

SwordPlots: Exploring Neuron Behavior within Dynamic Communities of Brain Networks

Chihua Ma and Angus G. Forbes

Department of Computer Science, University of Illinois at Chicago, Chicago, IL 60607

E-mail: cma6@uic.edu

Daniel A. Llano

Department of Molecular and Integrative Physiology, University of Illinois at Urbana-Champaign, Urbana, IL 61801

Tanya Berger-Wolf and Robert V. Kenyon

Department of Computer Science, University of Illinois at Chicago, Chicago, IL 60607

Abstract. Study of the behavior of individual members in communities of dynamic networks can help neuroscientists to understand how interactions between neurons in brain networks change over time. Visualization of those temporal features is challenging, especially for networks embedded within spatial structures, such as brain networks. In this article, the authors present the design of SwordPlots, an interactive multi-view visualization system to assist neuroscientists in their exploration of dynamic brain networks from multiple perspectives. Their visualization helps neuroscientists to understand how the functional behavior of the brain changes over time, how such behaviors are related to the spatial structure of the brain, and how communities of neurons with similar functionality evolve over time. To evaluate their application, they asked neuroscientists to use SwordPlots to examine four different mouse brain data sets. Based on feedback, their visualization design can provide neuroscientists with the ability to gain new insights into the properties of dynamic brain networks. © 2016 Society for Imaging Science and Technology. [DOI: 10.2352/J.ImagingSci.Technol.2016.60.1.010405]

INTRODUCTION

Application of the approach of social network analysis (SNA) to mining and analyzing neurobiology data sets can be valuable in helping neuroscientists to understand and explore their data from different perspectives. The application of SNA to the study of brain networks has nodes mapped to a single neuron or a group of neurons, while edges between nodes represent a functional connection between neuronal sites. Neurons communicate through electrochemical signals. They process incoming chemical and electrical signals and transmit these signals to other neurons and other types of cells. A neural network is a group of connected neurons that performs a certain function. A community is defined as a group of nodes that tend to interact with other nodes in the same group more often than the nodes outside the group. However, SNA is not adequate to measure brain functional connectivity over time

since it lacks the ability to capture dynamic changes in the connectivity of these networks. Recently, there have been attempts to ameliorate this problem, one of which is dynamic network analysis (DNA).¹ Detection of the evolution of neural communities using DNA can be helpful for understanding important changes in network structure and individual node behavior over time.

We worked closely with neuroscientists who use the DNA method to explore changes in functional behaviors and their resulting effects on community identities of neurons over time. The research goal of our domain experts is to utilize social network modeling algorithms to analyze mouse brain time-series imaging data to uncover interactions among neurons across time and space. The brain networks can be identified through the application of a linear correlation method and then passing these correlation metrics to a community analysis algorithm. Given the large numbers of potential nodes (>20,000) in a typical brain imaging experiment, effective visual representations are necessarily needed to help domain scientists to study the mouse brain using the model of DNA. Although a number of visualization tools have been developed during the last few years, the development of applications for the visualization and visual analytics of dynamic brain networks still requires further effort. A major challenge in visualizing dynamic brain networks is to reveal temporal features while at the same time accounting for additional constraints imposed by brain architecture, i.e., neuron location. Spatial and non-spatial visual integration challenges are common across domains, from geospatial applications to bioinformatics to neuroscience applications.²

In this article, we present SwordPlots, an interactive multi-view visualization application which assists the domain scientist to explore neuron behaviors within dynamic communities in the mouse brain. Our visualization design also provides the ability to filter data in a variety of ways in order to make it easier for the user to explore the data. The user can apply these capabilities to focus on a single node of interest. A single node encapsulates many neurons, but neuroscientists assume that the neurons within a single node

Received June 30, 2015; accepted for publication Nov. 4, 2015; published online Dec. 09, 2015. Associate Editor: Thomas Wischgoll.

1062-3701/2016/60(1)/010405/13/\$25.00

have similar functionality since they are all in close proximity to the same brain region. Once a node of interest has been identified, the user has the ability to highlight the node to gain more detailed information. We also create an additional visualization view which displays some characteristics of the community model structure in space. We use two case studies to illustrate the effectiveness of our visualization tool. The visualization has the overall effect of reducing the data load on the scientists by allowing them to pare down a large number of neural connections within an experiment to those areas of interest and examine their behavior in time.

The visualization has the contributions of allowing neuroscientists to choose any one of the nodes to track its behavior within dynamic brain communities during a selected time range, providing sufficient interactions at both the individual community level and the node level which can offer the user an overview of the network and the ability to focus on nodes of interest, and being designed for a very new research topic which combines the DNA with neuroscience. The remainder of this article is organized as follows. Background knowledge is introduced in the second section. The third section discusses the related work. In the fourth section we present the implementation of our visualization system in detail. The fifth section presents two case studies, followed by a qualitative evaluation in the sixth section. The seventh section discusses some limitations and future work, and the eighth section summarizes the contributions of the article.

BACKGROUND

Brain Network Analysis

Network analysis of brain imaging data has proven to be a particularly challenging computational problem. For example, functional magnetic resonance imaging (fMRI) data using blood oxygen level dependent signals involves the analysis of tens of thousands of voxels, each with their own time series. Analysis of pairwise correlations across all of these voxels produces millions of potential interactions between voxels, leading to difficulties in producing a coherent understanding of the patterns of activity. Early efforts to analyze these data sets were focussed on seed-based correlation approaches, whereby a time series in a region of interest was correlated to all other voxels in the brain.³⁻⁵ Such methods have uncovered multiple different large-scale brain networks, but are susceptible to biases based on the initial seeds chosen. Multivariate approaches, such as independent component analysis (ICA), have also been used and are free from the biases inherent in visualizing networks based on an initial seed.^{6,7} Most recently, graph-theory-based methods have been applied to brain imaging data and have revealed not only the presence of underlying networks but also the logic underlying the organization of networks by characterizing certain nodes as hubs, and by revealing small-world behavior in certain brain networks.⁸

A potential problem with these approaches to brain network analysis is their limited ability to extract the natural dynamics of a network's organization. That is, generally

these techniques are used to generate a single static network map using a window of many minutes of the obtained data. However, brain activity is intrinsically highly dynamic, with functional associations between neurons and brain regions that ebb and flow as the organism's level of arousal, focus of attention or topic of thought changes. Recently, neuroscientists have applied DNA tools initially developed for the study of social dynamics^{9,10} to the analysis of mouse brain imaging data. Among the analysis tools commonly used to understand the relationship between structure and function in social networks is community inference. The definition of "community" in the analysis of brain networks is analogous to that of a neural assembly,¹¹ which we define as a group of neurons that act collectively and are functionally connected.

Community Analysis

In network analysis, communities are generally defined as groups of nodes that tend to have more and stronger connections with each other than with nodes outside. This concept is derived from social networks, and can be applied in a variety of fields including biological networks.¹² Communities appear in networks where nodes join together in tight groups that have few connections between them. Several methods have been proposed to find the structure, of which the best known is called modularity optimization. Modularity is a standard metric for finding and evaluating communities.¹³ Community detection by modularity optimization on large networks is a computationally challenging problem. Therefore, it needs to use algorithms that find high modularity partitions within large networks in a short time. The Louvain algorithm¹⁴ is ideally suited to finding communities in such cases. Besides the advantage of being extremely fast, its steps are easy to implement. In each static network, the nodes within the same cluster have more edges connected to each other than to nodes outside the cluster.

When dealing with dynamic networks, we are interested in analyzing the evolution of communities across time that can be defined as dynamic communities. We use the method of Dynamic Community Inference (CommDy)^{9,15-17} to study how the interactions and the structure of clusters change over time. In this method, dynamic communities are essentially viewed as dynamic clusters, where the membership of the individual inside the cluster is determined by a total value of "social cost". The definition of social cost is based on two explicit assumptions about individual behavior, motivated by research in the social sciences.⁹ First, it assumes that individuals tend not to change their home community affiliation too often.¹⁸ Second, it assumes that individuals tend to interact with their respective home communities most of the time.¹⁹ These assumptions are translated into three cost parameters potentially incurred by an individual. First, it posits a cost for a switch from one community to another. Second, there is a cost of visiting a community of which one is not a member. Third, in data sets for which not all individuals are observed all of the time, there is a

cost of absence for an individual who is not observed at a gathering of the community of which it is a member. A dynamic community is then defined as a time series of sets of individuals among whom the overall social cost of interacting is minimized.^{9,15,17} Figure 2 shows an example data set of five individuals during six time steps. Circles are individuals labeled with their identification numbers, while rectangles are groups. Communities are identified by matching colors.

