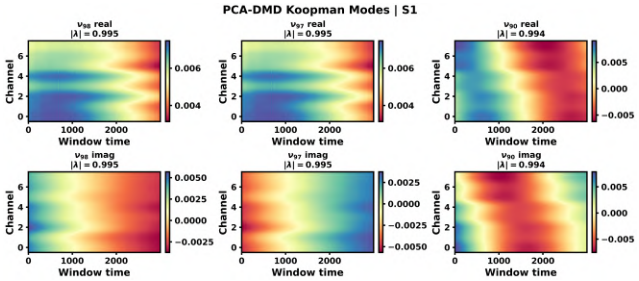

(A) PCA-DMD Koopman modes for S1.

(B) Cross-subject zero-shot summary.

Method	Corr ↑	HD ↓	KLD ↓
PCA-DMD	.950-.980	.001-.007	.0005-.002
DMD	∼.10	.10-.26	3-5
HODMD	∼.10	.10-.26	3-5
HankelDMD	∼.10	.10-.26	3-5
SpDMD	.15-.27	.07-.25	1.7-3.2

PCA-DMD gives the strongest reconstruction and distributional agreement.

Figure 3: Koopman mode structure and quantitative cross-subject reconstruction summary. Additional Koopman spectra and baseline mode visualizations are provided in Appendix E and D.

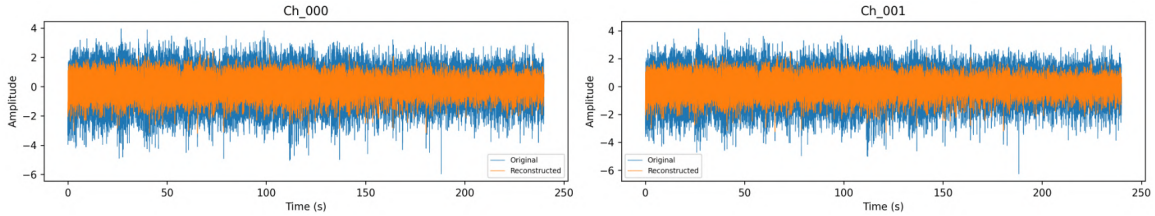

Figure 4: External validation on Allen Neuropixels LFP: representative reconstruction for the first four channels from a 93-channel, 300,000-sample recording.

onds. Each PCA-DMD input window has dimension $3000 \times 93 = 279{,}000$, making this a substantially higher-dimensional setting than the main 8-channel LFP experiments. Fig. 4 shows that PCA-DMD preserves the dominant temporal structure across representative channels while smoothing high-frequency fluctuations. Additional validation details and all remaining Allen Neuropixels reconstruction plots are provided in Appendix G.

5. Conclusion

We presented a systematic study of long-horizon neural signal reconstruction using PCA-DMD in a cross-subject zero-shot setting. Extending prior work, we showed that PCA-DMD scales to 300k samples while maintaining strong reconstruction fidelity across unseen subjects. Allen Neuropixels validation further confirms scalability from the original 8-channel LFP setup to substantially higher-dimensional 93-channel recordings, indicating that PCA-DMD can capture collective cross-channel dynamics while remaining computationally tractable. Across source-target pairs, PCA-DMD achieves high correlation and low distributional mismatch (HD and KLD), outperforming DMD, HODMD, SpDMD, and HankelDMD. These results show that combining PCA-based dimensionality reduction with Koopman-based latent evolution yields a scalable, robust, and interpretable framework for high-dimensional neural time-series reconstruction. Koopman spectral analysis further supports the stability and transferability of the learned latent dynamics.

Future Work. Future work will explore multi-scale and adaptive representations, localized spectral modeling, and hybrid dynamical systems formulations that preserve the scalability and interpretability of PCA-DMD while capturing richer temporal structure.

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] Allen Institute. Allen Institute - Visual Behavior - Neuropixels, 2024. URL https://doi.org/10.48324/dandi.000713/0.240702.1725. Dataset.

[2] Corbett Bennett, Sam Gale, Greggory Heller, Tamina Ramirez, Hannah Belski, Alex T. Piet, Omid Zobeiri, et al. Map of spiking activity underlying change detection in the mouse visual system. *bioRxiv*, 2025. doi: 10.1101/2025.10.17.683190. URL https://doi.org/10.1101/2025.10.17.683190.

[3] Manuel Brenner, Florian Hess, Jonas M Mikhaeil, Leonard F Bereska, Zahra Monfared, Po-Chen Kuo, and Daniel Durstewitz. Tractable dendritic RNNs for reconstructing nonlinear dynamical systems. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 2292–2320. PMLR, 17–23 Jul 2022.

[4] Maximilian Brenner, Florian Hess, Jonas M. Mikhaeil, Leonard Bereska, Zahra Monfared, Po-Chen Kuo, and Daniel Durstewitz. A scalable generative model for dynamical system reconstruction from neuroimaging data, 2024. URL https://arxiv.org/abs/2411.02949.

[5] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2022.

[6] Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, Eurika Kaiser, and J. Nathan Kutz. Chaos as an intermittently forced linear system. *Nature Communications*, 8:19, 2017. doi: 10.1038/s41467-017-00030-8.

[7] György Buzsáki. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10):1073–1188, 2015. doi: 10.1002/hipo.22488.

[8] 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. doi: 10.1038/nrn3241.

[9] John P. Cunningham and Byron M. Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, 2014. doi: 10.1038/nn.3776.

[10] Kamran Diba and György Buzsáki. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10(10):1241–1242, 2007. doi: 10.1038/nn1961.

[11] Daniel Durstewitz, Georgia Koppe, and Max Ingo Thurm. 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.

[12] David J. Foster. Replay comes of age. *Annual Review of Neuroscience*, 40:581–602, 2017. doi: 10.1146/annurev-neuro-072116-031538.

[13] 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.

[14] 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. doi: 10.1038/s41583-018-0077-1.

[15] Mihailo R. Jovanović, Peter J. Schmid, and Joseph W. Nichols. Sparsity-promoting dynamic mode decomposition. *Physics of Fluids*, 26(2):024103, 2014. doi: 10.1063/1.4863670.

[16] Bernard O. Koopman. 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.

[17] A. Kujur, Z. Monfared, and S. Safavi. Transient neural dynamics reconstruction. In *NeurIPS 2025 Workshop on Learning from Time Series for Health*, 2025.

[18] J. Nathan Kutz, Steven L. Brunton, Bingni W. Brunton, and Joshua L. Proctor. *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016. doi: 10.1137/1.9781611974508.

[19] 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.

[20] Soledad Le Clainche and José M. Vega. Higher order dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 16(2):882–925, 2017. doi: 10.1137/15M1054924.

[21] 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.

[22] Andrea Navas-Olive, Adrian Rubio, Saman Abbaspoor, 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.

[23] Chethan Pandarinath, Daniel J. O'Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D. Stavisky, Jonathan C. Kao, others, and David Sussillo. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10):805–815, 2018. doi: 10.1038/s41592-018-0109-9.

[24] Clarence W. Rowley, Igor Mezić, Shervin Bagheri, Philipp Schlatter, and Dan S. Henningson. Spectral analysis of nonlinear flows. *Journal of Fluid Mechanics*, 641:115–127, 2009. doi: 10.1017/S0022112009992059.

[25] Peter J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010. doi: 10.1017/S0022112010001217.

