



JIVE integration of imaging and behavioral data

Qunqun Yu^{a,b,*}, Benjamin B. Risk^b, Kai Zhang^{a,b}, J.S. Marron^a

^a Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, NC, USA

^b The Statistical and Applied Mathematical Sciences Institute, NC, USA

ARTICLE INFO

Keywords:

Human connectome project
Task-fMRI
Multivariate analysis
Partial least squares
Canonical correlation analysis

ABSTRACT

A major goal in neuroscience is to understand the neural pathways underlying human behavior. We introduce the recently developed Joint and Individual Variation Explained (JIVE) method to the neuroscience community to simultaneously analyze imaging and behavioral data from the Human Connectome Project. Motivated by recent computational and theoretical improvements in the JIVE approach, we simultaneously explore the joint and individual variation between and within imaging and behavioral data. In particular, we demonstrate that JIVE is an effective and efficient approach for integrating task fMRI and behavioral variables using three examples: one example where task variation is strong, one where task variation is weak and a reference case where the behavior is not directly related to the image. These examples are provided to visualize the different levels of signal found in the joint variation including working memory regions in the image data and accuracy and response time from the in-task behavioral variables. Joint analysis provides insights not available from conventional single block decomposition methods such as Singular Value Decomposition. Additionally, the joint variation estimated by JIVE appears to more clearly identify the working memory regions than Partial Least Squares (PLS), while Canonical Correlation Analysis (CCA) gives grossly overfit results. The individual variation in JIVE captures the behavior unrelated signals such as a background activation that is spatially homogeneous and activation in the default mode network. The information revealed by this individual variation is not examined in traditional methods such as CCA and PLS. We suggest that JIVE can be used as an alternative to PLS and CCA to improve estimation of the signal common to two or more datasets and reveal novel insights into the signal unique to each dataset.

1. Introduction

There has been an unprecedented increase in the volume of data in neuroscience, motivating the need for new large scale techniques that can simultaneously analyze datasets from different modalities. This article focuses on developing methods to integrate neuroimaging and behavioral data from the Human Connectome Project (HCP), whose primary goal is to characterize the neural pathways that underlie brain function and behavior in healthy young adults (van Essen et al., 2013). To elucidate the relationship between brain function and behavior, we simultaneously analyze both task functional magnetic resonance imaging (fMRI) and behavioral variables. Some of these behavioral variables are measured at the time of task performance. Others are from the NIH toolbox, Penn Computerized Neurocognitive Battery and other tests that characterize a range of motor, sensory, cognitive and emotional processes.

Two traditional approaches usually applied for exploring the relationships between two multivariate sets of variables are Partial

Least Squares (PLS, Wold et al., 1984) and Canonical Correlation Analysis (CCA, Hotelling, 1936). These approaches focus on pairs of linear combinations of features that have maximum covariance (PLS) or correlation (CCA). PLS was first introduced to the neuroimaging community by McIntosh et al. (1996) and a comprehensive review of the applications of PLS in neuroimaging is provided by Krishnan et al. (2011). One major application of PLS lies in relating brain activity to behavior, which successfully discovered brain features of patients with disorders such as schizophrenia and Alzheimer's disease (Nestor, 2002; Menzies et al., 2007; Tippett and Black, 2008). However, one major problem with PLS is that the individual variation is not considered. Furthermore, the variance in one of the datasets can dominate the decomposition and lead to poor recovery of the shared signal (Section 3.5.2 in Hastie et al. (2005)). This problem may be alleviated by scaling, but finding an appropriate scaling can be challenging when the numbers of features differ by several orders of magnitude (e.g., imaging and behavior). Another method used to study the relationship of imaging and behavior is CCA. For example, Smith et al. (2015) relate

* Correspondence to: UNC Department of Statistics and Operations Research, Chapel Hill, NC 27599-3260, USA.
E-mail address: qunyu@live.unc.edu (Q. Yu).

<http://dx.doi.org/10.1016/j.neuroimage.2017.02.072>

Received 8 November 2016; Received in revised form 23 February 2017; Accepted 24 February 2017

Available online 27 February 2017

1053-8119/ © 2017 Published by Elsevier Inc.

functional connectomes to behavior using CCA. They nicely link a specific pattern of the functional connectomes to behavioral and demographic measures, where a single CCA mode is a substantial part of the total variation. Their analysis uses an ICA-based parcellation to reduce the number of spatial locations (91,282) to 200 networks. It then requires SVD to reduce the number of features to be less than the number of subjects. This is because CCA tends to find spurious directions driven by sampling artifacts in high dimensional, low sample size (HDLSS) settings such as imaging data (Pezeshki et al., 2004).

Another challenge for studying activation patterns comes from the fact that different networks may be simultaneously activated by a task and consequently the effect of a network of interest may be confounded. For example, in a working memory designed task (Barch et al., 2013), participants are asked to look at pictures and remember them sequentially. This not only activates the working memory related system but also the vision system. The vision effect may be even stronger than that of the working memory system, which makes it hard to discover the working memory effect in the brain. Contrasts between tasks are typically used to isolate a task of interest, but this leads to a loss of information. For example, large activation contrasted with no activation can appear equivalent to small activation contrasted with negative activation. The HCP has made a large amount of imaging and behavioral data available, and jointly analyzing the data sources may improve our ability to distinguish functional networks. In particular, those behavioral measures containing the task-related variables are expected to be highly related with appropriate parts of the image signals. Then, integrating these two datasets can be helpful to deeply investigate the task-related effect.

The main purpose of this paper is an improved method for integrating imaging and behavioral data to study activation patterns. Motivated by recent results on integrating different datasets, we use the Joint and Individual Variation Explained (JIVE) method (Lock et al., 2013; Feng et al., 2015). JIVE is an approach to explore two different types of variation: *joint* variation across different data blocks (e.g. the variation appearing both in imaging and behavior components, such as the working memory related variation in Sections 4.2 and 4.3) and *individual* variation that is unique to each data block (e.g. image variation not related to behavioral variables, such as the motor variation in Section 4.4). The term individual here refers to data blocks, not to subjects nor participants. The individual subspace is conceptually different from the residuals of the joint subspace, as it represents the de-noised signal unique to a data block. By simultaneously extracting information common to multiple datasets (joint variation) and information unique to each dataset (individual variation), JIVE can be used to more effectively identify the relationship between behavioral and image data.

There are two generations of JIVE: the first one by Lock et al. (2013) and the second by Feng et al. (2015). Three major improvements in the second generation are better theoretical properties, a non-iterative approach that facilitates a faster algorithm, and automatic handling of wildly different numbers of features across the datasets. Thus, we propose to apply the second generation to integrate imaging and behavioral data. Compared to PLS, JIVE not only finds the shared information but also clearly indicates individual sources of variation, shown in the toy example in Section 2. Additionally, JIVE works on multiple data blocks (e.g., different modalities of neuroimaging). This article demonstrates the usefulness of JIVE by focusing on two types of data, imaging and behavioral data, in three examples. In the first two examples, the imaging data is the task fMRI from a working memory designed task, but with different levels of working memory related variation. In the third example, the imaging data is fMRI driven by a motor task. The behavioral data in all examples is the same and contains some working memory task-related variables, but none of the motor task-related variables. Thus, the imaging and behavioral data are very strongly, highly, and weakly related in these three examples and the results show the corresponding different levels of working memory

related signals in the joint component, indicating JIVE is a powerful multivariate statistical method for jointly analyzing imaging and behavior data.

2. Methods

In this section, we provide an introduction to JIVE and use a toy example to illustrate its usefulness to separate joint and individual signals. Suppose $X (d_1 \times n)$ and $Y (d_2 \times n)$ are two different data sources for a set of subjects, where each column is a common subject and each row is a variable. For MRI, d_1 is the number of voxels and d_2 is the number of behavioral variables. In addition, both X and Y consist of two parts: an underlying low rank signal (A_X and A_Y) and noise (E_X and E_Y). In our data application, the rank of the signal was estimated from the scree plot (Web Supplement Section S3). We also examined an automatic approach to rank selection (Dray, 2008) in the same supplementary section, but it resulted in unsatisfactory results due to gross overfitting. Each low rank signal A_X and A_Y can be decomposed into joint and individual variation, denoted as J_X and I_X , J_Y and I_Y respectively:

$$\begin{aligned} X &= A_X + E_X = J_X + I_X + E_X \\ Y &= A_Y + E_Y = J_Y + I_Y + E_Y \end{aligned} \quad (1)$$

JIVE aims to find patterns across subjects that are common for A_X and A_Y and patterns across subjects that are unique to each low rank signal. Those patterns are called *subject signatures* and they are singular eigenvectors of Singular Value Decompositions (SVDs) living in the row space, \mathbb{R}^n . This row space representation circumvents the challenge of choosing an appropriate normalization across data blocks. In particular, the joint matrices J_X and J_Y are assumed to have the same row space,

$$\text{row}(J_X) = \text{row}(J_Y) = \text{row}(J), \quad (2)$$

while the intersection of the individual row spaces is the zero vector space, i.e.,

$$\text{row}(I_X) \cap \text{row}(I_Y) = \{\vec{0}\}. \quad (3)$$

Note that $\text{row}(I_X)$ and $\text{row}(I_Y)$ are not necessarily orthogonal. Furthermore, the orthogonality of the joint and individual variation is imposed, i.e.,

$$\text{row}(J_X) \perp \text{row}(I_X), \quad \text{row}(J_Y) \perp \text{row}(I_Y). \quad (4)$$

It is shown in Feng et al. (2015) that the decomposition $A_X = J_X + I_X$, $A_Y = J_Y + I_Y$ subject to the constraints (2), (3) and (4) is unique.