Brain Network Data and Characteristics

The data of dynamic brain networks we use in the article are flavoprotein autofluorescence imaging data.^{20,21} The technique of flavoprotein autofluorescence permits imaging of subthreshold activity across broad areas of a living brain slice with high sensitivity without introducing potentially toxic dyes. The image data have dimensions of 172 (width) \times 130 (height) \times T (time). Each pixel (node) in the image represents a group of neurons. A weighted correlation network is generated with a list of weighted edges connecting pairs of nodes. Weighted edges $\omega(X, Y)$ represent the linear correlation coefficient between any pair of two nodes (X and Y) over a certain time window t (Eq. (1)). The threshold correlation coefficient and window size are both systematically varied. By sliding the window one time step (for each iteration) over the entire timeline (T frames), a time series of correlation networks is obtained.

$$\begin{aligned} \omega(X, Y) &= \text{corr}(X, Y) \\ &= \frac{1}{t-1} \sum_{i=1}^t \left(\frac{X_i - X_{\text{mean}}}{S_X} \right) \left(\frac{Y_i - Y_{\text{mean}}}{S_Y} \right), \\ \text{corr}(X, Y) &\in [-1, 1], \end{aligned} \quad (1)$$

where

$$\begin{aligned} X_{\text{mean}} &= \frac{1}{t} \sum_{i=1}^t X_i \quad \text{and} \\ S_X &= \sqrt{\frac{1}{t-1} \sum_{i=1}^t (X_i - X_{\text{mean}})^2}. \end{aligned}$$

Then we find snapshots of functional clusters using the Louvain clustering method, followed by the application of CommDy. To better understand the neuron behaviors, we need to identify some characteristics of nodes in dynamic brain networks based on the use of CommDy. The SwordPlots application encodes the characteristics listed below.

Raw Data Attributes. In our brain network data, the coordinates of pixels show the real locations (regions in the brain) of the signals, and the pixel intensities represent how strong/weak the signals at the corresponding regions are. Therefore, the pixel values as well as their coordinates within the image are critical features for analyzing the network.

Node Degree. The node degree for many real networks can yield insight into mechanisms underlying system growth.²² Thus, there are significant benefits to visualizing the overall

spatial shape of the degree distribution and its variation with time.

Home Community. Each individual has its own home community identification code to identify its membership. The individuals belonging to the same community have the same home community identification code. Neurons belonging to the same community have similar functionality.

Temporary Community. Each individual also has its temporary community identification code to identify the group where it stays at a certain time step. The individuals staying in the same community have the same temporary community identification code. Neurons have similar functionality to their current community, but tend to return to their original home community after a short time.

Consistency. Sometimes, the individual may have a stronger connection with members in other groups. Thus, it visits other communities for a short period of time, without changing its home community identification code. The consistency is a measurement of times when the temporary community of a node is the same as its home community during the entire activity period. For example, node 4 in Fig. 2 belonging to the blue community visits the pink community at T3. Its consistency is 0.8. Neurons behave anomalously if they have low consistency.

Switchingness. Switchingness measures the percentage of times when a node changes its membership during the entire activity period. For example, node 3 in Fig. 2 changes its home community identification twice, from T1 to T2 and from T4 to T5. Its switchingness is 0.5. Neurons with low consistency and high switchingness have unstable functionality.

Task Analysis

Based on in-depth discussions with domain experts from neuroscience and computational biology, we identified five main tasks that could be enabled using a visualization application. *Task 1:* observe multiple attributes of neurons over time to understand how neurons interact with each other (neuron behaviors). *Task 2:* interact with an individual node (i.e, a group of neurons), such as selecting interesting nodes at any region in the brain. *Task 3:* find relationships between the functional and spatial structures in the brain, to see whether two regions are close to each other or cross two sides of the brain when they are highly correlated. *Task 4:* identify patterns in how dynamic communities change over time in order to show functional connectivity in the brain. *Task 5:* generate hypotheses and make predictions about brain data sets.

RELATED WORK

Brain/Neuronal Connectivity Visualization

A wide range of visualization tools for biological network analysis has been discussed in a comprehensive survey by Pavlopoulos et al.²³ Cytoscape^{24,25} is a popular bioinformatics package for biological network visualization

and data integration. In general, the analysis and visualization of brain networks focus on either the cell level or the region of interest (ROI) level. At the cell level, Lin et al. developed the Neuron Navigator (NNG)²⁶ for analyzing, observing and discovering the connectivity within the neural maze. Sorger et al. implemented an interactive two-dimensional graph, *neuroMap*,²⁷ to render the brain and its interconnections. Al-Awami et al. designed *NeuroLines*²⁸ for scalable detailed analysis of neuronal structures and their connectivity at the nanoscale level. For the ROI level, Li et al.²⁹ presented a visual analytics approach for the identification of ROI and brain network exploration. Jianu et al.³⁰ developed a visualization of two-dimensional neural maps for connectivity exploration and analysis in the human brain. Irimia et al.³¹ implemented a circle-based visual representation of human connectomics for classifying neuron connectivity relationships in the brain. Margulies et al.³² provided an overview of various frameworks for visualizing anatomical and functional connectivity in the human brain. Christodoulou et al. presented *BrainNetVis*,³³ which effectively quantifies and visualizes brain networks. *BrainNet Viewer*,³⁴ developed by Xia et al., is a graph-theoretical network visualization toolbox which uses MATLAB to generate figures for brain connectomes. In addition, the comparison of different connectivity data is another important problem in the neurosciences. Alper et al.³⁵ discussed matrix-based visual representations to compare brain connectivity patterns. Forbes et al.³⁶ introduced the 3D+Time Brain View application for visualizing functional fMRI data in a stereoscopic environment; their system illustrates the temporal evolution of brain activity clusters in response to linguistic stimuli. Ma et al.³⁷ used an animated dual representation consisting of the enhanced node-link diagram and the distance matrix to visualize dynamic brain networks. However, it has the limitation on scaling to larger networks.

Dynamic Data Visualization

A good overview of dynamic data visualization techniques is presented by Beck et al.³⁸ Beck et al. state that the time dimension can be effectively mapped using animation or to a space dimension representing a timeline.^{39,40} Animation of a series of dynamic graphs has the advantage that the entire screen can be devoted to the drawing of the graph.⁴¹ *GraphAEL*⁴² supports the notion of time slice to visualize evolving graphs with a temporal component. Frishman and Tal developed an algorithm for drawing a sequence of clustered graphs,⁴³ and also created a method for drawing a sequence of graphs online based on a force-directed node-link layout.⁴⁴ Animation of transitions between time slices may also help users to understand how the graph structure evolves.⁴¹ Bach et al.⁴⁵ designed a visual interface called *GraphDiaries* to improve support for identifying, tracking and understanding changes in dynamic networks. Several other studies⁴⁶⁻⁴⁸ also support the idea that animation can be an effective way of presenting transitions that are beneficial for the purpose of visualization.

Although animation is useful to visualize the dynamic changes in networks, the exploration of animated graph diagrams leads to high cognitive efforts due to our limited short-term memory.⁴⁹ A better overview is provided if the time dimension is mapped onto a timeline and the dynamic graph is visualized in a single static image.⁵⁰ Compared with animation, the timeline approach makes it easier for users to identify changes in the context of the evolution, as they are still visualized on the screen. Classic examples of time-series data visualization techniques such as line graphs and bar charts focus on presenting univariate data sets. Visualization of multivariate time-series data is complex and requires special effort. A well-known technique dealing with this problem is the *ThemeRiver*^{51,52} developed for document visualization. It is an intuitive representation of the change between both variables and time steps. Placing node-link diagrams on a timeline is a simple way to visualize dynamic networks. Greilich et al. designed *TimeArcTrees*⁵³ for visualizing weighted, dynamic compound digraphs by drawing a sequence of node-link diagrams from left to right in a single view. When visualizing dynamic communities, line graphs are generally used to represent individual vertices, with color-coding presenting community identity. Reda et al.⁵⁴ focussed on revealing the community structure implied by the evolving interaction patterns between individuals in dynamic social networks. Vehlou et al.⁵⁵ combined a dynamic community structure with a dynamic graph in a single image to reveal typical life time phenomena of communities.

