



# Mental State Classification Using Multi-Graph Features

Guodong Chen<sup>1†</sup>, Hayden S. Helm<sup>2\*†</sup>, Kate Lytvynets<sup>2</sup>, Weiwei Yang<sup>2</sup> and Carey E. Priebe<sup>1</sup>

<sup>1</sup> Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD, United States, <sup>2</sup> Microsoft Research, Microsoft, Redmond, WA, United States

We consider the problem of extracting features from passive, multi-channel electroencephalogram (EEG) devices for downstream inference tasks related to high-level mental states such as stress and cognitive load. Our proposed feature extraction method uses recently developed spectral-based multi-graph tools and applies them to the time series of graphs implied by the statistical dependence structure (e.g., correlation) amongst the multiple sensors. We study the features in the context of two datasets each consisting of at least 30 participants and recorded using multi-channel EEG systems. We compare the classification performance of a classifier trained on the proposed features to a classifier trained on the traditional band power-based features in three settings and find that the two feature sets offer complementary predictive information. We conclude by showing that the importance of particular channels and pairs of channels for classification when using the proposed features is neuroscientifically valid.

**Keywords:** electroencephalogram, band-based features, multi-graph features, mental workload prediction, ablation study

## OPEN ACCESS

### Edited by:

Bradley Jay Edelman,  
Max Planck Institute of Neurobiology  
(MPIN), Germany

### Reviewed by:

Manaf Altaieb,  
University of Wasit, Iraq  
Yasue Mitsukura,  
Keio University Shonan Fujisawa  
Campus, Japan

### \*Correspondence:

Hayden S. Helm  
haydensh@msn.com

<sup>†</sup>These authors have contributed  
equally to this work

### Specialty section:

This article was submitted to  
Brain-Computer Interfaces,  
a section of the journal  
Frontiers in Human Neuroscience

**Received:** 27 April 2022

**Accepted:** 30 May 2022

**Published:** 08 July 2022

### Citation:

Chen G, Helm HS, Lytvynets K,  
Yang W and Priebe CE (2022) Mental  
State Classification Using Multi-Graph  
Features.  
*Front. Hum. Neurosci.* 16:930291.  
doi: 10.3389/fnhum.2022.930291

## 1. INTRODUCTION

Successful non-invasive Brain-Computer Interfaces (BCI) require solving a high dimensional, high frequency prediction problem. For EEG-based systems in particular, we have access to tens of streams of data that are drastically attenuated by the skull. Hence, mining relevant predictive signal is a serious challenge. This is particularly true for passive tasks such as cognitive load and stress prediction where no explicit action or evoked potential is used to differentiate conditions. Since passive tasks and prediction are not trial based, the “ground truth” labels are often weaker compared to active tasks (e.g., command and control *via* motor imagery; Miller et al., 2010; Kaya et al., 2018) and tasks with evoked responses (Spüler et al., 2012).

The most common approach to passive, non-invasive EEG-based BCI predictive tasks is to leverage neuroscientifically relevant features from the waveforms from each of the channels. For example, the alpha band (8–12 Hz) is known to be more active in stressed individuals and the theta (4–7 Hz) and low beta (13–20 Hz) bands are known to be active in fatigued persons (Nayak and Anilkumar, 2021). Simple functions of the relative masses of the sub-waveforms are also popular and useful features (Kamzanova et al., 2011). Similarly, functions of data from two channels of the EEG-device (such as frontal alpha asymmetry) have been shown to have different characteristics under different mental states and can be useful features for classification (Fischer et al., 2018). Both of these types of features rely on conventional neuroscience wisdom and are thus often hand-crafted for the particular predictive task.

Connectivity and correlation-adjacent features such as synchronization and EEG coherence between signals from different sensors have been explored in Wei et al. (2007) and Wu et al. (2021).

However, the choice of different sensors to measure coupling strength is highly subjective and require neurophysiological a priori knowledge. Recent research has also investigated representing EEG signals as matrices (Yger et al., 2016; Congedo et al., 2017). These approaches derive the kernel on the Riemannian space using the average correlation matrix for a particular class. For other EEG-based tasks, such as classifying motor imagery (Oikonomou et al., 2017), techniques such as Common Spatial Filtering (CSP) and its variants (Koles, 1991; Blankertz et al., 2007; Ang et al., 2008) are used to learn a supervised projection from the set of sensors to an optimal subspace. Supervised methods embeddings such as CSP, at least empirically, often fail in passive BCI applications due to the supervisory signal being “weak” and not trial based.

On the flip side of conventional neuroscientific features is deep learning. Deep learning has achieved a state-of-art performance in fields such as speech recognition, visual object recognition, and object detection (Bengio et al., 2003; LeCun et al., 2015). Many recent works have explored the use of convolutional neural network-based methods for automatic feature extractions in EEG-based BCIs (Lawhern et al., 2018; Siddharth et al., 2022). These approaches either use a pre-trained model such as VGG-16 (Simonyan and Zisserman, 2014) to extract features from the power spectral density heat maps of different bands (Siddharth et al., 2022) or encapsulated an optimal spatial filtering in the network structure (Lawhern et al., 2018). In our experience, as seems to be the experience of others (Lotte et al., 2018), current deep learning-based methods have not shown a significant improvement over methods utilizing conventional features in multiple BCI applications (Lotte et al., 2018) and are notably less interpretable.

In this paper we propose a method that jointly learns a set of features from relationships between pairs of channels. The method leverages the temporal and spatial relationships between channels and will learn a different projection from the space of statistical-dependence matrices for each task—yielding considerable flexibility compared to traditional hand-selected features. The process from going from EEG signal to feature vectors borrows heavily from recent developments in spectral-based multi-graph analysis (Wang et al., 2019; Gopalakrishnan et al., 2020). We demonstrate in three different experiments each on two datasets that the proposed set of features can improve performance over standard band-based features and that the combination of the two nearly always improves over either individually. We argue that this improvement is due to the two sets of features capturing complementary predictive information. Lastly, we analyze the importance of each channel and each pair of channels and argue that the proposed feature set is neuroscientifically valid.

## 2. MATERIALS AND METHODS

### 2.1. Band-Based Features

In this subsection, we detail the traditional band frequency (BF) power spectral density approach (Lotte et al., 2018). Consider a (potentially preprocessed) multi-variate EEG time series  $\{S^{(t)} \in \mathbb{R}^{n_c} : t = 1, 2, \dots, T\}$ , where  $S_j^{(t)}$  represents signal from the  $j$ -th channel at time  $t$  and  $n_c$  is the number of channels. We first window the multi-variate EEG time series into potentially overlapping windows. Defining the window size to be  $w$  and the overlap between windows to be  $h$ , the EEG segment of the  $k$ -th window of the  $j$ -th electrode  $X_j^{(k)} \in \mathbb{R}^w$  is given by

$$X_j^{(k)} = [S_j^{(k*(w-h)+1)}, S_j^{(k*(w-h)+2)}, \dots, S_j^{((k+1)*(w-h)+h)}].$$

We let  $n_w$  be the number of windows after splitting the original times series in the time domain. For each channel  $j$ , we estimate the power spectral density  $\{\hat{P}_i(X_j^{(k)}) : k = 1, 2, \dots, n_w\}_{i=1}^I$  and sum the powers in  $I$  non-overlapping bands. We then normalize the band powers such that for each channel the features sum to 1,

$$F_{ij}^{(k)} := \hat{f}_i(X_j^{(k)}) = \frac{\hat{P}_i(X_j^{(k)})}{\sum_{i=1}^I \hat{P}_i(X_j^{(k)})}.$$

Thus, for each time window  $k$ , we have  $n_c \times I$  features. We consequently flatten this matrix into a single,  $n_c \cdot I$  length vector  $F^{(k)}$ . Finally, because usually  $n_c \cdot I$ -dimensional space will be relatively large given the amount of data we have access to, we project the feature vector into a low-dimensional space. The projection is learned *via* PCA of the available training data.

