A SPARSE REDUCED RANK FRAMEWORK FOR GROUP ANALYSIS OF FUNCTIONAL NEUROIMAGING DATA

Mihye Ahn, Haipeng Shen, Weili Lin and Hongtu Zhu

University of North Carolina at Chapel Hill

Supplementary Material

S1 Simulation Studies

In this section, we use simulation studies to illustrate the performance of the proposed SRR, in comparison with three competing methods: K-SVD (Aharon, Elad, and Bruckstein (2006)), FastICA with spatial concatenation (Hyvärinen and Oja (2000)), and GIFT for group spatial ICA (Calhoun et al. (2001b)). We compare the methods from the following four aspects; (i) estimation of model rank, (ii) identification of frequencies of interest, (iii) estimation of frequency bases and spatial maps, and (iv) detection of group difference.

S1.1 Simulation Setup

We assume that there are two groups (e.g., disease and control groups), with fifty subjects in each group, indexed by $s = 1, \ldots, 50$ and $g = 1, 2$. For subject s of group g, we simulate the 200×400 image data $\mathbf{Z}_g^s$ according to

$$\mathbf{Z}_g^s = \mathbf{A}_g^s \mathbf{S}_g^s + \mathbf{E}_g^s, \tag{S1.1}$$

where the temporal matrix $\mathbf{A}_g^s$ is 200×5 containing five time courses across 200 time points and the spatial matrix $\mathbf{S}_g^s$ is 5×400 with five 20×20 two-dimensional spatial maps. See Figure 1 for a graphical illustration. We consider two scenarios of signal-to-noise ratio (SNR) as described below, and repeat the simulation 50 times under each scenario.

Simulating $\mathbf{A}_g^s$. For subject s of group g, the ith time course ($i = 1, \ldots, 5$) is generated as follows:

$$\mathbf{a}_{ti}^{s,g} = \tau_{i1}^{s,g} \cos(2\pi w_i t) + \tau_{i2}^{s,g} \sin(2\pi w_i t) \quad \text{for time} \quad t = 1, \ldots, 200,$$

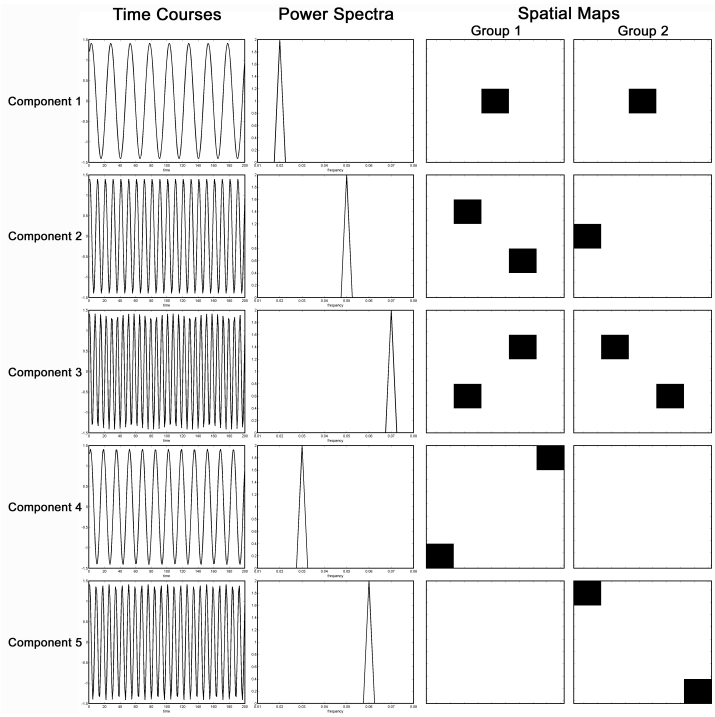

Figure 1: Exemplary time courses of simulation studies: power spectra and spatial maps. Components 1-5 are ordered from top to bottom.

where $\mathbf{a}_{ti}^{s;g}$ is the (t,i)th entry of $\mathbf{A}_g^s$ in (S1.1), $\tau_{11}^{s,g},\ldots,\tau_{52}^{s,g}$ are *i.i.d.* $N(0,\ 1)$, and the frequency parameters are $(w_1,\ldots,w_5) = (0.02, 0.05, 0.07, 0.03, 0.06)$. The random τ's enable subject-specific amplitude and phase for the time courses. The frequencies are chosen within the range of (0.009, 0.08), that is widely used in a band-pass filter for preprocessing resting-state fMRI data (Biswal (1995)). As an example, five time courses corresponding to one subject along with the corresponding power spectra are shown in the first two columns of Figure 1.

