Small MultiPiles: Piling Time to Explore Temporal Patterns in Dynamic Networks

Benjamin Bach, Nathalie Henry-Riche, Tim Dwyer, Tara Madhyastha, Jean-Daniel Fekete, Thomas Grabowski

Small MultiPiles: Piling Time to Explore Temporal Patterns in Dynamic Networks

B. Bach[1], N. Henry-Riche[2], T. Dwyer[3], T. Madhyastha[4], J-D. Fekete[5] and T. Grabowski[4]

[1]Microsoft Research-Inria Joint Centre, France
[2]Microsoft Research, Redmond, WA, USA
[3]Monash University, Victoria, Australia
[4]University of Washington, WA, USA
[5]Inria, France

Abstract
We introduce MultiPiles, a visualization to explore time-series of dense, weighted networks. MultiPiles is based on the physical analogy of piling adjacency matrices, each one representing a single temporal snapshot. Common interfaces for visualizing dynamic networks use techniques such as: flipping/animation; small multiples; or summary views in isolation. Our proposed 'piling' metaphor presents a hybrid of these techniques, leveraging each one's advantages, as well as offering the ability to scale to networks with hundreds of temporal snapshots. While the MultiPiles technique is applicable to many domains, our prototype was initially designed to help neuroscientists investigate changes in brain connectivity networks over several hundred snapshots. The piling metaphor and associated interaction and visual encodings allowed neuroscientists to explore their data, prior to a statistical analysis. They detected high-level temporal patterns in individual networks and this helped them to formulate and reject several hypotheses.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Information Interfaces and Presentation]: User Interfaces—Graphical user interfaces.

1. Introduction

The human brain is an example of a highly-connected, dynamically-reconfiguring, weighted network [Spo11, DLBDS*06]. Activity in individual brain regions over time is obtained through functional magnetic resonance imaging (fMRI) and the correlation of activity between regions yields a measure for functional connectivity. The characterization of changes and temporal states in these functional brain networks is an open research challenge [BMZ*10]. It is hoped that such characterization can support neuroscientists researching neuro-degenerative diseases such as Alzheimer's or Parkinson's disease [HWA*13].

Other networks of the same type can be found in a variety of natural phenomena such as migration and trade between countries, social and transportation networks, as well as food webs and network traffic. While statistical models are widely employed to quantify and analyze such networks [LSS13],

Figure 1: *Piling as hybrid between small multiples and aggregation + flipbook.*

visualization is crucial for exploration as well as discussion and communication of findings. Furthermore, data contains uncertainty and the choice of a particular statistic requires a good overview of the data. Currently, there is no visualization that scales to the required number of time steps and connection density while simultaneously showing the evolution of topological patterns in states and over time.

To address these problems, we introduce *MultiPiles*, a visual tool for the exploration of changes in dense and weighted networks with a sequence of hundreds of tempo-

Figure 2: *Fuzzy states in brain activity. [RHVA08]*

ral samples (snapshots). The technique introduces a 'piling' metaphor to a small multiples display of matrix time series, i.e. *MultiPiles* is based on the physical analogy of interactively piling adjacency matrices, each one representing a snapshot in the network (Figure 1). Piles represent an arbitrary period of time, for example a stable state or a set of transitions between states. Thus, piling helps to organize the data and explore temporal periods, as well as visually aggregate information about the snapshots contained in each pile. While usually time is aggregated according to equidistant units (e.g., days, hours), MultiPiles allows the analyst to create piles based on change in the data. As a consequence, topological states can be quickly spotted and compared.

We implemented *MultiPiles* for a group of neuroscientists who investigate brain connectivity in patients with Parkinson's disease. We gathered feedback with interviews, observations and questionnaires from the scientists as they used *MultiPiles* to analyze their own data. The feedback shows that interactive and automatic piling helps to: gain insight into network states on different levels; describe time variant network states; and compare connectivity changes across individuals and conditions. The researchers reported that the interface is simple to use and efficiently supports formation and discussion of hypotheses in the exploration phase prior to a statistical analysis.

Section 2 describes challenges in understanding brain connectivity and Section 3 refers to related work on visualizations of dynamic networks and the piling metaphor in user interfaces. Section 4 describes MultiPiles based on a walkthrough scenario and Section 5 reports on insights and feedback from neuroscientists analysing their data with MultiPiles.

2. Functional Brain Connectivity Exploration

Functional brain connectivity represents temporal correlations between the blood-oxygen-level-dependent (BOLD) signal in different regions of the brain, as they evolve over time. These correlations can be observed in subjects who are resting in the fMRI scanner, and provide crucial information on the functioning of systems in the brain. While network based methods [Spo11] are an important method for analyzing functional connectivity data, many factors pose significant challenges for such analysis. For example: (i) connectivity data is typically dense (highly connected); (ii) connection weights due to correlation between regions give only a

fuzzy indication of connectivity; (iii) a patient's movements or general body functions such as heartbeat and breathing introduce noise into the data and is difficult to remove with preprocessing methods; (iv) an fMRI scan results in several hundreds of snaphots; (v) though it holds great promise, the study of dynamic functional connectivity is relatively recent and there is not yet full agreement about the correct statistical methods to analyze the data [HWA*13].