[26] Peter J. Schmid. Dynamic mode decomposition and its variants. *Annual Review of Fluid Mechanics*, 54:225–254, 2022. doi: 10.1146/annurev-fluid-030121-015835.

[27] Matthew A. Wilson and Bruce L. McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679, 1994. doi: 10.1126/science.8036517.

Technical Appendices and Supplementary Material

This appendix provides supplementary material supporting the main paper, including methodological details, hyperparameter selection, full cross-subject results, additional Koopman mode and spectral analyses, Allen Neuropixels validation plots, and implementation details for baseline DMD variants.

Appendix A. Additional Methodological Details

The original LFP recordings contain approximately 22.3 million samples, corresponding to about 744 seconds at $f_s = 30{,}000$ Hz. To make very long-horizon reconstruction computationally tractable, we use up to $T = 300{,}000$ samples per subject in the main experiments. The window length $w = 3000$ provides sufficient temporal context, while the stride $\delta = 30$ creates strong overlap between neighboring windows and supports smooth reconstruction after overlap-add averaging.

The PCA step induces a low-rank representation of the windowed signal. Given $X = U\Sigma V^\top$, the rank-r approximation $X_r = U_r\Sigma_r V_r^\top$ retains dominant variance while reducing the effective state dimension from $d_x = wC$ to $r = 100$. This compression improves numerical conditioning and makes Koopman regression feasible for large multichannel neural recordings.

To characterize the learned dynamics, we analyze the Koopman spectrum of the latent operator. Let $Kv_k = \lambda_k v_k$ denote the eigenpairs of K, giving the latent expansion $z_t \approx \sum_{k=1}^{r} c_k \lambda_k^t v_k$. Eigenvalues near the unit circle indicate stable and persistent latent evolution, while the corresponding Koopman modes capture spatial–temporal structure across channels and window time.

Appendix B. Hyperparameter Selection and Experimental Settings

B.1. Hyperparameter Selection

We selected the final PCA-DMD configuration through an ablation over temporal window size, step size, and latent dimension. Window sizes $\{150, 200, 300, 500, 1000, 3000\}$, step sizes $\{10, 20, 30, 50, 60\}$, and latent dimensions $\{8, 20, 50, 100\}$ were evaluated. The final setting, $w = 3000$, $\delta = 30$, and $r = 100$, gave the best reconstruction performance while preserving numerical stability and was therefore used for all main long-horizon, cross-subject, and external validation experiments.

B.2. Experimental Settings

Table 1: 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	–

Appendix C. Full Cross-Subject Zero-Shot Results

In the cross-subject zero-shot setting, each model is trained on one source subject and directly applied to an unseen target subject without fine-tuning. This protocol tests whether the learned dynamics capture subject-invariant structure rather than subject-specific waveform statistics. PCA-DMD consistently achieves high correlation and low distributional mismatch across source–target pairs, whereas standard DMD variants degrade substantially under transfer.

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.

C.1. Metric Comparison Across DMD Variants

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.

Appendix D. Additional Koopman Mode Visualizations

Additional Koopman mode visualizations are provided for PCA-DMD and baseline DMD variants. Compared with classical DMD-based methods, PCA-DMD produces less redundant and more expressive mode structures, with both global patterns and higher-frequency variations across channels and time. This supports the interpretation that PCA-based dimensionality reduction improves both conditioning and expressiveness of the learned Koopman representation.

Appendix E. Koopman Spectral Analysis

We analyze the learned dynamics through the eigenvalue spectra and eigenvalue radius distributions of the Koopman operators. Across subjects, PCA-DMD eigenvalues are consistently organized near the unit circle, indicating stable and persistent latent evolution over long horizons. Baseline DMD variants show less coherent spectral organization, reflecting reduced reconstruction stability under long-horizon and cross-subject settings.

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.

Appendix F. Additional Baseline Reconstruction Results

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

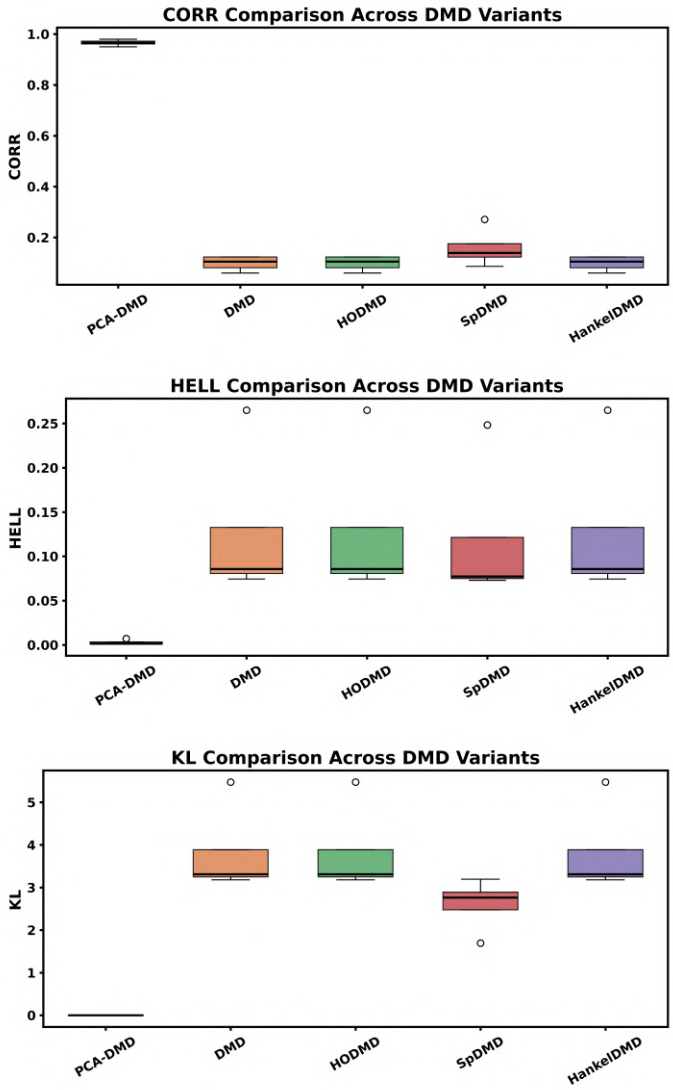

Figure 5: Comparison of reconstruction metrics across all methods in the cross-subject zero-shot setting: correlation (top), HD (middle), and KLD (bottom). PCA-DMD consistently outperforms the baseline DMD variants in both fidelity and robustness.

