Neuroinformatics (2019) 17:211–223
https://doi.org/10.1007/s12021-018-9395-8

ORIGINAL ARTICLE

An Uncertainty Visual Analytics Framework for fMRI Functional Connectivity

Michael de Ridder[1] · Karsten Klein[2] · Jean Yang[3] · Pengyi Yang[3] · Jim Lagopoulos[4] · Ian Hickie[4] · Max Bennett[4] · Jinman Kim[1]

Published online: 11 August 2018

Abstract

Analysis and interpretation of functional magnetic resonance imaging (fMRI) has been used to characterise many neuronal diseases, such as schizophrenia, bipolar disorder and Alzheimer's disease. Functional connectivity networks (FCNs) are widely used because they greatly reduce the amount of data that needs to be interpreted and they provide a common network structure that can be directly compared. However, FCNs contain a range of data uncertainties stemming from inherent limitations, e.g. during acquisition, as well as the loss of voxel-level data, and the use of thresholding in data abstraction. Additionally, human uncertainties arise during interpretation due to the complexity in understanding the data. While existing FCN visual analytics tools have begun to mitigate the human ambiguities, reducing the impact of data limitations is an open problem. In this paper, we propose a novel visual analytics framework with three linked, purpose-designed components to evoke deeper interpretation of the fMRI data: (i) an enhanced FCN abstraction; (ii) a temporal signal viewer; and (iii) the anatomical context. Each component has been specifically designed with novel visual cues and interaction to expose the impact of uncertainties on the data. We augment this with two methods designed for comparing subjects, by using a small multiples and a marker approach. We demonstrate the enhancements enabled by our framework on three case studies of common research scenarios, using clinical schizophrenia data, which highlight the value in interpreting fMRI FCN data with an awareness of the uncertainties. Finally, we discuss our framework in the context of fMRI visual analytics and the extensibility of our approach.

Keywords Visual Analytics · Functional Magnetic Resonance Imaging · Functional Connectivity · Uncertainty · Framework · Visualization

Introduction

Medical imaging modalities, such as functional Magnetic Resonance Imaging (fMRI), are fundamental to the growing understanding of the human brain. The technique images changes in blood oxygenation in a series of temporal scans to produce a proxy for neuronal activity (Jezzard et al. 2001; Massimo Filippi and Filippi 2009). The resulting fMRI are 4D images containing three spatial and one temporal dimensions. Analysis and interpretation of fMRI is used to characterise diseases, such as schizophrenia (Arbabshirani et al. 2014), bipolar disorder (Rashid et al. 2016), post-traumatic stress disorder (F. Liu et al. 2015) and Alzheimer's disease (Sarraf and Tofighi 2016). These insights rely on the researcher's ability to compare images, either across populations, e.g. schizophrenia and healthy (O. Sporns 2010), or for the same subject at different time points, e.g. for Alzheimer's disease staging (Sheline and Raichle 2013).

Functional connectivity networks (FCNs) are a foundational, widely used abstraction designed to convey the similarities and differences between fMRI images in a manageable and understandable format (O. Sporns 2010). FCNs are created by first registering an fMRI to a pre-segmented atlas – known as a parcellation – which contains multiple regions of interest

✉ Michael de Ridder
michael.deridder@sydney.edu.au

1 Biomedical and Multimedia Information Technologies (BMIT) research group, The University of Sydney, Sydney, Australia

2 Department of Computer Science and Information Science, The University of Konstanz, Konstanz, Germany

3 School of Mathematics and Statistics, The University of Sydney, Sydney, Australia

4 Brain and Mind Centre, The University of Sydney, Sydney, Australia

(ROIs). Second, a representative temporal signal for each ROI is constructed, often by averaging the temporal signals of all voxels within the ROI. Finally, a fully connected network is created that summarises how similar – or 'coactive – each pair of ROIs are over the temporal sequence, using methods such as correlation of the representative signals. FCNs are widely used because they greatly reduce the amount of data that needs to be interpreted and they provide a common network structure for all images that can be directly compared. Moreover, well-established graph analysis techniques, e.g. graph kernel analysis (Jie et al. 2016), can be applied to the networks as these can effectively summarise key features, such as network topology.

While FCNs have proven valuable in fMRI analysis, the abstraction results in limitations and inaccuracies, primarily due to the loss of voxel-level data that limit their interpretation (O. Sporns 2010; M Filippi 2016; O Sporns 2014). These ambiguities include data uncertainties, arising from inherent hardware and image processing limitations, that cause, e.g. imaging artifacts, and the abstraction process of grouping voxels and temporal signals into single data points; as well as human uncertainties, such as cognitive load and difficulty in understanding the data (M Filippi 2016; Ristovski et al. 2014; Carp 2012; O Sporns 2014). To mitigate some of the human issues, the fMRI community has actively investigated visual analytics tools, which combine automated processing with interactive visualisation for improved analysis (Margulies et al. 2013). These tools typically combine a graphical visualisation of the FCN with a link to the spatial location of ROIs, e.g. through node location (Irimia et al. 2012) or by overlaying the network on a 3D representation of the anatomy (Xia et al. 2013). This is done to reduce the ambiguity caused by high cognitive load as users require less mental reconstruction to connect the FCNs to their existing spatial knowledge of the brain (Margulies et al. 2013). Other systems combine visual analytics for FCN data with interactive visualisations of related brain data, such as structural and diffusion tensor imaging (Angulo et al. 2016). However, the data uncertainties are overlooked by existing visual analytics tools. These limitations are often considered one of the last major barriers preventing wider uptake of fMRI in research and clinical environments (Lee et al. 2013; Peeters and Sunaert 2007). As a result, there is a demand for an uncertainty visual analytics that considers the data imprecision alongside the human ambiguity, which will enable a deeper understanding of the fMRI.

In this paper, we propose a visual analytics framework for fMRI FCNs to address the inherent data and human uncertainties. This framework consists of three specifically designed, linked components: the FCN abstraction, the temporal signal viewer and the anatomical context viewer. Each component has purpose-designed visual cues and interactions to guide interpretation and evoke a deeper understanding of the data. These have been augmented with two methods of comparing subjects, a common task in fMRI analysis: small multiples and a marker approach. The framework has been developed in collaboration with domain experts, based on a set of design guidelines. We demonstrate that our framework enables a user to gain an understanding of uncertainties during analysis without disrupting existing workflows. Finally, we discuss our framework in the context of fMRI visual analytics.

Background and Related Work

Uncertainties in fMRI

Uncertainties in fMRI FCNs are created throughout an image analysis pipeline that consists of distinct phases: (i) image acquisition and processing; (ii) data abstraction to the FCN; and (iii) human interpretation and visualisation. Each phase in the pipeline adds and compounds the uncertainties. Image acquisition shortcomings arise primarily from hardware limitations, for example, low resolution, poor signal to noise ratio, and physiological differences between patients (Massimo Filippi and Filippi 2009; Jezzard et al. 2001). Image processing includes steps such as image reconstruction, normalisation, and registration. Each of these add ambiguity as the algorithms have inherent inaccuracies and require the manual selection of parameters that can vastly impact the output data (Massimo Filippi and Filippi 2009; Jezzard et al. 2001; Eklund et al. 2016). The abstraction to FCN phase add uncertainty through two forms of heterogeneity: spatial and temporal. Spatial heterogeneity is caused because ROIs in parcellations often contain a mixture of numerous heterogeneous voxels, since they are based on pre-defined anatomical locations, rather than temporal signals of the individual subject (O Sporns 2014; M Filippi 2016). Temporal heterogeneity, meanwhile, is a result of the temporal signal of each ROI potentially containing multiple subsets of activity states or other artifacts, such as noise, that are compared to create a single number for each coactivation (O Sporns 2014; M Filippi 2016; Ristovski et al. 2014). In addition, most FCN analysis approaches rely on thresholding the network based on the coactivation values to filter out potentially uninformative coactivations and make comparison easier. However, there are no accepted methods for selecting a good threshold, or threshold range (Stevens et al. 2013; Gorgolewski et al. 2016). Instead, thresholds that are simply considered 'high enough' or 'low enough' are commonly used (Stevens et al. 2013). Finally, uncertainties during human interpretation and visualisation arise mostly due to the complexity of the data and associated cognitive load on users. For example, it is important to provide anatomical context during interpretation as it is mentally taxing to relate ROIs back to anatomical location (Margulies et al. 2013). Similarly, visualisations must be easy to interpret and not present misleading or potentially spurious data, such as implied connections (Margulies et al. 2013). Many of these limitations and imprecisions are inherent to the steps in the analysis pipeline and therefore cannot be fully mitigated. As a result, it is paramount to ensure users are aware of the effects of the uncertainties on the data during interpretation.