### 2.2. Learning a Representation From a Time Series of Graphs

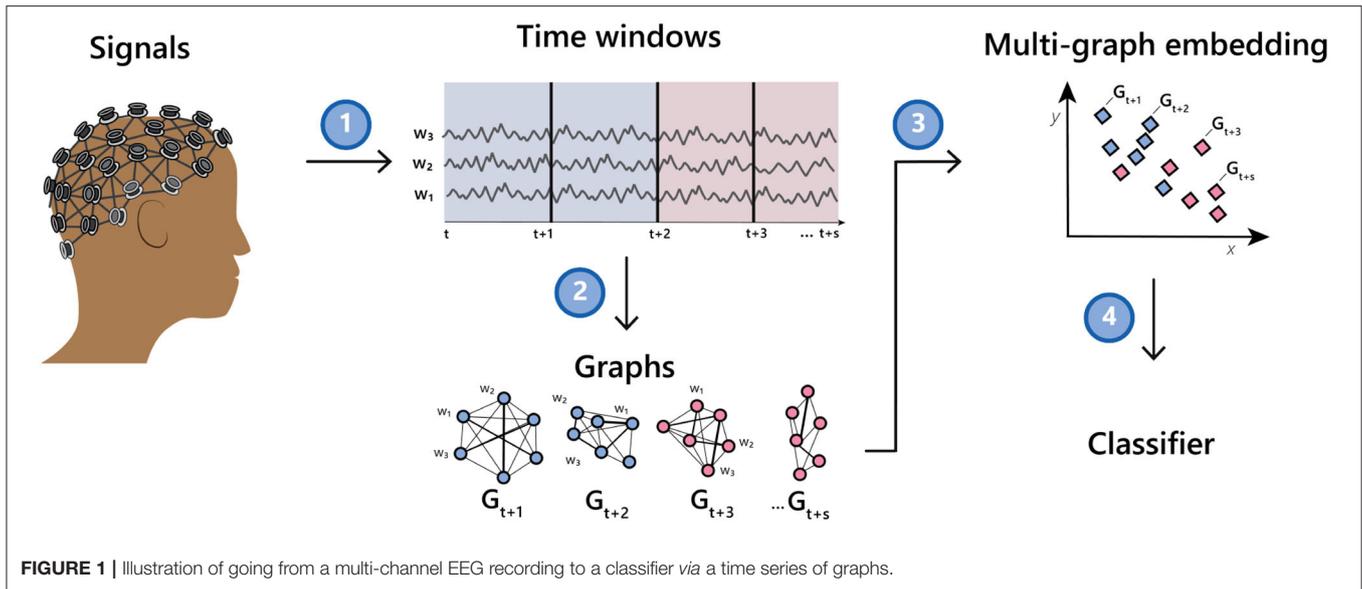
The method we describe herein takes as input a collection of multi-channel time series and induces a network or graph on the set of channels. It then induces a graph on the set of networks and finally learns a single vector representation for each multi-channel time series. The representation can then be used as input to a classifier or other tools to aid downstream inference. See **Figure 1** for an illustration of the method.

The described method is natively transductive and thus only learns a representation for the windows of the multi-channel time series that it has access to when learning the embedding. This can be limiting in applications where we want to apply the learned embedding from one session (or participant) to the data from another session (or participant). To alleviate this issue we describe an “out-of-sample” embedding that can take a previously unseen multi-channel time series and map it to the learned embedding space. We sometimes refer to the process of taking a windowed segment of EEG and projecting it into the appropriate embedding space as a feature mapping and the corresponding function as a feature map.

We consider the windowed multi-variate time series described in Section 2.1:

$$X_j^{(k)} = [S_j^{(k*(w-h)+1)}, S_j^{(k*(w-h)+2)}, \dots, S_j^{((k+1)*(w-h)+h)}].$$

Let  $s: \mathbb{R}^w \times \mathbb{R}^w \rightarrow \mathbb{R}$  be a similarity function on objects in  $\mathbb{R}^w$  and  $A^{(k)}$  be the similarity matrix



**FIGURE 1** | Illustration of going from a multi-channel EEG recording to a classifier via a time series of graphs.

between pairs of the windowed times series for window index  $k$ ,

$$A^{(k)} = \begin{bmatrix} s(X_1^{(k)}, X_1^{(k)}) & \dots & s(X_1^{(k)}, X_{n_c}^{(k)}) \\ \vdots & \ddots & \vdots \\ s(X_{n_c}^{(k)}, X_1^{(k)}) & \dots & s(X_{n_c}^{(k)}, X_{n_c}^{(k)}) \end{bmatrix}.$$

The matrix  $A^{(k)} \in \mathbb{R}^{n_c \times n_c}$  can be thought of as a graph where the vertices are channels and the edge weight is the similarity score between the two channels for window  $k$ .

Given the collection of graphs  $\{A^{(k)}\}_{k=1}^{n_w}$  there are numerous kernels that map a graph to a real-valued vector space (Kriege et al., 2020). We use a spectral approach because it is relatively scalable and statistically principled in general settings (Athreya et al., 2018). Hence, we induce a graph on the  $\{A^{(k)}\}$  by considering a pairwise similarity  $s' : \mathbb{R}^{n_c \times n_c} \times \mathbb{R}^{n_c \times n_c} \rightarrow \mathbb{R}$  and corresponding pairwise similarity matrix

$$B = \begin{bmatrix} s'(A^{(1)}, A^{(1)}) & \dots & s'(A^{(1)}, A^{(n_w)}) \\ \vdots & \ddots & \vdots \\ s'(A^{(n_w)}, A^{(1)}) & \dots & s'(A^{(n_w)}, A^{(n_w)}) \end{bmatrix}. \quad (1)$$

Recall from results in statistical network analysis that when a graph  $B = U\Sigma V^T$  is a realization of a random dot product graph (RDPG) (Athreya et al., 2018) that the appropriately truncated left singular vectors of  $B$  scaled by the square root of the corresponding singular values are consistent estimates for the underlying latent positions  $Z \in \mathbb{R}^{n_w \times d}$  where  $Z^{(k)} \in \mathbb{R}^d$  is the latent position corresponding to window  $k$ . In applications where the RDPG assumption is violated—as is likely in ours—reducing the dimensionality from  $n_w$  to dimension  $d < n_w$  in this way has shown to be useful for downstream inference tasks.

We let  $\hat{Z} \in \mathbb{R}^{n_w \times d}$  be the estimate of the latent positions after embedding  $B$ . These vectors will be the representations for the

multi-channel time series that we use for inference. They will also serve as the basis representation for embedding previously unseen EEG windows.

In our experiments we evaluate the effectiveness  $\hat{Z}$  as the representations used for mental state classification when letting  $s(\cdot, \cdot)$  be Pearson correlation and

$$s'(A^{(k)}, A^{(k')}) = 1 - \frac{\|A^{(k)} - A^{(k')}\|_2 - \min\{\|A^{(k)} - A^{(k')}\|_2\}_{k, k' \in \{1, \dots, n_w\}}}{\max\{\|A^{(k)} - A^{(k')}\|_2\}_{k, k' \in \{1, \dots, n_w\}} - \min\{\|A^{(k)} - A^{(k')}\|_2\}_{k, k' \in \{1, \dots, n_w\}}}.$$

We make some comments on potential improvements over these functions in Section 4.

As mentioned above, the process to go from a windowed time series to a jointly learned vector representation in  $\mathbb{R}^d$  does not natively extend to previously unseen data. Though one could learn a new embedding every time new data is collected, this approach is relatively unscalable in both time and computation and of little use for real-time systems. To overcome this problem, “out-of-sample” extensions of popular transductive embedding methods have been proposed to map new data to the space spanned by the original embeddings (Bengio et al., 2003; Trosset and Priebe, 2008; Tang et al., 2013; Levin et al., 2018). We now describe the out-of-sample extension to the adjacency spectral embedding analyzed by Tang et al. (2013).

Given the original set of windowed time series  $\{X^{(k)}\}_{k=1}^{n_w}$ , the learned embeddings  $\hat{Z} \in \mathbb{R}^{n_w \times d}$ , a similarity function on data from the channels of the time series (i.e.,  $s(\cdot, \cdot)$  from above), a similarity function on the graphs  $\mathbb{R}^{n_c \times n_c}$  [i.e.,  $s'(\cdot, \cdot)$ ], and a (potentially yet unseen) windowed time series  $X \in \mathbb{R}^{n_c \times w}$  with corresponding similarity matrix  $A$ , we define the out-of-sample mapping to be

$$T(X; \{X^{(k)}\}_{k=1}^{n_w}, \hat{Z}^k, s, s') := \sum_{k=1}^{n_w} s'(A, A^{(k)}) \left[ (\hat{Z}^T \hat{Z})^{-1} \hat{Z}^T \right]_k. \quad (2)$$

As discussed in Tang et al. (2013), in the context of graph embeddings the feature mapping  $T$  is considerably more time and computationally efficient than recalculating the singular value decomposition of the matrix defined in (1) at the cost of a small performance degradation on downstream inference. Hence, for real-time applications  $T$  is preferred to recalculation.

## 2.3. Datasets

Mental Math (Zyma et al., 2019) and MATB-II (proprietary) are representative of stress prediction and cognitive load prediction tasks, respectively, and are thus helpful in understanding the high-level mental state prediction capabilities of the three sets of features.