In practice, the noise matrices, E_X, E_Y , will affect the estimators of the signal matrices, A_X, A_Y , respectively, which may result in an improper decomposition of the joint and individual variation. The notion of an angle between 2 planes in 3-D space can be defined in terms of *principal angles* between hyperplanes in the high dimensional case. Let \tilde{A}_X and \tilde{A}_Y be estimates of A_X and A_Y . If some subspaces of $\text{row}(\tilde{A}_X)$ and $\text{row}(\tilde{A}_Y)$ have a small angle, they are considered as potential estimates of the joint row space. Principal Angle Analysis (PAA) and the generalized $\sin\theta$ theorem provide a bound for the angle between two subspaces so that they can be viewed as a single joint space. Practically, this bound cannot be directly used because the error matrices E_X and E_Y are not observable. Feng et al. (2015) proposed a bootstrap based method to provide estimates for this bound, which is then used as the threshold for finding the joint subspace. In particular, this subspace is found using a Principal Angle Analysis, which is an appropriate singular value decomposition of the combined row spaces of \tilde{A}_X and \tilde{A}_Y . JIVE is easily extended to multiple data blocks by using this SVD on the concatenation of multiple right singular vector matrices.

Fig. 1 shows the analytical steps involved in JIVE. The blue and orange matrices in the first column are the two data blocks of the

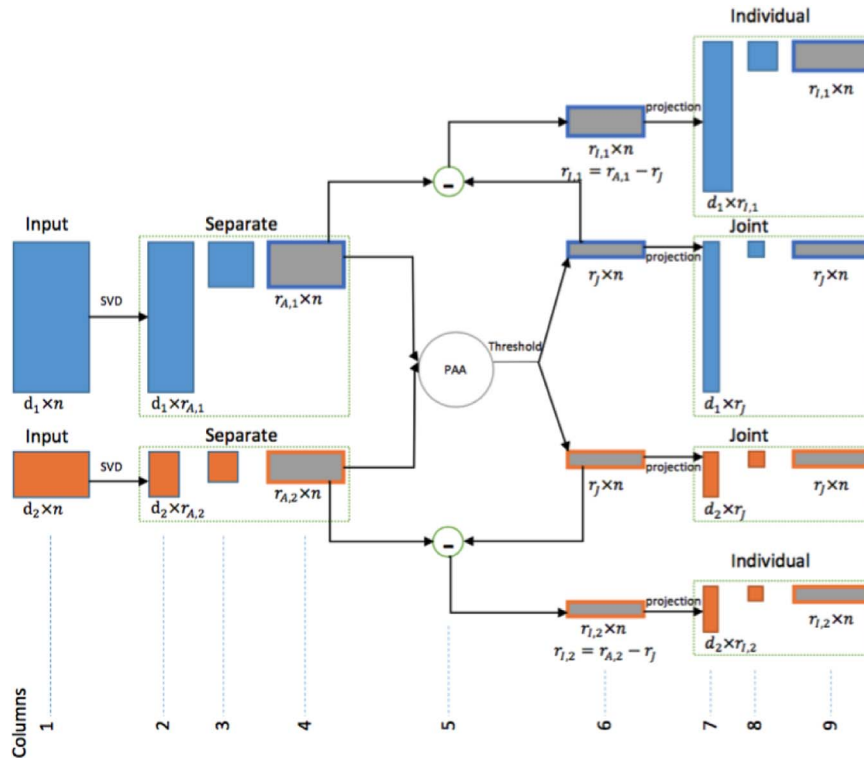


Fig. 1. Analytical steps of the JIVE algorithm. Raw data on the left are input to separate SVD low rank approximations (green boxes). Next steps use the separate SVD scores (all score matrices are shown as gray boxes), first thresholding the Principal Angle Analysis to obtain joint scores, then performing basis subtraction to obtain individual scores. Finally, projection gives the JIVE loadings.

original data. JIVE starts with their low rank approximations, \tilde{A}_X and \tilde{A}_Y , which is done through *separate* SVDs. All SVDs in this diagram are shown in green dotted boxes (columns 2, 3, 4, 7, 8 and 9), with resulting separate loadings (left boxes shown as columns 2 and 7, same as the PCA eigenvectors), singular values (center boxes shown as columns 3 and 8) and scores (right boxes shown as columns 4 and 9). The scores, i.e., the right singular matrices, are shown here as gray filled boxes in columns 4, 6 and 9. PAA is applied to the SVD scores to determine the row space (i.e., the scores space) for the joint component (the two middle gray boxes in the sixth column). Performing a basis-wise subtraction of this space from the row space of A_X and A_Y results in the row spaces for individual signals (first and fourth gray rectangles in the sixth column). The joint and individual signals can be derived by projecting \tilde{A}_X and \tilde{A}_Y to the corresponding row spaces.

A toy example is provided in Fig. 2. The matrices X and Y are 100×100 and 10000×100 simulated datasets, with each column representing a subject and each row a feature. Fig. 2 shows several heat map views, with each entry colored according to the color bar at the bottom of each matrix, which uses red as +, blue as -, and white for 0. The joint signals are rank one matrices generated by a row vector which contrasts the first half of the subjects with the second half. This is the only joint subject signature (joint signature for short). The individual component for X is also rank one with a signature vector orthogonal to the joint signature and similarly for the Y individual component. The only individual signature for X indicates three groups of data objects, while the only individual signature for Y indicates 4 groups. These two individual signatures are not orthogonal to each other. The noise matrices are random Gaussian with the same standard deviation for each entry. The left column (X and Y) is the summation of the right three columns.

Fig. 3 shows the JIVE approximation for the observed X and Y in the toy example as well as the JIVE estimation of the two types of signals. JIVE captures both the joint and individual signals and thus provides a good estimate of the true signals of X and Y in the presence of noise.

Figure S1 of the Web Supplement shows that PLS captures the rank-two overall structure, but poorly separates the joint and individual signals. Feng et al. (2015) provide another example where CCA and PLS fail to recover the true signals where there is a large difference in the number of features in the individual datasets. Furthermore, they also show the old version of JIVE (Lock et al., 2013) does not separate joint and individual signals appropriately in their example.

3. Data

In this section, we introduce the imaging and behavioral data from the HCP 500 subjects release, where the subjects are healthy adults aged from 22 to 36 years old. We utilize the subject-specific z-statistic maps provided in CIFTI (cortical surface plus subcortical gray matter, Glasser et al. (2013)) format of the HCP data from two different task categories: a working memory/category specific representation task and a motor task. A detailed description of the statistical model used to estimate the z-statistics is provided in Barch et al. (2013) and additional information about the data and CIFTI format is described in Glasser et al. (2013). The first task is a combination of 2 working memory tasks (0-back and 2-back) based on 4 different categories of pictures (body parts, faces, places and tools). There are 8 task blocks in total and in each task block, the participants were given pictures sequentially and asked to indicate if the current picture was the same as a target picture (0-back task blocks) or the same as the one that appeared 2 back (2-back task blocks). These are referred to here as the *working memory category tasks* for convenience. In particular, we use the z-statistic images (also called statistical parametric maps) of contrasting 2-back and 0-back task blocks, referred as *Case 1: 2-0 back contrast images*. We also use the z-statistics from the 2-back tool task, which are not contrasts. This is referred to as *Case 2: 2-back tool images*. In additional sessions, the subjects completed motor tasks, where participants were asked to either tap their left or right fingers, squeeze their left or right toes, or move their tongue. Here, we study the

