

DeepGraphDMD: Interpretable Spatio-Temporal Decomposition of Non-linear Functional Brain Network Dynamics

Md Asadullah Turja¹, Martin Styner², and Guorong Wu³

¹ Department of Computer Science, University of North Carolina at Chapel Hill,
mturja@cs.unc.edu

² Department of Computer Science, University of North Carolina at Chapel Hill,
styner@email.unc.edu

³ Department of Psychiatry, University of North Carolina at Chapel Hill,
guorong_wu@med.unc.edu

Abstract. Functional brain dynamics is supported by parallel and overlapping functional network modes that are associated with specific neural circuits. Decomposing these network modes from fMRI data and finding their temporal characteristics is challenging due to their time-varying nature and the non-linearity of the functional dynamics. Dynamic Mode Decomposition (DMD) algorithms have been quite popular for solving this decomposition problem in recent years. In this work, we apply GraphDMD — an extension of the DMD for network data — to extract the dynamic network modes and their temporal characteristics from the fMRI time series in an interpretable manner. GraphDMD, however, regards the underlying system as a linear dynamical system that is sub-optimal for extracting the network modes from non-linear functional data. In this work, we develop a generalized version of the GraphDMD algorithm — DeepGraphDMD — applicable to arbitrary non-linear graph dynamical systems. DeepGraphDMD is an autoencoder-based deep learning model that learns Koopman eigenfunctions for graph data and embeds the non-linear graph dynamics into a latent linear space. We show the effectiveness of our method in both simulated data and the HCP resting-state fMRI data. In the HCP data, DeepGraphDMD provides novel insights into cognitive brain functions by discovering two major network modes related to fluid and crystallized intelligence.

Keywords: Component Analysis · Dynamic Mode Decomposition · Koopman Operator · Non-linear Graph Dynamical Systems · Functional Network Dynamics · State-space models · Graph Representation Learning.

1 Introduction

The human brain has evolved to support a set of complementary and temporally varying brain network organizations enabling parallel and higher-order information processing [16,20,24]. Decoupling these networks from a non-linear mixture

of signals (such as functional MRI) and extracting their temporal characteristics in an interpretable manner has been a long-standing challenge in the neuroscience community.

Conventional mode/component decomposition methods such as Principal Component Analysis (PCA) or Independent Component Analysis (ICA) assume the modes to be static [7,15,26] and thus sub-optimal for the functional networks generated by time-varying modes. Dynamic Mode Decomposition (DMD) can be treated as a dynamic extension of such component analysis methods since it allows its modes to oscillate over time with a fixed frequency [18]. This assumption is appropriate for the human brain as the functional brain organizations are supported by oscillatory network modes [2,8,11,23,27]. An extension of DMD for network data called GraphDMD [4] preserves the graph structure of the networks during the decomposition. In our work, we extend GraphDMD to a sequence of sliding window based dynamic functional connectivity (dNFC) networks to extract independent and oscillatory functional network modes.

Under the hood, GraphDMD regards the network sequence as a linear dynamical system (LDS) where a linear operator shifts the current network state one time-point in the future. The LDS assumption, however, is not optimal for modeling functional brain networks that exhibit complex non-linearity such as rapid synchronization and desynchronization as well as transient events [6]. Articles [3,13] propose switching linear dynamical system (SLDS) to tackle the nonlinearity of spatiotemporal data by a piecewise linear approximation. While these models offer interpretability, their shallow architecture limits their generalizability to arbitrary nonlinear systems. On the other hand, the methods in [5,21,10] model the non-linearity with a deep neural network. While these models have more representation capabilities compared to SLDS, the latent states are not interpretable. More importantly, all of these methods consider the node-level dynamics instead of the network dynamics.

Here, we propose a novel Deep Graph Dynamic Mode Decomposition (DeepGraphDMD) algorithm that applies to arbitrary non-linear network dynamics while maintaining interpretability in the latent space. Our method uses Koopman operator theory to lift a non-linear dynamical system into a linear space through a set of Koopman eigenfunctions (Figure 1a). There has been a growing line of work that learns these measurement functions using deep autoencoder architectures [14,22]. Training these autoencoders for network data, however, has two unique challenges – 1. preserving the edge identity in the latent space so that the network modes are interpretable, 2. enforcing linearity in the latent space for the high dimensional network data. In DeepGraphDMD, we tackle the first challenge by indirectly computing the network embeddings by a novel node embedding scheme. For the second challenge, we introduce a sparse Koopman operator to reduce the complexity of the learning problem. We evaluate the effectiveness of our novel method in both simulated data and resting-state fMRI (rs-fMRI) data from Human Connectome Project.