### 2.3.1. Mental Math

In the Mental Math study there are two recordings for each participant—one corresponding to a resting state and one corresponding to a stressed state. For the resting state, participants counted mentally (i.e., without speaking or moving their fingers) with their eyes closed for three minutes. For the stressful state, participants were given a four digit number (e.g., 1,253) and a two digit number (e.g., 43) and asked to recursively subtract the two digit number from the four digit number for 4 min. This type of mental arithmetic is known to induce stress (Noto et al., 2005).

There were initially 66 participants (47 women and 19 men) of matched age in the study. Thirty of the participants were excluded from the released data due to poor EEG quality. Thus we consider the provided set of 36 participants analyzed by Zyma et al. (2019). The released EEG data was preprocessed *via* a high-pass filter and a power line notch filter (50 Hz). Artifacts such as eye movements and muscle tension were removed *via* ICA.

The EEGs were recorded monopolarly using Neurocom EEG 23-channel system (Ukraine, XAI-MEDICA). For our analysis we further apply high pass (0.5 Hz) and low pass (30 Hz) filters and window the data into two and a half second chunks with no overlap. In our experiments we consider the two-class classification task {stressed, not stressed} and note that some of the windowed data in the training will be adjacent to some of the windowed data in the test set by virtue of how the experiment was set up.

The TSG features that we use in the Mental Math experiments are based on the pairwise Frobenius norms between correlation matrices as described in Section 2.2. The number of components kept in the singular vectors is data-dependent and is the second “elbow” of the scree plot of singular values estimated by Zhu and Ghodsi’s method Zhu and Ghodsi (2006).

The BF features we use are derived from the low (4.1–5.8 Hz) and high (5.9–7.4 Hz) theta bands; low (7.4–8.9 Hz), middle (9.0–11.0 Hz), and high (11.1–12.9 Hz) alpha bands; and low (13.0–19.9 Hz), medium (20.0–25.0 Hz), and high (25.0–30.0 Hz) beta bands. Once the power of each of these bands is normalized by channel there are  $19(8) = 152$  features per window that we then project into a lower dimension *via* PCA. As with the TSG features, we select the number of components *via* Zhu and Ghodsi’s method. In our experiments this translates

to between a 3 and 8 dimensional feature space that we then use for classification.

### 2.3.2. MATB-II

We also consider data collected under NASA’s Multi-Attribute Task Battery II (MATB-II) protocol. MATB-II is used to understand a pilot’s ability to perform under varying cognitive load requirements (Santiago-Espada et al., 2011) by attempting to induce four different levels of cognitive load—no (passive), low, medium, and high—that are a function of how many tasks the participant must actively tend to.

The data we consider includes 50 healthy subjects with normal or corrected-to-normal vision. There were 29 female and 21 male participants and each participant was between the ages of 18 and 39 (mean 25.9, std 5.4 years). Each participant was familiarized with MATB-II and then actively underwent for two sessions containing three segments. The three segments were further divided into blocks with the four different levels of cognitive requirements. The sessions lasted  $\sim 50$  min and were separated by a 10 min break.

The EEG data was recorded using a 24-channel Smarting MOBI device and was preprocessed using high pass (0.5 Hz) and low pass (30 Hz) filters and segmented in ten second, non-overlapping windows. In our analysis we consider the two-class problem {no & low cognitive load, medium & high cognitive load} and use the same feature extraction methods as described in Section 2.3.1.

## 2.4. Classification Experiments

We study the three sets of features in in-session, non-constant querying, and transfer learning contexts derived from the two class classification problems {stressed, not stressed} and {no & low cognitive load, medium & high cognitive load}. The experiments attempt to mimic real-world use cases for passive BCIs: the in-session experiment will inform us on how much data we need to be performant for an average participant given no auxiliary data (e.g., data from additional sessions, tasks, or participants); the passive querying experiment gives us an idea of how often a system must query the participant for a label when passively recording EEG signal; and the transferability experiment will quantitatively (*via* balanced accuracy) measure the utility of pre-trained and lightly fine tuned models from other sessions and other subjects. In all, we think this set of experiments provides a useful suite of results to evaluate the considered feature sets and to inspire more nuanced experiments.

For all the experiments we use scikit-learn’s (Pedregosa et al., 2011) implementation of a random forest as the classifier after learning the appropriate set of features. The default hyperparameters (e.g., number of trees, maximum depth of each tree, impurity, etc.) from version 0.22 are used. We report the balanced accuracy for all experiments except for the sensor importance analysis for which we report standard accuracy. Error bars correspond to the standard error of the average accuracy across subjects. We provide a summary table in **Table 1** to more easily compare results across experiments.

**TABLE 1** | Summary table of the in-session, non-constant querying, and transfer results for the Mental Math (“MM”) and MATB-II (“MATB”) datasets.

	MM ( $\rho = 0.0$ )	MM ( $\rho = 0.1$ )	MM ( $\rho = 0.8$ )	MATB ( $\rho = 0.0$ )	MATB ( $\rho = 0.1$ )	MATB ( $\rho = 0.8$ )
TSG-1	n/a	59.2 (1.2)	82.0 (2.4)	n/a	62.6 (1.2)	69.0 (1.5)
TSG-2	n/a	72.1 (2.0)	82.1 (2.4)	n/a	65.9 (1.3)	69.0 (1.4)
TSG-3	n/a	n/a	n/a	66.6 (0.1)	65.6 (0.8)	68.2 (0.8)
TSG-4	53.4 (0.7)	62.0 (1.5)	66.5 (1.8)	58.2 (0.4)	63.2 (0.6)	65.9 (0.6)
BF-1	n/a	58.0 (1.2)	70.9 (2.5)	n/a	68.3 (1.2)	77.4 (1.3)
BF-2	n/a	56.2 (1.0)	70.7 (2.5)	n/a	68.9 (1.2)	77.4 (1.3)
BF-3	n/a	n/a	n/a	71.8 (0.8)	70.3 (0.7)	75.3 (0.8)
BF-4	52.8 (0.5)	64.2 (1.2)	72.9 (1.5)	55.7 (0.2)	63.6 (0.5)	68.8 (0.5)
TSG+BF-1	n/a	59.5 (1.0)	85.1 (2.2)	n/a	70.4 (1.2)	80.3 (1.3)
TSG+BF-2	n/a	61.9 (1.2)	85.2 (2.2)	n/a	72.4 (1.2)	80.4 (1.3)
TSG+BF-3	n/a	n/a	n/a	75.3 (0.9)	74.2 (0.8)	78.8 (0.8)
TSG+BF-4	53.6 (0.7)	66.5 (1.3)	76.5 (1.7)	59.0 (0.4)	68.1 (0.6)	74.3 (0.7)
ANOVA-1	n/a		***	n/a	***	***
ANOVA-2	n/a	***	***	n/a	***	***

Each column corresponds to a different proportion ( $\rho \in \{0.0, 0.1, 0.8\}$ ) of in-distribution data available per the experiments described in the first three subsections of Section 2.4. For each method (TSG, BF, TSG+BF) there are four variations (1 = in-session, 2 = non-constant querying, 3 = cross-session transfer, 4 = cross-subject transfer). Values are average balanced accuracies across subjects. Standard errors are given in parentheses. An “n/a” indicates that the method is not valid for a setting. For example, “BF-1” and “MM ( $\rho = 0$ )” is “n/a” because the method “BF-1” cannot be trained in settings without in-task data. The bottom two rows present one-way repeated measure ANOVAs results comparing three methods (TSG, BF, TSG+BF) at two variations (1 = in-session, 2 = non-constant querying). \*\*\* denotes that  $p < 0.001$  according to one-way repeated measure ANOVAs.

### 2.4.1. In-session Classification

The first experiment we consider is a standard classification experiment: we split the in-session data into a training set and a testing set, learn the feature map and classifiers using only the training set, and evaluate the classifier using the testing set. We estimated the performance using 20 random train-test splits with varying proportions of the data used for training for each subject.

### 2.4.2. Non-constant Querying

The next experiment we consider is inspired by the scenario where a participant is wearing an EEG headset throughout the day and the system has a the capability to query the user for a self-reported {stress, not stressed} or {no & low load, medium & high load} label for the most recent time windows. In these scenarios there is often a non-trivial amount of in-distribution but unlabeled data that can be leveraged to improve the representations ultimately used for inference.