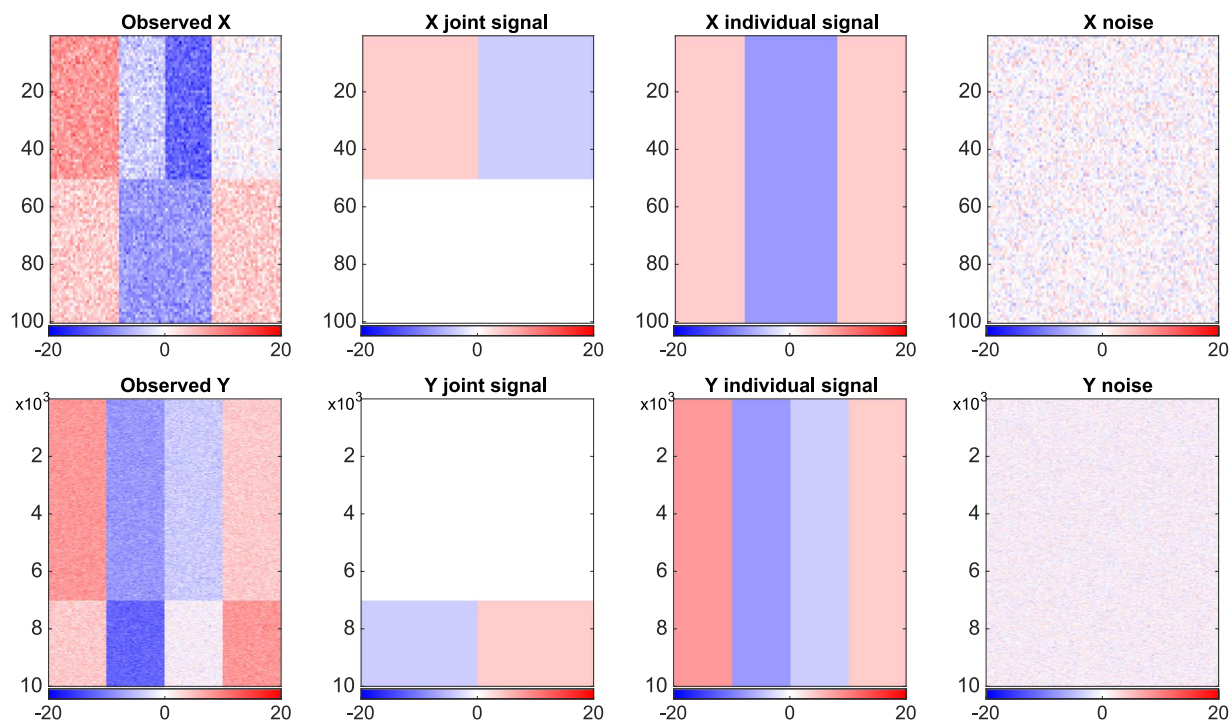


Fig. 2. Data construction of the toy example. The first column presents the observed data matrices X and Y , which are sums of the joint and individual signals and the noise, shown in the remaining columns. This is a heat map view with entries colored according to the color bar at the bottom of each panel. These structures are difficult to capture using conventional methods due to different orders of magnitude of the numbers of features.

z -statistic images contrasting the right finger-tapping task and the average of all motor tasks, abbreviated as *Case 3: motor right-hand*, in our JIVE analysis. In total, we use three different types of imaging data: 2-0 back contrast images, 2-back tool images and the motor right-hand images. These cases are shown here because they result in three different interesting JIVE analyses. Image data are vectorized for each

subject and then the subjects' data are concatenated to form a data matrix, shown in the left part of Fig. 4.

HCP also collected a set of behavioral measurements that are helpful for understanding the relationship between brain activity and human function. Fig. 5 shows the behavioral variables for one participant. Their full description is provided in the Web Supplement

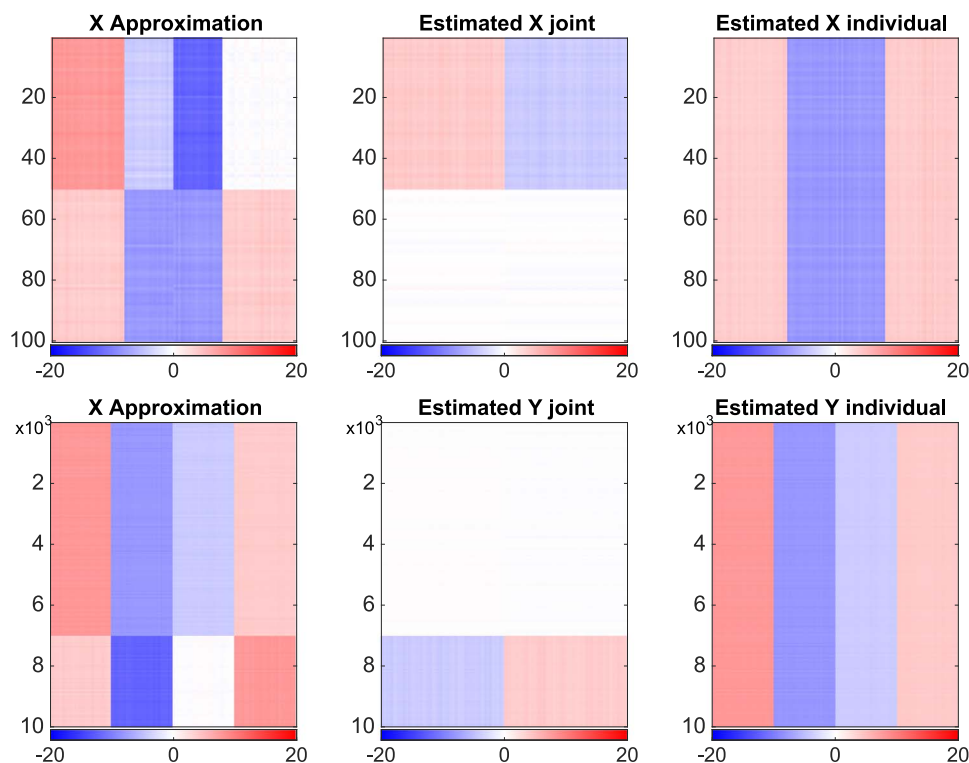


Fig. 3. JIVE estimation of the toy example. The first column shows the rank-2 approximations of X and Y , the second for the estimated joint signals and the third for the individual signals. The color scheme is the same as that in Fig. 2. JIVE captures these diverse types of variation for each data block.

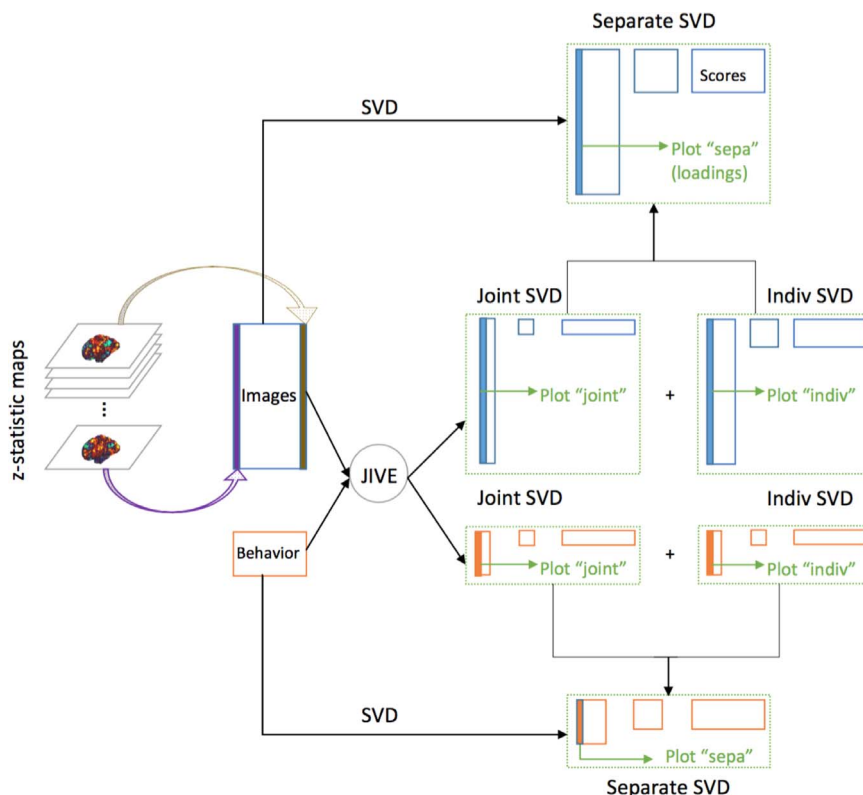


Fig. 4. JIVE analysis diagram of the HCP data, showing sources of loading vectors used in the plots.

