

BrainCove: A Tool for Voxel-wise fMRI Brain Connectivity Visualization

A.F. van Dixhoorn¹, J. Milles², B. van Lew² and C.P. Botha^{1,2}

¹Computer Graphics and Visualization, Dept. of INSY, Delft University of Technology, Delft

²LKEB, Dept. of Radiology, Leiden University Medical Center, Leiden
The Netherlands

Abstract

Functional brain connectivity from fMRI studies has become an important tool in studying functional interactions in the human brain as a complex network. Most recently, research has started focusing on whole brain functional networks at the voxel-level, where fMRI time-signals at each voxel are correlated with every other voxel in the brain to determine their functional connectivity. For a typical 4mm isotropic voxel resolution, this results in connectivity networks with more than twenty thousand nodes and over 400 million links. These cannot be effectively visualized or interactively explored using node-link representations, and due to their size are challenging to show as correlation matrix bitmaps. In this paper, we present a number of methods for the visualization and interactive visual analysis of this new high resolution brain network data, both in its matrix representation as well as in its anatomical context. We have implemented these methods in a GPU raycasting framework that enables real-time interaction, such as network probing and volume deformation, as well as real-time filtering. The techniques are integrated in a visual analysis application in which the different views are coupled, supporting linked interaction. Furthermore, we allow visual comparison of different brain networks with side-by-side and difference visualization. We have evaluated our approach via case studies with domain scientists at two different university medical centers.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—functional brain connectivity, brain mapping, raycasting, GPU volume rendering

1. Introduction

With functional MRI (fMRI) connectivity, the functional connections between different parts of the brain can be measured non-invasively, *in vivo* and in 3D, down to the voxel level. This can be done during the performance of a task, in order to determine the brain networks involved in completing the task, or during resting state, in order to shed light on the intrinsic connectivity networks of the brain.

Functional brain connectivity has already proven to be a valuable tool for research in areas related to cognitive psychology, neuroscience and behavioral studies. Traditional approaches in fMRI connectivity research used a seed-based approach, where only the brain regions connected to a selected seed region are derived, or independent component analysis to describe the functional connectivity networks. Recently, researchers have begun to focus on whole-brain networks, applying concepts from graph theory that enable

more complete studies of brain networks than the aforementioned traditional methods. The first studies focused on inter-regional connectivity, where the properties of the brain network were explored by measuring the connectivity between all anatomical brain regions, such as the 90 cortical and sub-cortical regions of the AAL template [BYHH95]. The resulting networks can be visualized effectively with matrix bitmaps or node-link diagrams [DVF*10]. More recently, research has also started to focus on functional brain connectivity at the voxel level [vdHSBHP08]. The resulting connectivity networks are several orders of magnitude larger than the region-based connectivity networks. In a typical 4mm isotropic resolution, the raw BOLD-fMRI image contains about 20,000 voxels, and the resulting network thus consists of 20,000 nodes and 400,000,000 links (including symmetrical links). The interactive visualization of such large connectivity networks using traditional matrix visualizations is computationally challenging, and using node-link

diagrams to represent the network is not really feasible. We present methods for visualizing large brain networks both as a highly interactive matrix representation as well as in their anatomical context. The implementation uses a raycasting framework that enables interactive exploration of the data. One of the unique aspects of our technique is the side-by-side coupled visualization of two of these voxel-based brain networks, enabling their direct visual comparison. Furthermore, we employ a flat-map representation for showing the connectivity data in spatial context with minimal occlusion, as well as real-time correlation volume splitting to enable visualization of and interaction also with interior volumes of the brain between the two lobes.

