

The Effect of Graph Frequencies on Dynamic Structures in Graph Signal Processing

Goerttler, S., Wu, M. & He, F

Author post-print (accepted) deposited by Coventry University's Repository

Original citation & hyperlink:

Goerttler, S, Wu, M & He, F 2023, The Effect of Graph Frequencies on Dynamic Structures in Graph Signal Processing. in 2022 IEEE Signal Processing in Medicine and Biology Symposium, SPMB 2022. 2022 IEEE Signal Processing in Medicine and Biology Symposium, SPMB 2022 - Proceedings, Institute of Electrical and Electronics Engineers Inc., pp. 1-6, 2022 IEEE Signal Processing in Medicine and Biology Symposium, Philadelphia, Pennsylvania, United States, 3/12/22.

<https://doi.org/10.1109/SPMB55497.2022.10014954>

DOI 10.1109/SPMB55497.2022.10014954

ISSN 2372-7241

ESSN 2473-716X

Publisher: IEEE

© 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Copyright © and Moral Rights are retained by the author(s) and/ or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This item cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder(s). The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

This document is the author's post-print version, incorporating any revisions agreed during the peer-review process. Some differences between the published version and this version may remain and you are advised to consult the published version if you wish to cite from it.

The Effect of Graph Frequencies on Dynamic Structures in Graph Signal Processing

S. Goerttler^{1,2}, M. Wu² and F. He¹

1. Centre for Computational Science and Mathematical Modelling,
Coventry University, Coventry, United Kingdom
2. Institute for Infocomm Research, Agency for Science,
Technology and Research (A*STAR), Singapore
goerttlers@uni.coventry.ac.uk, wumin@i2r.a-star.edu.sg, fei.he@coventry.ac.uk

Abstract— Multivariate signals are signals consisting of multiple signals measured simultaneously over time and are most commonly acquired by sensor networks. The emerging field of graph signal processing (GSP) promises to analyse dynamic characteristics of multivariate signals, while at the same time taking the network, or spatial structure between the signals into account. To do so, GSP decomposes the multivariate signals into graph frequency signals, which are ordered by their magnitude. However, the meaning of the graph frequencies in terms of this ordering remains poorly understood. Here, we investigate the role the ordering plays in preserving valuable dynamic structures in the signals, with neuroimaging applications in mind. In order to overcome the limitations in sample size common to neurophysiological data sets, we introduce a minimalist simulation framework to generate arbitrary amounts of data. Using this artificial data, we find that lower graph frequency signals are less suitable for classifying neurophysiological data than higher graph frequency signals. We further introduce a baseline testing framework for GSP. Using this framework, we conclude that dynamic, or spectral structures are poorly preserved in GSP, highlighting current limitations of GSP for neuroimaging.

Keywords— multivariate signals, neurophysiological signals, graph signal processing, graph Fourier transform.

I. INTRODUCTION

Multivariate signals arise in areas as diverse as biomedical imaging and geophysical signal analysis, where multiple spatially distributed sensors measure signals simultaneously. Unlike regular signals, multivariate signals can capture both temporal and spatial characteristics of the underlying system. The temporal characteristics can be analysed with conventional signal processing, e.g., by simply analysing each signal independently and collating the results across the signals. On the other hand, the spatial characteristics can be analysed using *spectral graph theory*, where the connectivity structure between the signals is determined and subsequently analysed [1]. However, neither strategy analyses the temporal and spatial characteristics jointly, possibly disregarding valuable interactions between the two.

Graph signal processing (GSP) is a recent method that allows to analyse the temporal characteristics while simultaneously taking the underlying graph structure of the signals into account [2]. Specifically, GSP uses the underlying graph structure to transform the multivariate

signal by decomposing the graph in *graph Fourier modes*. This transformation is called the *graph Fourier transform* (GFT).

GSP has been employed for neurophysiological imaging, such as functional magnetic resonance imaging (fMRI) [3–6] or electroencephalography (EEG) [7–9]. These applications include a wide variety of GSP techniques, such as filtering of signals [4], spectral decomposition and dimensionality reduction [3], or total variation analysis [7]. GSP can also be used in an inverse manner to retrieve the graph structure from a multivariate signal [8]. This is achieved by finding the graph that minimises the total variation, following the assumption that the multivariate signal is smooth across the graph. GSP techniques, such as graph spectral decomposition, can be leveraged to improve classification of neurophysiological signals [3], which commonly uses temporal characteristics to derive biomarkers. However, this means that the effective use of GSP for the classification of neurophysiological signals hinges on how well the GFT applied on multivariate signals preserves these temporal structures, such as spectral features. In this study, we therefore systematically evaluate the classification performance restricted to each graph frequency, giving insight into which graph frequencies preserve these temporal structures best.