Due to this complexity and uncertainty, neuroscientists require tools to explore the data prior to a formal statistical analysis. This can help to formulate hypotheses, reject some of them early in the analysis process, and then to iterate over the results. Thus, statistical methods need to be carefully supervised through interactive techniques. Ideally, we would like to provide rapid exploration capabilities that enable neuroscientists to take advantage of their visual perception to quickly identify patterns of interest in their data. Such a tool should also allow them to investigate if these patterns appear with statistically significant frequency in different patient groups . While several visualization tools exist for fMRI data [XWH14, fls, LWX*14], they mostly rely on 3-dimensional brain representations and have not been designed for exploration purposes but rather to allow neuroscientists to illustrate their findings. Alper et al. [ABHR*13] demonstrated that visualization can have a strong potential to assist neuroscientists in performing high-level tasks on brain connectivity data. They propose a matrix encoding to compare two networks, enabling the comparison of two weighted networks, for example, pre- and post-surgery snapshots. In this paper, we tackle the problem of identifying patterns in entire dynamic networks and a large number of snapshots.

For our research in supporting neuroscientists to explore brain connectivity, we used a user-centered design methodology with T. Madhyastha and T. Grabowski, co-authors of the present paper, at the Integrated Brain Imaging Center (IBIC) at the University of Washington. After a series of four one hour interviews, we formalized the following set of tasks they would like to perform on functional connectivity network exploration. In terms of characterizing states, these tasks go beyond tasks that have previously been reported [APS14].

T1 **Identify temporal states**: i.e., distinct stable topological configurations indicated by a period of time with similar topological features (Figure 2). In many cases, such periods can be subdivided into finer granules.
T2 **Identify transitions between states**: i.e., periods of time where the state is changing (Figure 2).
T3 **Characterize topological features** in states and their evolution in transitions (trends, clusters, etc.).
T4 **Summarize** functional brain connectivity to a sequence of states and transitions (identify the ordering and recurrence of states and transitions)
T5 **Compare states, transitions** and their ordering across multiple conditions/individuals.

3. Related work

3.1. Visualizing Temporal Changes in Networks

Many techniques exist to visualize temporal changes in networks [BBDW14, BDA*14, KKC14]. The most common ones comprise *flipping* (or animating), *juxtaposing* (small multiples), and *summary views*. Each technique targets different user tasks and suits different datasets, but none of them effectively scales to networks with many snapshots or allows for the identification of temporal states. Small multiples are images of the network's topology at every snapshot placed juxtaposed on the screen for a quick overview and rough comparison of all snapshots [BBL12]. For dynamic networks, matrices [PS12], node-link diagrams [BBL12], and a technique related to node-link diagrams [BW14] have been used. However, flipping through snapshots, often accompanied by animated transitions [RM14, BPF14a], supports pre-attentive perception of subtle changes [Ren02, APP11, BBL12]. Flipping is helpful in cases where the data does not change much and identity of data elements is important. However, users must memorize previously seen data and navigate back and forth for comparison [BPF14a]. This results in longer observation times than with juxtaposing [APP11, FQ11] and requires users to memorize previously seen snapshots. Similarly, when juxtaposing the user is required to pan or scroll, if networks have a large number of snapshots. Alternatively, shrinking images in size results in lower resolution and less detail. In cases where comparison of individual snaphots is hard, using flipping or juxtaposition, summary views can show direct differences [ABHR*13], aggregate information from multiple snaphots [CKN*03] and encode temporal evolution into glyphs [SWS10, YEL10, BN11, FFM*13]. Summary views can though give a high-level and task-specific overview, if the chosen visual encoding is appropriate.

Despite the complementariness of these techniques, most interfaces either rely on one of these techniques or integrate them in a loose way through coordinated views (e.g., [RM13, BPF14a]). Bach et al. [BDA*14] describe how all of these techniques are related within the conceptual model of a space-time cube, exemplified for dynamic networks in Cubix [BPF14b]. While Cubix implements flipping, juxtaposing and a summary view, they are independent of each other and a user can only see one a time, e.g., the juxtaposing view shows all matrices, but does not allow for aggregation. In contrast, MultiPiles focuses on a subset of views and creates hybrid views, providing: i) interaction to create and edit periods of time (piles), ii) interactive techniques to explore and filter piles, iii) visual encodings summarizing changes within piles, iv) pile-specific node ordering that can be propagated to other piles, and v) automatic piling based on the amount of change in the data.

Following Beck's et al. taxonomy [BBDW14], MultiPiles can be categorized as *Layered Matrices* with a focus on networks with many snapshots (>100), where tasks involve identifying topological patterns and their recurrence. To that end, summary views summarize periods of the network, while flipping reveals subtle changes. Screen area is effectively used to show multiple juxtaposed summary views for overview and rapid comparison (Figure 1).

3.2. The Piling Metaphor

The piling metaphor has been used to organize documents in virtual desktop environments [MSW92], similar to documents on physical desktops. Users could manipulate and browse these piles as well as organize documents into piles depending on the documents' attributes. Users appreciated that metaphor and the ability to group object spatially, while memorizing their piles. BumpTop [AB06] adds interaction techniques and physical behaviour to piles of documents, when manipulated on a tabletop display. The operations include lasso selection, fanning-out, and sorting piles.

