Intrinsic Geometry Visualization for the Interactive Analysis of Brain Connectivity Patterns

Giorgio Conte, Allen Ye, Kyle Almryde, Olusola Ajilore, Alex Leow, and Angus Graeme Forbes

Abstract

Understanding how brain regions are interconnected is an important topic within the domain of neuroimaging. Advances in non-invasive technologies enable larger and more detailed images to be collected more quickly than ever before. These data contribute to create what is usually referred to as a connectome, that is, a comprehensive map of neural connections. The availability of connectome data allows for more interesting questions to be asked and more complex analyses to be conducted. In this paper we present a novel web-based 3D visual analytics tool that allows user to interactively explore the intrinsic geometry of the connectome. That is, brain data that has been transformed through a dimensionality reduction step, such as multidimensional scaling (MDS), Isomap, or t-distributed stochastic neighbor embedding (t-SNE) techniques. We evaluate our tool through a series of real-world case studies, demonstrating its effectiveness in aiding domain experts for a range of neuroimaging tasks.

Introduction

Providing a deeper understanding of the interconnectedness of the human brain is a primary focus in the neuroimaging community. Imaging techniques, such as *functional Magnetic Resonance Imaging* (fMRI), *diffusion tensor imaging* (DTI) and high angular resolution diffusion imaging (HARDI), enable neuroimagers to collect and derive data about how different brain regions connect from both a structural and a functional point of view [20]. Analogous to the *genome* for genetic data, the *connectome* is a map of neural connections [30].

Complex functional and structural interactions between different regions of the brain have necessitated the development and growth of the field of connectomics. The brain connectome at the macro-scale is typically mathematically represented using connectivity matrices that describe the interaction among different brain regions. Most current connectome study designs use brain connectivity matrices to compute summarizing statistics of either a global or a nodal level [31].

In the current work, we introduce the potential utility of deriving and analyzing the *intrinsic geometry* of brain data, that is, the topological space defined using derived connectomic metrics rather than anatomical features. The utility of this intrinsic geometry could lead to a greater distinction of differences not only in clinical cohorts, but possibly in the future to monitor longitudinal changes in individual brains in order to better deliver individualized precision medicine.

Specifically, we introduce a prototype application that provides researchers with the ability to perform visual analytics tasks related to the exploration of the intrinsic geometry of a dataset and the comparison of how the dataset looks when embedded within different topological spaces. The *intrinsic geometry* represents the brain connectome after non-linear multidimensional data reduction techniques are applied. Dimensionality reduction techniques remap the brain according to network properties. This means that the position of each node does not correspond to its anatomical location, as it does in the original brain geometry. Instead, its position is based on the strength of the interaction that each region has with the others, whether structural or functional. The stronger the connectivity between two regions, the closer they are in the intrinsic geometry. In the intrinsic geometry we are more interested in the shape the brain connectome assumes independent of the anatomical distances between nodes. Thus, the space in which the intrinsic geometry is plotted in is called *topological space* [5].

Through using a variety of dimensionality reduction techniques such as Isomap [32] and t-distributed stochastic neighbor embedding (t-SNE) [33], a brain's connectivity matrix can be directly embedded into topographical spaces. Linear dimensionality reduction techniques such as multidimensional scaling (MDS) [3] and principal component analysis (PCA) [19] have been previously used in unrelated fields of medicine as a way to distinguish clinical cohorts through biomarkers, although it can be argued that they are not suitable for complex high-dimensional connectome data [16, 35]. To our knowledge our approach represents the first comprehensive application of dimensionality reduction techniques in the ever-expanding field of brain connectomics.

This intrinsic geometry concept provides an underlying connectomic visualization that is not obscured by the standard anatomical structure. That is, visualizing connectivity information within an anatomical representation of the brain can potentially limit one's ability to clearly understand the complexity of a human brain connectome; some meaningful structural patterns may be much easier to see in topological space. The use of intrinsic geometry relies on the intuition that the brain's intrinsic geometry should reflect graph properties of the corresponding brain connectivity matrix, rather than the inter-regional Euclidean distances in the brain's physical space [7]. Fig. 1 shows the shape of the intrinsic geometry for example connectomes using Isomap, MDS, and t-SNE reduction techniques.

Related Work

Many approaches to visualizing the connectome have been presented and, broadly speaking, three main types have emerged: node-link diagrams, matrix representations, and circular layout [25, 26]. The *Connectome Visualization Utility* [22], the *Brain Net Viewer* [37], and the *Connectome Viewer Toolkit* [14] each provide a 3D node-link representation. In these tools, the dimension of the nodes is bound to a graph-based metrics, like nodal strength or nodal degree, while the weight of the edges is displayed using different colors or by changing diameter of the cylin-