The contributions of this paper are: 1.) We present a technique with which large voxel-based fMRI connectivity matrices of around twenty-thousand by twenty-thousand correlations can be interactively visualized on a desktop PC, both directly and in their anatomical context. 2.) We introduce a method that allows for the interactive visual comparison of multiple of these large connectivity matrices in a side-by-side or difference visualization, which, to the best of our knowledge, has not been shown before. 3.) We evaluate our approach by performing a case study with two independent groups of domain scientists. 4.) Our complete implementation is available under a permissive open source license[†].

visualization is the design of a technique that is able to render at interactive speeds, enabling the user to interact with the data in real-time. To accomplish this, we utilize the GPU architecture and the increasingly greater amounts of texture memory available on recent graphics cards. The rest of this paper is organized as follows: Section 2 presents related work done on this topic. Section 3 presents the method starting with a general overview. Section 4 discusses some implementation details and technical challenges that were addressed. In Section 5, we evaluate our method, including feedback from expert users, structured according to a case study evaluation. Section 6 completes this paper with conclusions and future work.

2. Related work

The visualization of region-wise functional connectivity networks is most commonly done with pixmaps that directly represent the correlation matrix, or using node-link diagrams. The pixmap is a pixel-based representation that adheres to the layout of the raw correlation matrix, directly mapping each correlation value to a color using a pre-defined color scale. For an $N \times N$ connectivity matrix, this results in a $N \times N$ bitmap image. For effective visualization, the pixmap should be reordered such that similar items are grouped [ME03]. For functional connectivity brain networks, the ordering is typically derived from anatomical lo-

cation, such as grouping voxels together if they are in the same anatomical region or brain lobe [DVF*10, HCG*08] or from hierarchical clustering (in which the leaves of the dendrogram are used for the ordering), such that the pixmap groups highly connected hubs together [HWG*11].

To see the functional network in its spatial context, the correlation matrix is typically represented as a node-link diagram, inter-connecting the N nodes with a straight line, whose thickness or color is based on the connectivity strength. The visual analysis of region-wise whole-brain functional connectivity networks has been studied before by Van Dixhoorn et al. [DVF*10], where the problem of visual clutter that arises when rendering node-link networks with twenty or more nodes [GFC05] was addressed by allowing the user to interactively filter on the connection strength. The node-link representation has also been used to visualize voxel-wise connectivity by Zuo et al. [ZEM*11], but here the links are drawn between twenty functional communities instead of between each voxel pair. Furthermore, the visualization procedures were carried out on a graphics workstation, rather than on a standard desktop computer.

Instead of the node-link representation, we employ a method typically used to visualize brain activation data from fMRI studies. An approach for this was described by Jainek et al. [JBB*08], where illustrative techniques are used to visualize functional data in anatomical context. To represent the activation data, the metaphor is used of activated regions emitting light. However, instead of a single activation map, our data consists of a large number of networks, one for each voxel in the underlying fMRI images. Rendering all these networks at once in would result in an ambiguous visualization. Instead, our method includes an interaction component in which the user indicates in which network he is interested by interactively selecting a seed voxel. Our method certainly shows parallels with the work recently published by Eklund et al. [EFAK11] and Böttger et al. [BMH*11], as well as with the interactive tool InstaCorr for the visualization of functional connectivity in AFNI [RW11]. The method of Böttger et al. allows the user to place a cross-hair on the desired seed voxel on orthogonal 2D slices of an anatomical scan, rendering the resulting correlation map on top as an overlay. The tools presented by Eklund et al. and InstaCorr provide similar functionality, but in addition they are able to visualize the correlation maps in their 3-D spatial layout. Our work is similar to those methods, but differs with respect the following points: 1.) The correlation matrix is computed in a pre-processing step such that we can provide a pixmap visualization that allows for detection of groups of voxels that are correlated. 2.) We provide a picking tool that allows the user to interactively and dynamically select a seed voxel on the cortical surface, directly in the 3-D representation. 3.) Our tool allows for interactive side-by-side comparison and difference visualization of multiple datasets.

[†] <https://bitbucket.org/avandixhoorn/braincove>

3. Method

The tool is called *BrainCove*. Its main interface is shown in Fig. 1. It contains one or more child windows, each of which consists of three main components: the correlation matrix at the top left, the orthogonal slice views at the top right and the anatomical visualization spanning the full width at the bottom. In this case, two datasets have been loaded and can be compared by use of linked interaction.