Appendix G. Allen Neuropixels External Validation

For external validation, PCA-DMD was applied to Allen Neuropixels LFP data with 93 jointly modeled channels and 300,000 samples. The signal was represented as $X \in \mathbb{R}^{93\times 300000}$, where rows correspond to channels and columns correspond to time samples. Each temporal window was flattened into a 279,000-dimensional vector, projected to a 100-dimensional PCA latent space, evolved using a learned Koopman operator, and reconstructed through inverse projection and overlap-add averaging. Importantly, channels were not modeled independently; all 93 channels were treated

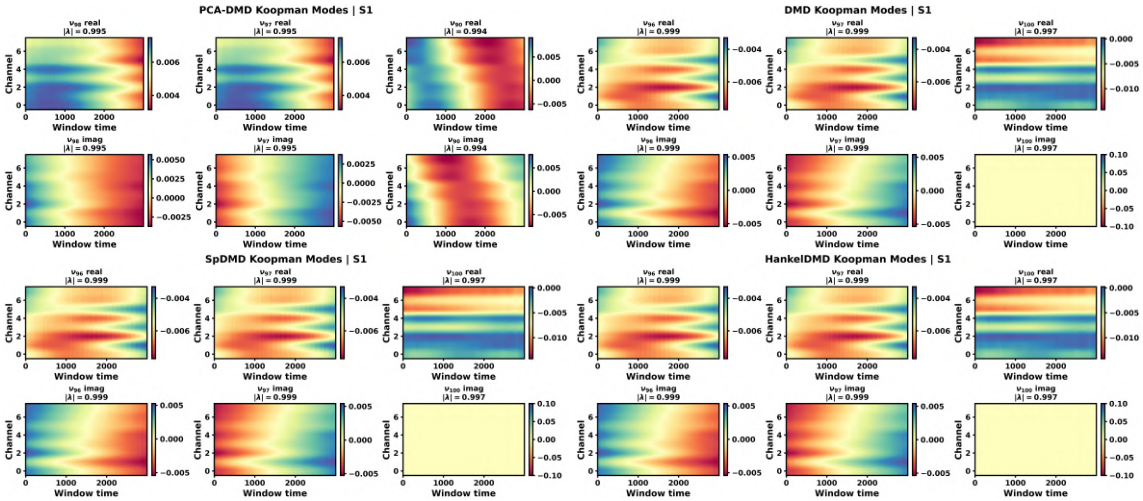

Figure 6: Koopman modes for DMD variants for Subject S1. The real and imaginary parts are shown separately.

jointly as a single high-dimensional multivariate dynamical system. This validation demonstrates that PCA-DMD remains computationally feasible beyond low-channel LFP recordings and can capture collective cross-channel temporal structure in higher-dimensional Neuropixels data.

G.1. Additional Allen Neuropixels Reconstruction Plots

Appendix H. Koopman Operator and DMD Variants

H.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

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

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

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

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)*:

$$_t = \sum_k \lambda_k^t \, \phi_{\lambda_k}(_0) \, \overset{\Phi}{_k}, \tag{2}$$

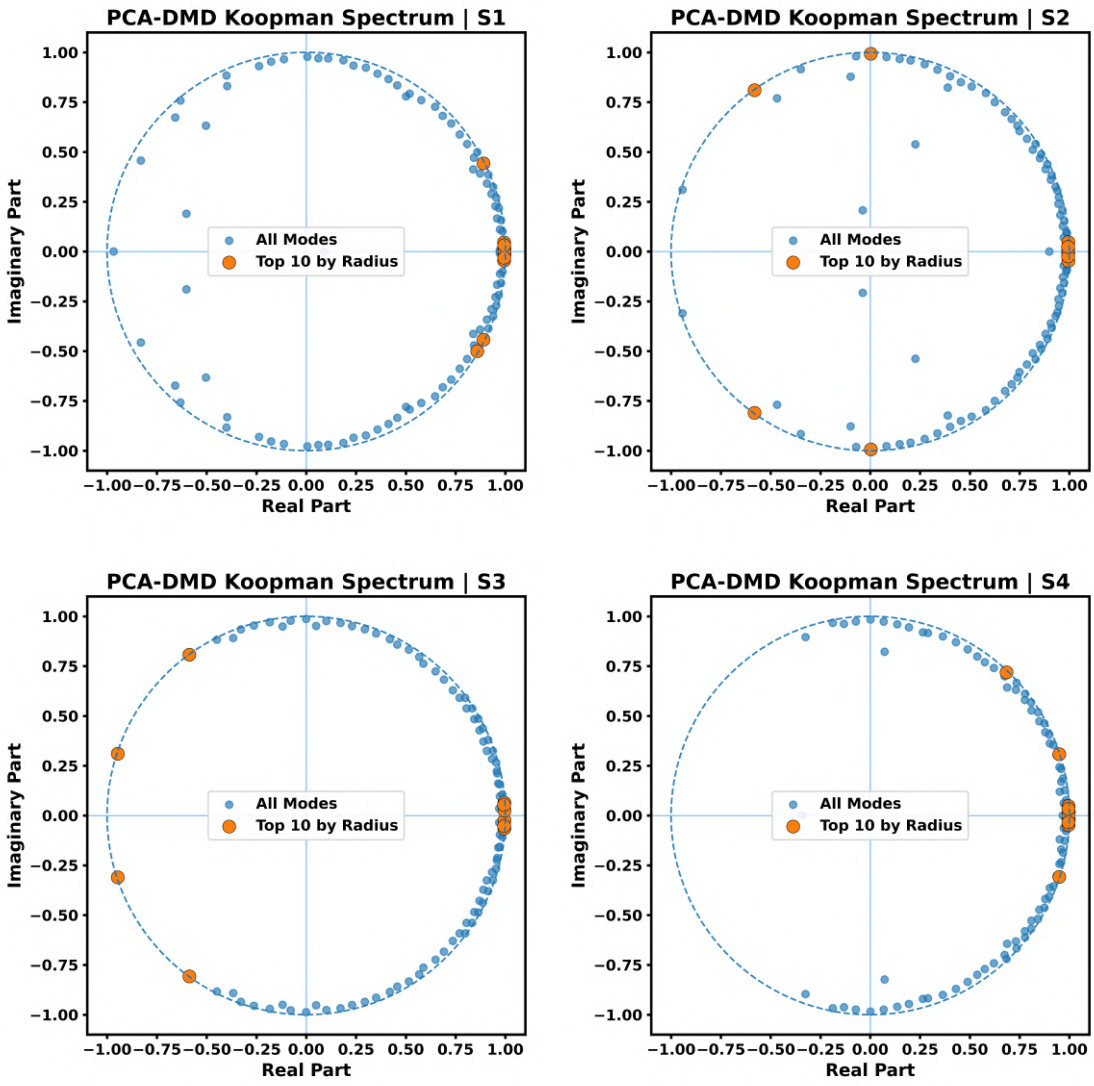

Figure 7: PCA-DMD Koopman eigenvalue spectra across four subjects. Dominant eigenvalues cluster near the unit circle, indicating stable latent dynamics.

where $\frac{\Phi}{k}$ 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).

H.2. Classical DMD

The Classical DMD implementation mirrors the windowing approach, utilizing PyDMD's DMD with an SVD rank of 100 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