For this experiment, as in the in-session experiment, we randomly split the data into a training set (here always 80%) and a testing set (20%). We then further split the training set into an unlabeled set and labeled set and evaluate systems for varying ratios of labeled data to training data. This means that when the ratio is 0.125 the feature extraction methods have access to 80% of all of the data and the classifier has access to only 10 of the 80%. When the ratio is equal to 1 then 80% of all of the data is used for feature extraction and to train the classifier. The performance is estimated by 20 random train-test splits. Note that here we do not need to change our feature extraction techniques to accommodate unlabeled data because both methods are entirely unsupervised. More involved methods may require treating the labeled and unlabeled data differently (Zhou and Belkin, 2014).

### 2.4.3. Transferability

The last classification experiment we consider investigates the transferability of feature maps and classifiers from one

session to another. Successful zero-shot or few-shot transfer for all participants would enable rapid deployment of real-time prediction systems for a variety of tasks. Unfortunately, as with many biosignals, the inter-session and inter-subject variability of EEG signals means that pre-trained models are unsuited for prediction related to a new participant (Wei et al., 2021). Our transfer experiments are designed to help understand the limitations and possibilities of single-session transfer in the context of mental state prediction.

We consider two transfer paradigms: zero-shot learning (Xian et al., 2017) and classifier fine-tuning. For zero-shot learning we simply take the feature maps and classifiers learned from another session and apply it to a new session. These approaches are rigid and only work well when the distribution shift between sessions is small. For classifier fine-tuning we use the feature map from another session and transform the data from the new session into the appropriate feature space. We then use data from the new session to update the classifier trained on the other session. Since we are using random forests, this simply means taking the existing structure of the trees and updating the posteriors in each of the leaf nodes for each tree with new labeled data. The posteriors of empty leaf nodes are set to 0.5 for simplicity. In the transferability results, the zero-shot performance corresponds to the x-tick 0.0 and the rest of the figure corresponds to classifier fine-tuning. The performance is estimated *via* 20 monte carlo simulations.

## 3. RESULTS

We evaluate three sets of features—band-power based features (“BF”) of Section 2.1, time series of graphs-based features (“TSG”) of Section 2.2, and the concatenation of the two (“TSG + BF”)—in three different settings (in-session classification, non-constant querying, and transfer) in the context of two datasets (“Mental

Math” and “MATB-II”). Our analysis demonstrates the value of including both band-based and graph-based features for high level mental state prediction tasks across multiple experimental settings. After discussing the classification experiments we argue that the graph-based features are neuroscientifically reasonable *via* a channel and pair of channels importance study.

### 3.1. In-session Classification Results

We study the three sets of features in in-session setting. The top panels of **Figure 2** show the average performance across subjects of the three sets of features for varying amounts of data available for training. For the Mental Math study (left), we see the graph-based features outperform the band-based features and their concatenation equal to the best throughout. For the MATB-II study (right), we see the band-based features outperform the graph-based features and their concatenation outperform both for the entirety of the studied regime. The result that TSG > BF for Mental Math but BF > TSG for MATB-II combined with the result that their concatenation is always preferred to either one suggests that the predictive information contained in the TSG and BF features is somewhat complementary—there exists information in one set that is not present in the other.

The bottom panels of **Figure 2** show the histograms of the accuracies for each participant when the feature maps and classifiers have access to 80% of the data for training. The left tail of the histograms shows that the performance is highly dependent on the subject. For example, in the Mental Math study (left) we see a non-trivial number of participants at chance level for each method even with 80% of the data available for training. On the other hand, we also see numerous participants with a balanced accuracy well above 90%. This effect is less exaggerated for the MATB-II data (right) but is still present for the band-based features and the concatenated features.

While not particularly informative of how real passive BCI predictive systems work, the performance of the systems in this experiment is somewhat of an upper bound for what we can expect from a real-time system when there is no additional data from other sessions or subjects. In particular, in this experiment the standard assumption of independence between the training and testing sets is likely violated due to correlations between adjacent windows. This can inflate classification performance (Hand, 2006).

The lines shown in **Figure 2** are estimated averages over all of the subjects. For each subject we estimated their corresponding average using 45 random train-test splits for each amount of training data. The histograms on the bottom row of **Figure 2** show the distribution of these averages when the system has access to 80% of the data.

### 3.2. Non-constant Querying Results

Next, we compare the three sets of features in non-constant querying setting. **Figure 3** shows the average performance across subjects for the three semi-supervised methods. Of note is the performance of TSG on the Mental Math study in the regime where the majority of the data is unlabeled relative to the same amount of labeled data for the in-session experiment (the x-tick

0.125 in **Figure 3** corresponds to the x-tick 0.1 in **Figure 2**). The boost is less impressive in the MATB-II study but still present.

The band-based features are not meaningfully improved with the availability of a set of unlabeled data which likely implies that the principal components of the data are relatively stable across windows. The gain in performance of the concatenated features is somewhere in between the two composite features. We see a small (1–3%) improvement over the corresponding in-session result for the first half of the regime.

The boosts in performance for the graph-based features and the concatenation of the features implies that the data used between queries in real-world systems can be used to improve predictive performance. This means that fewer queries can be made and users will be less likely to tire from self-reporting labels when leveraging the collected but unlabeled data.

This experiment likewise suffers from the effect of adjacent windows being in different splits of the data. The averages in **Figure 3** are estimated using 45 random train-test splits.

