

Immersive Analytics for Clinical Neuroscience

Johnson J.G. Keiriz* Olusola Ajilore† Alex D. Leow‡ Angus G. Forbes§
University of Illinois at Chicago, IL

ABSTRACT

This paper introduces *NeuroCave*, an open-source immersive analytics software tool for exploring and making sense of both structural and functional connectomes describing the interconnectivity between brain regions. The software supports visual analysis tasks for clinical neuroscientists, facilitating explorations of connectome datasets used in comparative group studies. Researchers can enter and leave virtual reality as desired in order to visualize and reason about connectome datasets from a range of different perspectives.

1 INTRODUCTION

Modern, noninvasive neuroimaging techniques provide a means with which to understand structural and functional brain networks, or connectomes [18]. Diffusion MRI derived white matter interconnectivity between different brain regions yields the structural connectome, and BOLD signal correlations generate the functional connectome. Mathematically, a connectome can be modeled as a graph by representing the different brain regions as nodes. Such models enable neuroscientists to apply network-theoretic methods and metrics, revealing important properties of the brain, such as small-worldness [1], clustering and modularity [15], and rich-club configuration [20], among others.

Through ongoing collaborations with The Collaborative Neuroimaging Environment for Connectomics (CoNECT Lab) at the University of Illinois Chicago, we identify four main tasks relevant for clinical neuroscientists:

- T1** Identify regions responsible for specific cognitive functions and study their interactions with other regions.
- T2** Compare individual networks to the mean or group average connectome, or compare differences between two group average connectomes. In group studies, individual variations as well as joint network characteristics are studied in order to identify commonalities or differences.
- T3** Identify the effect of structural connectivity on the functional activity of the brain. Comparing both structural and functional at the same time to reveal the complex mappings between them.
- T4** Identify individual or group changes occurring on the structural or functional connectivities due to the onset of disease or aging, and assess connectome restoration in drug studies.

To support these tasks, we explored a range of visual encodings and layout strategies, both in 2D and 3D, using an iterative design process in which we gathered feedback from expert users (clinical neuroscientists, psychiatry professors, medical students, and post-doctoral researchers) and made changes based on their feedback.

Our neuroscience collaborators wanted to take advantage of the greater sense of immersion that virtual reality (VR) provides, but became aware that in some cases it was easier to make comparative analyses within a traditional desktop environment. This was true even for 3D layouts, and our visualization software, *NeuroCave*, enables a workflow in which researchers can enter and leave VR environments as desired. In practice, users tended to use the VR environments to make initial explorations of the data and to generate hypotheses about the connectomes, and then switch to a desktop view once more nuanced investigations were required. However, we believe this is partly due to users not being as familiar with navigating in VR (and especially with making fine-grained selections), and also because of the need to use additional applications during an analysis (i.e., for web search, taking notes, etc.) that are not readily available when wearing a portable VR headset.

2 NEUROCAVE

NeuroCave enables a wide range of interactive methods to support these tasks, including: flexible data loading and data transformations, enabling comparisons within and between connectome datasets; user defined coloring scheme (based on lobar information, modular and/or community affiliation, etc); a coordinate system that can be defined interactively by users or automatically determined via a modular layout; adjustable glyph size and transparency of nodes and edges; adjustable connectivity threshold for displaying edges; shortest path between two nodes; on-demand edge bundling and edge coloring; and on-demand labeling of nodes and edges.

NeuroCave is inspired by a previous project, *BRAINtrinsic* [4, 5], which provides a VR environment in which to explore connectome datasets in different topological spaces after being transformed by dimensionality reduction algorithms [24]. *NeuroCave* provides richer interaction and a more responsive interface, supports a wider range of VR hardware, and, unlike other connectome analytics tools, is specifically designed to support comparison tasks for group studies. *NeuroCave* is implemented as a web-based application that runs on all major web browsers and makes use of *three.js*, a JavaScript graphics library, for real-time rendering of 3D scenes. The default view is formed of two side-by-side rendering views (see the left panel of Fig. 1 and both panels of Fig. 2). Each view enables the interactive visualization of a connectome as a node-link diagram.