Besides mapping the time to the 2D space, visualizing temporal data based on the space-time cube has become another popular method for dynamic graph visualization. Gatsky et al.⁵⁶ presented some exploratory analysis of spatio-temporal data by using the space-time cube. Groh et al.⁵⁷ used an interactive three-dimensional model to visualize activity and social proximity in streaming event data during a given time period. *SocioScape*⁵⁸ is an interactive tool for the visual exploration of spatially referenced, dynamic social networks. The *Wakame* visualization system⁵⁹ can support discovering anomalies and comparing performance across multiple time series for multi-dimensional, spatio-temporal data. Bach et al. introduced *Matrix Cubes*⁶⁰ to represent dynamic networks based on the space-time cube metaphor.

Although a large number of both neuronal connectivity visualizations and dynamic graph visualizations exist, attempting to integrate them is a challenge since the former focus on spatial data and the latter focus on non-spatial data. Maries et al. developed a framework called *GRACE*² for the visual integration, comparison and exploration of correlations between spatial and non-spatial data sets. However, our goal is to visualize the dynamic change of multiple abstract attributes (non-spatial) of mouse brain networks (spatial) during a period of time. We found that none of the existing visualizations enabled all of the tasks defined in the section above. Thus, we introduce a novel visualization application to assist domain experts in

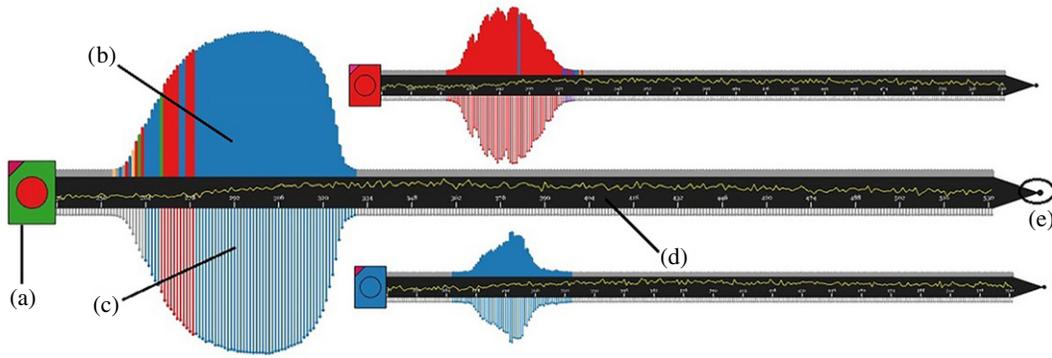


Figure 1. A SwordPlot includes multiple integrated components for representing time-series data and the community membership of a node in a dynamic network. This figure shows the main elements of a SwordPlot: (a) the sword pommel for representing the community identifications (the circle for home community and the rectangle for temporary community) of the node at the current time step; (b) the upper sword cross-guard for representing the node’s temporary community identifications over time; (c) the lower sword cross-guard for representing the node’s home community identifications over time; (d) the sword blade for representing the raw data (pixel value); and (e) the sword point used for changing the size of a SwordPlot.

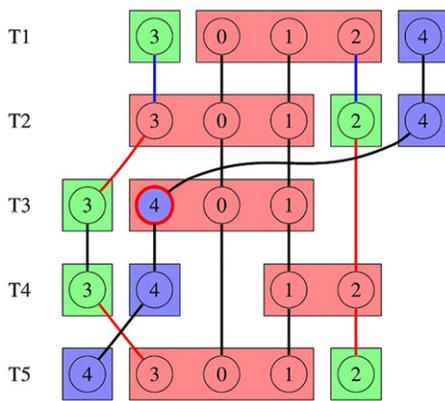


Figure 2. Visualization of CommDy on an example data set including five members across five time steps.¹⁵ Rectangles represent temporary communities, while circles represent home communities. The same colors indicate the same community identifications. At T1, there are three communities: green, pink and blue. At T2, member 3 switches to the pink community and member 2 switches to the green community. At T3, member 2 is absent and member 4 visits the pink community without changing its home community identification. At T4, member 0 is absent, member 2 switches to the pink community, and member 4 comes back to its home community. At T5, member 3 switches to the pink community and member 2 switches to the green community.

exploring the dynamics of mouse brain networks at the neuron level without losing their spatial reference. Figure 3 shows an overview of the pipeline to process the mouse brain data.

VISUAL DESIGN

The main goals of our SwordPlots application are to explore neuron behavior over time and discover the relationships among the resulting dynamic communities in both space and time. Our visualization technique uses a combination of different visual representations, and customized interactions to support exploration of the domain data. According to the requirements of the domain experts, we implemented our

integrated multi-view visualization system with the following techniques. We choose to use timeline-based representations to observe neuron behaviors over time (*Tasks 1* and *4*). Animation may be helpful for *Task 1* as well. Sufficient interactions are provided for our domain scientists to interact with individual nodes (*Task 2*). The space-time-cube-based representation is used to find relationships between spatial structures and network properties (*Task 3*). The SwordPlots application includes four coordinated views: the image control panel; a panel that shows recommendations and enables the filtering of data; the Space-Attribute cube; and a viewport that presents a number of SwordPlot visualizations to represent individual nodes. Each of these views is discussed below.

SwordPlot

In our design, we evaluated different visual encodings to better represent important attributes of dynamic brain networks. Although the representation of circles and rectangles shown in Fig. 2 is effective to visualize the community identifications, it has limitations: it does not scale well to more than a few nodes or a few time steps; it does not clearly track how the home and temporary community identifications of individual nodes change over a long time period; it lacks any spatial references for the nodes. Thus, we created a SwordPlot model (Figure 1), a timeline-based plot for each node in the entire network. A detailed explanation of the sword parts is shown in Figure 4. A SwordPlot provides an overview of a node’s behaviors across an entire timeline (*Task 1*) and makes it easy to identify interesting behaviors at certain time steps, such as when it changes its memberships or temporarily visits other communities. In addition, displaying multiple SwordPlots makes it possible to compare different nodes at once to find patterns within the data (*Task 4*). We also use a brain image slice view to display the real locations (*Task 3*) of selected nodes represented by SwordPlots.