Simulating $\mathbf{S}_g^s$. The last two columns of Figure 1 depict the five 20-by-20 spatial maps for two groups, respectively. We code the activated regions as 1 (colored as black) and the non-activated regions as 0 (colored as white). We comment that: for group 1, the first four spatial maps remain the same across all 50 subjects, while for the fifth spatial map, we randomly select 6% of the entries as activated for each subject; for group 2, a similar subject-to-subject variation is introduced to the fourth spatial map, with the other four spatial maps shared by all 50 subjects in the group. For each subject, we then vectorize the five 20-by-20 image maps into the 5×400 spatial map matrix $\mathbf{S}_g^s$ in (S1.1).

Simulating $\mathbf{E}_g^s$. We then simulate the noise terms $\mathbf{E}_g^s$ under two scenarios of SNR: SNR=1 or 4. SNR is defined as the ratio of σ_S to σ_N, where σ_S is the standard deviation

(SD) of the true signal, computed according to Allen et al. (2012), and σ_N is the SD of the noise. We then generate the noise terms from a zero-mean normal distribution with the SD being $\sigma_N = \sigma_S/\text{SNR}$.

S1.2 Competing Methods

We now describe how the four competing methods are implemented. Note that our SRR is performed in the frequency domain, while the other three methods, K-SVD/FastICA/GIFT, are carried out in the time domain.

For SRR, we first convert the simulated data matrices $\mathbf{Z}_g^s$ into the frequency domain, and compute the scaled periodogram (Shumway and Stoffer (2011)). To focus on low frequency fluctuations, we extract the power spectra corresponding to the frequency range of (0.009, 0.08), which results in the 29×400 data matrices, $\mathbf{Y}_g^s$, as in (2.1). SRR is then applied to $\mathbf{Y}_g^s$ to obtain the estimates $\widehat{\mathbf{U}}$ and $\widehat{\mathbf{M}}_g^s$.

To apply K-SVD and FastICA, we first concatenate $\mathbf{Z}_g^s$ along the spatial direction to obtain the concatenated matrix $\mathbf{Z}$. Both methods can decompose $\mathbf{Z}$ into a common temporal signal matrix $\widehat{\mathbf{A}}$ and a concatenated spatial map matrix $\widehat{\mathbf{M}}$. We can then compute the power spectra matrix $\widehat{\mathbf{U}}$ from the estimate $\widehat{\mathbf{A}}$, and divide $\widehat{\mathbf{M}}$ into the subject-specific spatial maps $\widehat{\mathbf{M}}_g^s$.

Finally, GIFT is applied on the temporally concatenated matrix to obtain subject-specific temporal signal matrices and a common spatial map. Through the built-in back-reconstruction step, GIFT also provides subject-specific spatial maps $\widehat{\mathbf{M}}_g^s$. The estimated temporal signal matrices $\widehat{\mathbf{A}}_g^s$ are transformed to the frequency domain in order to derive $\widehat{\mathbf{U}}_g^s$.

For all the methods, we match and sort the resulting components using the correlations between true and estimated power spectra. Below we only report the results under SNR=4 because similar performances are observed when SNR=1.

S1.3 Results when the model rank is estimated

Performance of Estimating Model Rank. We first assess the performance of our data-driven BIC_R statistic (2.18) for estimating the model rank. For comparison, we consider the minimum description length (MDL) criterion implemented in GIFT (Li, Adali, and Calhoun (2007)). Both K-SVD and FastICA do not have the capability for estimating model rank.