The input to our application is a pre-processed correlation matrix in raw binary format, combined with a list of coordinates in the well-known standardized MNI space that correspond to the elements in the correlation matrix.

3.1. Data pre-processing

Our pre-processing follows Ferrarini et al. [FVvL*11] and includes motion-correction, removal of non-brain tissue, grand mean intensity normalization, registration to MNI-152 standard space and downsampling to 4mm isotropic resolution.

After pre-processing, the correlation matrix is calculated. First, white matter (WM) and cerebrospinal fluid (CSF) time series and the average whole-brain signal are extracted from the fMRI data sets. The data is then masked with the Automated Anatomical Labeling (AAL, [TLP*02]) atlas and using regression analysis the influence of artifacts is reduced [FVvL*11]. The final correlation matrix is computed by calculating Pearson's correlation coefficient between each pair of neural activation timeseries. This raw correlation map and a list of MNI coordinates corresponding to the rows and columns of the matrix are the input to our application.

3.2. Pixmap view

The pixmap representation is a direct visualization of the correlation matrix, where the correlation values are mapped to colors using a continuous colorscale. The elements on the rows and columns are reordered so that anatomically close voxels are grouped. This reordering is computed by looking up the AAL atlas region for each voxel, and then grouping according to the atlas region indices. In a second pass, voxels are further grouped by hemisphere.

Directly mapping each correlation in the matrix to a colored pixel results in a bitmap of 400 megapixels. If each color is to be represented with three 8-bit channels, the complete bitmap would require 1.2 gigabytes of storage on top of the 800 megabytes that is required to store the raw correlation matrix. In order to display such an image on the screen at interactive frame rates, we employ GPU raycasting.

A plane is used as a proxy object to convert from screen coordinates to indices in the correlation matrix. Since there are only values on the plane, the raycasting algorithm is relatively simple, as no ray-marching is required. The algorithm




Figure 2: *Brushing the matrix representation highlights the selected voxels in the anatomical view, where it can be seen from any view. The brushed selection appears twice because of the symmetry in the matrix. Voxels along the horizontal axis of the matrix are shown in yellow, voxels on the vertical axes in red.*

is implemented in OpenCL, with the entire correlation matrix uploaded to a GPU buffer as a float array. For each view ray, the correlation value is extracted and mapped to its final pixel color.

3.2.1. Interaction and Filtering

Zooming and panning of the correlation matrix are highly efficient, as this is built-in in the transformation pipeline. In addition, applying a filter is nothing more than an extra statement in the raycasting algorithm, which means that filtering can be done on-the-fly. The tool currently only supports thresholding on absolute correlation as a filter, but other filters can be implemented easily.

To improve the integration of functional connectivity and spatial location, we linked the pixmap representation to a 3-D anatomical visualization. By hovering with the cursor over the pixmap, the 3-D visualization highlights in real-time the voxels corresponding to indicated connection. Furthermore, we implemented a *brushing* technique that allows the user to select a group of links using the mouse. The corresponding voxels are then highlighted in the anatomical view (see Fig. 2).

3.3. Anatomical view

In the anatomical view, the connectivity network is rendered in its spatial context. Due to the size and connectedness of the network, the node-link representation would not be suitable. Instead, we visualize in real-time the connection strengths of the whole brain to the voxel or region currently indicated. Correlations are mapped to a perceptually linear blue-to-yellow colorscale, with bright blue and bright yellow representing strong negative and strong positive correlations respectively. The correlation can also be interactively




Figure 1: An overview of the application window with two datasets. Each dataset is opened in its own child window. Each child window contains of three views: the pixmap view on the top-left, the slice views on the top right and the anatomical view on the bottom.