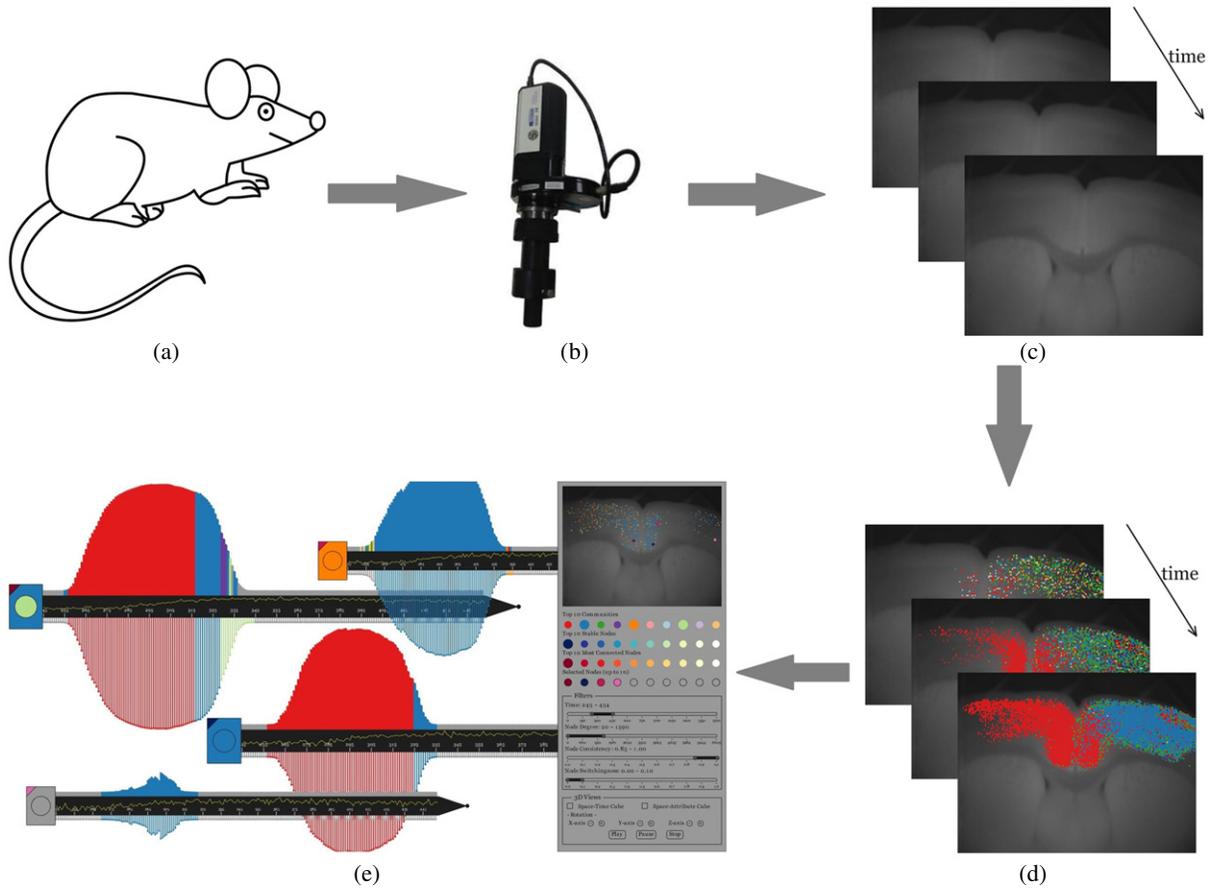


Figure 3. Flow diagram: (a) mouse brain, (b) camera, (c) time-series imaging data of mouse brain slices, (d) generate dynamic communities using DNA and (e) present the analysis results using the SwordPlots visualization.

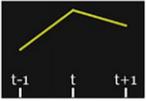
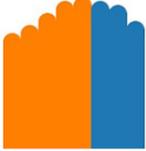
Sword Parts	Sword Pommel	Sword Body	Upper Cross-guard	Lower Cross-guard	Sword Point
Representation	Current Time Status	Raw pixel's value	Temporary Community	Home Community	—
Visual Encoding					
Interaction	change location	—	—	display detail panel	change size

Figure 4. A detailed explanation of the various visual encodings within the different components of a SwordPlot, showing the change in community over time for a single node.

Sword Pommel. We still keep the visual representation of the community identification shown in Fig. 2 using the color of the circle to represent the node's home community and the color of the rectangle to show its temporary community. This works well for a single node at a particular time step. The triangle at the top left corner is used for mapping the node to

its location in the image of the mouse brain slice. Dragging the sword pommel to move the SwordPlot to a new location avoids overlapping by others.

Sword Body. The sword body is used to display the raw pixel's value represented by the yellow line during the

selected time range. With the pixel values embedded, the relationships between neuron signals and their community identities/structures can be easily explored. This has the benefit of providing users with a view to see the correlations between these two attributes by visualizing the raw data pixel values. The time labels are drawn on the sword body as well. The time slider is used for zooming in/out on the SwordPlot, so that the neuroscientists can focus on the time range of interest.

Sword Cross-guard. The sword cross-guard is the bar-chart-based representation. The heights of the bars in both the upper and lower cross-guards are mapped to a statistic, e.g., node degree in this article. Since the neuroscientists are only interested in the node degree distributions in time qualitatively, but not the exact values at certain particular time points, we decided not draw the axes for the cross-guards. As we discussed in the Brain Network Data and Characteristics section, each individual should have two identification codes: the home community represented by the color of the upper cross-guard and the temporary community represented by the color of the lower cross-guard. We display the same color above and below if the member stays in its home community. If the individual is visiting its neighboring community at time t then we display a different color above (home) and below (temporary). The reason why we use two types of bar charts to represent the upper and lower cross-guards is to distinguish the two community identifications in a more obvious way. Missing cross-guards mean that the node degree is zero. In other words, the node is not active at the corresponding time. Clicking on the end point of the thin bar in the lower cross-guard pops up a detail panel showing some statistical information on the node at the corresponding time step. The statistical information includes some quantitative data, e.g., the pixel value, node degree, time point, etc.

Sword Point. To have more SwordPlots displayed and to avoid overlaps of multiple SwordPlots in the main view, the user is allowed to adjust the size of swords interactively by dragging the sword point. Dragging the sword point toward to the sword pommel decreases the size of the SwordPlot, while dragging it away from the sword pommel increases the size.

Figure 5 is a demonstration of the sequence of changes of a node's individual community identifications. The node with its membership in blue stays in the blue community until the time 319, and then joins the red community with change of its home community identification as well. However, it comes back to the blue community at time 326 without changing its membership. At time 329, it leaves the red community and becomes a member of the purple community. In the next time steps, it visits the orange community and then comes back to its home community.

Image Control Panel

The image of a mouse brain slice shown in Figure 6 is used as a control panel to select an interesting node. The user clicks on a node (pixel) in the image to draw a SwordPlot (Task 2) in the main view and add it to the list

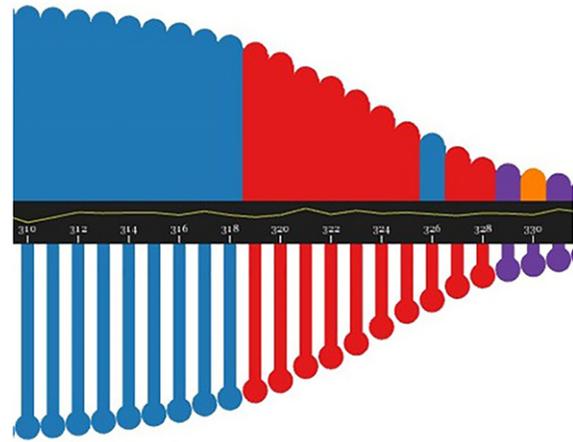


Figure 5. Transitions of the individual memberships of a node's home community (the lower part) and its temporary community (the upper part). The node changes its home community identity from blue to red, and then to purple. The node visits the blue community once when it belongs to the red community, and visits the orange community once when it belongs to the purple one.

of selected nodes as well. They click on the node in the selected list to remove its SwordPlot from the main view. Meanwhile, the image displays the corresponding locations of the selected nodes, which builds the connections between time-dependent attributes and the spatial information. In addition, nodes in the selected communities are also drawn on the image (Fig. 6). Playing animation helps the domain scientists to observe the community structures at different time steps and to track the evolution of communities. Theoretically, the SwordPlots for all nodes can be displayed in the main view simultaneously. However, due to the limitations of the distinguishable colors and human's working memory capacity, the number of selected nodes will not exceed a certain value. We discussed this with our domain experts and decided to use ten as the maximum number.

Recommendation and Filtering

Besides giving the ability for the user to explore some interesting regions randomly, our visualization provides recommendations as well to make the exploration more efficient. The three recommendations are calculated based on three characteristics: the community size, the node stability and the node degree. The community size is based on the total number of members the community has during the entire time period. The stability is calculated by two factors: consistency and switchingness. The higher consistency and lower switchingness the node i has, the more stable it is (Eq. (2)). The most connected nodes have high accumulated node degrees during the entire time period T (Eq. (3)). Again, as we indicated in the subsection above, we chose to use the top ten recommendations due to the scalability of the color encoding. We used ten categorical colors for the top ten communities from the D3 color ordinal scales,^{61,62} and used Cynthia Brewer⁶³ for the top ten stable and top ten most connected nodes. The user clicks on the circle buttons of the top ten communities to draw all of the

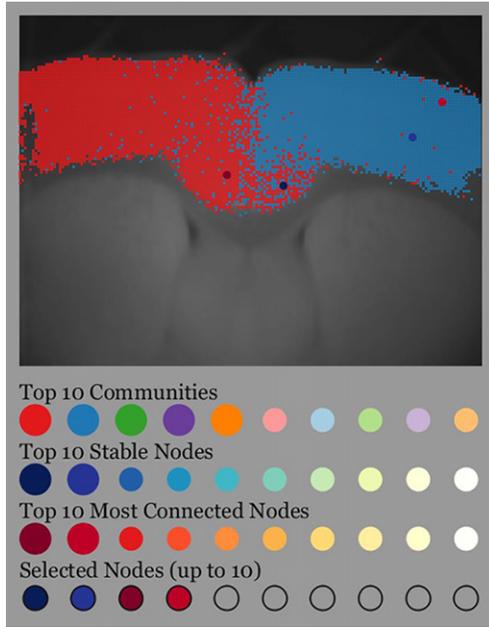


Figure 6. The image control panel is used for selecting interesting nodes and displaying community structures and identities. The recommendation panel includes the top ten communities, top ten stable nodes and top ten most connected nodes. The selected nodes are added into the list in the last row.