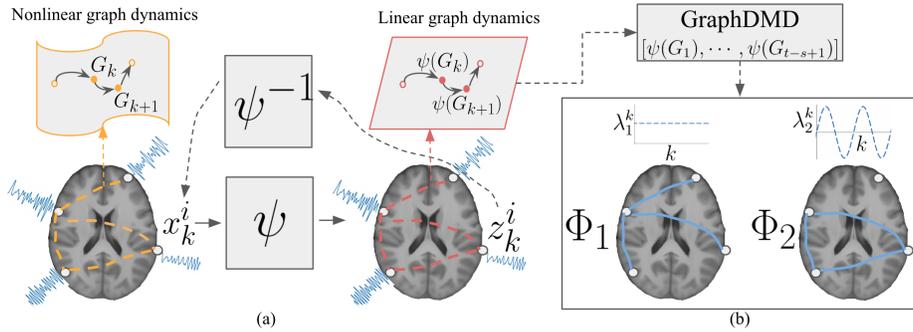


Fig. 1. (a) Illustration of the DeepGraphDMD model that embeds a nonlinear graph dynamical system into a linear space, and, (b) interpretable dynamic modes and their temporal characteristics after applying GraphDMD in the linear space.

2 Methodology

Let's assume $X \in \mathbb{R}^{n \times t}$ is a matrix containing the BOLD (blood-oxygen-level-dependent) signal of n brain regions (ROIs) in its rows at t time frames sampled at every $k\Delta t$ time points, where Δt is the temporal resolution. To compute the dynamic connectivity matrix at time point $k\Delta t$, a snapshot $X_k = X_{:,k:k+s}$ is taken in a sliding window of s time frames. A correlation matrix $G_k \in \mathbb{R}^{n \times n}$ is then computed from X_k by taking the pearson correlation between the rows of X_k , i.e., $G_k^{ij} = \text{pearson}(x_k^i, x_k^j)$ where x_k^i, x_k^j are the i^{th} and j^{th} row of X_k respectively. This yields a sequence of graphs $G = [G_1, G_2, \dots, G_{t-s+1}]$. Let's also assume that $g_k \in \mathcal{R}^{n^2}$ is a vectorized version of G_k , i.e. $g_k = \text{vec}(G_k)$ and $g \in \mathcal{R}^{n^2 \times (t-s+1)}$ is a matrix containing g_k in its columns. The goal is to decouple the overlapping spatiotemporal modes from the network sequence G using 1. Graph Dynamic Mode Decomposition algorithm, and 2. a novel Deep Learning-based Graph Dynamic Mode Decomposition algorithm.

2.1 Graph Dynamic Mode Decomposition

GraphDMD [4] assumes that g_k follows an LDS:

$$g_{k+1} = Ag_k \quad (1)$$

where $A \in \mathbb{R}^{n^2 \times n^2}$ is a linear operator that shifts the current state g_k to the state at the next time frame g_{k+1} . To extract the low dimensional global network dynamics, GraphDMD projects A into a lower dimensional space \hat{A} using tensor-train decomposition, applies eigendecomposition of \hat{A} , and projects the eigenvectors back to the original space which we refer to as dynamic modes (DMs). GraphDMD uses tensor-train decomposition to maintain the network structure of g_k and thus, the DMs from GraphDMD can be reshaped into $n \times n$ adjacency matrix forms. Let's assume these DMs are $\Phi_1, \Phi_2, \dots, \Phi_r$ where $\Phi_p \in \mathbb{C}^{n \times n}$ and

the corresponding eigenvalues are $\lambda_1, \lambda_2, \dots, \lambda_r$ where $\lambda_p \in \mathbb{C}$ (Figure 1b). Here, r is the total number of DMs. Φ_p corresponds to the coherent spatial mode and λ_p defines its temporal characteristics (growth/decay rate and frequencies). We can see this by unrolling equation 1 in time:

$$g_{k+1} = A^k g_1 = \sum_{p=1}^r \Phi_p \lambda_p^k b_p = \sum_{p=1}^r \Phi_p a_p^k \exp(\omega_p k \Delta t) b_p \quad (2)$$

where $\lambda_p = a_p \exp(\omega_p \Delta t)$, Φ^\dagger is the conjugate transpose of Φ , $b_p = \text{vec}(\Phi_p^\dagger) g_1$ is the projection of the initial value onto the DMD modes, $a_p = \|\lambda_p\|$ is the growth/decay rate and $\omega_p = \text{Im}(\ln \lambda_p) / \Delta t$ is the angular frequency of Φ_p .

