

Graph Frequency Analysis of Brain Signals

Weiyou Huang, Leah Goldsberry, Nicholas F. Wymbs, Scott T. Grafton, Danielle S. Bassett, and Alejandro Ribeiro

Abstract—This paper presents methods to analyze functional brain networks and signals from graph spectral perspectives. The notion of frequency and filters traditionally defined for signals supported on regular domains such as discrete time and image grids has been recently generalized to irregular graph domains and defines brain graph frequencies associated with different levels of spatial smoothness across the brain regions. Brain network frequency also enables the decomposition of brain signals into pieces corresponding to smooth or rapid variations. We relate graph frequency with principal component analysis when the networks of interest denote functional connectivity. The methods are utilized to analyze brain networks and signals as subjects master a simple motor skill. We observe that brain signals corresponding to different graph frequencies exhibit different levels of adaptability throughout learning. Further, we notice a strong association between graph spectral properties of brain networks and the level of exposure to tasks performed and recognize the most contributing and important frequency signatures at different levels of task familiarity.

Index Terms—Functional brain network, fMRI, filtering, graph signal processing, motor learning, network theory.

I. INTRODUCTION

THE study of brain activity patterns has proven valuable in identifying neurological disease and individual behavioral traits [1]–[3]. The use of functional brain networks describing the tendency of different regions to act in unison has proven complementary in the analysis of similar matters [4]–[7]. It is not surprising that signals and networks prove useful in similar problems since the two are closely related. In this paper we advocate an intermediate path in which we interpret brain activity as a signal supported on the graph of brain connectivity. We show how the use of graph signal processing tools can be used to glean information from brain signals using the network as an aid to identify patterns of interest. The benefits of incorporating network information into signal analysis has been demonstrated

Manuscript received November 1, 2015; revised May 1, 2016 and July 26, 2016; accepted July 29, 2016. Date of publication August 16, 2016; date of current version September 23, 2016. This work was supported by NSF CCF-1217963, PHS NS44393, ARO ICB W911NF-09-0001, ARL W911NF-10-2-0022, ARO W911NF-14-1-0679, NIH R01-HD086888, NSF BCS-1441502, NSF BCS-1430087, and the John D. and Catherine T. MacArthur and Alfred P. Sloan foundations. Content of this paper does not necessarily represent official views of any of the funding agencies. The guest coordinating the review of this paper and approving it for publication was Dr. Dimitri Van De Ville.

W. Huang, L. Goldsberry, D. S. Bassett, and A. Ribeiro are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: whuang@seas.upenn.edu; goleah@seas.upenn.edu; dsb@seas.upenn.edu; aribeiro@seas.upenn.edu).

N. F. Wymbs is with the Department of Physical Medicine and Rehabilitation, Johns Hopkins University, Baltimore, MD 21218 USA (e-mail: nwymbs1@jhmi.edu).

S. T. Grafton is with the Department of Psychological and Brain Sciences, University of California at Santa Barbara, Santa Barbara, CA 93106 USA (e-mail: scott.grafton@psych.ucsb.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSTSP.2016.2600859

in multiple domains. Notable examples of applications include video compression [8], breast cancer diagnostics [9], movie recommendation [10], and semi-supervised learning [11].

The fundamental GSP concepts that we utilize to exploit brain connectivity in the analysis of brain signals are the graph Fourier transform (GFT) and the corresponding notions of graph frequency components and graph filters. These concepts are generalizations of the Fourier transform, frequency components, and filters that are used in regular domains such as time and spatial grids [12]–[14]. As such, they permit the decomposition of a graph signal into components that represent different modes of variability. We can define low graph frequency components representing signals that change slowly with respect to brain connectivity networks in a well defined sense and high graph frequency components representing signals that change fast in the same sense. This is important because low and high *temporal* variability have proven important in the analysis of neurological disease and behavior [15], [16]. GFT based decompositions permit a similar analysis of variability across regions of the brain for a fixed time—a sort of *spatial* variability measured with respect to the connectivity pattern. We demonstrate here that it is useful in a similar sense; see e.g. Figs. 6–8 and 10.

The GSP studies in this paper are related to principal component analysis (PCA), which has been used with success in the analysis of brain signals [17], [18]. The difference with the GSP analysis we present here is that PCA implicitly assumes the brain network to be a correlation matrix and the signals to be drawn from a stochastic model. More importantly, whereas the GFT can be used to, e.g., decompose the signal into low, medium, and high frequency components, PCA is mostly utilized for dimensionality reduction; which in the language of this paper is tantamount to analyzing a few low graph frequency components. Another important difference is that PCA focuses on identifying variability across *different* realizations of brain signals, but the GFT identifies spatial variability of a *single* realization. GSP is also related to the spectral analysis of networks in general and Laplacians in particular [19], [20]. The difference in this case is that these spectral analyses yield properties of the networks. In GSP analyses, the network provides an underlying structure, but the interest is on signals expressed on this stratum.

Recent studies have already demonstrated that exploiting information from the underlying graph networks of structural brain connectivity to filter functional signals such as fMRI [21] and EEG [22] result in increased localization accuracy. The notion of investigating specific graph frequencies has just recently been introduced to capture key features of structural resting state networks (RSNs). Individual low frequencies have been identified to represent with low reconstruction error resting state networks related to visual and motor regions, and higher frequencies have been seen to well reconstruct more complex RSNs [23]. The research we present here extends these

concepts by using GSP not only for investigating the underlying network as a whole, or individual frequencies on the underlying network, but to study how different ranges of frequencies capture significant information about the signal activation patterns, and applying this to dynamic, task-associated signals. These previous findings give high motivation for such an investigation.

The goal of this paper is to introduce GSP notions that can be used to analyze brain signals and to demonstrate their value in identifying patterns that appear when monitoring activity as subjects learn to perform a visual-motor task. Specifically, the contributions of this paper are: (i) To explain tools from the emerging field of GSP and show how they can be applied in analyzing brain signals. (ii) To evaluate the graph spectrum of brain functional network and to define artificial network construction methods that replicate the features of the graph spectrum of functional networks. (iii) To examine the temporal variation of brain signals corresponding to different graph frequencies when participants perform visual-motor learning tasks. (iv) To investigate the contribution of brain signals associated with different graph frequencies to the learning success at different stages of visual-motor learning.

We begin the paper with the introduction of basic notions of graphs and graph signals. Particular emphasis goes into the definition of the graph Fourier transform and the interpretation of graph frequency components as different modes of spatial variability measured with respect to the brain network (Section II-A). We also introduce the notion of graph filters and discuss the interpretation of graph low-pass, band-pass and high-pass filters as a local averaging operation (Section II-B). We point out that the discussion here is more extensive than necessary for readers familiar with GSP so that the paper is accessible to readers that are not necessarily familiar with the subject.

We then move on to describe two different experiments involving the learning of different visual-motor tasks by different sets of participants (Section III). We visualize the decomposed graph frequencies relating to the functional brain network (Section IV). We find that high graph frequencies of functional networks concentrate on visual and sensorimotor modules of the brain—the two brain areas well-known to be associated with motor learning [24], [25]. This motivates us to consider graph frequencies other than low frequency components, whereas the PCA-oriented approach has been focusing on low frequencies. We also describe the construction of a simple model to establish artificial networks with a few network descriptive parameters (Section IV-A). We observe that the model is able to mimic the properties of actual functional brain networks and we use them to analyze spectral properties of the brain networks (Section IV-B). The paper then utilizes graph frequency decomposition to visualize and investigate brain activities with different levels of spatial variation (Section V). It is noticed that the decomposed signals associated to different graph frequencies exhibit different levels of temporal variation throughout learning (Section V-A). Finally, we also define learning capabilities of subjects, and examine the importance of brain frequencies at different task familiarity by evaluating their respective

correlation with learning performance at different task familiarities (Section VI). We find as learning progresses, we favor different levels of graph frequency components.

II. GRAPH SIGNAL PROCESSING

The interest of this paper is to study brain signals in which we are given a collection of measurements x_i associated with each cortical region out of n different brain regions. An example signal of this type is an fMRI reading in which x_i estimates the level of activity of brain region i . The collection of n measurements is henceforth grouped in the vector signal $\mathbf{x} = [x_1, x_2, \dots, x_n]^T \in \mathbb{R}^n$. A fundamental feature of the signal \mathbf{x} is the existence of an underlying pattern of structural or functional connectivity that couples the values of the signal \mathbf{x} at different brain regions. Irrespective of whether connectivity is functional or structural, our goal here is to describe tools that utilize this underlying brain network to analyze patterns in the neurophysiological signal \mathbf{x} .

We do so by modeling connectivity between brain regions with a network that is connected, weighted, and symmetric. Formally, we define a network as the pair $\mathcal{G} = (\mathcal{V}, \mathbf{W})$, where $\mathcal{V} = \{1, \dots, n\}$ is a set of n vertices or nodes representing individual brain regions and $\mathbf{W} \in \mathbb{R}^{n \times n}$ represents weights of edges in the network with $w_{ij} \geq 0$ being the weight of the edge (i, j) , in which $i, j \in \mathcal{V}$. Since the network is undirected and symmetric we have that $w_{ij} = w_{ji}$ for all (i, j) . The weights $w_{ij} = w_{ji}$ represent the strength of the connection between regions i and j , or, equivalently, the proximity or similarity between nodes i and j . In terms of the signal \mathbf{x} , this means that when the weight w_{ij} is large, the signal values x_i and x_j tend to be related. Conversely, when the weight w_{ij} is small, the signal values x_i and x_j are not directly related except for what is implied by their separate connections to other nodes.

We adopt the conventional definitions of the degree and Laplacian matrices [26, Ch. 1]. The degree matrix $\mathbf{D} \in \mathbb{R}^{n \times n}$ is a diagonal matrix with its i th diagonal element $D_{ii} = \sum_{j=1}^n w_{ij}$. The Laplacian matrix is defined as the difference $\mathbf{L} := \mathbf{D} - \mathbf{W} \in \mathbb{R}^{n \times n}$. The components of the Laplacian matrix are explicitly given by $L_{ij} = -w_{ij}$ and $L_{ii} = \sum_{j=1}^n w_{ij}$. Observe that the Laplacian is real, symmetric, diagonal dominant, and with strictly positive diagonal elements. As such, the matrix \mathbf{L} is positive semidefinite. The eigenvector decomposition of \mathbf{L} is utilized in the following section to define the graph Fourier transform and the associated notion of graph frequencies.

We note that brain networks, irrespective of whether their connectivity is functional [27] or structural [28], tend to be stable for a window of time, entailing associations between brain regions during captured time of interest. Brain activities can vary more frequently, forming multiple samples of brain signals supported on a common underlying network.

A. Graph Fourier Transform and Graph Frequencies

Given that the graph Laplacian \mathbf{L} is real symmetric, it can be decomposed into its eigenvalue components,

$$\mathbf{L} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^H, \quad (1)$$

such that for the set of eigenvalues $\{\lambda_k\}_{k=0,1,\dots,n-1}$, the diagonal eigenvalue matrix is defined as $\mathbf{\Lambda} := \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{n-1})$, and

$$\mathbf{V} := [\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{n-1}] \quad (2)$$

is the eigenvector matrix. \mathbf{V}^H represents the Hermitian (conjugate transpose) of the matrix \mathbf{V} . We assume the set of eigenvalues of the Laplacian \mathbf{L} are ordered so that $0 = \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_{n-1}$. The validity of (1) follows because the eigenvectors of symmetric matrices are orthogonal so that the definition in (2) implies that $\mathbf{V}^H \mathbf{V} = \mathbf{I}$. The eigenvector matrix \mathbf{V} is used to define the Graph Fourier Transform of the graph signal \mathbf{x} as we formally state next; see, e.g., [14].