nodes in the corresponding community. In Fig. 6, the top five communities are selected to display on the image of a mouse brain slice. However, only the top two are being observed at the current time point. Fig. 6 also shows the two most stable nodes with larger blue dots and the two most connected nodes with larger red dots. Filter sliders are used to detect interesting features including time, node degree, node consistency and node switchingness. The main contribution of the filtering operation is to offer the user a way to explore the relationships between different attributes. For example, nodes that are located near the boundary of the activation area may have higher switchingness than nodes in the center. Nodes with higher degree have slightly lower consistency compared with nodes with lower degree.

$$\text{stability}(i) = \text{consistency}(i) \times (1 - \text{switchingness}(i)), \quad (2)$$

$$\text{TotalConnections} = \sum_{t=1}^T \text{nodeDegree}(t). \quad (3)$$

Space-Attribute Cube

We create a 3D model to represent a network statistic with the spatial reference in the brain networks. This is similar to but has important differences from the idea of the space-time cube. According to the technique of the space-time cube, the vertical dimension represents the positions of an object at different moments in time. In our 3D model, nodes can be represented in a cube as dots placed vertically according to one property value of the dynamic networks, such as node degree. We call this 3D model the Space-Attribute Cube. Color in the Space-Attribute Cube, like the SwordPlot, is used to represent the community identifications. The Space-

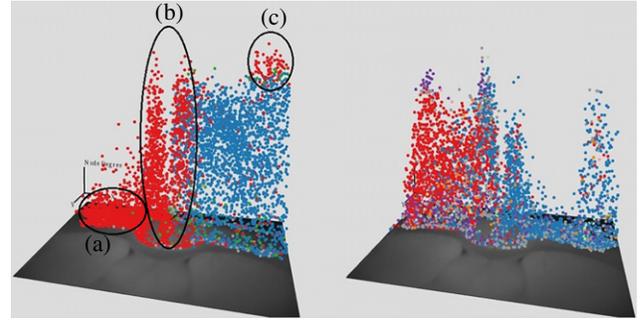


Figure 7. Space-Attribute Cubes at two time steps. Each individual node is represented by a colorful dot. The color indicates the community that the node belongs to. The X and Y coordinates show the real location of the node in the brain. The height of the node represents its corresponding node degree.

Attribute Cube reveals the distribution of node degrees of all active individuals in space, and how the distribution is related to the community structure. The left cube in Figure 7 demonstrates such a case, with the nodes belonging to the same community (red) located in different regions of the brain (a)–(c) with very different distributions of node degrees. One explanation is that the red nodes in (a) have strong connections with the nodes at the top of (b) and (c), but few connections with nodes that are also located in (a).

To overcome the problem of occlusion⁵⁹ in 3D representations (perspective projection) of data, the visualization provides three options for the user to rotate the Space-Attribute Cube along the X, Y, and Z axes in either the clockwise or the anti-clockwise direction. To discover how the property measured in the third dimension changes over time, animation is applied to the Space-Attribute Cube. The animation starts playing on pressing the play button and pauses on pressing the pause button. Once the animation is paused, the user can press the left or right arrow key on the keyboard to advance the animation one frame forward or in reverse. Animation used in the Space-Attribute Cube provides a spatial form representation for multivariate time-dependent data and a temporal representation for multivariate spatial data. Fig. 7 presents different distributions of node degrees at two time points. The left cube at earlier time indicates that the right side of the brain has stronger connections than the left side, while the right cube at later time shows the opposite situation.

CASE STUDY

We used the imaging data of $172 \times 130 \times T$, as described in the Brain Network Data and Characteristics section, to evaluate the SwordPlots application. Here, T is the number of frames, which is different (between 1000 and 2000) for different data sets. We build correlation networks based on these 22,360 pixels per frame captured during T frames, and then apply the CommDy method, as discussed in the Community Analysis section, to the time series of these networks.

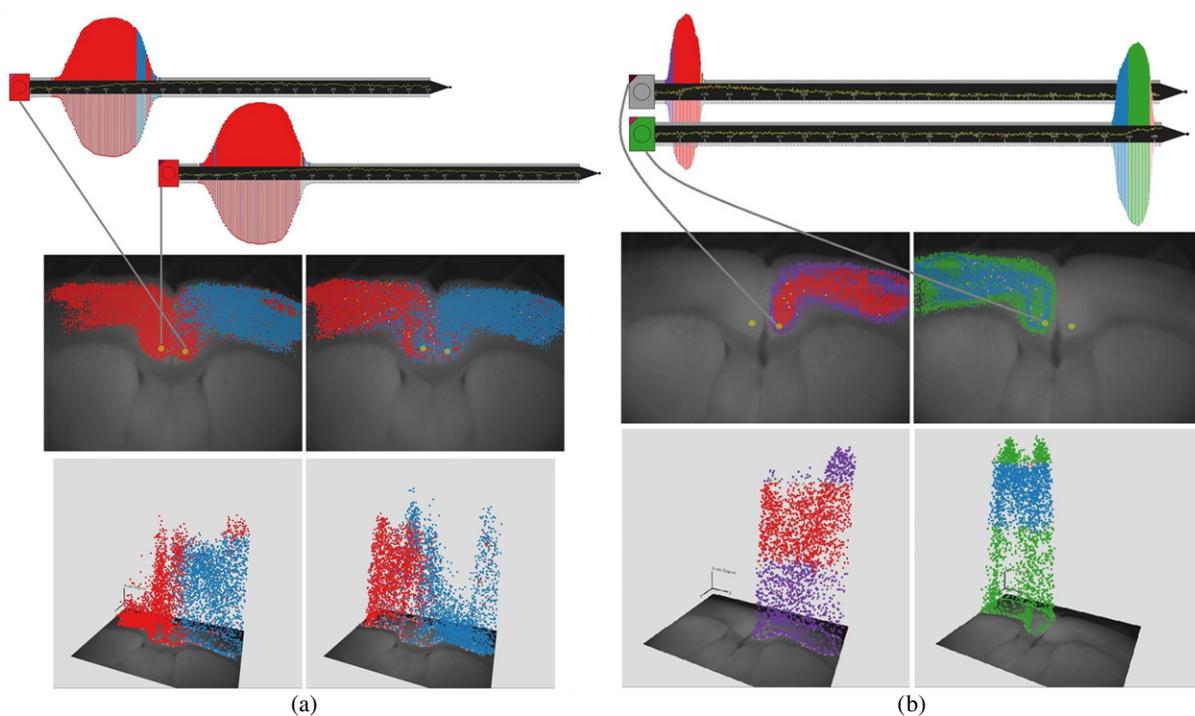


Figure 8. The behaviors of two nodes in the coronal slices of the mouse brain (a) before and (b) after the corpus callosum cut are displayed using the SwordPlots, with the snapshots of 2D and 3D views at two time steps.

Case Study I

For the first case study we looked at a coronal slice preparation of a mouse brain. The data contain mirror-image connectivity where the two cerebral hemispheres are connected by a single easily manipulable conduit of information flow: the corpus callosum. One hemisphere of the brain is stimulated, and the activation of correlated areas in the contralateral hemisphere is measured. To assess whether the dynamic communities identified by CommDy correspond to known neuronal networks, after trans-hemispheric networks are characterized, the corpus callosum is cut using microscissors. We then optimize the CommDy analytical parameters based on known changes to the network induced by callosotomy. All data are collected at 70 frames per second using a $2.5\times$ objective and a Retiga Exi camera with 8×8 hardware binning and StreamPix software for image collection. This produces pixel sizes of approximately $24\times 24\ \mu\text{m}$.