2.1 Group Visualization

NeuroCave loads connectome data from a user-specified folder. This folder must contain all adjacency matrices as well as the corresponding topological and clustering information of the subjects within the study. An index file states the subject ID and its corresponding data files. Each study or analysis session requires a predefined Atlas that provides numerical labels and their corresponding anatomical names to each node. *NeuroCave* currently supports three Atlases by default: FSL-based parcellation, which consists of 82 labels from FreeSurfer [8]; the Brain Hierarchical Atlas (BHA), comprised of 2514 labels [7], and the Harvard-Oxford Atlas, which uses 177 labels [12]. Additional Atlases can be created and existing ones can be customized easily, simply by using a pre-existing Atlas as a template for defining a new one. (See the online instruction manual available at <https://github.com/CreativeCodingLab/NeuroCave> for more details of how to load in datasets and customize Atlases.)

*jgadel2@uic.edu

†oajilore@psych.uic.edu

‡aleow@psych.uic.edu

§aforbes@uic.edu




Figure 1: The left panel shows the user interface of *NeuroCave*, presenting multiple views to investigate connectome data. A researcher can compare different datasets or the same dataset from different perspectives. Here, (A) shows a high resolution functional connectome, and (B) shows the same connectome from a different orientation. When viewing multiple copies of the same dataset, actions can be synchronized so that interacting with one connectome updates the other. Users can choose an atlas to label brain regions. (C) shows the the color/glyph atlas, and a user can toggle on or off specific brain regions or classifications by clicking next to items in the atlas. Here, the right parietal lobe has been turned off, making it easier for the researcher to interactively explore the specific areas of interest. In (D) a slider controls the opacity of all selected connectivities; in (E), users can turn edge bundling on or off for selected brain regions; and in (F), users can set the minimum or maximum threshold values and number of hops to determine which edges to display. Users can change the glyph size of individual nodes or selected brain regions in order to highlight relevant information. Here, the user has selected and enlarged a node (G), and is investigating edges emanating from the left parietal lobe above a threshold of 1.35518. *NeuroCave* is a web application that runs in both desktop and mobile environments, and users can switch seamlessly between the standard and VR modes on demand. The inset image (bottom center of the left panel) shows a neuroscientist exploring a 3D dataset in virtual reality using Oculus Rift VR with Touch controllers. The right panel shows a modular representation on *NeuroCave*'s platonic solid layout, which can be used to effectively visualize multi-level hierarchical clustering.

A common task in disease studies involves the comparison of two groups of subjects— e.g., a healthy control group versus a disease group— in order to derive conclusions about alterations due to the disease. To the best of our knowledge, no existing connectome visualization application facilitates real-time simultaneous comparison for two or more clinical datasets (although immersive environments have been used to provide comparisons between individuals, including in neuroscience contexts [6, 16]). To address this, *NeuroCave* enables neuroscientists to visualize connectome datasets via a synchronized “side-by-side” layout, making it easier to explore differences between groups of subjects, or the same group represented using different spaces, modalities, or in different coordinate systems.

2.2 Topology Visualization

NeuroCave positions nodes according to the provided topological information. Available topologies include the anatomical positioning or any of number of applied transformations that reformulate this positioning into an abstract space. These topologies are automatically identified by the application, and ongoing development aims to enable the transformation of anatomical datasets into a range of topological spaces on-demand.

Currently, we have applied a range of dimensionality reduction techniques to connectome datasets, including Isomap [19] and t-SNE [11], and we make use of these methods to help identify patterns in the “intrinsic geometry” (i.e., the geometry as determined by the brain connectivity itself, either structural or functional) of a connectome dataset [24].

Users can switch between anatomical and abstract topological spaces as needed to support particular analyses, making it possible to see the same data transformed in various ways in order to investigate the connectome from a range of different perspectives. Both panels in Fig. 2 show a comparison of the same connectome dataset in an

anatomical versus a topological space (for functional and structural connectomes, respectively).