Definition 1: Given a signal $\mathbf{x} \in \mathbb{R}^n$ and a graph Laplacian $\mathbf{L} \in \mathbb{R}^{n \times n}$ accepting the decomposition in (1), the Graph Fourier Transform (GFT) of \mathbf{x} with respect to \mathbf{L} is the signal

$$\tilde{\mathbf{x}} := \mathbf{V}^H \mathbf{x}. \quad (3)$$

The inverse (i)GFT of $\tilde{\mathbf{x}}$ with respect to \mathbf{L} is defined as

$$\mathbf{x} := \mathbf{V} \tilde{\mathbf{x}}. \quad (4)$$

We say that \mathbf{x} and $\tilde{\mathbf{x}}$ form a GFT pair.

Observe that since $\mathbf{V}\mathbf{V}^H = \mathbf{I}$, the iGFT is, indeed, the inverse of the GFT. Given a signal \mathbf{x} we can compute the GFT as per (3). Given the transform $\tilde{\mathbf{x}}$ we can recover the original signal \mathbf{x} through the iGFT transform in (4).

There are several reasons that justify the association of the GFT with the Fourier transform. Mathematically, it is just a matter of definition that if the vectors \mathbf{v}_k in (1) are of the form $\mathbf{v}_k = [1, e^{j2\pi k/n}, \dots, e^{j2\pi k(n-1)/n}]^T$, the GFT and iGFT in Definition 1 reduce to the conventional time domain Fourier and inverse Fourier transforms. More deeply, if the graph \mathcal{G} is a cycle, the vectors \mathbf{v}_k in (1) are of the form $\mathbf{v}_k = [1, e^{j2\pi k/n}, \dots, e^{j2\pi k(n-1)/n}]^T$. Since cycle graphs are representations of discrete periodic signals, it follows that the GFT of a time signal is equivalent to the conventional discrete Fourier transform; see, e.g., [29].

An important property of the GFT is that it encodes a notion of variability akin to the notion of variability that the Fourier transform encodes for temporal signals. To see this, define $\tilde{\mathbf{x}} := [\tilde{x}_0, \dots, \tilde{x}_{n-1}]^T$ and expand the matrix product in (4) to express the original signal \mathbf{x} as

$$\mathbf{x} = \sum_{k=0}^{n-1} \tilde{x}_k \mathbf{v}_k. \quad (5)$$

It follows from (5) that the iGFT allows us to write the signal \mathbf{x} as a sum of orthogonal components \mathbf{v}_k in which the contribution of \mathbf{v}_k to the signal \mathbf{x} is the GFT component \tilde{x}_k . In conventional Fourier analysis, the eigenvectors $\mathbf{v}_k = [1, e^{j2\pi k/n}, \dots, e^{j2\pi k(n-1)/n}]^T$ carry a specific notion of variability encoded in the notion of frequency. When k is close to zero, the corresponding complex exponential eigenvectors are smooth. When k is close to n , the eigenfunctions fluctuate more rapidly in the discrete temporal domain. In the graph setting, the graph Laplacian eigenvectors provide a similar notion of frequency. Indeed, define the total variability of the graph signal

\mathbf{x} with respect to the Laplacian \mathbf{L} as

$$\text{TV}(\mathbf{x}) = \mathbf{x}^H \mathbf{L} \mathbf{x} = \sum_{i \neq j} w_{ij} (x_i - x_j)^2, \quad (6)$$

where in the second equality we expanded the quadratic form. It follows that the total variation $\text{TV}(\mathbf{x})$ is a measure of how much the signal changes with respect to the network. For the edge (i, j) , when w_{ij} is large we expect the values x_i and x_j to be similar because a large weight w_{ij} is encoding functional similarity between brain regions i and j . The contribution of their difference $(x_i - x_j)^2$ to the total variation is amplified by the weight w_{ij} . If the weight w_{ij} is small, activities at brain regions i and j tend to be uncorrelated, and therefore the difference between the signal values x_i and x_j makes little contribution to the total variation. We can then think of a signal with small total variation as one that changes slowly over the graph and of signals with large total variation as those that change rapidly over the graph.

Consider the total variation of the eigenvectors \mathbf{v}_k and use the facts that $\mathbf{L}\mathbf{v}_k = \lambda_k \mathbf{v}_k$ and $\mathbf{v}_k^H \mathbf{v}_k = 1$ to conclude that

$$\text{TV}(\mathbf{v}_k) = \mathbf{v}_k^H \mathbf{L} \mathbf{v}_k = \lambda_k. \quad (7)$$

It follows from (7) and the fact that the eigenvalues are ordered as $0 = \lambda_0 \leq \dots \leq \lambda_{n-1}$, that the total variations of the eigenvectors \mathbf{v}_k follow the same order. Combining this observation with the discussion following (6), we conclude that when k is close to 0, the eigenvectors \mathbf{v}_k vary slowly over the graph, whereas for k close to n the eigenvalues vary more rapidly. Therefore, from (5) we see that the GFT and iGFT allow us to decompose the brain signal \mathbf{x} into components that characterize different levels of variability. The GFT coefficients \tilde{x}_k for small values of k indicate how much these slowly varying signals contribute to the observed brain signal \mathbf{x} . On the other hand, the GFT coefficients \tilde{x}_k for large values of k describe how much rapidly varying signals contribute to the observed brain signal \mathbf{x} .

B. Graph Filtering and Frequency Decompositions

Given a graph signal \mathbf{x} with GFT $\tilde{\mathbf{x}}$ we can isolate the frequency components corresponding to the lowest K_L graph frequencies by defining the filtered spectrum $\tilde{\mathbf{x}}_L := \tilde{\mathbf{H}}_L \tilde{\mathbf{x}}$ satisfying $\tilde{x}_{Lk} = \tilde{x}_k$ for $k < K_L$ and $\tilde{x}_{Lk} = 0$ otherwise. The filter $\tilde{\mathbf{H}}_L$ can be written as the diagonal matrix $\tilde{\mathbf{H}}_L := \text{diag}(\tilde{\mathbf{h}}_L)$ where the vector $\tilde{\mathbf{h}}_L$ takes value 1 for frequencies smaller than K_L and is otherwise null,

$$\tilde{h}_{Lk} = \mathbb{I}[k < K_L]. \quad (8)$$

Utilizing the definitions of the GFT in (3) and the iGFT in (4), the spectral operation $\tilde{\mathbf{x}}_L = \tilde{\mathbf{H}}_L \tilde{\mathbf{x}}$ is equivalent to performing the following operations in the graph vertex domain

$$\mathbf{x}_L = \mathbf{V} \tilde{\mathbf{x}}_L = \mathbf{V} \tilde{\mathbf{H}}_L \tilde{\mathbf{x}} = \mathbf{V} \tilde{\mathbf{H}}_L \mathbf{V}^{-1} \mathbf{x} =: \mathbf{H}_L \mathbf{x}. \quad (9)$$

From the equality in (9), we can see that the signal \mathbf{x}_L contains the low graph frequency components of \mathbf{x} , and so we say the matrix \mathbf{H}_L is a *graph low-pass filter*.

The filter $\mathbf{H}_L := \mathbf{V} \tilde{\mathbf{H}}_L \mathbf{V}^{-1}$ admits an alternative representation as the expansion $\mathbf{H}_L = \sum_{k=0}^{K_L-1} h_{Lk} \mathbf{L}^k$ in terms of Laplacian

powers [29]. The coefficients h_{Lk} in this expansion are elements of the vector $\mathbf{h}_L := \Psi^{-1} \tilde{\mathbf{h}}_L$ where Ψ is the Vandermonde matrix defined by the eigenvalues of \mathbf{L} , i.e.,

$$\Psi := \begin{bmatrix} 1 & \lambda_0 & \cdots & \lambda_0^{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \lambda_{n-1} & \cdots & \lambda_{n-1}^{n-1} \end{bmatrix}. \quad (10)$$

Since the eigenvalues are ordered in (10), the coefficients h_{Lk} tend to be concentrated in small indexes k , and the expansion $\mathbf{H}_L = \sum_{k=0}^{n-1} h_{Lk} \mathbf{L}^k$ is therefore dominated by small powers \mathbf{L}^k . From this fact it follows that we can think of the graph low-pass filtered signal \mathbf{x}_L as resulting from a localized averaging of the elements of \mathbf{x} . To understand this interpretation, simply note that $\mathbf{L}^0 \mathbf{x} = \mathbf{x}$ coincides with the original signal, $\mathbf{L} \mathbf{x}$ is an average of neighboring elements, $\mathbf{L}^2 \mathbf{x}$ is an average of elements in nodes that interact via intermediate common neighbors, and, in general, $\mathbf{L}^k \mathbf{x}$ describes interactions between k -hop neighbors. The fact that \mathbf{x}_L can be considered as a signal that follows from local averaging of \mathbf{x} implies that \mathbf{x}_L has smaller total variation than \mathbf{x} and is consistent with the interpretation of low graph frequencies presented in Section II-A. We point out that the definition $\mathbf{h}_L = \Psi^{-1} \tilde{\mathbf{h}}_L$ assumes the inverse matrix Ψ^{-1} exists. This holds true if the graph Laplacian does not have duplicate eigenvalues, which is the case for all functional brain networks examined in the paper.

Other types of graph filters can be defined analogously to study interactions between signal components other than the local interactions captured in \mathbf{x}_L . Apart from the graph low-pass filter \mathbf{H}_L , we also consider a graph band-pass filter \mathbf{H}_M and a graph high-pass filter \mathbf{H}_H , whose graph frequency responses are defined as

$$\tilde{h}_{Mk} = \mathbb{I}[K_L \leq k < K_L + K_M], \quad (11)$$

$$\tilde{h}_{Hk} = \mathbb{I}[K_L + K_M \leq k]. \quad (12)$$

The definitions in (8), (11), and (12) are such that the low-pass filter takes the lowest K_L graph frequencies, the band-pass filter captures the middle K_M graph frequencies, and the high-pass filter the highest $n - K_L - K_M$ frequencies. The three filters are defined such that the graph frequencies of their respective interest are mutually exclusive yet collectively exhaustive. As a result, if we use $\mathbf{x}_M := \mathbf{H}_M \mathbf{x}$ and $\mathbf{x}_H := \mathbf{H}_H \mathbf{x}$ to respectively denote the signals filtered by the band-pass and high-pass filters, we have that the original signal can be written as the sum $\mathbf{x} = \mathbf{x}_L + \mathbf{x}_M + \mathbf{x}_H$. This gives a decomposition of \mathbf{x} into low, medium, and high frequency components which respectively represent signals that have slow, medium, and high variability with respect to the connectivity network between brain regions. This decomposition is utilized in this paper to analyze brain activity patterns associated with the learning of visual-motor tasks.

III. BRAIN SIGNALS DURING LEARNING

We considered two experiments in which subjects learned a simple motor task [30]–[32]. In the experiments, forty-seven

	Session 1	Session 2	Session 3	Session 4
MIN Sequences	50	110	170	230
MOD Sequences	50	200	350	500
EXT Sequences	50	740	1430	2120

Fig. 1. Relationship between training duration, intensity, and depth for the first experimental framework. The values in the table denote the number of trials (i.e., “depth”) of each sequence type (i.e., “intensity”) completed after each scanning session (i.e., “duration”) averaged over the 20 participants.

right-handed participants (29 female, 18 male; mean age 24.13 years) volunteered with informed consent in accordance with the University of California, Santa Barbara Internal Review Board. After exclusions for task accuracy, incomplete scans, and abnormal MRI, 38 participants were retained for subsequent analysis.

Twenty individuals participated in the first experimental framework. The experiment lasted 6 weeks, in which there were 4 scanning sessions, roughly at the start of the experiment, at the end of the 2nd week, at the end of the 4th week, and at the end of the experiment, respectively. During each scanning session, individuals performed a discrete sequence-production task in which they responded to sequentially presented stimuli with their dominant hand on a custom response box. Sequences were presented using a horizontal array of 5 square stimuli with the responses mapped from left to right such that the thumb corresponded to the leftmost stimulus. The next square in the sequence was highlighted immediately following each correct key press; the sequence was paused awaiting the depression of the appropriate key if an incorrect key was pressed. Each participant completed 6 different 10-element sequences. Each sequence consists of two squares per key. Participants performed the same sequences at home between each two adjacent scanning sessions, however, with different levels of exposure for different sequence types. Therefore, the number of trials completed by the participants after the end of each scanning session depends on the sequence type. There are 3 different sequence types (MIN, MOD, EXT) with 2 sequences per type. The number of trials of each sequence type completed after each scanning session averaged over the 20 participants is summarized in Fig. 1. During scanning sessions, each scan epoch involved 60 trials, 20 trials for each sequence type. Each scanning session contained a total of 300 trials (5 scan epochs) and a variable number of brain scans depending on how quickly the task was performed by the specific individual.