Table 1 reports the estimated ranks, along with standard errors in parenthesis, under SNR=∞, 1, and 4. SNR=∞ corresponds to the noise-less cases, i.e. setting $\mathbf{E}_g^s = 0$ in (S1.1) and no subject-specific activation in the fifth spatial map for group 1 and the fourth spatial map for group 2 in Figure 1. In all the 50 simulation replications, BIC_R and MDL consistently choose the rank as 6 and 3, respectively. MDL for GIFT

Table 1: Performance of model rank estimation.

	SNR=∞	SNR=1	SNR=4
True Rank	5	5	5
SRR (BIC$_R$)	5	6 (0)	6 (0)
GIFT (MDL)	1	3 (0)	3 (0)

underestimates the model rank under all SNRs. Our BIC$_R$ works quite well, only slightly overestimating the rank for noisy data with no variability. Unshown results suggest that the additional component obtained by SRR is the overall average.

Influence of Misspecifying Model Rank. First, we compare the reconstructed images for each group by calculating the correlations between the true images and the corresponding estimates. Figure 2 compares the boxplots of such correlations given by SRR and GIFT for each group, respectively. SRR shows much higher correlations than those by GIFT for both groups with less variability. As described in Section S1.1, there are five frequencies of interest: 0.02, 0.05, 0.07, 0.03, and 0.06 (in order of the temporal components). We next compare the spatial maps corresponding to each frequency of interest. For each simulation replication, we first compute the average of the reconstructed images within each group, and then extract the resulting average spatial map corresponding to each of the five frequencies, which are further averaged over all replications. Figure 3 compares the true spatial maps with the estimates from SRR and GIFT for the five frequencies of interest from top to bottom. The results suggest that SRR gives clear spatial maps that are close to the true image maps, whereas GIFT only performs well when there is no group difference (at frequency 0.02, first row).

S1.4 Results using the True Model Rank

From now on, we assume that the true model rank 5 is known. We apply all methods to estimate the rank-5 model, and compare the performance of estimation from various perspectives.

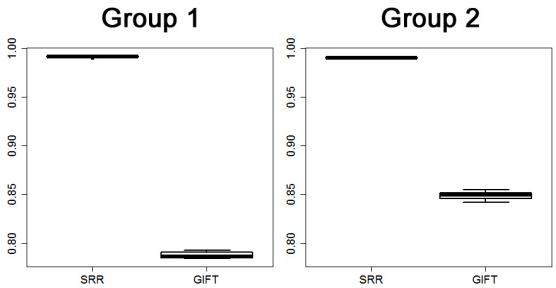

Figure 2: Correlation between reconstructed and true images when the model rank is estimated.

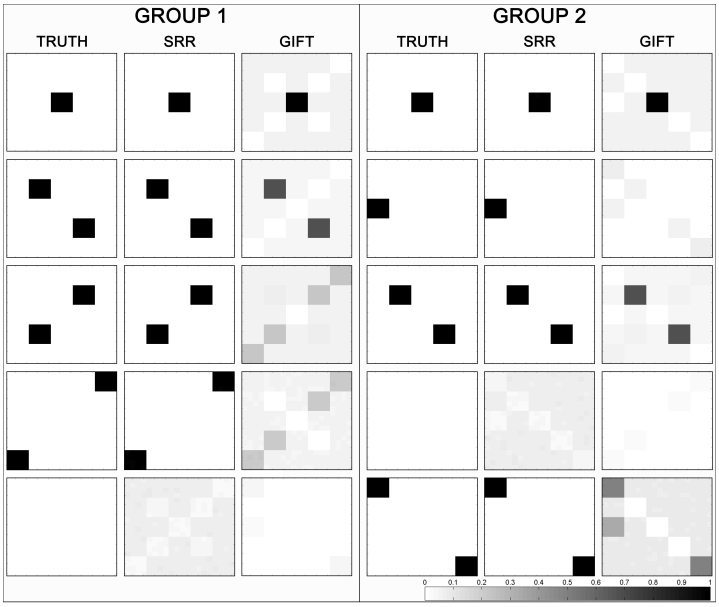

Figure 3: Spatial maps at frequencies of interest when the model rank is estimated (left panels: Group1, right panels: Group2, from top to bottom: $f = 0.02, 0.05, 0.07, 0.03, 0.06$)

