Blockwise Human Brain Network Visual Comparison Using NodeTrix Representation

Xinsong Yang, Lei Shi, Madelaine Daianu, Hanghang Tong, Qingsong Liu and Paul Thompson



Fig. 1: NodeTrix visual comparison between the brain network of high v.s. low Composite Creativity Index (CCI) group. Significant difference is found on the connectivity of prefrontal lobes bilaterally. Diagonal cells (intra-block edges) are greyed out.

Abstract—Visually comparing human brain networks from multiple population groups serves as an important task in the field of brain connectomics. The commonly used brain network representation, consisting of nodes and edges, may not be able to reveal the most compelling network differences when the reconstructed networks are dense and homogeneous. In this paper, we leveraged the block information on the Region Of Interest (ROI) based brain networks and studied the problem of blockwise brain network visual comparison. An integrated visual analytics framework was proposed. In the first stage, a two-level ROI block hierarchy was detected by optimizing the anatomical structure and the predictive comparison performance simultaneously. In the second stage, the NodeTrix representation was adopted and customized to visualize the brain network with block information. We conducted controlled user experiments and case studies to evaluate our proposed solution. Results indicated that our visual analytics method outperformed the commonly used node-link graph and adjacency matrix design in the blockwise network comparison tasks. We have shown compelling findings from two real-world brain network data sets, which are consistent with the prior connectomics studies.

Index Terms-Brain Network, Visual Comparison, Hybrid Representation.

1 INTRODUCTION

The study of human brain networks stands at the heart of brain connectomics, an increasingly prominent field of computational neuroscience to understand the landscape of neural connections in our brain [37]. In this paper, we focus on the macroscopic brain network defined by high-level Region Of Interests (ROI), versus the microscopic brain networks composed of cells, synapses and voxels [4] [15]. The structural brain network in the macroscopic level is typically constructed in two steps. First, magnetic resonance imaging (MRI) is used to accurately

- Xinsong Yang, Lei Shi and Qingsong Liu are with SKLCS, Institute of Software, Chinese Academy of Sciences. Lei Shi is the corresponding author, E-mail: shijim@gmail.com. Xinsong Yang is also with Beijing University of Posts and Telecommunications.
- Madelaine Daianu and Paul Thompson are with Imaging Genetics Center, Mark & Mary Stevens Institute for Neuroimaging & Informatics, University of Southern California.
- Hanghang Tong is with School of Computing, Informatics and Decision Systems Engineering, Arizona State University.

Manuscript received 31 Mar. 2016; accepted 11 July 2016; date of publication Jan. 2017; date of current version xx xxx 2016. For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org. segment the human cortex into ROIs¹, which serve as nodes in the brain network. Second, diffusion MRI or Diffusion Weighted Imaging (DWI) is conducted to capture the water diffusion along white matter pathways on the cortex, which detects weighted edges (i.e., fiber connections) among ROI nodes on the brain network.

One of the most interesting questions in brain networks relates to the anatomical difference in connectivity that may be influenced by subject-level characteristics, such as age, gender, genetic profile and disease status. For instance, among Alzheimer's disease (AD) patients, the connectivity level in areas related to the memory center of the brain, is significantly lower than those in the matched healthy controls [11]. Extracting and understanding these discriminative connectivity features can significantly boost our understanding of the brain, and the application in disease prediction, gene engineering and many others. Although much progress has been made, the detection and validation of discriminative brain connectivity features remains a daunting task. First, there is still a need for large-scale and high-quality brain network measurements. Connectivity features that describe the network among brain regions can be very high-dimensional, i.e., consisting of thousands of individual connections. The studies that use

¹ROIs are highly specialized brain regions. In this work, we apply the Desikan-Killiany parcellation [14] by FreeSurfer to define 68 or 70 ROIs on the brain cortex. Other popular cortex parcellations include Desikan-Killiany-Tourville Protocol (62 ROIs) and Destrieux Atlas (148 ROIs).

advanced methods can cause overfitting when the number of subjects is small, possibly leading to less reliable results. Second, it is costly to validate numerous computational models proposed in sophisticated brain network studies for which subject-matter experts are required.

Visualization tools are important for addressing some of these challenges to compare high-dimensional connectivity patterns in the brain. On one hand they provide an interactive platform to fuse multiple sources of brain networks and measurements to test various computational models and parameters. On the other hand, the visual interface allows domain experts to work seamlessly with the data sets, analysis algorithms and comparison results for efficient evaluation of the data. A natural question arises regarding which visual representation works the best for comparing human brain networks. The pioneering work by Alper et al. [5] contributed towards answering this question. They found that the node-link representation, though popular in illustrating social and information networks, is not the best choice to visually compare weighted brain networks. Meanwhile an adjacency matrix design performs better in low-level brain network comparison tasks such as finding connectivity patterns and neighboring elements. In this work, we implement some of their user studies on real-world brain networks (versus the synthetic data set in Alper et al. [5]). While most of our results align well with Alper's in that the matrix design outperforms the node-link graph, it is critical to notice that on real brain networks, both methods bear very low task accuracies and high completion times (e.g. < 72% and > 34s in Table 2). We argue that this is because the human brain network is dense and might be quite similar among population groups (Section 3.1). It can be infeasible, or at least difficult, to ask users to visually compare dense and homogeneous human brain networks with respect to their weighted connectivity patterns.

In this paper, we describe a new way to visually compare human brain networks. Instead of directly comparing the traditional ROIbased brain networks, we exploit the inherent block structure of the human brain by considering the blockwise brain network comparison problem (Section 3.2). The rationale behind the blockwise comparison can be explained as follows: from an anatomical standpoint, the human brain is composed of multiple blocks - meaning that ROIs can be divided into left and right hemispheres, and can be further broken down into several cerebral lobes, often linked to highly specialized brain functions. For example, the temporal lobe is responsible for forming long-term memory, while the insula is involved in human consciousness. Furthermore, due to the physical block structure, domain users generally compare brain network connectivity patterns in relation to this anatomical breakdown. It has been shown that Alzheimer's disease may lead to alterations in fiber connectivity in the temporal and parietal lobes primarily, predominantly in the left hemisphere [40]. More blockwise comparison studies in clinical practice are surveyed in Section 3.2. Computationally, we show in Section 4.2 that incorporating the block structure of ROIs can improve the predictive performance in classifying human brain networks while not interfering with the extraction of comparative patterns. In summary, the contribution of this work can be described as follows.

- We propose the blockwise brain network comparison problem over the ROI-based network comparison studied previously (Section 3.2), and present an integrated visual analytics framework to address this new problem (Section 3.3). The effectiveness of our framework is demonstrated through case studies over two real-world data sets: one comparing the average brain network of a group of AD patients with that of healthy controls; another comparing between two groups of normal people with high and low creativity scores (Section 6).
- In this framework, we introduce an adaptive clustering algorithm on the ROIs of human brain networks to maximize the predictive classification performance with the resulting blockwise connectivities (Section 4). The key challenge is to overcome local maxima in optimizing the feature selection process.
- We adapt a hybrid visualization design, aka NodeTrix by Henry et al [20] as shown in Figure 1, to optimize the blockwise network visual comparison task, which takes the ROI block struc-

Table 1: Results on the Trend task (blue = best outcome).

Measure	Acc.	Sig.	Time	Sig.
Node-Link Overlaid (NO)	0.50	< 0.001	25.63	>0.05
Matrix Overlaid (MO)	0.86		27.16	
Node-Link Side-by-Side (NS)	0.92		24.02	
Matrix Side-by-Side (MS)	0.81		24.61	

Table 2: Results on the Connectivity task (blue = best outcome).