Information Visualisation and Visual Analytics

Existing information visualisation and visual analytics systems aim to simplify the presentation of FCNs and provide anatomical context to minimise the cognitive load and improve understanding of the data. Several different visualisation metaphors have been used in FCN presentation; these include node-link diagrams, matrix views, radial graphs, and more complex metaphors. Notable examples of these include: Irimia et al. (Irimia et al. 2012) proposed a radial visualisation of connectivity data, which implied anatomical location through the order of nodes on the radial, i.e. right hemisphere on the right, left hemisphere on the left; de Ridder et al. (de Ridder et al. 2015) combined the node ordering idea with a linked anatomical reference; Böttger et al. (Böttger et al. 2014) used edge bundling and glyphs in node-link diagrams to minimise clutter when overlaid on a 3D view of the anatomy; Bach et al. (B. Bach et al. 2015) extended the concept of heatmap matrices for exploring temporal patterns in FCNs; Cui et al. (Cui et al. 2014) proposed GraphFlow, a visualisation that summarises the evolution of graph metrics over time; Angulo et al. (Angulo et al. 2016) integrated other related data types, such as DTI images, with various graphing techniques; and Bach et al. used an extension on connected scatterplots for FCNs (B Bach et al. 2016). This research has led to widespread use of some systems, e.g. BrainNet Viewer (Xia et al. 2013). While newer tools have begun to look at comparing subjects, e.g. (Fujiwara et al. 2017) which performs dimensionality reduction on the FCN data to project multiple subjects onto a plane, indicating the overall FCN patterns by their locations. However, while these tools satisfy the human uncertainties, they do not tackle the underlying data limitations.

Design Guidelines

In this section, we present key design guidelines (DG) that steered the development of our framework. These were created from discussions with the co-authored domain experts, principally MB and JY, from literature analysis of fMRI uncertainties and from studying existing fMRI visual analytics tools.

DG1: Enhance an Established FCN Visualisation for Data Uncertainty The basis of our visual analytics framework should adopt an established FCN visualisation to ensure we leverage the benefits it provides to human understanding and enhance it by integrating with uncertainty visualisations.

DG2: Facilitate Comparison of FCNs Comparison among different subjects or between subject populations is a fundamental step in many FCN analyses that most existing visual analytics solutions do not address.

DG3: Provide Interactive FCN Thresholding with Visual Cues One of the major uncertainties in fMRI analysis arises due to thresholding. Interactivity with visual cues to highlight possible threshold values may allow users to mitigate some of the imprecision about the impact of their threshold selection.

DG4: Highlight Spatial Heterogeneity Heterogeneity of voxels in ROIs is a major uncertainty. Visual cues can be used to highlight when a region is highly heterogeneous, both in relation to other ROIs, and as an absolute heterogeneity value.

DG5: Expose Temporal Heterogeneity ROIs contain both representative and voxel-level temporal signals. Imaging artifacts, short bursts of activity and trends in these signals create ambiguity regarding whether the summarisation to a coactivation value is accurate.

Framework Design

In this section, we present our visual analytics framework. The overall design of our framework is illustrated in Fig. 1 and is detailed in the sections below. Figure 2 gives an overview of the data processing workflow that is employed by our framework.

DG1: Functional Connectivity Network Abstraction and Anatomical Context

Abstraction of FCNs is a well-studied area with presentation generally taking one of three forms (Margulies et al. 2013): matrix representations; radial graphs; and node-link diagrams, often on a 3D rendering of the anatomy. While there are advantages and disadvantages in all these approaches, we eliminated the use of a node-link abstraction as there is no inherent spatial information and when the spatial context is added by rendering on the anatomy, there are issues, such as implied connections, occlusion and visual clutter. In our framework, we allow a user to select between radial and matrix views. For both, ROIs are named nodes around the outside, while coactivation values are indicated by edge opacity. We utilised principles that have been presented in previous works, e.g., to minimise the effects of human uncertainties during analysis. The two radial graphs on Fig. 1 right illustrate the following concepts.

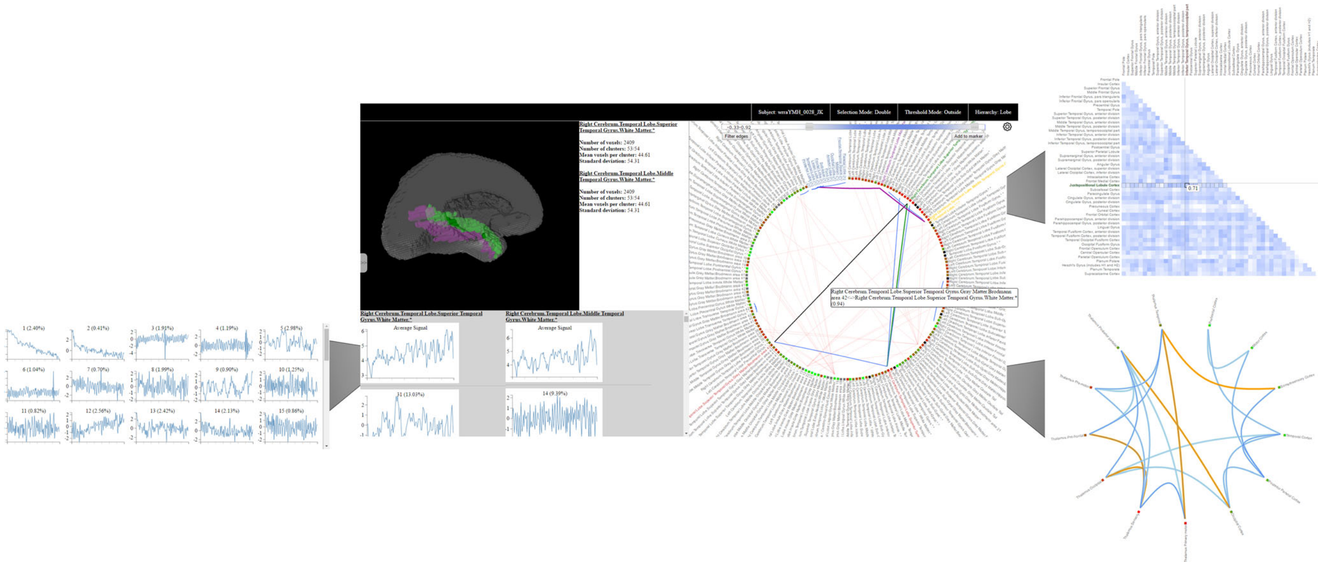

Fig. 1 Overview of our visual analytics framework. The core interface shows the three proposed components for: **a**) functional connectivity, with indicators for thresholding and heterogeneity issues; **b**) anatomical context; and **c**) temporal sequences, which can show artifacts in the imaging data. The popouts to the side show how these components can be replaced in the framework

Node ordering is exploited to indicate an approximation of the anatomical location. The default node ordering has ROIs from the right-hemisphere on the right of the graph, ROIs from the left-hemisphere on the left, and ROIs from the central parts of the brain in the middle. This technique has been argued as a method to reduce cognitive load (Margulies et al. 2013), and hence human uncertainties during analysis. Nodes can be reordered based on a user-input sequence.

Edge bundling is applied to the radial layout to minimise visual clutter and to indicate the overall organisation of the network. We perform edge bundling based on the ratio of edges to nodes, so graphs with fewer edges have less bundling and graphs with more edges have more bundling.