Identification of Frequencies of Interest. We first evaluate how well the four methods can detect the five frequencies of interest. Figure 4 displays the average power spectra matrices as images in gray scale, along with the truth. The y-axis highlights the five frequencies of interest from top to bottom, while the x-axis indexes the five extracted frequency components, i.e. the columns in $\widehat{\mathbf{U}}$. As one can see, SRR performs the best. Note that the 4th and 5th components are mixed, which makes sense as the frequencies 0.03 and 0.06 are aliasing. K-SVD fails to detect the first component ($f = 0.02$) and also shows some noise on the other components. GIFT and FastICA perform comparably well, but many components are confounded with each other.

Accuracy of Estimating Frequency Bases and Spatial Maps. The first panel of Figure 5 compares the boxplots of correlations between the true power spectra

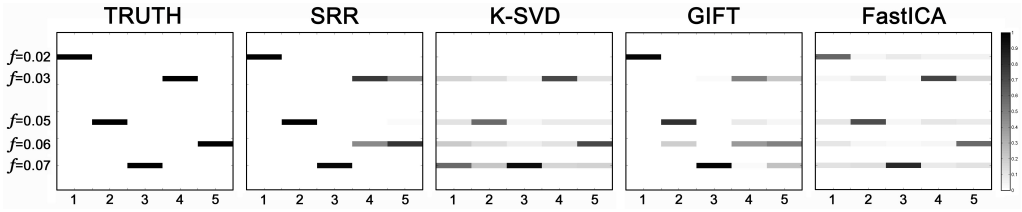

Figure 4: Average power spectra matrices.

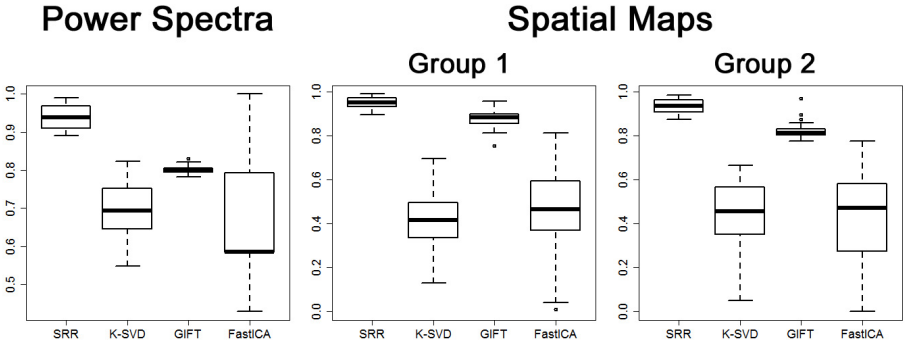

Figure 5: Power spectra and spatial maps: correlation between estimates and truth.

matrix $\mathbf{U}$ and the estimate $\widehat{\mathbf{U}}$ obtained from each method. SRR performs the best, with the correlations ranging from 0.9 to 1.0. FastICA gives very variable results with the lowest median correlation value. K-SVD yields the second smallest correlations with large variation. GIFT gives quite stable performance, but the median correlation is only around 0.8. Note that GIFT does not estimate common temporal factors, hence we instead compute the correlation between the true power spectra and $\widehat{\mathbf{U}}_g^s$ for all subjects, which explains its small variability.

The middle and right panels of Figure 5 compare the boxplots of the correlation between the true spatial maps and the corresponding estimates $\widehat{\mathbf{M}}_g$ for each group g, respectively. K-SVD and FastICA are inferior to SRR and GIFT in terms of variation and accuracy. SRR outperforms GIFT, although both methods have small variability.

Figure 6 compares each component of the spatial maps. For each method, the displayed $\widehat{\mathbf{M}}_g$'s are averaged over all replications. Again, SRR is the winner, although there is some slight confounding between the last two components, due to the aliasing of the corresponding frequencies (0.03 and 0.06). For most of the components, K-SVD fails to identify the signals in both groups. The results of FastICA show some noise. GIFT shows weak signals and some noise, especially in the second and fifth components in group 2.