Measure	Acc.	Sig.	Time	Sig.
Node-Link Overlaid (NO)	0.33	< 0.001	62.65	>0.05
Matrix Overlaid (MO)	0.72		34.93	
Node-Link Side-by-Side (NS)	0.56		56.92	
Matrix Side-by-Side (MS)	0.72		70.15	

ture into account when displaying the human brain network (Section 5.1). To further improve the performance of NodeTrix in real-life comparison tasks, we propose to apply force-directed edge bundling algorithms according to the block representation of brain networks (Section 5.2). The NodeTrix visualization is evaluated in the user study by comparing to the classical nodelink and adjacency matrix visualization on blockwise brain network comparison tasks.

2 RELATED WORK

Brain network (aka the connectome) visualization is a popular interdisciplinary research topic. In the neuroscience community, Margulies et al. summarized three classes of visualizations based on the various steps in their data transformation pipeline [31]. On the low-level DWI data, in particular Diffusion Tensor Imaging (DTI), tensor glyphs in different shapes were invented to display the rich data dimensionality. For example, the latest design by Prčkovska et al. [34] augmented spherical polar ploy glyphs [21] with optimized shape and color scheme. Their design can highlight multiple maxima in order to make the "peaks" of glyphs more distinct in the visualization. On the streamline reconstructed by tractography [25], traditional visualizations display a deterministic fiber tract without considering its uncertainty and dimensionality. Berres et al. [8] introduced nested surface layers to visualize probabilistic tractograms which can indicate the connectivity score between voxels. On the connectome constructed from whole-brain fiber tracts, which is more related to our work, both graph, matrix and chordmap [32] metaphors have been applied for visualization. For a more extended literature review on connectome visualization, we refer to this survey [31].

Most recently in the visualization community, Al-Awami et al. proposed Neurolines [4], a novel visual metaphor to examine the neuronal connectivities at nano-scale. Neurolines designed a multiscale subway-like visualization which can be interactively rendered and scaled up to display thousands of neurites for the detailed analysis of neuronal structures and their connectivity. Everts et al. [15] studied the visual abstraction of brain fiber tracts from Diffusion Tensor Imaging (DTI) data. Through arranging the fiber tracts by their local similarity at multiple scales, the global white matter architecture inside the brain can be constructed from fiber tracts, with less visual clutters than previous approaches. Overall, these compelling methods on the visual abstraction of brain connectivity mostly work at the microscopic level of human brains, e.g., the neuron-level [4] or the voxel-level [15]. On the comparison of human brain networks across population groups, the difference at the microscopic level can be less significant, due to their higher uncertainty, than the macroscopic connectivities, i.e., the ROI-level brain networks.

The task of visually representing and comparing high-level brain networks is relatively less explored, though the general topic of visual comparison has been a focus of the information visualization research [19] [30]. Bach et al. presented Matrix Cube [6], an effective 3D cube visual representation to display the functional brain networks. Their focus is on the dynamic pattern of brain networks, but not on the comparison between groups of average brain networks. The work by

Table 3: List of comparison studies using brain networks; blue text indicates the analysis result related to the block structure on human brains.

Paper	Topic	Data & Subject	Method	Analysis Results	GS #Cite
Wernicke 1874 [42]	aphasia	10 subjects	clinical	disconnection of the sensory speech zone from the	1429
				motor speech area in the left hemisphere	
Geschwind 1965 [18]	Disc.	animals and man	clinical	disconnection syndrome between the visual-limbic	3146
				system, non-limbic associations	
Horwitz 1987 [24]	AD	42 subjects (21	PET	fewer correlation between frontal and parietal lobes,	157
		AD, 21 control)		larger between the cerebellum and temporal lobe	
Buckner 2005 [9]	AD	764 subjects (5	mixed	convergence of effects in posterior regions, interac-	1138
		studies)		tions with medial temporal lobe	
Daianu 2013 [12]	AD	111 subjects (83	3T MRI,	k-core breakdown in the left hemisphere, left/right	40
		AD, 28 control)	DTI	asymmetry intensifies	
Ho 2003 [22]	SZ	96 subjects (73	MRI	reduction in frontal lobe white matter volume, in-	499
		SZ, 23 control)		crease in frontal lobe cerebrospinal fluid volume	
Kubicki 2005 [28]	SZ	47 subjects (21	DTI, MRI	decreased diffusion anisotropy in fiber bundles con-	356
		SZ, 26 control)		necting higher-level regions	
Olabi 2011 [33]	SZ	1795 subjects	MRI	decreases over time in frontal, parietal and temporal	177
		(27 studies)		white matter volume	
Chiang 2009 [10]	Gene	92 twins	4T DTI	FA heritable in frontal, parietal and occipital lobes	281
Scott 2010 [35]	Gene	71 children	fMRI	CNTNAP2 expressed in frontal lobe circuits	125
Liu 2010 [29]	Gene	57 subjects	fMRI	decreased prefrontal-related connectivities between	72
				prefrontal regions and the posterior cortices	

Alper et al. [5] might come closest to our study. They evaluated the effectiveness of the node-link and matrix representation in weighted graph comparison. Their results favored the matrix design with an overlaid comparison method. In this work, we study the same research topic on visual comparison, but consider a new problem of blockwise brain network comparison after a survey of neuroscience literature on their high-level human brain connectivity analysis requirement. Notably, we introduce the computational method for the ROI clustering and integrate this method into a coherent visual analytics framework that can guide the effective discovery of brain network biomarkers across population groups.

3 OVERVIEW

3.1 Preliminary User Study

We start from a controlled user experiment to compare two basic visualization designs for the brain network comparison: the node-link graph and the adjacency matrix. In addition, we consider two visual comparison methods: the overlaid comparison [5] and the straightforward side-by-side comparison. When combined, we study four candidate interfaces: node-link overlaid (NO), matrix overlaid (MO), nodelink side-by-side (NS), matrix side-by-side (MS).

Alper et al. [5] first conducted this comparative study on NO and MO visualizations over a group of synthetic brain networks. In contrast, our study is carried out on real-world brain networks, comparing 42 AD patients with 50 healthy controls. Real-world brain networks have two significant differences from synthetic networks that can affect the user study result. First, *high network density*. For instance, the brain network of each of 92 subjects in our data set has 454~1021 edges (i.e., fiber connectivity with nonzero strength) between 68 ROI nodes, leading to a network density of 0.2~0.45, whereas the synthetic network in [5] has a 5% or 10% density. Second, *similar network topology* across subjects. We have built a Pearson's correlation coefficient matrix by the weighted topology vector of 92 subjects in our data set, each has a vector length of 2278. More than 90% pairs of subjects have a weighted topology correlation greater than 0.8, and the average correlation is close to 0.9.