Figure 3: A seed voxel on the cortical surface can be selected by hovering with the mouse over the desired position.

thresholded, in which case only correlations higher than the threshold are colormapped. An example of this visualization is shown in Fig. 3.

This is also implemented using a GPU-based raycasting approach. Each frame is rendered using two stages. In the first stage, a correlation volume is computed in which each voxel contains its correlation with the currently selected seed voxel. This correlation volume is then rendered by the raycaster, together with a high resolution anatomical volume that provides context. Both volumes are stored in GPU texture memory and sent to the OpenCL raycasting kernel. Along with each volume, a transformation matrix is sent to the kernel that defines for each volume the mapping from world position to voxel position.

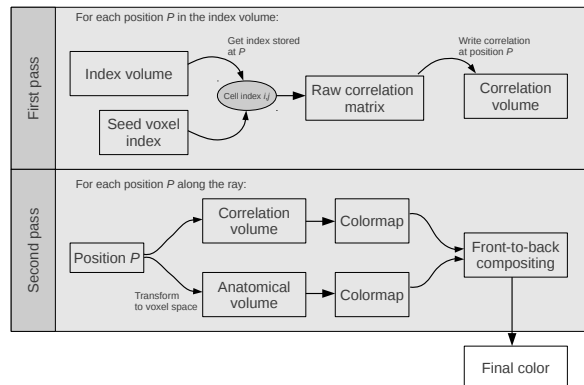


Figure 4: The pipeline for the two-pass raycasting of the correlation volume.

At each sample position along each view ray, the correlation volume is sampled. If the correlation value at the current sample position is below the threshold, or the sample position is out of the correlation volume bounds, the raycaster samples the anatomical volume instead. The value is then mapped to a color for the current voxel using two different transfer functions, one for the correlation volume and one for the anatomical volume. The resulting color values along the ray are then composited to form the final pixel color. The complete process is represented in Fig. 4.

3.3.1. Interaction

The 3-D window with the volume rendered correlation map allows for interactive rotation, zoom and pan, allowing the user to see the correlation from every angle.

Besides hovering over a single cortical voxel, the size of the seed region can also be increased, such that a group of voxels is used. In addition, it is also possible to use the AAL region to which the selected voxel belongs as a seed region. Using the ray casting selection, it is not possible to select voxels behind the outer cortical layer of the brain. To overcome this limitation, two alternative selection methods have been implemented.

Brain cleaving To be able to select structures that are located in-between the two hemispheres, such as the thalamus or the hippocampus, we employed a volume deformation method with which the user is able to spread the two brain hemispheres apart, similar to the Hinge Spreader proposed by McGuffin et al. [MTB03]. Once the two brain halves are spread apart, the user can use the mouse to select a voxel on the inner side of each hemisphere, in the same way as selecting a voxel on the cortical surface. The deformation is implemented in the raycasting algorithm using ray deformation, and uses a sagittal plane, passing through a point P that is located between the two hemispheres. This plane can be freely adjusted as well.

While stepping through the volume the algorithm determines for every sample point Q whether the point is to the left or to the right of the sagittal plane, using its plane normal N . Once the point Q has been classified as either left or right from the split plane, its transformed position is found by multiplying the voxel position with the corresponding rotation matrix for left or right rotation.

$$\vec{Q}' = \begin{cases} A_L * \vec{Q} & \text{if } (\vec{Q} - \vec{P}) \cdot \vec{N} < 0 \\ A_R * \vec{Q} & \text{if } (\vec{Q} - \vec{P}) \cdot \vec{N} > 0 \end{cases} \quad (1)$$

where A_L and A_R are the rotation matrices for the left and right hemispheres respectively. See Fig. 5 for a schematic representation of this method. When implementing this method in a raycasting algorithm, the process is actually reversed: the sample points visited by the raycaster should be considered to be already in deformed space. The sample value for that position can then be found by projecting the sample position back to non-deformed space using the inverse transformation. A special case occurs when the sample position is in the region between the two rotated hemispheres (the hatched region in Fig. 5). During the ray traversal, positions in this region are skipped.