Testing Group Difference Finally, we conduct hypothesis testing for group difference based on the estimated spatial maps $\widehat{\mathbf{M}}_g^s$, according to Model (2.4). We adjust the entry-wise p-values using the Bonferroni correction for multiple comparisons, before computing the average over all replications. Figure 7 shows the true spatial group difference maps, along with the maps of $-\log_{10}(p\text{-value})$ for GIFT and SRR. Note that the p-value maps are thresholded at $\alpha = .0001$. Both K-SVD and FastICA yield very large p-values for all of five components, which means no significant difference between two groups, so we omit their p-value maps here.

According to the first row of Figure 7, there is no group difference in the first component, while the other components show some differences between two groups. The p-values maps in the 2nd and 3rd rows suggest that SRR performs much better than

GIFT, and correctly detects the spatial regions in which activation is different between two groups.

S2 Proof of Equations (2.10), (2.12), and (2.13)

Minimizing (2.11) is equivalent to minimizing

$$\mathcal{L}_{ij} = \|\mathbf{k}_{i,j} - \widehat{\mathbf{m}}_i u_{ji}\|^2 + \lambda_i |u_{ji}| \quad \text{for} \quad i = 1, \ldots, q \quad \text{and} \quad j = 1, \ldots, T.$$

When $\lambda_i = 0$, we have the OLS estimate of u_{ji}, given by $\widehat{u}_{ji}^{OLS} = < \mathbf{k}_{i,j}, \widehat{\mathbf{m}}_i > / \|\widehat{\mathbf{m}}_i\|_2^2$. It then follows that

$$\mathcal{L}_{ij} = \|\mathbf{k}_{i,j}\|_2^2 - 2\widehat{u}_{ji}^{OLS} \|\widehat{\mathbf{m}}_i\|_2^2 u_{ji} + \|\widehat{\mathbf{m}}_i\|_2^2 u_{ji}^2 + \lambda_i |u_{ji}|.$$

Taking derivative with respect to u_{ji}, we have $-2\widehat{u}_{ji}^{OLS} \|\widehat{\mathbf{m}}_i\|_2^2 + 2\|\widehat{\mathbf{m}}_i\|_2^2 u_{ji} + \lambda_i \text{sgn}(u_{ji}) = 0$. When $\widehat{u}_{ji}^{OLS} > 0$, we have $u_{ji} \geq 0$. Also, when $\widehat{u}_{ji}^{OLS} \leq 0$, we have $u_{ji} \leq 0$. Thus, it follows that

$$u_{ji} = \widehat{u}_{ji}^{OLS} - \frac{\lambda_i \text{sgn}(u_{ji})}{2\|\widehat{\mathbf{m}}_i\|_2^2}.$$

The minimizer of (2.11) is achieved at $\widetilde{u}_{ji} = \text{sgn}(\widehat{u}_{ji}^{OLS})(|\widehat{u}_{ji}^{OLS}| - \lambda_i / 2\|\widehat{\mathbf{m}}_i\|_2^2)_+$ for $i = 1, \ldots, q$ and $j = 1, \ldots, T$.

Equation (2.10) is a special case of Equation (2.13) when $i = 1$.

References

Aharon, M., Elad, M., and Bruckstein, A. (2006). K-SVD: an algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.* **54**, 4311-4322.

Allen, E., Erhardt, E., Wei, Y., Eichele, T., and Calhoun, V. (2012). Capturing inter-subject variability with group independent component analysis of fMRI data: a simulation study. *NeuroImage* **59**, 4141-4159.

Biswal, B., Yetkin, F., Haughton, V., and Hyde, J. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* **34**, 537–541.

Calhoun, V., Adali, T., Pearlson, G., and Pekar, J. (2001b). A method for making group inferences from functional MRI data using independent component analysis. *Hum. Brain Mapp.* **14**, 140–151.

Hyvärinen, A. and Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural Networks* **13**, 411–430.

Li, Y.-O., Adali, T., and Calhoun, V. (2007). Estimating the number of independent components for functional magnetic resonance imaging data. *Hum. Brain Mapp.*, **28**, 1251-1266.

Shumway, R. and Stoffer, D. S. (2011). *Time Series Analysis and Its Application with R Examples* (3rd ed.). Springer, New York.