2.2 Adaptation of Graph-DMD for Nonlinear Graph Dynamics

Since the dynamics of the functional networks are often non-linear, the linearity assumption of equation 1 is sub-optimal. In this regard, we resort to Koopman operator theory to transform the non-linear system into an LDS using a set of Koopman eigenfunctions ψ , i.e., $\psi(g_{k+1}) = A\psi(g_k)$ [9]. We learn ψ using a deep autoencoder-based architecture — DeepGraphDMD — where the encoder and the decoder are trained to approximate ψ and ψ^{-1} , respectively. We enforce $\psi(g_k)$ to follow an LDS by applying Latent Koopman Invariant Loss [22] in the form:

$$\mathcal{L}_{lkis} = \|Y' - (Y'Y^\dagger)^{-1}Y\|_F^2 \quad (3)$$

where $Y = \begin{pmatrix} \psi(g_1) & \psi(g_2) & \dots & \psi(g_{t-s}) \end{pmatrix}$, $Y' = \begin{pmatrix} \psi(g_2) & \psi(g_3) & \dots & \psi(g_{t-s+1}) \end{pmatrix}$ are

two matrices with columns stacked with $\psi(g_k)$ and Y^\dagger is the right inverse of Y . After training, we reshape $\psi(g_k)$ into a $n \times n$ network $\psi(G_k)$ and generate the latent network sequence $\psi(G_1), \dots, \psi(G_{t-s+1})$. We then apply GraphDMD (described in section 2.1) on this latent and linearized network sequence to extract the DMs Φ_p and their corresponding λ_p .

However, there are two unique challenges of learning network embeddings using the DeepGraphDMD model: 1. the edge identity and, thus, the interpretability will be lost in the latent space if we directly embed g_k using ψ , and 2. Y^\dagger doesn't exist, and thus \mathcal{L}_{lkis} can't be computed because Y is low rank with the number of rows $\frac{n(n-1)}{2} \gg$ the number of columns $t - s + 1$.

To solve the first problem, instead of learning $\psi(g_k)$ directly, we embed the BOLD signal x_k^i of each ROI independently using the encoder to learn the latent embeddings z_k^i (Figure 1a). We then compute the pearson correlation between the latent embeddings of the ROIs to get the Koopman eigenfunctions of g_k i.e., $\psi(g_k^{ij}) = \text{pearson}(z_k^i, z_k^j)$. The weights of the encoder and decoder are shared across the ROIs.

The second problem arises because the Koopman operator A regresses the value of an edge at the next time-point as a linear combination of all the other edges at the current time-point, i.e., $g_{k+1}^{ij} = \sum_{p,q=1}^N w_{pq} g_k^{pq}$. This results in $\mathcal{O}(n^2)$

covariates with $t - s + 1 \ll \mathcal{O}(n^2)$ samples making the regression ill-posed. We propose a sparse Koopman operator where each edge g_k^{ij} is regressed using only the edges that share a common end-point with it, i.e., $g_{k+1}^{ij} = \sum_{p=1, p \neq i, j}^n w_{ip} g_k^{ip} + \sum_{q=1, q \neq i, j}^n w_{qj} g_k^{qj} + w_{ij} g_k^{ij}$ (Supplementary Figure 1). Since there are only $\mathcal{O}(n)$ such edges, it solves the ill-posedness of the regression.

Other than \mathcal{L}_{lkis} , we also train the autoencoder with a reconstruction loss \mathcal{L}_{recon} which is the mean-squared error (MSE) between x_k^i and the reconstructed output from the decoder \hat{x}_k^i . Moreover, a regularizer \mathcal{L}_{reg} in the form of an MSE loss between g_k and the latent $\psi(g_k)$ is also added. The final loss is the following:

$$\mathcal{L} = \mathcal{L}_{recon} + \alpha \mathcal{L}_{lkis} + \beta \mathcal{L}_{reg} \quad (4)$$

where α and β are hyper-parameters. We choose α , β , and other hyperparameters using grid search on the validation set. The network architecture and the values of the hyper-parameters of DeepGraphDMD training are shown in Supplementary Figure 1. The code is available in ⁴.

2.3 Window-based GraphDMD

We apply GraphDMD in a short window of size 64 time frames with a step size of 4 time frames instead of the whole sequence G because, in real-world fMRI data, both the frequency and the structure of the DMs can change over time. We then combine the DMs across different sliding windows using the following post-processing steps:

Post-processing of the DMs: We first group the DMs within the frequency bins: 0-0.01 Hz, 0.01-0.04Hz, 0.04-0.08Hz, 0.08-0.12Hz, and 0.12-0.16Hz. We then cluster the DMs within each frequency bin using a clustering algorithm and select the cluster centroids as the representative DMs (except for the first bin where we average the DMs). We chose the optimal clustering algorithm to be Spherical KMeans [1] (among Gaussian Mixture Model, KMeans, Spherical KMeans, DBSCAN, and, KMedoids) and the optimal number of clusters to be 3 for every frequency bin based on silhouette analysis [17] (Supplementary Figure 2). We use this frequency binning technique to allow for slight variations of ω of a DM over the scanning session. To align these representative DMs across subjects, we apply another round of spherical clustering on the DMs from all subjects and align them based on their cluster memberships.