Section S8. There are 139 behavioral variables that can be categorized into eight domains: cognition in the NIH toolbox (colored red; shortened as cognition), emotion in the NIH toolbox (orange; emotion), motor in NIH toolbox (brown; motor), personality (green), delay discounting (cyan; ddis), working memory accuracy (blue; wm_acc), working memory reaction time (magenta; wm_rt), and others (e.g., spatial orientation and fluid intelligence; black). These eight domains are straightforward to understand except perhaps delay discounting (ddis). As described in Shamosh and Gray (2008) and Odum (2011), delay discounting is a tendency to undervalue larger later rewards relative to smaller earlier rewards and is used to measure self-regulation/impulsivity. In the HCP study, a higher score indicates

greater self-control. The performance variables in the domains wm_acc and wm_rt are thought to be highly related to the working memory category task, which suggests JIVE is a promising method to study the brain areas which are working memory and task-related. Additionally, the motor right-hand images are expected to be only weakly related to the behavioral data and thus serve as a control case.

Some necessary data preprocessing steps handling missing data and the dramatically differing scales of behavioral variables are carried out before JIVE analysis. There were only 2 subjects with substantial missing data who were excluded. The few remaining missings were handled by imputations as discussed in section S2 in the Web Supplementary material. Furthermore, the scales of the behavioral

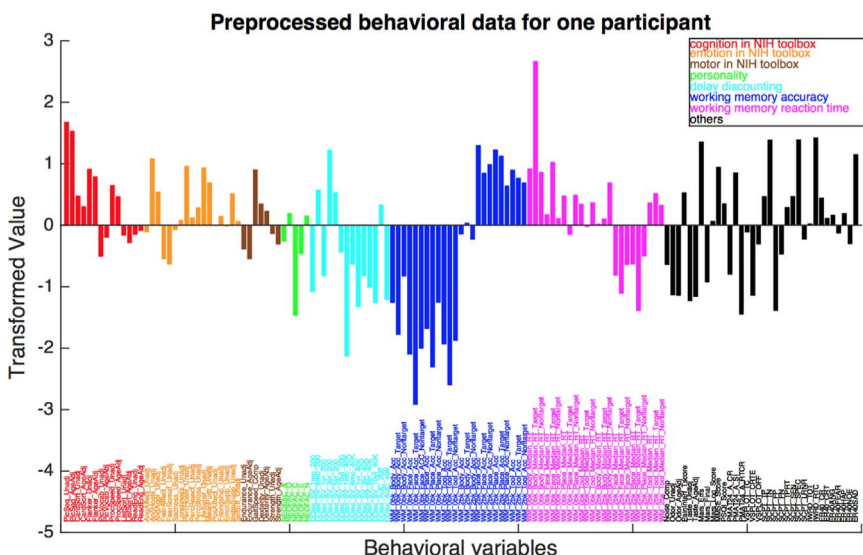


Fig. 5. Preprocessed behavioral data for one participant. The names of the behavioral variables are displayed as the vertical text. They are grouped into eight domains and colored according to the legend. Detailed information is provided in the Web Supplement Section S8.

variables vary by several orders of magnitude and some variables are very skewed, suggesting a strong need for transformation and normalization. Behavioral data are transformed using the shifted logarithm transformation, winsorisation and standardization developed by Feng et al. (2016), which improves the closeness of each behavioral variable distribution to normality. More specifically, the transformation is shown as Eq. (2) in the Web Supplement Section S2 and further details are also provided there. After data cleaning, the dimensions for the imaging and behavioral data matrices are $91,282 \times 487$ and 139×487 , respectively.

4. Results

This section applies JIVE to the preprocessed behavioral data and activity images described in Section 3 in the three cases of: 1) 2-0 back contrast images; 2) 2-back tool images; and 3) motor right-hand images. These three cases were carefully chosen in terms of how much related signal is expected to be found in the joint component. In the first case, the contrast images are designed to selectively eliminate the activations that are common to both 0-back task blocks and 2-back task blocks, e.g., the activation in the visual network. Hence, the signals associated with working memory in this case should be the strongest among the three. In the second case, the effect of the visual network is expected to be an important driver of the image signals and may be a distraction from discovering the effect of the working memory system. The 2-back tool task block is chosen over the other 7 task blocks because it was seen in Barch et al. (2013) to have less category specific activation. In the third case, none of the behavioral variables were collected during the motor task and thus the behavioral data are thought to be at most weakly related to the activity images. Hence, no or weak signals are expected in the joint component.

An overview of the JIVE analysis of the HCP data appears in Fig. 4. The analysis for the $91,282 \times 487$ image data and 139×487 behavioral datasets took approximately 2.6 minutes in Matlab with a 2.7 GHz processor. The initial vectorization of the z-maps is shown on the left. Conventional separate SVD analyses are shown at the top and bottom. Recall that the detailed JIVE steps are depicted in Fig. 1. The low rank approximation in the first step of the JIVE analysis (column 4 in Fig. 1) is based on an inspection of the scree plots (Web Supplement Section S3).

After the application of JIVE in the three cases, we study the SVD view of the joint and individual signals. The visualization of scores, which is useful for understanding relationships between participants, are not shown because they are approximately Gaussian (Figure S6 in the Web Supplement) and thus not very informative. This may be due to the fact that only healthy young adults are included in the HCP data, and thus the subjects are relatively homogeneous. Score plots tend to be very useful when there are important subgroups of participants. More insightful for these data are loadings plots, which are the entries of the left singular vectors (denoted as SVD1, SVD2, etc.) indicated with green arrows in Fig. 4. Then the image loadings are displayed on the cortical surface (e.g., Fig. 8) and in volume space for the subcortical structures (e.g., Web Supplement Figure S9). The behavioral data loadings are plotted as colored bar plots as in Fig. 6. The SVD loadings are especially useful for finding the corresponding brain areas responsible for the joint and individual signals as well as indicating involved behavioral variables. For comparison, SVD is first applied to separate behavioral and image data to establish a baseline. Here, we focus on interpreting the results for the cortical surface. We provide a figure for the subcortical results for Case 1 in the Web Supplement Section S4. The terminologies *separate SVD*, *joint SVD* and *individual SVD* are used to label these three analyses. In particular, these SVD results are the loadings shown in Fig. 4. The three cases all share the same behavior data, and consequently the separate SVD of the behavior is identical for all cases (Section 4.1). The following three subsections show the main results for the three cases, respectively. The significance

assessment of our JIVE results is shown in the last subsection.

4.1. Separate SVD on behavioral data

Bar plots for the behavioral variables, which show the entries, i.e., loadings, of the SVD1, SVD2 and SVD3 direction vectors are shown in Figs. 6 and 7. Variables are colored as in Fig. 5. We observe the following: 1) the SVD1 direction captures the variation of good or bad overall in-task working memory performance; 2) the SVD2 reveals a mode of variation which seems to be related to thoughtfulness and deliberation; 3) the SVD3 characterizes positive versus negative emotion and personality. Each of these points is discussed in detail as follows.

(1) *SVD1 direction reveals the variation associated with overall in-task working memory performance.* The SVD1 loadings in Fig. 6(a) show that the blue working memory accuracy (wm_acc) variables and the magenta working memory reaction time (wm_rt) variables are important drivers with opposite signs. (Note that the sign of each loading vector is not identifiable, so the pertinent information is whether variables in different domains have the same or opposite signs.) This makes intuitive sense because shorter reaction time is related to better in-task performance, and better performance results in higher accuracy which explains a large part of the variation in the sample of subjects. Furthermore, the variables in cognition (red) listed in the Web Supplement S8 measure picture sequence memory, attention and inhibitory control, speed of processing, etc. The higher cognitive variables are similarly related to better in-task performance. Greater delay discounting (cyan) values imply an individual discounts future gains less, which indicates greater self-control. Hence, higher values are related with better in-task performance. Thus, the SVD1 direction corresponds well with good or bad working memory task performance.

(2) *SVD2 reveals a mode of variation which seems to be related to thoughtfulness and deliberation.* Fig. 6(b) shows the SVD2 loadings. The variables in the domains of delay discounting (cyan) and working memory reaction time (magenta) are the major contributors to this mode of variation and work in the same direction. Interestingly, SVD2 reveals a mode of variation in which greater self-control is associated with slower reaction time, which may be related to more thoughtful and deliberate subjects.