Our experiment result with 12 subjects over the real brain network data set is summarized in Table 1 (Trend task) and Table 2 (Connectivity task). The task design follows those in Alper et al. and the implementation details are the same with our second user study explained in Section 6.1. There are two observations with respect to the results by Alper and colleagues that motivate our follow-up research. First, the comparison between MO and NO methods leads to the same outcome: when the overlaid method is applied, the matrix design is

significantly better in accuracy and completion time when compared to the node-link visualization, with the exception of the completion time in the Trend task. However, the actual task accuracy deteriorates greatly on the real data set: on the Trend task, it drops from 0.96 (MO) and 0.85 (NO) in Alper et al. to 0.86 and 0.50 in our study; on the Connectivity task, it drops from 0.90 (MO) and 0.71 (NO) to 0.72 and 0.33. The completion time measure has a similar pattern. These effects can be attributed to the dense and homogeneous nature of real brain networks in the ROI level, which prevents users from accomplishing visual comparison tasks. Our second finding, the side-by-side comparison design, though not considered in the previous study, leads to comparable performance to the overlaid design when the network density is high. This is especially true judging from subjective ratings. The scores for the node-link side-by-side design in user experience and usability (3.83 and 3.92 in the 7-point likert scale, higher is better) are only slightly below those of the matrix overlaid design (4.07 and 4.08). Note that the side-by-side comparison bears an additional advantage in much shorter training time than the overlaid comparison for ordinary users.

3.2 Blockwise Brain Network Comparison

The preliminary user study result indicates that it may not be effective to visually compare real brain networks solely at the macroscopic ROI level. The cortical connectivity at that level is adversely dense and little visual difference is presented between individual subjects. In fact, from the domain expert's view (neuroscientists, e.g., co-authors of this paper; doctors, our collaborators), rarely do they examine the brain network difference only from the ROI level between individual subjects. In this work, motivated by the inherent block structure of human brains (e.g., left/right hemispheres, anatomical lobe classifications), we consider a new problem of blockwise network comparison over the existing ROI-based brain network. At the ROI block level, the abstracted connectivity pattern can be more salient in visual comparison and far more useful for domain experts.

To validate the importance and characterize user tasks on the new problem, we did a literature review of 33 neuroscience studies on the comparison of human brain networks (aka the white matter connectivity) between diagnostic groups (AD, Schizophrenia, genetic disorders). We examined each study by the criterion that whether or not the comparison analysis had detected significant patterns in the ROI block level, rather than only detected between particular ROIs. The representative literature is given in Table 3 and a full list is documented in the supplemental material. In summary, 22 out of 33 studies (66.7%) in the review met our criterion. Dating back to the 19th



Fig. 2: Visual analytics framework for blockwise brain network comparison.

Edge Bundling

century, Wernicke described the disconnection syndrome as a loss of connectivity between the sensory speech and the motor speech areas in the left hemisphere [42], which was among the first block-level connectivity studies. On the group-based comparison studies between AD patients and controls, we examined 12 relevant studies cited by Daianu et al. [12]. Eight studies meet our criteria, details can be found in Table 3. On the comparison study of Schizophrenia (SZ) patients, we reviewed 7 highly cited papers on the white matter deficiency - 6 of which reported results based on higher-order brain subnetworks [33], aka block-level networks considered here. We also reviewed 12 brain network studies related to genetics, cited by Thompson and colleagues [39]. 6 of which fitted our criterion.

Based on the literature review, we conclude that the ROI block level difference among diagnostic groups are crucial for the classification and prediction of certain disease or symptom. We summarize a set of blockwise brain network comparison tasks based on the practice in the literature and feedbacks from domain experts.

B1: Given one ROI block, identify noticeable changes of its local connectivity between diagnostic groups.

B2: Given one ROI block, identify noticeable changes of its global connectivity (i.e., between the ROIs of the current block and the ROIs in other connected blocks) between the diagnostic groups.

B3: Identify ROI blocks that have local/global connectivity changes between the diagnostic groups.

B4: Identify pairs of ROI blocks that present connectivity changes between the diagnostic groups.

3.3 Visual Analytics Framework

Consider the design space for the blockwise comparison task defined above. The user study result in Section 6.1 shows that applying the classical node-link or matrix design leads to poor visual comparison performance, especially on the inter-block connectivities. In a few ROI definitions, two ROIs within the same block can be geometrically separate in the projected brain cortex, making it difficult for users to visually compare the block-level connectivities.

In this work, we propose a visual analytics framework to integrate the algorithmic ROI block generation with the interactive visualization to help identify and analyze blockwise brain network connectivity differences. As shown in Figure 2, in the first stage, we apply clustering algorithms to detect a two-level ROI block hierarchy (Section 4). In the top level, the functional blocks divide all ROIs based on an anatomical or functional classification (e.g., the lobe classification); and in the bottom level, each functional block is partitioned into several comparison blocks, in a way to maximize the group-level difference in the network comparison. In the second stage, the raw ROI connectivity and its hierarchical block structure are integrated and represented by a hybrid NodeTrix visualization (Section 5). The functional blocks are shown as matrices, the comparison blocks are shown as merged rows/columns inside the matrix, and each ROI is represented by the individual row/column. To optimize the inter-block connectivity comparison across diagnostic groups, we introduce the force-directed edge bundling algorithm to visually aggregate the connections on the same pair of ROI blocks.

Table 4: Notations for the brain network classification.

SYMBOL	DESCRIPTION
N, G_i	# of subjects and their brain networks
n, p, e _j	# of nodes, # of edges, each edge in the network
$\mathscr{X}, X, x_i, x_{ij}$	edge weight variable, weight matrix on all sub-
	jects, weight vector on G_i , the component on e_j
\mathcal{Y}, y, y_i	outcome variable, its value on all subjects and G_i

4 ROI CLUSTERING

The top-level ROI hierarchy is composed of functional ROI blocks. In this work, we adopt the cerebral lobe classification that partitions each hemisphere into six lobes: frontal lobe, parietal lobe, occipital lobe, temporal lobe, limbic lobe, and insula. Other functional cortex parcellations are also applicable, depending on the targeted comparison task. In the second-level ROI hierarchy, we apply clustering algorithms to detect comparison blocks that best reveal the connectivity difference among diagnostic groups. The main idea is to link the accuracy in the predictive analysis, i.e., the classification of brain networks into diagnostic groups, with the visual comparison performance among groups of brain networks. The problem of optimizing the blockwise brain network visual comparison is converted to finding the best ROI clustering that maximizes the brain network classification accuracy.

4.1 Brain Network Classification

1

The brain network classification problem is defined as predicting the label of a subject (outcome) using the features on the subject's brain network. Here the label indicates the diagnostic group. Throughout this paper, we consider the binary classification, i.e., the label has two levels, e.g., diseased or healthy subject. The brain network features are limited to the set of connectivity strengths, i.e., the number of fibers going through each pair of ROIs. These connectivity strengths become the edge weight in the brain network, so we also call them edge features interchangeably.

The formal notations in this definition are listed in Table 4. We consider N subjects and their brain networks, denoted by G_1, \dots, G_N . Each brain network has n nodes (i.e., ROIs) and p edges (i.e., fiber connections) between pairs of nodes, denoted by e_1, \dots, e_p . The edge weight (i.e., connectivity strength) is defined by the variable \mathscr{X} . On G_i , the network of the *i*th subject, the edge weights are denoted by the vector $X_i = (x_{i1}, \dots, x_{ip})'$. At the network level, each subject is associated with a binary outcome variable \mathscr{Y} (i.e., subject label). The value of \mathscr{Y} on N subjects is denoted by the vector $Y = (y_1, \dots, y_N)'$. Using a basic logistic regression model, the objective function to best fit the model given data is written as minimizing the Negative Log Likelihood (NLL).

$$Min \, NLL(W) = \sum_{i=1}^{N} log(1 + e^{-y_i W^T X_i})$$
(1)

where $W = (w_1, \dots, w_p)'$ denotes the weight vector for all p edge features. The edge with a larger weight means that it has a higher influence on the outcome.