There are multiple choices not only to retrieve the graph structure for a data set, but also to compute the graph Fourier modes from this graph, making the GFT ambiguous. Despite this ambiguity, recent applications of GSP have paid little attention to validating the method with suitable baseline models [4, 7]. To close this gap, we introduce a set of baseline models, which either distort or exclude the graph structure. Testing the model against the set of baseline models allows to isolate and estimate the specific value of taking the graph structure into account.

Our contributions in this study can be summarised as follows:

- We develop a minimalist simulation framework, which captures the essential characteristics of multivariate signals;
- introduce a baseline testing framework for GSP; and

- determine to what extent dynamic structures are preserved in GSP in dependence on the graph frequency.

II. THEORY

II-A. Multivariate signals in biomedical imaging

Many brain imaging techniques, such as fMRI, EEG, or magnetoencephalography (MEG), record multiple signals simultaneously in time at different spatial locations. In fMRI, the signal locations are called *voxels*, whereas in MEG and EEG they are called *channels*. For example, an EEG measurement setup with N_c channels which record the brain over a time period T at a sampling rate f yields a data matrix $\mathbf{X} \in \mathbb{R}^{N_c \times N_t}$, where $N_t = \lfloor f \cdot T \rfloor$ is the number of time samples. The i -th row of \mathbf{X} correspond to the *temporal signal*, or simply signal, at location i and can be denoted as \mathbf{x}_{i*} . On the other hand, the j -th column corresponds to the spatial, or *graph signal* at time sample j , which we denote as \mathbf{x}_{*j} .

II-B. Graphs

Formally, networks can be represented by a *weighted graph* $\mathcal{G} := (\mathcal{V}, \mathbf{A})$, which consists of N vertices $\mathcal{V} = \{1, 2, 3, \dots, N\}$ and the *weighted adjacency matrix* $\mathbf{A} \in \mathbb{R}^{N \times N}$, whose entries a_{ij} represent the strength of the connectivity between node i and node j . If the connectivities between each pair of nodes are symmetric, or $a_{ij} = a_{ji}$, then \mathbf{A} is symmetric and the graph is said to be undirected. Furthermore, if the diagonal elements a_{ii} are all zero, meaning that the nodes are not connected to themselves with loops, the graph is called a *simple graph*.

A second, useful representation of a graph \mathcal{G} is its *Laplacian matrix*. The Laplacian can be directly computed from the adjacency matrix \mathbf{A} as $\mathbf{L} = \mathbf{D} - \mathbf{A}$. Here, $\mathbf{D} = \text{diag}(\sum_j a_{1j}, \dots, \sum_j a_{Nj})$ denotes the *degree matrix*. While the weighted adjacency matrix can be viewed as a graph shift operator, the Laplacian can be associated with the negative discrete Laplace operator.

II-C. Graph retrieval

GSP transforms the multivariate signals using the graph, or spatial structure which underlies the data. The retrieval of the graph is not unique in neuroimaging, but can be based on the functional connectivity, the structural connectivity, or the geometric location of the sensors [7]. Even though some of these graph retrieval methods can yield similar graphs in neuroimaging [10], the choice of the graph retrieval method can have a significant effect on the results [3] as well as on the interpretation thereof.

In this study, the graph is based on the *functional connectivity* between the signals. This method is purely data-driven and can therefore be used on all data sets. Common choices to build the graph include computing

pairwise Pearson correlations or covariances, but nonlinear measures such as mutual information can be used as well. As shown in subsection II-E, there is a link between GFT using functionally retrieved graphs and principal component analysis.

