SI2-SSI: Collaborative Research: Cognition-aware Visual Analytics of Brain Circuits

David H. Laidlaw (PI), Caroline Ziemkiewicz, Brown University Computer Science;
Steven Sloman (Co-PI), David Badre (Co-PI),
Brown University Cognitive, Linguistic, and Psychological Sciences;
Mark Schnitzer (Co-PI), Stanford Biological Sciences and Applied Physics and
Investigator of the Howard Hughes Medical Institute;
Jeff Chi-Tat Law, Stanford Biological Sciences;
Fritz Drury, Rhode Island School of Design

We propose to develop, test, and deploy software tools for scientific study of brain circuits. The driving scientific application for our research is the study of human brain function and architecture. Major advances in our understanding of how brains work have occurred in recent decades, yet much remains unknown. Network models of the brain are natural because they often reflect both the behavior and the anatomy of the brain. They also provide a valuable abstraction for the huge quantities of imaging data that arise from experimentation. The target user community is brain scientists studying architecture and function of the human brain. Two of the four labs involved in this project are a part of this community.

Our aim is to blend good design and software engineering with research to extend and augment the process using principles derived from perceptual and cognitive psychology. The tools we propose will help brain researchers pose and test hypotheses about networks using input from multiple sources. These sources will include databases of published results about networks, data from their own experiments, and the scientists' hypotheses and working assumptions. Experimental data will include 3D imaging data from functional and diffusion MRI as well as light microscopy.

The software we propose will support the reasoning process by gathering and managing data about networks and connections; gathering and managing imaging and experimental data; providing interactive mechanisms for visually selecting and analyzing portions of the data; and explicitly capturing, recording, and documenting users' scientific reasoning.

Intellectual Merit The intellectual merit of this project is threefold. First, the proposed infrastructure will enable brain scientists to advance their research agendas more efficiently and more quickly by incorporating information from a broader set of sources into their scientific reasoning. Second, computer scientists will advance their understanding of how humans interact with computer systems at the cognitive level and how that can improve those interactions. Third, we anticipate advances in the understanding of human cognition from our cognitive modeling and experimentation.

Broader Impact. This highly interdisciplinary project will demonstrate to undergraduates, graduate students, and postdoctoral scholars how such research can be done. Because the tools will be made widely available, they will potentially benefit of the entire brain science research community. Many other disciplines study linked types of data – gene regulation, protein signaling and even crime and terrorism analysis – and all have the potential to benefit. Any human computer interface that involves reasoning has the potential to be improved by results from this research. Finally, the process of designing and developing effective interfaces for humans to use computers may well be broadly improved.

SI2-SSI: Collaborative Research: Cognition-aware Visual Analytics of Brain Circuits

David H. Laidlaw (PI), Caroline Ziemkiewicz, Brown University Computer Science;
Steven Sloman (Co-PI), David Badre (Co-PI),
Brown University Cognitive, Linguistic, and Psychological Sciences;
Mark Schnitzer (Co-PI), Stanford Biological Sciences and Applied Physics and
Investigator of the Howard Hughes Medical Institute;
Jeff Chi-Tat Law, Stanford Biological Sciences;
Fritz Drury, Rhode Island School of Design

a Multidisciplinary Research Agenda

Understanding the human brain is a daunting enterprise. The thousands of brain researchers each face a difficult problem: they have to interpret their voluminous data in the context of everyone else's voluminous data. More details illustrating this pressing need, including specific examples from our labs, are in Sec. b.

We propose to develop software tools to make this enterprise less daunting and more productive. Our experience in the visualization research laboratory within the computer science department at Brown has involved successfully developing software tools in support of numerous scientific users [JDL09a; JYN^{+}09; KOB^{+}08; CJPS^{+}06; Zha06; MCL09]. We have also considered deeply the design and evaluation process of such tools[AJLD08; KAM^{+}08; AJLD05; KL07]. We are well poised to attack this problem.

A significant part of the proposed research will involve incorporating principles of perception and cognition into the design and evaluation of the proposed software. Most software is developed to address relatively low level workflows, but we believe that significant improvements in productivity can be gained by optimizing tools using such principles. Scientists within Brown's Department of cognitive, linguistic, and psychological sciences have a deep understanding of these areas of knowledge and will be instrumental in our software development project. They have knowledge about how scientists think that are developed and specific enough to provide guidance about what functionality to include in the tools and how to design them. This knowledge concerns how scientists test hypotheses, how they represent data, and what kind of neural processing systems they have.