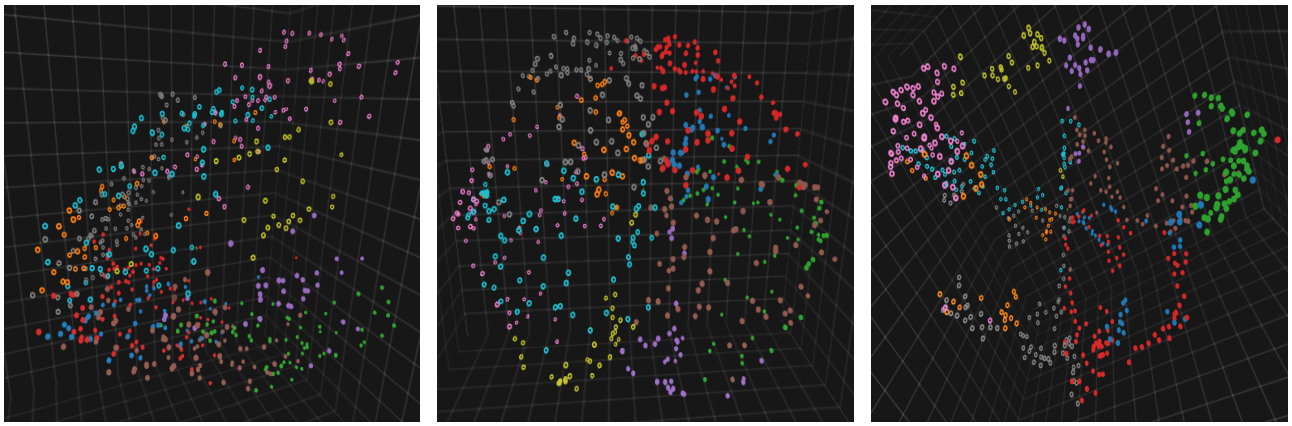

(a) MDS Space (b) Isomap Space (c) tSNE Space

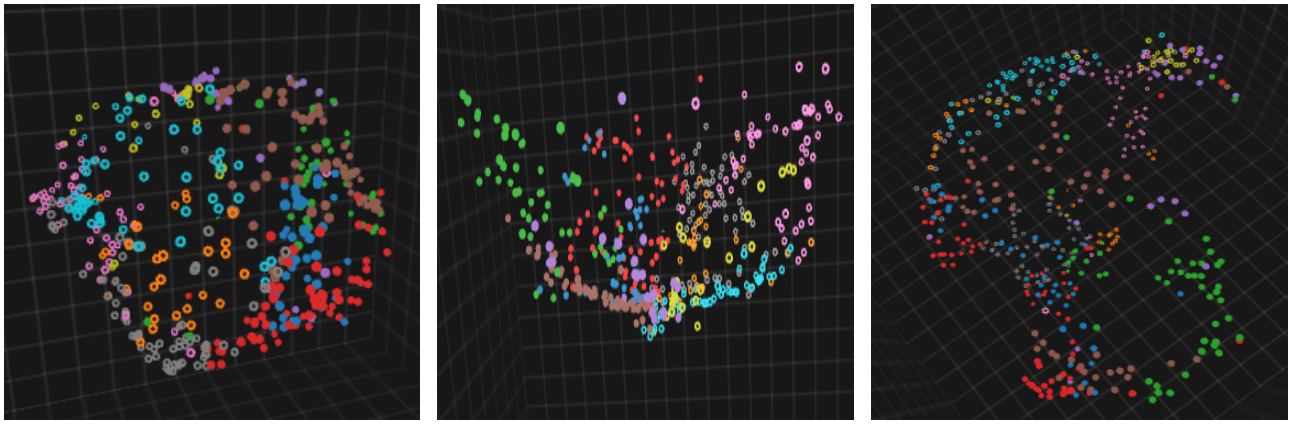

(d) MDS Space (e) Isomap Space (f) tSNE Space

Figure 1. *In these figures we present the different shapes the structural connectome assumes in topological space when different dimensionality reduction techniques are applied. With any reduction of data from a higher dimension to a lower one some information is lost, and so our tool enables users to investigate intrinsic geometries resulting from different techniques in order to explore the different topological spaces that may shed light on a particular connectome dataset. The screenshots are taken from different points of views and the colors represent different lobes of the brain.*

drical link. Other functional brain connectivity visualization tools instead utilize 2D node-link diagrams, such as a technique introduced by Jianu et al. [18] that explores brain connectivity using 2D neural maps and a frequency-dependent approach by Salvador et al. [29] that displays functional connectivity via an undirected graph layout. The main advantage of using node-link diagrams is that they provide an overview of the entire graph that makes it easy to understand which nodes are indirectly connected. The 3D repreeentations additionally provide meaningful spatial information, and the tools that utilize 3D node-link diagrams locate the nodes relative to the real anatomical position. However, excessive visual clutter is introduced when the number of nodes and edges increase, affecting the readability of the graph. Ma et al. [23] explore a dual representation of a dataset, providing both an adjacency matrix and an anatomically grounded node-link diagram in order to provide distinct, yet interrelated views of the same data. This approach also includes a temporal dimension and uses glyphs to indicate community membership of nodal regions, but does not extend to 3D data.

Both the *Connectome Visualization Utility* [22] and the *Connectome Viewer Toolkit* [14] cam make use of adjacency matrices to represent connections; additionally, the former also includes a circle layout view. This view, also known as a *connectogram* was first introduced by Irimia et al. [17]. A connectogram displays brain regions around a circle, and the interconnections between them are represented as edges that connect regions together inside the circle. Transparency is used to represent the weight of the edges, which reduces the visual clutter and highlights only on strong links, while weak edges fade into the background. By using connectograms it is possible to incorporate additional information adding more than one nested circles; the outermost circle can be used to represent the cortical parcelations, while the inner circles could use heat maps to display different structural measures associated with the corresponding regions. This approach mitigates the clutter that can occur in other approaches when the number of edges and nodes increase.