II-D. Graph Fourier Transform

The GFT is an extension of the Fourier transform to graphs and is at the heart of GSP. Let \mathcal{G} be a weighted, undirected graph without loops. The *eigendecomposition* of the Laplacian \mathbf{L} of graph \mathcal{G} is given by $\mathbf{L} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^\top$. The eigenvalues λ_i of \mathbf{L} are the diagonal elements of $\mathbf{\Lambda}$ and real-valued, sorted in ascending order, whereas the eigenvectors \mathbf{v}_i are the orthogonal columns \mathbf{q}_{*i} of \mathbf{Q} . Alternatively, the weighted adjacency matrix or the symmetrically normalised Laplacian can be used instead of the Laplacian [2, 11, 12].

The GFT, which transforms a spatial signal $\mathbf{x} \in \mathbb{R}^{N_c}$ into its graph spectral representation $\tilde{\mathbf{x}}$, is then simply defined as:

$$\tilde{\mathbf{x}} = \mathbf{GFT} \mathbf{x} := \mathbf{Q}^\top \mathbf{x}. \quad (1)$$

Here, \mathbf{GFT} denotes the graph Fourier transform matrix. The i -th row of \mathbf{GFT} corresponds to the *graph Fourier mode* at graph frequency i , which is given by the eigenvector \mathbf{v}_i . Likewise, a multivariate signal $\mathbf{X} \in \mathbb{R}^{N_c \times N_t}$, where each column \mathbf{x}_{*j} corresponds to a spatial signal at time sample j , can be transformed column-wise: $\tilde{\mathbf{X}} = \mathbf{GFT} \mathbf{X}$. The rows $\tilde{\mathbf{x}}_{i*}$ of the transformed multivariate signal $\tilde{\mathbf{X}}$ are the transformed signals, which are associated with the eigenvalues λ_i .

For the Laplacian, the graph Fourier modes are eigenvectors of the discrete Laplace operator, which is a second-order differential operator in space. In the same way, classical Fourier modes $e^{i2\pi ft}$ with frequency f are eigenfunctions of the second partial derivative with respect to time t . This analogy links the classical Fourier transform to the graph Fourier transform. If the adjacency matrix is used instead, the analogy becomes less obvious theoretically and may only hold for special cases such as the cyclic shift graph, because the adjacency matrix does not correspond to a differential operator.

It can be shown that the ordering of the eigenvalues can be linked to the ordering of the graph frequencies [12], enabling us to refer to the transformed signals as ordered *graph frequency signals*. If the Laplacian matrix is used in the GFT, higher eigenvalues correspond to higher graph frequencies, whereas if the adjacency matrix is used, higher eigenvalues correspond to lower graph frequencies. Note that in the latter case, eigenvectors with higher graph frequencies have lower weights in the eigendecomposition, meaning that these frequencies contribute less to the adjacency matrix. However, this suggests that transformed signals with higher graph

frequencies may only weakly depend on the adjacency matrix, and thereby the spatial structure.

II-E. Link to principal component analysis (PCA)

Principal component analysis (PCA) projects multivariate samples onto a set of orthogonal components, which are the eigenvectors of the covariance matrix of the data set. Crucially, if the covariance matrix is used for the GFT, then GFT and PCA are mathematically equivalent. Note that for normalised signals with standard deviation one, which is common for some neurophysiological signals such as EEG signals, the correlation and the covariance matrix are the same. Note also that the diagonal entries of the signal correlation matrix are one and will not affect the eigenvectors. In other words, using the correlation matrix of normalised signals with or without diagonal elements for GFT is equivalent to PCA.

II-F. Baseline models

We compose a set of three baseline models, which either distort the graph in GSP, or use conventional signal processing. The goal of all three models is to isolate the boost in performance as a result of using the graph structure underlying the data.

The first baseline model is the *permuted nodes GSP model*, which uses the graph underlying the data, but randomly permutes the nodes of this graph. In this way, the eigenmodes have the same weights as the actual model, but at different locations. Therefore, the model distorts the graph structure, while retaining the effect of the weight magnitudes. The second, related baseline model is the *random graph GSP model*, which uses GSP with a randomly generated graph. While this model transforms the data using GFT, these transformations are not based on the actual graph structure. The third baseline model is the *single channel model*, which does not transform the data and is equivalent to conventional signal processing. It can also be viewed as GSP with the identity transform, which is given by the identity matrix $\mathbb{1}$. Note that the "graph Fourier modes" are given by the single channels and have no natural ordering, which can also be seen by the fact that the eigenvalues of the identity matrix are degenerate. This model is the only model that excludes the graph structure, making it an indispensable baseline model.