### 3.3. Transferability Results

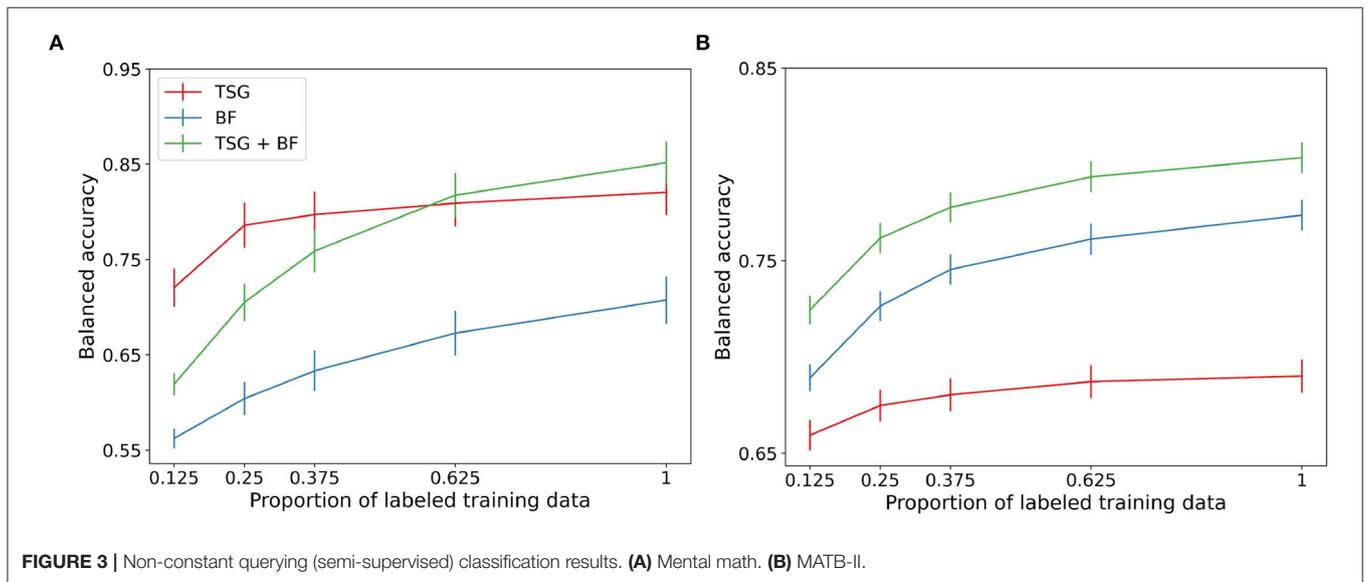
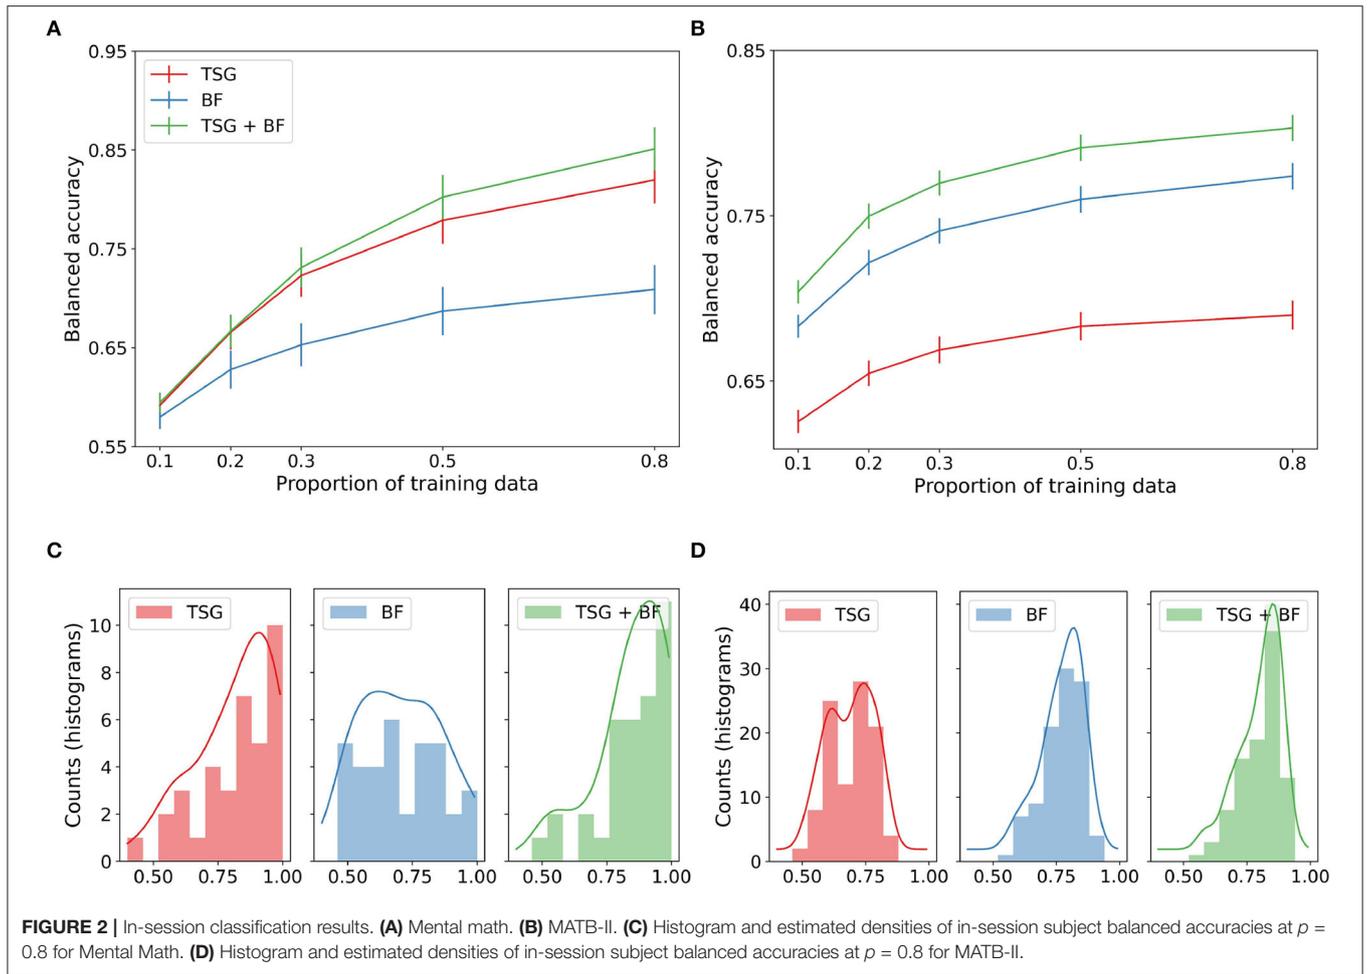
We next present the transferability results of feature maps and classifiers from one session to another. The top panel of **Figure 4** shows the zero-shot and classifier fine tuning results when transferring from one MATB session to another for a given subject. We again plot the average across subjects as in the previous two experiments.

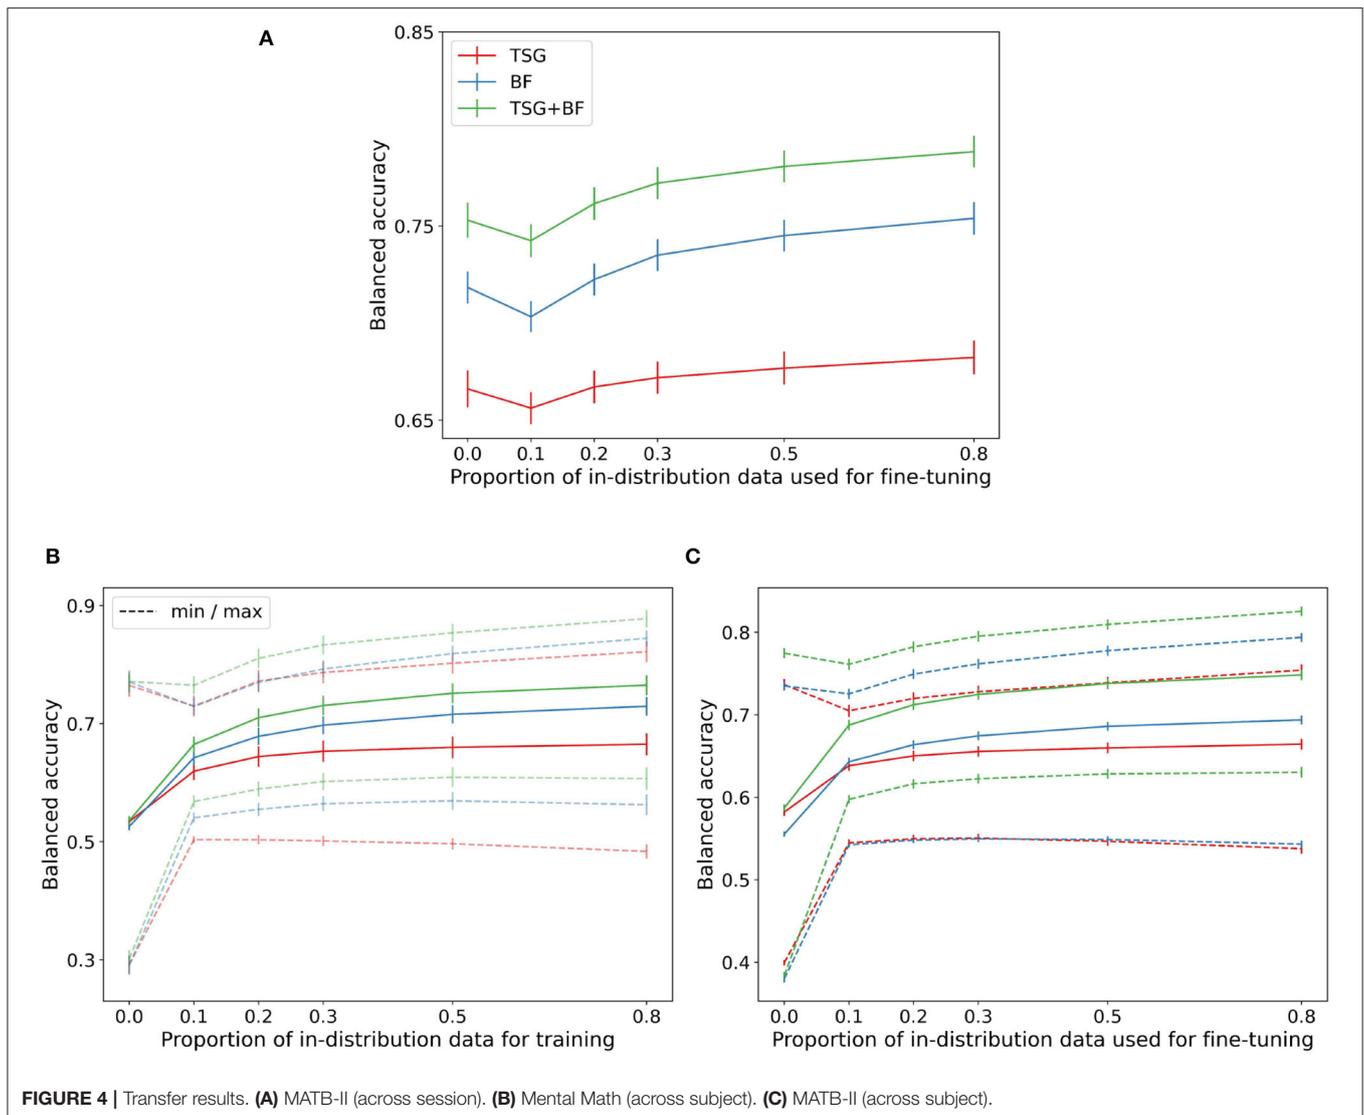
In the MATB experimental design the two sessions are close in time and thus the distribution shift is relatively small as is evident by the good zero-shot performance across methods. For TSG the zero-shot is effectively as good as the fine tuning for the majority of the regime. For all methods the zero-shot performance and is approximately as good as having access to about 30% of the in-session data. Recall that in the MATB experiment each session is  $\sim 50$  min, meaning the zero-shot approach across sessions can save  $\sim 50(0.3) = 15$  min of calibration time and data collection.