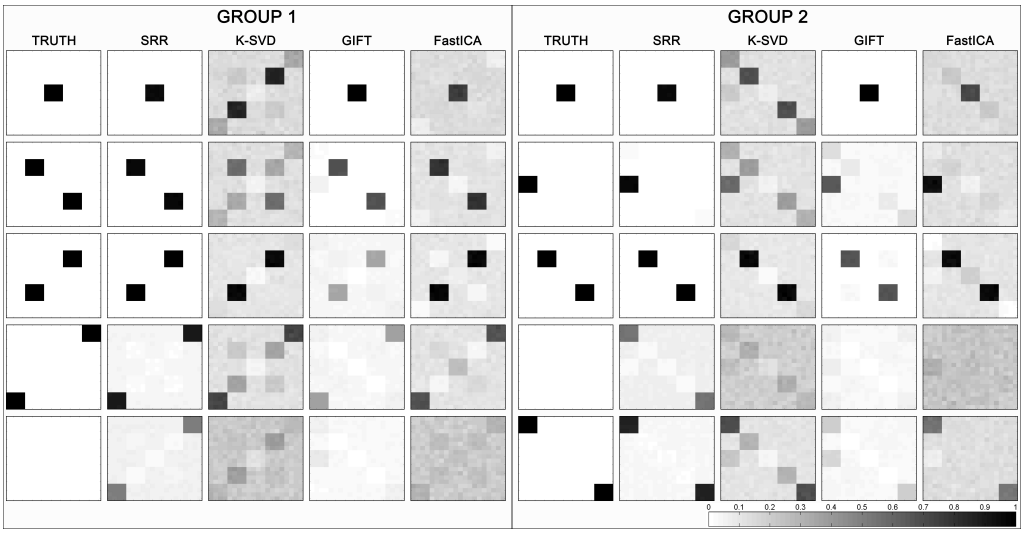

Figure 6: Group-level spatial maps. Components 1-5 are ordered from top to bottom.

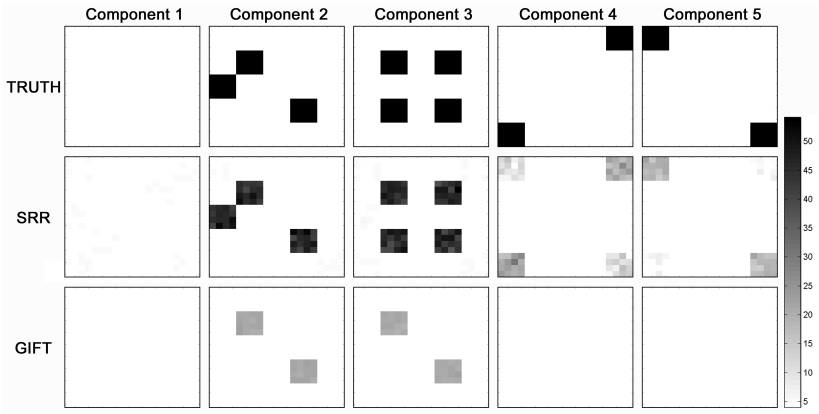

Figure 7: Testing group difference: true difference maps and the $-\log_{10}(p\text{-value})$ maps of SRR and GIFT.

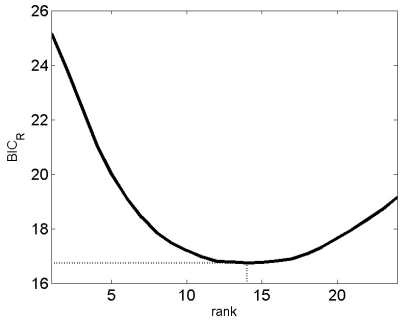

Figure 8: BIC_R curve for rank determination.