While MultiPiles also provides interaction to create and explore piles, the purpose of a pile is to reduce visual complexity of the data. A pile is meant to visually aggregate selections of information rather than just organizing objects. Recently, Sadana *et al.* [SMDS14] proposed a technique that superposes representations of sets to compare which elements are common across sets. Their *MultiLayers* are formed by direct manipulation using drag and drop. MultiLayers are similar to our piles, however we do not use piles for set logical comparison, but as organizational structure, visual aggregation, and to detect changes over time.

4. MultiPiles

While designing MultiPiles, we faced the following questions: 1) *What is the set of interactions required for navigation, exploration and pile management?* 2) *Which visual encodings are possible and useful?* and 3) *How can we create piles automatically, based on changes in the data?* In this section, we describe our design rationale for the questions above and illustrate the set of features offered by MultiPiles via a scenario in which a neuroscientist explores functional brain connectivity data extracted from an fMRI scan. We give examples of patterns found with the system, and show how MultiPiles supports the tasks listed in the end of Section 2. For the sake of simplicity and space in this article, we used an example network of 29 nodes evolving only 95 time points, out of 300.

After loading a network, MultiPiles shows all matrices as small multiples (Figure 3). The overview gives an idea about the different periods in the network. MultiPiles' interface consists of three main components that we present below: the *piles plot* (d), the *piling slider* (b) and the *timeline* (c). A menu bar allows for saving and retrieving a user's sessions and applying operations to all piles (a). While the piling metaphor can potentially support any visualization including node-link diagrams [BBL12, BW14], we selected matrices in this article for three reasons: i) no visual clutter for

Figure 4: *Eight piles indicating topological states in brain connectivity. The piles have been created by our temporal clustering heuristic. Users can adapt the piling slider (right) to define the number of piles they want to obtain.*

Figure 3: *MultiPiles user interface: a) Main menu, b) Piling slider, c) Timeline, d) Piles plot showing adjacency matrices.*

dense networks (such as brain connectivity networks), ii) efficient encoding of edge weight [ABHR*13] with the ability to use temporal glyphs (e.g., [BN11]), and iii) easier visual comparison of topological patterns.

4.1. Piles

Our fictional neuroscientist, Emma, is going to explore this data attempting to complete tasks described earlier. Many of the matrices in Figure 3 look similar to each other, potentially indicating states of the network. To create piles of similar matrices, Emma adjusts the piling slider to 8 as shown in Figure 4, resulting in 8 piles calculated by our temporal cluster heuristic (see Section 4.2). These are 8 temporal periods in the network (T1), where matrices are most similar, i.e., the mutual euclidean distance between matrices of a pile is below a given threshold. The temporal order of matrices is always preserved to identify sub states and transitions.

The adjacency matrix "top-most" on each pile is called *cover matrix* and summarizes the snapshots in that pile. For example, cover matrices in Figure 4 show the mean value for any connection in the period represented by the pile; darker values indicate a higher mean weight (strong correlation between regions). Immediately above the cover matrix *matrix previews* summarize the content of the pile. Similar to the

edge of a book, a preview represents one matrix inside the pile as a horizontal bar in which gray values indicate node-degree for every node. Stacking matrix previews above each pile is in accordance with the metaphor of piling and enables quick pile size comparison. We tried stacking matrix previews on the right hand side of each pile, which is consistent with the direction of time, but we found that option harder to navigate and compare piles.

The automatic clustering creates homogeneous piles that can be interpreted as stable periods of the network. However, Emma can manually create or split piles according to other semantics. For example, she may decide to create piles according to tasks the subject was performing during the scan. As Emma looks at these piles, she notices that states represented in Figure 4 all exhibit a different topology (T5). Four of these piles are rather dense (piles 1,2,5,7) while others are sparser and exhibit clusters (4, 6). Pile 2 represents the longest state (largest pile) and Emma notes that dense states seem to generally last longer in this data. Piles 4, 8, and possibly 6 look very similar, which may indicate a recurring state in functional brain connectivity for this individual (T4).

4.2. Piling Time

The *timeline* (Figure 5) complements the piles plot by providing visual clues about changes in the network. This compact overview can serve to manually segment the network into a sequence of piles. We describe its encoding and interactions below

The timeline is a heatmap presenting a summary of every node's connectivity over time, but other centrality metrics and node attributes are applicable. Each node in the network is encoded as a row and each snapshot as a column. Thus, a single cell at row i and column t encodes the weighted degree of the ith node at the tth time step. This visual encoding is the same as used in the matrix previews of the piles. Thus, from left to right, a row in the timeline shows the evolution of connectivity of one specific node to all other nodes. Figure 6 shows a detail from the timeline in Figure 5 where one can identify that the brain region labeled *SALLFIC* is strongly connected to others only sporadically while region *SALRFIC*[†] is connected through most of the timespan.

---

† SALLFRIC and SALRFIC are the left and right frontal insular cortices, part of the brain's *salience network* [SMS*07].