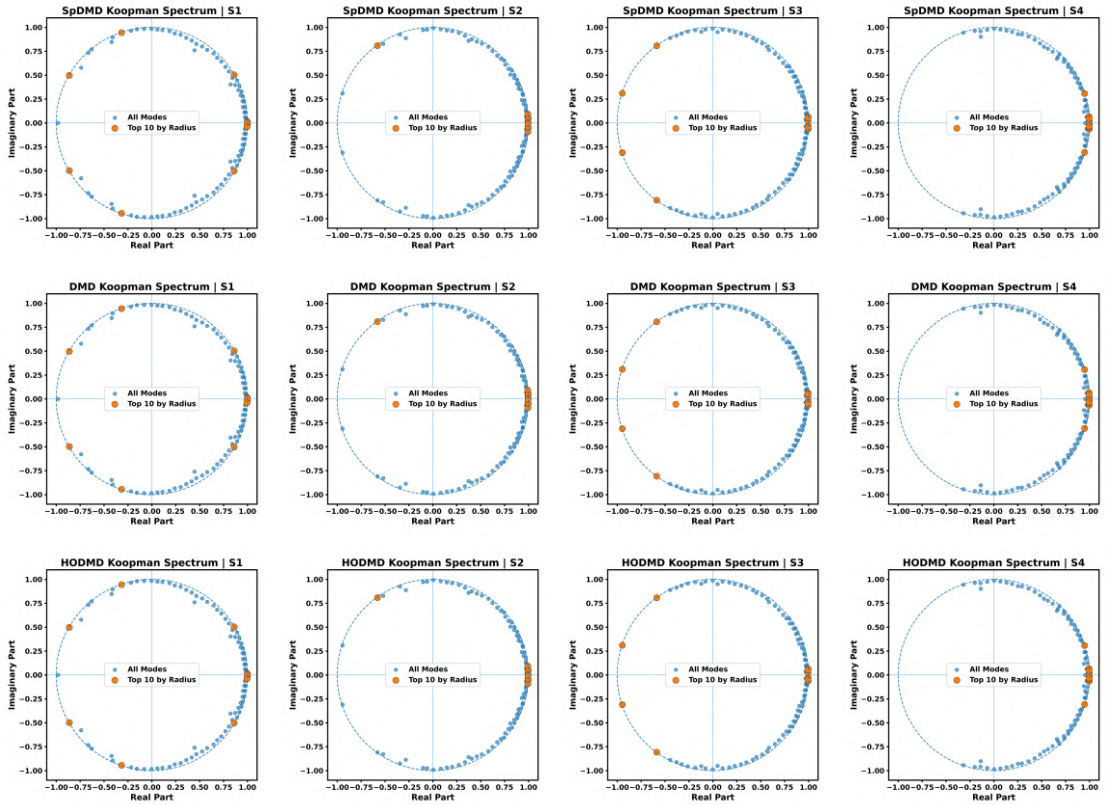

Figure 8: Additional Koopman eigenvalue spectra across for DMD variants.

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 $\mathbf{\Phi}$ and eigenvalues $\mathbf{\Lambda}$:

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

with $\mathbf{b}$ as the initial amplitude. The predicted snapshots are $\mathbf{X}_{\text{pred}} = \text{Re}(\mathbf{\Phi}\mathbf{\Lambda}\mathbf{b})$. 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 $t \in [\delta, \min(N, T - \delta)]$.

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 ↑	HD ↓	KLD ↓
S1→S2	0.9504	0.0020	0.0019
S1→S3	0.9671	0.0011	0.0006
S1→S4	0.9778	0.0013	0.0007
S2→S1	0.9623	0.0072	0.0020
S2→S3	0.9693	0.0011	0.0007
S2→S4	0.9798	0.0013	0.0005
S3→S1	0.9624	0.0071	0.0022
S3→S2	0.9553	0.0017	0.0010
S3→S4	0.9800	0.0012	0.0005
S4→S1	0.9643	0.0071	0.0020
S4→S2	0.9589	0.0016	0.0013
S4→S3	0.9707	0.0010	0.0007
DMD/HODMD/HankelDMD (avg.)	∼0.10	∼0.10–0.26	∼3–5
SpDMD (avg.)	∼0.15–0.27	∼0.07–0.25	∼1.7–3.2

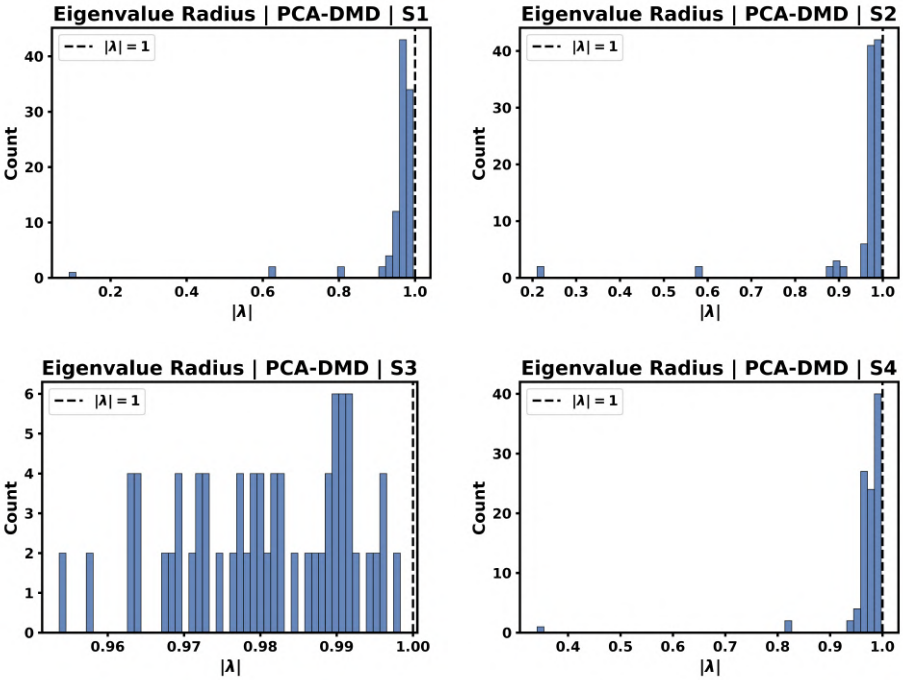

Figure 9: Eigenvalue radius distribution ($|\lambda|$) for PCA-DMD across all subjects.

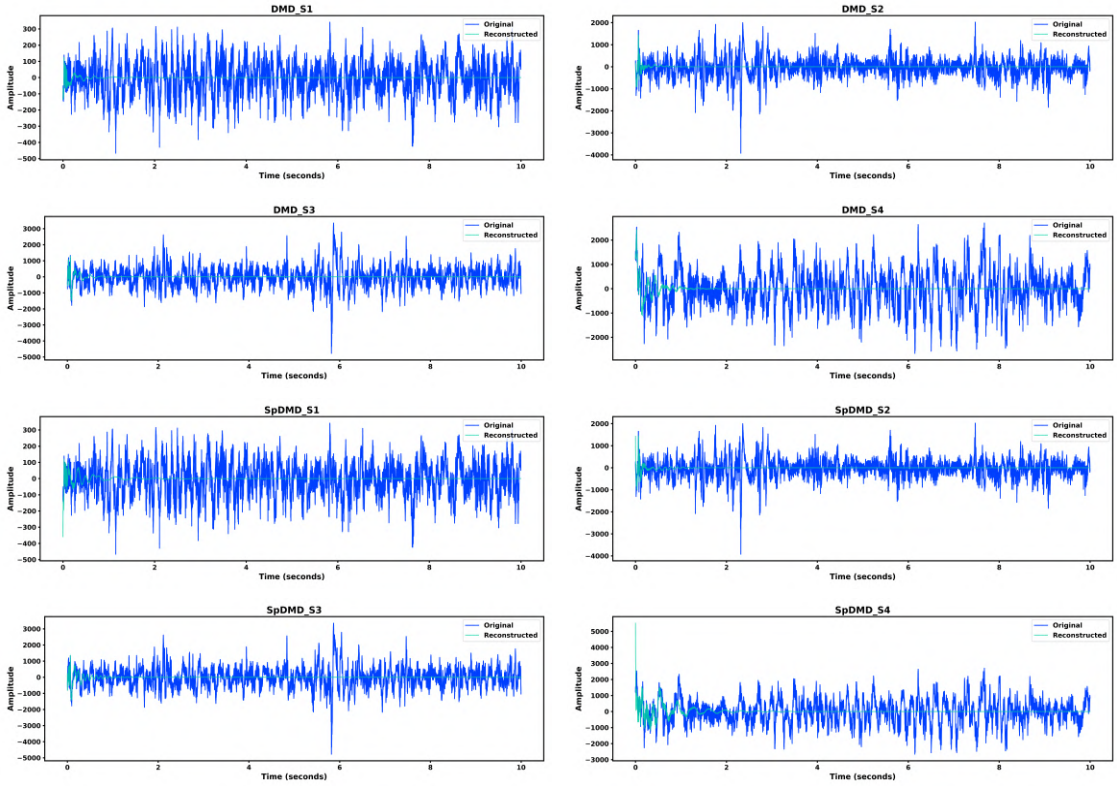

Figure 10: Additional cross-subject reconstruction examples for DMD variants.