Alongside the FCN, providing anatomical context is crucial in minimising human uncertainties due to cognitive load during interpretation. We separate the anatomy from the FCN abstraction – rather than overlaying the network on the anatomy in a node-link diagram or by placing brain lobe diagrams on the radial – to keep the anatomy whole, thus reducing mental reconstruction, and to minimise visual clutter as shown in Fig. 1 top-left. The anatomical representation consists of a surface rendering of the brain, and a surface rendering of each ROI. The surface renderings are created using atlas images, ensuring they are accurate representations in size, shape and location. The use of an atlas, rather than a patient-specific mesh, was requested by the co-authored domain experts because it is a common, repeatable structure; however, this is an optional feature as patient-specific meshes can be created. We have provided this ability because both atlas and patient-specific meshes have drawbacks in their use regarding anatomical accuracy against ability to directly compare subjects. In keeping with the design of the framework, as long as a user is aware of the limitations, they can better adjust their analysis and interpretation; potentially using both approaches.

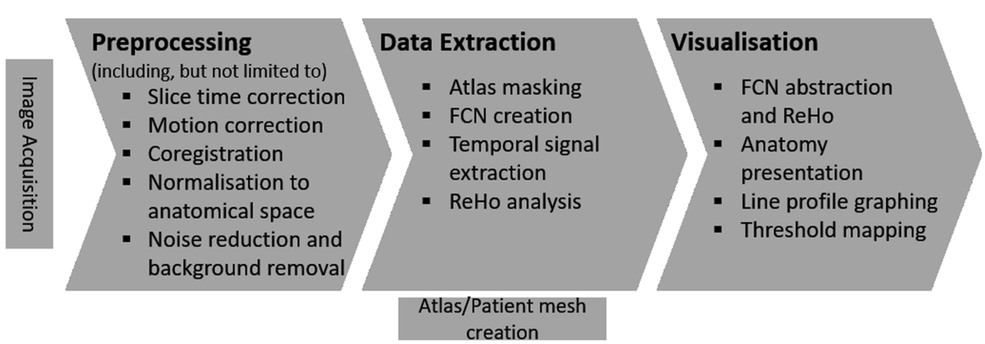

Fig. 2 An overview of the data processing workflow used by our framework. Specific techniques and algorithms are not provided as these are designed to be interchangeable instead, the high level steps are given. Preprocessing is performed to convert the raw scanner data into standard fMRI images. Data extraction then handles the data from multiple angles to extract the information required in our uncertainty visualisations. This is augmented with optional mesh creation (from an atlas or patient specific). Finally, the visualisation step outlines the features that are discussed in detail in this paper

DG2: Comparing Functional Connectivity Networks

We provide two key extensions to the FCN visualisations to meet the design goal of facilitating comparison of FCNs by proposing a small multiples view, and a marker, which indicates differences on a single radial graph.

Small multiples has been implemented as a comparison method that does not require aggregation which may be susceptible to imprecisions, such as from outliers. Figure 3 illustrates our small multiples view of FCN networks. In this view, interaction with any of the multiples is mirrored on the other graphs.

Markers are a novel contribution to this domain, created as shown in Fig. 4: multiple FCNs are selected by a user to create a 'marker network' which is used to highlight the similarities and differences with any other FCN. The marker network is the average of all the coactivation values, normalised with respect to the maximum coactivation. By default, we provide ± half a standard deviation to broaden the range of possible similarity overlap. Due to individual baselines in fMRI coactivity, comparisons are based on normalised FCNs with respect to the individual maximum, and not the raw coactivity measures. That is, during the comparison, the maximum coactivation in both the marker network and the individual

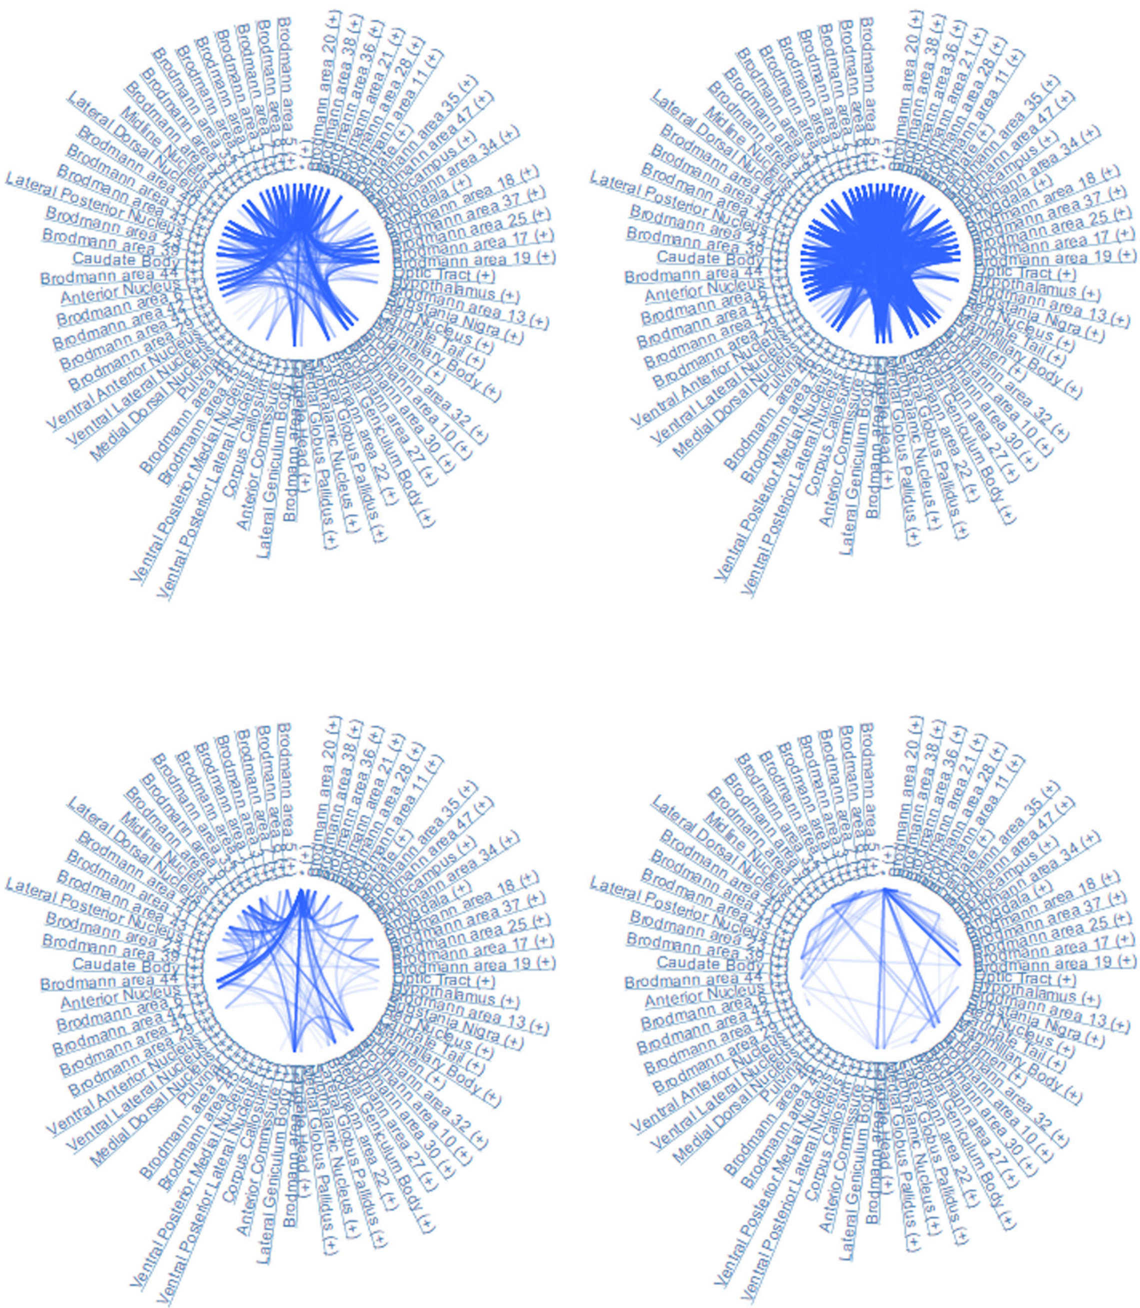

Fig. 3 Small multiples view showing four subject FCNs at the same threshold clearly indicating the high-level differences that are visible in this view