Fig. 3: Tree search in the agglomerative clustering of ROIs. Each node represents an ROI clustering setting where the node label gives the resulting classification accuracy. Each link indicates a feasible binary combination of two ROI clusters.

When classifying brain networks, it is known that using all the edge features for training can lead to severe overfitting, which downgrades the classification accuracy in the testing set. The general approach is to apply feature selection that only picks the most informative features to predict the label. Here we adopt the lasso model [41] which adds a L_1 regularization term ($||W||_1$, i.e., L_1 norm of the weight vector W) for the edge feature selection. The objective function of the lasso model is given by

$$Min \sum_{i=1}^{N} log(1 + e^{-y_i W^T X_i}) + \lambda ||W||_1$$
(2)

where the parameter $\lambda > 0$ controls the degree of overall model sparsity. A larger λ will lead to less selected edge features. The weight vector W determines the feature selection result. The *i*th edge feature will be selected if $w_i \neq 0$, and unselected if $w_i = 0$. To determine the sparsity parameter λ , we iterate over a list of logarithmically spaced parameter choices within the feasible range for nonzero weight vectors. The best λ is chosen as the one leading to the highest classification accuracy in 10-fold cross-validations.

The lasso model is popular because of its effectiveness to achieve good classification accuracy. However, lasso does not capture the interaction effect among features, nor does it take the feature group information as indicated by the block structure in our work. In our final classification model, we apply a recent variant of lasso, namely the sparse group lasso (SGL) [36] which incorporates feature grouping information to optimize the classification. These edge feature groups can be directly linked to the ROI clustering in that when network nodes are grouped, their edges are aggregated correspondingly. The objective function of the SGL model is defined by

$$Min \sum_{i=1}^{N} log(1 + e^{-y_i W^T X_i}) + \alpha \lambda ||W||_1 + (1 - \alpha) \lambda \sum_{m=1}^{M} ||W^{(m)}||_2$$
(3)

where *M* denotes the number of feature groups, $W^{(m)}$ is the partial weight vector of the *m*th feature group. The parameter α controls the the group-wise sparsity. The remaining question is to determine which ROI clustering provides the best feature grouping information for the brain network classification.

4.2 Optimal ROI clustering

We start from a trivial setting of n ROI clusters, i.e., one cluster per ROI. In this setting, the objective function of the SGL model (Eq.



Fig. 4: ROI clustering performance by the roll-back greedy scheme.

(3)) degenerates to that of the basic lasso model (Eq. (2)) because $\sum_{m=1}^{M} ||W^{(m)}||_2 = ||W||_1$. Any nontrivial *k*-ROI-clustering (k < n) can be seen as a *k*-partition of the *n* initial clusters. The number of possible nontrivial ROI clusterings amounts to the Bell number [7]. Therefore, finding the best ROI clustering that yields the highest classification accuracy leads to an NP hard problem requiring $O(n^n)$ computation time, when the exhaustive search on all possible ROI clusterings is applied.

The search-based solution can be improved by only combining ROIs within each functional block, but the complexity is still exponential. A more feasible approach is to apply the greedy search by clustering ROIs in an agglomerative manner. An illustration of this process, i.e., the search tree, is given in Figure 3(a). In this tree, each node represents an ROI clustering setting and the node label indicates the classification accuracy when this ROI clustering is applied, i.e., a_0 for the initial clustering and a_{ij} for the *j*th feasible clustering in the *i*th step. The directed link indicates the process to merge two ROI clusters into a new cluster, moving from one clustering setting to another. The search starts from the root of the tree where the trivial *n*-ROI-clustering is applied. In the next step, all feasible binary combinations of trivial ROI clusters are examined by classifying with the SGL model under the new clustering. These clusterings are represented as the nodes in the second hierarchy. The greedy search scheme will pick the clustering setting leading to the highest accuracy in each step, e.g., in this figure, the clustering that combines ROI cluster {#3} and {#4} which yields a classification accuracy of $a_{14} \ge a_{1j}$. After that, the search process is repeated one step further and combines the ROI cluster {#3,#4} with $\{\#5\}$. In n-1 steps, all ROIs are combined into one cluster and the search process terminates.

The greedy search scheme leads to a cubic computational complexity assuming the SGL solver has a constant cost, which is feasible for computation because the number of ROIs (n) is small. However, due to the non-concave and combinatorial nature of the search space, the greedy search is highly likely to get stuck in the local maxima. To address this issue, we propose an improved greedy search scheme by adding roll-back operations. As shown in Figure 3(b), instead of using the classification accuracy as the search heuristic, we turn to maximize the accuracy improvement from the last step, denoted by I_{ii} for the *j*th combination in the *i*th step. This allows us to introduce a slope constraint S. In case the best accuracy improvement in certain step of the greedy search falls below S, e.g., $I_{23} < S$ in Figure 3(b), we drop this search branch by overriding the accuracy improvement of its upper node (i.e., I_{14}) to I_{23} . This branch will not be searched any more because $I_{14} < S$. The search process restarts with a roll-back one step further, i.e., to the root node in this figure. As the previous best choice of I_{14} has been dropped, the second best clustering indicated by the improvement I_{12} is picked and the greedy search continues from this new branch. Except for the slope constraint S, we control another two parameters to keep the search process going: the maximal roll-back depth D_{max} and the slope constraint step ΔS . When the search rolls back more than D_{max} hierarchies consecutively, we loosen the slope constraint by subtracting one step of $|\Delta S|$ from S. This also fits the actual curve of the accuracy improvement. As shown in Figure 4, the best accuracy appears in the middle of the clustering process, and after that drops almost monotonically.

The greedy search with roll-back scheme can lead to significantly



Fig. 5: The NodeTrix visualization design for blockwise visual comparison of brain networks.

increased computational complexity as more than one path on the search tree can be traversed. We introduce the caching strategy to reduce the computation requirement. As shown in Figure 3(b), the same ROI cluster combination can appear in different steps of the search tree, e.g., I_{11}/I'_{21} and I_{13}/I'_{23} . The strategy is to assume the accuracy improvement by the same combination does not change much across search steps. Under this assumption, every new search in later steps can first fetch the accuracy improvement from the cache established in earlier steps. This greatly reduces the cost in repeating SGL modelings for every new clustering.

We have applied the greedy search and the roll-back scheme in the 113-subject high-CCI v.s. low-CCI brain network data set. The performance curves are depicted in Figure 4. The ROI clustering will mildly increase the classification accuracy before a half of the ROIs are grouped together, while the number of features selected remains constant. This demonstrates that the inherent block structure of human brain networks can be leveraged in its classification to better abstract, but not undermine, the comparison pattern, if only an adequate number of clusters is specified. When the clustering proceeds to generate over-simplified block structure, the classification accuracy will drop quickly and it is hard to discover good blockwise comparison patterns any more. We have integrated the classification accuracy measure during the search process into the visualization interface to guide the interactive visual comparison of brain networks.