Currently, there are no effective tools that enable a user to interactively investigate the intrinsic geometry of connectome data, and none that allow a user to apply and visualize complex transformations to connectome datasets. To address these issues, we incorporate an interactive 3D node-link diagram to visualize connectome data, described below. Moreover, we support viewing

via virtual reality displays, as it has been shown that stereoscopic 3D visualization can make it easier for user to the comprehension of the anatomical location of regions as well as the functional connections [27]. Nonetheless, the visualization of 3D networks can outperform 2D static ones, especially when considering complex tasks [1]; for example, initial investigations by Forbes et al. demonstrate an effective a stereoscopic system to visualize temporal data of the brain activity [13].

Tasks

Due to the high complexity of the human connectome, simply being able to *explore* the data to support sensemaking is a fundamental task. However, a researcher typically has a well defined assumption about the contents of their data and a clear idea about which aspects of the data he or she wants to explore. Effective exploration involves being able to find relevant information quickly by filtering the data in order to identify patterns, to assist in the generation of new hypotheses, or to confirm or invalidate expected results.

Researchers often need to examine multiple datasets in order to *compare* the structure or activity of one region of a brain with another, or to compare different populations or experimental conditions. For example, a psychiatric researcher could be interested in understanding the differences between the functional connectivity of healthy controls versus depressed participants. Individuals with depression show a higher functional connectivity between the regions within the default mode network than in healthy controls [15]. Being able to visually distinguish details about the different activity levels within specific brain regions is necessary to support a deeper understanding of these pathologies, as well as to verify experimental results and enable the generation of new hypotheses.

Neuroimagers may also need to *identify* the importance of regions that can be directly or indirectly affected by damage to the brain, such as in the case of traumatic brain injury [21]. It is also important to understand the structural and functional implications of neurosurgical interventions such as temporal lobotomy [2], or as a predictive measure towards behavioral therapy outcomes for use in aphasia treatment [24, 34]. Lastly, experts are interested in identifying culpable regions when investigating neurodegenerative disease and neuro-psychological disorders such as Alzheimer's [8, 39] and schizophrenia [6]. Including both neuroanatomical and topological representations allows researchers to more comprehensively address these complex issues.

The Visualization Application

The primary layout for our prototype application is a 3D node-link diagram, motivated by interest researchers have in understanding the brain's *intrinsic geometry*. The position of each region in the *topological space* is highly relevant in this context. Although many visualization researchers have noted some potential pitfalls in making use of 3D representations for visual analysis tasks, the importance of being able to compare the anatomical geometry with the different intrinsic geometries necessitates this layout. When viewing the intrinsic geometry, the individual nodes represent different brain regions and are represented with circular glyphs, while edges representing a functional or a structural connection between these regions are displayed using lines. Fig. 2 shows an example view in our application.

A main concern with the use of node-link diagrams is the potential for visual clutter when displaying a highly interconnected graph, such as the human brain connectome. Instead of showing all the connections simultaneously, by default our tool only shows nodes, hiding all links unless explicitly required. Through interaction, users are able to display or hide connections according to their preferences and current needs. We also allow the user to choose to view the connections only within a particular subgraph that is relevant for a particular task. This edges-on-demand technique allows exploration tasks to be performed by showing only the connections starting from a specific region that is currently being interrogated. The user can pin the connections in the scene just by clicking on the node itself. We use varying degrees of transparency to visually encode the strength of edge weights. Stronger connections are then represented using opaque lines, while weaker edges are more transparent. Transparency is scaled relative to only the currently displayed edges.

Information about which hemisphere particular nodes belong to can be meaningful for certain tasks. Being able to understand quickly whether or not global right/left symmetric patterns are still recognizable in the intrinsic geometry also helps the domain experts by providing an anatomical reference during the exploration of the intrinsic geometry. We represent nodes from hemispheres using two different glyphs, circular and toroidal.

Colors are used to highlight the neuroanatomical membership of each node in the brain. In our application, each glyph belongs to one of the 82 neuroanatomical regions as defined by Freesurfer [11]. However, the data structure is flexible enough to accept any membership or affiliation structure. Currently, these affiliations are hard-coded by default, but our application has the ability to compute affiliations on the fly according to specific graph metrics.

Analytics Features

Our application facilitates a range of user interactions to support visual analysis, including the ability to:

- create the shortest path tree rooted in the node selected by the user;
- visualize the shortest path between the two nodes;
- let the user turn on and off particular regions;
- let the user quickly switch between different geometries;
- compute the nodal strength for each node of the graph.

We use Dijkstra's algorithm [9] to create the shortest path tree. In structural connectomes, since the adjacency matrix defines the number of reconstructed white matter tracts connecting two regions, the edge length is set to the inverse of the fiber count (the higher the number of tracts, the more coupled two nodes are and thus the shorter their distance is). From a mathematical point of view:

$$d(i,j) = \frac{1}{w_{ij}} \tag{1}$$

where $d(i,y)$ is the distance between node i and node j and w_{ij} is the weight of the edge which links i and j contained in the adjacency matrix. The user can filter the shortest path tree according to two different measures: graph distance and number of intermediate nodes, or "hops." In the first case the user can filter the