Eighteen subjects participated in the second experimental framework. The experiment had 3 scanning sessions spanning the three days. Each scanning session lasted roughly 2 hours and no training was performed at home between adjacent scanning sessions. Subjects responded to a visually cued sequence by generating responses using the four fingers of their nondominant hand on a custom response box. Visual cues were presented as a series of musical notes on a pseudo-musical staff with four lines such that the top line of the staff mapped to the leftmost key pressed with the pinkie finger. Each 12-note sequence randomly ordered contained three notes per line. Each training epoch involved 40 trials and lasted a total of 245 repetition times (TRs),

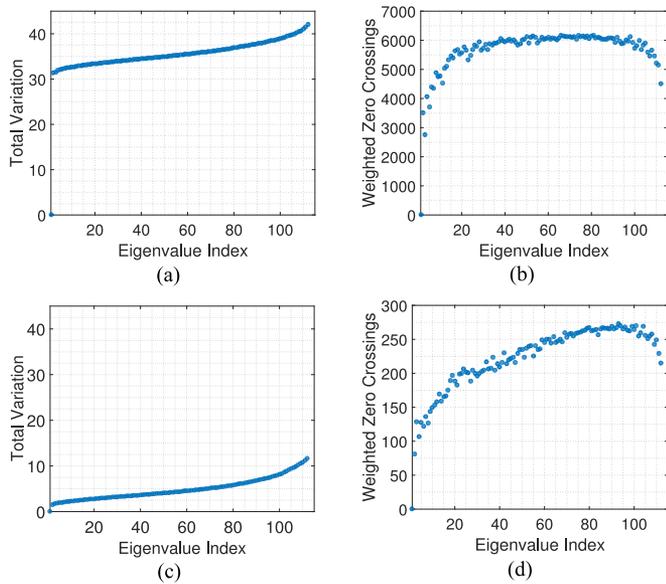


Fig. 2. (a) Total variation $TV(\mathbf{v}_k)$ and (b) weighted zero crossings $ZC(\mathbf{v}_k)$ of the graph Laplacian eigenvectors for the brain networks averaged across participants in the 6 week training experiment. (c) and (d) present the values for the 3 day experiment. In both cases, the Laplacian eigenvectors associated with larger indexes vary more on the network and cross zero relatively more often, confirming the interpretation of the Laplacian eigenvalues as notions of frequencies. Besides, note that total variation increases relatively linearly with indexes.

with a TR of 2,000 ms. Each training session contained 6 scan epochs (240 trials) and lasted a total of 2,070 scan TRs.

In both experiments participants were instructed to respond promptly and accurately. Repetitions (e.g., “11”) and regularities such as trills (e.g., “121”) and runs (e.g., “123”) were excluded in all sequences. The order and number of sequence trials were identical for all participants. Participants completed the tasks inside the MRI scanner for scanning sessions.

Reordering with fMRI was conducted using a 3.0 T Siemens Trio with a 12-channel phased-array head coil. For each functional run, a single-shot echo planar imaging sequence that is sensitive to blood oxygen level dependent (BOLD) contrast was utilized to obtain 37 (the first experiment) or 33 (the second experiment) slices (3 mm thickness) per repetition time (TR), an echo time of 30 ms, a flip angle of 90° , a field of view of 192 mm, and a 64×64 acquisition matrix. Image preprocessing was performed using the Oxford Center for Functional Magnetic Resonance Imaging of the Brain (FMRIB) Software Library (FSL), and motion correction was performed using FMRIB’s linear image registration tool. The whole brain is parcellated into a set of $n = 112$ regions of interest that correspond to the 112 cortical and subcortical structures anatomically identified in FSL’s Harvard-Oxford atlas. The choice of parcellation scheme is the topic of several studies in resting-state [33], and task-based [34] network architecture. The question of the most appropriate delineation of the brain into nodes of a network is open and is guided by the particular question one wants to ask. We use Harvard-Oxford atlas here because it is consistent with previous studies of task-based functional connectivity during

learning [30], [31]. The threshold in probability cutoff settings of Harvard Oxford atlas parcellation is 0 so that no voxels were excluded.

For each individual fMRI dataset, we estimate regional mean BOLD time series by averaging voxel time series in each of the n regions. We evaluate the magnitude squared spectral coherence [35] between the activity of all possible pairs of regions to construct $n \times n$ functional connectivity matrices \mathbf{W} . Besides, for each pair of brain regions i and j , we use t -statistical testing to evaluate the probability $p_{i,j}$ of observing the measurements by random chance, when the actual data are uncorrelated [36]. In the 3 day dataset, the value of all elements with no statistical significance ($p_{i,j} > 0.05$) [37] are set to zero; the values remain unchanged otherwise. In the 3 day experiment, a single brain network is constructed for each participant. Thresholding is applied because the networks are for the entire span of the experiment and many entries in \mathbf{W} would be close to zero without threshold correction. In the 6 week experiment, due to the long duration of the experiment, we build a different brain network per scanning session, per sequence type for each subject. Because each network describes the functional connectivity for one training session given a subject, not many entries will be removed even in the presence of threshold correction; consequently, no thresholding is applied for the 6 week dataset. We normalize the regional mean BOLD observations $\hat{\mathbf{x}}(t)$ at any sample time t and consider $\mathbf{x}(t) = \hat{\mathbf{x}}(t)/\|\hat{\mathbf{x}}(t)\|_2$ such that the total energy of activities at all structures is consistent at different t to avoid extreme spikes due to head motion or drift artifacts in fMRI.

IV. BRAIN NETWORK FREQUENCIES

In this section, we analyze the graph spectrum brain networks of the dataset considered. For the brain network \mathbf{W} of each subject, we construct its Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{W}$, and evaluate the total variation $TV(\mathbf{v}_k)$ [cf. (7)] for each eigenvector \mathbf{v}_k . Fig. 2(a) and (c) plot the total variation of all graph eigenvectors averaged across participants of the 6 week training experiment and 3 day experiment, respectively. In both experiments, the Laplacian eigenvectors associated with larger indexes fluctuate more on the network. Another observation is that with respect to graph frequency indices $0 < k < 100$, the total variation increases almost linearly.

Besides total variation, the number of zero crossings is used as a measure of the smoothness of signals with respect to an underlying network [14]. Since brain networks are weighted, we adapt a slightly modified version—weighted zero crossings—to investigate the given graph eigenvector \mathbf{v}_k

$$ZC(\mathbf{v}_k) = \frac{1}{2} \sum_{i \neq j} w_{ij} \mathbb{I} \{v_k(i)v_k(j) < 0\}. \quad (13)$$

In words, weighted zero crossings evaluate the weighted sum of the set of edges connecting a vertex with a positive signal to a vertex with a negative signal. Fig. 2(b) and (d) demonstrate the weighted zero crossings of all graph eigenvectors averaged across subjects of the 6 week and 3 day experiments, respectively. The weighted zero crossings increase almost

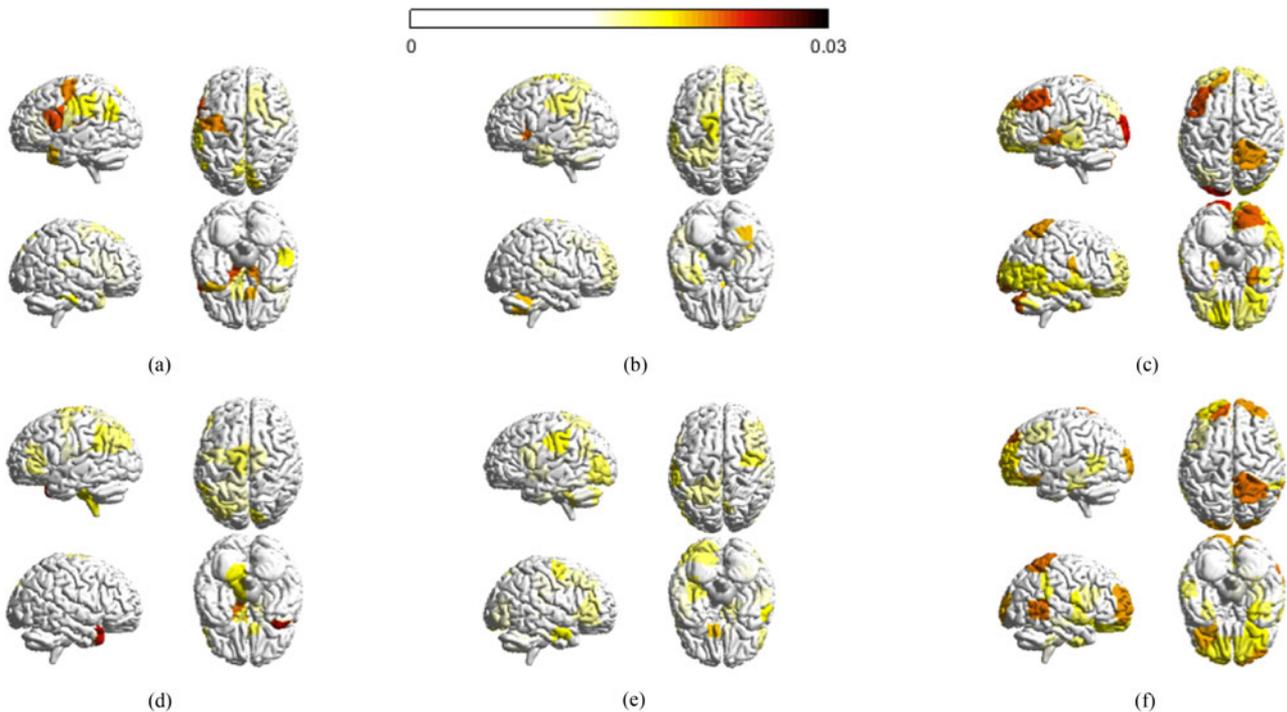


Fig. 3. Absolute magnitude at each of the n cortical structures averaged across participants in the 6 week experiment and averaged across all frequency components in (a) the set of low graph frequencies $\{\mathbf{v}_k\}_{k=0}^{K_L-1}$, (b) the set of middle graph frequencies $\{\mathbf{v}_k\}_{k=K_L}^{K_L+K_M-1}$, and (c) the set of high graph frequencies $\{\mathbf{v}_k\}_{k=K_L+K_M}^{n-1}$. (d)–(f) presents the average absolute magnitudes for the 3 day experiment. Only brain regions with absolute magnitudes higher than a fixed threshold (0.015) are colored. The magnitudes at different brain regions across the datasets are significantly similar in the low and high graph frequencies (correlation coefficients 0.5818 and 0.6616, respectively). The brain regions with high magnitude values significantly overlap with the visual and sensorimotor modules, in which more than 60% of values greater than the threshold belong to the visual and sensorimotor modules.

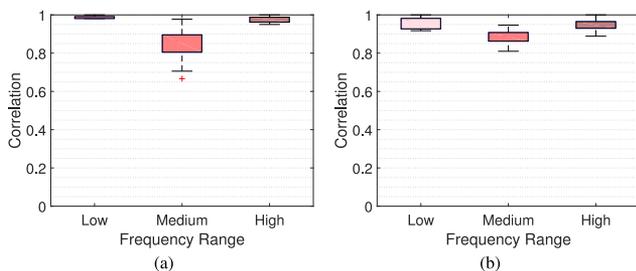


Fig. 4. Box plots demonstrating the robustness of parameters chosen for different frequency ranges for the absolute magnitudes across brain regions for (a) 6 week and (b) 3 day experiments. Each box plot presents the correlation coefficients between the frequency range selected for this paper and all frequency ranges for K_L and K_M between 32 and 42, inclusive.

proportionally with graph frequency index k until they eventually level off for $0 \leq k \leq 100$. For k greater than 100, though, eigenvectors associated with higher graph frequencies exhibit lower weighted zero crossings.