5 VISUALIZATION DESIGN

We consider two design goals for the brain network visualization in our scenario: 1) optimize the user performance in the visual comparison of brain networks; 2) incorporate the two-level ROI block hierarchy into the visualization for blockwise comparisons. It has been shown in [5] that the node-link representation is not the best choice for visual comparison. Meanwhile, the adjacency matrix can be better than the node-link design in user performance, but it is difficult to embed multiple ROI block hierarchies into the matrix design while highlighting the blockwise comparison patterns. In this work, we explore the hybrid approach, that uses the NodeTrix representation to visualize the blockwise brain network connectivity for comparison between population groups. The NodeTrix metaphor was invented by Henry et al [20] to display social networks having inherent clustering nature. In their design, the dense social communities can be represented by matrices, while the sparser inter-community connections are drawn in the node-link diagram. We echo the same design rationale on brain networks. The weighted graph clustering result (k=12) on our brain network data set shows a very close structure to the 12-lobe brain parcellation: 46 out of 68 ROIs have the same membership in the graph clustering and the lobe classification. This demonstrates that the intralobe connectivity is on average denser than their inter-lobe counterpart, which advocates the use of the NodeTrix design.

5.1 NodeTrix Representation

Figure 5 depicts an example of using NodeTrix to display the blockwise brain network comparison in a side-by-side design. For each brain network, the high-level visual representation is composed of several matrices, each corresponding to a functional block in the brain, i.e., one of the anatomical lobes in our setting. Each row/column in the matrix represents a single ROI whose index is drawn as labels on the border of the matrix, e.g., the largest frontal lobe contains 11 ROIs in each hemisphere. The rows and columns can be combined together to represent the comparison blocks in the ROI clustering hierarchy, as shown in Figure 1. Inside each matrix, the entries drawn as colorcoded cells indicate the intra-lobe ROI connections. The red hue is selected as the default fill color, while the color saturation is used to represent the fiber strength of each ROI connectivity. The more saturated the color, the higher connectivity strength. The detailed strength-color mapping is indicated and controlled by the color map slider in the center of the control panel in Figure 5. User interactions are supported to tune the color map and optimize the visual comparison (Section 5.3). Note that the diagonal cells in each matrix are greyed out as the fiber loops on each ROI are unimportant for the comparison of brain networks. For the inter-lobe ROI connections, curved edges are drawn between the matrices where the two endpoints are on the brim of the source and target ROI column/row. Both the edge thickness and color saturation indicate their connectivity strength.

Our design supports three comparison views by different projections of the brain network: sagittal view (Figure 9), axial view (Figure 5), coronal view (supplemental video). They are also distinguished by separate background brain sketches. On the layout of each view, we consider optimizing both the topographical proximity and the performance for visual comparison. The matrices representing the functional blocks are placed in the center of their anatomical lobes. The ROIs within each functional/comparison block are initially ordered by their ROI indices. We also provide an interaction to reorder ROIs according to their magnitude of difference in the comparison.

5.2 Brain Connectivity Edge Bundling

The brain connectivities between ROI blocks are drawn as B-spline curves between matrices. These edges are bundled together by selecting appropriate control points for each curve. We propose two edge bundling algorithms for separate objectives: the semantic bundling to optimize the visual comparison and the geometric bundling to reveal the physical routing path of the underlying fiber pathways.

Semantic edge bundling. In this method, the objective is to reduce the visual clutter raised by the dense ROI-level connections while preserving the blockwise connectivity patterns. We achieve this by a three-step force-directed edge bundling algorithm.

In the first step, we cluster all the ROI connectivity edges into groups according to their blockwise semantic information. The edges connecting the same source and destination matrices are grouped together and share the same set of control points. In other words, edges will be bundled to represent blockwise connectivities. In the second step, we determine the initial placement of control points in each edge bundle. To keep their original connectivity structure, the control points are uniformly sampled from the straightline connecting the center of their source and destination matrices. In the third step, we apply the force-directed edge bundling algorithm to adjust these control points to further alleviate visual clutter. The basic idea is to move the control points of all edge bundles towards each other to reduce the number of spatially distributed control points. In more detail, all control points are considered as nodes in a proximity graph. Then the force-directed layout [17] is conducted on this graph to refine the placement of all control points. The key choice is to decide the optimal distance (i.e., proximity) between each pair of control points. We apply Holten and Van Wijk's algorithm in [23]

$$OptDist = Dist_0 \cdot \left(\frac{C_e}{Min(C_e)}\right)^k \quad where \quad C_e = C_a \cdot C_s \cdot C_p \quad (4)$$

where $Dist_0$ denotes the initial distance between the two control points. C_a , C_s , C_p indicates the adjustment based on the angle, length and proximity of the matrix connection lines of these two control points.

Geometric edge bundling. In another method, we start by aligning the brain network connectivity with the fiber tracts reconstructed through tractography. These fiber tracts can be geometrically clustered into bundles [26] where each bundle represents the connectivity between several ROI blocks. The control points for each blockwise connection can be selected from these bundles to reveal the geometric routing of fibers in the human brain.

In more detail, we introduce a two-step algorithm to compute the control points on each fiber track bundle. In the first step, a streamline is extracted from every track in the selected fiber bundle by uniform sampling. Then we use the sampled points in all tracks to compute a centerpiece streamline. The fiber track staying closest to this centerpiece streamline is selected to represent the entire fiber bundle. The second step is to re-sample the representative fiber track to extract a small number of control points for the edge bundling. Each fiber track is composed of a large number of points. We start from the first point and keep track of the accumulative curvature in iterating all the points, which is known as the winding angle. Once the winding angle is larger than a threshold, the current point is sampled and the winding angle is reset to zero. In the final algorithm, we choose a winding angle threshold of 360 degrees according to Tao et al. [38]. In case the winding angle does not reach the threshold after scanning the entire track, we apply uniform sampling. Mathematically, the curvature k_i at point p_i of a fiber track streamline is computed by

$$k_i = \cos^{-1}(\overrightarrow{p_{i-1}p_i} \cdot \overrightarrow{p_ip_{i+1}})$$
(5)

where p_{i-1} and p_{i+1} are neighboring points of p_i on the streamline.

5.3 Interaction for Visual Comparison

Following the visualization design, we support basic network interactions as well as a suite of customized interactions to facilitate the user's visual comparison task on brain networks. At the ROI level, we design an interactive color map, as shown in the range slider of the central panel in Figure 5. The left range selector controls the edge filtering. All the edges below a specified strength threshold will be removed in the display. The right range selector controls the maximal edge strength. The edges above this maximum will be drawn in the largest color saturation. The edges with strength inside the range will have a linear mapping to the color saturation spectrum. When the binary mode is selected, the range selector becomes the point selector. All remaining edges after the filtering will be drawn in the same maximal color saturation. At the ROI block level, we introduce a matrix reorder interaction. All ROIs inside each lobe are first sorted by the difference between the two compared groups of networks. The ROI with the largest difference is placed in the matrix center and all ROIs are arranged from the center to the border by this order. At the

network level we design a network contrast interaction which, when applied, each connectivity strength in the network is subtracted by the minimal strength of the two networks in comparison. Subsequently, only the connectivity difference is shown after this interaction.

6 EVALUATION

6.1 User Experiment on Blockwise Comparisons

We conducted a controlled experiment to understand the effect of visual designs on the user's performance in blockwise brain network comparison tasks. We compared three visualizations: node-link networks, adjacency matrices, and the proposed NodeTrix design (Figure 5). NodeTrix applies a one-level blocked structure mapping brain lobes to the matrix. Within each lobe, the comparison blocks are not used because in our pilot study, users need additional time and training to locate the best clustering parameter. In all three designs, the side-by-side comparison is adopted.

Experiment design. We recruited 16 subjects (10 male, 6 female) who were all computer science graduate students and have experience with data visualization. The experiment applied a full-factorial within-subject design that every subject was tested on including all visualization methods. The experiment was divided into the training session and the test session for each visualization. In the training session, users warmed up by completing the same suite of tasks on a much easier setting (e.g., less nodes and edges in the brain network to compare). The organizer checked the result of each training task and addressed all questions before proceeding to the next step.