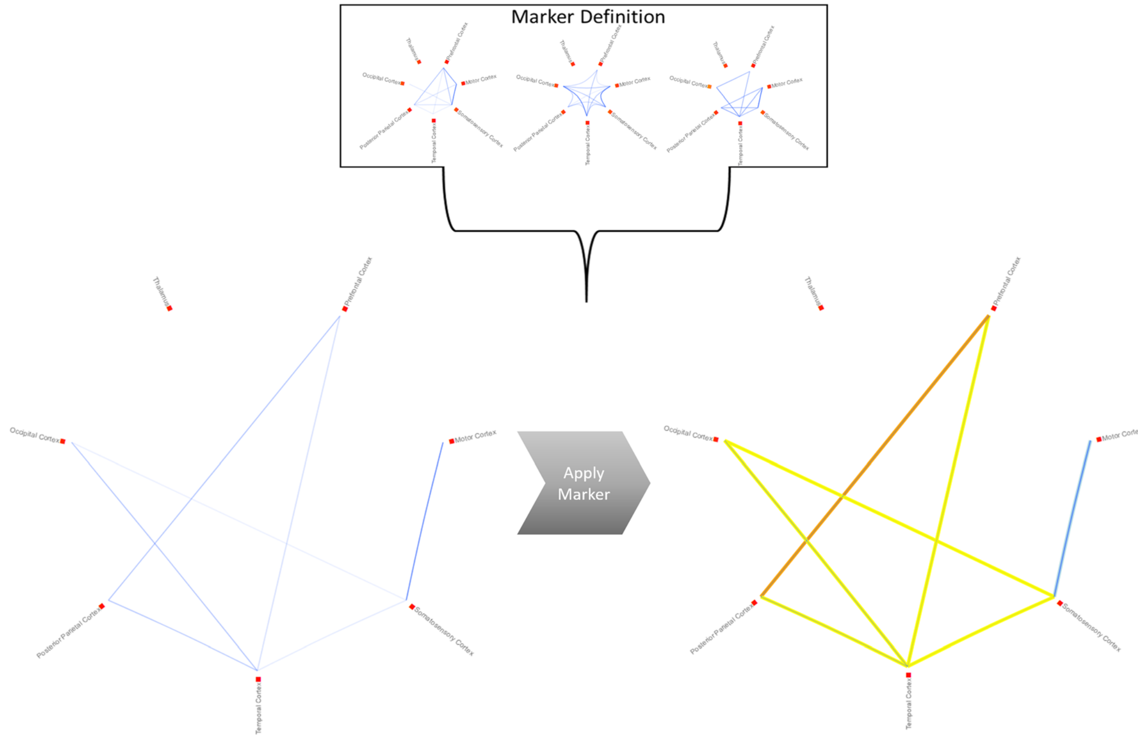

Fig. 4 An example of applying a marker to a network. Edges are highlighted depending on whether they are above (orange), below (blue) or within (yellow) the marker range. The colours are based on user selection

subject is always 1, with all lower coactivations a ratio of the maximum. Therefore, shifts in maximum coactivation value are accounted for, and the shape of the network is then compared. Markers networks can be used to represent a population, such as diseased FCNs, for comparison to a single subject. Moreover, before each FCN is added to the marker definition, it is viewed against current marker network. In this way, we aim to mitigate some of the selection bias by providing more information before selection.

DG3: Thresholding

Interactive thresholding is enabled using a threshold bar, which is set above the FCN abstraction. On this bar, the proportion of edges within a set of threshold buckets is indicated using a heatmap, to give a visual cue for selecting a threshold as shown in Fig. 5. The range of possible coactivation values are divided into twenty buckets of equal size, e.g. bucket size of 0.1 for Pearson correlation. For each bucket, the proportion of edges determines the colour of the heat map. When in small multiples mode (Section "DG2: Comparing Functional Connectivity Networks"), these heatmaps are stacked within the threshold bar, providing a further visual cue for comparison, Fig. 5.

DG4: Heterogeneity

To satisfy the heterogeneity design goal, we provide visual cues on our FCN abstraction. To determine the heterogeneity of each ROI, we use the established ReHo measure (Zang et al. 2004) in the AFNI software package (National Institute of Health 2016). This measure is widely used in fMRI analysis, e.g. (Y. Liu et al. 2008; Zeng et al. 2015), to measure the temporal homogeneity of labelled regions. It is independent of ROI size or location, making it ideal for our purpose, as different ROIs can be effectively compared. We provide a novel indication of this value in our framework using coloured boxes placed alongside each FCN node, that minimises clutter

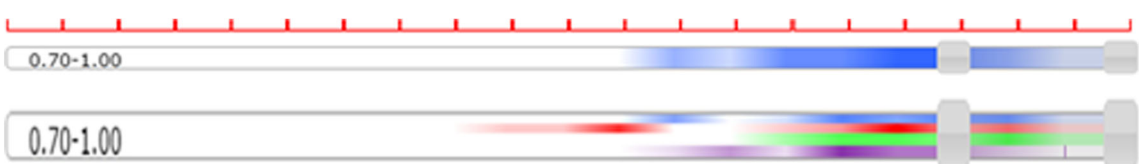

Fig. 5 The two versions of threshold bar with a grid added in red above the bars to indicate the division of the buckets. Top: in single subject mode; and bottom: in small multiples mode with 4 subjects selected. The bottom bar has been stretched vertically for more visual clarity and is best used with fewer subjects in the framework

and ensures users can see the details while exploring the FCN. Red boxes indicate highly heterogeneous ROIs, and green boxes, highly homogeneous ROIs, with a colour transition between. Further, we allow users to view the relative heterogeneity of all ROIs, where the most homogeneous ROI is pure green, and the most heterogeneous is pure red, allowing for associations to be made within a single image. These two modes are shown in Fig. 6.

DG5: Temporal Signals

Exposing the temporal signals across the whole brain, independent of FCN ROIs, is important for highlighting temporal inaccuracies. However, because fMRI data contains tens to hundreds of thousands of voxels, it is not practical to display all temporal signals. Hence, in our framework, we perform data reduction using independent components analysis (ICA), which is widely used in fMRI analysis. In our presented implementation, we used FSL (FMRIB Analysis Group 2016) to perform ICA and automatically select an image-specific number of components to demonstrate how the visual component works and interacts with the rest of the framework. The data reduction method is fully interchangeable. Users can perform manual grouping or apply their own clustering if they wish to better control for possible errors; such as limitations with automatic selection or selection bias. We class each voxel by its most representative component. As shown in Fig. 1 bottom-left, we display the temporal line graphs of every component, ordered by the proportion of voxels they represent, when no node is selected. Upon selection, we list the ordered components within the region and allow comparison on multiple selection.

Interaction

Interaction is crucial in FCN analysis as the users' needs are often changing with different subjects or studies. In the framework, the visual components are linked such that interaction with one component directly impacts what is visualised in another component. In the below descriptions, all colours can be set by the user.

Functional Connectivity Network

Node Selection Users can select nodes in two modes: single and double. In single select mode, selection of a node results in all connected edges within the threshold to be highlighted. The list of temporal signals within a region are also displayed and the ROI will be shown on the anatomy with connected ROIs displayed in a different colour. Double select mode is shown in Fig. 1 middle. The edges for both selected nodes are highlighted, with the temporal signals listed side-by-side and the two ROIs displayed on the anatomy. Connected ROIs are not shown on the anatomy to prevent occlusion and clutter.

Edge Hover When a user hovers over an edge on the FCN visualisation, a tooltip is shown that names the two ROIs and shows the exact coactivation value.