(a) All time steps in a *single* pile

(b) 8 piles

(c) 29 piles

Figure 5: *Timeline with different automatic pilings based on the similarity between matrices. White vertical lines separate piles.*

Figure 6: *Node degree encoding on the timeline.*

Vertical columns of the timeline summarize the degree of all nodes at a given time. Thus, it is possible to identify changes in the grayscale of series of vertical columns. For example, Figure 5(a) shows periods of higher connectivity (dark columns) alternating with periods of lower connectivity (bright columns).

Using the *piling slider* (Figure 4, right), Emma can adapt the distance threshold that is used to create piles. As illustrated in Figure 7, a high threshold increases the level of aggregation and results in fewer piles, each one containing many matrices on average. Likewise, decreasing this threshold results in more and smaller piles (fewer matrices per pile), but with a higher similarity within each pile. Our hierarchical clustering is similar to [vWvS99]. Shown in Figure 5(a), Emma first creates a single pile, containing all matrices. Then she adjusts the piling slider down to split that large pile into smaller piles, indicated by vertical white lines in the timeline (Figure 5(b)). These piles are the same as shown in Figure 4. Emma can investigate finer grain piling as in Figure 5(c) by adjusting the slider to a lower threshold.

Figure 7: *Hierarchically aggregating matrices into piles. Each horizontal layers show the piles and matrices shown at the same moment. Moving the piling up results in fewer piles, moving the slider down, results in more piles/matrices.*

Automatic piling in MultiPiles is based on the Euclidean distance d between pairs of matrices. It offers two clustering heuristics each parameterized by a threshold p controllable with the slider:

- **Sequential piling** considers only distances between *adjacent* snaphots t_i and t_{i+1} and yields piles where maximal distance between adjacent matrices is less than p. If $d < p$, t_{i+1} is added to the pile of t_i. If $d \geq p$, a separate pile for t_{i+1} is created.
- **Clustered piling** considers distances between *all* matrices and yields piles where distances between all contained matrices are below the threshold p, i.e., for a set T of sequential snapshots $\{t_i, t_{i+1}, t_{i+2}, ...\}$, T becomes a pile, if $d < p$ for all pairs of snaphots in T. Clustered piling is more restrictive and therefore highlights coherent states with little change.

4.3. Exploring Piles

The timeline view and the piles plot are coordinated by brushing-and-linking. Hovering over any column in the timeline brings the corresponding matrix to the top of its pile. That is, the cover matrix summarizing the pile is replaced by the actual matrix at the selected time step. Using this interaction over the timeline allows Emma to quickly browse through all matrices in the piles plot.

For detailed browsing of matrices within a pile, Emma can interact with the pile directly. Hovering a pile/matrix while pressing shift shows node labels. Hovering over the horizontal lines above the pile (i.e., the matrix previews) flips through matrices within the pile and highlights the corresponding column in the timeline (Figure 8). Flipping gives an impression of the trend in connectivity, e.g. stable connections or emerging clusters (T2). Moreover, it can reveal when a pile is not as homogeneous as expected and may require further splitting. For example, see Figure 8, Emma

Figure 8: *Flip through pile by hovering matrix previews above the cover matrix. From left to right: cover matrix showing mean values, hover matrices 59, 60, and 61.*

observes that between snapshots 59 and 61, the network changes faster than in the rest of the pile and she decides to split the pile (Section 4.4).

Emma now wishes to further explore how connections change within each pile to see *which nodes are involved in patterns* and *compare these patterns between piles*. Cover matrices convey the network topology present in an entire pile via four visual representations described below and illustrated in Figure 9. These options can be changed interactively via the pile contextual menu that appears on every pile on mouse over.

| (a) Mean | (b) Trend | (c) Variation | (d) Bar Charts |

Figure 9: *Examples of visualizing aggregated information in cover matrices, for the same pile. (a) Mean value, (b) Trend (red=increase, blue=decrease), (c) variation, (d) bar charts.*

- **Mean**—Each cell in the cover matrix is shaded based on the mean of the values for that cell in all the matrices in the pile. The pile in Figure 9(a) shows two strong clusters during that time. Cover matrices showing the mean values use a light blue shading, in order to differentiate them from an individual matrix (black).
- **Trend**—A blue cell indicates a decrease in strength for that correlation within the time period of the pile, red means increase (i.e., linear regression). The same pile shows two different sets of regions, one that increase connectivity, and one that decreases (Figure 9(b)).
- **Variation**—The variation cover matrix shows the standard deviation for each connection. Its standard color is dark green to distinguish it from other cover matrices. Darker cells indicate higher deviation. Figure 9(c) shows that one node's connectivity varies more than others.
- **Bar Charts**—While the previous three methods aggregate data, we provide one example of a method that nests timelines in each matrix cell, similar to glyphs [YEL10, BN11]. Figure 9(d) shows a detail from a pile and shows a bar chart in each cell that represents the connection strength at any time point (from left to right).

4.4. Interactive Piling

Automatic pile aggregation is an easy way for users to get a first overview. Browsing piles and observing cover matrices allows for quick exploration. However, automatic piling relies on a specific distance metric that may be *sensitive to noise*, and that users must *understand and interpret*. Introducing a "human in the loop" can alleviate these caveats. For example, the human can spot an outlier causing the algorithm to split a pile in two and decide to ignore this noise in a subsequent analysis.