Figure 8 shows, in an *in-vitro* preparation, an example of the behaviors of two nodes at analogous locations across the two hemispheres in (a) the pre-cut and (b) the post-cut conditions. From the SwordPlots in Fig. 8(a), we find that the two nodes have similar trends of raw pixel value, similar distributions of node degrees in time, and the same community identifications (red) during most of the time. This indicates the similar behaviors of the two nodes across the two hemispheres. From the observation of the image slice view using animation, we see the partially symmetrical structure of the red community in both hemispheres where the two nodes are located. We then find more nodes

with such similar patterns in that region using SwordPlots, illustrating the connection between the left and right sides of the mouse brain. The Space-Attribute Cube is used to explore the community structure, which reveals some information that cannot be retrieved from the image slice view. For example, we may not be able to see any difference between two time steps through the image slice views, but we can obviously see the variation in node degree using the Space-Attribute Cube.

After the corpus callosum is cut (see Fig. 8(b)), we can easily see that the activation in the two hemispheres happens at different times via the SwordPlots, and moreover that there is no community across the two hemispheres. Since there is a long period of silence between two activations, animation is not an efficient way to identify when the activation occurs. However, we can target the interesting time periods directly by using the SwordPlot view. The views of the Space-Attribute Cube also show an interesting phenomenon, where a community (purple on the right side of the brain or green on the left side) has a distinctly different distribution in node degree, which causes the ring structures shown in the image slice views. One explanation is that nodes in the outer ring have strong connections with the nodes in the upper right region, but few connections with each other.

Case Study II

For our second case study, we investigated the impact of aging and/or peripheral hearing loss on auditory cortex (AC) activity. We examined network activity in the AC in slice preparations^{20,64} taken from young and aged mice. We used

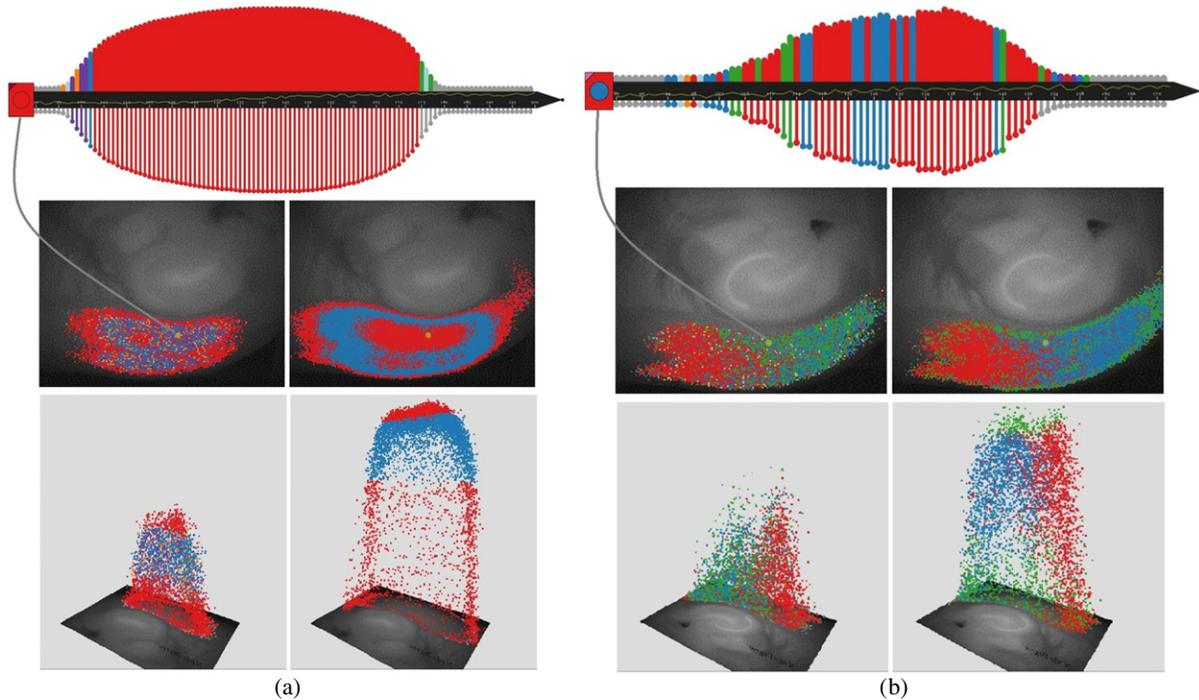


Figure 9. The behavior of one node in the auditory thalamocortical slices of (a) a young mouse and (b) an aged mouse is displayed using the SwordPlots, with the snapshots of 2D and 3D views at two time steps.

the auditory thalamocortical slice, which contains large areas of the AC. For this study we had two main hypotheses. 1. Aging is associated with diminished network associations within the AC. 2. Changes in network activity with aging are caused, at least in part, by peripheral hearing loss. Using the SwordPlots visualization tool, we were able to assess the likelihood of the validation of these hypotheses.

In a different animal, using an *in-vivo* preparation, Figure 9(a) shows the behavior of a node in the AC in a young mouse. Its SwordPlot indicates its consistency within the red community during the active period. We can also see an outer ring structure in this data set. However, the SwordPlot of a node at a similar location in Fig. 9(b) of an aged mouse shows frequent changes in colors within the sword cross-guard area as well as inconsistencies between the upper and lower cross-guard colors, indicating a high switchingness and a low consistency. This demonstrates the instability of the aged mouse brain.

We also found contrasting community structures in the young and aged mice. In the young mouse, the red community is mainly in the center of the AC surrounded by the blue one, while in the aged mouse the red community stays in the left part of the AC and the green/blue one is located on the right side. The views of the Space-Attribute Cubes also show that the young mouse has only one peak of the node degrees, located in the center red community, but the aged mouse has two peaks that are apparent on the left (the red community) and right sides (the green/blue community). These differences between the young and aged mice, shown clearly in the SwordPlots application, provide an initial validation for our hypotheses.

EVALUATION

To further evaluate our approach, we also interviewed two graduate students who each have extensive experience in visualizing functional neuroimaging data. The first is a graduate student with over five years' experience working with MRI machines in a Speech and Hearing laboratory at a research university. The second is a graduate student in computer science who is working with a laboratory in the Department of Psychiatry at another university to develop techniques to analyze properties of the human connectome. After a brief introduction to the system and an explanation of its features and primary visual encodings, we conducted a cognitive walkthrough, asking each of them to carry out tasks using an example data set. During the walkthrough we encouraged them to ask questions and to give us feedback. We were interested primarily in the qualitative response of experts who were not familiar with the specific design goals of the technique, and to see what its utility might be to experts in a different area, but who make use of similar data.

The feedback was generally positive; they both stated that they liked using the system and that they found the visualization approach interesting and potentially useful. One of them said that the SwordPlot representation provides an “easy way” to understand the “changing of community identifications,” and that it “packages a lot of information well.” However, they were split on the novelty of SwordPlot's visual encodings. One said that he had not seen anything like it, while the other said that “the combination of timelines is not novel,” but admitted that it did seem like it would “give an intuitive and effective way” for him to explore his data. One participant mentioned that he liked the way in which

raw data were linked to the community identifications and node degrees, and noted that it could be useful for revealing potential correlations between those attributes.

In addition to these general comments, each of the participants provided more detailed feedback about some of our specific design choices. One of the participants wanted to know why we limited users to the selection of a maximum of ten simultaneous SwordPlots. We acknowledged that, due to limits of screen space, there is a tradeoff between the ability to show detailed information and displaying more models, which could enable a user to find similar patterns across multiple nodes. The participant then agreed that more than five or six would probably provide an overwhelming amount of information and would be difficult to perceive. The other participant said that he would rather have more detailed information provided as he could not imagine the need to ever focus on comparing more than two nodes at the same time. One of the participants liked the recommendation operation, as it provides a useful starting point for investigating the data. He told us that “it would be nice to have the top ten least connected nodes available as well,” and that he would like to have the ability to choose a node with a specific rank. The other participant told us that he liked the ability to apply a combination of filters, which seemed like a helpful way to discover patterns, for example, to see how certain nodes are sensitive to a certain feature given a fixed range of other features.