2.3 Clustering Visualization

NeuroCave also supports the visualization of clusters of nodes (i.e., modular or community structure), either embedded within a topological space, or simply as groups of related points (where the spatial positioning of a cluster of nodes within the cluster has no meaning). Clustering information is input as a vector of integer values, where each value represents a different module or cluster. When there is no meaningful spatial positioning provided for clusters (or when we choose to exclude this information), *NeuroCave* makes use of a novel layout technique that exploits the geometrical properties of platonic solids. In brief, a platonic solid is a regular, convex polyhedron constructed by congruent regular polygonal faces with the same number of faces meeting at each vertex. Five platonic solids exist: tetrahedron, cube, octahedron, dodecahedron and icosahedron, with four, six, eight, twelve, and twenty faces, respectively. Based on how many clusters are generated, a suitable platonic solid is chosen such that its number of faces is greater than the number of these clusters, with glyphs for each cluster covering the corresponding face of a platonic solid embedded in a sphere. This enables the user to “enter” into the geometry (i.e. into the “NeuroCave”) via one of the unpopulated faces, providing a more immersive experience of the data, especially when he or she toggles the display to activate the virtual reality mode.

Users can interactively rearrange the position of the clusters within the platonic solid in order to more easily see particular clusters, for example, those that are densely interconnected, or that are relevant for a particular analysis session. When applicable, clustering can be recomputed on demand, with the user specifying the number of clusters, which in turn updates the platonic solid that is generated. Connections between individual nodes within clusters or between clusters can be visualized as well, as described below. The




Figure 2: Left panel: Connectivity emerging from the anterior (red ring) and posterior (yellow ring) parts of the precuneus in anatomical space (leftmost connectome) and an “intrinsic” space generated using Isomap (second connectome from left). The color code represents the modular structure of the connectome consisting of 4 modules. Note that the orange community contains the default mode network. The bottom-left inset shows another view of the left panel with all nodes enlarged in order to better see the modular structure. The top-center inset plot shows the residual geodesics for the first 10 dimensions of the Isomap dimensionality reduction algorithm. Right panel: *NeuroCave* also supports explorations of structural connectomes, here we see both intrinsic (second connectome from the right) and anatomical (rightmost connectome) geometry, for the same connectome dataset gathered the same participants as the functional connectome shown in the left panel, using a different color atlas.

right panel of Fig. 1 shows an example of visualizing multi-level hierarchical clustering using our platonic solids approach.

2.4 Node visualization

By default, we utilize two different glyphs (spheres and cubes) to differentiate between left and right hemisphere affiliation. Nodes can be colored according to lobar or modular information. Controlling nodal transparency is also possible according to their color scheme and colors can be interactively assigned to different amounts of transparency modes as desired. For example, a brain region that has been assigned a particular color can be toggled on or off (made visible or invisible), or the transparency of one or more regions or clusters can be increased or decreased in order to emphasize or de-emphasize them (see Fig. 1, left). The glyph size of individual nodes or groups of nodes is also interactively adjustable by the user via our interface, both in desktop mode or in VR mode. Text labels identifying the nodes can be displayed for all nodes or for user-selected nodes on demand.

2.5 Edge Visualization

NeuroCave introduces a range of features to visualize edges efficiently in order to show the connectivity between brain regions. Network visualizations that have an excessive amount of overlapping edges, common in dense node-link diagrams, can introduce unwanted visual clutter, which makes it more difficult to read and interpret the network. We provide different ways to mitigate this problem. First, we provide the option to hide all edges by default (i.e., to show only the nodes), and then enable a user interactively add edges as desired. In this mode, a user can select any node as a “root” node, causing all connected edges stemming from this node to be displayed. Second, to minimize the clutter occurring from edge crossings, even when all edges are displayed, we use the force directed edge bundling (FDEB) algorithm to group edges going in the same direction [10]. FDEB creates these bundles through an iterative algorithm that consists of a series of subdivision cycles. In each cycle, we subdivide an edge into a specified number of points (by default, we use 6 cycles, and we double the number of points each cycle, ending up with 64 subdivision points plus the two original points of the edge). After the subdivision, we iteratively move each subdivision point in an update step to a new position determined by modeling the forces among the points.