Thus, MultiPiles provides a set of operations, summarized and explained in Figure 10, to allow for the interactive creation and exploration of piles, using the timeline or the piles plot. These operations are designed to support two major strategies, summarized in Figure 1: i) gain a quick overview over all time steps arranged as small multiples, then gradually create piles by grouping similar matrices together (**bottom-up**), or ii) start from a single pile and, while flipping through the matrices, split it into smaller piles when specific changes occur (**top-down**).

Automatic and interactive piling are designed to work in concert. However, while interactive piling can serve to adjust the automatic piling, adjusting the piling slider after a manual piling has been performed will replace all current piles by the calculated piles (see Figure 7). To avoid losing potentially precious piling adjustments, MultiPiles provides a history tracking mechanism allowing users to undo any operation performed. A user's current piling can also be saved on the local machine for later recovery, as well as exported as plain text to be recorded, loaded on a different machine, or send to a collaborator.

4.5. Finding States in Individual Regions

While automatic piling takes into account all regions and connections equally, with manual piling Emma is able to pay attention to particular brain regions and subnetworks. For example, she wants to create piles according to the activity in a single region which she knows is involved in a particular cognitive task. She wants to know (a) *when is this region active?*, (b) *which other regions have similar patterns of activity?*, and (c) *which topologies are present during the region's active periods?*.

Below is an example workflow to answer these questions: (1) Emma creates a single pile using the piling slider, (2) she clicks on the label of the required region in the timeline and obtains the upper picture in Figure 11. The timeline no-longer shows node degree in rows, but the connection strength between the selected brain region to all other regions in the network; the picture shows *how much* the selected region is correlated with *which* other region and *when*. The selected region is strongly correlated to other regions at various times. (3) Emma can create piles interactively on the timeline (Section 4.4), or if she moves the piling slider, the automatic piling only takes into account the states

- **Push (1)** and **Pop (2)**—Matrices can be dragged and drop onto each other to form a pile.

- **Split pile (3)**—Flipping through a pile helps perceive suble changes and the user can split the pile at this time point by a simple click on the matrix preview on the pile or the corresponding column in the timeline (similar to the cutting operation in video editing software). The new pile, containing the selected matrix and successors, is inserted to the right of the current pile using an animated transition.

- **Combine (4)**—Two piles are combined if, in the timeline, the user clicks on the first time step that belongs to the second pile, i.e. the snaphot to the right of any white bar. On the piles plot, clicking on the gap after a pile, piles both previous piles.

- **Pile (5)**—Lasso over a sequence of matrices or piles merges them into a single pile. Alternatively, the user can click in between two matrices which piles all matrices between the gap and the last pile.

- **De-pile (6)**—All a pile's matrices are released if the user double clicks on a pile.

- **Flip-through (7)**—Mousing over the matrix previews as illustrated in Figure 8 shows each individual matrix.

- **Pile preview (8)**—Comparing patterns in the first and last matrix of the piles for example may prove difficult without actually splitting the pile. MultiPiles supports a temporarily depiled mode activated by a simple click on the cover matrix. This places all matrices contained in the pile side-by-side, as small multiples. Clicking on any of these matrices restores the initial pile.

- **Zoom (9)** and **Pan (10)**—Matrices can be zoomed and panned to show a set of node-pairs and their relations only. Pan works by dragging on the mouse, zoom with the mouse wheel on the matrix. In both cases, Shift must be held.

- **Cover Matrices (11)**—Cover matrices are switched through the pile contextual menu (individually for each pile), or by buttons in the interface (for all piles).

- **Order (12)**—Applies different row and column orderings to matrices in a pile, available through the pile contextual menu.

- **Propagate (13)**—Propagates the exact order of rows and columns from one pile, to all other piles.

Figure 10: *Operations on piles including formation, exploration, aggregation, filtering.*

Figure 11: *Hovering over a node on the left side of the timeline shows only this node's connections (large node label). Top and bottom timelines show two different nodes exhibiting (a) asynchronous, (b) inverse or (c) synchronous activity.*

present in the selected node. (4) On the piles plot, Emma can now switch to a different cover matrix type (e.g., variation) and see whether states identified for the selected region match states of the entire network (Figure 12(a)). Low variance shows where general connectivity is correlated with the states identified for the selected region, high variance shows where the general connectivity is not correlated.

Emma can further use the timeline to compare correlation across two regions by moving the mouse over the node labels and flip through the nodes. The lower timeline in Figure 11 shows a different brain region and its correlations. These figures reveal periods where both regions are asynchonous (a), synchonous (c), and inverted (b), i.e., one region is generally active, while another one is inactive. Yet, both regions are correlated to the same subset of regions during the same period (d).

4.6. Identifying Subnetworks

A set of regions and their connectivity is called a subnetwork. Subnetworks can describe anatomically close regions (e.g., regions in the same cortex), regions responsible for the same task, or clusters of regions that are connected by a high correlation.