It would be interesting to examine where the associated eigenvectors lie anatomically, and the relative strength of their values. To facilitate the presentation, we consider three sets of eigenvectors, $\{\mathbf{v}_k\}_{k=0}^{K_L-1}$, $\{\mathbf{v}_k\}_{k=K_L}^{K_L+K_M-1}$ and $\{\mathbf{v}_k\}_{k=K_L+K_M}^{n-1}$, and compute the absolute magnitude at each of the n cortical and subcortical regions averaged across participants and across all graph frequencies belonging to each of the three sets. Fig. 3 presents the average magnitudes for the two experimental frameworks considered in the paper using BrainNet [38],

where brain regions with absolute magnitudes lower than a fixed threshold are not colored. Throughout the paper, the parameter K_L is set as 40 and K_M is set as 32. This combination yields three roughly equally-sized components with one piece corresponding to the 40 lowest graph frequencies and another piece corresponding to the 40 highest frequencies. The results presented in the paper are robust with the choice of parameters: we examined the results for K_L and K_M in the range of 32 to 42, inclusive, and found similar observations as the ones presented. To demonstrate that, Fig. 4 presents the range of correlation coefficients calculated between the frequency ranges selected for this paper and all frequency ranges for K_L and K_M between 32 and 42, inclusive, giving 120 correlation values for each box plot. The correlation coefficients reported are a quantification of similarity measure when we examine the similarity between two vectors, given as the absolute magnitudes passing the given threshold across all brain regions. An investigation of cosine similarity gives high similarities as well. We have also conducted robustness testing in our analysis of learning rate (Section VI) and have plotted our results obtained from using all 121 possible frequency ranges (Fig. 12) and have quantified the robustness of parameters in Fig. 11.

A. Artificial Functional Brain Networks

An approach to analyze the complex networks is to define a model to generate artificial networks [39], [40]. The main motivation of an artificial network model is to use them to analyze

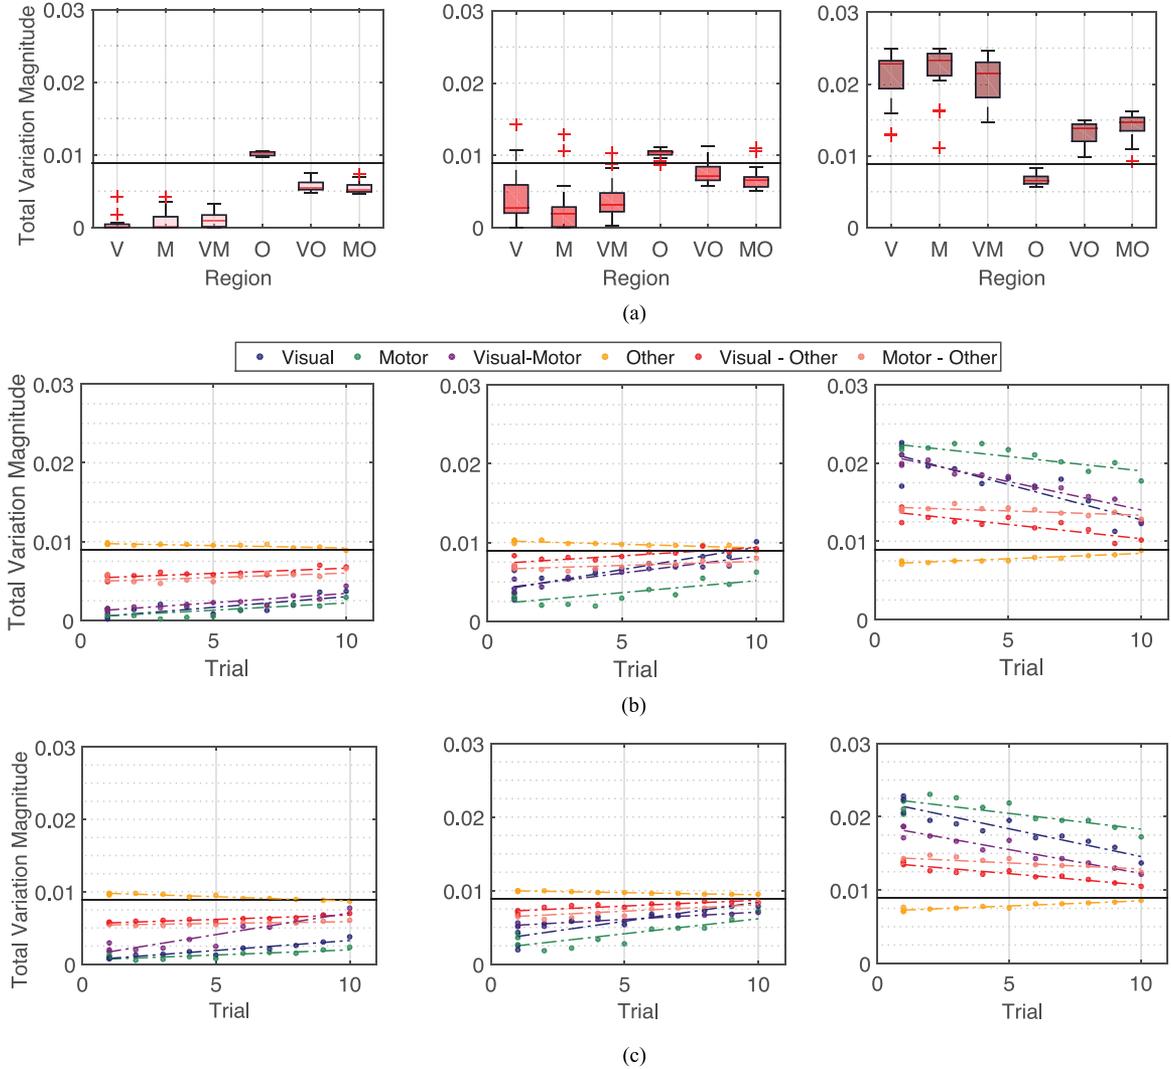


Fig. 5. Spectral property of brain networks in the 6 week experiment. (a) Left: Averaged total variation of eigenvectors \mathbf{v}_k for 6 different types of connections of the brain averaged across all eigenvectors associated with low graph frequencies $\mathbf{v}_k \in \{\mathbf{v}_k\}_{k=0}^{K_L-1}$, across all participants and scan sessions. Middle: Across all eigenvectors associated with mid-range graph frequencies $\mathbf{v}_k \in \{\mathbf{v}_k\}_{k=K_L}^{K_L+K_M-1}$. Right: Across all eigenvectors with high graph frequencies $\mathbf{v}_k \in \{\mathbf{v}_k\}_{k=K_L+K_M}^{n-1}$. (b) Median total variations of brain networks across participants of different scanning sessions and different sequence types with respect to the level of exposure of participants to the sequence type at the scanning session. Relationship between training duration, intensity, and depth is summarized in Fig. 1. Value of 1 on the x -axis in the figures refers to minimum exposure to sequences (all 3 sequence types of the first session), and value of 10 on the x -axis denotes the maximum exposure to sequences (EXT sequence types of the fourth session). An association between spectral property of brain networks and the level of exposure is clearly observed (average correlation coefficient 0.8164). (c) Median total variations evaluated upon artificial networks. Spectral properties of actual brain networks can be closely simulated using a few parameters. The main text gives all correlation values for similarity between variance among subjects and between correlations of training intensity.

complex brain networks. Examples of such models include the Barabási-Albert model for scale free networks [41] and recent developments and insights on weighted network models [40]. Here we present a framework to construct artificial networks that can be used to mimic the functional brain networks with only a few parameter inputs. The model is related to weighted block stochastic model [42], but involves more aspects like individual variance and analyzes links independent of their connectivity strength to other brain regions. The output of the method would be a symmetric network with edge weights between 0 and 1 without self-loops.

To begin, suppose the desired network has two clusters of nodes \mathcal{V}_1 and \mathcal{V}_o . The algorithm requires the average edge

weight μ_1 for connections between nodes of the first cluster \mathcal{V}_1 , average edge weight μ_o for links between nodes of the other cluster \mathcal{V}_o , and average edge weight μ_{1o} for inter-cluster connections. To reflect the fact that the edge weights on some links are independent of their joining vertices, for each edge within \mathcal{V}_1 , with probability $p_e < 1$, its weight is randomly generated with respect to uniform distribution $\mathcal{U}[0, 1]$ between 0 and 1, and with probability $1 - p_e$, its weight is randomly generated with respect to uniform distribution $\mathcal{U}[\mu_1 - \delta, \mu_1 + \delta]$. The parameter p_e determines the percentage of edges whose weights are selected irrespective of their actual locations. To further simulate the observation that different participants may possess distinctive brain networks, if the edge weight is randomly generated from

a uniform distribution $\mathcal{U}[\mu_1 - \delta, \mu_1 + \delta]$, it is then perturbed by $w_u \sim \mathcal{U}[-u_\epsilon/2, u_\epsilon/2]$ where u_ϵ controls the level of perturbation. The edge weights for connections within cluster \mathcal{U}_o are generated similarly: with probability $1 - p_\epsilon$, the edge weight is randomly chosen from the uniform distribution $\mathcal{U}[\mu_o - \delta, \mu_o + \delta]$ before being contaminated by $w_u \sim \mathcal{U}[-u_\epsilon/2, u_\epsilon/2]$. The edge weights for connections between clusters \mathcal{U}_1 and \mathcal{U}_o are formed analogously using μ_{1o} . The method presented here can be easily generalized to analyze brain networks with more regions of interest, i.e. by specifying sets of regions of interest and by detailing the expected correlation values on each type of connection between different regions.

Remark 1: At one extreme we can make each node i belonging to a different set $\mathcal{V}_i = \{i\}$. Then the method requires the inputs of expected weights for all nodes, or alternatively speaking, the expected network. At the other extreme, there is only one set of nodes \mathcal{V} , and then the method is highly akin to a network with edge weights completely randomly generated. Any construction of interest would have some prior knowledge regarding the community structure. Therefore, the method proposed here can be used to see if the network constructed with the specific choice of community structure highly simulate the key properties of the actual network, and can be used to examine the evolution of community structure in the brain throughout the process to master a particular task.

B. Spectral Properties of Brain Networks

In this section, we analyze graph spectral properties of brain networks. Given the graph Laplacian, we examine the fluctuation of its eigenvectors on different types of connections in the brain network [25]. More specifically, given an eigenvector \mathbf{v}_k , its variation on the visual module is defined as

$$\text{TV}_{\text{visual}}(\mathbf{v}_k) = \frac{\sum_{i,j \in \mathcal{V}_v, i \neq j} w_{ij} (v_k(i) - v_k(j))^2}{\sum_{i,j \in \mathcal{V}_v, i \neq j} w_{ij}}, \quad (14)$$

where \mathcal{V}_v denotes the set of nodes belonging to the visual module. The measure $\text{TV}_{\text{visual}}(\mathbf{v}_k)$ computes the difference for signals on the visual module for each unit of edge weight. To facilitate interpretation, we only consider three sets of eigenvectors $\{\mathbf{v}_k\}_{k=0}^{K_L-1}$, $\{\mathbf{v}_k\}_{k=K_L}^{K_L+K_M-1}$, and $\{\mathbf{v}_k\}_{k=K_L+K_M}^{n-1}$. We then compute the visual module total variation $\text{TV}_{\text{visual}}^L$ averaged over eigenvectors $\{\mathbf{v}_k\}_{k=0}^{K_L-1}$, and $\text{TV}_{\text{visual}}^M$ as well as $\text{TV}_{\text{visual}}^H$ similarly. Besides $\text{TV}_{\text{visual}}$, we also examine the level of fluctuation of eigenvectors on edges within the motor module, denoted as TV_{motor} , and on connections belonging to brain modules other than the visual and motor module $\text{TV}_{\text{others}}$. Further, there are links between two separate brain modules, and to assess the variation of eigenvectors on those links, we define total variations between the visual and motor modules

$$\text{TV}_{\text{visual-motor}}(\mathbf{v}_k) = \frac{\sum_{i \in \mathcal{V}_v, j \in \mathcal{V}_m} w_{ij} (v_k(i) - v_k(j))^2}{\sum_{i \in \mathcal{V}_v, j \in \mathcal{V}_m} w_{ij}}, \quad (15)$$

where \mathcal{V}_m denotes the set of nodes belonging to the motor module. Total variations $\text{TV}_{\text{visual-others}}$ between the visual and other modules, and total variations $\text{TV}_{\text{motor-others}}$ between the motor and other modules are defined analogously. We chose to study

visual and motor modules separately from other brain modules because of their well-known associations with motor learning [24], [25].