The performance across methods degrades non-trivially when going from zero-shot to parts of the regime where only a small amount of training data is available. This dip in performance is again indicative of the proximity of the distributions between sessions and can be smoothed out by continual or lifelong learning techniques (Thrun, 1995; Vogelstein et al., 2020; Geisa et al., 2021). Unfortunately there is no concept of transferring across sessions for the Mental Math study and so there is no corresponding figure.

The bottom row of **Figure 4** contains the cross-subject transfer results for both Mental Math (left) and MATB-II (right). For each subject there are  $n_{subjects} - 1 = 35$  corresponding accuracies for the Mental Math data and  $n_{subjects} \cdot n_{sessions} - 2 = 98$  corresponding accuracies for the MATB-II data. From these accuracies we record the average, minimum accuracy, and maximum accuracy when transferring to each subject. The average across subjects of these three statistics for both datasets are shown in the respective figures.

The range between the averages of the minimums and maximums in both datasets is particularly interesting—if you attempt to use zero-shot or few-shot classification from the wrong subject the performance is severely hampered; on the





other hand, if you can correctly identify which feature map and classifier is “best” you can get a serious improvement gain over the average performance. Selecting which model or set of models to transfer from is an active area of research (Helm et al., 2020) and successful techniques will likely be of unusual importance in biosensing settings to keep on-device models as light, private, and effective as possible.

The dip present in the cross-session figure is not present in the average cross-subject transfer curves (it is present in the “max” curve) and is indicative of the average out of subject distribution being sufficiently far from the distribution we are attempting to transfer to.

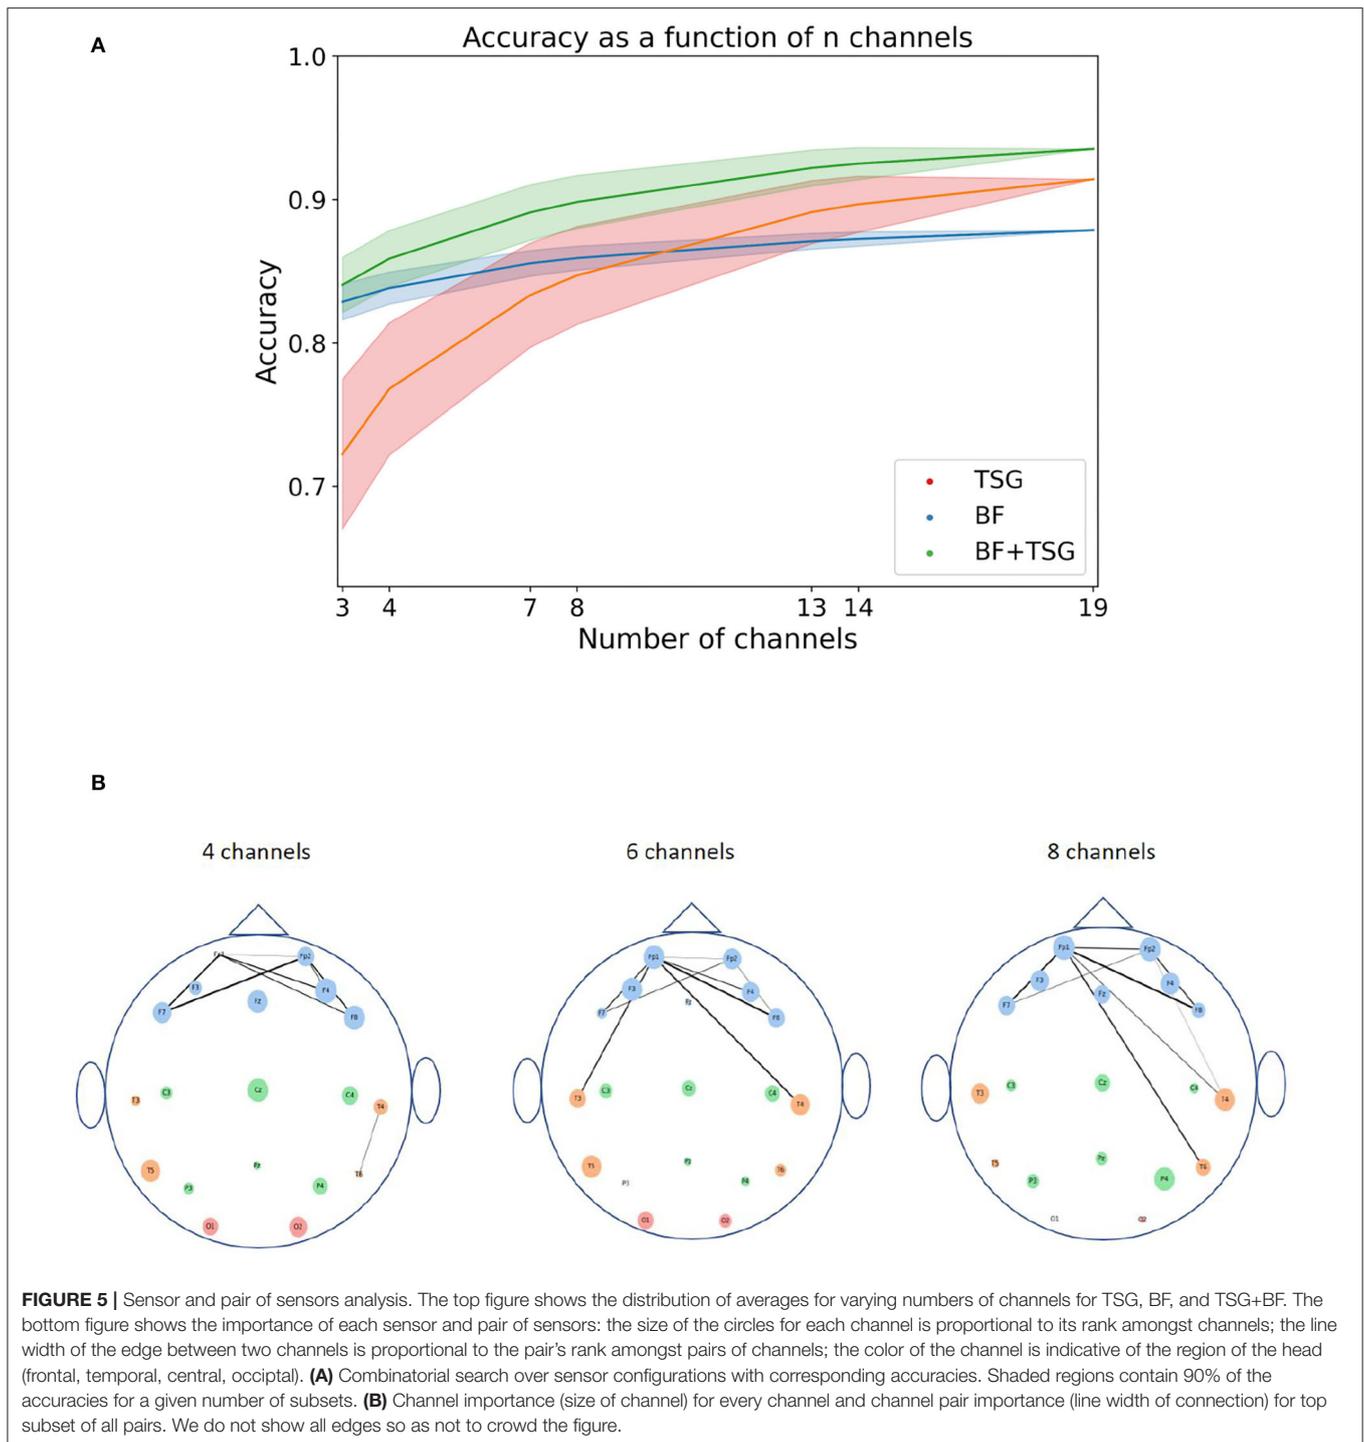
### 3.4. Summary Table

**Table 1** summarizes the results across the three experiments for easier comparison across the different settings. Note that for non-constant querying the x-ticks in **Figure 3** are multiplied by 0.8 to get to the appropriate amount of labeled training data [e.g.,  $0.125 (0.8) = 0.1$ ].