III. METHODOLOGY

III-A. Simulated data set

We developed Algorithm 1 to generate multivariate signal samples $\mathbf{X}_{sim} \in \mathbb{R}^{N_c \times N_t}$ with N_c channels and N_t time samples. The generated samples simulate neuroimaging measurements, such as EEG, with which they share two crucial characteristics: Firstly, the samples have an underlying graph structure, such that each pair of signals

Algorithm 1 Dynamic neuroimaging data generation

Input $N_t, \mathbf{A}_s, h, \alpha, \beta, \gamma$

Output \mathbf{X}

- 1: $\varepsilon_{ij}, \tilde{\varepsilon}_{ij} \sim \mathcal{N}(0, 1^2), \quad i = 1, \dots, N_c, j = 1, \dots, N_t$
 - 2: $\hat{\varepsilon}_{i*} \leftarrow (\mathcal{F}^{-1} \circ h \circ \mathcal{F})(\tilde{\varepsilon}_{i*}), \quad i = 1, \dots, N_c$
 - 3: $\mathbf{x}_{*1} \leftarrow \beta \hat{\varepsilon}_{*1}$
 - 4: **for** $t \leftarrow 2, N_t$ **do**
 - 5: $\hat{\mathbf{x}}_{*t-1} \leftarrow \mathbf{x}_{*t-1} - \bar{\mathbf{x}}_{*t-1} + \gamma \boldsymbol{\varepsilon}_{*t}$
 - 6: $\mathbf{x}_{*t} \leftarrow \alpha \mathbf{A}_s \hat{\mathbf{x}}_{*t-1} + \beta \hat{\varepsilon}_{*t}$
 - 7: **end for**
-

has a specific connectivity associated to it. Secondly, the temporal signals have a specific spectral profile.

The algorithm takes as input the number of generated time samples N_t , the weighted adjacency matrix $\mathbf{A}_s \in \mathbb{R}^{N_c \times N_c}$, the filter function h , as well as parameters α , β and γ . While \mathbf{A}_s controls the connectivity, or spatial structure, the filter function h controls the spectral profile. Specifically, using two filter functions h_1 and h_2 allows to generate data samples with two conditions. The similarity between the two filter functions controls the difficulty of classifying the conditions. This simulates, for example, EEG data where multiple conditions, such as Alzheimer's disease or healthy control, each have a specific spectral profile [13]. Parameter α controls the strength of the correlation structure, while parameter β controls the strength of the spectral structure. Parameter γ controls the self-amplification of the signals during the simulation.

In each time step of Algorithm 1, we firstly centre the graph signal of the previous time step \mathbf{x}_{*t-1} around zero and add Gaussian noise $\gamma \boldsymbol{\varepsilon}_{*t}$ to this signal (line 5). Secondly, we use the adjacency matrix \mathbf{A}_s as a graph shift operator to translate the graph signal in time, which enforces the structural connectivity \mathbf{A}_s in our data (line 6). Thirdly, we scale the translated signal by α and add normalised coloured noise $\hat{\varepsilon}_{*t}$ scaled with β , whose spectral density profile is controlled by the filter h (line 6). Finally, the simulated multivariate signal is labelled with its condition.

We used three different filter functions h_2 for condition 2 (orange lines in Figure 1(c)), resulting in overall three data sets which are either easy, medium or difficult to classify. The connectivity structure matrix $\mathbf{A}_s \in \mathbb{R}^{23 \times 23}$ was generated by drawing weights $a_{ij}^s \sim \mathcal{U}_{[-0.1, 0.4]}$ from a uniform distribution for $i > j$, and setting $a_{ii}^s = 0.4$ and $a_{ji}^s = a_{ij}^s$. We set the simulation parameters to $\alpha = 0.5$, $\beta = 1$, and $\gamma = 1$. For each simulated participant, we generated $N_c = 23$ time-series signals with $N_t = 2048$ time samples each.