Our development approach will follow a typical spiral software engineering design process, with fast iteration on requirements and prototypes in the beginning and gradually slower iterations solidifying the software as we incorporate feedback and evaluation at multiple levels. Feedback and evaluation will occur at multiple levels including local feedback on sketched designs, focus groups, analysis of video and tracking data, and formal experiments. This evaluation will help us refine not only the tools but also some of the underlying cognitive principles.

The tools will be distributed both in binary and source form. We will endeavor to build a community of users and developers who will help to motivate and realize sustainability.

This effort is clearly an ambitious one. But we have gathered a small group with a large breadth of experience to attack this important problem. While we acknowledge that there is significant risk in what we propose, we believe that the potential rewards balance that risk. Even if we are not completely successful, we believe that the tools will still accelerate brain science as well as adding to human knowledge about how the brain works.

b Need for Visual Analysis of Brain Networks

The human brain is one of the most complex organ in our body. It contains billions of neurons that form interconnected networks at different scales, from small networks with tens of neurons to large networks with long-range connections that span across multiple brain areas. These neural networks are thought to subserve many important functions, from perception, cognition to learning and memory. However, our understanding of the neural basis underlying these brain functions are limited for the following reasons. First, a brain area typically receives both external inputs from other brain areas and internal inputs from within the same area. The inputs fibers from these two sources are both dense and intermingle. Therefore, it is very difficult to study connectivity and dynamics of neural circuits in intact brain. Second, because the brain wiring is extremely complex, researchers often finds it difficult to interpret the data without the help of software tools.

Recent advances in molecular and imaging techniques have opened new possibilities in studying neural connectivity in intact brain. These studies use magnetic resonance imaging or fluorescence microscopy to trace neural connections and provide increasingly detailed information about neuronal wiring from within and between brain areas. Such information is extremely valuable because it allows researchers to study how information from external areas is transformed by intrinsic circuits in a brain area to generate new outputs. However, the amount of data generated from these studies only exacerbates the problems in data analysis. Furthermore, most neuroanatomy software systems display only either external (between brain areas) or internal (within a brain area) connection data. In order to fully understand the function of cortical circuits, we need an interactive visualization tool that displays both internal and external connectivity and allows the user to reduce the complexity by filtering out irrelevant information.

Understanding neural circuits and their function has potentially far-reaching implications in science, medicine, and engineering. The proposed visual analysis tools would have an immediate impact on neuroscience research in interpreting circuits connections that might underlie different aspects of brain function. They also provide a convenient framework for comparing neural connectivity in normal and diseased brains, thus could be used to diagnose neurological diseases and evaluate treatments. In addition, understanding the computational principles that underlie higher brain functions, such as attention, decision formation, learning, and memory, would aid design of artificial intelligence systems. The annual neuroscience meeting draws tens of thousands of researchers studying the brain. A substantial portion of them study networks within the brain directly or would benefit from better understanding collectivity in the context of their research problems.

We illustrate the needs with examples from two of the four labs involved with this research. These examples represent some of the different scales of study that are relevant, and they illustrate some of the different types of analyses and features that are likely to be important.

b.1 Schnitzer Lab Brain Circuit Analysis A research direction in the Schnitzer laboratory at Stanford is the study of neural circuits that underlie sensorimotor learning. A brain area that we are interested in is the prefrontal cortex, which is a polymodal area that receives sensory and reward signals from multiple sources and sends outputs to many cortical and subcortical motor structures to guide behavior. It is suggested that sensorimotor learning occurred as a result of changes in sensory-motor mapping that occurs within the local circuits in the prefrontal cortex.