Slice view To allow also for the easy selection of seed voxels at arbitrary positions in the brain, a view window is available that contains three orthogonal slice views. Using the mouse, a seed point can be selected on any of the three slices. The current voxel selection is linked between the slice views, the anatomical view and the pixmap view.




Figure 5: A schematic representation of the brain split method (left) and the resulting visualization (right). The hatched region in the left drawing indicates the region in which the raycaster can skip the voxels.




Figure 6: The Lambert's Cylindrical flatmap representation of the brain, viewing from the anterior in the middle to posterior at the two sides.

3.4. Flat mapping

The anatomical view described in the previous section shows the correlation map in its correct spatial context, but suffers from the occlusion problem inherent to any 3-D visualization. To address this problem, our tool can also show a flat map of the cerebral cortex.

We have implemented two cylindrical projections to flat map the cortex: Lambert's Cylindrical Equal-Area and Braun's Stereographic Cylindrical projection. Lambert's projection causes severe distortion in the poles, but has the advantage that the shape of the brain can be easily recognized. Braun's projection on the other hand, distributes the distortion over the full height of the map, but results in brain maps that are harder to recognize. Figure 6 shows an example.

4. Implementation

The visualization techniques described in the previous sections are implemented in a prototype application in C++, using Qt and VTK. The raycasting algorithm is implemented in OpenCL. The OpenCL raycaster renders to a texture that is shared with OpenGL using the OpenCL/GL interoperability. This texture is then mapped to a quad in the VTK render

window, a method commonly used in the VTK volume ray-cast mappers.

To enable real-time visualization, the raw correlation matrices are transferred to the GPU using an OpenCL buffer object. The size of a matrix depends on the resolution of the input fMRI volume. For a typical resolution of 4mm, the resulting correlation matrix is about 800-900 megabytes in size. Thus, to visualize two datasets concurrently, a graphics card with at least two gigabytes of onboard memory is required.

5. Case Study

The prototype application with the presented visualization techniques was evaluated with domain scientists in order to investigate the possible role of the application in the existing pipeline of fMRI connectivity research. The evaluation was set up as an exploratory case study following the guidelines set out by Yin [Yin09]. The main study question was formulated as: *How can the functional connectivity visualization tool, called BrainCove, assist domain scientists in studying patterns in functional brain connectivity, the relation with brain anatomy and in studying inter-subject or inter-group differences?* and the case was defined as the use of our application by external domain experts who were targeted as prospective end users. An evaluation session was held with two groups, one with a group of four neuroscientists from the Leiden University Medical Center (LUMC) and a second at the Amsterdam Medical Center (AMC), in which the tool was first presented to a group of 30 people from the neuroimaging and neuroscience domain in an informal and interactive presentation and then continued as case study with a smaller group of researchers (8 in total) specialized in fMRI connectivity.

In the following sections, we discuss the user feedback structured according to the case study propositions.

5.1. Matrix Visualization

The matrix visualization gives an overview of the data and allows for the detection of groups of voxels that are correlated. This proposition was confirmed by all participants. One user specifically mentioned that the matrix visualization is useful to quickly select peaks of highly correlated groups of voxels. Participants added that feedback about the ordering is important and that rendering of labels would make this representation more comprehensible, although it was also noted that the lack of labels is compensated by highlighting the corresponding voxels in the anatomical visualization.

The matrix visualization allows for a quick check on the quality of the data, such that errors can be identified before using the data further in the analysis pipeline. The use of the matrix visualization for quality checks was not directly apparent to the participants. They agreed that this representation could be used to detect large artifacts in the data, but

noted that these artifacts would have already been detected earlier in the analysis pipeline.

Being able to see the spatial context for correlations or groups of correlations using the view linking aids in interpreting the FC matrix. The participants agreed that the view linking helps in interpreting the FC matrix, even claiming that without the linked interaction, they would not have a clue on how to interpret the matrix. One of the users further stated that the matrix visualization would be most useful to find large scale differences between subjects. The linking would then help to see where the differences are in spatial context.