H.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 100 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:

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

where $\mathbf{\Phi}$ contains SpDMD modes, and $\mathbf{\Lambda}$ is a diagonal matrix of eigenvalues. The predicted snapshots are obtained as $\mathbf{X}_{\text{pred}} = \text{Re}(\mathbf{\Phi}\mathbf{\Lambda}\mathbf{b})$, with $\mathbf{b}$ derived from the initial snapshot projection. Using

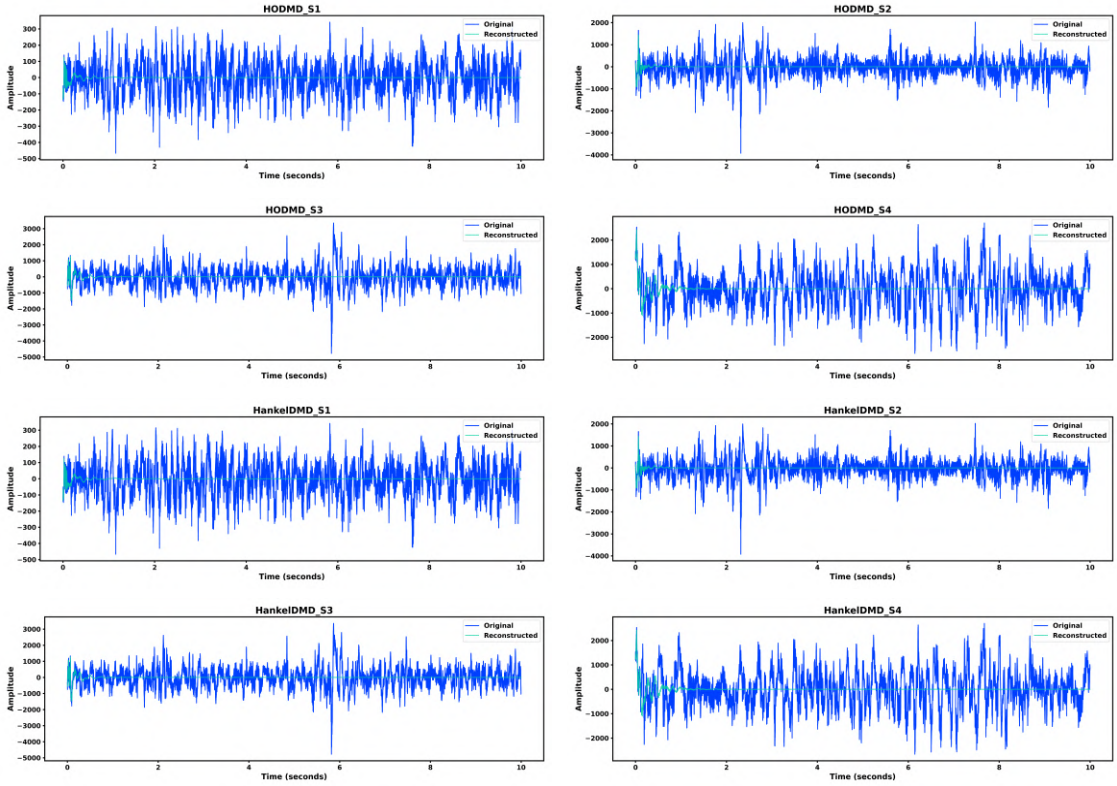

Figure 11: Additional cross-subject reconstruction examples for DMD variants. Each pair of panels corresponds to one method across subjects.

$\mathbf{X} = \mathbf{X}_{\text{full}}[: -1]$ 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 $t \in [\delta, \min(N, T - \delta)]$.

H.4. HODMD

HODMD for LFP reconstruction incorporates data scaling with StandardScaler before fitting PyDMD's HODMD with SVD rank 100 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 $\tilde{\mathbf{X}}_{\text{full}}^T \in \mathbb{R}^{w \times M}$,