To form hypothesis about the precise location at which learning occurs, it is important to consider simultaneously the external and internal connections. The visual analysis tools proposed would be particularly useful for this reasoning process because, as opposed to traditional brain visualization tools, it will allow simultaneous visualization of connectivity at multiple levels. The user can, as a result, visualize interactions between external and internal circuits in a common visual framework. The reasoning process can further be aided by selective filtering of irrelevant information (e.g. inputs from auditory cortex in a visual task), and post-hoc reevaluation of the visual reasoning process. This workflow was captured in [BS09] and is shown in Figure 1.

b.2 Badre Lab Brain Circuit Analysis David Badre's lab studies cognitive control, which refers to our ability to plan and guide our behavior based on internally maintained goals. Human cognitive control function is classically associated with the prefrontal cortex (PFC), as damage to this region impairs goal-directed behavior. However, cognitive control function is modulatory rather than transmissive in that the route from sensation to action does not pass obligatorily through PFC. Rather, the distributed representation of goal information in PFC neurons modulates the mappings between inputs and outputs represented elsewhere in the brain. As a consequence, understanding cognitive control function requires studying how PFC operates dynamically within systems-level networks. At least two projects in the lab offer examples of this brain network-level approach to understanding cognitive control function.

(1) Memory requires cognitive control. Consider, for example, trying to remember a specific person's name or an event from your past. In these cases, it is necessary to modulate ongoing retrieval process, supported by regions of medial temporal lobes, in order to increase the probability of retrieving only the name or event you want and not other information. PFC is necessary for such directed retrieval, and it is widely held that PFC supports control of memory via modulation of the MTL. However, the specific

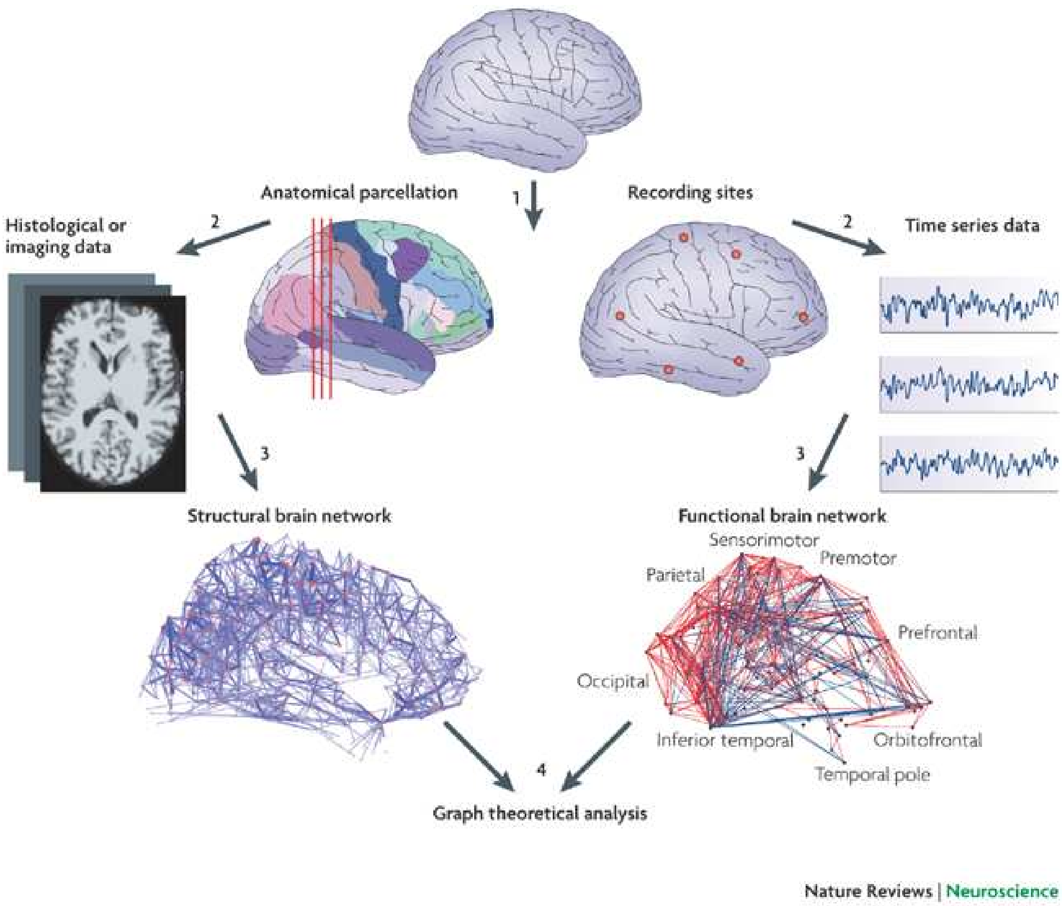

Figure 1: An example of the multi-level workflow of brain network analysis (from [BS09]).

pathways and dynamics by which this modulation occurs remain unknown. Tracer studies in the non-human primate have motivated two candidate polysynaptic pathways from PFC to MTL, one dorsal and one ventral. We are currently using a combined fMRI, effective connectivity, and DTI tractography approach to (a) locate anatomical evidence of dorsal and ventral PFC-MTL pathways in the human brain, (b) assess their differential functional connectivity and dynamics during a strategic retrieval task, and (c) study how certain neurotransmitter systems, like dopamine, may be critical signalers for this circuit.