In our scenario, Emma is interested in the regions of a certain cortex, which all start with the letters *DAN*. From the 6 top most rows in Figure 5(a) she sees that those regions are generally active during the entire period of the fMRI scan, with some certain drops in activity around the middle of the scan. She wants to have a closer look on those regions and how they are connected over time (T3). She selects the regions from the cortex in the timeline by dragging over them with the mouse. Piles and matrices in the piles plot get filtered and show only the selected regions. Likewise, any automatic piling will now take only the selected regions into account, based on the matrices shown on the piles plot. Figure 12 shows the matrices for the selected regions only, after an automatic piling. Emma decides to show variation and then trends in the cover matrix, which shows alternating phases of activity. To investigate regions currently filtered, Emma *pans* inside the matrix, by pressing *shift* while dragging inside the matrix.

4.7. Node Ordering in Matrices

By default, all matrices and piles have the same ordering of nodes. Emma already provided an ordering with the data, a grouping of brain regions into higher level cortices. This ordering helps her to preserve some characteristics about her data. If no such ordering is given, or if the given ordering does not show any patterns of interest, MultiPiles can calculate a global ordering which tries to be find a topological clustering across all matrices. Yet, this ordering may not be optimal for any individual matrix or pile (Figure 13(a)).

Figure 12: *Zooming into the matrices, showing only 8 regions and their connections. Cover matrices show standard deviation a) and alternating trends in red and blue b).*

| (a) Pile 4 | (b) Pile 4 | (c) Pile 8 | (d) Pile 8 |

Figure 13: *Local ordering in pile 4 and 8 from Figure 3. (a) global ordering for pile 4, (b) optimal (local) ordering for pile 4, (c) optimal (local) ordering for pile 8, and (d) optimal ordering from pile 4, applied to pile 8.*

After having observed one subnetwork, Emma wants to better understand how different subnetworks are connected; especially, the similarity between states 4 and 8 in Figure 4. In the pile menu, she clicks on the command for "Local Ordering" which tries to optimize the ordering for all matrices in this pile. Figures 13(b) shows pile 4 with its optimal node ordering. Any pile can be ordered locally to spot individual patterns, but it is then impossible to compare the topologies across piles (See Figures 13(b) and 13(c)). To compare pile 4 and 8, she selects 'Propagate Ordering' from the pile menu for pile 4 (figure 13(d)). Node ordering is now the same in both piles and she immediately sees which regions differ from the two main clusters.

5. Evaluation

During the six months design and iterative implementation of MultiPiles, we had regular sessions for feedback and discussion with our two co-author neuroscientists to refine the prototype. After we converged towards a working prototype, we performed a 1.5 hour session of collaborative analysis with Multiples using functional brain connectivity data from one of their studies on Parkinson's disease. Our collaborators then used the tool for several days in close collaboration with us and in conjunction with statistical packages to pursue their exploration.

5.1. Collaborative Analysis

The collaborative analysis sessions involved guided use of MultiPiles with data from one of their current studies. It only took a few interactions to investigate if a given network exhibited states and whether these states tend to re-occur during the scan. Although not directly interacting herself, our collaborators quickly grasped the idea behind piling and MultiPiles and were able to instruct us which operations to perform. The series of interactions consisted in using the slider to 1) automatically create piles of matrices, 2) browse through these piles to visually assess their homogeneity, 3) split or merge piles when judged necessary, 4) re-order one of the bigger piles to exhibit a clearer pattern in the matrix and, finally 5) propagate this ordering to other piles to assess if a similar pattern would emerge in the rest of the scan as well as in the data of other individuals.

As we explored and compared data from and across several other individuals (T5), our collaborators formed (and rejected) several hypotheses. One of them pointed out that this tool would possibly save her a lot of time, explaining that probably a large majority of the hypotheses she formed on the fly would be rejected right away, early in the workflow, as compared to existing analysis pipelines. Most comments we gathered on the usefulness of the system during these sessions centered around the simple mechanisms MultiPiles offers to create states (piles) and perform an initial assessment on the plausibility of these states.

5.2. Preliminary Findings

Following this first collaborative session, the neuroscientists used MultiPiles over several days to more systematically compare 24 subjects with Parkinson's disease and 21 healthy controls. Figure 14 shows two representative networks from each group: patients on the top, controls on the bottom. Row and column ordering is the same across piles and individuals. The same piling threshold was applied to every individual.

Based on their knowledge, our collaborators hypothesized that controls would exhibit more changes in functional connectivity throughout the scan, and thus, that MultiPiles would produce a larger number of piles for controls than subjects for a given similarity threshold. Exploring different similarity thresholds using the slider, they were surprised that the number of states in both groups did not seem to differ significantly. As they browsed through the piles, however, they could identify that piles looked more homogeneous in Parkinsons' disease subject data. To support their investigation further and enable statistical analysis, we augmented MultiPiles to output a number of statistics from each piling configuration. Using a piling threshold of 5.5 (euclidean distance), determined a priori as representing a mean change of 25% in all correlations within a matrix, they identified significant differences in the piling patterns between participants with Parkinson's disease and controls. As hypothesized earlier, controls did show significant patterns in their

piling, if not in their number. They had higher variability in the size of their piles as evidenced by greater standard deviation (1.71 versus 1.44, p < .001) and a larger maximum pile size (9.29 versus 8.25, p =.002). By augmenting MultiPiles with a minimal set of confirmatory features (i.e. outputting a set of measures on piles), we could support our analysts in rejecting/validating their hypotheses.