(3) *SVD3 characterizes positive versus negative emotion and personality.* Fig. 7(a) shows the SVD3 loadings. In this mode, relatively fewer variables stand out than in SVD1 and SVD2. In particular, the variables in the emotion (orange) and personality (green) domains have the largest loadings for SVD3. Fig. 7(b) studies the variables in those two domains in more detail. The variables highlighted by the red rectangle are the negative emotion variables measuring anger, hostility, fear, sadness, etc. These variables work in one direction. Others highlighted by the cyan rectangle measure positive emotion variables, e.g., life satisfaction, positive affect and friendship, which as expected work in the other direction. A similar pattern shows up for the variables in the personality domain where four of the five personality variables show substantial variation. Among those four, the variables measuring agreeableness, conscientiousness and extraversion point in the opposite direction of the variable measuring neuroticism. SVD3 is thus about variation in terms of positive versus negative emotion and personality.

4.2. Case 1: JIVE on behavioral and working memory 2-0 back image data

In this section, we compare the separate SVDs with the full JIVE analysis in Case 1 using the first two SVD loadings, shown in Fig. 8. We observe that: 1) separate SVD (the first four rows in the first column) splits the working memory network signal between the first two components; 2) the JIVE joint SVD1 component (the second row,




Fig. 6. SVD1/SVD2 loadings of the behavioral data. (a): SVD1, the variables in the domain of cognition in the NIH toolbox (red), delay discounting (cyan; ddis), working memory accuracy (blue; *wm_acc*) and working memory reaction time (magenta; *wm_rt*) showing that the SVD1 direction corresponds well with overall in-task memory. (b): SVD2, the variables in the domains of *ddis* (cyan) and *wm_rt* (magenta) work together, suggesting variation in thoughtfulness and deliberation. Same color scheme as in Fig. 5 is used.

second column) uses behavioral information to capture the full activation in the working memory related regions in a single component and the JIVE individual component captures the task unrelated overall activation; 3) PLS1 (the fifth row, second column) gives a similar result, one in which there is less contrast between regions than the JIVE joint SVD1 component (see Web Supplement S12 for a larger version of PLS1 versus JIVE joint SVD1).

(1) *The separate SVD splits the signal.* The separate SVD1 and SVD2 loadings of the image data displayed on the cortical surface are shown as the second and fourth rows in the first column of Fig. 8, respectively. Both panels are colored according to the color bar on their right. In all such displays in this section, the range of the corresponding loadings is used as the range of the color and 0 is colored as black. The overall red pattern in the SVD1 loadings indicates this mode of variation is driven by different levels of overall activation among the subjects. This SVD1 pattern is similar to that for the 2-back tool and motor right-hand tasks studied in later sections, and thus does not seem to be task specific. SVD1 also shows some hot spots indicated by the yellow circles. In the SVD2 loadings, these same red spots contrast with adjacent regions. To examine whether these regions are working memory related, we compare the separate SVD2 loadings with the *group level analysis* of the 500-subject release of the HCP data (lower right panel), which is an updated version of the analysis of Barch et al. (2013). Independent studies (Satterthwaite et al., 2013) have found similar working memory related areas in the brain. Relative to this group level result, the separate SVD2 discovers some of the working memory related signals. However, they are not as strong as those in the group level result because this important signal is split between SVD1 and SVD2 as shown using the circles in the first column.

(2) *JIVE focuses the working memory related signal into one component.* The JIVE joint SVD1 loadings of the imaging data (the second row, second column) show the full extent of the working memory related regions. This is because the JIVE joint signals in the image data share a common row space with those in the behavioral

data, thus concentrating that signal. A major part of this common information is working memory 2-0 back task-related. Thus, the joint SVD in the joint image block captures much more of the working memory associated signals than does separate SVD, and focuses on this instead of splitting it between SVD1 and SVD2. Additionally, the joint variation in the subcortical gray matter regions is also similar to the patterns observed in the group analysis (Web Supplement Figure S9). For the behavioral SVD1 in the top row, the joint and separate are very similar. This shows that much of the natural behavioral variation is associated with this task. More specifically, the working memory task-related behavioral variables (blue and magenta, first row and second column) are important drivers for the joint signals and appear in the joint, in which the *wm_acc* (blue) and *wm_rt* (magenta) variables have opposite signs. In the second row, the overall red pattern in the separate image SVD1 is captured by the JIVE individual component, showing that the variation of overall brain activity across participants is not associated with these behavior variables. On the other hand, the working memory related signal (highlighted with circles) appears in the joint component. Hence, through the JIVE integration with behavioral data, the task unrelated signal goes to the individual variation and task-related signal is enhanced in the joint signals in the JIVE result. The third row shows the separate, joint and individual SVD2 loadings for the behavioral data. Again, the separate loadings are very similar to the joint loadings. This suggests that the possible link between the *ddis* (cyan) variables and the *wm_rt* (magenta) variables noticed in the separate SVDs is also associated with this task. In the fourth row, both separate and joint components appear to be somewhat associated with working memory. However, in the Euclidean space of the 91,282 image variables, the joint SVD1 and SVD2 direction vectors are orthogonal, so there is an important contrast between these, apparently in terms of working memory regions shown in the upper part of the display. This is a potential new discovery about how behavior variables relate to brain activity. Again, the individual signals are not or are only weakly working memory related. Thus, JIVE has proven to be very effective

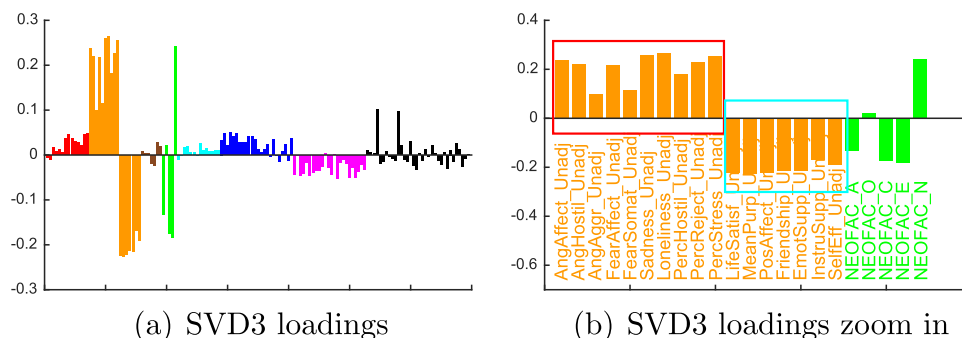


Fig. 7. (a): SVD3 loadings of the behavioral data. Variables in domains of emotion (orange) and personality (green) are the major factors for SVD3. (b): A zoom in of these variables. Note that the variables measuring positive emotion or personality work in one direction while the negative variables work in the other direction. Same color scheme as in Fig. 5 is used.