One-way repeated measure ANOVAs are performed to test whether there was a significant difference among the band-power based features (“BF”), time series of graphs-based features (“TSG”), and the concatenation of the two (“TSG + BF”). The results of 20 random train-split Monte Carlo simulations at two different settings are shown in **Table 1**. In the majority of contexts explored above the combination of the band-based and graph-based features out performed either on their own and, again, suggests that both should be used to capture as much predictive information as possible.

### 3.5. Sensor Ablation Analysis

While the classification results of Section 2.4 are interesting in their own right, the sensor locations and device set ups used in the experimental procedures are relatively arbitrary and are not subjected to the same design constraints as a real-world system. Indeed, commercially viable EEG hardware needs to be performant with a relatively minimal head



coverage and number of sensors. To understand the effect of removing sensors on classification performance, we re-ran the in-session classification experiment from Section 2.4.1 with all subsets of {3,4,5,6,7,8,13,14,18,19} channels and recorded their accuracy on the Mental Math data. For each subset we estimated the average accuracy using 30 random train-test splits with 80% of the data used for training for each subject.

**Figure 5A** shows the average accuracy<sup>1</sup> across subsets as a function of the number of channels in the subset. Shaded regions

<sup>1</sup>We performed this analysis before the classification experiments of Section 2.4 and, importantly, before we began considering balanced accuracy instead of standard accuracy. Since the analysis is combinatorial and thus computationally expensive we did not re-run the analysis. While the quantitative results change with the switch to standard accuracy, the qualitative results still hold.

indicate the 90% interval around the mean subset accuracy. The subsets in the right tail of the distribution of subsets are notably close in performance to the full suite of channels for the concatenated features with access to only three or four channels. This implies that a performant system is possible with only a small selection of channels.

We can understand which channels and pairs of channels are driving the high accuracy with such a small subset by estimating their importance. For each channel we first calculate the average accuracy on the stress prediction task for subsets that it is an element of. We then rank the channels based on their average accuracy. We can do the same for pairs of channels. **Figure 5B** is a pictorial description of these rankings for subsets of size 4, 6, and 8: the size of the circles for each channel is proportional to its rank amongst channels and the line width of the edge between two channels is proportional to the pair's rank amongst pairs of channels. The color of the channel is indicative of the region of the head (frontal, temporal, central, occipital).

The channel importance and pair of channel importance for each 4, 6, and 8 channel subsets are well aligned with conventional stress-related neuroscience wisdom. In particular, the average importance of individual channels in the frontal region (blue channels) is noticeably larger than the average importance of the other regions (Yuen et al., 2009). Further, there are a considerable amount of heavily weighted edges between frontal channels on one side of the brain to frontal or temporal channels on the other side of the brain (Fischer et al., 2018). Hence, we think that including the pairwise correlation matrix as the basis for a set of features is neuroscientifically reasonable.

## 4. DISCUSSION AND CONCLUSION

In this note we proposed a set features, TSG, that are jointly learned from the correlation matrices of different windows. We showed in multiple classification paradigms that the TSG features can be used to significantly improve performance over canonical band-based features, BF, in both stress and cognitive load contexts. We noted that neither the TSG features nor the BF features dominate the other across the different experiments and thus argued that these two sets of features are “complementary” or that one set of features contains predictive information that the other does not. Given these results we think that graph-based features will be useful in pushing the community closer to performant non-invasive BCIs.

We note that the effectiveness of prediction results is restrained to the context of our pipeline. We note that though we interpret the results below as if the labels used to train the classifiers are only related to stress and cognitive load, there is a possibility of non-trivial confounding variables that could artificially improve performance on the intended task. We do not pursue a formal analysis of these confounding variables herein.

We also want to emphasize that the proposed set of features are but one instance of a class of features based on the implicit graph structure on the channels. We plan to investigate other natural similarities on the channels beyond just Pearson correlation, e.g., mutual information

and conditional entropy, and think that these more nuanced statistical dependence measures could capture more meaningful class-conditional information.

Further, the similarity that we use to induce a graph on the collection of statistical-dependence matrices is Frobenious norm based and as simple as possible. It is well known in the network statistics community that the Frobenious norm between matrices is too crude of a metric to uncover particular asymptotic behaviors of the graph (Athreya et al., 2018; Kriege et al., 2020). Indeed, often dimensionality reduction techniques are used on the collection of graphs before calculating the distance between the objects. We think that including a multi-graph embedding step such as the omnibus embedding (Levin et al., 2017) or the COSIE framework (Arroyo et al., 2021) could non-trivially improve performance over the results presented here. Similarly, we think applying different kernels on the networks, such as the Gaussian kernel, could be a fruitful direction.

## DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found at: <https://physionet.org/content/eegmat/1.0.0/>.

## ETHICS STATEMENT

The studies involving human participants were reviewed and approved by the Bioethics Commission of Educational and Scientific Centre Institute of Biology and Medicine, Taras Shevchenko National University of Kyiv for the Mental Math dataset, and by the Ethical Committee, Faculty of Medical Sciences University of Kragujevac, Kragujevac, Serbia for the MATB-II dataset. The patients/participants provided their written informed consent to participate in this study.

## AUTHOR CONTRIBUTIONS

GC, HH, WY, and CP contributed to the conception and design of the study and methodology development. HH, KL, and WY organized the datasets. GC and HH performed the formal analysis and visualization and wrote the manuscript. WY and CP supervised the project. All authors contributed to manuscript revision, read, and approved the submitted version.

## FUNDING

This work was supported in part by Microsoft Research.

## ACKNOWLEDGMENTS

We would like to thank Ben Cutler, Steven Dong, Amber Hoak, Tzzy-Ping Jung, Bogdan Mijovic, Ben Pedigo, Siddharth Siddharth, and Christopher White for helpful comments and suggestions throughout our investigations.