Task. All four tasks listed below were conducted on the brain networks from 92 subjects (including controls vs. AD patients) to compare differences at anatomical network level.

T1 (Local connectivity): does the overall edge weight within the right frontal lobe subnetwork decrease or increase in controls when compared to AD patients?

 T_2 (Lobe-Lobe connectivity): does the overall edge weight between the right frontal and the right temporal lobe decrease or increase in controls when compared to AD patients?

T3 (Lobe-ROI connectivity): does the overall edge weight between the frontal lobe and all other ROIs decrease or increase in controls when compared to AD patients?

T4 (Lobe connectivity retrieval): identify one lobe-lobe pair whose connectivity has the largest overall edge weight difference between brain networks of the two diagnostic groups?

Note that for each task×visualization pair, each user was tested on three difficulty levels by controlling the difficulty for detecting differences in the group comparison. This level is shuffled for different users and visualizations to eliminate learning and ordering effect. For each task, we recorded the subject's answer and completion time. The task completion time was measured after the subject had read the question, so that the reading skill variation was excluded. On *T*1 through *T*3, the relevant lobes in the task were highlighted using color coding in the matrix (node) outline of each visualization. On *T*4, all lobe names are labeled on the visualization. This was done to minimize the visual search time for the lobe in the task. During the test session, subjects were also asked to respond to three subjective questions immediately after they completed all the tasks for each visualization. Answers were selected from a 0~6 Likert scale (larger is better).

Q1 (Usability): How much does this visualization help you in completing the tasks and finding the correct answers?

Q2 (Cognitive Performance): How much does this visualization help you to understand the brain network data?

Q3 (User Experience): How would you rank your experience with this visualization design?

Result. Experiment results were analyzed separately for each task. Significant level was set at 0.05 throughout the analysis.

Task Accuracy: The user's accuracy in completing each task is summarized in Figure 6(a). For T1, the intra-lobe local connectivity task, all three visualizations led to rather high accuracies on average (Node-Link: 0.93, Matrix: 0.88, NodeTrix: 0.91), and there were small differences in between. For T2, the lobe-lobe connectivity task, NodeTrix (0.93) had better accuracy than Matrix (0.86) and Node-Link (0.67).



Fig. 6: The user performance distribution of three visual designs.



Fig. 7: Subjective scores for three visual designs in the user study.

A repeated-measure ANOVA test on their average task accuracy (normality and sphericity held, the same below) shows that there were significant group-level differences (F(2,26) = 8.943, p = 0.001). Posthoc test using the Bonferroni correction reveals a significant difference between Node-Link and NodeTrix (p = 0.001), but the difference between Node-Link and Matrix was not significant (p = 0.078). For T3, the lobe-ROI connectivity, Node-Link (0.91) and NodeTrix (0.91) have better average accuracy than Matrix (0.76), but the difference was not significant (p = 0.12). For T4, the lobe connectivity retrieval task, most subjects failed to complete it using the Node-Link visualization deisgn, so we only compared Matrix and NodeTrix designs. In average, NodeTrix (0.98) had a much higher accuracy than Matrix (0.36), and the difference was significant based on a paired samples t-test (t(13) = -13.0, p < 0.001).

Completion Time: The user's completion time in each task is summarized in Figure 6(b). For T1, the intra-lobe local connectivity task, Matrix (5.31s) and NodeTrix (5.17s) took slightly shorter time than Node-Link (6.01s) to complete, but the difference was not significant (p = 0.27). For T2, the lobe-lobe connectivity task, again Matrix (5.43s) and NodeTrix (5.12s) took shorter time than Node-Link (7.75s). A repeated-measure ANOVA test (normality and sphericity held, the same below) showed a significant group-level difference (F(2,26) = 6.598, p = 0.005). Post-hoc tests using the Bonferroni correction revealed a significant difference between Node-Link and NodeTrix (p = 0.045), but not between Node-Link and Matrix (p = 0.057). For T3, the lobe-ROI connectivity, all visualizations took a similar amount of time: Node-Link (4.96s), Matrix (5.62s), Node-Trix (5.18s), and the differences were not significant (p = 0.53). For T4, the lobe connectivity retrieval task, NodeTrix (8.87s) had a much shorter completion time than Matrix (20.21), and the difference was significant under a paired samples t-test (t(13) = 4.11, p = 0.001).

Subjective Questions: The user's subjective scores are summarized in Figure 7. It is shown that Node-Link and NodeTrix have better average scores than Matrix in all three subjective measures. We then applied the Friedman test to analyze their group-level difference, which does not require a normality assumption. Results indicate that while there was no significant difference in the usability score (p = 0.41, the average score, Node-Link: 4.36, NodeTrix: 4.36, Matrix: 3.79), the cognitive performance and user experience were significantly different across groups. On the cognitive performance, Node-Link (4.93) and NodeTrix (4.79) were significantly better than Matrix (4.21) under the Friedman test ($\chi^2(2) = 9.46, p = 0.009$). The post-hoc Wilcoxon signed rank test results were p = 0.004, 0.046. On user experience, Node-Link (3.79), NodeTrix (4.36) and Matrix (3.21) have significant group-level difference by the Friedman test ($\chi^2(2) = 9.33, p = 0.009$), but the post-hoc Wilcoxon signed rank test only detected significant difference between NodeTrix and Matrix (p = 0.004).

Discussion. The user study results show that, on the block-block connectivity tasks for both low-level examination (T2) and high-level retrieval (T4), the NodeTrix design performs significantly better than the Node-Link design, and better than the Matrix design, for both task accuracy and performance time. Especially for the high-level retrieval task of T4, NodeTrix was significantly better than Matrix. This demonstrates the superiority of the NodeTrix design in revealing the block-level brain network structure, which is extremely important in exploratory tasks (like T4) where noisy data is present. On the localized single-block tasks (T1 and T3), while the completion time is close because of the low difficulty level, NodeTrix achieved a similar task accuracy to the Node-Link design, and both were better than the Matrix design, though not significantly. This indicates that the hybrid block representation (NodeTrix) does not affect the local connectivity retrieval, while a full block design (Matrix) could downgrade the local comparison performance when no customized technique (e.g., overlaid design) is applied. For the user's subjective rating, we noticed that all users disliked comparing matrices side-by-side. Though the Matrix design can still work (no significance in usability), user's cognitive gain and experience were significantly poorer than NodeTrix and Node-Link design. This can be attributed to the counter-intuitive nature of the adjacency matrix. NodeTrix compensates this deficiency by keeping a high-level node-link representation.

6.2 Case Study

We evaluated the NodeTrix visualization design on two sets of real-life brain networks. The first is on Alzheimer's Disease patients. More recently, as studied in Thompson et al. [40] and Daianu et al. [12], the AD patient's brain network tends to degenerate in an asymmetric fashion predominantly in the left hemisphere. Furthermore, the metaanalysis in [9] from five studies of 1138 subjects showed a converged effect of AD to the posterior cortical regions in the left hemisphere (mainly in the parietal lobe).