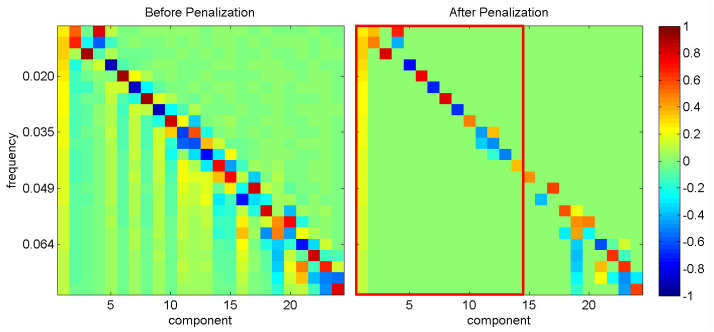

Figure 9: Heat maps of the estimated $\mathbf{U}$ matrices. Left: $\widehat{\mathbf{U}}$ before penalization, Right: $\widetilde{\mathbf{U}}$ after penalization. The red square shows the final estimate of the $\mathbf{U}$ matrix of rank 14 determined by BIC_R.

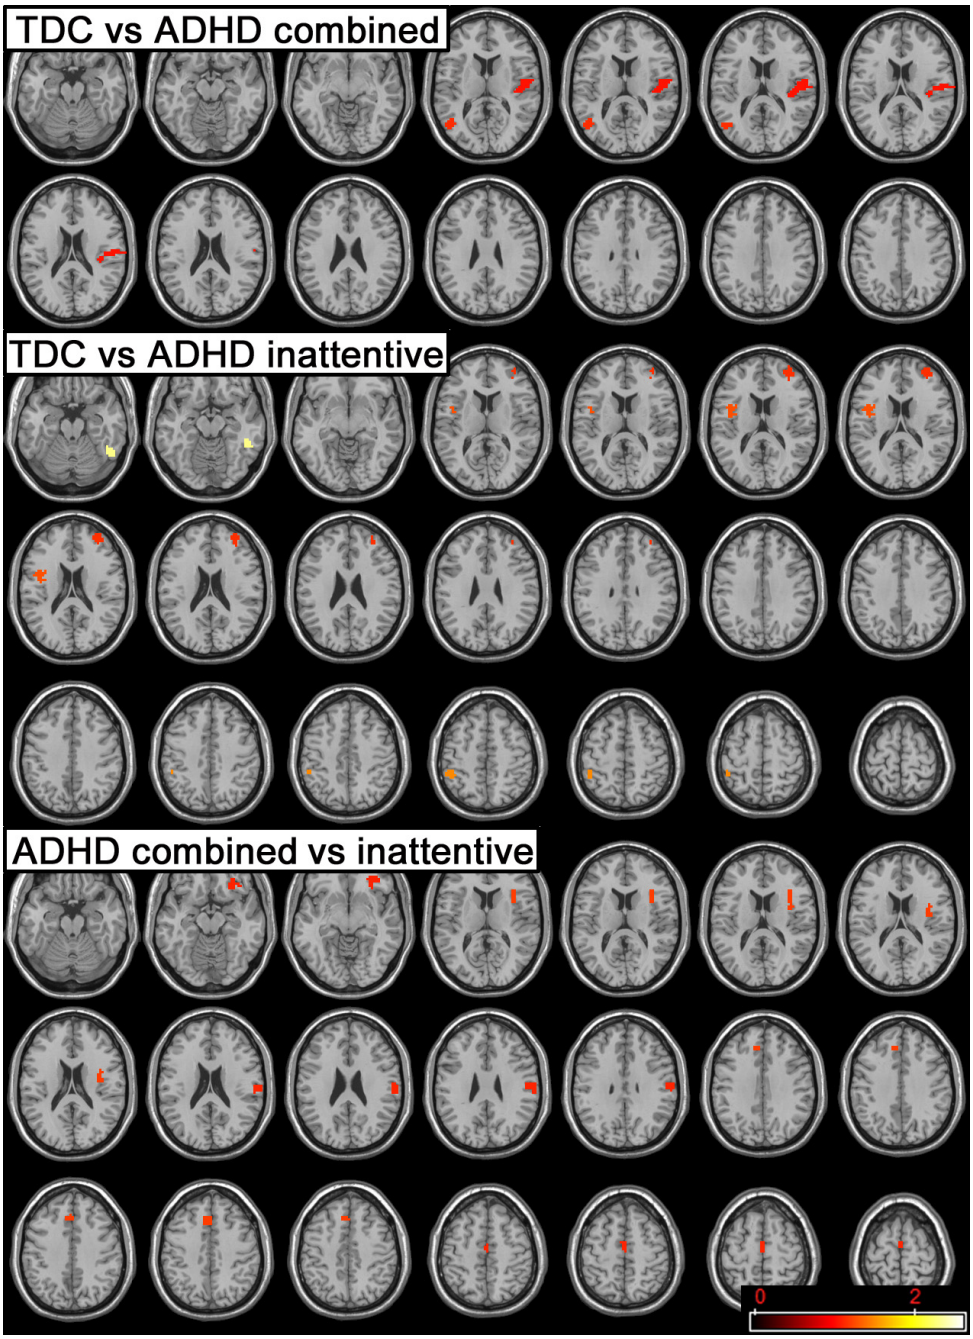

Figure 10: Pairwise comparison among three groups. Top: TDC vs ADHD combined subtype, Middle: TDC vs ADHD inattentive subtype, Bottom: ADHD combined vs inattentive subtypes. (scale$=-\log_{10}(p\text{-value})$)