Filtering During analysis users often focus on a subset of ROIs or edges. We facilitate this by allowing users to filter ROIs from the FCN abstraction, removing them from the presented graph. This allows users to reduce visual complexity and focus on the subset of interest.

Anatomical Context

On the 3D anatomy viewer, users can toggle whether all the ROIs are shown or just the surface is shown. In the case where all ROIs are shown, users can select an ROI on the anatomy. This has the same effect as the single node select mode on the FCN.

Thresholding

Users can adjust the threshold bar to set an upper and lower bound in two modes using the two grey boxes on each bar in Fig. 5: inside mode and outside mode. Inside mode updates the FCN to show all edges between the lower and upper bounds, and hide all edges outside these values. Outside mode shows all edges below the lower bound and above the upper bound and hides all edges between the two values. In outside mode, the edges below the lower threshold are shown in a different colour to the edges above the upper threshold.

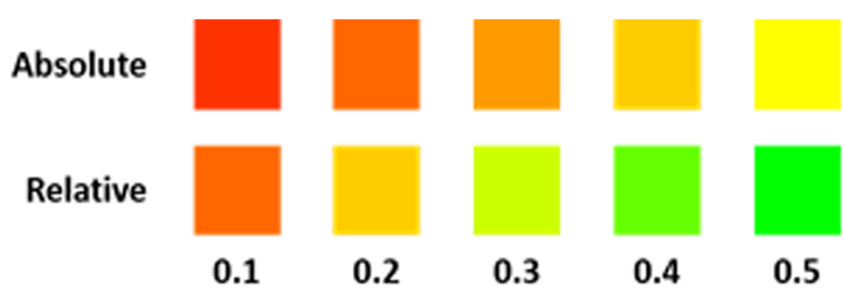

Fig. 6 Comparison of the heterogeneity indicators in absolute and relative mode for a range of ReHo values

Temporal Signals

The list of temporal signals, in single select, double select, and when nothing is selected can be reordered by drag and drop to improve visual comparison. Each signal can also be clicked on to show the spread of voxels across the brain for the selected class. Further, regions on the FCN abstraction that contain voxels of the selected class are highlighted.

Case Studies

We present the following case studies as examples of three common research scenarios where an understanding of FCN uncertainties adds value to the current workflow. The dataset used consisted of a cohort of 91 clinical fMRI subject studies who are either healthy ($n = 38$) or have been diagnosed with schizophrenia ($n = 53$). The schizophrenia cohort was selected because it is a challenging and active research area where the effects of imprecisions in the data may look similar to the sparse expected patterns resulting from functional dysconnectivity (Wang et al. 2017; Liang et al. 2006; Friston et al. 2016). In these case studies parcellations that relate cortical regions to the thalamus were used as these connections are of interest when studying schizophrenia, e.g. (Woodward et al. 2012; Giraldo-Chica and Woodward 2016). The case studies were performed based on discussions and feedback from the co-authored domain experts.

Case Study 1: Analysis of Uncertainties in a Subject

In this case study, we demonstrate how exploring uncertainty through our framework can be used to expose underlying issues in the data. This case study replicates a research scenario where a user has found something uncharacteristic about their data. Schizophrenia is known to be characterised by functional dysconnectivity when compared to healthy patients (Friston et al. 2016). Visually, this means schizophrenia FCNs will have fewer connections at high thresholds and there will be no clear patterns between FCNs on the radial, while healthy patients will have dominant patterns and more connections at higher thresholds. However, upon initial inspection some of the FCNs do not meet these expectations.

Figure 7(c) shows the radial of a healthy subject that is sparse at the default high threshold and has no clear connectivity pattern and a very small threshold range on the inidcator. The initial visual feature that stands out on the radial is the oddly high connectivity between the motor cortex and the somatosensory cortex, indicated by the small red arrow. A user can easily drill-down into this connection to ascertain where this may have come from. The anatomy viewer is shown in Fig. 7(a), while the average signals and top temporal signals are in 7(b). These clearly indicate that the regions are adjacent to one another and that there is a common peak towards the end of the average temporal signals that is out of scale when compared to the rest of the average signals. This is likely to indicate patient movement resulting from external stimuli

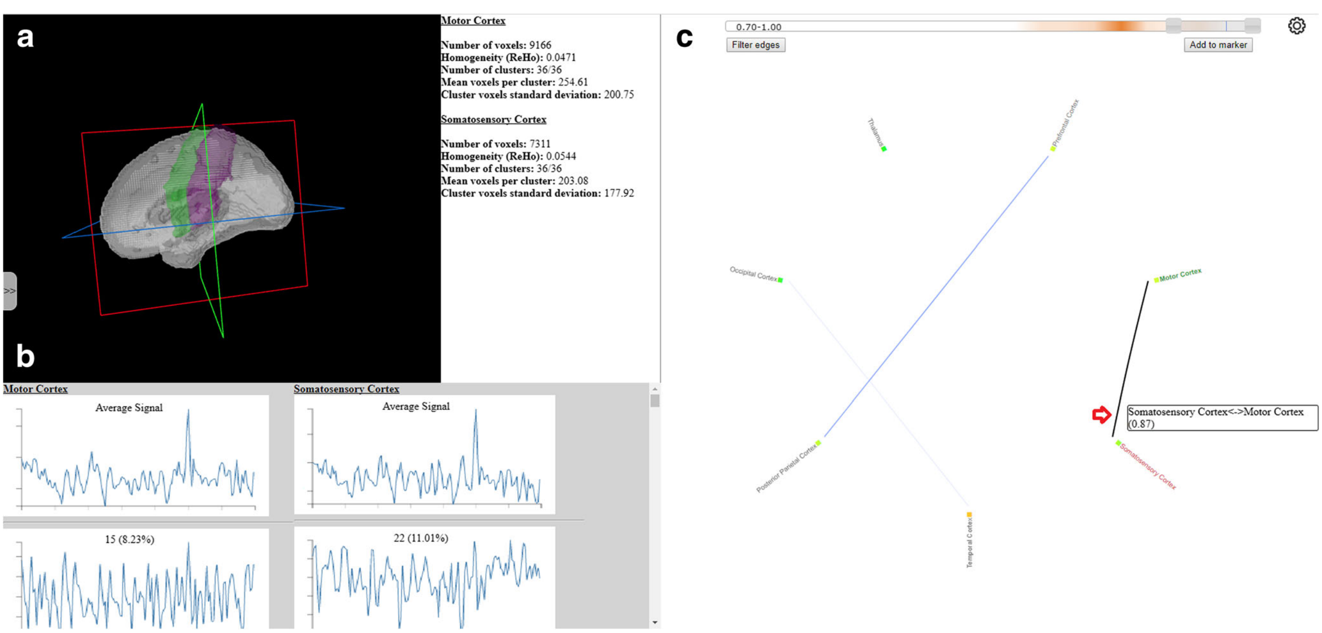

Fig. 7 Overview for a healthy subject that is unexpectedly sparse at the high default threshold with two regions selected that have an incongruously high coactivation. In (**a**) the two regions are shown to be adjacent; (**b**) shows that the two average signals have out-of-scale peaks near the end; and (**c**) shows the sparse FCN and the high coactivation between the motor cortex and somatosensory cortex (indicated by the red arrow)

during the scan, as the somatosensory cortex is involved in processing external stimuli, such as auditory noises, and the motor cortex is involved in movement (Swenson 2006).

Case Study 2: Assessing Average FCNs

One of the most common methods for studying fMRI FCNs is to average the matrices from the population and then to compare the averaged FCNs. This practice, while convenient, is known to be sub-optimal due to the within-population variability across individuals that is obscured by the averaging (Fujiwara et al. 2017). In this case study, we show that our framework makes it easy to assess the quality of population averages. To do so, users can begin by interacting with the small multiples view to compare the average healthy and average schizophrenia FCNs to individual healthy and schizophrenia patients.