We analyzed the data set from ADNI - a publicly available AD Consortium [2], and applied NodeTrix to reproduce the above brain network patterns. The ADNI data we used for this study contained DWIs from 202 subjects - 50 healthy controls, 72 with early mild cognitive impairment (eMCI), 38 with late mild cognitive impairment (IMCI) and 42 AD patients. The participants underwent whole-brain MRI at 16 different site across North America [12] and their ages ranged between 55 to 90 years. For each subject, an ROI-level network was computed using an automated segmentation method, FreeSurfer [3] based on the Desikan-Killiany atlas [14]. The edge weights among ROIs described the fiber density between pair of ROIs (the number of fibers passing through a pair of ROIs). To focus on the alterations patterns in the network, we greyed out the connections from each ROI to itself (i.e., the diagonal cell in the adjacency matrix). We started the analysis by comparing the average brain network of 42 AD patients with those of 50 controls. The initial view was rather cluttered because of the high number of edges and complexity of the average brain networks (Figure 5). Therefore, we focused on the network differences by manipulating the edge color map slider. First, we applied filters to eliminate many low-density connections that overwhelm the network view. Then we switched to the binary color mapping mode to only compare the unweighted connectivity controlled by a predefined density threshold. The result was illustrated in Figure 8(a). It can be quickly identified that the high-density fiber connectivity network is more symmetrical (between left and right hemisphere) in the control group than in the AD group. The AD patients have relatively sparse connectivity in the left hemisphere (the right part of the AD network in Figure 8(a) due to the projection method). It appears that the connectivity of the left parietal lobe breaks down the most. We further confirmed this hypothesis by using the network contrast interaction. As



Fig. 8: AD case: (a) The brain network asymmetry is shown; (b) When the contrast interaction is used, detailed lobe differences are revealed.



(b) Right sagittal view

Fig. 9: Sagittal views in the AD case: (a) Left; (b) Right.



Fig. 10: CCI case using the contrast interaction.

shown in Figure 8(b), most of the large differences between the two population groups were located in the intra- or inter-parietal lobe connectivity subnetworks. A matrix reorder operation was also applied to place the ROIs with the largest differences in the center of each lobe matrix. This became clearer when we switch to the left sagittal view (Figure 9(a)). Another observations was that the AD patients appeared to have thicker connectivity pattern between the right insula (ROI #36) and the right parietal/frontal lobes, though this finding needs to be validated from a statistical standpoint (Figure 8(b), Figure 9(b)).

In the second case study, we deployed our visualization on the 113subject brain networks from the Open Connectome project [1]. The raw brain networks were also defined using 70 ROIs in FreeSurfer based on the Desikan-Killiany atlas. Due to the higher resolution DWIs and possibly, difference in the processing methods, both the scale and distribution of the connectivity patterns appeared different from those in the ADNI data set (which is expected). In this case, we compared the subject's brain network using their Composite Creativity Index (CCI), which assesses the creativity of subjects in ten different domains (e.g., visual arts, music, etc.) through a questionnaire. The CCI score is normally distributed in population and we classified the scores into two classes for the comparison tasks: the high class with CCI higher than or equal to 100 (the average across populations) and the low class smaller than 100. Our goal was to validate the results from the seminal work by Flaherty [16] who showed that the frontal and temporal lobes are important areas involved in creative thinking and expression. Again, we started the analysis by displaying the brain networks of the CCI high vs. low population groups. The comparison view is further optimized by applying the network contrast and matrix reordering operation. As illustrated in Figure 10, it is visible that the frontal and temporal lobe connectivity patterns are stronger in the high-CCI group than in the low-CCI group.

In recent work [27] [13], scientists found linkage between the prefrontal lobe and the creative mind and problem solving tendency. We reproduced this discovery by incorporating the optimal ROI clustering into the visualization design. As shown in Figure 1, the sub-cluster corresponding to the prefrontal lobe in the frontal matrix reveals much stronger connectivity vs. the other lobes in the contrast view.

7 CONCLUSION AND FUTURE WORK

In this paper, we have presented an integrated visual analytics method for comparing blockwise brain network differences among population groups, a highly relevant task for clinical research aimed at visualizing the network of human brains. Existing visual representations such as node-link graph and adjacency matrix suffer from low accuracy when comparing real-life brain networks that are homogeneous and densely interconnected. We tackle this problem by proposing: 1) a greedy ROI clustering algorithm driven by the predictive classification performance; 2) an improved NodeTrix design for displaying the blockwise brain network connectivity patterns; and 3) a suite of customized interactions to accelerate the visual comparison between groups of interest. Visual clutter reductions and the linkage to the underlying fiber tract data are achieved by two edge bundling algorithms considering the separate semantic and geometric information. We demonstrate the effectiveness of our method in blockwise brain network visual comparison by both controlled user experiment in contrast to state-of-the-art visualization methods, and case studies on real-world brain networks.

Despite the success of our method in visually comparing brain networks within the studied scope, there are several unsolved challenges due to the enormous complexity and uncertainty of real-world human brain networks. First, the current NodeTrix design supports up to two levels of network hierarchy, while the high-resolution human connectome can carry many more hierarchies, ranging from block and ROI level to voxel and neuronal level. One aspect of our future work is to explore new visualization and interaction designs to address this practical challenge. Second, in this work we only consider the pairwise side-by-side comparison of brain networks. In fact, many diagnostic classes are not bilateral, and comparisons are conducted among more than two groups. We plan to study how the proposed method can be extended to work with a multiple-comparison scenario. Finally, our method can be extended to support the visual comparison of many other geospatial networks (e.g., dynamic traffic and migration networks), on which there are inherent blockwise connectivity patterns and the network node positions are geospatially fixed.

ACKNOWLEDGMENTS

This work is supported by China National 973 project 2014CB340301 and Natural Science Foundation of China No. 61379088. Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (NIH Grant U01 AG024904) and DOD ADNI (award number W81XWH-12-2-0012).