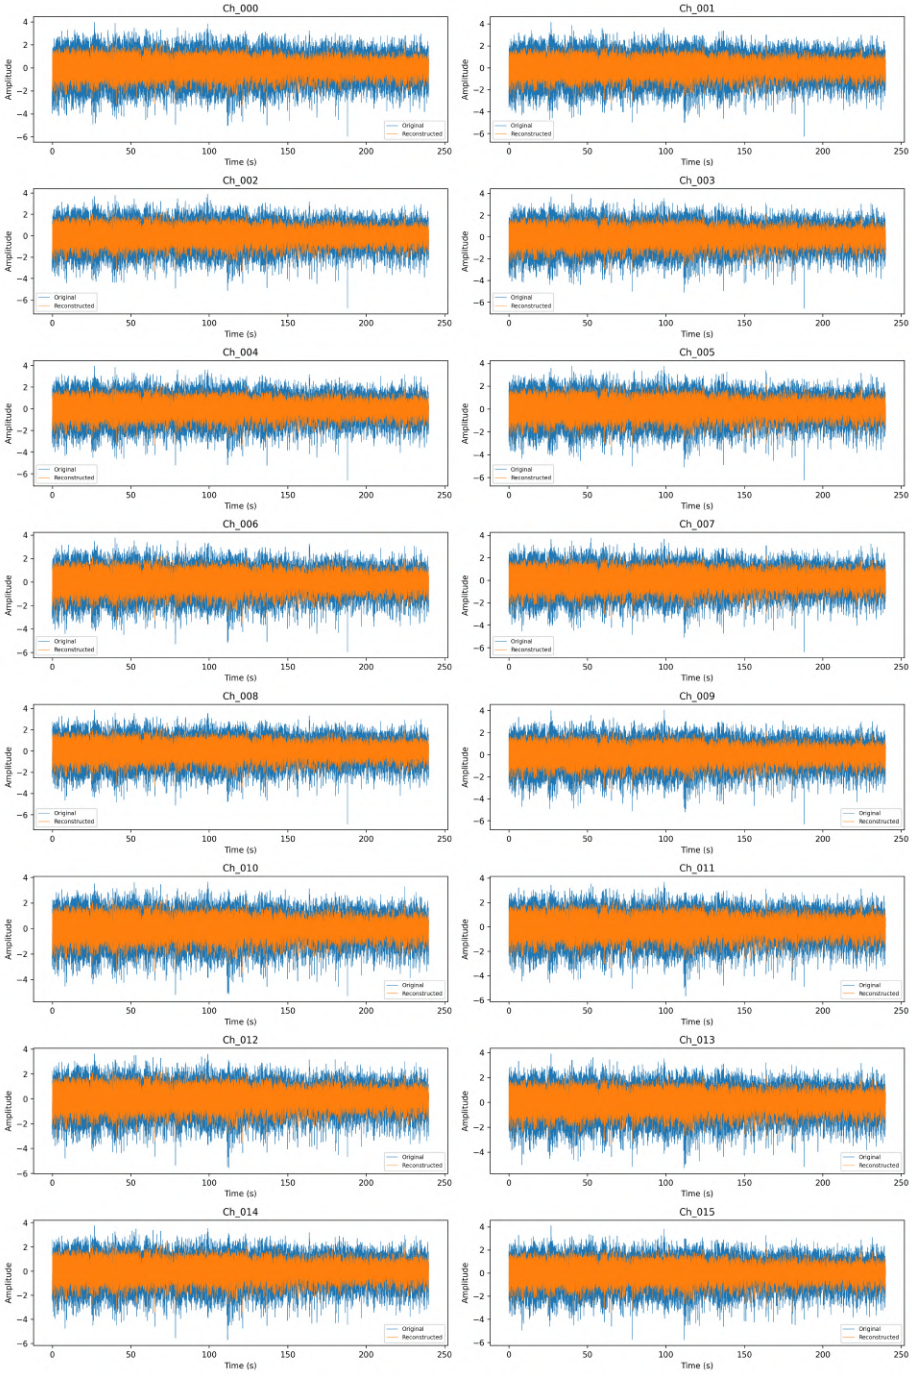

Figure 12: Original and reconstructed Allen Neuropixels LFP signals for channels 0-15.

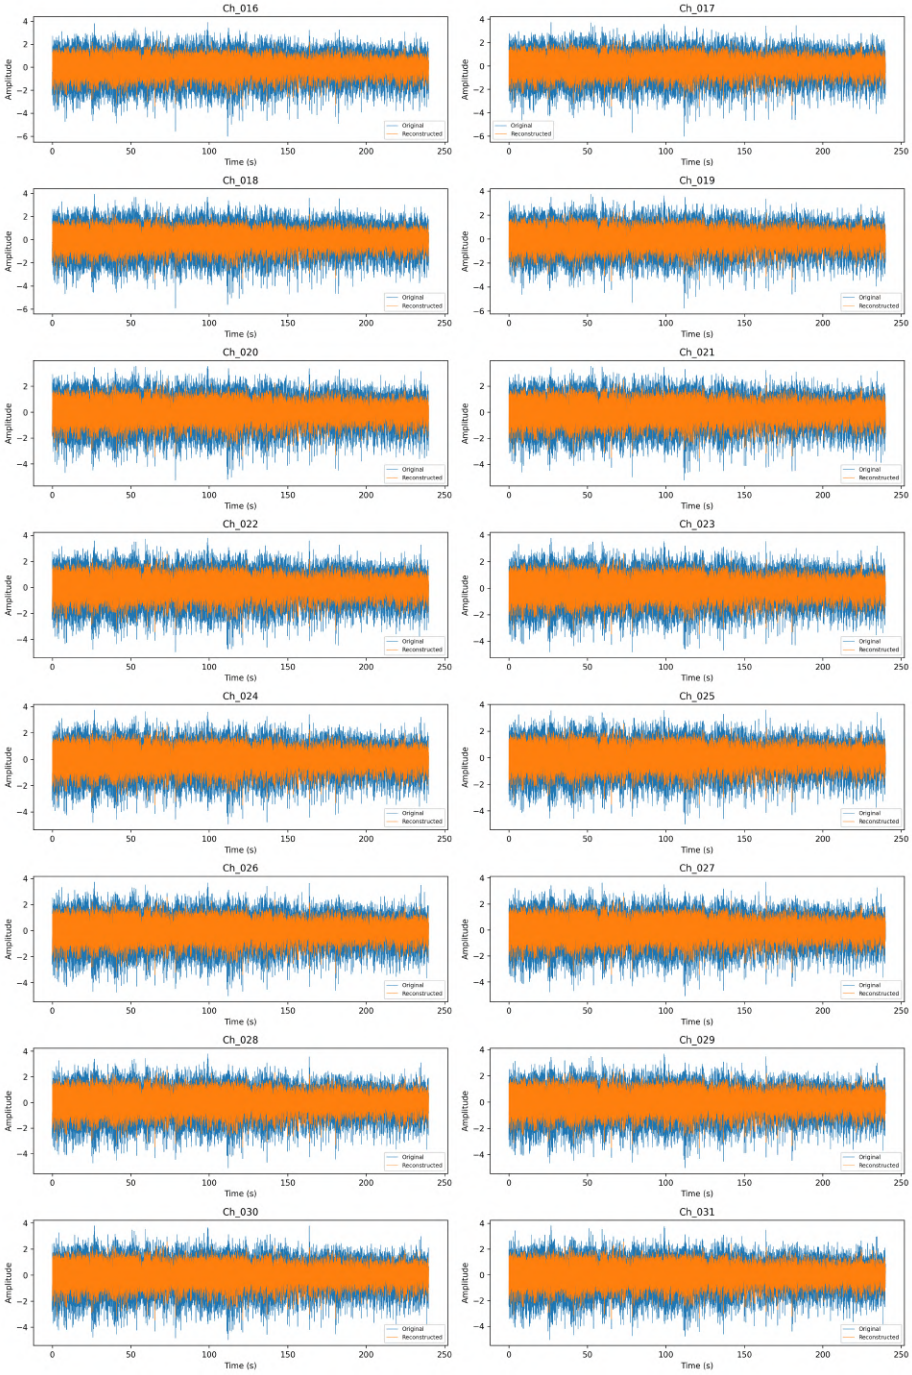

Figure 13: Original and reconstructed Allen Neuropixels LFP signals for channels 16-31.