For example, Fig. 8 shows the application of our small multiples visualisation in comparing the average healthy FCN, which was made using a randomly selected half of the healthy patients ($n = 19$), with 5 individual healthy patients that were not used to build the average. In the figure, it is clear from the radial visualisations that, while 2 of the individual FCNs resemble the average, the other 3 are quite different. A user can quickly group the individuals by viewing the small multiples into these subgroups to assess the average. By using the marker function, a user can then create a marker of one group and directly compare the differences where they can adjust the threshold to visualise specific radial edges. This is shown in Fig. 9, which contains a marker created with 10 similar schizophrenia patients applied to another schizophrenia patient with a different pattern. The threshold has been lowered to compare the bottom connections which may be used to understand functional dysconnectivity. Using the marker, it is clear that the FCN for the individual schizophrenia patient contains coactivations much lower than the marker (blue highlight).

These distinct groups that are clearly visible using our framework indicate the impact of inter-subject variability that

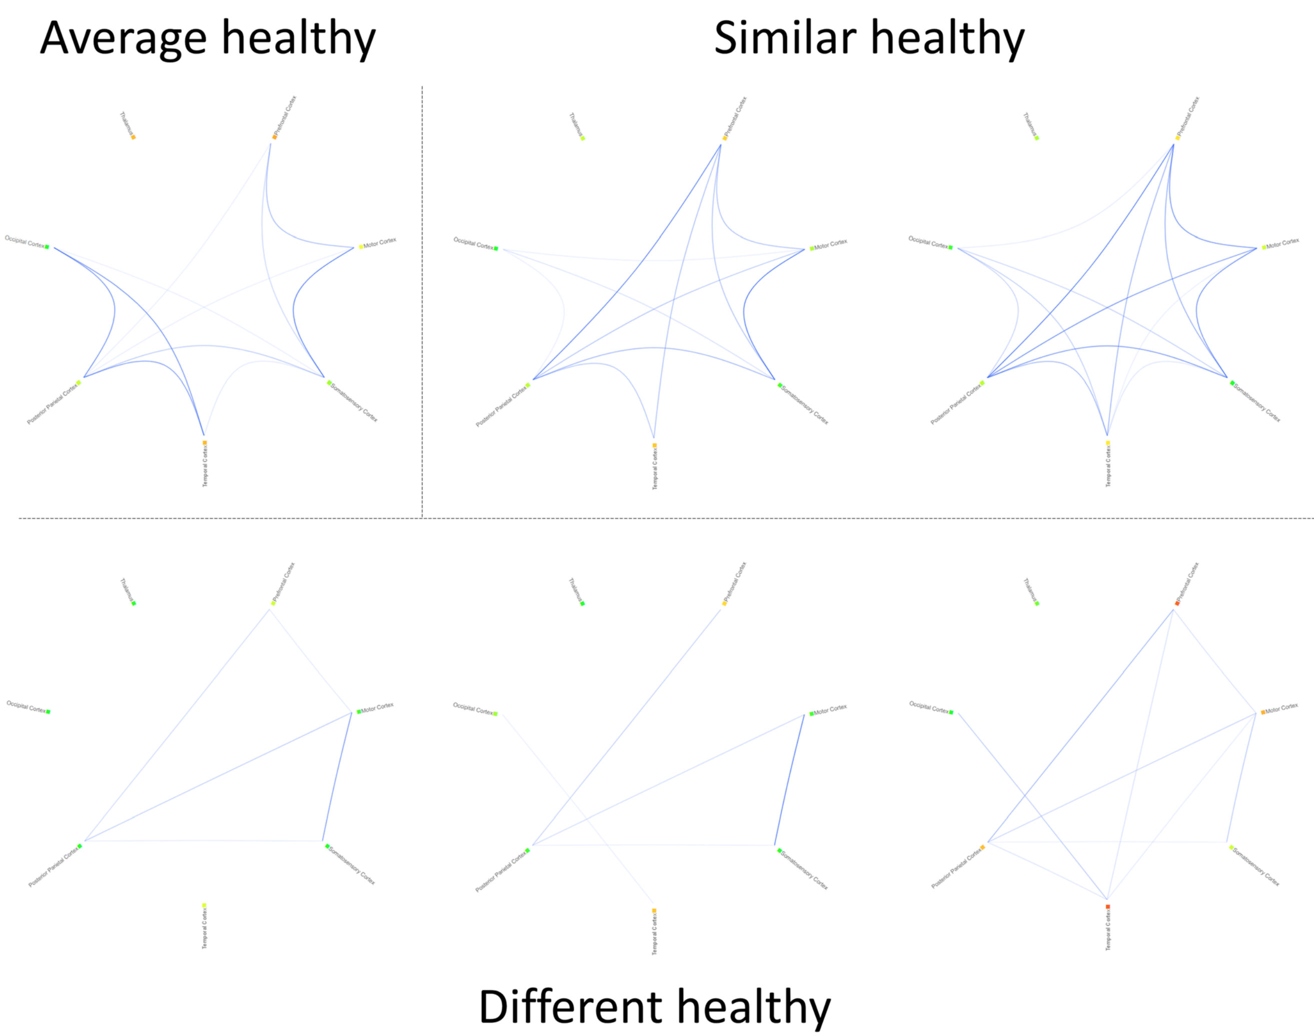

Fig. 8 Small multiples visualisation of the average healthy FCN (top-left) and five individual health patients used to assess the accuracy and quality of the average FCN. The top two have a similar pattern to the average, while the bottom three have a distinctly different pattern at the same threshold range

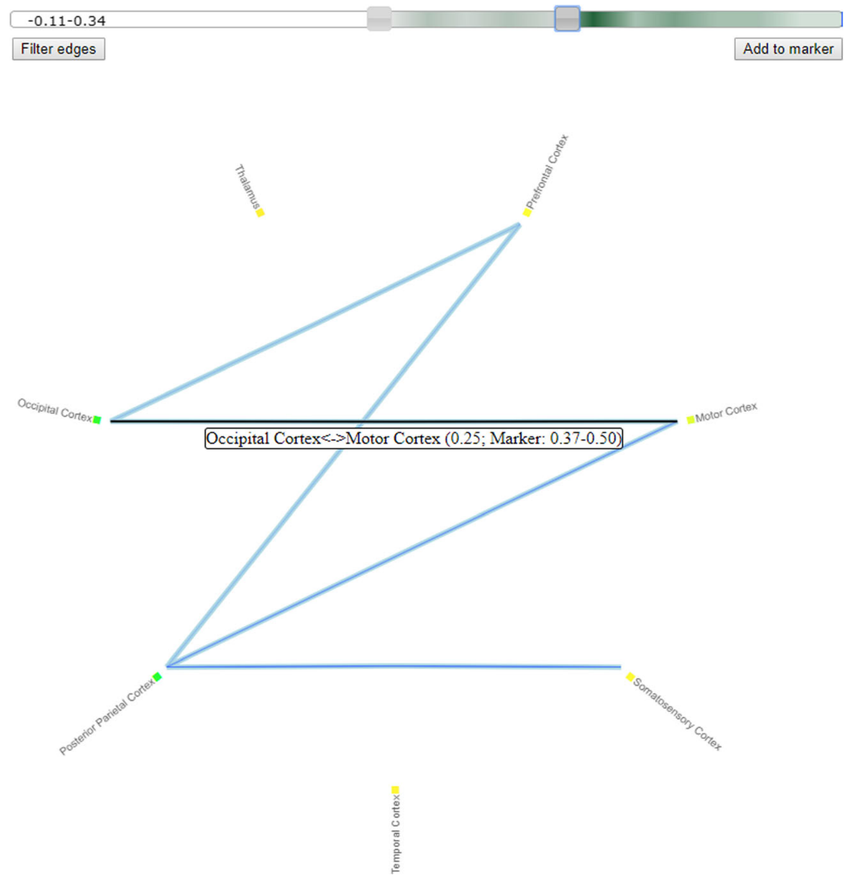

Fig. 9 Using the marker function to investigate a single schizophrenia subject. The marker was created using ten other schizophrenia patients, all with similar FCNs. The threshold bar has been set to a low range to observe the less coactive regions

can have on average FCNs. Moreover, users can also assess whether the differences may be a result of issues due to limitations, such as from high heterogeneity of regions. This process is demonstrated in Case Study 3.