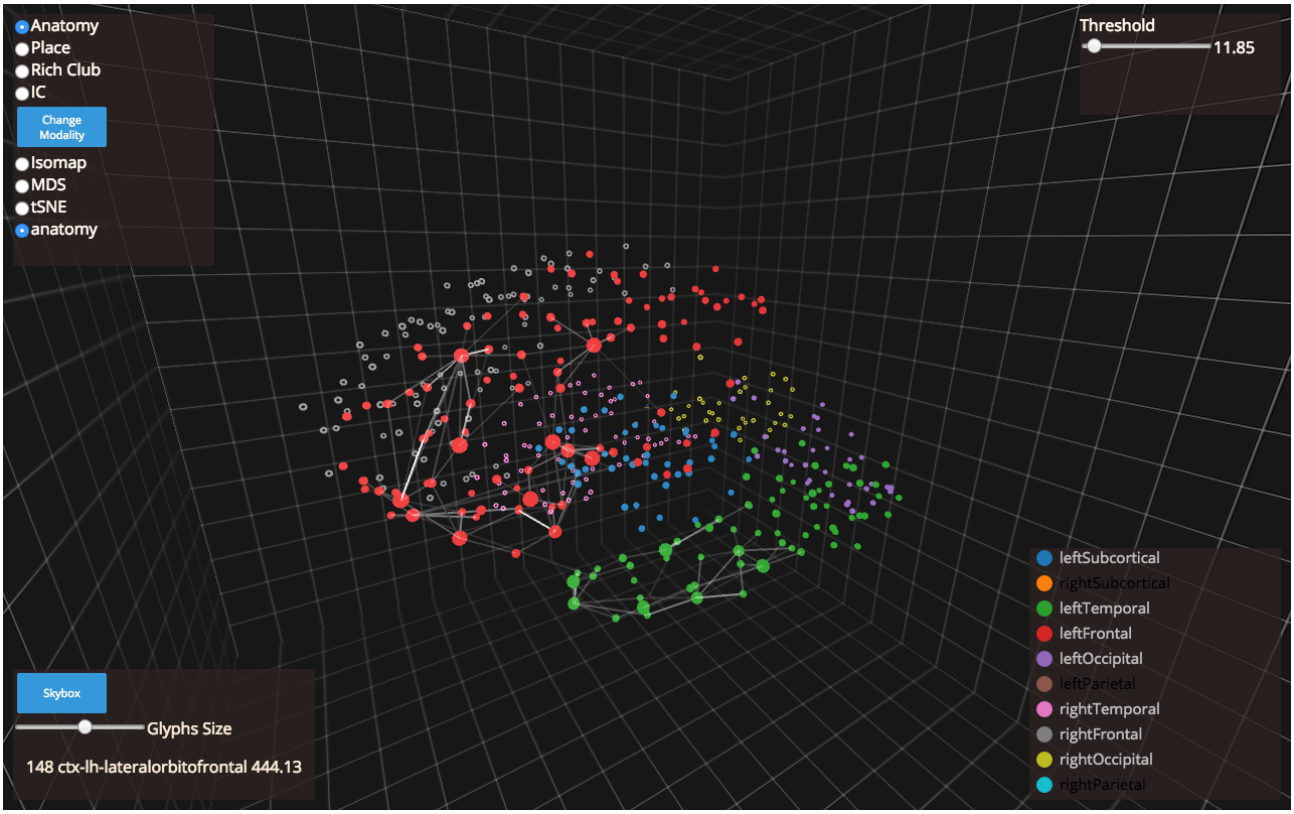

Figure 2. *This figure shows the main view of our tool. On the upper left a simple menu allows users to choose the way nodes are grouped (color encodings) and the topological space in which we visualize connectome data. On the upper right a slider sets the minimum edge weights for an edge to be visible. On the bottom left the name of the node selected and its nodal strength are showed. We also allow users to change the size of the glyphs and to hide/show the 3D grid in the background. On the lower right, a figure legend mapping each color to its neuroanatomical label is shown. The connectome visualization is displayed at the center of the scene.*

tree according to the relative distance with respect to its farthest node. Given a threshold t, all the nodes that satisfy the following inequality are drawn:

$$\begin{cases} d(r,i) \leq maxDistance(r) \cdot t \\ 0 \leq t \leq 1 \end{cases} \tag{2}$$

where r is the root node, i is the node considered, $maxDistance(r)$ is the distance between the root node and the farthest node, and t is the threshold chosen by the user. If $t = 0$ then only the root node is displayed, while if $t = 1$ the entire shortest path tree is drawn. In the latter case, the user is able to filter out nodes that are not reachable within a certain number of nodes from the root.

We enable the user to interactively select two specific nodes in order to show the **shortest route** between them. We can also display all the nodes in the network to provide the overall context of this sub-graph.

Being able to **select regions** is also important. The number of nodes displayed could affect the visual clutter of the display. We let the user choose whether to display or not groups of nodes depending on their affiliations. Thus, neuroimagers can explore only the regions that are strictly relevant to their research goals.

A main feature of our tool is the ability to **switch between geometries**. We provide a menu to select the space they want to explore. Switching geometries can be done with just one click, allowing the users to see how the connectome data appears embedded in different topological spaces.

Nodal strength is a graph-based metric which defines the *centrality* of a node. The nodal strength is defined as follows:

$$Nodal\ strength_i = \sum_{j=0}^{N} w_{ij} \tag{3}$$

where N is the number of nodes in the graph and w_{ij} is the weight of the link between i and j [28]. This graph-based metric helps experts to understand the relevance of a node in the network. This value is presented in a numerical label that appears when a node is interacted with; the user can filter out nodes below a particular threshold according to this measure.