3 Experiments

3.1 Dataset

We use rs-fMRI for 840 subjects from the HCP Dense Connectome dataset ⁵ [25]. Each fMRI image was acquired with a temporal resolution (Δt) of 0.72 s and a

⁴ <https://github.com/mturja-vf-ic-bd/DeepGraphDMD.git>

⁵ <https://www.humanconnectome.org/storage/app/media/documentation/s1200/HCP1200-DenseConnectome+PTN+Appendix-July2017.pdf>

2 mm isotropic spatial resolution using a 3T Siemens Skyra scanner. Individual subjects underwent four rs-fMRI runs of 14.4 min each (1200 frames per run). Group-ICA using FSL’s MELODIC tool [7] was applied to parcellate the brain into 50 functional regions (ROIs). To find the correlation between cognition with the rs-fMRI data, we select two behavioral measures related to fluid intelligence: `CogFluidComp`, `PMAT24_A.CR` and one measure related to crystallized intelligence: `ReadEng`, and, the normalized scores of the fluid and crystallized cognition measures: `CogTotalComp`. We regress out the confounding factors: age, gender, and head motion from these behavioral measures using ordinary least squares [12].

3.2 Baseline Methods

We compare GraphDMD and DeepGraphDMD against three decomposition methods: Principal Component Analysis (PCA), Independent Component Analysis (ICA), and standard Dynamic Mode Decomposition (DMD) [18]. We use the `sklearn` decomposition library for PCA ⁶ and ICA ⁷ and the `pyDMD`⁸ library for standard DMD. We apply PCA and ICA on g , and DMD directly on the bold signal X instead of g (for reasons described in section 4.2). We choose the number of components (`n_components`) to be three for these decomposition methods, (Results for other `n_components` values are shown in Supplementary Table 1). The components are aligned across subjects using spherical clustering similar to the GraphDMD modes (Section 2.3). We also compare with static functional connectivity (sFC), which is the pairwise pearson correlation between brain regions across all time frames.

3.3 Simulation Study

We generate a sequence of dynamic adjacency matrices G using equation 2 from three time-varying modes Φ_1, Φ_2, Φ_3 with corresponding frequencies $\omega_1 \sim \mathcal{N}(0.1, 0.05)$, $\omega_2 \sim \mathcal{N}(1, 0.1)$, $\omega_3 \sim \mathcal{N}(2.5, 0.1)$ (Hz). Each Φ_p is a 32×32 block diagonal matrices with block sizes 16, 8, and 4. We choose $a_1 = 1.01$, $a_2 = 0.9$, $a_3 = 1.05$ and $b_1 = b_2 = b_3 = 1$. We simulate the process for $k = 1, \dots, 29$ time-points yielding a sequence of 30 matrices of shape 32×32 . We repeat the process ten times with different $\omega_1, \omega_2, \omega_3$ and generate ten matrix sequences. We apply PCA, ICA, and GraphDMD (section 2.1) on G to extract three components and compare them against the ground truth modes using pearson correlation.

3.4 Application of GraphDMD and DeepGraphDMD in HCP data

Comparison of DMs with sFC: The ground truth DMs are unknown for the HCP dataset; however, we can use the sFC as a substitute for the ground truth DM with $\omega = 0$ (static DM). sFC offsets the DMs with $\omega > 0$ as they have both

⁶ `sklearn.decomposition.PCA`

⁷ `sklearn.decomposition.FastICA`

⁸ <https://mathlab.github.io/PyDMD/dmd.html>

positive and negative cycles, and thus only retain the static DM. For comparison, we compute the pearson correlation between the DM within the frequency bin 0-0.01 Hz and sFC for both GraphDMD and DeepGraphDMD. For PCA and ICA, we take the maximum value of the correlation between sFC and the (PCA or ICA) components.

Regression Analysis of Behavioral Measures from HCP: In this experiment, we regress the behavioral measures with the DMs within each frequency bin (section 2.3) using Elastic-net. As an input to the Elastic-net, we take the real part of the upper diagonal part of the DM and flatten it into a vector. We then train the Elastic-net in two ways — 1. single-band: where we train the Elastic-net independently with the DMs in each frequency bin, and 2. multi-band: we concatenate two DMs in the frequency bins: 0-0.01Hz and 0.08-0.12Hz and regress using the concatenated vector. For evaluation, we compute the correlation coefficient r between the predicted and the true values of the measures.

4 Results

4.1 Simulation Study

In Figure 2a, we show the results after applying PCA, ICA, and, GraphDMD on the simulated data described in section 3.3. Since the DMs in this data are oscillating, the data generated from this process are more likely to be overlapping compared to when the modes are static. As a result, methods that assume static modes, such as PCA and ICA, struggle to decouple the DMs and discover modes in overlapping high-density regions. For example, in Mode 2 of ICA, we can see the remnants of Mode 3 in the blue boxes and Mode 1 (negative cycle) in the orange boxes. We observe similar scenarios in the Mode 3 of ICA and Mode 1, Mode 2, and Mode 3 of PCA (the red boxes). On the other hand, the DMs from GraphDMD have fewer remnants from other modes and closely resemble the ground truth. To empirically compare, the mean (\pm std) pearson correlation for PCA, ICA, and, GraphDMD are $0.81(\pm 0.04)$, $0.88(\pm 0.03)$, and $0.98(\pm 0.01)$.