Case Study 3: Ensuring Data Integrity

Ensuring the integrity of fMRI datasets is a necessary practice that involves discarding subjects with issues, such as imaging artifacts, which may negatively impact the analysis. Typically it involves complex statistical procedures and the use of other datasets, e.g. (Woodward et al. 2012) used a five-point (signal-to-noise ratio, percent drift, percent fluctuation, radius of decorrelation, percent standard deviation) quality assurance test, reliant on a training dataset and pre-defined thresholds without any detailed subject-level inspection. Using our framework, a user can ensure data integrity at a much more granular and detailed level using only one data set. This allows users to gain an understanding of the imprecisions that led to issues in the data so they can, e.g. remove subsets of the temporal sequence for a subject, rather than entirely discarding the subject.

To perform this, a user can take their knowledge of the population averages, gained in Case Study 2, and assess each subject against it using our framework. In particular, any FCN that does not resemble any of the averages can be marked as a candidate for low data integrity. Using the thresholding, heterogeneity (ReHo) and temporal signal visualisations and functions, a user can easily see areas of concern. For example, FCNs with highly homogeneous regions are clearly visible in small multiples, Fig. 8 bottom-left; disproportionately high threshold heatmaps stand out, Fig. 10; and out-of-scale spikes in temporal signals from the same subject can be seen, as in Fig. 11. These are likely caused by a machine or physiological artifact that was not removed during processing. Due to this level of detail, a user can remove the temporal frames impacted by the spike, rather than discarding the whole image, resulting in an FCN that is quite similar to what is expected. Moreover, this level of inspection can be combined with statistical approaches, meaning that more information can be gathered from a single cohort.

Discussion

Analysis of fMRI data is a challenging task. One of the prevailing methods, using functional connectivity networks, has been widely researched by the visual analytics

Fig. 10 A threshold heatmap from a schizophrenia subject that is very tightly clustered at the high range

field. However, the proposed solutions all focus on addressing uncertainties surrounding human understanding and cognitive load. This results in other limitations, deriving from artifacts in the source images and those that arise during the abstraction process. Thus, in this paper we proposed a framework that can holistically address the range of human and data uncertainties. We demonstrated how our novel integration of components can collectively enhance the fMRI analysis pipeline. We innovate by designing specialised visual analytics components and demonstrating that our purpose-designed integration, interaction and comparison techniques for exposing and exploring uncertainties in FCN abstraction and the source fMRI images is both a novel application and framework that provides new, necessary insight into the data. The identification of the required components and the interrelation between them is also not a straightforward task. These shortcomings are commonly thought to be one of the last major hurdles to wider research and clinical uptake of fMRI (Lee et al. 2013; Peeters and Sunaert 2007). The co-authored domain experts note that this presented framework intuitively presents a lot of information on the fMRI and that the exposure to the temporal aspects of the data through the line profiles is useful. Similarly, they commented that we have managed to integrate a lot of information into the one integrated interface well. Moreover, by interactively allowing the presentation and comparison of multiple subjects and single subjects to population markers, we address two aspects of fMRI analysis that many previous visual analytics tools have not addressed.

In planning and design of our framework, we adapted elements from prior visual analytics solutions as they provide strong benefits despite their inherent limitations. Principally, we adapted the combined combination of anatomy and FCN network into split views that integrate with one another for minimised cognitive load. In parallel, our node ordering on the radial is appropriated from Irimia et al. (Irimia et al. 2012); we have added the ability for users to upload their own ordering. Similarly edge bundling (Böttger et al. 2014), dimensionality reduction (Fujiwara et al. 2017) and the use of multiple graph types (Angulo et al. 2016) have been extended for use in uncertainty analysis. By adapting and refining these techniques for our specific development, we have ensured our proposed framework was built on state-of-the-art approaches that have been shown to be useful in minimising human ambiguities, e.g. cognitive load. This allowed our design to focus on exposing and mitigating data limitations in combination with the human issues.

Additionally, our framework was designed to support future developments in the field. All the pre-processing steps: the parcellation chosen and 3D mesh generation, the data reduction technique, the coactivation measure, and the heterogeneity measure can be interchanged with other techniques and selected depending on the task at hand. This also allows for the integration of newer techniques as more is discovered about the brain and fMRI analysis. In our case studies, we have demonstrated how our framework already allows a user to gain a deeper understanding of the limitations and inaccuracies, and how this benefits analysis. By showing the process of creating and viewing schizophrenia markers, we illustrate how visualising the uncertainties on top of the standard FCN approaches, as well as allowing direct comparison of subjects, presents a potentially clinical application of our framework. This, along with our design that can adapt to improvements in visualisation techniques, will allow our framework to reveal further information as more is understood about the brain. In

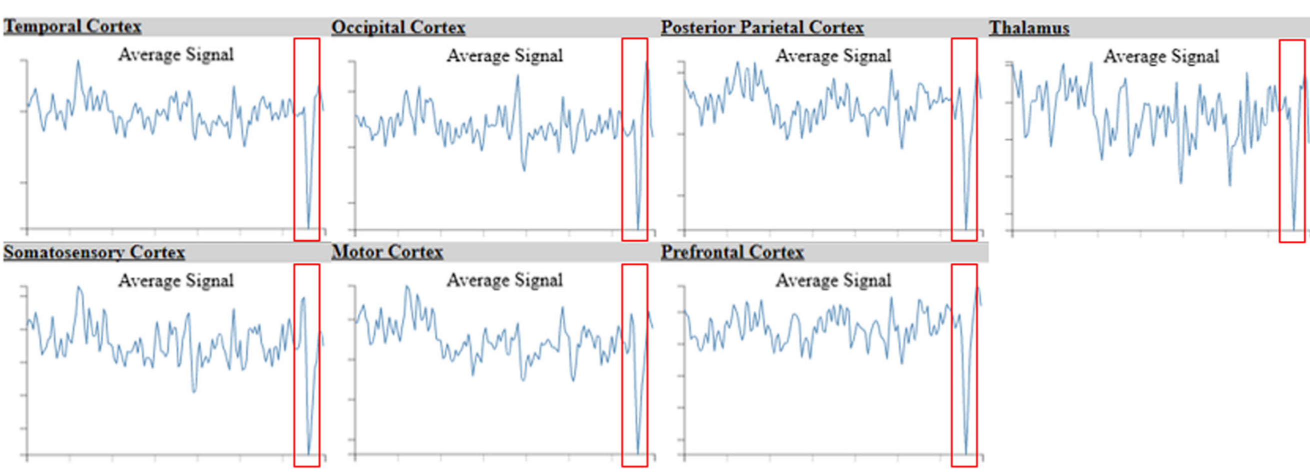

Fig. 11 The average activity profiles of the seven regions in the parcellation for one subject. All profiles contain similar drops at the end of the temporal sequence indicating an issue in the data integrity that can be explored in the other interface components, such as the homogeneity cues

future studies, we will investigate the impact and usability of our framework with a user study. In this study, we will also assess the usability for novices in order to improve the design for those beginning in the fMRI field. This will be part of our future work.

Conclusion

We have presented a new visual analytics framework for fMRI FCNs which integrates multiple visual components to address data and human uncertainties during interpretation. Our framework consists of three linked views: the FCN abstraction, the temporal signal viewer and the anatomical context, each with specifically designed visual cues and interactions to guide interpretation. We augment this with two methods of comparing subjects, a common task in fMRI analysis, by using small multiples and a marker approach. In our case studies, we have demonstrated not only how an understanding of the limitations and ambiguities adds value to the interpretation, but also that it is our framework that exposes these uncertainties.

Information Sharing Statement

A software implementation of the framework has been made open source and available at https://github.com/mderidder-usyd/CereVA. The available implementation was uncoupled from the ethics protected image data used in the case studies. Two example simulation patients have been created instead.

Compliance with Ethical Standards

Conflict of Interest None declared.