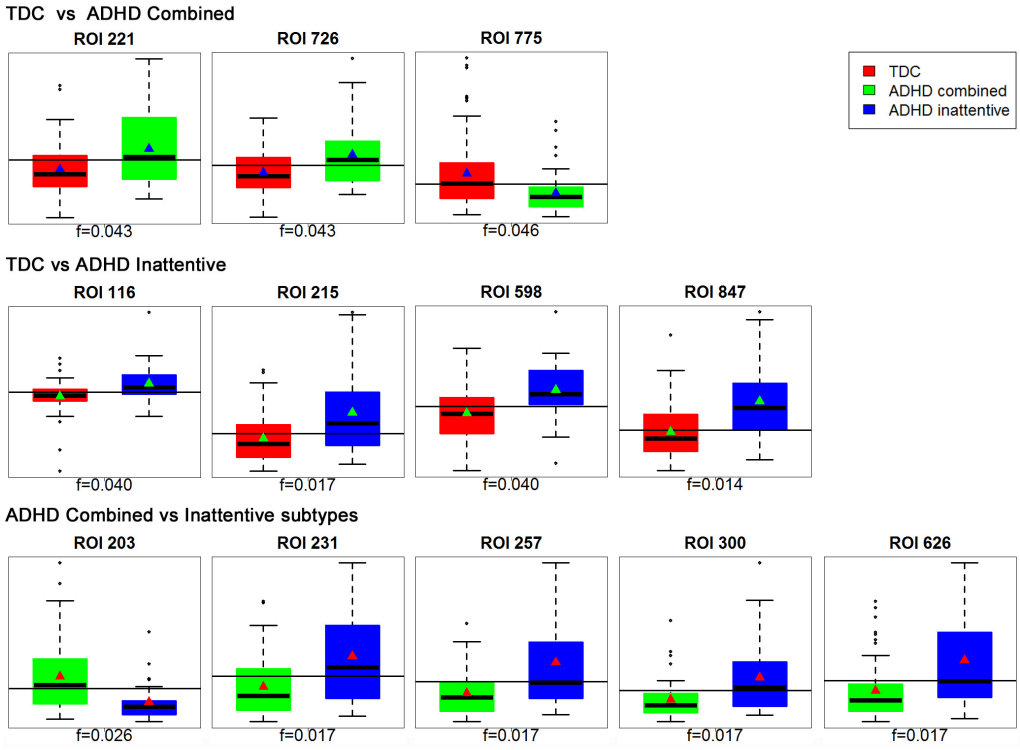

Figure 11: Boxplots of $\widetilde{\mathbf{u}}_i \widetilde{\mathbf{m}}_{ij}^{g,s}$ when the jth ROI shows significant group differences at the ith frequency component. We present the ith frequency component on the x-axis of each plot. The horizontal line is drawn at zero. Small triangles are added to show the (i, j)th component of $\widetilde{\mathbf{U}}\widetilde{\mathbf{M}}_g$ for each group g. (Red=TDC, Green=ADHD combined, Blue=ADHD inattentive)