(2) Ongoing behavior can be expressed at multiple levels of abstraction, from a general goal to a concrete sequence of motor responses. Cognitive control can be required at any of these levels, and growing evidence suggests that rostro-to-caudal frontal cortex may support cognitive control at progressively concrete levels en route to a motor response. Recent data and modeling work in our lab suggests that this hierarchical control architecture may arise from dynamic interactions among a series of nested loops between the basal ganglia and lateral frontal cortex. However, anatomical evidence from non-human primates indicates that frontal-basal ganglia connectivity may differ rostro-caudally, being notably absent at the more rostral extent of PFC. We are currently using a combination of fMRI localization methods, functional connectivity, and DTI tractography in order to more fully characterize the dynamic interactions between basal ganglia and PFC during hierarchical cognitive control tasks.

The proposed tool will be valuable in both of these cases because it will help us mentally merge our circuit diagrams with the intrinsically 3D region and connectivity data that fMRI and diffusion MRI tractography provide. This is a particularly challenging cognitive task, and support for removing as many distractions and other obstacles from the analysis process will help us move forward more quickly and efficiently.

b.3 Some Requirements for Tools These examples, combined with our understanding of visualization tools, of scientific support tools, and of cognition, suggest some specific requirements for the tools we are proposing. First, they need to be able to support reasoning about brain regions and their connections. This support needs to include access to brain region and connectivity knowledge that has already been published. These published data are typically available via several curated databases. Examples of brain connectivity databases are the Brain Architecture Management System (BAMS) [BS07; BDS05], Collations of Connectivity Data on the Macaque Brain (CoCoMac) [K"ot04; SKB$^+$01] and the Functional Anatomy of the Cerebro-Cerebellar System (FACCS) [BLL$^+$05]. Second, this reasoning requires diagramming the regions and connectivity. Third, the diagrams need to handle multiple levels of scale and abstraction. Fourth, users need to be able to interact with imaging data and time course experimental data to support the reasoning. Our examples of imaging data include functional MRI, diffusion MRI, and optical microscopy. Reasoning with these data will involve being able to interact with them in 3-D and understand networks and brain regions both in their anatomical 3-D space and in abstract representations. Similar, but simpler, support will be needed for time course data. Fifth, they need to be able to track and refine their scientific analyses over weeks to years.

c Related Systems and Their Limitations

Relevant existing research and systems span several research areas, research systems, and software products. In the following subsections we break these up into systems for assisted analysis, systems for circuitry exploration, and cognitive-modeling approaches to human-computer interaction,

c.1 Systems for Assisted Analysis Traditionally, visualization enables humans to understand data by representing it visually. More recently, interest in other fundamental visualization issues has emerged. Among others, *Illuminating the Path* [TCoE$^+$05] was written in the context of growing intelligence needs after 2001 and defines the new field of visual analytics as the science of analytical reasoning facilitated by interactive visual interfaces. In addition to defining the term, the booklet establishes the value of such research in the context not only of homeland security but also of any field that involves analysis of complicated and large data.

Recent work in visual analytics can be broadly separated into two categories: theoretical research based on existing cognitive studies and applied work. In the theoretical domain, [Bod] and [PC05] present a five-stage sensemaking model derived through Cognitive Task Analysis (CTA) and verbal protocol experiments with analysts to identify leverage points for visualization. Authors in [RC08] and [ITC08] analyze how users synthesize multiple collections of evidence in a collaborative setting, using a physical, visual medium. Their results, a break-down of analysis tasks with observed frequency/duration and valuable insight into the workflows of collaborative sense making, are useful for deciding which analysis tasks to support. Most of these efforts are not specifically targeted at visualization aided scientific reasoning or, more specifically, visual network analysis but are rather concerned with general tasks pertaining to intelligence analysis. We hypothesize that there are particularities to supporting analysis using visualization systems that are tightly connected to how visualization systems are actually used – a topic current under explored. As such, this proposal plans to complement existent research with low-level task information specific to the proposed domain. A second limitation of current theoretical work is that it was generally restricted to laboratory settings and artificial tasks. The collaborative nature of this proposal will enable us to gather and interpret data from real-life scientific analysis which can both link theoretical research to real applications, and capture and quantify analysis dimensions that are hard to track in artificial settings (e.g., psychological factors such as drive or boredom or long-term temporal effects on analysis).