Standard implementations of edge bundling are too slow for the large numbers of edges that can appear in some connectome datasets

visualized in *NeuroCave*, reducing the frame-rate of the application and preventing an effective real-time experience. Therefore, we introduce an enhanced WebGL texture-based implementation, extending previous work by Wu et al. [22]. Since the division and update operations are performed on each point independently, the FDEB algorithm is parallelizable and can be optimized for the GPU. The texture-based method stores the subdivision points in a 2D GPU texture, where each row represents the 3D coordinates of points belonging to the same edge. Since write operations are generally unavailable to GPU textures in most WebGL implementations, a ping-pong algorithm is used to first render results to an offscreen framebuffer object. Two shaders are utilized: the first performs the subdivision operation, and the second executes the update steps. GPU textures possess a limitation on their sizes. Hence, a large number of edges cannot fit inside one texture. In Wu et al.’s implementation, this limitation on the texture size limits is not addressed. We enhanced the algorithm using a tiling approach in which we can extend the maximum number of edges by using multiple textures. Since the total number of points of each edge after all cycles will be known ahead (64 subdivision + 2 end points), we tile the edges when the maximum number of possible rows per texture is achieved. Our texture-based implementation can bundle the closest 1000 edges to the selected node at interactive rates on a desktop computer. While this is sufficient for the datasets we explored, we also enable users to choose threshold values that limit only connections above or below specified strengths to be computed, both to improve performance in situations where the data contains very dense interconnections, and to assist in analyses focused on particular connectivity weights.

Each edge can be colored using a gradient, whose two colors are chosen according to the colors of the source and target nodes that it connects. The gradient is skewed towards the node possessing the higher nodal strength, that is, the sum of weights of links connected to the node. This enables the user to quickly recognize the strength of the selected node with respect to its interconnected neighbors, which can help in identifying important nodes or clusters of nodes, as well as to highlight the reason for modular changes when they occur in group studies. The left panel of Fig. 2 shows an example of this edge coloring approach.

2.6 Virtual Reality

Ware and Mitchell [21] show that stereographic visualization reduces the error rate in graph perception for large graphs. Alper et al. [2] observe that, when coupled with a highlighting technique, stereoscopic representations of 3D graphs outperform their non-immersive coun-

terpart. *NeuroCave* harnesses the capabilities of VR environments, which can facilitate spatial manipulation, identification, and classification of objects and imagery, and can aid users in understanding complex scenes [3, 9, 13], extending the immersive functionality available in previous connectome analytics software [4, 14].

NeuroCave can be viewed on a normal desktop or mobile environment, or via a VR system. Currently we support the Oculus Rift and the Samsung Gear VR platforms, with support for additional platforms planned in the near future. In addition to the standard 3D manipulations of panning, rotating, and zooming, *NeuroCave* supports the advanced interaction features available on the Oculus Rift via the use of Oculus Rift Touch controllers. The Touch controllers are a pair of VR input devices that track each hand, enabling an effective gesture-based manipulation of the VR environment. The user selects the preview area to be explored in VR and then uses the “thumbsticks” on the Touch devices to navigate the visualized connectome through panning, rotating, and zooming. Nodal selection is enabled via a two step procedure: first, pressing the grip button lets the user point at and highlight a node; second, pressing the index button selects the highlighted node. We can mimic some of this functionality in other VR platforms (e.g., any platform that supports WebVR, such as Google Cardboard or Daydream), but node selection is not as effortless if the controllers do not contain tracking sensors. Users can enter and leave VR mode as often as they like in order to support investigations of connectome datasets.



Figure 3: In this photo, a user is shown wearing a Neuroelectrics Starstim HeadCap, which reads EEG signals and can also direct transcranial current stimulation to particular brain regions. Future work will explore the live representation of a patient’s connectome as a component of a biofeedback application that allows a patient to modulate his or her own brain activity.

3 DISCUSSION

Through using the various features in combination in order to interact with connectome datasets, *NeuroCave* supports a range of analysis tasks. The rich set of visualization features provided by *NeuroCave* makes it possible for users to explore connectome datasets in flexible

manner, to make observations about connectome data, to generate hypotheses about these observations, and then to dive in more deeply to support or invalidate hypotheses. That is, *NeuroCave* supports the process of generating and querying visual representations in order to answer task-specific questions and to facilitate sensemaking [17].