Both of the participants indicated that they found the 3D view informative. One of them told us that it would be “totally useful” to see the coordinated image slice view and SwordPlots together. Both of the participants opened the 3D view of the Space–Attribute Cube after they noticed interesting patterns in the image slice view. One told us that they would like to have the ability to interact further with the 3D view of the Space–Attribute Cube. He also wanted to know how the Space–Attribute Cube, and our tool in general, would handle volumetric data rather than image slices. Overall, both of the experts told us that they saw this as a compelling tool overall, and indicated that they would be interested in trying it out on their own data sets.

DISCUSSION AND FUTURE WORK

The two case studies indicated that our interactive multi-view visualization application with animation can be of significant help in the exploration of neuron behavior within dynamic communities of mouse brain networks. Further user studies will be necessary to document the degree to which this visualization can be of help. While the individual visual encoding is not novel, the combination of visual encoding in a tool to handle spatial and non-spatial information with the application of dynamic network analysis to neuroscience research has not been performed before. Our use of the multiple coordinated views through interactive filtering and color-coding provided us with new insight into this domain problem. In addition, it also steered our investigation, and allowed us to identify changes in analysis parameters that aided our visualization of these data

sets. Future work will include incorporating an option to vertically align the SwordPlots to make comparison tasks easier. We also plan to more clearly indicate links between each SwordPlot and its corresponding node in the image. Furthermore, we are exploring ways to extend the scalability of the application, such as by providing an option to condense the SwordPlots by turning off either the upper or the lower part of the cross-guard if it is not currently required.

CONCLUSION

In this article, we presented SwordPlots, an interactive application for visualizing neuron behavior within dynamic communities of brain networks. The design of the SwordPlots application coordinates multiple visualization approaches, including the use of animation to show change over time, the introduction of a novel 3D representation called the Space–Attribute Cube, and the SwordPlot itself, which provides details about individual neuron behavior. We tested the application using real-world image data on mouse brain networks, showing that it aids domain experts in fulfilling a range of visual analysis tasks, and we showed that the application shows promise in other domains through a qualitative study with two experts in neuroimaging. While the response from the neuroscientists and the neuroimagers was encouraging, we plan to improve the application both to better fulfill the needs of the neuroscientists that we collaborated with and to make it a more generally useful tool for researchers with different types of dynamic data. Specifically, we expect that our visualization would be useful for any data set whose elements have a meaningful spatial location and that can belong to and change communities, such as data sets collected from studying animal social interactions.

REFERENCES

- 1 K. M. Carley, *Dynamic Network Analysis* (Citeseer, 2003).
- 2 A. Maries, N. Mays, M. Hunt, K. F. Wong, W. Layton, R. Boudreau, C. Rosano, and G. E. Marai, “Grace: a visual comparison framework for integrated spatial and non-spatial geriatric data,” *IEEE Trans. Vis. Comput. Graphic* **19**, 2916–2925 (2013).
- 3 B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde, “Functional connectivity in the motor cortex of resting human brain using echoplanar MRI,” *Magn. Reson. Med.* **34**, 537–541 (1995).
- 4 M. D. Fox, M. Corbetta, A. Z. Snyder, J. L. Vincent, and M. E. Raichle, “Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems,” *Proc. Natl. Acad. Sci. USA* **103**, 10046–10051 (2006).
- 5 J. Vincent, G. Patel, M. Fox, A. Snyder, J. Baker, D. Van Essen, J. Zempel, L. Snyder, M. Corbetta, and M. Raichle, “Intrinsic functional architecture in the anaesthetized monkey brain,” *Nature* **447**, 83–86 (2007).
- 6 V. Calhoun, T. Adali, V. McGinty, J. Pekar, T. Watson, and G. Pearlson, “fMRI activation in a visual-perception task: network of areas detected using the general linear model and independent components analysis,” *NeuroImage* **14**, 1080–1088 (2001).
- 7 C. F. Beckmann, M. DeLuca, J. T. Devlin, and S. M. Smith, “Investigations into resting-state connectivity using independent component analysis,” *Phil. Trans. R. Soc. Lond. B* **360**, 1001–1013 (2005).
- 8 E. Bullmore and O. Sporns, “The economy of brain network organization,” *Nature Rev. Neurosci.* **13**, 336–349 (2012).
- 9 T. Berger-Wolf, C. Tantipathananandh, and D. Kempe, “Dynamic community identification,” *Link Mining: Models, Algorithms, and Applications* (Springer, 2010), pp. 307–336.