Case Studies

In the following section, we present real-world case studies that use our visualization tool to help understand patterns in both the structural and functional connectome.

Case Study 1: Understanding Rich Club Nodes

We wanted to understand how the connectivity of the brain changes when specific regions of the brain are removed. In particular, we wanted the see the differences between a complete

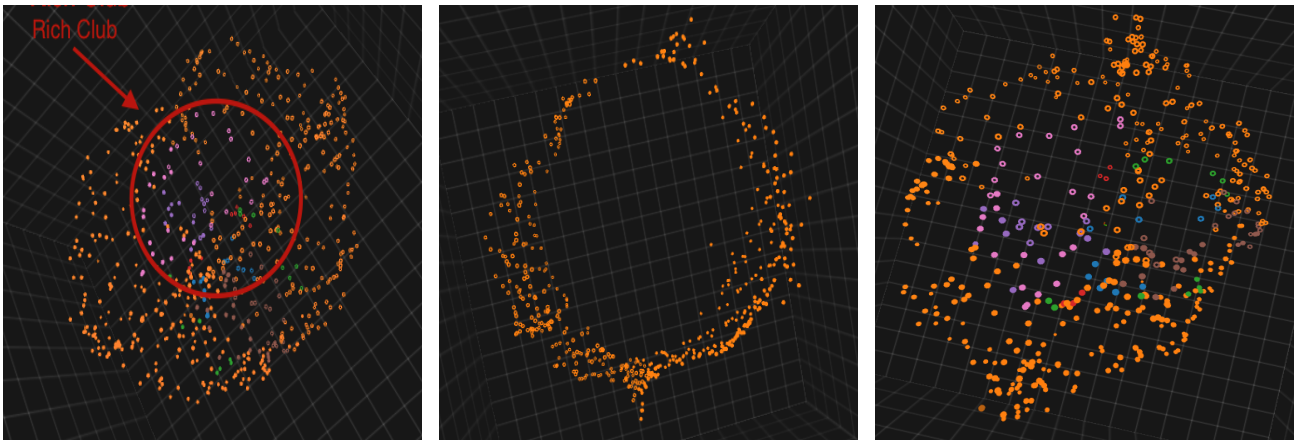

(a) Complete Structural Connectome (b) Structural Connectome Without Rich Club (c) Random Regions Removed

Figure 3. *This figure compares the complete structural connectome, Figure (a), and the structural connectome when nodes with the colored rich club property are removed, Figure (b). By comparing (a) and (b), it is very clear that without the rich club nodes, the intrinsic geometry of the brain becomes diffuse and nodes are less coupled to each other. Rich club regions form the core of the brain's structural connectome. These results are put in context when we consider Figure (c). Figure (c) shows a connectome after an equivalent number of nodes to rich club nodes were randomly selected and removed. It is clear there are subtle differences between the (a) and (c) but no gross changes to the structure as with targeted rich club removal (b). Put together, these simulated region removal analyses confirm the importance of rich club nodes.*

structural connectome and a connectome in which nodes belonging to the *rich club* were removed. The basic concept behind the rich club property is the tendency for nodes with high nodal strengths to form tightly interconnected groups [38]. Mathematically speaking, given a graph N and the parameter k which defines a nodal strength cut off, the rich club property is defined as

$$\phi(k) = \frac{2E_{>k}}{N_{>k}(N_{>k}-1)} \tag{4}$$

where $E_{>k}$ is the number of edges in N between the nodes of nodal strength greater or equal to k and $N_{>k}$ is the number of nodes in N with nodal strength greater or equal then k. This metric could also be seen as follows:

$$\phi(k) = \frac{E_{>k}}{\binom{N_{>k}}{2}} \tag{5}$$

Given that, $\phi(k)$ is the number of realized edges ($E_{>k}$) normalized with respect all the possible edges there could be between these nodes in a complete graph.

From Fig. 3 it is possible to see that the complete structural connectome forms a shape similar to a "bowl," while the connectome without rich club nodes shows a big "hole" in the middle. It is clear that those rich club nodes keep the entire network tightly interconnected. When they are missing, the remaining brain regions are more distant from each other, becoming less correlated and less coupled together.

Those differences gain particular relevance as we consider a different simulation. Instead of removing the nodes that were shown to have a particular characteristic (i.e. the rich club property), we also performed a random nodes removal using a uniform probability distribution and removing an equivalent number of nodes. As we can see from Figs. 3(a) and 3(c), the differences between the complete structural connectome and the one with random removal are not significant. Thus, this result validates the importance of rich club nodes.

Case Study 2: Comparing Depressed and Healthy Connectomes

Functional MRI (fMRI) has been widely used to study neural tasks, but a growing subset of fMRI is being dedicated to the *default mode network* (DMN) or how the brain responds when no external stimuli or task is given. The main feature of DMN function is strong interregional coordination of baseline oscillatory activities between its participant nodes. Core DMN brain regions are the ventromedial and dorsomedial prefrontal cortex, the posterior cingulate, the precuneus, the lateral temporal cortex and the hippocampus. Our results show that in depression (Fig. 4) there is strong coupling between the precuneus and hippocampus in the resting-state topological space, in line with extensive evidence supporting the involvement of these two regions in DMN functional organization [4, 36].