5.2. Anatomical Visualization

The visualization of FC in spatial context supports mental integration of FC and anatomy. During the session at the LUMC, the participants mentioned that they find it difficult to navigate in the current 3-D visualization, because of a lack of reference. Scientists usually navigate and orient themselves within the data using orthogonal slices of a structural brain or by manually typing in the MNI coordinates. They strongly suggested that we should integrate a set of simple orthogonal slice views linked to the 3-D anatomical view that would serve as an anatomical reference. With such an anatomical reference, they would confirm this proposition.

We implemented this before conducting the second session at the AMC, where the users confirmed this hypothesis, stating that this integration of functional connectivity and anatomy is an essential element.

The visualization of FC data in which the voxels emit "light" when they are functionally connected is an intuitive representation of the correlation maps. All users generally confirmed this proposition. It was remarked that this method corresponds to the metaphor used in other tools used in the field: "what lights up is active". Users from both groups noted that the colormap used to represent the correlation could be improved, such that the scaling is more diverse. They indicated that the current color map shows the correlation map too brightly, which makes it hard to see differences. One of the users further remarked that the color map was not warm enough. Participants in both groups further suggested to add a colormap legend in the scene and to make the range of the colormap adjustable.

Interactively selecting a seed voxel by hovering with the mouse over the voxel of interest facilitates the detection of interesting networks and abnormalities. This proposition was confirmed by all participants. They indicated that this interaction technique is the biggest difference with other tools. One participant stated that the tool could make a significant contribution to the procedure of selecting a seed voxel in seed-based analysis. The current method includes a priori selection of a seed voxel and computation of the correlations

with all other voxels. According to the participant, loading the whole-brain correlation matrix into our application would allow for better comprehension in selecting the seed voxel because the effect of choosing a specific seed voxel is immediately visualized.

The brain split approach is useful for selecting voxels inside the brain volume (for instance, in the cingulate cortex). The participants did not readily confirm this proposition. One participant from the first group remarked that he had trouble finding the cingulate between the brain lobes due to the lack of anatomical reference. The participants generally agreed that orthogonal slice views would be preferred for selecting voxels inside the brain volume. In the second session, the three orthogonal views with the MNI structural brain was considered a better technique for selecting a seed voxel in the brain volume, which confirmed the findings from the first session. Interestingly, one of the participants in the AMC group considered the brain split approach an effective method for reducing occlusion in the 3-D visualization, but not so much for probing.

Context visualization using a high resolution MRI head volume and a coloured and outlined anatomical atlas aids in relating FC to anatomical regions. The utility of the colored anatomical atlas was not directly apparent to the participants in the first session. One of the participants even noted that the coloring was more confusing than helpful. Again, the suggestion was made to integrate a linked view with orthogonal slices of a structural brain or a semi-transparent surface rendering of the brain (usually referred to as a “glass brain”) for providing anatomical context. Following these suggestions, we removed the anatomical atlas coloring from the volume rendering and integrated orthogonal slice views in the tool. During the second session, participants confirmed that the slices with a structural MNI brain are the preferred way for navigation purposes and provide sufficient anatomical context.

5.3. Flat map

Visualizing FC in a 2D projection of the spatial locations facilitates in forming a mental map of the complete connectivity network in a single view. Remarkably, researchers from both groups were not familiar with the use of “flat maps” for two-dimensional representation of functional connectivity in anatomical context. They generally found it difficult to orient themselves in the cylindrical flat map representation without structural anatomical context and mentioned that a learning curve would be involved to get accustomed to such a representation. One participant further commented that the presented flat map is limited to functional connectivity studies of the cortex, which makes it unsuitable for use in groups that focus on sub-cortical regions. In general, participants did, however, see potential in the use of mappings that are able to represent the complete connectivity network in a single view.