4.2 Application of GraphDMD and DeepGraphDMD in HCP data

Comparison of DMs with sFC: The average pearson correlations with sFC across all the subjects are $0.6(\pm 0.09)$, $0.6(\pm 0.09)$, $0.84(\pm 0.09)$, and $0.86(\pm 0.05)$ for PCA, ICA, GraphDMD, and, DeepGraphDMD (Figure 2b) respectively. This shows that the DMD-based methods can robustly decouple the static DM from time-varying DMs. In comparison, the corresponding PCA and ICA component has significantly lower correlation due to the overlap from the higher frequency components.

Regression Analysis of Behavioral Measures from HCP: We show the values of r across different methods in Table 1. We only show the results

for two frequency bins 0-0.01Hz and 0.08-0.12 Hz, as the DMs in the other bins are not significantly correlated ($r < 0.2$) with the behavioral measures (Supplementary Table 2). The table shows that multi-band training with the DMs from the DMD-based methods significantly improves the regression performance over the baseline methods. Compared to sFC, GraphDMD improves r by 22%, 6%, 0.7%, and, 3% for *CogFluidComp*, *PMAT24_A_CR*, *ReadEng*, *CogTotalComp*, respectively and DeepGraphDMD further improves the performance by 5%, 2.2%, 0.7%, and, 1.5%, respectively. Significant performance improvement for *CogFluidComp* can be explained by the DM within the bin 0.08-0.12 Hz. This DM provides additional information related to fluid intelligence ($r = 0.227$ for GraphDMD) to which the sFC doesn't have access. By considering non-linearity, DeepGraphDMD extracts more robust and less noisy DMs (Figure 2b-c), and hence, it improves the regression performance by 8% compared to GraphDMD in this frequency bin. By contrast, the standard DMD algorithm yields unstable modes with $a_p \ll 1$ when applied to the network sequence G . These modes have no correspondence across subjects and thus can't be used for regression. We instead apply DMD on the BOLD signal X , but the DMD modes show little correlation with the behavioral measures. PCA and ICA perform significantly worse than the baseline sFC method for all behavioral measures.

Traditional dynamical functional connectivity analysis methods (such as sliding window-based techniques) consider a sequence of network states. However, our results show that these states can be further decomposed into more atomic network modes. The importance of decoupling these network modes from non-linearly mixed fMRI signals using DeepGraphDMD has been shown in regressing behavioral measures from HCP data.

5 Conclusion

In this paper, we proposed a novel algorithm — DeepGraphDMD — to decouple spatiotemporal network modes in dynamic functional brain networks. Unlike

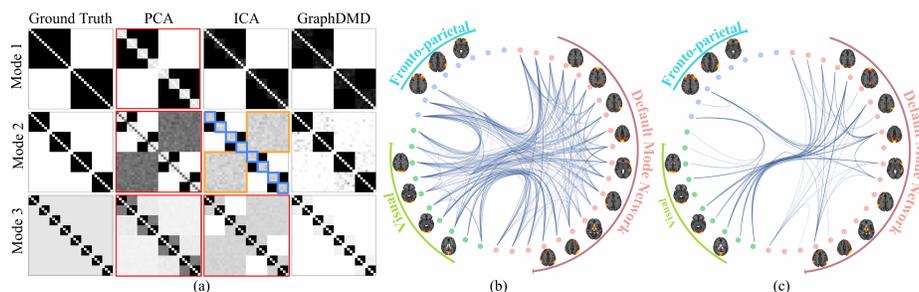


Fig. 2. (a) Ground truth network modes from simulated data (column 1) and extracted network modes from PCA (2nd column), ICA (3rd column), and, GraphDMD (4th column), (b) Circle plot of the average DMs with $\omega \approx 0$, (c) $\omega \in [0.08 - 0.12]$ from DeepGraphDMD organized based on common resting-state networks [19].

	Frequency (Hz)	CogFluidComp	PMAT24_A_CR	ReadEng	CogTotalComp
sFC	N/A	0.253 ± 0.003	0.294 ± 0.004	0.407 ± 0.004	0.440 ± 0.004
PCA	N/A	0.109 ± 0.003	0.126 ± 0.003	0.224 ± 0.003	0.238 ± 0.003
ICA	N/A	0.148 ± 0.005	0.158 ± 0.004	0.239 ± 0.005	0.266 ± 0.006
DMD	0-0.01, 0.08-0.12	0.064 ± 0.002	0.169 ± 0.002	0.132 ± 0.003	0.138 ± 0.006
Graph DMD	0-0.01	0.254 ± 0.003	0.289 ± 0.004	0.402 ± 0.004	0.438 ± 0.003
	0.08-0.12	0.227 ± 0.004	0.193 ± 0.004	0.145 ± 0.004	0.248 ± 0.004
	0-0.01, 0.08-0.12	0.308 ± 0.004	0.312 ± 0.004	0.410 ± 0.003	0.454 ± 0.004
Deep	0-0.01	0.259 ± 0.003	0.290 ± 0.002	0.404 ± 0.002	0.439 ± 0.002
Graph	0.08-0.12	0.245 ± 0.002	0.201 ± 0.004	0.144 ± 0.003	0.251 ± 0.004
DMD	0-0.01, 0.08-0.12	0.325 ± 0.003	0.319 ± 0.003	0.413 ± 0.002	0.461 ± 0.003