Using our application, clinicians explored the DMN in the averaged brain of 10 healthy controls and the averaged brain of 15 subjects with depression. They were able to gain insight the functional connectome by investigating differential patterns of interaction between the hippocampus, thalamus, and putamen. They found that these regions tended to be mixed together within the control participants, suggesting functional coupling. By contrast, in the depressed participants cohort, the hippocampus and the thalamus demonstrated a tend towards separating from the putamen, creating a separate cluster that is closer to the precuneus. This behavior is clearly visible in Fig. 4. Previous literature supports this idea of tight clustering of the hippocampus and thalamus regions in depressed patients, something not observed as often in healthy controls [15]. This could be a unique signature for major depression and offers some insight into the behavior of depressed subjects, especially their tendency to ruminate on their past and unsuccessful life events.

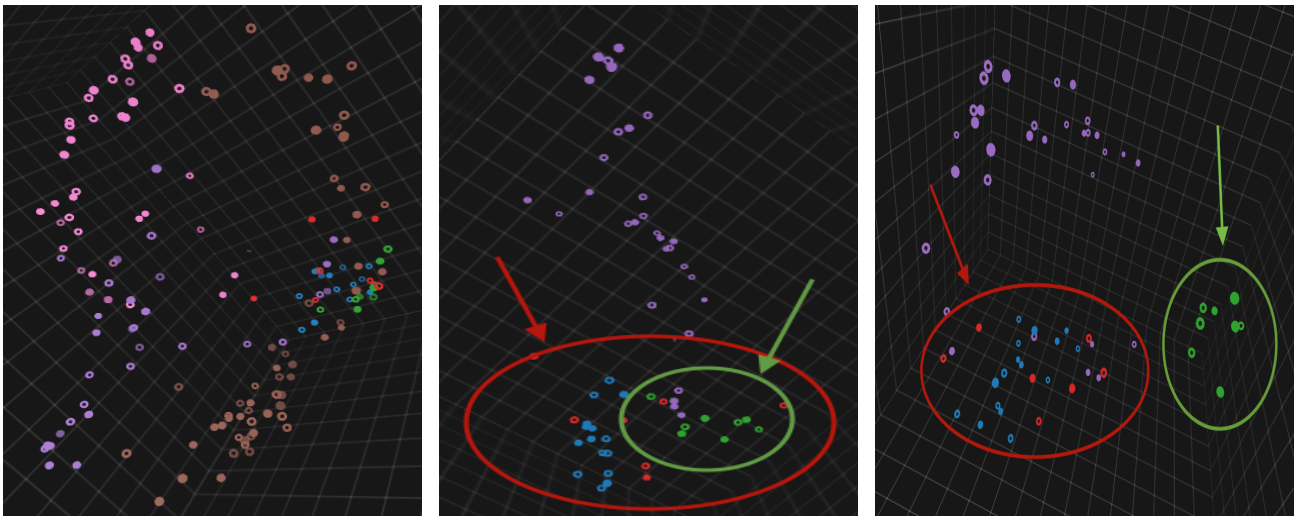

(a) Functional Connectome (b) Healthy Subjects (c) Depressed Subjects

Figure 4. *This figure compares the intrinsic topology of the functional connectome in healthy versus depressed subjects. In Figure (b) it is clear that the hippocampus and thalamus regions (blue and red nodes) tend to create a cluster apart from other regions. On the contrary, in healthy subjects the same regions are mixed together with putamen regions (green nodes).*

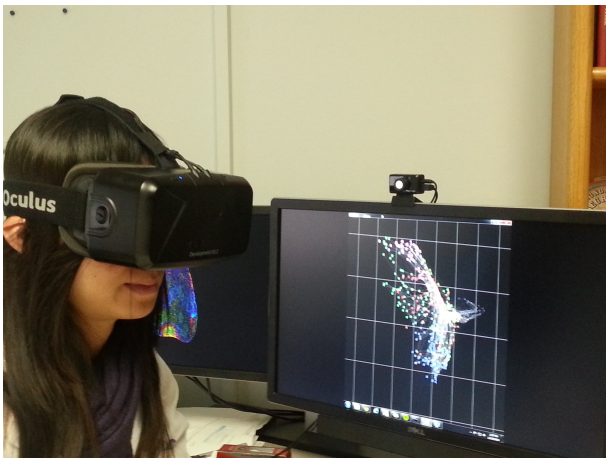

Figure 5. *A photo of a neuroimaging researcher exploring the intrinsic geometry of the brain with the our visualization application in immersive 3D using an Oculus Rift headset.*

Discussion

Since spatial vicinity equates to stronger connectivity in the intrinsic space, the user is able to freely and easily explore the terrain of either functional or structural brain connectivity. Indeed, the real advantage of exploring in the intrinsic space (especially when coupled with virtual-reality technology), is the ability to display the connectivity relationship among a number of brain regions, as neuroimagers unfold complex high-dimensional connectivity data into easily understandable and relatable configurations in 3D. This is evident in the resting state fMRI case study where the default mode network connectivity alterations in regions including precuneus and hippocampus can be easily appreciated. Experts also found that this tool transforms connectivity matrices in an engaging way that does not require much of a learning curve to understand and to use. By converting fiber count or functional connectivity into a distance measure, this visualization software creates a "road map" of the human brain. While the actual connectivity matrix can be parsed— much like knowing the distance to any stop of a road trip— it is hard to comprehend these strict numerical quantities without a map to help guide relative locations. Our visualization application allows for such an appreciation to occur and provides methods for interacting with individual nodes to discover highly integrated circuits in both functional and structural connectomes. Moreover, by inducing virtual lesions, one can compare the relative importance of certain brain regions and graph theoretical metrics by the subsequent changes in the topographical shape of the connectome.