Figure 1 shows the simulated signals (a) along with their spatial (b) and spectral (d) structure. The spectral

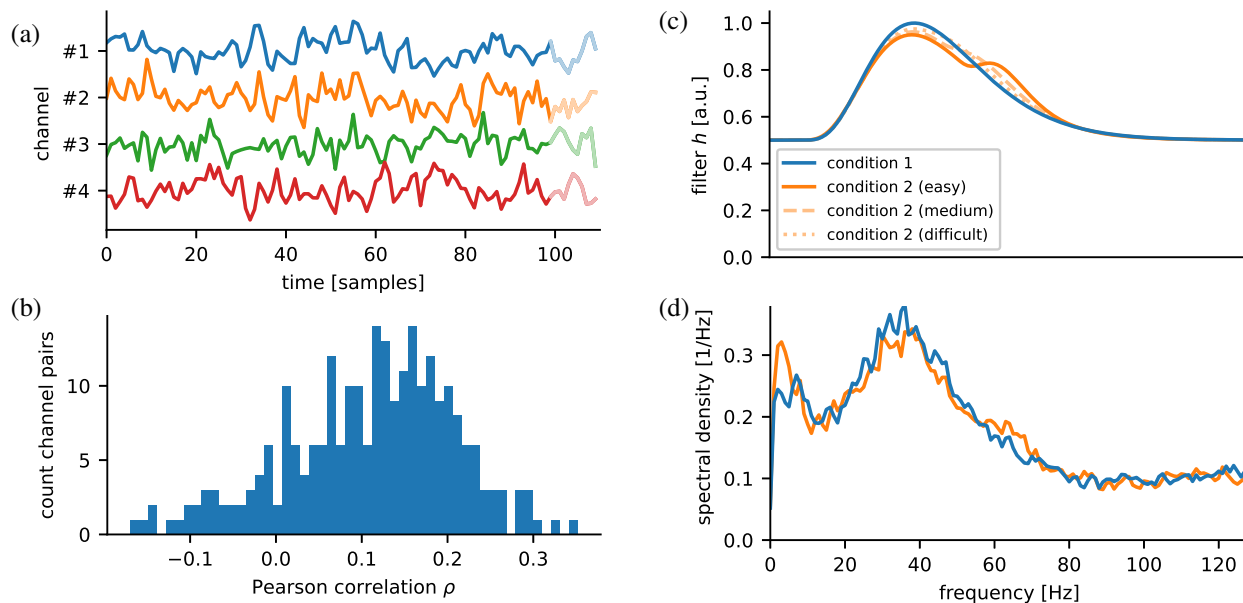


Figure 1. Simulated neurophysiological signals. (a) Time signals for four selected channels across the first 100 samples. Channels #2, #3 and #4 are positively, weakly and negatively correlated to channel #1, respectively. (b) Distribution of correlations of channel pairs, exhibiting the spatial connectivity structure. In the simulation, this structure is controlled by the matrix \mathbf{A}_s , but it is also affected by the simulation parameters. (c,d) Demonstration of power spectral density control and adjustment of classification difficulty. (c) The Fourier filter function h used to colour the noise in Algorithm 1, assuming a sampling frequency of 256 Hz. For each data set, two conditions are simulated. Condition 2 can be varied to make it easy (solid), medium (dashed), or difficult (dotted) to distinguish from condition 1. The difficulty depends on the similarity between the two conditions. (d) Welch power spectral density of simulated signals averaged across all channels for two easily distinguishable conditions. Figures (c) and (d) clearly show that the shape of the power spectral density profile of the simulated signals can be controlled. Parameter α in Algorithm 1 can be used to reduce the power density at lower frequencies

structure (d) of the simulated signals is enforced by the spectral density profile of the coloured noise (c), demonstrating that the spectral structure can be controlled.

III-B. Analysis

The goal of the analysis is to GFT-transform the multivariate signal into graph frequency signals and subsequently classify the samples into the two conditions, using only one graph frequency signal at a time. This approach allows us to determine the classification performance for each graph frequency, which we use to assess how well dynamical structures are preserved at this frequency.