5.4. Visual Comparison

Coordinated visualizing multiple datasets side-by-side supports the finding of differences between subjects. All participants confirmed this proposition, noting that the visual comparison would be a powerful tool mainly for visually comparing single subjects or patients to a group mean, since this is mostly a visual task. This would also enable the use of resting state fMRI connectivity as a disease marker. One of the attending medical doctors further remarked that the visual comparison could also be employed in intervention measurement in a clinical setting. For example, patients with obsessive-compulsive disorder are increasingly being treated with deep brain stimulation (DBS). Being able to visualize the functional network of the brain pre-DBS and post-DBS side-by-side would allow clinicians to see changes in the network, which is helpful in judging whether the current treatment is effective or should be changed.

Visualizing the difference between two datasets supports the finding of differences between subjects or the impact of preprocessing on FC networks. The value of visualizing the absolute correlation difference between different subjects was not directly confirmed. It was remarked that absolute difference in correlation between subjects could be attributed to noise or to differences in the strength of the measured signal. Participants did see potential in using the difference visualization within one subject to compare the influence of using different preprocessing pipelines. However, participants did see the greatest potential in connecting the tool to a database with group averages such as for healthy subjects and for different pathologies, such that single subjects can be visually compared with the group average. Another suggestion was to create the average on-demand from a group study.

5.5. General remarks

The participants in the first group were generally impressed by the visualizations and saw the potential in the tool, but stated that due to the lack of anatomical reference (by means of structural data) and the inability to select voxels that are at a distance from the cortical surface, they would not readily use the tool for visual analysis. They suggested the use of a high resolution brain surface rendering, such as the one generally used in the SPM toolbox, or orthogonal slices of the MNI structural brain volume. Once the anatomical reference could be dealt with, they saw potential in using the tool especially to compare individual subjects to a group average, such as in comparing pathologies with healthy subjects with patients with different pathologies or with the effects of drugs on functional connectivity. They furthermore suggested to add a feature that makes it possible to select different types of connectivity (such as hub voxels) on-the-fly, using a pipeline that is running in the background and a method that enables the creation of high-quality pictures for publications.

The domain scientists in the second session were enthusiastic about the possibilities the tool offers for the visualization of the connectivity data. We attribute the difference in enthusiasm between the first and second evaluation to our addition of anatomical reference using the structural MNI slice views before conducting the second evaluation. Especially the visual comparison of individuals and group averages was considered an important contribution.

Both groups independently considered the ability to visually compare connectivity networks to be the major contribution of the tool.

6. Conclusions

In this paper, we presented a tool that couples a number of visualizations in order to facilitate the visual analysis of voxel-wise fMRI connectivity. Using our tool, the analyst is able to quickly identify interesting patterns in the functional network of the brain and differences in connectivity patterns between subjects or groups by visually comparing multiple datasets side-by-side. Currently, three different visualizations are implemented, including a pixmap representation and direct volume rendering of the correlation map for a given seed voxel in both anatomical context and a flat-map layout that shows the correlation map in pseudo-anatomical spacing. We evaluated our tool in case studies with groups of domain scientists at two different academic medical centers.

Currently, the flat-map representation uses cylindrical projections that result in a distorted projection. We plan to review other types of cortical maps that produce less distorted projections. On the longer term, functionality to calculate the connectivity measures on-the-fly will be added, such that the tool can directly read pre-processed 4D NIFTI files. Finally, we plan to extend the comparative visualization by allowing the import and on-the-fly generation of group mean datasets, that allows the analyst to visually compare individual patients with an overall group.