Table 1. Comparison of r for the behavioral measures across different methods.

other decomposition methods, DeepGraphDMD accounts for both the non-linear and the time-varying nature of the functional modes. As a result, these functional modes from DeepGraphDMD are more robust compared to their linear counterpart in GraphDMD and are shown to be correlated with fluid and crystallized intelligence measures.

References

- Banerjee, A., Dhillon, I.S., Ghosh, J., Sra, S.: Clustering on the unit hypersphere using von mises-fisher distributions. *Journal of Machine Learning Research* **6**(46), 1345–1382 (2005), <http://jmlr.org/papers/v6/banerjee05a.html>
- Casorso, J., Kong, X., Chi, W., Van De Ville, D., Yeo, B.T., Liégeois, R.: Dynamic mode decomposition of resting-state and task fmri. *NeuroImage* **194**, 42–54 (2019). <https://doi.org/10.1016/j.neuroimage.2019.03.019>, <https://www.sciencedirect.com/science/article/pii/S1053811919301922>
- Fox, E., Sudderth, E., Jordan, M., Willsky, A.: Nonparametric bayesian learning of switching linear dynamical systems. In: Koller, D., Schuurmans, D., Bengio, Y., Bottou, L. (eds.) *Advances in Neural Information Processing Systems*. vol. 21. Curran Associates, Inc. (2008), https://proceedings.neurips.cc/paper_files/paper/2008/file/950a4152c2b4aa3ad78bdd6b366cc179-Paper.pdf
- Fujii, K., Takeishi, N., Hojo, M., Inaba, Y., Kawahara, Y.: Physically-interpretable classification of biological network dynamics for complex collective motions. *Scientific Reports* **10** (02 2020). <https://doi.org/10.1038/s41598-020-58064-w>
- Gao, Y., Archer, E.W., Paninski, L., Cunningham, J.P.: Linear dynamical neural population models through nonlinear embeddings. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016), https://proceedings.neurips.cc/paper_files/paper/2016/file/76dc611d6ebaafc66cc0879c71b5db5c-Paper.pdf
- He, B.J.: Robust, Transient Neural Dynamics during Conscious Perception. *Trends in cognitive sciences* **22**(7), 563–565 (7 2018). <https://doi.org/10.1016/J.TICS.2018.04.005>, <https://pubmed.ncbi.nlm.nih.gov/29764721/>
- Hyvarinen, A.: Fast and robust fixed-point algorithms for independent component analysis. *IEEE transactions on Neural Networks* **10**(3), 626–634 (1999)