Conclusion and Future Work

This paper introduced a novel visualization application that enables a user to interactively explore the human brain connectome. Being able to select edges on demand allows users to explore the entire connectivity network while limiting the visual clutter typical of highly connected node-link diagrams. Moreover, the analytics features provided by our tool make it easier to investigate the most relevant parts of the network according to the users current goals. Additionally, it enables users to view and compare the intrinsic geometry of connectome datasets in a number of different topological spaces in order to enable new understandings of the data.

Although our application has already enabled us to explore the intrinsic geometry of the brain in specific use cases, there are many useful features that we plan to introduce. For instance, it would be useful to allow users to add or remove nodes interactively and then apply dimensionality reduction techniques directly to only the currently visible nodes. We also plan to continue our investigation into the use of virtual reality systems. One current exploration involves combining the Oculus Rift and Leap Motion devices together to enable gesture interaction within an im-

mersive environment to enable a more engaging exploration with the brain's intrinsic geometry (see Fig. 5). We are also interested in exploring the use of motion as a means to augment identification of and reasoning about individual nodes and node clusters [10, 12]. Finally, we are excited about the potential utility of immersive connectome visualization as part of a biofeedback process that can provide users with the ability to see and control their connectome. Currently, we can use the DTI and fMRI data to track long term changes in the brain and follow macroscale neuroplastic changes within the brain. A future goal is to also map changes to the connectome using electroencephalography (EEG) and provide feedback in real-time to patients to promote changes in the brain through treatment.

References

[1] B. Alper, T. Hollerer, J. Kuchera-Morin, and A. G. Forbes. Stereoscopic highlighting: 2d graph visualization on stereo displays. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):2325–2333, 2011.

[2] L. Bonilha, J. A. Helpern, R. Sainju, T. Nesland, J. C. Edwards, S. S. Glazier, and A. Tabesh. Presurgical connectome and postsurgical seizure control in temporal lobe epilepsy. *Neurology*, 81(19):1704–1710, 2013.

[3] I. Borg and P. J. Groenen. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2005.

[4] S. J. Broyd, C. Demanuele, S. Debener, S. K. Helps, C. J. James, and E. J. Sonuga-Barke. Default-mode brain dysfunction in mental disorders: A systematic review. *Neuroscience & Biobehavioral Reviews*, 33(3):279–296, March 2009.

[5] E. T. Bullmore and D. S. Bassett. Brain graphs: graphical models of the human brain connectome. *Annual review of clinical psychology*, 7:113–140, 2011.

[6] V. D. Calhoun, J. Sui, K. Kiehl, J. Turner, E. Allen, and G. Pearlson. Exploring the psychosis functional connectome: aberrant intrinsic networks in schizophrenia and bipolar disorder. *Frontiers in psychiatry*, 2, 2011.

[7] G. Conte, A. Q. Ye, A. G. Forbes, O. Ajilore, and A. D. Leow. BRAINtrinsic: A virtual reality-compatible tool for exploring intrinsic topologies of the human brain connectome. In Y. Guo, K. Friston, A. Faisal, S. Hill, and H. Peng, editors, *Brain Informatics and Health*, volume 9250 of *Lecture Notes in Artificial Intelligence*, pages 67–76. Springer, 2015.

[8] Z. Dai, C. Yan, Z. Wang, J. Wang, M. Xia, K. Li, and Y. He. Discriminative analysis of early alzheimer's disease using multi-modal imaging and multi-level characterization with multi-classifier (m3). *Neuroimage*, 59(3):2187–2195, 2012.

[9] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.

[10] R. Etemadpour and A. G. Forbes. Density-based motion. *Information Visualization*, 2015. Published online before print: http://ivi.sagepub.com/content/early/2015/10/06/1473871615606187.

[11] B. Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.

[12] A. G. Forbes, C. Jette, and A. Predoehl. Analyzing intrinsic motion textures created from naturalistic video captures. In *Proceedings of the International Conference on Information Visualization Theory and Applications (IVAPP)*, pages 107–113, 2014.

[13] A. G. Forbes, J. Villegas, K. Almryde, and E. Plante. A stereoscopic system for viewing the temporal evolution of brain activity clusters in response to linguistic stimuli. In A. J. Woods, N. S. Holliman, and G. E. Favalora, editors, *Stereoscopic Displays and Applications XXV*, volume 9011 of *Proceedings of SPIE-IS&T Electronic Imaging*, pages 90110I–1–7. San Francisco, California, February 2014.

[14] S. Gerhard, A. Daducci, A. Lemkaddem, R. Meuli, J.-P. Thiran, and P. Hagmann. The Connectome Viewer Toolkit: An open source framework to manage, analyze and visualize connectomes. *Frontiers in Neuroinformatics*, 5(3), 2011.