Fig. 5(a) presents boxplots of the variation for eigenvectors of different graph frequencies measured over different types of connections across participants, at the start of the six week training. Despite that total variation of eigenvectors should increase with their frequencies, the variation on the other module $\text{TV}_{\text{others}}^L$ of eigenvectors associated with low frequencies are higher than $\text{TV}_{\text{others}}^H$ (pass t-test with $p < 0.0001$). This observation is discussed in detail in Section IV-C.

Next we study how the graph spectral properties of brain networks evolve as participants become more familiar with the tasks. Fig. 5(b) illustrates the median of the variation for eigenvectors of different graph frequencies measured over different types of connections across subjects, at 10 different levels of exposure in the six week training. As participants become more acquainted with the assignment, their brain networks display lower variation in the visual and motor modules and higher variation in the other modules for low and middle graph frequencies, and the exact opposite is true for high graph frequencies. The association with training intensity is statistically significant (average correlation coefficient $r = 0.8164$).

C. Discussion

Firstly, we examine why we see a decrease in zero crossings of graph frequencies when k is greater than 100 in Fig. 2. A detailed analysis shows this is because the functional brain networks are highly connected with nearly homogeneous degree distribution, and consequently each high graph frequency tends to have a value with high magnitude at one vertex of high degree and similar values at other nodes, resulting in a smaller global zero crossings for eigenvectors associated with very high frequencies.

Secondly, in terms of the visualization of graph frequencies in Fig. 3, the most interesting finding relates to the eigenvectors associated with high graph frequencies. The magnitudes at different brain regions for high frequencies are significantly similar across the two datasets investigated (correlation coefficient 0.6616). There are very few noticeable brain regions in which the absolute magnitude highlighted in the first dataset is not likewise highlighted in the second. Given the different experiment setups, it would not be uncommon to observe large variations across datasets. However, the fact that we see the majority of brain regions similarly highlighted in the two experiments solidifies our understanding that eigenvector decomposition captures general signatures, as opposed to task-specific realizations. Additionally, brain regions with high magnitude values are highly alike (greater than 60% overlap) to the visual and sensorimotor cortices [43]. This is likely to be a consequence of the fact that visual and motor regions are more strongly connected with other structures, and hence an eigenvector with a high magnitude on visual or motor structures would result in high global spatial variation. The eigenvectors of low graph frequencies are more spread across the networks, resulting in low global variations. The middle graph frequencies are less

interesting—the magnitudes at most regions (greater than 90%) do not pass the threshold, and little associations (correlation coefficient 0.3529) can be found between the eigenvectors of the 6 week and 3 day experiments.

Thirdly, to better interpret the meaning of variations for specific types of connections, we construct artificial networks as described in Section IV-A with visual and motor modules as regions of interest, and consider other modules to be brain regions other than visual and motor modules. We observe that there are three contributing factors that cause the variation within a specific module to become higher for higher eigenvectors and to become lower for lower eigenvectors: (i) Increases in the average edge weight for connections within the module, (ii) Increments in the average edge weight for links between this module to other module, and (iii) Escalation in the average edge weight for associations within the other module. This can also be observed by analyzing closely the definition of total variation. If a module is highly connected, in order for the eigenvector associated with a low graph frequency to be smooth on the entire network, it has to be smooth on the specific module, resulting in a low value in the variation of an eigenvector associated with a low graph frequency with respect to the module of interest. Similarly, the increase in the variation of connections between two modules, e.g. between visual and other modules are resulted from: (i) The growth in the average edge weight for connections between visual and other modules, or (ii) The augmentation of average weight for links within the other module. The graph spectral properties as in Fig. 5(a) are observed because (i) visual and motor modules are themselves highly connected, and (ii) visual module is also strongly linked with motor module.

Finally, in analyzing the evolution of graph spectral properties as participants become more familiar with the tasks, following the interpretations based on artificial network analysis, this evolution in graph spectral properties of brain networks is mainly caused by the decrease in values of connections within visual and motor modules and between the visual and motor modules. An interesting observation is that the values in the variation of eigenvectors associated with high frequencies decline with respect to the visual module much faster than that of motor module, even though the visual module is more strongly connected throughout training compared to the motor module. A deep analysis using artificial networks shows that this results from the following three factors: (i) Though more strongly connected compared to the motor module, connections within the visual module weaken very quickly, (ii) The motor module is more closely connected with the other module than the link between the visual module to the other module, and (iii) Association levels within the other module stay relatively constant. Therefore, as participants become more exposed to the tasks, compared to the visual module, the motor module becomes more strongly connected. The graph spectral properties of actual brain networks and their evolution can be closely imitated using artificial networks as plotted in Fig. 5(c). The artificial network created for our analysis best imitated the real brain networks with parameters p_ϵ of 0.10, u_ϵ of 0.10, and δ of 0.01. The average edge weights μ for visual (v), motor (m), other (o), and inter-connecting regions are $\mu_v = 0.6028$, $\mu_m = 0.4902$, $\mu_o = 0.3098$, $\mu_{vm} = 0.3985$,

$\mu_{vo} = 0.3181$, and $\mu_{mo} = 0.3271$. The correlation coefficients of association with training intensity between real and artificial networks for low, medium, and high graph frequencies are 0.6436, 0.7187 and 0.8457, respectively. Additionally, the variation among participants in real dataset can be closely mimicked using artificial network model we proposed, with correlation coefficients 0.9338, 0.9660, and 0.9486 for low, medium, and high graph frequencies, respectively. The analysis for the three day training dataset is highly similar (correlation coefficients 0.9834, 0.9186, and 0.9674 for low, medium, and high graph frequencies, respectively) and for this reason we do not present and analyze it separately here.

V. FREQUENCY DECOMPOSITION OF BRAIN SIGNALS

The previous sections focus on the study of brain networks and their graph spectral properties. In this section, we investigate brain signals from a GSP perspective, and analyze the brain signals by examining the decomposed graph signals \mathbf{x}_L , \mathbf{x}_M , and \mathbf{x}_H with respect to the underlying brain networks. We compute the absolute magnitude of the decomposed signal \mathbf{x}_L for each brain region averaged across all sample signals for each individual during a scan session and then averaged across all participants. Similar aggregation is applied for \mathbf{x}_M and \mathbf{x}_H .

Fig. 6 presents the distribution of the decomposed signals corresponding to different levels of spatial variations for the first scan session (top row) and the last scan session (bottom row) in the 6 week experiment. Fig. 7 exhibits how the decomposed signals are distributed across brain regions in the 3 day experiment. Brain regions with absolute magnitudes lower than a fixed threshold are not colored.

A. Temporal Variation of Graph Frequency Components

We analyze temporal variation of decomposed signals with respect to different levels of spatial variations. To this end, we evaluate the variance of the decomposed signals over multiple temporal scales—over days and minutes—for the two experiments. We describe the method specifically for \mathbf{x}_L for simplicity and similar computations were conducted for \mathbf{x}_M and \mathbf{x}_H . At the macro timescale, we average the decomposed signals \mathbf{x}_L for all sample points within each scanning session with different sequence type, and evaluate the variance of the magnitudes of the signals [15] across all the scanning sessions and sequence types. For the 6 week experiment, there are 4 scanning sessions and 3 different sequence types, so the variance is with respect to 12 points. For the 3 day experiment, there are 3 scanning sessions and only 1 sequence type, so the variance is for 3 points. As for the micro or minute-scale, we average the decomposed signals \mathbf{x}_L for all sample points within each minute, and evaluate the variance of the magnitudes of the averaged signals across all minute windows for each scanning session with different sequence types. The evaluated variance is then averaged across all participants of the experiment of interest.

Fig. 8 displays the variance of the decomposed signals \mathbf{x}_L , \mathbf{x}_M and \mathbf{x}_H at two different temporal scales of the two experiments. For the 6 week dataset, 3 session-sequence combinations, with the number proportional to the level of exposure of participants

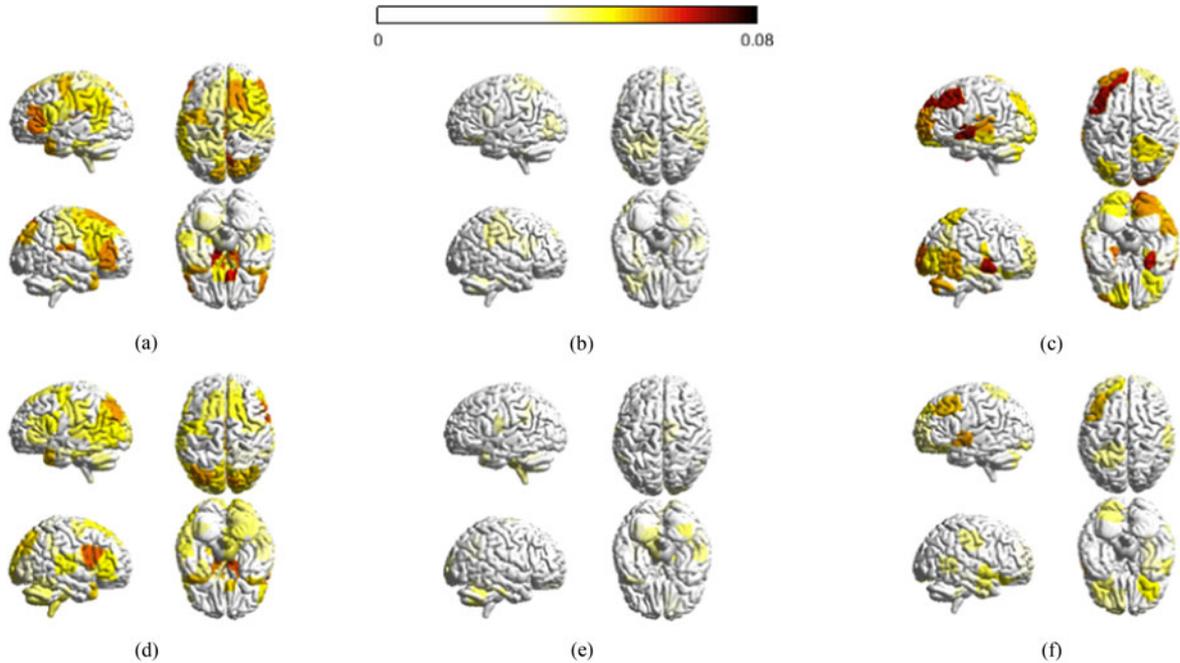


Fig. 6. Distribution of decomposed signals for the 6 week experiment. (a) Absolute magnitudes for all brain regions with respect to x_L —brain signals varying smoothly across the network—averaged across all sample points for each individual and across all participants at the first scan session of the 6 week dataset. (b) With respect to x_M and (c) with respect to x_H —signals rapidly fluctuating across the brain. (d), (e), and (f) are averaged x_L , x_M and x_H at the last scan session of the 6 week dataset, respectively. Only regions with absolute magnitudes higher than a fixed threshold are colored.