8. Ikeda, S., Kawano, K., Watanabe, S., Yamashita, O., Kawahara, Y.: Predicting behavior through dynamic modes in resting-state fmri data. *NeuroImage* **247**, 118801 (2022). <https://doi.org/https://doi.org/10.1016/j.neuroimage.2021.118801>, <https://www.sciencedirect.com/science/article/pii/S1053811921010727>
9. Koopman, B.O.: Hamiltonian Systems and Transformation in Hilbert Space. *Proceedings of the National Academy of Sciences of the United States of America* **17**(5), 315 (5 1931). <https://doi.org/10.1073/PNAS.17.5.315>, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1076052/>
10. Krishnan, R.G., Shalit, U., Sontag, D.A.: Structured inference networks for non-linear state space models. In: *AAAI Conference on Artificial Intelligence* (2016)
11. Kunert-Graf, J.M., Eschenburg, K.M., Galas, D.J., Kutz, J.N., Rane, S.D., Brunton, B.W.: Extracting reproducible time-resolved resting state networks using dynamic mode decomposition. *Frontiers in Computational Neuroscience* **13** (2019). <https://doi.org/10.3389/fncom.2019.00075>, <https://www.frontiersin.org/articles/10.3389/fncom.2019.00075>
12. Liégeois, R., Li, J., Kong, R., Orban, C., Van De Ville, D., Ge, T., Sabuncu, M.R., Yeo, B.T.: Resting brain dynamics at different timescales capture distinct aspects of human behavior. *Nature communications* **10**(1), 2317 (2019)
13. Linderman, S., Johnson, M., Miller, A., Adams, R., Blei, D., Paninski, L.: Bayesian Learning and Inference in Recurrent Switching Linear Dynamical Systems. In: Singh, A., Zhu, J. (eds.) *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research*, vol. 54, pp. 914–922. PMLR (20–22 Apr 2017), <https://proceedings.mlr.press/v54/linderman17a.html>
14. Lusch, B., Nathan Kutz, J., Brunton, S.L.: Deep learning for universal linear embeddings of nonlinear dynamics
15. McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J.: Analysis of fmri data by blind separation into independent spatial components. *Human brain mapping* **6**(3), 160–188 (1998)
16. Osada, T., Ogawa, A., Suda, A., Nakajima, K., Tanaka, M., Oka, S., Kamagata, K., Aoki, S., Oshima, Y., Tanaka, S., Hattori, N., Konishi, S.: Parallel cognitive processing streams in human prefrontal cortex: Parsing areal-level brain network for response inhibition. *Cell Reports* **36**(12), 109732 (2021). <https://doi.org/https://doi.org/10.1016/j.celrep.2021.109732>, <https://www.sciencedirect.com/science/article/pii/S2211124721011815>
17. Rousseeuw, P.J.: Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* **20**, 53–65 (1987). [https://doi.org/https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/https://doi.org/10.1016/0377-0427(87)90125-7), <https://www.sciencedirect.com/science/article/pii/0377042787901257>
18. Schmid, P.J.: Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics* **656**, 5–28 (2010)
19. Seitzman, B.A., Gratton, C., Marek, S., Raut, R.V., Dosenbach, N.U., Schlaggar, B.L., Petersen, S.E., Greene, D.J.: A set of functionally-defined brain regions with improved representation of the subcortex and cerebellum. *NeuroImage* **206**, 116290 (2020). <https://doi.org/https://doi.org/10.1016/j.neuroimage.2019.116290>, <https://www.sciencedirect.com/science/article/pii/S105381191930881X>
20. Sigman, M., Dehaene, S.: Brain mechanisms of serial and parallel processing during dual-task performance. *Journal of Neuroscience* **28**(30), 7585–7598 (2008)
21. Sussillo, D., Jozefowicz, R., Abbott, L., Pandarinath, C.: Lfads - latent factor analysis via dynamical systems (08 2016)

22. Takeishi, N., Kawahara, Y., Yairi, T.: Learning koopman invariant subspaces for dynamic mode decomposition. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. p. 1130–1140. NIPS’17, Curran Associates Inc., Red Hook, NY, USA (2017)
23. Turja, M.A., Wu, G., Yang, D., Styner, M.A.: Learning the latent heat diffusion process through structural brain network from longitudinal β -amyloid data. In: Heinrich, M., Dou, Q., de Bruijne, M., Lellmann, J., Schläfer, A., Ernst, F. (eds.) Proceedings of the Fourth Conference on Medical Imaging with Deep Learning. Proceedings of Machine Learning Research, vol. 143, pp. 761–773. PMLR (07–09 Jul 2021), <https://proceedings.mlr.press/v143/turja21a.html>
24. Turja, M.A., Zsembik, L.C.P., Wu, G., Styner, M.: Constructing consistent longitudinal brain networks by group-wise graph learning. In: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.T., Khan, A. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2019. pp. 654–662. Springer International Publishing, Cham (2019)
25. Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., et al.: The human connectome project: a data acquisition perspective. *Neuroimage* **62**(4), 2222–2231 (2012)
26. Viviani, R., Grön, G., Spitzer, M.: Functional principal component analysis of fmri data. *Human brain mapping* **24**(2), 109–129 (2005)
27. Xiao, J., Uddin, L.Q., Meng, Y., Li, L., Gao, L., Shan, X., Huang, X., Liao, W., Chen, H., Duan, X.: A spatio-temporal decomposition framework for dynamic functional connectivity in the human brain. *NeuroImage* **263**, 119618 (2022). <https://doi.org/https://doi.org/10.1016/j.neuroimage.2022.119618>, <https://www.sciencedirect.com/science/article/pii/S1053811922007339>

Supplementary Material — DeepGraphDMD: Interpretable Spatio-Temporal Decomposition of Non-linear Functional Brain Network Dynamics

Md Asadullah Turja¹, Martin Styner², and Guorong Wu³

¹ Department of Computer Science, University of North Carolina at Chapel Hill,
mturja@cs.unc.edu

² Department of Computer Science, University of North Carolina at Chapel Hill,
styner@email.unc.edu

³ Department of Psychiatry, University of North Carolina at Chapel Hill,
guorong_wu@med.unc.edu