From our initial qualitative observations, we found that users were very engaged in exploring the data while in VR mode, and enjoyed switching between the different available layouts based on the different clustering and dimensionality reduction techniques. Users also indicated that they appreciated the ability to bring up different datasets on demand while immersed in the VR mode, especially to see if patterns discovered in one connectome (e.g., a connectome dataset representing average healthy subjects) were present in another (e.g., a connectome dataset representing average diseased subjects). Somewhat surprisingly, users also readily moved between the desktop display and the VR display without complaint. Future work will quantitatively assess the impact of VR mode on analysis tasks and empirically investigate the current *NeuroCave* workflow, which encourages moving between VR and desktop modes.

NeuroCave has been utilized in a wide range of contexts to explore a variety of connectome datasets. These include a resting-state fMRI high-resolution dataset consisting of 2514 regions-of-interest,¹ as well as sex-specific resting-state functional connectomes in the F1000 repository, a large 986 subjects resting-state fMRI connectome dataset,² both of which are publicly available at the NITRC neuroimaging data repository. For the latter analysis, subtle differences were discovered between male and female connectomes related to language and emotion/affect processing, as were differences in self-referential/autobiographical information retrieval. Future work will build upon this preliminary insight, and we plan to extend *NeuroCave* to support visual analyses to validate the discoveries that were first observed in an immersive environment.

We plan to extend *NeuroCave* to support temporal connectome datasets, and also to integrate real-time EEG data, which will introduce novel visual analytics challenges [23]. Ongoing research at CoNECT Lab³ investigates the use of noninvasive transcranial current stimulation for cognitive enhancement and therapeutic purposes, such as neuro-rehabilitation and emotional regulation. In one pilot study, subjects wear a Neuroelectrics Starstim HeadCap,⁴ and direct stimulation to particular brain regions in order to mitigate anxiety. Future immersive analytics applications could present an interactive feedback loop in which clinicians could see a visual representation of a patient’s connectome, and stimulate the patient to modulate brain activity in order to improve cognitive functionality.

The *NeuroCave* application and open-source code, along with detailed instructions, video documentation, and examples, are freely available at our GitHub repository: <https://github.com/CreativeCodingLab/NeuroCave>. All panels in Fig. 1 and Fig. 2 show screenshots of connectome visualizations created using *NeuroCave*, and a range of connectome datasets can be explored via the online version of *NeuroCave* at <https://creativecodinglab.github.io/NeuroCave>.

¹http://www.nitrc.org/frs/?group_id=964

²https://www.nitrc.org/projects/fcon_1000/

³The home page for CoNECT Lab at University of Illinois at Chicago can be found at <http://brain.uic.edu/>