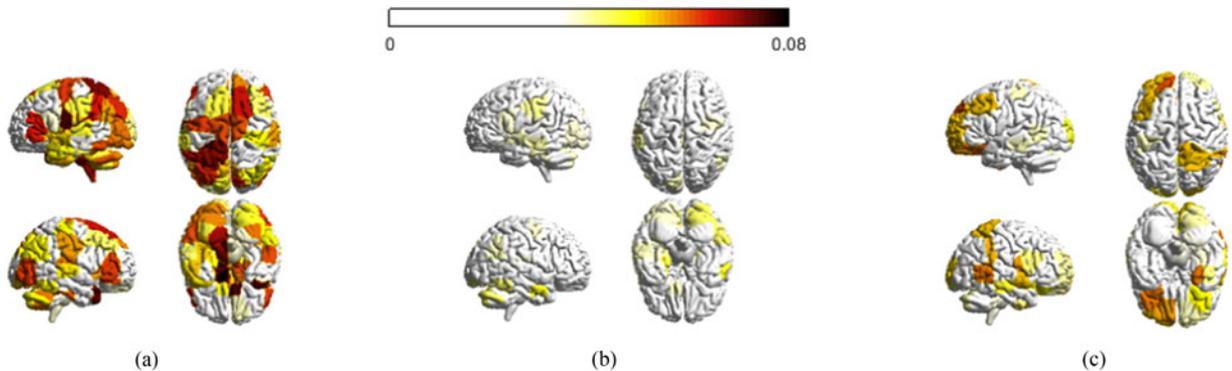


Fig. 7. Distribution of decomposed signals for the 3 day experiment. (a), (b), and (c) are the absolute magnitudes for all brain regions with respect to x_L , x_M and x_H , averaged across all sample points for each subject and across participants in the 3 day experiment, respectively. Regions with absolute value less than a threshold are not colored.

to the sequence (1-MIN refers to MIN sequence at session 1, 5 denotes MIN sequence at session 4, 9 entails EXT sequence at session 3) are selected out of the 12 combinations in total for a cleaner illustration, but all the other session-sequence combinations exhibit similar properties.

B. Discussion

A deep analysis of Figs. 6 and 7 yields many interesting aspects of graph frequency decomposition. First, for x_L , the magnitudes on adjacent brain regions tend to possess highly similar values, resulting in a more evenly spread brain signal distribution, where as for x_H , neighboring signals can exhibit highly dissimilar values; this corroborates the motivation to use graph

frequency decomposition to segment brain signals into pieces corresponding to different levels of spatial fluctuations. Second, decomposed signals for a specific level of variation, notably x_H , are highly similar with respect to different scan sessions in an experiment as well as with respect to the two experiments with different sets of participants. The correlation coefficient between datasets for high graph frequencies is 0.6469. Third, recall that we normalize the brain signals at every sample point for all subjects, and for this reason signals x_L , x_M and x_H would be similarly distributed across the brain if nothing interesting happens at the decomposition. However, in both Figs. 6 and 7, it is observed that many brain regions possess magnitudes higher than a threshold in x_L ($\sim 60\%$ pass) and x_H ($\sim 20\%$ pass) while not many brain regions pass the thresholding with respect to x_M

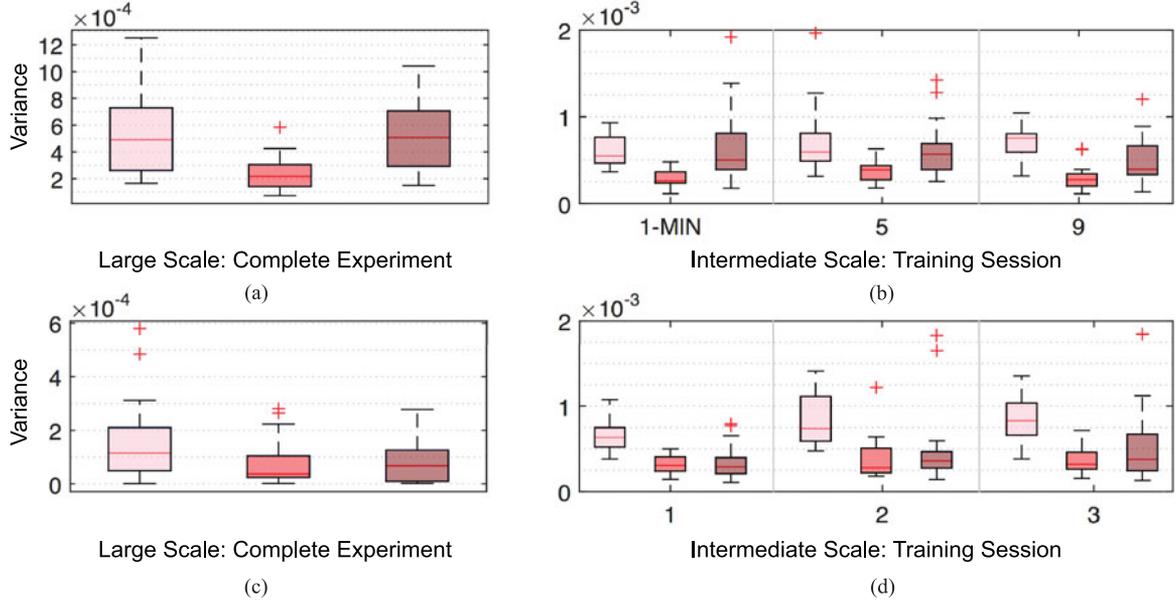


Fig. 8. Temporal adaptations of spatial variations. Boxplots showing differences in temporal adaptabilities between brain activities with smooth (pink), moderate (red) and rapid (maroon) spatial variations, measured over the complete experiment for 6 week (a) and 3 day (c), and individual training sessions for 6 week (b) and 3 day (d) experiments. We measured the temporal adaptations using the variance of the averaged activities over the complete experiment or with individual training sessions. Compared to activities with moderate spatial variations, smooth (95% sessions pass t-test with $p < 0.01$) and rapid (65% sessions pass t-test with $p < 0.005$) spatial variations have significantly higher temporal adaptations.

($\sim 3\%$ pass). It has long been understood that the brain combines some degree of disorganized behavior with some degree of regularity and that the complexity of a system is high when order and disorder coexist [44]. \mathbf{x}_L varies smoothly across the brain network and therefore can be regarded as regularity (order), whereas \mathbf{x}_H fluctuates rapidly and consequently can be considered as randomness (disorder). This evokes the intuition that graph frequency decomposition segments a brain signal \mathbf{x} into pieces \mathbf{x}_L and \mathbf{x}_H , which reflect order and disorder (and are therefore more interesting), as well as the remaining \mathbf{x}_M .

For the variance analysis, it is expected for the low graph frequency components (smooth spatial variation) to exhibit the smallest temporal variations, exceeded by medium and then high counterparts. Nonetheless, it is observed that brain activities with smooth spatial variations exhibit the most rapid temporal variation. Because it has been shown that temporal variation of observed brain activities is associated with better performance in tasks [15], this indicates a stronger contribution of low graph frequency components during the learning process. Furthermore, since the measurements were normalized such that the total energy of overall brain activities stayed constant at different sampling points, the rapid temporal changes of low graph frequency components should be accompanied by fast temporal variation of some other components, which are found to be high frequency components in all cases. Because these results were consistent for all of the temporal scales and datasets that we examined, and the association between temporal variability and positive performance has been established [16], we concluded that brain activities with smooth or rapid spatial variations offer greater contributions during learning. The graph frequency signatures at different stages of learning is analyzed in the next section.

	$\ \mathbf{x}_L\ _2$	$\ \mathbf{x}_M\ _2$	$\ \mathbf{x}_H\ _2$
6 week experiment (linear scale)	-0.3155	0.0897	0.4125
6 week experiment (logarithm scale)	-0.5409	0.3992	0.3565
3 day experiment	-0.9873	0.8443	0.9605

Fig. 9. Pearson correlation coefficients between the number of trials (level of task familiarity) and R values, defined as correlations between learning rate parameters and the norm of the decomposed signal of interest. More obvious adaptability between decomposed signals and learning across training is observed for \mathbf{x}_L and \mathbf{x}_H , with decreasing association with exposure to tasks for the former and increasing importance for the latter.

VI. FREQUENCY SIGNATURES OF TASK FAMILIARITY

Given that the decomposed signals exhibit interesting perspectives, it is natural to probe whether the signals corresponding to different levels of spatial variations associate with learning. To this end, we first describe how learning rate is evaluated. Given a participant, for each sequence completed, we defined the movement time M as the difference between the time of the first button press and the time of the last button press during a single sequence. We then estimate the participant’s learning rate by fitting an exponential function (plus a constant) using the robust outlier correction [45] to the sequence of movement times \mathbf{M}

$$\mathbf{M} = c_1 e^{\mathbf{t}/\kappa} + c_2. \quad (16)$$

where \mathbf{t} is a sequence representing the time index, κ is the exponential drop-off parameter (which we call the “learning rate parameter”) used to describe the early and fast rate of improvement, and c_1 and c_2 are nonnegative constants. Their sum $c_1 + c_2$ is an estimation of the starting speed of the participant of interest prior to training, while the parameter c_2 entails the fastest

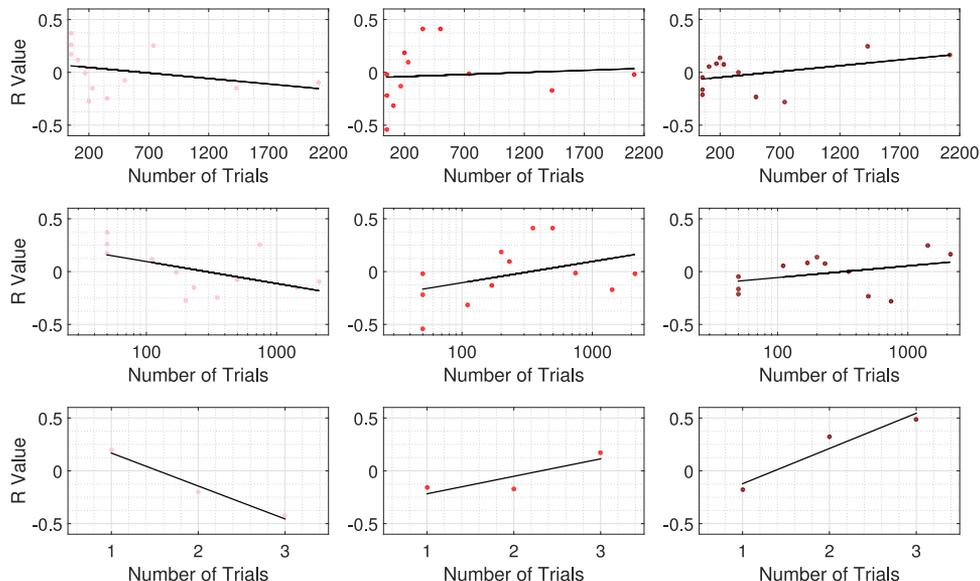


Fig. 10. Scatter plots in which each point is for a specific training session (level of task familiarity), depicting the R value defined here as correlations between learning rate parameters and the norm of the decomposed signal of interest (Pink points in the Left: \mathbf{x}_L , Red points in the Middle: \mathbf{x}_M , and Maroon points in the Right: \mathbf{x}_H). Top row: 6 week experiment with number of trials described in linear scale. Middle row: 6 week experiment with the number of trials evaluated in logarithm scale. We examine 6 week experiment by ordering the number of trials in both linear and logarithm scales to alleviate the fact that number of trials are densely distributed towards small values. Bottom row: 3 day experiment in which the number of trials is represented by the 3 scanning sessions in the experiment.

speed to complete the sequence attained by that participant after extended training. A negative value of κ indicates a decrease in movement time $M(t)$, which is thought to indicate that learning is occurring [46]. We chose exponential because it is viewed as the most statistically robust choice [47]. Further, the approach that we used has the advantage of estimating the rate of learning independent of initial performance or performance ceiling.