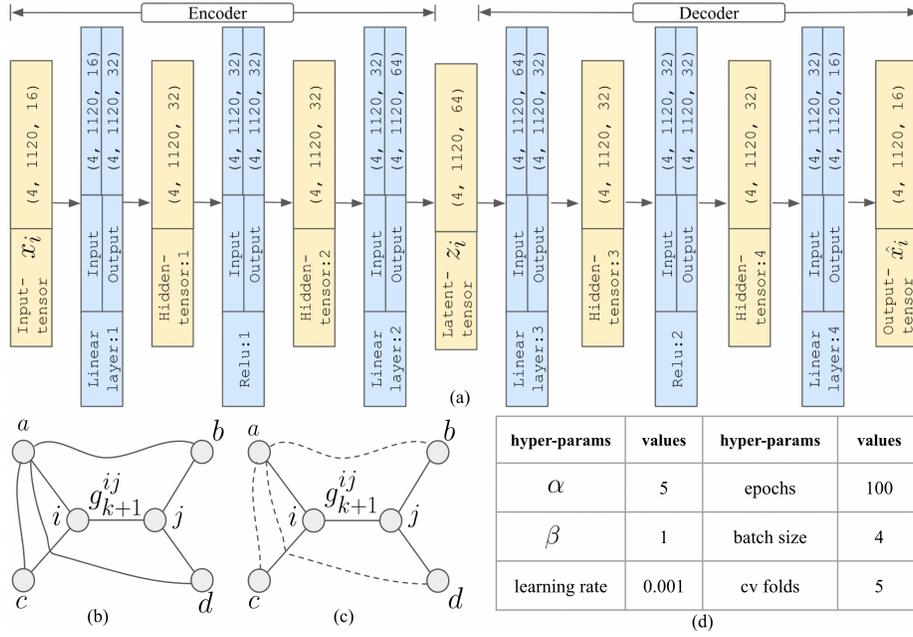


Fig. 1. (a) The network architecture of DeepGraphDMD, (b) regular Koopman Operator represents g_{k+1}^{ij} as a linear combination of previous values of all the other edges ($g_k^{ia}, g_k^{jb}, g_k^{ab}, g_k^{ic}, g_k^{jd}, g_k^{ad}, g_k^{ij}$) of the graph whereas (c) Sparse Koopman Operator only considers neighboring edges: $g_k^{ia}, g_k^{jb}, g_k^{ic}, g_k^{jd}, g_k^{ij}$ (doesn't consider the dashed edges), and (d) table of hyper-parameters for the DeepGraphDMD training.

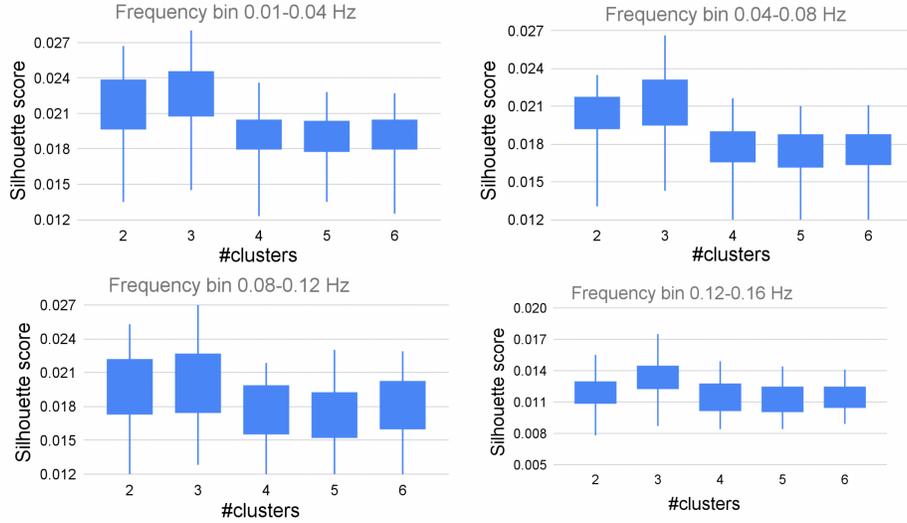


Fig. 2. Boxplot of the silhouette scores from all the subjects with respect to the number of clusters across different frequency bins (higher is better).

n_components	2	3	4	5	6
PCA	0.238	0.238	0.238	0.238	0.238
ICA	0.248	0.266	0.217	0.165	0.159

Table 1. The value of r of CogTotalComp for different n_components in PCA and ICA. r values for PCA don't change because the first PCA component always gives the best r value which doesn't change with n_components.

	Frequency (Hz)	CogFluidComp	PMAT24_A_CR	ReadEng	CogTotalComp
Graph DMD	0.01-0.04	0.184 ± 0.012	0.175 ± 0.009	0.186 ± 0.017	0.219 ± 0.021
	0.04-0.08	0.095 ± 0.010	0.180 ± 0.012	0.182 ± 0.021	0.184 ± 0.021
	0.12-0.16	0.156 ± 0.015	0.167 ± 0.011	0.183 ± 0.019	0.192 ± 0.017
Deep Graph DMD	0.01-0.04	0.181 ± 0.011	0.182 ± 0.013	0.182 ± 0.015	0.207 ± 0.018
	0.04-0.08	0.112 ± 0.014	0.181 ± 0.011	0.185 ± 0.015	0.187 ± 0.019
	0.12-0.16	0.155 ± 0.011	0.167 ± 0.016	0.186 ± 0.012	0.193 ± 0.014

Table 2. Comparison of r for the behavioral measures for the other frequency bins.