The first step of the analysis consists of retrieving the graph structure from the training samples, for which we used the functional connectivity. Note that we need to use the same graph for all samples to keep the graph Fourier modes constant. Specifically, we computed the correlation matrix for each sample in the training set and subsequently averaged all matrices, yielding a common weighted adjacency matrix. Secondly, we carried out the GFT, yielding $N_c = 23$ graph frequency signals. We used the weighted adjacency and the Laplacian matrix for the GFT in two separate experiments. Thirdly, we extracted

spectral features from those signals for each sample. To this end, we computed the Welch power spectral density with a window of 128 for each transformed signal, removed the last 28 values, and downsampled the remaining 100 values by averaging five values each, yielding 20 features per graph frequency. Lastly, we trained a support vector machine classifier separately for each graph frequency to classify the labelled samples, using only the 20 features calculated from the graph frequency signal. The baseline models are analysed following the same steps, except that the graph in the GFT was replaced as described in subsection II-F.

III-C. Testing

Generating the data matrices is time-consuming, as one simulation step is needed for each time sample. We therefore used a modified, perpetual version of cross-validation, illustrated in Figure 2. Initially, $N_s = 100$ samples are generated, divided evenly in condition 1 and condition 2. This data set is split into $k = 10$ folds, nine of which are used for training. This results in 90 training samples, mimicking the sparsity of samples in neuroimaging. The remaining fold is used for testing, resulting in 10 testing samples per iteration. For the




Figure 2. Illustration of the perpetual cross-validation used during the analysis. The whole data set is split into $k = 10$ folds. One fold is assigned to be the testing set, while the remaining $k - 1$ folds comprise the training set. For the subsequent iteration, a fold with newly generated data is added to the training set, while the testing fold is shifted by one fold

subsequent iteration, data is generated to form a new fold with $N_s/k = 10$ samples, which is added to the training set, whereas the testing fold is shifted by one fold. The final accuracy scores are averaged across all samples in the testing set and across all iterations. As a result of varying the number of iterations per model, our main model was tested on overall 5000 testing samples, whereas the permuted nodes GSP, random graph GSP, and single channel baseline model were tested on 5000, 4000, and 2000 testing samples, respectively.

The advantage of this perpetual version of cross-validation is that it can have arbitrarily many iterations, while at the same time reusing each fold k -times. It further does not require much storage, as used samples can be overwritten with newly generated samples.

IV. RESULTS

Figure 3 illustrates the results of our analysis of the artificially generated data. Specifically, it shows the classification accuracy in dependence on the graph frequency of the transformed signal for three simulated difficulty levels. The main model, shown in blue, is compared against the three baseline models. The accuracy scores in the figures are averaged across all testing samples, as described in subsection III-C.

The accuracy of the graph DC component is by far the lowest. This can be understood by considering the case of the Laplacian matrix: When using this matrix for the GFT, the graph DC component is equivalent to the average across all signals. However, when averaging the signals, some temporal structures in the signals are also averaged out.

The model performs worse at lower graph frequencies than all GSP-based baseline models, and only equally well than the single channel baseline model. The relatively poor performance of the single channel baseline model can be attributed to the fact that, as opposed to all other models, this model does not gather spectral structures from more than one signal, giving the




Figure 3. Classification accuracy as a function of the graph frequency of the transformed signals. Three data sets with varying classification difficulty were simulated, whereby the difficulty was controlled by modifying the filter h as shown in Figure 1(c). Easy, medium, and difficult classification difficulties result in the high, medium, and low accuracies overlaid in the figure, respectively. The model is shown in blue, while the baseline models are shown in orange, green and red. Note that the single channel baseline model in red does not have a graph frequency ordering. Models using the GFT with the Laplacian (adjacency) matrix are shown in solid (dotted) lines

classifier less information to use. For easy to classify data sets, our model perform slightly better than the baseline models at higher graph frequencies.

The choice of the specific graph representation for the GFT does not have a strong impact on the results, which has been reported earlier[4]. Further, we show that our results are replicated for various difficulty levels.

V. DISCUSSION AND CONCLUSION

In this study, we have developed a neuroimaging data generation algorithm, which dynamically generates arbitrarily many samples. We have demonstrated that this data has a spatial connectivity structure, as well as a controllable spectral structure. However, our simulated

data may differ from real data in crucial ways. For example, the spectral structure is enforced by adding coloured noise, and does not primarily arise out of the network interaction, which however may be the case for real-life neuroimaging data. Nevertheless, we uphold that the generated data capture the essential characteristics of neurophysiological signals relevant for GSP, as demonstrated in Figure 1, and hence that their analysis allows to draw some general conclusions about GSP applied to real-life neurophysiological signals.