References

Angulo, D. A., Schneider, C., Oliver, J. H., Charpak, N., & Hernandez, J. T. (2016). A Multi-facetted Visual Analytics Tool for Exploratory Analysis of Human Brain and Function Datasets. *Frontiers in neuroinformatics, 10*.

Arbabshirani, M., Castro, E., & Calhoun, V (2014). Accurate classification of schizophrenia patients based on novel resting-state fmri features. In *EMBC,* 6691–6694.

Bach, B., Henry-Riche, N., Dwyer, T., Madhyastha, T., Fekete, J. D., & Grabowski, T. (2015). Small MultiPiles: Piling time to explore temporal patterns in dynamic networks. *Computer Graphics Forum, 34*(3), 31–40.

Bach, B., Shi, C., Heulot, N., Madhyastha, T., Grabowski, T., & Dragicevic, P. (2016). Time curves: Folding time to visualize patterns of temporal evolution in data. *IEEE Transactions on Visualization and Computer Graphics, 22*(1), 559–568.

Böttger, J., Schäfer, A., & Lohmann, G. (2014). Three-dimensional mean-shift edge bundling for the visualization of functional connectivity in the brain. *IEEE Transactions on Visualization and Computer Graphics, 20*(3), 471–480.

Carp, J. (2012). On the plurality of (methodological) worlds: Estimating the analytic flexibility of FMRI experiments. *Frontiers in Neuroscience, 6*, 149.

Cui, W., Wang, X., & Riche, N. H. (2014). Let It Flow : a Static Method for Exploring Dynamic Graphs. 121–128, doi:https://doi.org/10.1109/PacificVis.2014.48.

de Ridder, M., Klein, K., & Kim, J (2015). CereVA-Visual Analysis of Functional Brain Connectivity. In *IVAPP,* 131–138.

Eklund, A., Nichols, T., & Knutsson, H. (2016). Can parametric statistical methods be trusted for fMRI based group studies? *PNAS, 113*(28), 7900–7905.

Filippi, M. (2016). *fMRI Techniques and Protocols*: Springer.

Filippi, M., & Filippi (2009). *fMRI techniques and protocols*: Springer.

FMRIB Analysis Group, O. U. (2016). FSL. http://fsl.fmrib.ox.ac.uk/.

Friston, K., Brown, H. R., Siemerkus, J., & Stephan, K. E. (2016). The dysconnection hypothesis (2016). *Schizophrenia Research, 176*(2), 83–94.

Fujiwara, T., Chou, J.-K., McCullough, A. M., Ranganath, C., & Ma, K.-L (2017). A visual analytics system for brain functional connectivity comparison across individuals, groups, and time points. In *Pacific Visualization Symposium (PacificVis), IEEE, 2017* (pp. 250-259): IEEE.

Giraldo-Chica, M., & Woodward, N. D. (2016). Review of thalamocortical resting-state fmri studies in schizophrenia. *Schizophrenia Research, 6*.

Gorgolewski, K. J., Varoquaux, G., Rivera, G., Schwartz, Y., Sochat, V. V., Ghosh, S. S., et al. (2016). NeuroVault. Org: A repository for sharing unthresholded statistical maps, parcellations, and atlases of the human brain. *Neuroimage, 124*, 1242–1244.

Irimia, A., Chambers, M. C., Torgerson, C. M., & Van Horn, J. D. (2012). Circular representation of human cortical networks for subject and population-level connectomic visualization. *Neuroimage, 60*(2), 1340–1351. https://doi.org/10.1016/j.neuroimage.2012.01.107.

Jezzard, P., Matthews, P., & Smith, S. (2001). *Functional MRI: an introduction to methods*: Oxford University Press.

Jie, B., Liu, M., Jiang, X., & Zhang, D. (2016) Sub-network Based Kernels for Brain Network Classification. In *Proceedings of the 7th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics,* (622–629): ACM.

Lee, M. H., Smyser, C. D., & Shimony, J. S. (2013). Resting-state fMRI: A review of methods and clinical applications. *AJNR. American Journal of Neuroradiology, 34*(10), 1866–1872. https://doi.org/10.3174/ajnr.A3263.

Liang, M., Zhou, Y., Jiang, T., Liu, Z., Tian, L., Liu, H., et al. (2006). Widespread functional disconnectivity in schizophrenia with resting-state functional magnetic resonance imaging. Neuroreport*,* 17(2), 209-213.

Liu, Y., Wang, K., Chunshui, Y. U., He, Y., Zhou, Y., Liang, M., et al. (2008). Regional homogeneity, functional connectivity and imaging markers of Alzheimer's disease: A review of resting-state fMRI studies. Neuropsychologia*,* 46(6), 1648-1656.

Liu, F., Xie, B., Wang, Y., Guo, W., Fouche, J.-P., Long, Z., et al. (2015). Characterization of post-traumatic stress disorder using resting-state fMRI with a multi-level parametric classification approach. Brain Topography*,* 28(2), 221-237.

Margulies, D. S., Böttger, J., Watanabe, A., & Gorgolewski, K. J. (2013). Visualizing the human connectome. *NeuroImage, 80*, 445–461. https://doi.org/10.1016/j.neuroimage.2013.04.111.

National Institute of Health (2016). AFNI. https://afni.nimh.nih.gov/.

Peeters, R., & Sunaert, S. (2007). Clinical BOLD fMRI: artifacts, tips and tricks. In *Clinical Functional MRI* (pp. 227-249): Springer.

Rashid, B., Arbabshirani, M. R., Damaraju, E., Cetin, M. S., Miller, R., Pearlson, G. D., et al. (2016). Classification of schizophrenia and bipolar patients using static and dynamic resting-state fMRI brain connectivity. Neuroimage*,* 134, 645-657.

Ristovski, G., Preusser, T., Hahn, H. K., & Linsen, L. (2014). Uncertainty in medical visualization: Towards a taxonomy. *Compters and Graphics, 39*, 60–73.

Sarraf, S., & Tofighi, G. (2016). Classification of alzheimer's disease using fmri data and deep learning convolutional neural networks. *arXiv Preprint arXiv, 1603*, 08631.

Sheline, Y. I., & Raichle, M. E. (2013). Resting state functional connectivity in preclinical Alzheimer's disease. *Biological Psychiatry, 74*(5), 340–347.

Sporns, O. (2010). *Networks of the Brain*: MIT Press.

Sporns, O. (2014). Contributions and challenges for network models in cognitive neuroscience. *Nature Neuroschience, 17*, 652–660.

Stevens, M. T. R., Darcy, R. C., Stroink, G., Clarke, D. B., & Beyea, S. D. (2013). Thresholds in fmri studies: Reliable for single subjects? *Journal of Neuroscience Methods, 219*(2), 312–323.

Swenson, R. (2006). Chapter 11: The Cerebral Cortex. In *Review of Clinical and Functional Neuroscience* (Vol. 1): Dartmouth Medical School.

Wang, S., Zhang, Y., Lv, L., Wu, R., Fan, X., Zhao, J., et al. (2017). Abnormal regional homogeneity as a potential imaging biomarker for adolescent-onset schizophrenia: A resting-state fMRI study and support vector machine analysis. *Schizophrenia Research*.

Woodward, N. D., Karbasforoushan, M. S., & Heckers, S. (2012). Thalamocortical dysconnectivity in schizophrenia. *American Journal of Psychiatry, 169*(10).

Xia, M., Wang, J., & He, Y. (2013). BrainNet viewer: A network visualization tool for human brain connectomics. *PLoS One, 8*(7), e68910.

Zang, Y., Jiang, T., Lu, Y., He, Y., & Tian, L. (2004). Regional homogeneity approach to fMRI data analysis. *Neuroimage, 22*(1), 394–400.

Zeng, H., Ramos, C. G., Nair, V. A., Hu, Y., Liao, J., La, C., et al. (2015). Regional homogeneity (ReHo) changes in new onset versus chronic benign epilepsy of childhood with centrotemporal spikes (BECTS): A resting state fMRI study. Epilepsy Research*,* 116, 79-85.