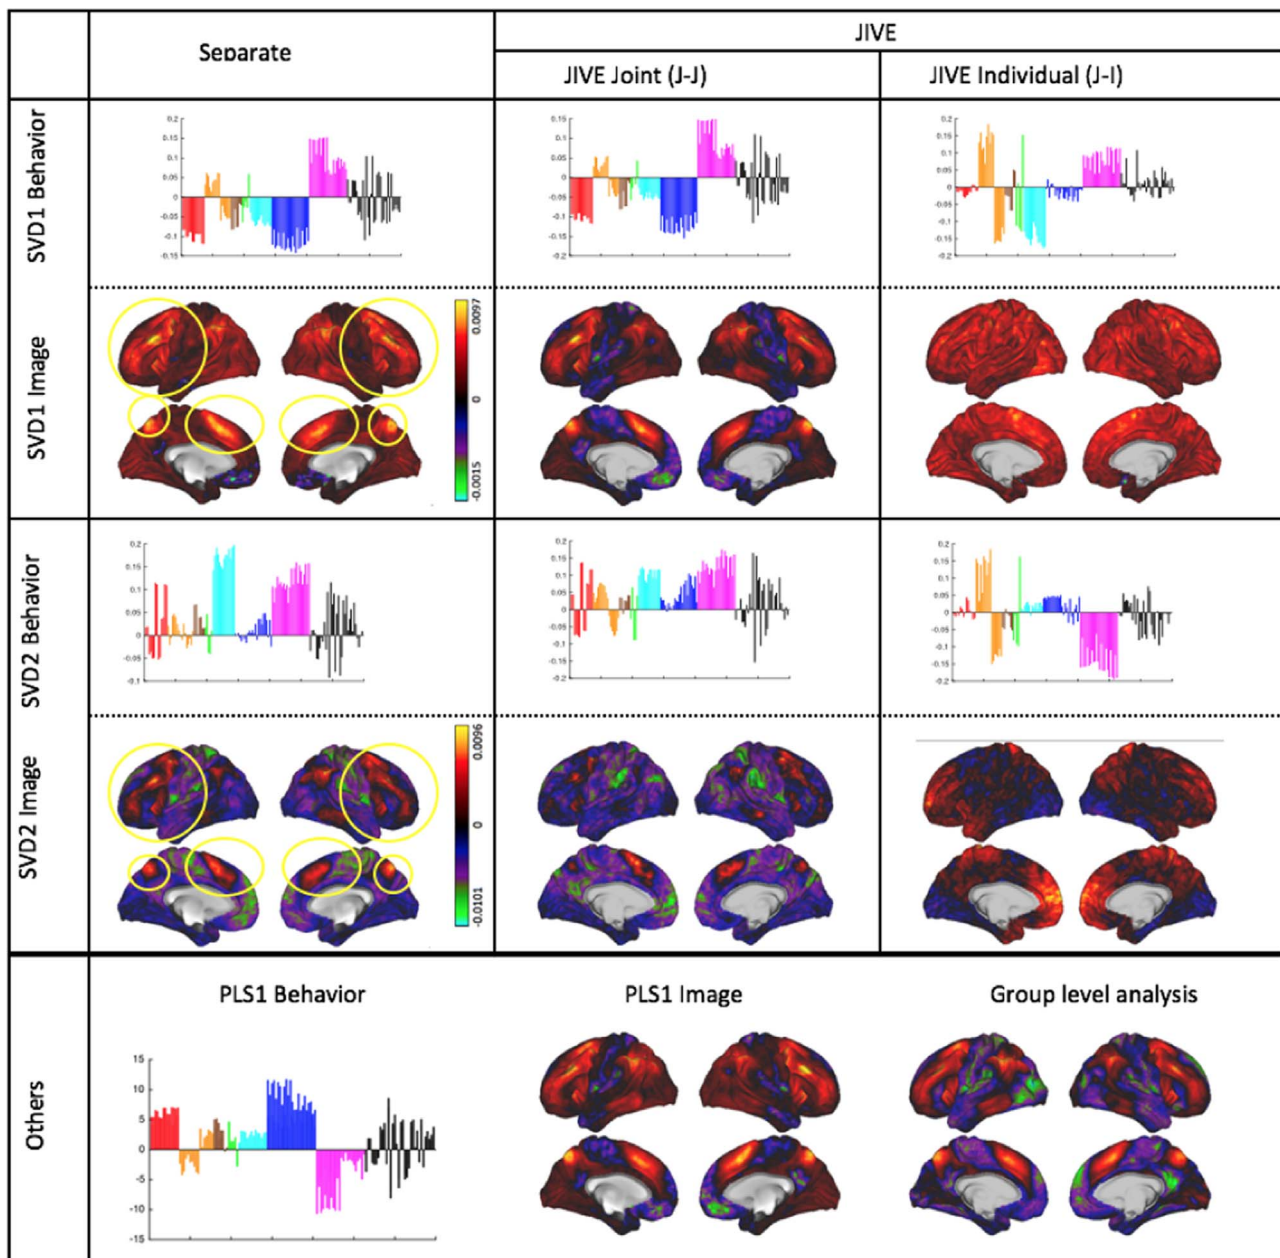


Fig. 8. Overview of Case 1 (2-0 back). Separate SVD1/SVD2 loadings on the behavior are similar to joint SVD1/SVD2 loadings. The joint image SVD1 strengthens the working memory related signals because the strong task unrelated signals go to the individual component. JIVE separates out intuitively sensible individual components and concentrates on the important activation in the joint component. The strong variation in overall activation appearing in the individual component shows that is not associated with behavior.

at finding the activation patterns responsible for working memory in the human brain.

(3) *PLS gives a result similar to the JIVE joint SVD1 component but with less contrast.* An alternative to JIVE is to use CCA or PLS to relate imaging and behavioral data. The result for CCA is in the Web Supplementary Materials Section S5. Classical CCA is not defined because of a non-invertible covariance matrix when the number of locations is greater than the number of subjects. CCA can still be computed for example using a generalized inverse approach, but this is useless for data analysis because of gross overfitting. The first PLS (PLS1) loadings of the behavior and image are shown as the first and second column in the fifth row. They are quite similar to the first JIVE joint loadings, but a careful look suggests that the color contrast between the working memory associated regions (precuneus and superior frontal gyrus) and the rest (e.g. paracentral lobule) are sharper in the joint SVD1 loadings, which are also more similar to the group

level analysis (see also Web Supplement Figure S12). This can be justified by the fact that the correlation between the JIVE joint SVD1 and the group level analysis is 0.9008, slightly higher than PLS1 whose correlation with the group level analysis is 0.8888. Hence, JIVE provides an improvement over PLS for finding the common signals in this case. We will also see in Case 2 that JIVE outperforms PLS.

4.3. Case 2: JIVE on behavioral and 2-back tools image data

In this section, we compare the separate SVD with the JIVE joint component in Case 2, shown in Fig. 9. We observe the following: 1) the separate SVD reveals some visual effects in SVD2 and working memory related variation in SVD4; 2) the JIVE joint component finds variation associated with both vision and working memory in a single component; 3) the JIVE individual components capture variation in overall activation and the default mode network; 4) PLS1 poorly highlights the

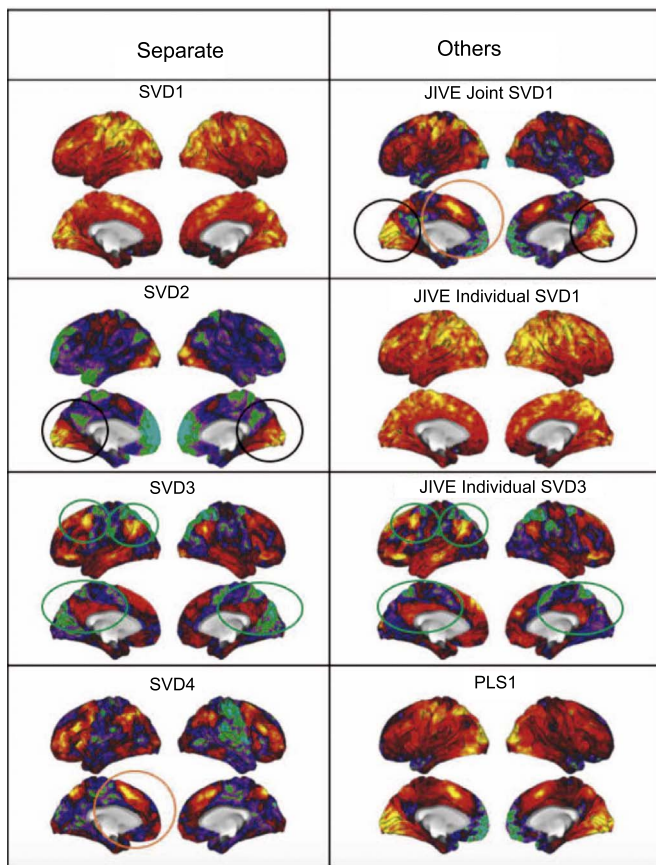


Fig. 9. Overview of Case 2 (2-back tool). The first column shows the separate SVDs. SVD1 is again the overall red pattern, SVD2 finds important variation in the visual cortex, SVD3 reveals an important variation related to the default mode network and SVD4 feels some working memory related variation. The JIVE joint SVD1 captures both the vision effect and working memory effect. The JIVE individual component feels the overall activation in SVD1 and the variation in default mode network in SVD3. The PLS1 loadings show less contrast in memory related regions versus other regions, e.g., precuneus versus paracentral lobule.

working memory network.

(1) *Separate SVD detects the visual effect and the working memory effect in SVD2 and SVD4, respectively.* The first 4 separate SVD loadings of the image data are shown in the first column of Fig. 9. As in Case 1, the overall red pattern shows up in the SVD1 loadings, which reflects differing levels of overall activation among the participants. In the SVD2 loadings, the hot spots in the black circles are the signals in the occipital lobe, which is the visual processing center of the brain as discussed in Sereno and Allman (1991), Kosslyn et al. (1995) and Howard et al. (1998). In SVD3, the red spots in the green ovals are effects in the default mode network. This network is found in at-rest individuals and is associated with mental processes such as day-dreaming and mind-wandering (Buckner et al., 2008). This mode of variation may capture the activation pattern during the intertask intervals. The SVD4 loadings appear to find some signals associated with working memory as shown in the joint SVD1 image in Fig. 8. But some activation (e.g., the hot spot in the brown circle) appears to extend to the prefrontal cortex. Hence, there is no clear boundary of the working memory related regions. Comparing Fig. 9 with the separate image in Fig. 8, the working memory associated effect is much weaker in Case 2 (SVD4) than Case 1 (SVD1 and SVD2). This confirms our expectation that the visual effect is a major effect in Case 2, which tends to dominate the effect associated with working memory. More importantly, unlike JIVE, the separate SVD gives no information about which component is working memory related because it does not integrate the imaging data with behavior.