During these explorations, MultiPiles also helped our collaborators identify new hypotheses. Browsing the piling configuration of different datasets and experimenting with different pile orderings revealed a specific topological pattern, highlighted in Figure 14(a), that seemed to occur more in Parkinson's subjects than controls. As a step towards confirming this observation, three independent coders visually compared all 45 scans and coded piles for the presence of this topological pattern. They used the exact same piling setup as described earlier and were given a reference matrix exhibiting the pattern. Statistical analysis of this coding session (t(40)=2.28,p=.027); rater ICC=.78, indicates the pattern can be identified by individual coders and that it appears more in subjects than controls. We are currently working towards distance metrics capable of characterizing such patterns. While this finding warrants further research to understand its significance, we demonstrated that MultiPiles can lead to discoveries in dense weighted dynamic networks.

6. Conclusion and Future Work

We presented MultiPiles, a system to visualize dense dynamic networks with changing edge weight and many snapshots. We reduce visual complexity by segmenting dynamic networks into more manageable parts (piles) and through a set of novel interactive matrix views, users are able to explore temporal states. States can refer to relatively stable periods, transitions between stable periods, or any other kind of period. MultiPiles has been tested with networks of 35 nodes and up to 300 time steps. The implementation is in WebGL and user interactions are responsive. Networks with more nodes result in larger matrices, but pan, zoom and filtering helps users to find a tradeoff between topological detail and temporal overview.

MultiPiles is used at the Integrated Brain Imaging Center at the University of Washington, allowing neuroscientists to explore their data interactively for the first time. We believe it is technically easy to integrate MultiPiles into a general workflow with better support for analytical measures. Multi-Piles is available online[‡] and is not limited to brain connectivity networks but can visualize any kind of dynamic network data such as migration and trading flows, social networks or other bio-medical networks. We believe that the piling metaphor and mechanism can be useful to other temporal data sets such as video (frames) or geo-spatial data.

[‡] `www.visualizingbrainconnectivity.org/multipiles`

(a) Patients

(b) Controls

Figure 14: *2 Representative scans from each group: patients (a) and controls (b). We observe the occurence of the black framed pattern across individuals.*

While there is room for extending each of the techniques presented in this paper (e.g., cover matrices, direct manipulation, timeline design), we intend to focus on three major aspects of future work: i) exploring alternative automatic clustering heuristics and providing a better rendition of temporal states in dynamic networks, ii) finding appropriate ways to visualize similarity between multiple snapshots and piles, and iii) comparing multiple dynamic networks (e.g., patients) with MultiPiles.

Acknowledgements

References

[AB06] AGARAWALA A., BALAKRISHNAN R.: Keepin' it real: Pushing the desktop metaphor with physics, piles and the pen. In *Proc. of CHI* (2006), ACM, pp. 1283–1292. 3

[ABHR*13] ALPER B., BACH B., HENRY RICHE N., ISENBERG T., FEKETE J.-D.: Weighted graph comparison techniques for brain connectivity analysis. In *Proc. of CHI* (2013), pp. 483–492. 2, 3, 4

[APP11] ARCHAMBAULT D., PURCHASE H. C., PINAUD B.: Animation, small multiples, and the effect of mental map preservation in dynamic graphs. *IEEE TVCG 17*, 4 (2011), 539–552. 3

[APS14] AHN J.-w., PLAISANT C., SHNEIDERMAN B.: A task taxonomy for network evolution analysis. *IEEE TVCG 20*, 3 (2014), 365–376. 2

[BBDW14] BECK F., BURCH M., DIEHL S., WEISKOPF D.: The State of the Art in Visualizing Dynamic Graphs. In *Proc. of EuroVis* (2014). 3

[BBL12] BOYANDIN I., BERTINI E., LALANNE D.: A Qualitative Study on the Exploration of Temporal Changes in Flow Maps with Animation and Small-Multiples. *Computer Graphics Forum 31*, 3pt2 (2012), 1005–1014. 3

[BDA*14] BACH B., DRAGICEVIC P., ARCHAMBAULT D., HURTER C., CARPENDALE S.: A Review of Temporal Data Visualizations Based on Space-Time Cube Operations. In *Proc. of EuroVis* (2014). 3

[BMZ*10] BISWAL B. B., MENNES M., ZUO X.-N., GOHEL S., KELLY C., SMITH S. M., BECKMANN C. F., ADELSTEIN J. S., BUCKNER R. L., COLCOMBE S., ET AL.: Toward discovery science of human brain function. *Proc. of National Academy of Sciences 107*, 10 (2010), 4734–4739. 1

[BN11] BRANDES U., NICK B.: Asymmetric Relations in Longitudinal Social Networks. *IEEE TVCG 17*, 12 (Dec. 2011), 2283–2290. 3, 4, 6

[BPF14a] BACH B., PIETRIGA E., FEKETE J.-D.: GraphDiaries: Animated Transitions and Temporal Navigation for Dynamic Networks. *IEEE TVCG 20*, 5 (2014), 740 – 754. 3