We evaluate the learning rate for all participants at each scanning session, and then compute the correlation between the norm $\|\mathbf{x}_L\|_2$ of the decomposed signal corresponding to low spatial variation and the learning rates across subjects. The correlation (R value) between the norms $\|\mathbf{x}_M\|_2$ as well as $\|\mathbf{x}_H\|_2$ and learning rates are also calculated. Fig. 10 plots the Pearson correlation coefficients at all scanning sessions of the two experiments considered. The horizontal axis denotes the level of exposure of participants to the sequence—which day in the 3 day experiment and how many number of trials participants have completed at the end of the scanning session in the 6 week experiment. Points are densely distributed for small number of trials in the 6 week experiment, so to mitigate this effect, we also plot the points by taking the logarithm of numbers of trials completed. We emphasize that due to normalization at each sampling point, the correlation values would all be 0 if graph frequency decomposition segments brain signals into three equivalent pieces. There are scan sessions where the correlation is of particular interest, however the most noteworthy observation is the change of correlation values with the level of exposure for participants.

In general, for \mathbf{x}_L corresponding to smooth spatial variation, we see a gradually decreasing trend in correlation with learning as training progresses. Although not all training sessions can be fit to this pattern (i.e. trials 500 and 740), it is still visible that the correlation with learning is above zero (≈ 0.25) at the start

	$\ \mathbf{x}_L\ _2$	$\ \mathbf{x}_M\ _2$	$\ \mathbf{x}_H\ _2$
6 week experiment (linear scale)	-0.2944	0.0486	0.3808
6 week experiment (logarithm scale)	-0.5334	0.3872	0.3189
3 day experiment	-0.9880	-0.9238	0.9976

Fig. 11. Pearson correlation coefficients for robustness testing, as comparable to Fig. 9. Each correlation coefficient is between the number of trials (level of task familiarity) and the average R value obtained at each trial. As such, each trial contains 121 R values for the different frequency ranges considered for K_L and K_M between 32 and 42, inclusive. Each R value is defined as the correlation between learning rate parameters and the norm of the decomposed signal of interest for a given frequency range.

of the training when participants perform the task for the first time and gradually shifts to below zero (≈ -0.25) at the end of the experiment when individuals are highly familiar with the sequence. For \mathbf{x}_H corresponding to vibrant spatial variation, its correlation with learning is below zero (≈ -0.2) at the start of the training, and gradually increases throughout training until it is above zero (≈ 0.25) at the end of the experiment, with the exception of trials 500 and 740. This is the exact opposite of \mathbf{x}_L . For \mathbf{x}_M , correlation between its norm $\|\mathbf{x}_M\|_2$ with learning rate generally increases with the intensity of training. However, this trend is not as obvious compared to other decomposition counterparts, and there are a greater number of sessions that cannot be fit to this pattern. The correlation between the number of trials and R values is summarized in Fig. 9. For robustness testing, we conduct similar analysis using the 120 other sets of parameters described in Section IV. The plots (similar to Fig. 10) for the R values resulting from all parameter choices are presented in Fig. 12 and the correlation between the number of trials and the average R value from considering all parameter

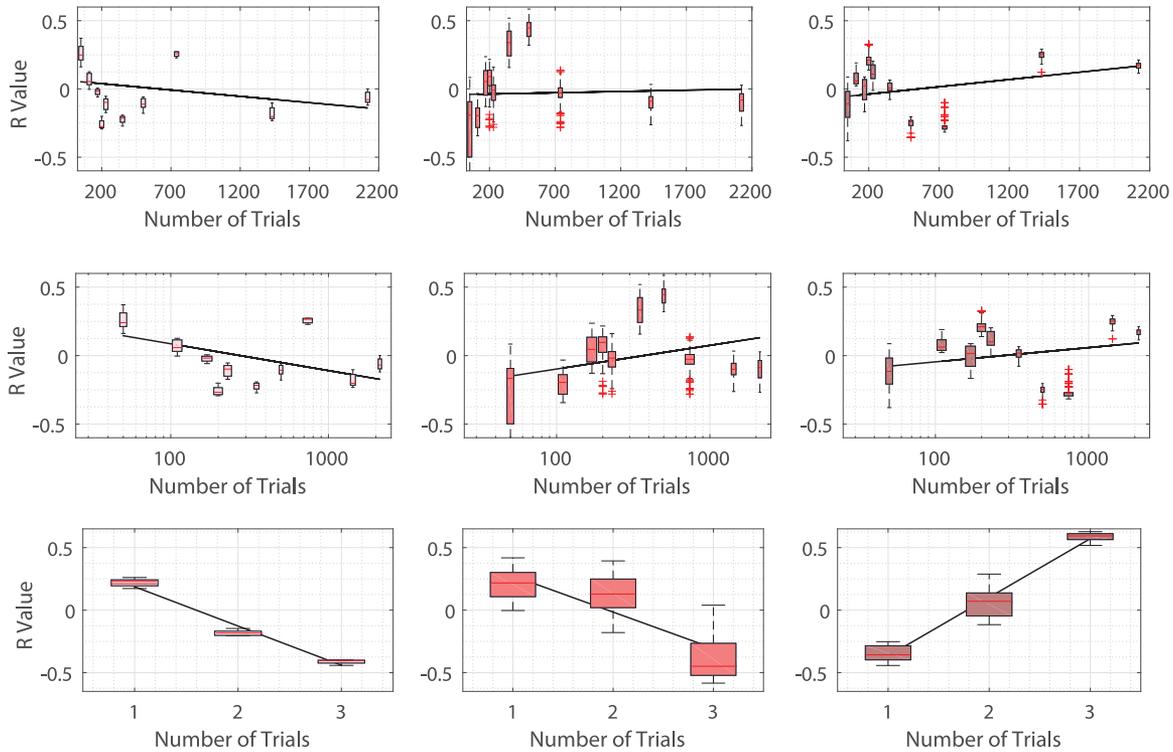


Fig. 12. Robustness testing to show similar trends as observed in Fig. 10. Each box is for a specific training session (level of task familiarity), depicting the R values obtained from changing the frequency ranges of K_L and K_M between 32 and 42, inclusive. As such, the R value is defined here as correlations between learning rate parameters and the norm of the decomposed signal of interest for a specific frequency range. Each box contains R values for 121 different combinations of frequency ranges.

choices is summarized in Fig. 11. Again, similar observations are found in different experiments involving different learning tasks and different sets of participants.

A. Discussion

This result further implies that the most association between learning or adaptability during the training process comes from the brain signals that either vary smoothly (x_L , regularity) or rapidly (x_H , randomness) with respect to the brain network. Therefore, the graph frequency decomposition could be used to capture more informative brain signals by filtering out non-informative counterparts, most likely associated with middle graph frequencies. Besides, the positive association between $\|x_L\|_2$ and learning rates as well as the negative association between $\|x_H\|_2$ and learning rates at the start of training indicates that it favors learning to have more *smooth*, *spread*, and *cooperative* brain signals when we face an unfamiliar task. As we gradually become familiar with the task, the smooth and cooperative signal distribution becomes less and less important, and there is a level of exposure when such signal distribution becomes destructive instead of constructive. We note that the task in the 3 day experiment is more difficult compared to that of the 6 week experiment, and therefore the time when the cooperative signal distribution starts to become detrimental (the point where the regression line intercepts the horizontal line of R value equaling 0) is also comparable in the two experiments, describing a certain level of familiarity to the task. When we

become highly familiar with the task, it is better and favors further learning to have *varied*, *spiking*, and *competitive* brain signals.

In the dataset evaluated here, we utilize the average coherence between time series at pairs of brain cortical and subcortical regions during the training as the network. Hence, a concentration of brain activities towards low graph frequencies would imply that activities on brain regions that are generally cooperative are indeed similar. Simultaneously, the interpretation of concentration of brain activities towards high graph frequencies is that brain activities on brain regions that are generally cooperative are in fact dissimilar. In terms of learning, one possible explanation is that there are two different stages in learning: we start by grasping the big picture of the task to perform relatively well, and then we refine the details to perform better and to approach our limits.

Because the graph frequency analysis method presented in this paper applies to any setting where signals are defined on top of a network structure representing proximities between nodes, it would be interesting in future to use this method to investigate other types of signals and networks in neuroscience problems. As an example, in situations given fMRI measurements on structural networks, concentration of signals in low graph frequency components would imply functional activities do behave according to the structural networks.

Besides, it has been understood that learning is different when one is unfamiliar or familiar with a particular task—it is easy to improve performance at first exposure due to the fact that one is

far from their performance ceiling. It would therefore be interesting to utilize graph frequency decomposition to further analyze the difference between learning scenarios at different stages of familiarity, e.g. adaptability at first exposure and creativity when one fully understands the components of the specific tasks.

VII. CONCLUSION

We used graph spectrum methods to analyze functional brain networks and signals during simple motor learning tasks, and established connections between graph frequency with principal component analysis when the networks of interest denote functional connectivity. We discerned that brain activities corresponding to different graph frequencies exhibit different levels of adaptability during learning. Further, the strong correlation between graph spectral property of brain networks with the level of familiarity of tasks was observed, and the most contributing frequency signatures at different task familiarity was recognized.