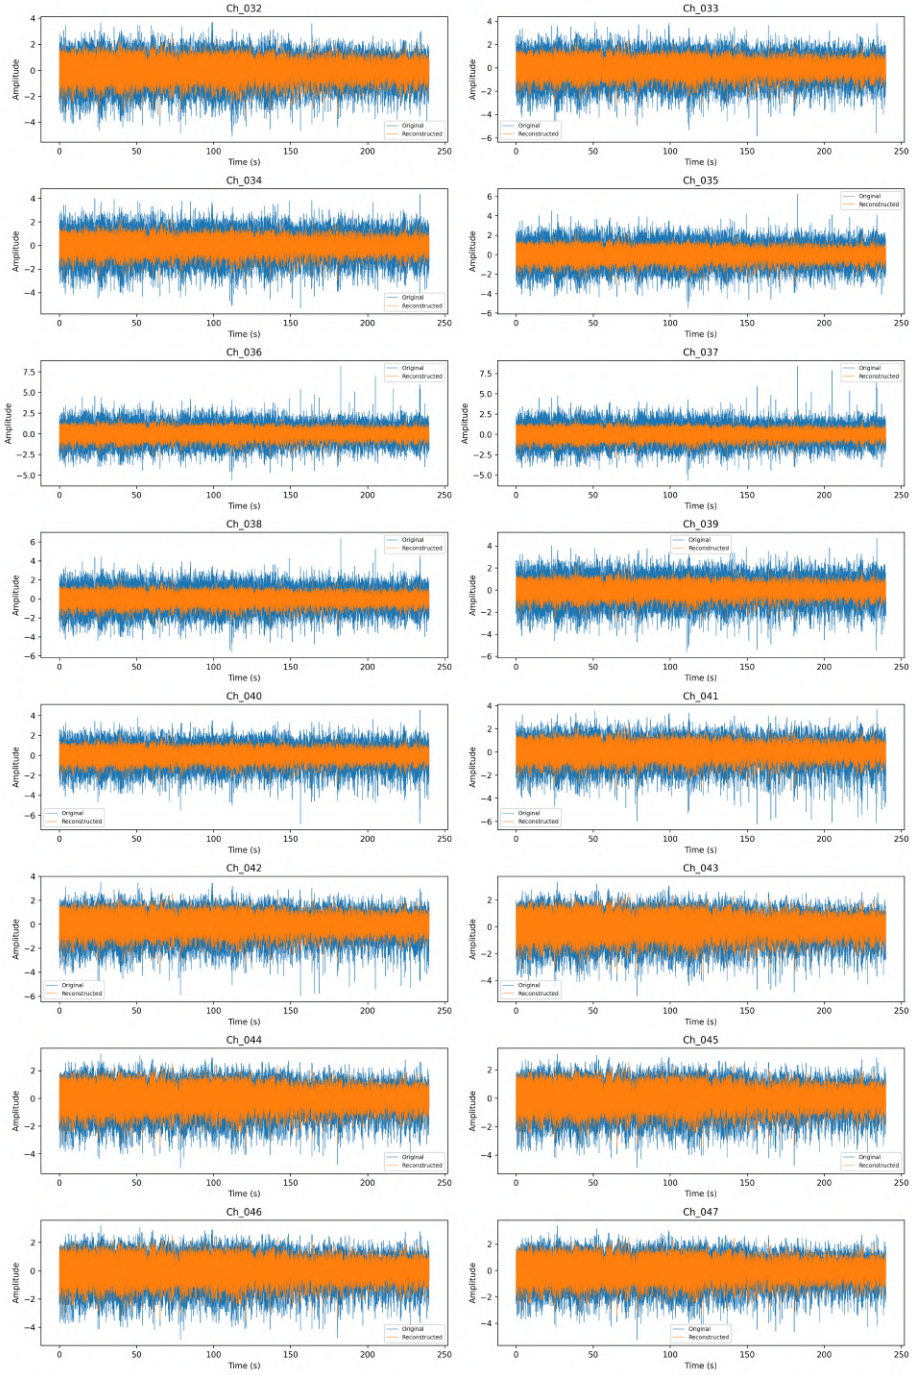

Figure 14: Original and reconstructed Allen Neuropixels LFP signals for channels 32-47.

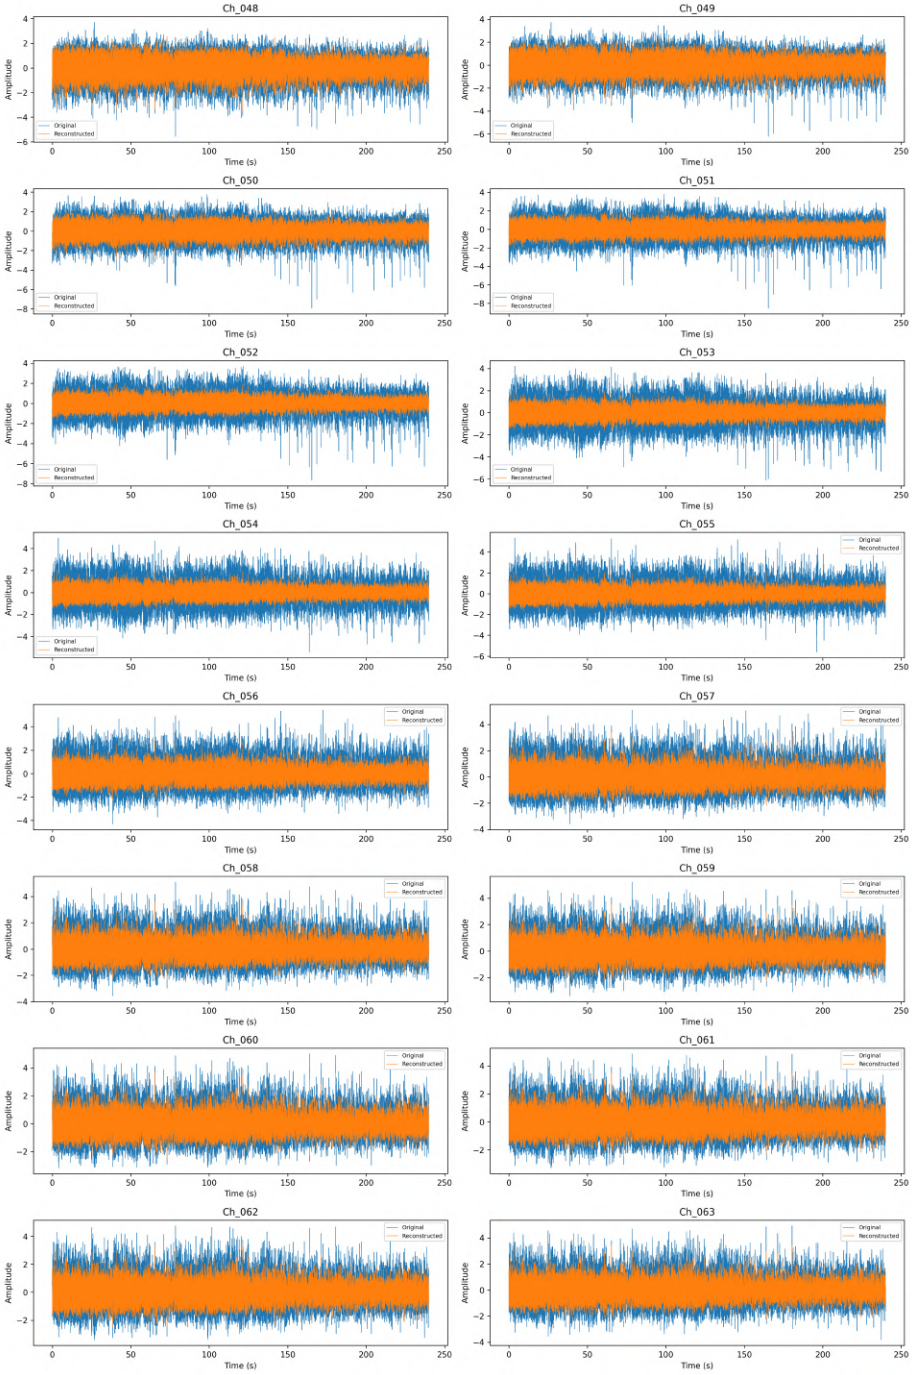

Figure 15: Original and reconstructed Allen Neuropixels LFP signals for channels 48-63.