(2) *JIVE reveals visual and working memory effects using one joint component.* The first row second column shows the JIVE joint SVD1 loadings, which detects the variation associated with working memory and gives a much cleaner boundary of the working memory area, highlighted with brown circle, than that in the separate SVD4. It also captures the signals in the visual network found in the separate SVD2 shown in the black circles. In the 2-back tool task block, vision and working memory effects are two important task-related effects and are found in the first JIVE joint component. Although the working memory related signal is weaker in this case than that in Case 1, JIVE is able to recover it in the joint component.

(3) *The JIVE individual components capture variation in overall activation and the default mode network.* As in Case 1, the variation of the overall activation in the separate SVD1 goes to the JIVE individual SVD1. Additionally, the variation in the default mode network, highlighted with green circles, is expected to be task unrelated, which is reflected by the fact that it is captured in the JIVE individual SVD3. The individual SVD2 reflects a different aspect of variation in the vision system, which is unrelated to behavior variation. It is shown in the Web Supplement Section S6 to save space. This valuable information from the JIVE individual component is not accessible through PLS.

(4) *PLS1 is a mixture of JIVE joint SVD1 and individual SVD1.* The PLS1 direction finds some activation in the working memory network. However, comparing the JIVE joint SVD1 to PLS, the first joint component in JIVE has greater contrast in the activated regions of the precuneus and superior frontal gyrus and the inactivated regions of the paracentral lobule. This may be because PLS1 does not separate the individual signals from the task related signals. For example, the right lateral view in the right cerebral hemisphere (the top right one in each panel) in PLS1 is a mixture of JIVE joint SVD1 and individual SVD1.

4.4. Case 3: JIVE on behavioral and motor image data

This case serves as a control because the behavioral data is predicted to be less related to the image data. We observe that 1) there are only weak task-related signals, i.e., signals associated with the motor system, in the joint component; and 2) these weak signals are most related to the out-of-task motor behavioral variables.

(1) *Motor related variation is largely captured by the JIVE individual component.* The first column in Fig. 10 shows the first three SVDs in the separate analysis. The SVD1 finds the overall activation across participants, the SVD2 detects the effect of the default mode network, and the SVD3 reveals some new signal highlighted in the yellow oval which is associated with the motor cortex. The first two modes of variation are not behavior related and thus appear in the individual SVD1 and SVD2 in the JIVE analysis in the third column. The signal related to the motor cortex in the separate SVD3 is also found in the JIVE individual component shown in the lower right panel. This is appropriate because the behavioral data contains very few explicit motor related variables.

(2) *The out-of-task motor variables contribute to the weak motor signals in the JIVE joint component.* Compared with the individual SVD3 loadings, the joint SVD1 and SVD2 loadings (the second column, the joint space in this case is rank-two) contain a very weak signal in the motor cortex, which is consistent with our expectation. Fig. 11 shows the joint SVD1 loadings of the behavioral data, where the brown variables highlighted in the blue rectangle are those out-of-task motor variables including dexterity, gait speed, endurance, strength, etc. (detailed information can be found in the Web Supplement Section S8). Relative to the joint SVD1 loadings in Case 1, shown in the first row and second column in Fig. 8, the brown behavioral variables in Case 3 have a much greater impact on the joint component than in Case 1. This suggests that they have a relatively stronger contribution to the rather weak motor cortex related signal found in the joint component.

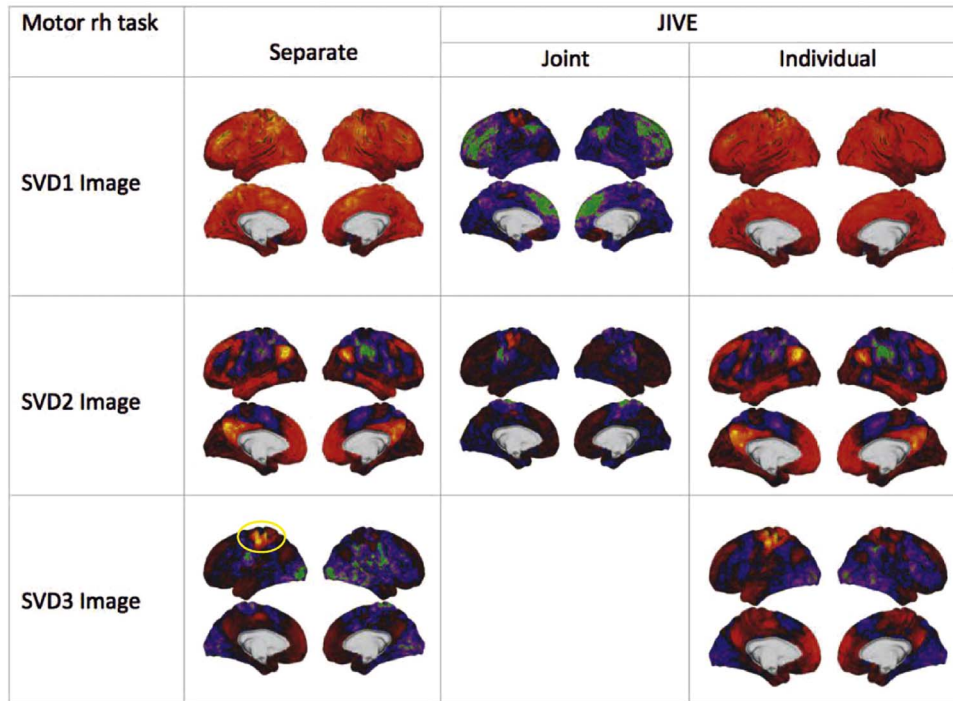


Fig. 10. Overview of Case 3 (motor right-hand). The behavior unrelated activation patterns found in the separate SVD analysis including the overall red pattern in SVD1 and the activation in the default mode network in SVD2 are also found in the JIVE individual SVD1 and SVD2 loadings. Additionally, the separate SVD3 captures the signal associated with the motor cortex shown in the highlighted region in the yellow oval, which also appears in individual SVD3 loadings. The JIVE joint component in the second column only feels a very weak signal in the motor cortex.




Fig. 11. Joint SVD1 of behavior. The variables in the motor domain (brown) contribute in a relatively stronger way to the weak signal associated with the motor cortex in the JIVE joint component.

4.5. Assessment of significance

A measure of the JIVE performance is the joint image energy (the sum of squares of the entries of the joint image matrix) shown in each panel as a vertical green line. To assess the statistical significance, we propose to use a permutation procedure. In particular, we permute the

columns in the behavioral matrices 5000 times in each of the three cases and Fig. 12 shows the results. For each panel, a black point is the joint image energy from one permutation and the black curve represents a density estimation. The p value and z scores measure the statistical significance. The empirical p value is the proportion of the black dots to the right of the green line which is a measure of statistical significance. Further comparison comes from the z scores which record how many standard deviations above the mean the green line is with respect to black dots. They show that JIVE results in Case 1 and Case 2 are significant, with stronger significance for Case 1, while the results in Case 3 are not significant. This provides a statistical validation of our joint SVD results.

5. Discussion

In this study, we demonstrate that JIVE is a powerful new method for integrating behavioral and imaging data. Through explicit separation of variation into joint and individual components, JIVE leads to new scientific insights of activation patterns related and unrelated to human behavior. In addition to the two-block version analyzed here,




Fig. 12. Permutation tests for statistical significance. The green vertical lines show the JIVE joint image energies, the black dots show the joint image energies from permutations, and the black curves show density estimations. The JIVE joint signals found in Case 1 and Case 2 are significant, while the joint signal in Case 3 is insignificant.

multi-block JIVE is available for the study of more than two data types, such as integrating genomic, behavioral, and imaging data.

We illustrated the usefulness of JIVE with three analyses. All analyses used the same behavioral dataset comprising the out-of-task HCP behavioral data as well as variables measuring the in-task memory performance. In Case 1, we analyzed the z-statistics from the 2-back versus 0-back memory contrast, and found that the joint SVD1 component in the imaging data captured the working memory network and the joint component in the behavioral data had comparatively large loadings on the in-task memory performance variables, while the individual image SVD1 component captured an overall activation that was spatially homogeneous. In Case 2, we analyzed the z-statistics from the 2-back tool task and found that the working memory network was more clearly delineated in the joint SVD1 component from the JIVE analysis than a comparable PLS analysis. Moreover, the JIVE analysis revealed that parts of the default mode network were captured by the individual SVD3 component. In Case 3, we analyzed the z-statistics from the right finger-tapping motor task. A weak signal in the right-hand motor cortex was found in the joint components SVD1 and SVD2, which coincided with larger loadings in the joint SVD1 of the motor-domain out-of-task HCP behavioral variables, while a strong signal in the right-hand motor cortex was found in the individual SVD3 component, which is consistent with the fact that no in-task motor performance variables were included in the analysis. The individual SVD1 component identified an overall activation that was relatively spatially homogeneous, and individual SVD2 clearly delineated the default mode network.