REFERENCES

- [1] H. Haken, *Principles of Brain Functioning: A Synergetic Approach to Brain Activity, Behavior and Cognition*, vol. 67. New York, NY, USA: Springer, 2013.
- [2] M. D. Fox and M. E. Raichle, "Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging," *Nature Rev. Neurosci.*, vol. 8, no. 9, pp. 700–711, 2007.
- [3] L. J. Cole, M. J. Farrell, E. P. Duff, J. B. Barber, G. F. Egan, and S. J. Gibson, "Pain sensitivity and fMRI pain-related brain activity in Alzheimer's disease," *Brain*, vol. 129, no. 11, pp. 2957–2965, 2006.
- [4] J. D. Medaglia, W. Huang, S. Segarra, C. Olm, J. Gee, M. Grossman, A. Ribeiro, C. T. McMillan, and D. S. Bassett, "Brain network efficiency is influenced by pathological source of corticobasal syndrome," *Neurology*, 2016, submitted for publication.
- [5] S. Achard, R. Salvador, B. Whitcher, J. Suckling, and E. Bullmore, "A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs," *J. Neurosci.*, vol. 26, no. 1, pp. 63–72, 2006.
- [6] E. Bullmore and O. Sporns, "The economy of brain network organization," *Nature Rev. Neurosci.*, vol. 13, no. 5, pp. 336–349, 2012.
- [7] J. Richiardi, S. Achard, H. Bunke, and D. Van De Ville, "Machine learning with brain graphs: predictive modeling approaches for functional imaging in systems neuroscience," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 58–70, May 2013.
- [8] H. Q. Nguyen *et al.*, "Compression of human body sequences using graph wavelet filter banks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.* 2014, pp. 6152–6156.
- [9] S. Segarra, W. Huang, and A. Ribeiro "Diffusion and superposition distances for signals supported on networks," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 1, no. 1, pp. 20–32, Mar. 2015.
- [10] J. Ma, W. Huang, S. Segarra, and A. Ribeiro "Diffusion filtering for graph signals and its use in recommendation systems," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Shanghai, China Mar. 20–25, 2016, pp. 4563–4567.
- [11] A. Gadde, A. Anis, and A. Ortega "Active semi-supervised learning using sampling theory for graph signals," in *Proc. 20th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining* 2014, pp. 492–501.
- [12] A. Sandryhaila and J. M. Moura "Discrete signal processing on graphs," *IEEE Trans Signal Process.*, vol. 61, no. 7, pp. 1644–1656, Apr. 2013.
- [13] A. Sandryhaila and J. M. Moura "Discrete signal processing on graphs: Frequency analysis," *IEEE Trans Signal Process.*, vol. 62, no. 12, pp. 3042–3054, Jun. 2014.
- [14] D. Shuman *et al.*, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process Mag.*, vol. 30, no. 3, pp. 83–98, May 2013.
- [15] D. D. Garrett, N. Kovacevic, A. R. McIntosh, and C. L. Grady "The modulation of bold variability between cognitive states varies by age and processing speed," *Cerebral Cortex*, vol. 23, pp. 684–693, 2012.
- [16] J. J. Heisz, J. M. Shedden, and A. R. McIntosh "Relating brain signal variability to knowledge representation," *Neuroimage*, vol. 63, no. 3, pp. 1384–1392, 2012.
- [17] N. Leonardi *et al.*, "Principal components of functional connectivity: A new approach to study dynamic brain connectivity during rest," *NeuroImage*, vol. 83, pp. 937–950, 2013.
- [18] R. Viviani, G. Grön, and M. Spitzer "Functional principal component analysis of fMRI data," *Human Brain Mapping*, vol. 24, no. 2, pp. 109–129, 2005.
- [19] L. M. Harrison, W. Penny, J. Daunizeau, and K. J. Friston "Diffusion-based spatial priors for functional magnetic resonance images," *Neuroimage*, vol. 41, no. 2, pp. 408–423, 2008.
- [20] J. Shi, and J. Malik "Normalized cuts and image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 888–905, Aug. 2000.
- [21] H. Behjat, N. Leonardi, L. Sörnmo, and D. Van De Ville "Anatomically-adapted graph wavelets for improved group-level fMRI activation mapping," *NeuroImage*, vol. 123, pp. 185–199, 2015.
- [22] D. Hammond, B. Scherrer, and S. Warfield "Cortical graph smoothing: A novel method for exploiting DWI-derived anatomical brain connectivity to improve EEG source estimation," *IEEE Trans. Med. Imag.*, vol. 32, no. 10, pp. 1952–1963, Oct. 2013.
- [23] S. Atasoy, I. Donnelly, and J. Pearson, "Human brain networks function in connectome-specific harmonic waves," *Nature Commun.*, vol. 7, 2016, Art. no. 10340.
- [24] J. A. Kleim *et al.*, "Motor learning-dependent synaptogenesis is localized to functionally reorganized motor cortex," *Neurobiol. Learning Memory*, vol. 77, no. 1, pp. 63–77, 2002.
- [25] D. S. Bassett, M. Yang, N. F. Wymbs, and S. T. Grafton "Learning-induced autonomy of sensorimotor systems," *Nature Neurosci.*, vol. 18, no. 5, pp. 744–751, 2015.
- [26] F. Chung, *Spectral Graph Theory*, vol. 92 Providence, RI, USA: Amer. Math. Soc., 1997.
- [27] W.-T. Zhang *et al.*, "Relations between brain network activation and analgesic effect induced by low vs. high frequency electrical acupoint stimulation in different subjects: a functional magnetic resonance imaging study," *Brain Res.*, vol. 982, no. 2, pp. 168–178, 2003.
- [28] J. Grutzendler, N. Kasthuri, and W.-B. Gan "Long-term dendritic spine stability in the adult cortex," *Nature*, vol. 420, no. 6917, pp. 812–816, 2002.
- [29] S. Segarra, A. G. Marques, G. Leus, and A. Ribeiro, "Reconstruction of graph signals through percolation from seeding nodes," *IEEE Trans. Signal Process.*, vol. 64, no. 16, pp. 4363–4378, Aug. 2016.
- [30] D. S. Bassett, N. F. Wymbs, M. P. Rombach, M. A. Porter, P. J. Mucha, and S. T. Grafton "Task-based core-periphery organization of human brain dynamics," *PLoS Comput. Biol.*, vol. 9, no. 9, p. e1003171, 2013.
- [31] D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton "Dynamic reconfiguration of human brain networks during learning," *Proc. Nat. Acad. Sci. USA*, vol. 108, no. 18, pp. 7641–7646, 2011.
- [32] N. F. Wymbs, D. S. Bassett, P. J. Mucha, M. A. Porter, and S. T. Grafton "Differential recruitment of the sensorimotor putamen and frontoparietal cortex during motor chunking in humans," *Neuron*, vol. 74, no. 5, pp. 936–946, 2012.
- [33] J. Wang *et al.*, "Parcellation-dependent small-world brain functional networks: A resting-state fMRI study," *Human Brain Mapping*, vol. 30, no. 5, pp. 1511–1523, 2009.
- [34] J. D. Power *et al.*, "Functional network organization of the human brain," *Neuron*, vol. 72, no. 4, pp. 665–678, 2011.
- [35] F. T. Sun, L. M. Miller, and M. D'Esposito "Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data," *Neuroimage*, vol. 21, no. 2, pp. 647–658, 2004.
- [36] Y. He, Z. J. Chen, and A. C. Evans "Small-world anatomical networks in the human brain revealed by cortical thickness from MRI," *Cerebral Cortex*, vol. 17, no. 10, pp. 2407–2419, 2007.
- [37] C. R. Genovese, N. A. Lazar, and T. Nichols "Thresholding of statistical maps in functional neuroimaging using the false discovery rate," *Neuroimage*, vol. 15, no. 4, pp. 870–878, 2002.
- [38] M. Xia *et al.*, "Brainnet viewer: A network visualization tool for human brain connectomics," *PLoS One*, vol. 8, no. 7, p. e68910, 2013.
- [39] F. Klimm, D. S. Bassett, J. M. Carlson, and P. J. Mucha "Resolving structural variability in network models and the brain," *PLOS Comput. Biol.*, vol. 10, no. 3, p. e1003491, 2014.
- [40] C. Lohse, D. S. Bassett, K. O. Lim, and J. M. Carlson "Resolving anatomical and functional structure in human brain organization: Identifying mesoscale organization in weighted network representations," *PLoS Comput. Biol.*, vol. 10, no. 10, p. e1003712, 2014.

- [41] R. Albert and A.-L. Barabási “Statistical mechanics of complex networks,” *Rev. Modern Phys.*, vol. 74, no. 1, p. 47, 2002.
- [42] C. Aicher, A. Z. Jacobs, and A. Clauset, “Learning latent block structure in weighted networks,” *J. Complex*, vol. 3, no. 2, pp. 221–248, 2015.
- [43] V. Braitenberg and A. Schüz, *Cortex: Statistics and Geometry of Neuronal Connectivity*, New York, NY, USA: Springer, 2013.
- [44] O. Sporns, *Networks of the Brain*, Cambridge, MA, USA: MIT Press, 2011.
- [45] D. A. Rosenbaum, *Human Motor Control*, New York, NY, USA: Academic, 2009.
- [46] E. Dayan and L. G. Cohen “Neuroplasticity subserving motor skill learning,” *Neuron*, vol. 72, no. 3, pp. 443–454, 2011.
- [47] A. Heathcote, S. Brown, and D. Mewhort “The power law repealed: The case for an exponential law of practice,” *Psychonomic Bull. Rev.*, vol. 7, no. 2, pp. 185–207, 2000.



Weiyu Huang received the B.Eng. (Hons.) degree in electronics and telecommunication from Australian National University (ANU), Canberra, Australia, in 2012. He is currently working toward the Ph.D. degree in electrical and systems engineering at the University of Pennsylvania (Penn), Philadelphia, PA, USA. From 2011 to 2013, he was a Telecommunication Engineer and Policy Officer with the Australian Communication and Media Authority. His research interests include network theory, pattern recognition, graph signal processing, and the study of networked

data arising in human, social, and technological networks. He received the National Scholarship offered by the Ministry of Education, China, University Medal (best graduates) and H.A. Johns Award (best academic achievement) offered by ANU, as well as the Department Graduate Award for Best Ph.D. Colloquium by Penn for academic year 2015–2016.



Leah Goldsberry is working toward the B.S.E. degree in systems science and engineering with the University of Pennsylvania (Penn), Philadelphia, PA, USA. She is a Rachleff Scholar of the University of Pennsylvania. Her research interests include application of graph signal processing to the study of networks and behavioral and cognitive performance.



Nicholas F. Wymbs received his B.A. in psychology from the University of Notre Dame in 2002 and awarded his Ph.D. in psychology from the University of California, Santa Barbara in 2011. Prior to his doctoral training, from 2002–2005, Nicholas studied fMRI of language and motor systems as a staff member of the Brain Research Imaging Center (BRIC) at the University of Chicago. He is currently a post-doctoral research fellow at Johns Hopkins University and the Kennedy Krieger Institute. His research interests concern the behavioral features and neural mechanisms of sensorimotor behavior. Currently, his focus concerns the application of neuroimaging and non-invasive brain stimulation techniques to understand motor control and motor skill learning in healthy and neurologically impaired human participants.



Scott T. Grafton received the M.D. degree from the University of Southern California and completed residencies in neurology at the University of Washington and in nuclear medicine at the University of California, Los Angeles (UCLA). He holds the Bedrosian-Coyne Presidential Chair in Neuroscience at the University of California, Santa Barbara (UCSB). He is recognized for developing multimodal brain mapping techniques for accelerating discovery and diagnosis of the nervous system. He is the Co-Director at the Institute for Collaborative

Biotechnologies, which draws on bioinspiration and innovative bioengineering solutions for both nonmedical and medical challenges posed by the defense and medical communities. He directs a research program at the interface of learning and skilled behavior, network science, and multimodal imaging. After a research fellowship at UCLA in brain imaging, he developed brain-imaging programs at the University of Southern California, Emory University, and Dartmouth College before joining the UCSB faculty in 2006, where he directs the UCSB Imaging Center. He uses fMRI, magnetic stimulation, and high-density EEG to characterize the neural basis of goal-directed behavior and mechanisms of brain plasticity using an approach grounded in 20 years of clinical experience.



Danielle S. Bassett received the B.S. degree in physics from Pennsylvania State University and the Ph.D. degree in physics from the University of Cambridge, U.K. She is the Eduardo D. Glandt Faculty Fellow and Associate Professor in the Department of Bioengineering, University of Pennsylvania. She is most well known for her work blending neural and systems engineering to identify fundamental mechanisms of cognition and disease in human brain networks. Following a postdoctoral position at the University of California, Santa Barbara, she was a

Junior Research Fellow at the Sage Center for the Study of the Mind. In 2012, she was named American Psychological Association’s “Rising Star” and given an Alumni Achievement Award from the Schreyer Honors College at Pennsylvania State University for extraordinary achievement under the age of 35. In 2014, she was named an Alfred P Sloan Research Fellow and received the MacArthur Fellow Genius Grant. In 2015, she received the IEEE Engineering in Medicine and Biology Society Early Academic Achievement Award, and was named an Office of Naval Research Young Investigator. In 2016, she received a National Science Foundation (NSF) CAREER Award. She is the founding Director of the Penn Network Visualization Program, a combined undergraduate art internship and K-12 outreach program bridging network science and the visual arts. Her work has been supported by the National Science Foundation, the National Institutes of Health, the Army Research Office, the Army Research Laboratory, the Alfred P Sloan Foundation, the John D and Catherine T MacArthur Foundation, and the Office of Naval Research.



Alejandro Ribeiro received the B.Sc. degree in electrical engineering from the Universidad de la Republica Oriental del Uruguay, Montevideo, in 1998, and the M.Sc. and Ph.D. degrees in electrical engineering from the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, in 2005 and 2007, respectively. From 1998 to 2003, he was a member of the technical staff at Bellsouth Montevideo. After his M.Sc. and Ph.D studies, in 2008, he joined the University of Pennsylvania (Penn), Philadelphia, where he is currently the Rosenbluth

Associate Professor in the Department of Electrical and Systems Engineering. His research interests include the applications of statistical signal processing to the study of networks and networked phenomena. His focus is on structured representations of networked data structures, graph signal processing, network optimization, robot teams, and networked control. He received the 2014 O. Hugo Schuck Best Paper Award, the 2012 S. Reid Warren, Jr., Award presented by Penn’s undergraduate student body for outstanding teaching, the National Science Foundation CAREER Award in 2010, and paper awards at the 2016 IEEE Workshop on Statistical Signal Processing, 2016 IEEE Sensor Array and Multichannel Signal Processing Workshop, 2015 Asilomar Conference on Signals, Systems, and Computers, 2013 American Control Conference, 2006 IEEE International Conference on Acoustics, Speech, and Signal Processing, and 2005 IEEE International Conference on Acoustics, Speech, and Signal Processing. He is a Fulbright scholar and a Penn Fellow.