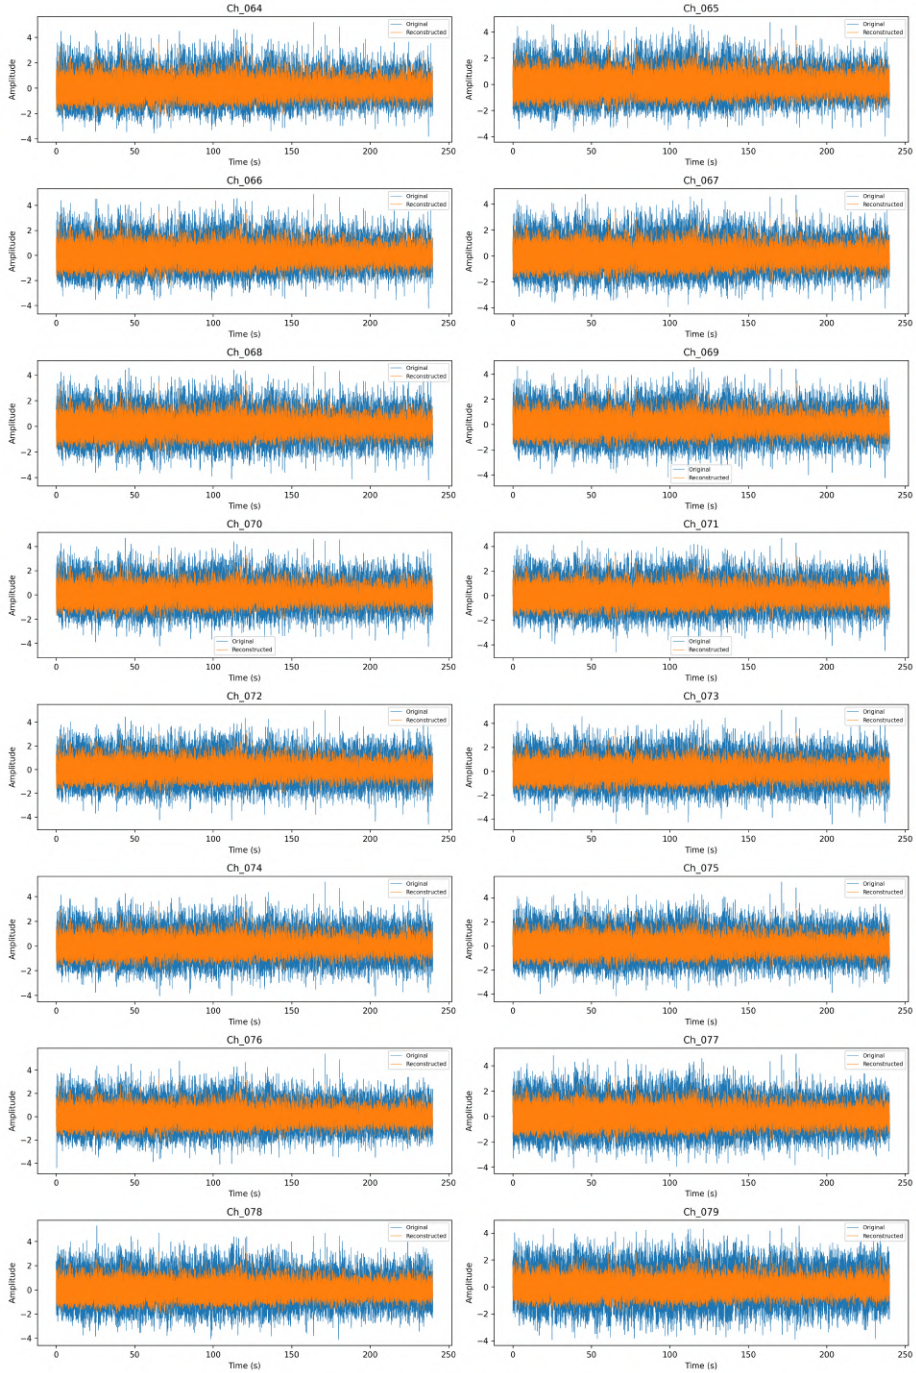

Figure 16: Original and reconstructed Allen Neuropixels LFP signals for channels 64-79.

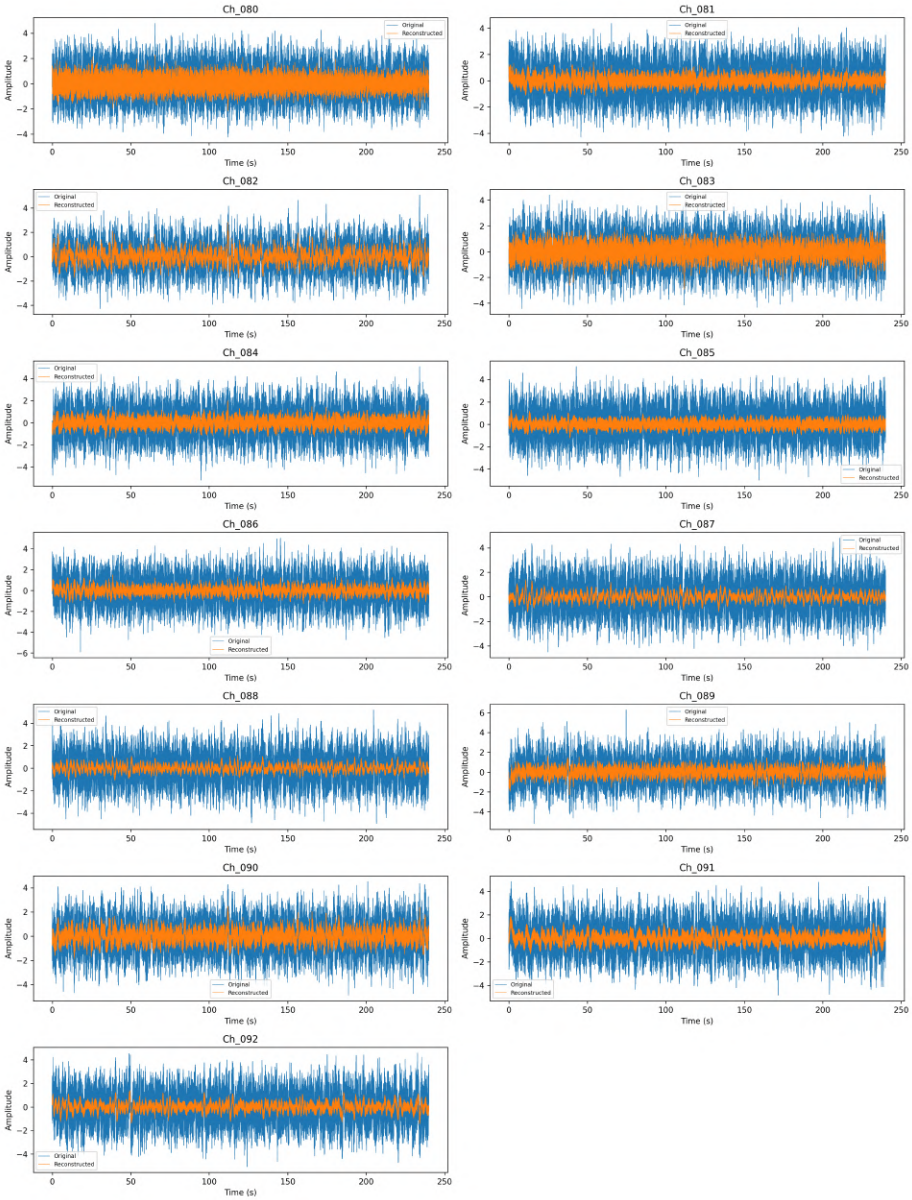

Figure 17: Original and reconstructed Allen Neuropixels LFP signals for channels 80-92.

constructing an augmented Hankel matrix and decomposing it as:

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

where $\mathbf{\Phi}$ and $\mathbf{\Lambda}$ 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 $\mathbf{X}_{\text{rec}}$. 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 $t \in [\delta, \min(N, T - \delta)]$.