[BPF14b] BACH B., PIETRIGA E., FEKETE J.-D.: Visualizing Dynamic Networks with Matrix Cubes. In *Proc. of CHI* (2014), ACM, pp. 877–886. 3

[BW14] BURCH M., WEISKOPF D.: A flip-book of edge-splatted small multiples for visualizing dynamic graphs. In *Proc. of Visual Information Communication and Interaction* (2014), ACM, pp. 29–38. 3

[CKN*03] COLLBERG C., KOBOUROV S., NAGRA J., PITTS J., WAMPLER K.: A system for graph-based visualization of the evolution of software. In *Proc. of SoftVis* (2003), ACM, pp. 77–ff. 3

[DLBDS*06] DE LUCA M., BECKMANN C., DE STEFANO N., MATTHEWS P., SMITH S. M.: fmri resting state networks define distinct modes of long-distance interactions in the human brain. *Neuroimage 29*, 4 (2006), 1359–1367. 1

[FFM*13] FUCHS J., FISCHER F., MANSMANN F., BERTINI E., ISENBERG P.: Evaluation of alternative glyph designs for time series data in a small multiple setting. In *Proc. of CHI* (2013), ACM, pp. 3237–3246. 3

[fls] fslview - Viewer for (f)MRI and DTI data. http://neuro.debian.net/pkgs/fslview.html. online, acessed Dec. 4th, 2014. 2

[FQ11] FARRUGIA M., QUIGLEY A.: Effective temporal graph layout: A comporative stydy of animations versus status display methods. *Information Visualization* (2011), 47–64. 3

[HWA*13] HUTCHISON R. M., WOMELSDORF T., ALLEN E. A., BANDETTINI P. A., ET AL.: Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage 80*, 0 (2013), 360 – 378. 1, 2

[KKC14] KERRACHER N., KENNEDY J., CHALMERS K.: The Design Space of Temporal Graph Visualisation. In *Proc. of EuroVis* (2014), Elmqvist N., Hlawitschka M., Kennedy J., (Eds.), Eurographics Association. 3

[LSS13] LEE M., SMYSER C., SHIMONY J.: Resting-state fMRI: a review of methods and clinical applications. *Americal Journal of Neuroradiology* (2013), 1866–1872. 1

[LWX*14] LIAO W., WU G., XU Q., JI G., ZHANG Z., ZANG Y., LU G.: DynamicBC: A MATLAB Toolbox for Dynamic Brain Connectome Analysis. *Brain Connectivity* (2014). 2

[MSW92] MANDER R., SALOMON G., WONG Y. Y.: A pile metaphor for supporting casual organization of information. In *Proc. of CHI* (1992), ACM, pp. 627–634. 3

[PS12] PERER A., SUN J.: MatrixFlow: temporal network visual analytics to track symptom evolution during disease progression. In *Proc. of AIMA annual symposium* (2012), American Medical Informatics Association. 3

[Ren02] RENKSINK R.: Change Detection. *Annual Review of Psycholoy* (2002). 3

[RHVA08] RABINOVICH M., HUERTA R., VARONA P., AFRAIMOVICH V.: Transient cognitive dynamics, metastability, and decision making. *PLoS Computational Biology*, 4 (2008). 2

[RM13] RUFIANGE S., McGUFFIN M. J.: DiffAni: Visualizing Dynamic Graphs with a Hybrid of Difference Maps and Animation. *IEEE TVCG 19*, 12 (2013), 2556–2565. 3

[RM14] RUFIANGE S., MELANCON G.: AniMatrix: A Matrix-Based Visualization of Software Evolution. In *Proc. of Working Conference on Software Visualization* (2014). 3

[SMDS14] SADANA R., MAJOR T., DOVE A., STASKO J.: On-Set: A Visualization Technique for Large-scale Binary Set Data. In *Proc. of VIS* (2014). 3

[SMS*07] SEELEY W. W., MENON V., SCHATZBERG A. F., KELLER J., GLOVER G. H., KENNA H., REISS A. L., GREICIUS M. D.: Dissociable intrinsic connectivity networks for salience processing and executive control. *The Journal of neuroscience 27*, 9 (2007), 2349–2356. 4

[Spo11] SPORNS O.: *Brain Networks*. MIT Press, 2011. 1, 2

[SWS10] STEIN K., WEGENER R., SCHLIEDER C.: Pixel-oriented visualization of change in social networks. *Proc. of International Conference on Advances in Social Networks Analysis and Mining 6* (2010), 233–240. 3

[vWvS99] VAN WIJK J., VAN SELOW E.: Cluster and calendar based visualization of time series data. In *Proc. of InfoVis* (1999), pp. 4–9, 140. 5

[XWH14] XIA M., WANG J., HE Y.: BrainNet Viewer: A Network Visualization Tool for Human Brain Connectomics. *PLoS ONE*, 8 (2014). 2

[YEL10] YI J.-S., ELMQVIST N., LEE S.: TimeMatrix: Visualizing Temporal Social Networks Using Interactive Matrix-Based Visualizations. In *International Journal of Human-Computer Interaction* (2010), vol. 26, pp. 1031–1051. 3, 6