[15] M. D. Greicius, B. H. Flores, V. Menon, G. H. Glover, H. B. Solvason, H. Kenna, A. L. Reiss, and A. F. Schatzberg. Resting-state functional connectivity in major depression: Abnormally increased contributions from subgenual cingulate cortex and thalamus. *Biological psychiatry*, 62(5):429–437, 2007.

[16] S. Howells, R. Maxwell, A. Peet, and J. Griffiths. An investigation of tumor 1H nuclear magnetic resonance spectra by the application of chemometric techniques. *Magnetic resonance in medicine*, 28(2):214–236, 1992.

[17] A. Irimia, M. C. Chambers, C. M. Torgerson, M. Filippou, D. A. Hovda, J. R. Alger, G. Gerig, A. W. Toga, P. M. Vespa, R. Kikinis, and J. D. Van Horn. Patient-tailored connectomics visualization for the assessment of white matter atrophy in traumatic brain injury. *Frontiers in neurology*, 3, 2012.

[18] R. Jianu, Ç. Demiralp, and D. H. Laidlaw. Exploring brain connectivity with two-dimensional neural maps. *IEEE Transactions on Visualization and Computer Graphics*, 18(6):978–987, 2012.

[19] I. Jolliffe. *Principal component analysis*. Wiley Online Library, 2002.

[20] D. K. Jones. *Diffusion MRI: Theory, methods, and applications*. Oxford University Press, 2010.

[21] J. Kim, D. Parker, J. Whyte, T. Hart, J. Pluta, M. Ingalhalikar, H. Coslett, and R. Verma. Disrupted structural connectome is associated with both psychometric and real-world neuropsychological impairment in diffuse traumatic brain injury. *Journal of the International Neuropsychological Society*, 20(9):887–896, 2014.

[22] R. A. LaPlante, L. Douw, W. Tang, and S. M. Stufflebeam. The Connectome Visualization Utility: Software for visualization of human brain networks. *PLoS ONE*, 9(12):e113838, 12 2014.

[23] C. Ma, R. V. Kenyon, A. G. Forbes, T. Berger-Wolf, B. J. Slater, and D. A. Llano. Visualizing dynamic brain networks using an animated dual-representation. In *Proceedings of the Eurographics Conference on Visualization (EuroVis)*, pages 73–77, Cagliari, Italy, May 2015.

[24] K. Marcotte, D. Adrover-Roig, B. Damien, M. de Praumont, S. Gnreux, M. Hubert, and A. I. Ansaldo. Therapy-induced neuroplasticity in chronic aphasia. *Neuropsychologia*, 50(8):1776 – 1786, 2012.

[25] D. S. Margulies, J. Böttger, A. Watanabe, and K. J. Gorgolewski. Visualizing the human connectome. *NeuroImage*, 80:445–461, 2013.

[26] H. Pfister, V. Kaynig, C. P. Botha, S. Bruckner, V. J. Dercksen, H.-C. Hege, and J. B. Roerdink. Visualization in connectomics. In *Scientific Visualization*, pages 221–245. Springer, 2014.

[27] G. M. Rojas, M. Gálvez, N. V. Potler, R. C. Craddock, D. S. Margulies, F. X. Castellanos, and M. P. Milham. Stereoscopic three-dimensional visualization applied to multimodal brain images: Clinical applications and a functional connectivity atlas. *Frontiers in neuroscience*, 8, 2014.

[28] M. Rubinov and O. Sporns. Complex network measures of brain connectivity: Uses and interpretations. *Neuroimage*, 52(3):1059–1069, 2010.

[29] R. Salvador, J. Suckling, C. Schwarzbauer, and E. Bullmore. Undirected graphs of frequency-dependent functional connectivity in whole brain networks. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):937–946, 2005.

[30] O. Sporns. The human connectome: A complex network. *Annals of the New York Academy of Sciences*, 1224(1):109–125, 2011.

[31] O. Sporns, G. Tononi, and R. Kötter. The human connectome: A structural description of the human brain. *PLoS computational biology*, 1(4):e42, 2005.

[32] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[33] L. Van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(2579-2605):85, 2008.

[34] S. van Hees, K. McMahon, A. Angwin, G. de Zubicaray, S. Read, and D. A. Copland. A functional mri study of the relationship between naming treatment outcomes and resting state functional connectivity in post-stroke aphasia. *Human Brain Mapping*, 35(8):3919–3931, 2014.

[35] S. Vujovic, S. Henderson, N. Presneau, E. Odell, T. Jacques, R. Tirabosco, C. Boshoff, and A. Flanagan. Brachyury, a crucial regulator of notochordal development, is a novel biomarker for chordomas. *The Journal of pathology*, 209(2):157–165, 2006.

[36] S. Whitfield-Gabrieli and J. M. Ford. Default mode network activity and connectivity in psychopathology. *Annual Review of Clinical Psychology*, 8(1):49–76, Apr. 2012.

[37] M. Xia, J. Wang, and Y. He. BrainNet Viewer: A network visualization tool for human brain connectomics. *PloS one*, 8(7):e68910, 2013.

[38] S. Zhou and R. J. Mondragón. The rich-club phenomenon in the internet topology. *Communications Letters, IEEE*, 8(3):180–182, 2004.

[39] D. Zhu, K. Li, D. P. Terry, A. N. Puente, L. Wang, D. Shen, L. S. Miller, and T. Liu. Connectome-scale assessments of structural and functional connectivity in MCI. *Human brain mapping*, 35(7):2911–2923, 2014.