- 10 P. Holme and J. Saramäki, *Temporal Networks* (Springer, 2013).
- 11 M. Abeles and G. L. Gerstein, "Detecting spatiotemporal firing patterns among simultaneously recorded single neurons," *J. Neurophysiol.* **60**, 909–924 (1988).
- 12 M. Girvan and M. E. Newman, "Community structure in social and biological networks," *Proc. Natl. Acad. Sci. USA* **99**, 7821–7826 (2002).
- 13 M. E. Newman and M. Girvan, "Finding and evaluating community structure in networks," *Phys. Rev. E* **69**, 026113 (2004).
- 14 V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *J. Stat. Mech. Theory Exp.* **2008**, P10008 (2008).
- 15 C. Tantipathananandh, T. Berger-Wolf, and D. Kempe, "A framework for community identification in dynamic social networks," *Proc. 13th ACM SIGKDD Int'l. Conf. on Knowledge Discovery and Data Mining* (ACM, 2007), pp. 717–726.
- 16 C. Tantipathananandh and T. Berger-Wolf, "Constant-factor approximation algorithms for identifying dynamic communities," *Proc. 15th ACM SIGKDD Int'l. Conf. on Knowledge Discovery and Data Mining* (ACM, 2009), pp. 827–836.
- 17 C. Tantipathananandh and T. Y. Berger-Wolf, "Finding communities in dynamic social networks," *IEEE 11th Int'l. Conf. on Data Mining (ICDM)* (Piscataway, NJ, 2011), pp. 1236–1241.
- 18 L. Backstrom, D. Huttenlocher, J. Kleinberg, and X. Lan, "Group formation in large social networks: membership, growth, and evolution," *Proc. 12th ACM SIGKDD Int'l. Conf. on Knowledge Discovery and Data Mining* (ACM, 2006), pp. 44–54.
- 19 S. Wasserman, *Social Network Analysis: Methods and Applications*, Vol. 8 (Cambridge University Press, 1994).
- 20 D. A. Llano, B. B. Theyel, A. K. Mallik, S. M. Sherman, and N. P. Issa, "Rapid and sensitive mapping of long-range connections in vitro using flavoprotein autofluorescence imaging combined with laser photostimulation," *J. Neurophysiol.* **101**, 3325–3340 (2009).
- 21 K. Shibuki, R. Hishida, H. Murakami, M. Kudoh, T. Kawaguchi, M. Watanabe, S. Watanabe, T. Kouuchi, and R. Tanaka, "Dynamic imaging of somatosensory cortical activity in the rat visualized by flavoprotein autofluorescence," *J. Physiol.* **549**, 919–927 (2003).
- 22 S. Kairam, D. MacLean, M. Savva, and J. Heer, "Graphprism: compact visualization of network structure," *Proc. Int'l. Working Conf. on Advanced Visual Interfaces* (ACM, 2012), pp. 498–505.
- 23 G. A. Pavlopoulos, A.-L. Wegener, and R. Schneider, "A survey of visualization tools for biological network analysis," *BioData Min* **1**, 12 (2008).
- 24 P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage, N. Amin, B. Schwikowski, and T. Ideker, "Cytoscape: a software environment for integrated models of biomolecular interaction networks," *Genome Res.* **13**, 2498–2504 (2003).
- 25 M. E. Smoot, K. Ono, J. Ruscheinski, P.-L. Wang, and T. Ideker, "Cytoscape 2.8: new features for data integration and network visualization," *Bioinformatics* **27**, 431–432 (2011).
- 26 C.-Y. Lin, K.-L. Tsai, S.-C. Wang, C.-H. Hsieh, H.-M. Chang, and A.-S. Chiang, "The neuron navigator: exploring the information pathway through the neural maze," *IEEE Pacific Visualization Symposium (PacificVis)* (IEEE, Piscataway, NJ, 2011), pp. 35–42.
- 27 J. Sorger, K. Buhler, F. Schulze, T. Liu, and B. Dickson, "Neuromap: interactive graph-visualization of the fruit fly's neural circuit," *IEEE Symposium on Biological Data Visualization (BioVis)* (IEEE, Piscataway, NJ, 2013), pp. 73–80.
- 28 M. Hadwiger, A. Al-Awami, J. Beyer, H. Strobelt, H. Pfister, N. Kasthuri, and J. Lichtman, "Neurolines: a subway map metaphor for visualizing nanoscale neuronal connectivity," *IEEE Trans. Vis. Comput. Graphics* **1** (2014).
- 29 K. Li, L. Guo, C. Faraco, D. Zhu, H. Chen, Y. Yuan, J. Lv, F. Deng, X. Jiang, and T. Zhang, "Visual analytics of brain networks," *NeuroImage* **61**, 82–97 (2012).
- 30 R. Jianu, C. Demiralp, and D. H. Laidlaw, "Exploring brain connectivity with two-dimensional neural maps," *IEEE Trans. Vis. Comput. Graphics* **18**, 978–987 (2012).
- 31 A. Irimia, M. C. Chambers, C. M. Torgerson, and J. D. Van Horn, "Circular representation of human cortical networks for subject and population-level connectomic visualization," *NeuroImage* **60**, 1340–1351 (2012).
- 32 D. S. Margulies, J. Böttger, A. Watanabe, and K. J. Gorgolewski, "Visualizing the human connectome," *NeuroImage* **80**, 445–461 (2013).
- 33 E. G. Christodoulou, V. Sakkalis, V. Tsiaras, and I. G. Tollis, "Brainnetvis: an open-access tool to effectively quantify and visualize brain networks," *Comput. Intell. Neurosci.* **2011**, 2011.
- 34 M. Xia, J. Wang, and Y. He, "Brainnet viewer: a network visualization tool for human brain connectomics," *PLoS ONE* **8**, e68910 (2013).
- 35 B. Alper, B. Bach, N. Henry Riche, T. Isenberg, and J.-D. Fekete, "Weighted graph comparison techniques for brain connectivity analysis," *Proc. SIGCHI Conf. on Human Factors in Computing Systems* (ACM, 2013), pp. 483–492.
- 36 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," *Proc. SPIE* **9011** (2014), p. 90110I.
- 37 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," *Proc. Eurographics Conf. on Visualization (EuroVis 15)—Short Papers* (2015), pp. 73–77.
- 38 F. Beck, M. Burch, S. Diehl, and D. Weiskopf, "The state of the art in visualizing dynamic graphs," *Proc. Eurographics Conf. on Visualization (EuroVis 14)—State of The Art Reports* (2014).
- 39 F. Beck, M. Burch, and S. Diehl, "Towards an aesthetic dimensions framework for dynamic graph visualisations," *IEEE 13th Int'l. Conf. on Information Visualisation* (IEEE, Piscataway, NJ, 2009), pp. 592–597.
- 40 F. Beck, M. Burch, and S. Diehl, "Matching application requirements with dynamic graph visualization profiles," *IEEE 17th Int'l. Conf. on Information Visualisation (IV)* (IEEE, Piscataway, NJ, 2013), pp. 11–18.
- 41 D. Archambault, H. C. Purchase, and B. Pinaud, "Animation, small multiples, and the effect of mental map preservation in dynamic graphs," *IEEE Trans. Vis. Comput. Graphics* **17**, 539–552 (2011).
- 42 C. Erten, P. J. Harding, S. G. Kobourov, K. Wampler, and G. Yee, "Graphael: graph animations with evolving layouts," *Graph Drawing* (Springer, 2004), pp. 98–110.
- 43 Y. Frishman and A. Tal, "Dynamic drawing of clustered graphs," *IEEE Symposium on Information Visualization (INFOVIS)* (IEEE, Piscataway, NJ, 2004), pp. 191–198.
- 44 Y. Frishman and A. Tal, "Online dynamic graph drawing," *IEEE Trans. Vis. Comput. Graphics* **14**, 727–740 (2008).
- 45 B. Bach, E. Pietriga, and J.-D. Fekete, "Graphdiaries: animated transitions and temporal navigation for dynamic networks," *IEEE Trans. Vis. Comput. Graphics* **20**, 740–754 (2014).
- 46 M. G. Beiró, J. R. Busch, and J. I. Alvarez-Hamelin, "Visualizing communities in dynamic networks," *LAWDN-Latin-American Workshop on Dynamic Networks* (2010), 4-p.
- 47 L. Shi, C. Wang, and Z. Wen, "Dynamic network visualization in 1.5 d," *Pacific Visualization Symposium (PacificVis)* (IEEE, Piscataway, NJ, 2011), pp. 179–186.
- 48 M. Trier and A. Bobrik, "Analyzing the dynamics of community formation using brokering activities," *Communities and Technologies 2007* (Springer, 2007), pp. 463–477.
- 49 M. Burch, C. Vehlou, F. Beck, S. Diehl, and D. Weiskopf, "Parallel edge splatting for scalable dynamic graph visualization," *IEEE Trans. Vis. Comput. Graphics* **17**, 2344–2353 (2011).
- 50 M. Burch, F. Beck, and D. Weiskopf, "Radial edge splatting for visualizing dynamic directed graphs," *GRAPP/IVAPP* (2012), pp. 603–612.
- 51 S. Havre, B. Hetzler, and L. Nowell, "Themeriver: visualizing theme changes over time," *IEEE Symposium on Information Visualization (InfoVis)* (IEEE, Piscataway, NJ, 2000), pp. 115–123.
- 52 S. Havre, E. Hetzler, P. Whitney, and L. Nowell, "Themeriver: visualizing thematic changes in large document collections," *IEEE Trans. Vis. Comput. Graphics* **8**, 9–20 (2002).
- 53 M. Greilich, M. Burch, and S. Diehl, "Visualizing the evolution of compound digraphs with timearctrees," *Computer Graphics Forum* (Wiley Online Library, 2009), Vol. 28, pp. 975–982.

- ⁵⁴ K. Reda, C. Tantipathananandh, A. Johnson, J. Leigh, and T. Berger-Wolf, "Visualizing the evolution of community structures in dynamic social networks," *Computer Graphics Forum* (Wiley Online Library, 2011), Vol. 30, pp. 1061–1070.
- ⁵⁵ C. Vehlou, F. Beck, P. Auwärter, and D. Weiskopf, "Visualizing the evolution of communities in dynamic graphs," *Computer Graphics Forum* (Wiley Online Library, 2014).
- ⁵⁶ P. Gatalisky, N. Andrienko, and G. Andrienko, "Interactive analysis of event data using space-time cube," *Proc. 8th Int'l. Conf. on Information Visualisation, 2004. IV 2004* (IEEE, Piscataway, NJ, 2004), pp. 145–152.
- ⁵⁷ G. Groh, H. Hanstein, and W. Wörndl, "Interactively visualizing dynamic social networks with dyson," *Proc. of the IUI09 Workshop on Visual Interfaces to the Social and the Semantic Web* (2009), Vol. 2.
- ⁵⁸ K. Reda, C. Tantipathananandh, T. Berger-Wolf, J. Leigh, and A. Johnson, "Socioscape—a tool for interactive exploration of spatio-temporal group dynamics in social networks," *Proc. IEEE Information Visualization Conf. (INFOVIS)* (IEEE, Piscataway, NJ, 2009).
- ⁵⁹ C. Forlines and K. Wittenburg, "Wakame: sense making of multidimensional spatial-temporal data," *Proc. Int'l. Conf. on Advanced Visual Interfaces* (ACM, 2010), pp. 33–40.
- ⁶⁰ B. Bach, E. Pietriga, and J.-D. Fekete, "Visualizing dynamic networks with matrix cubes," *Proc. 32nd Annual ACM Conf. on Human Factors in Computing Systems* (ACM, 2014), pp. 877–886.
- ⁶¹ D3 Color Ordinal Scales, <https://github.com/mbostock/d3/wiki/Ordinal-Scales>.
- ⁶² M. Bostock, V. Ogievetsky, and J. Heer, "D³ data-driven documents," *IEEE Trans. Vis. Comput. Graphics* **17**, 2301–2309 (2011).
- ⁶³ Cynthia Brewer, <http://colorbrewer2.org/>.
- ⁶⁴ D. A. Llano, J. Turner, and D. M. Caspary, "Diminished cortical inhibition in an aging mouse model of chronic tinnitus," *J. Neurosci.* **32**, 16141–16148 (2012).