REFERENCES

- [1] Open connectome project. http://www.openconnectomeproject.org/.
- [2] The Alzheimer's Disease Neuroimaging Initiative (ADNI). http://adni.loni.usc.edu/.
- [3] FreeSurfer. http://surfer.nmr.mgh.harvard.edu/, 2012.
- [4] A. K. Al-Awami, J. Beyer, H. Strobelt, N. Kasthuri, J. W. Lichtman, H. Pfister, and M. Hadwiger. Neurolines: a subway map metaphor for visualizing nanoscale neuronal connectivity. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2369–2378, 2014.
- [5] B. Alper, B. Bach, N. Henry Riche, T. Isenberg, and J.-D. Fekete. Weighted graph comparison techniques for brain connectivity analysis. In *CHI*, pages 483–492, 2013.
- [6] B. Bach, E. Pietriga, and J.-D. Fekete. Visualizing dynamic networks with matrix cubes. In CHI, pages 877–886, 2014.
- [7] E. T. Bell. Exponential numbers. *The American Mathematical Monthly*, 41(7):411–419, 1934.
- [8] A. Berres, M. Goldau, M. Tittgemeyer, G. Scheuermann, and H. Hagen. Tractography in context: Multimodal visualization of probabilistic tractograms in anatomical context. In VCBM, pages 9–16, 2012.
- [9] R. L. Buckner, A. Z. Snyder, B. J. Shannon, G. LaRossa, R. Sachs, A. F. Fotenos, Y. I. Sheline, W. E. Klunk, C. A. Mathis, J. C. Morris, et al. Molecular, structural, and functional characterization of alzheimer's disease: evidence for a relationship between default activity, amyloid, and memory. *The Journal of Neuroscience*, 25(34):7709–7717, 2005.
- [10] M.-C. Chiang, M. Barysheva, D. W. Shattuck, A. D. Lee, S. K. Madsen, C. Avedissian, A. D. Klunder, A. W. Toga, K. L. McMahon, G. I. De Zubicaray, et al. Genetics of brain fiber architecture and intellectual performance. *The Journal of Neuroscience*, 29(7):2212–2224, 2009.
- [11] M. Daianu, N. Jahanshad, T. M. Nir, C. R. Jack, M. W. Weiner, M. A. Bernstein, and P. M. Thompson. Rich club analysis in the alzheimer's disease connectome reveals a relatively undisturbed structural core network. *Human brain mapping*, 36(8):3087–3103, 2015.
- [12] M. Daianu, N. Jahanshad, T. M. Nir, A. W. Toga, C. R. Jack Jr, M. W. Weiner, and P. M. Thompson, for the Alzheimer's Disease Neuroimaging Initiative. Breakdown of brain connectivity between normal aging and alzheimer's disease: a structural k-core network analysis. *Brain connectivity*, 3(4):407–422, 2013.
- [13] L. C. de Souza, H. C. Guimarães, A. L. Teixeira, P. Caramelli, R. Levy, B. Dubois, and E. Volle. Frontal lobe neurology and the creative mind. *Frontiers in psychology*, 5, 2014.
- [14] R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, et al. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage*, 31(3):968–980, 2006.
- [15] M. H. Everts, E. Begue, H. Bekker, J. B. Roerdink, and T. Isenberg. Exploration of the brains white matter structure through visual abstraction and multi-scale local fiber tract contraction. *IEEE Transactions on Visualization and Computer Graphics*, 21(7):808–821, 2015.
- [16] A. W. Flaherty. Frontotemporal and dopaminergic control of idea generation and creative drive. *Journal of Comparative Neurology*, 493(1):147– 153, 2005.
- [17] E. R. Gansner, Y. Koren, and S. North. Graph drawing by stress majorization. In GD, pages 239–250, 2004.
- [18] N. Geschwind. Disconnexion syndromes in animals and man. *Brain*, 88(3):585–585, 1965.
- [19] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, 2011.
- [20] N. Henry, J.-D. Fekete, and M. J. McGuffin. Nodetrix: a hybrid visualization of social networks. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1302–1309, 2007.
- [21] M. Hlawitschka and G. Scheuermann. Hot-lines: Tracking lines in higher order tensor fields. In *IEEE Visualization*, pages 27–34, 2005.
- [22] B.-C. Ho, N. C. Andreasen, P. Nopoulos, S. Arndt, V. Magnotta, and M. Flaum. Progressive structural brain abnormalities and their relationship to clinical outcome: a longitudinal magnetic resonance imaging study early in schizophrenia. *Archives of general Psychiatry*, 60(6):585– 594, 2003.
- [23] D. Holten and J. J. Van Wijk. Force-directed edge bundling for graph visualization. In *Computer Graphics Forum*, volume 28, pages 983–990, 2009.

- [24] B. Horwitz, C. L. Grady, N. Schlageter, R. Duara, and S. Rapoport. Intercorrelations of regional cerebral glucose metabolic rates in alzheimer's disease. *Brain research*, 407(2):294–306, 1987.
- [25] Y. Jin and H. E. Cetingül. Tractography-embedded white matter stream clustering. In *ISBI*, pages 432–435, 2015.
- [26] Y. Jin, Y. Shi, L. Zhan, B. A. Gutman, G. I. de Zubicaray, K. L. McMahon, M. J. Wright, A. W. Toga, and P. M. Thompson. Automatic clustering of white matter fibers in brain diffusion mri with an application to genetics. *NeuroImage*, 100:75–90, 2014.
- [27] S. W. Kleibeuker, P. C. M. Koolschijn, D. D. Jolles, M. A. Schel, C. K. De Dreu, and E. A. Crone. Prefrontal cortex involvement in creative problem solving in middle adolescence and adulthood. *Developmental cognitive neuroscience*, 5:197–206, 2013.
- [28] M. Kubicki, H. Park, C. Westin, P. Nestor, R. Mulkern, S. Maier, M. Niznikiewicz, E. Connor, J. Levitt, M. Frumin, et al. DTI and MTR abnormalities in schizophrenia: analysis of white matter integrity. *Neuroimage*, 26(4):1109–1118, 2005.
- [29] B. Liu, M. Song, J. Li, Y. Liu, K. Li, C. Yu, and T. Jiang. Prefrontalrelated functional connectivities within the default network are modulated by COMT val158met in healthy young adults. *The Journal of Neuroscience*, 30(1):64–69, 2010.
- [30] S. Liu, W. Cui, Y. Wu, and M. Liu. A survey on information visualization: recent advances and challenges. *The Visual Computer*, 30(12):1373– 1393, 2014.
- [31] D. S. Margulies, J. Böttger, A. Watanabe, and K. J. Gorgolewski. Visualizing the human connectome. *NeuroImage*, 80:445–461, 2013.
- [32] J. McGonigle, A. L. Malizia, and M. Mirmehdi. Visualizing functional connectivity in fMRI using hierarchical edge bundles. In *OHBM*, 2011.
- [33] B. Olabi, I. Ellison-Wright, A. M. McIntosh, S. J. Wood, E. Bullmore, and S. M. Lawrie. Are there progressive brain changes in schizophrenia? a meta-analysis of structural magnetic resonance imaging studies. *Biological psychiatry*, 70(1):88–96, 2011.
- [34] V. Prčkovska, T. H. J. M. Peeters, M. van Almsick, B. ter Haar Romeny, and A. V. i Bartroli. Fused DTI/HARDI visualization. *IEEE Transactions* on Visualization and Computer Graphics, 17(10):1407–1419, 2011.
- [35] A. A. Scott-Van Zeeland, B. S. Abrahams, A. I. Alvarez-Retuerto, L. I. Sonnenblick, J. D. Rudie, D. Ghahremani, J. A. Mumford, R. A. Poldrack, M. Dapretto, D. H. Geschwind, et al. Altered functional connectivity in frontal lobe circuits is associated with variation in the autism risk gene CNTNAP2. *Science translational medicine*, 2(56):56ra80, 2010.
- [36] N. Simon, J. Friedman, T. Hastie, and R. Tibshirani. A sparse-group lasso. *Journal of Computational and Graphical Statistics*, 22(2):231– 245, 2013.
- [37] O. Sporns. The human connectome: a complex network. Annals of the New York Academy of Sciences, 1224(1):109–125, 2011.
- [38] J. Tao, C. Wang, C.-K. Shene, and R. A. Shaw. A vocabulary approach to partial streamline matching and exploratory flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(5):1503–1516, 2016.
- [39] P. M. Thompson, T. Ge, D. C. Glahn, N. Jahanshad, and T. E. Nichols. Genetics of the connectome. *Neuroimage*, 80:475–488, 2013.
- [40] P. M. Thompson, M. S. Mega, R. P. Woods, C. I. Zoumalan, C. J. Lindshield, R. E. Blanton, J. Moussai, C. J. Holmes, J. L. Cummings, and A. W. Toga. Cortical change in alzheimer's disease detected with a disease-specific population-based brain atlas. *Cerebral Cortex*, 11(1):1– 16, 2001.
- [41] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal* of the Royal Statistical Society. Series B, pages 267–288, 1996.
- [42] C. Wernicke. Der aphasische Symptomencomplex: eine psychologische Studie auf anatomischer Basis. Springer-Verlag, 1874.