⁴<http://www.neuroelectrics.com/products/starstim/starstim-tcs/>

REFERENCES

- [1] S. Achard, R. Salvador, B. Whitcher, J. Suckling, and E. Bullmore. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *Journal of Neuroscience*, 26(1):63–72, 2006.
- [2] B. Alper, T. Hollerer, J. Kuchera-Morin, and A. G. Forbes. Stereoscopic highlighting: 2d graph visualization on stereo displays. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2325–2333, 2011.
- [3] C. J. Bohil, B. Alicea, and F. A. Biocca. Virtual reality in neuroscience research and therapy. *Nature Reviews Neuroscience*, 12(12):752–762, 2011.
- [4] G. Conte, A. Q. Ye, K. Almryde, O. Ajilore, A. Leow, and A. G. Forbes. Intrinsic geometry visualization for the interactive analysis of brain connectivity patterns. In *Proceedings of Electronic Imaging, Visualization and Data Analysis*, pp. 481–1–8, 2016.
- [5] G. Conte, A. Q. Ye, A. G. Forbes, O. Ajilore, and A. Leow. BRAINtrinsic: A virtual reality-compatible tool for exploring intrinsic topologies of the human brain connectome. In Y. Guo, K. Friston, F. Aldo, S. Hill, and H. Peng, eds., *Brain Informatics and Health*, chap. 7, pp. 67–76, 2015.
- [6] M. de Ridder, Y. Jung, R. Huang, J. Kim, and D. D. Feng. Exploration of virtual and augmented reality for visual analytics and 3d volume rendering of functional magnetic resonance imaging (fMRI) data. In *Proc. IEEE Symposium on Big Data Visual Analytics (BDVA)*, 2015.
- [7] I. Diez, P. Bonifazi, I. n. Escudero, B. Mateos, M. A. Muñoz, S. Stramaglia, and J. M. Cortes. A novel brain partition highlights the modular skeleton shared by structure and function. *Scientific Reports*, 5, 2015.
- [8] B. Fischl. Freesurfer. *NeuroImage*, 62(2):774–781, 2012.
- [9] A. Forbes, J. Villegas, K. R. Almryde, and E. Plante. A stereoscopic system for viewing the temporal evolution of brain activity clusters in response to linguistic stimuli. In *Proc. SPIE*, vol. 9011, p. 90110I, 2014.
- [10] D. Holten and J. J. Van Wijk. Force-directed edge bundling for graph visualization. In *Proc. Eurographics Conference on Visualization (EuroVis)*, pp. 983–998, 2009.
- [11] L. v. d. Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- [12] N. Makris, J. W. Meyer, J. F. Bates, E. H. Yeterian, D. N. Kennedy, and V. S. Caviness. MRI-based topographic parcellation of human cerebral white matter and nuclei: II. Rationale and applications with systematics of cerebral connectivity. *NeuroImage*, 9(1):18–45, 1999.
- [13] G. E. Marai, A. G. Forbes, and A. Johnson. Interdisciplinary immersive analytics at the Electronic Visualization Laboratory: Lessons learned and upcoming challenges. In *Proc. IEEE VR Workshop on Immersive Analytics*, pp. 54–59, 2016.
- [14] D. S. Margulies, J. Böttger, A. Watanabe, and K. J. Gorgolewski. Visualizing the human connectome. *NeuroImage*, 80:445–461, 2013.
- [15] D. Meunier, R. Lambiotte, and E. T. Bullmore. Modular and hierarchically modular organization of brain networks. *Frontiers in Neuroscience*, 4, 2010.
- [16] C. Nowke, M. Schmidt, S. J. van Albada, J. M. Eppler, R. Bakker, M. Diesmann, B. Henschel, and T. Kuhlen. VisNEST — Interactive analysis of neural activity data. In *Proc. IEEE Symposium on Biological Data Visualization (BioVis)*, pp. 65–72, 2013.
- [17] D. M. Russell, M. J. Stefik, P. Pirolli, and S. K. Card. The cost structure of sensemaking. In *Proc. ACM INTERACT and CHI Conference on Human Factors in Computing Systems*, pp. 269–276. ACM, 1993.
- [18] O. Sporns, G. Tononi, and R. Kötter. The human connectome: a structural description of the human brain. *PLoS Computational Biology*, 1(4):e42, 2005.
- [19] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [20] M. P. van den Heuvel and O. Sporns. Rich-club organization of the human connectome. *Journal of Neuroscience*, 31(44):15775–15786, 2011.
- [21] C. Ware and P. Mitchell. Visualizing graphs in three dimensions. *ACM Transactions on Applied Perception (TAP)*, 5(1):2, 2008.
- [22] J. Wu, L. Yu, and H. Yu. Texture-based edge bundling: A web-based approach for interactively visualizing large graphs. In *Proc. IEEE International Conference on Big Data*, pp. 2501–2508, 2015.
- [23] M. Xing, O. Ajilore, O. E. Wolfson, C. Abbott, A. MacNamara, R. Tadayonnejad, A. G. Forbes, K. L. Phan, H. Klumpp, and A. Leow. Thought Chart: Tracking dynamic EEG brain connectivity with unsupervised manifold learning. In *Proc. International Conference on Brain and Health Informatics (BHI)*, pp. 149–157. Springer, 2016.
- [24] A. Q. Ye, O. Ajilore, G. Conte, J. GadElkarim, G. Thomas-Ramos, L. Zhan, S. Yang, A. Kumar, R. L. Magin, A. G. Forbes, and A. D. Leow. The intrinsic geometry of the human brain connectome. *Brain Informatics*, 2(4):197–210, 2015.