One limitation of GSP for neuroimaging is the ambiguity of the GFT, which is due to the number of choices for how to retrieve the graph (see subsection II-C) and for which graph representation to use. These limitations, in combination with the novelty of the method, highlight the need for a thorough validation procedure for GSP, as opposed to basing it mainly on theoretical justifications. To this end, we have introduced a baseline testing framework for GSP in subsection II-F, consisting of a set of overall three baseline models. Note that the GFT can be linked to the PCA, as shown in subsection II-E, making the PCA unsuitable as a baseline model.

Using the artificial data, we have systematically evaluated the classification performance of single graph frequency signals in terms of their graph frequency. We found that higher graph frequencies outperform lower ones, which is in line with the results obtained in [3]. This specific result may not generalise to all neurophysiological data sets, simulated or real-life, because it may depend on the precise interaction between the spatial and the spectral structure. The result is nevertheless surprising, as we expected lower graph frequencies to preserve temporal structures better, given that they are composed of more closely related channels. We further found that even at higher graph frequencies our model performs only marginally better than the baseline models. Note also that higher graph frequencies may not relate much to the spatial structure, as pointed out in subsection II-D. Taken together, our results do not suggest that GSP leverages the spatial structure to improve spectral feature-based classification.

One possible explanation for this result is that in order to GFT-transform the signals, the signals are mixed and therefore interfere with each other, which leads to an attenuation of the spectral features. This shortcoming may limit the application of GSP for EEG, where the dynamic structure, i.e., the spectral features, is crucial to characterise the data, but it may also impose limitations for applying GSP to neurophysiological data in general.

REFERENCES

- [1] S. Atasoy, I. Donnelly, and J. Pearson, "Human brain networks function in connectome-specific harmonic waves," *Nature communications*, vol. 7, no. 1, pp. 1–10, 2016.
- [2] A. Ortega, P. Frossard, J. Kovačević, J. M. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 808–828, 2018.
- [3] M. Ménoiret, N. Farrugia, B. Padeloup, and V. Gripon, "Evaluating graph signal processing for neuroimaging through classification and dimensionality reduction," *2017 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, 2017, pp. 618–622.
- [4] W. Huang, T. A. Bolton, J. D. Medaglia, D. S. Bassett, A. Ribeiro, and D. Van De Ville, "A graph signal processing perspective on functional brain imaging," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 868–885, 2018.
- [5] S. Itani and D. Thanou, "A graph signal processing framework for the classification of temporal brain data," *2020 28th European Signal Processing Conference (EUSIPCO)*. IEEE, 2021, pp. 1180–1184.
- [6] H. Behjat and M. Larsson, "Spectral characterization of functional mri data on voxel-resolution cortical graphs," *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2020, pp. 558–562.
- [7] S. Mortaheb, J. Annen, C. Chatelle, H. Cassol, G. Martens, A. Thibaut, O. Gosseries, and S. Laureys, "A graph signal processing approach to study high density eeg signals in patients with disorders of consciousness," *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2019, pp. 4549–4553.
- [8] S. S. Saboksayr, G. Mateos, and M. Cetin, "Eeg-based emotion classification using graph signal processing," *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 1065–1069.
- [9] K. Glomb, J. R. Queralt, D. Pascucci, M. Defferrard, S. Tourbier, M. Carboni, M. Rubega, S. Vulliemoz, G. Plomp, and P. Hagmann, "Connectome spectral analysis to track eeg task dynamics on a subsecond scale," *NeuroImage*, vol. 221, p. 117137, 2020.
- [10] J. S. Damoiseaux and M. D. Greicius, "Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity," *Brain structure and function*, vol. 213, no. 6, pp. 525–533, 2009.
- [11] A. Gavili and X.-P. Zhang, "On the shift operator, graph frequency, and optimal filtering in graph signal processing," *IEEE Transactions on Signal Processing*, vol. 65, no. 23, pp. 6303–6318, 2017.
- [12] A. Sandryhaila and J. M. Moura, "Discrete signal processing on graphs: Frequency analysis," *IEEE Transactions on Signal Processing*, vol. 62, no. 12, pp. 3042–3054, 2014.
- [13] J. Jeong, "Eeg dynamics in patients with alzheimer's disease," *Clinical neurophysiology*, vol. 115, no. 7, pp. 1490–1505, 2004.