References

- [BMH*11] BÖTTGER J., MARGULIES D. S., HORN P., THOMALE U. W., PODLIPSKY I., SHAPIRA-LICHTER I., CHAUDHRY S. J., SZKUDLAREK C., MUELLER K., LOHMANN G., HENDLER T., BOHNER G., FIEBACH J. B., VILLRINGER A., VAJKOCZY P., ABBUSHI A.: A software tool for interactive exploration of intrinsic functional connectivity opens new perspectives for brain surgery. *Acta Neurochirurgica* 153, 8 (Apr. 2011), 1561–1572. 2
- [BYHH95] BISWAL B., YETKIN F. Z., HAUGHTON V. M., HYDE J. S.: Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine / Official Journal of the Society of Magnetic Resonance in Medicine* 34, 4 (Oct. 1995), 537–541. PMID: 8524021. 1
- [DVF*10] DIXHOORN A. F. V., VISSERS B. H., FERRARINI L., MILLES J., BOTHA C. P.: Visual analysis of integrated resting state functional brain connectivity and anatomy. In *VCBM* (2010), pp. 57–64. 1, 2
- [EFAK11] EKLUND A., FRIMAN O., ANDERSSON M., KNUTSSON H.: A GPU accelerated interactive interface for exploratory functional connectivity analysis of fMRI data. In *ICIP* (Brussels, Belgium, Sept. 2011). 2
- [FVvL*11] FERRARINI L., VEER I. M., VAN LEW B., OEI N. Y. L., VAN BUCHEM M. A., REIBER J. H. C., ROMBOUTS S. A. R. B., MILLES J.: Non-parametric model selection for subject-specific topological organization of resting-state functional connectivity. *NeuroImage* 56, 3 (June 2011), 1453–1462. PMID: 21338693. 3
- [GFC05] GHONIEM M., FEKETE J., CASTAGLIOLA P.: On the readability of graphs using node-link and matrix-based representations: a controlled experiment and statistical analysis. *Information Visualization* 4, 2 (July 2005), 114–135. 2
- [HCG*08] HAGMANN P., CAMMOUN L., GIGANDET X., MEULI R., HONEY C. J., WEDEEN V. J., SPORNS O.: Mapping the structural core of human cerebral cortex. *PLoS Biol* 6, 7 (July 2008), e159. 2
- [HWG*11] HUTCHISON R., WOMELSDORF T., GATI J. S., LEUNG L., MENON R. S., EVERLING S.: Resting-State connectivity identifies distinct functional networks in macaque cingulate cortex. *Cerebral Cortex* (2011). 2
- [JBB*08] JAINEK W. M., BORN S., BARTZ D., STRASSER W., FISCHER J.: Illustrative hybrid visualization and exploration of anatomical and functional brain data. *Computer Graphics Forum* 27, 3 (May 2008), 855–862. 2
- [ME03] M. F., E. K.: Effect ordering for data displays. *Computational Statistics and Data Analysis* 43, 4 (2003), 509–539. 2
- [MTB03] MCGUFFIN M. J., TANCAU L., BALAKRISHNAN R.: Using deformations for browsing volumetric data. In *IEEE Visualization, 2003. VIS 2003* (Oct. 2003), IEEE, pp. 401–408. 5
- [RW11] ROBERT W. C.: AFNI: what a long strange trip it's been. *NeuroImage* (2011). 2
- [TLP*02] TZOURIO-MAZOYER N., LANDEAU B., PAPANASSIOU D., CRIVELLO F., ETARD O., DELCROIX N., MAZOYER B., JOLIOT M.: Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* 15, 1 (Jan. 2002), 273–289. PMID: 11771995. 3
- [vdHSBHP08] VAN DEN HEUVEL M. P., STAM C. J., BOERSMA M., HULSHOFF POL H. E.: Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. *NeuroImage* 43, 3 (Nov. 2008), 528–539. PMID: 18786642. 1
- [Yin09] YIN R. K.: *Case study research: design and methods*, 4 ed. Sage Publications, 2009. 6
- [ZEM*11] ZUO X., EHMKE R., MENNES M., IMPERATI D., CASTELLANOS F. X., SPORNS O., MILHAM M. P.: Network centrality in the human functional connectome. *Cerebral Cortex* (Oct. 2011). 2