The current implementation of JIVE relies on a scree plot visual tuning to estimate the ranks of the signal subspaces A_X and A_Y . Hence, an interesting future direction is to objectively determine these initial ranks and assess their significance in the JIVE analysis.

Because the computational core of JIVE is essentially singular value decompositions, JIVE is also capable of tackling larger scale data challenges. For example, instead of summarizing the time series of MRI images by the z-statistic images, a more direct extraction of behavior related information from the time series could come from taking the full time series as one JIVE data block. In this way, the imaging matrix becomes $Td \times n$ and the loadings would also be a time series of images. Incorporating the temporal information can be useful for studying the change in brain activation patterns.

Another potential JIVE application is to use different task data to isolate a signal of interest, such as separating the working memory signal from the visual signal. In Case 1, we relied upon predefined contrasts to define the imaging data, but one could pursue a data-based approach for contrasting conditions. Future work could explore including the 2-back z-statistics as one dataset and the 0-back z-statistics as a second dataset. Then the working memory related variation is expected to be captured in a component of the *individual* variation of the 2-back task, since the working memory related regions of the brain are expected to be activated in the 2-back task but not the 0-back task.

Interpretation of the insights found by JIVE can be a tricky issue, for example driven by potential confounders. There are several approaches to dealing with this. One is to consider the confounders to be responses in a linear regression, then apply JIVE to the resulting residuals. A conceptual weakness of this, and other regression based approaches (such as defining an individual component using residuals from a regression with canonical variates/latent variables in CCA/PLS) is that regression works in an asymmetric predictive fashion, which goes in the wrong direction. A major strength of JIVE in this context is that it treats all variables in a symmetric fashion (with respect to prediction), so we recommend a multi-block approach to JIVE, adding a third block containing the confounders. The appropriateness of symmetric vs. asymmetric analyses is context dependent. When there is clear causation happening, an asymmetric regression approach is sensible. Our behavioral variables such as personality traits perhaps are

best viewed as driving the image observations. On the other hand, our task performance variables may be viewed as being caused by what is seen in the image or both may be driven by some unknown other factor. Overall, these indicate to us that a symmetric analysis is more appropriate for initial exploratory analyses as performed here. The individual component of JIVE also differs from regression residual methods through noise reduction, accomplished in the first JIVE step.

6. Conclusion

This article demonstrates JIVE integration of imaging and behavioral data from the Human Connectome Project. Compared with traditional methods such as CCA and PLS, JIVE appears to more effectively identify the signal shared between two datasets. Unlike traditional methods, JIVE also allows an examination of the signal that is unique to each dataset. Thus our analyses suggest JIVE can be used to improve our understanding of brain function and behavior.

Acknowledgment

We thank the reviewers for their valuable feedback which led to significant improvement of the paper. This research started during the SAMSI (Statistical and Applied Mathematical Sciences Institute) program on Challenges in Computational Neuroscience, 2015–2016. We thank Timothy Johnson and Martin Lindquist for their valuable comments and suggestions. We are also very grateful to Qing Feng for providing the JIVE software. In addition, we thank the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. This work was partially supported by the National Science Foundation under Grant DMS-1127914 to SAMSI. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.neuroimage.2017.02.072>.

References

- Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L., Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., Bryant, E., Hartley, T., Footer, O., Bjork, J.M., Poldrack, R., Smith, S., Johansen-Berg, H., Snyder, A.Z., van Essen, D.C., Consortium, W.-M.H., 2013. Function in the human connectome: task-fMRI and individual differences in behavior. *Neuroimage* 80 (October), 169–189.
- Buckner, R.L., Andrews-Hanna, J.R., Schacter, D.L., 2008. The brain's default network. *Ann. New Y. Acad. Sci.* 1124 (1), 1–38.
- Dray, S., 2008. On the number of principal components: a test of dimensionality based on measurements of similarity between matrices. *Comput. Stat. Data Anal.* 52 (4), 2228–2237.
- Feng, Q., Hannig, J., Marron, J.S., 2015. Non-iterative joint and individual variation explained. *arXiv:1512.04060*.
- Feng, Q., Hannig, J., Marron, J.S., 2016. A note on automatic data transformation. *Stat. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Essen, D.C.V., Jenkinson, M., 2013. The minimal preprocessing pipelines for the human connectome project. *NeuroImage* 80, 105–124. (mapping the Connectome).*
- Hastie, T., Tibshirani, R., Friedman, J., Franklin, J., 2005. The elements of statistical learning: data mining, inference and prediction. *Math. Intell.* 27 (2), 83–85.
- Hotelling, H., 1936. Relations between two sets of variates. *Biometrika* 28 (3/4), 321–377.
- Howard, R., Brammer, M., David, A., Woodruff, P., Williams, S., et al., 1998. The anatomy of conscious vision: an fmri study of visual hallucinations. *Nat. Neurosci.* 1 (8), 738–742.
- Kosslyn, S.M., Thompson, W.L., Kim, I.J., Alpert, N.M., et al., 1995. Topographical

- representations of mental images in primary visual cortex. *Nature* 378 (6556), 496–498.
- Krishnan, A., Williams, L.J., McIntosh, A.R., Abdi, H., 2011. Partial least squares (pls) methods for neuroimaging: a tutorial and review. *Neuroimage* 56 (2), 455–475.
- Lock, E.F., Hoadley, K.A., Marron, J., Nobel, A.B., 2013. Joint and individual variation explained (jive) for integrated analysis of multiple data types. *Ann. Appl. Stat.* 7 (1), 523.
- McIntosh, A., Bookstein, F., Haxby, J.V., Grady, C., 1996. Spatial pattern analysis of functional brain images using partial least squares. *Neuroimage* 3 (3), 143–157.
- Menzies, L., Achard, S., Chamberlain, S.R., Fineberg, N., Chen, C.-H., del Campo, N., Sahakian, B.J., Robbins, T.W., Bullmore, E., 2007. Neurocognitive endophenotypes of obsessive-compulsive disorder. *Brain* 130 (12), 3223–3236.
- Nestor, P.G., O'Donnell, B.F., McCarley, R.W., Niznikiewicz, M., Barnard, J., Shen, Z.J., Bookstein, F.L., Shenton, M.E., 2002. A new statistical method for testing hypotheses of neuropsychological/mri relationships in schizophrenia: partial least squares analysis. *Schizophr. Res.* 53 (1), 57–66.
- Odum, A.L., 2011. Delay discounting: I'm ak, you're ak. *J. Exp. Anal. Behav.* 96 (3), 427–439.
- Pezeshki, A., Scharf, L.L., Azimi-Sadjadi, M.R., Lundberg, M., 2004. Empirical canonical correlation analysis in subspaces. In: *Signals, Systems and Computers, 2004. Conference Record of the Thirty-Proceedings of the Eighth Asilomar Conference on*, vol. 1. IEEE, pp. 994–997.
- Satterthwaite, T.D., Wolf, D.H., Erus, G., Ruparel, K., Elliott, M.A., Gennatas, E.D., Hopson, R., Jackson, C., Prabhakaran, K., Bilker, W.B., et al., 2013. Functional maturation of the executive system during adolescence. *J. Neurosci.* 33 (41), 16249–16261.
- Sereno, M.I., Allman, J., 1991. Cortical visual areas in mammals. *Neural basis Vis. Funct.* 4, 160–172.
- Shamosh, N.A., Gray, J.R., 2008. Delay discounting and intelligence: a meta-analysis. *Intelligence* 36 (4), 289–305.
- Smith, S.M., Nichols, T.E., Vidaurre, D., Winkler, A.M., Behrens, T.E., Glasser, M.F., Ugurbil, K., Barch, D.M., van Essen, D.C., Miller, K.L., 2015. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nat. Neurosci.* 18 (11), 1565–1567.
- Tippett, W.J., Black, S.E., 2008. Regional cerebral blood flow correlates of visuospatial tasks in alzheimer's disease. *J. Int. Neuropsychol. Soc.* 14 (06), 1034–1045.
- van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.-M.H., et al., 2013. The wu-minn human connectome project: an overview. *Neuroimage* 80, 62–79.
- Wold, S., Ruhe, A., Wold, H., Dunn, W., III, 1984. The collinearity problem in linear regression. the partial least squares (pls) approach to generalized inverses. *SIAM J. Sci. Stat. Comput.* 5 (3), 735–743.