## REFERENCES

- Ang, K. K., Chin, Z. Y., Zhang, H., and Guan, C. (2008). "Filter bank common spatial pattern (FBCSP) in brain-computer interface," in *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)* (Hongkong), 2390–2397.
- Arroyo, J., Athreya, A., Cape, J., Chen, G., Priebe, C. E., and Vogelstein, J. T. (2021). Inference for multiple heterogeneous networks with a common invariant subspace. *J. Mach. Learn. Res.* 22, 1–49.
- Athreya, A., Fishkind, D. E., Tang, M., Priebe, C. E., Park, Y., Vogelstein, J. T., et al. (2018). Statistical inference on random dot product graphs: a survey. *J. Mach. Learn. Res.* 18, 1–92. doi: 10.48550/arXiv.1709.05454
- Bengio, Y., Paiement, J.-F., Vincent, P., Delalleau, O., Roux, N., and Ouimet, M. (2003). Out-of-sample extensions for LLE, Isomap, MDS, Eigenmaps, and spectral clustering. *Adv. Neural Inform. Process. Syst.* 16, 177–184. doi: 10.1162/0899766041732396
- Blankertz, B., Tomioka, R., Lemm, S., Kawanabe, M., and Muller, K.-R. (2007). Optimizing spatial filters for robust EEG single-trial analysis. *IEEE Signal Process. Mag.* 25, 41–56. doi: 10.1109/MSP.2008.4408441
- Congedo, M., Barachant, A., and Bhatia, R. (2017). Riemannian geometry for eeg-based brain-computer interfaces; a primer and a review. *Brain Comput. Interfaces* 4, 155–174. doi: 10.1080/2326263X.2017.1297192
- Fischer, N. L., Peres, R., and Fiorani, M. (2018). Frontal alpha asymmetry and theta oscillations associated with information sharing intention. *Front. Behav. Neurosci.* 12:166. doi: 10.3389/fnbeh.2018.00166
- Geisa, A., Mehta, R., Helm, H. S., Dey, J., Eaton, E., Dick, J., et al. (2021). Towards a theory of out-of-distribution learning. *arXiv [Preprint]*. arXiv: 2109.14501. Available online at: <https://arxiv.org/pdf/2109.14501.pdf> (accessed January 6, 2022).
- Gopalakrishnan, V., Chung, J., Bridgeford, E., Pedigo, B. D., Arroyo, J., Upchurch, L., et al. (2020). Multiscale comparative connectomics. *arXiv [Preprint]*. arXiv: 2011.14990. Available online at: <https://arxiv.org/pdf/2011.14990.pdf> (accessed April 13, 2022).
- Hand, D. J. (2006). Classifier technology and the illusion of progress. *Stat. Sci.* 21, 1–14. doi: 10.1214/088342306000000060
- Helm, H. S., Mehta, R. D., Duderstadt, B., Yang, W., White, C. M., Geisa, A., et al. (2020). A partition-based similarity for classification distributions. *arXiv [Preprint]*. arXiv: 2011.06557. Available online at: <https://arxiv.org/pdf/2011.06557.pdf>
- Kamzanova, A., Matthews, G., Kustubayeva, A., and Jakupov, S. (2011). EEG indices to time-on-task effects and to a workload manipulation (cueing). *Int. J. Psychol. Behav. Sci.* 5, 928–931. doi: 10.5281/zenodo.1071802
- Kaya, M., Binli, M. K., Ozbay, E., Yanar, H., and Mishchenko, Y. (2018). A large electroencephalographic motor imagery dataset for electroencephalographic brain computer interfaces. *Sci. Data* 5, 1–16. doi: 10.1038/sdata.2018.211
- Koles, Z. J. (1991). The quantitative extraction and topographic mapping of the abnormal components in the clinical EEG. *Electroencephalogr. Clin. Neurophysiol.* 79, 440–447. doi: 10.1016/0013-4694(91)90163-X
- Kriege, N. M., Johansson, F. D., and Morris, C. (2020). A survey on graph kernels. *Appl. Netw. Sci.* 5, 1–42. doi: 10.1007/s41109-019-0195-3
- Lawhern, V. J., Solon, A. J., Waytowich, N. R., Gordon, S. M., Hung, C. P., and Lance, B. J. (2018). EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces. *J. Neural Eng.* 15:056013. doi: 10.1088/1741-2552/aace8c
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. doi: 10.1038/nature14539
- Levin, K., Roosta, F., Mahoney, M., and Priebe, C. (2018). "Out-of-sample extension of graph adjacency spectral embedding," in *International Conference on Machine Learning* (Stockholm: Stockholmsträskö, 2975–2984).
- Levin, K., Athreya, A., Tang, M., Lyzinski, V., Park, Y., and Priebe, C. E. (2017). A central limit theorem for an omnibus embedding of random dot product graphs. *arXiv/1705.09355*. doi: 10.1109/ICDMW.2017.132
- Lotte, F., Bougrain, L., Cichocki, A., Clerc, M., Congedo, M., Rakotomamonjy, A., et al. (2018). A review of classification algorithms for EEG-based brain-computer interfaces: a 10 year update. *J. Neural Eng.* 15:031005. doi: 10.1088/1741-2552/aab2f2
- Miller, K. J., Schalk, G., Fetz, E. E., Den Nijs, M., Ojemann, J. G., and Rao, R. P. (2010). Cortical activity during motor execution, motor imagery, and imagery-based online feedback. *Proc. Natl. Acad. Sci. U.S.A.* 107, 4430–4435. doi: 10.1073/pnas.0913697107
- Nayak, C. S., and Anilkumar, A. C. (2021). *EEG Normal Waveforms*. Treasure Island, FL: StatPearls Publishing.
- Noto, Y., Sato, T., Kudo, M., Kurata, K., and Hirota, K. (2005). The relationship between salivary biomarkers and state-trait anxiety inventory score under mental arithmetic stress: a pilot study. *Anesth. Anal.* 101, 1873–1876. doi: 10.1213/01.ANE.0000184196.60838.8D
- Oikonomou, V. P., Georgiadis, K., Liaros, G., Nikolopoulos, S., and Kompatsiaris, I. (2017). "A comparison study on EEG signal processing techniques using motor imagery EEG data," in *2017 IEEE 30th International Symposium on Computer-Based Medical Systems (CBMS)* (Thessaloniki), 781–786. doi: 10.1109/CBMS.2017.113
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). SCIKIT-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830. doi: 10.48550/arXiv.1201.0490
- Santiago-Espada, Y., Myer, R. R., Latorella, K. A., and Comstock, J. R. (2011). *The Multi-Attribute Task Battery II (MATB-II) Software for Human Performance and Workload Research: A User's Guide (NASA/TM-2011-217164)*. Hampton, VA: National Aeronautics and Space Administration; Langley Research Center.
- Siddharth, Jung, T. -P., and Sejnowski, T. J. (2022). "Utilizing deep learning towards multi-modal bio- sensing and vision-based affective computing" in *IEEE Transactions on Affective Computing* (IEEE), 96–107. doi: 10.1109/TAFFC.2019.2916015
- Simonyan, K., and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv [Preprint]* arXiv:1409.1556. doi: 10.48550/arXiv.1409.1556
- Spüler, M., Rosenstiel, W., and Bogdan, M. (2012). Online adaptation of a C-VEP Brain-computer interface (BCI) based on error-related potentials and unsupervised learning. *PLoS ONE* 7:e51077. doi: 10.1371/journal.pone.0051077
- Tang, M., Park, Y., and Priebe, C. E. (2013). Out-of-sample extension for latent position graphs. *arXiv [Preprint]* arXiv:1305.4893. doi: 10.48550/arXiv.1305.4893
- Thrun, S. (1995). "Is learning the n-th thing any easier than learning the first?," in *Proceedings of the 8th International Conference on Neural Information Processing Systems*, eds P. S. R. Diniz, J. A. K. Suykens, R. Chellappa, and S. Theodoridis (Oxford: Academic Press), pp. 640–646.
- Trosset, M. W., and Priebe, C. E. (2008). The out-of-sample problem for classical multidimensional scaling. *Comput. Stat. Data Anal.* 52, 4635–4642. doi: 10.1016/j.csda.2008.02.031
- Vogelstein, J. T., Dey, J., Helm, H. S., LeVine, W., Mehta, R. D., Geisa, A., et al. (2020). Representation ensembling for synergistic lifelong learning with quasilinear complexity. *arXiv [Preprint]*. arXiv: 2004.12908. Available online at: <https://arxiv.org/pdf/2004.12908.pdf> (accessed May 14, 2022).
- Wang, N., Anderson, R. J., Ashbrook, D. G., Gopalakrishnan, V., Park, Y., Priebe, C. E., et al. (2019). Node-specific heritability in the mouse connectome. *bioRxiv*. doi: 10.1101/701755
- Wei, C.-S., Keller, C., Li, J., Lin, Y.-P., Nakanishi, M., Wagner, J., et al. (2021). Inter- and intra-subject variability in brain imaging and decoding. *Front. Comput. Neurosci.* 15:791129. doi: 10.3389/fncom.2021.791129
- Wei, Q., Wang, Y., Gao, X., and Gao, S. (2007). Amplitude and phase coupling measures for feature extraction in an EEG-based brain-computer interface. *J. Neural Eng.* 4:120. doi: 10.1088/1741-2560/4/2/012
- Wu, C.-T., Huang, H.-C., Huang, S., Chen, I., Liao, S.-C., Chen, C.-K., et al. (2021). Resting-state EEG signal for major depressive disorder detection: a systematic validation on a large and diverse dataset. *Biosensors* 11:499. doi: 10.3390/bios11120499
- Xian, Y., Schiele, B., and Akata, Z. (2017). "Zero-shot learning-the good, the bad and the ugly," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Honolulu, HI), 4582–4591. doi: 10.1109/CVPR.2017.328
- Yger, F., Berar, M., and Lotte, F. (2016). Riemannian approaches in brain-computer interfaces: a review. *IEEE Trans. Neural Syst. Rehabil. Eng.* 25, 1753–1762. doi: 10.1109/TNSRE.2016.2627016
- Yuen, E. Y., Liu, W., Karatsoreos, I. N., Feng, J., McEwen, B. S., and Yan, Z. (2009). Acute stress enhances glutamatergic transmission

- in prefrontal cortex and facilitates working memory. *Proc. Natl. Acad. Sci. U.S.A.* 106, 14075–14079. doi: 10.1073/pnas.0906791106
- Zhou, X., and Belkin, M. (2014). “Semi-supervised learning,” in *Academic Press Library in Signal Processing, Vol. 1*, eds P. S. R. Diniz, J. A. K. Suykens, R. Chellappa, S. Theodoridis (Waltham, MA: Elsevier), 1239–1269. doi: 10.1016/B978-0-12-396502-8.0022-X
- Zhu, M., and Ghodsi, A. (2006). Automatic dimensionality selection from the scree plot via the use of profile likelihood. *Comput. Stat. Data Anal.* 51, 918–930. doi: 10.1016/j.csda.2005.09.010
- Zyma, I., Tukaev, S., Seleznov, I., Kiyono, K., Popov, A., Chernykh, M., et al. (2019). Electroencephalograms during mental arithmetic task performance. *Data* 4:14. doi: 10.3390/data4010014

**Conflict of Interest:** HH, KL, and WY were employed by the company Microsoft. GC and CP's research was partially funded by Microsoft.

**Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Chen, Helm, Lytvynets, Yang and Priebe. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.