At the opposite end of the spectrum from theoretical to applied, new applications probe the feature and design space of analysis-support software. Several applications for thought mapping and evidence management use the paradigm of laying out reasoning artifacts on a canvas, either freely or as a tree/graph structure. Such systems include: The Concept Map [CCH$^+$05], MindManager [Min], The Analyst's Notebook [Not07], Visual Links [Vis], The Scalable Reasoning System (SRS) [PMB$^+$07] and The nSpace Sandbox Component [WSP$^+$06]. While useful, the canvas approach is limited: visual vocabulary and interactions are sparse, an automatic layout reduces flexibility while manual layout is cumbersome, usability decreases as the volume of data increases, and, more importantly, problem solving or analysis techniques are minimally supported. Several systems depart from the canvas paradigm. Entity Workspace [BIC06] operates only on textual evidence and uses grouping and linking as an organizing paradigm in a highly structured medium. In HARVEST [GZA06] users can not only visualize existing information, but also construct new analytical knowledge from existing information and apply visualization to it. In [YRW07] and [YXRW] the authors apply similar principles to multi-dimensional visualizations and relate to this proposal by using specific visualization characteristics to drive the organization of evidence. Finally, [EKHW08] departs from conventional methods by structuring analysis as short stories hyperlinked to evidence, a paradigm based on a narrative theory [Fis99] suggesting that people are storytellers and excel at evaluating a story for consistency, detail and structure.

Many of the systems mentioned before are developed and evaluated in the context of intelligence analysis, and as such deal with textual information, large volumes of documents, and temporal or geographic data. Also, few of the systems explore the full range of cognitive principles, ground their design in theoretical and empirical evidence, or aim for thorough evaluation. Most systems don't go beyond the concepts of hypothesis and confirming or disconfirming evidence to structure analysis; they don't aim to convincingly demonstrate that the employed techniques facilitate better analysis either in terms of results or an improved analysis behavior in accordance to normative guidelines. The proposed work has the opportunity to bridge the gap between theory and practice by joining cognitive scientists, visualization experts, and domain specialists ensuring that relevant principles from cognitive psychology are used to design visualization mediums that can be evaluated in concrete scientific settings, with measurable results.

Finally, several approaches have investigated the possibility of recording users' interactions and workflows to facilitate undo operations, next-step recommendations or work-flow reproducibility. VisTrails [CFS$^+$06] records how data sources, filters, visualization methods and visual operations are linked together to produce a useful end-visualization. In [JKMG07] the authors introduce a model of visualization exploration process and a framework to encapsulate, share, and analyze visual explorations. Herr, et al., [HMSA08] describe a taxonomy of architectural and interface issues, identifying design decisions and associated trade-offs. None of these approaches aim to link a visualization workflow to a decisional workflow, which is what we propose in this process.

c.2 Cognitive Modeling in Human Computer Interaction A goal of this project is to improve user performance on brain circuit analysis and other analysis tasks by identifying design principles for computer interfaces that are well aligned with user workflow, including cognitive reasoning and decision-making. If we can predict how users will engage with our software, we can effectively refine its design. A major consideration when undertaking this predictive modeling task is that human actions must be analyzed across many orders of magnitude of time.

In [And02], Anderson argues that we can build successively longer "bridges" across these time scales in understanding, for example, how low-level actions of a student (e.g. eye-tracking across a sheet of paper or computer display) cascade into long-term educational influences. We will support this kind of cascading for analysts using our tool, and will draw on previous research in cognitive modeling architectures in addressing interaction analysis at these multiple time scales. Much of this past research has attempted to create models that predict the performance time of an average user completing a unit task using a proposed user interface. In Project Ernestine [GJA93], for example, researchers showed that a CPM-GOMS cognitive analysis using explicit hierarchical knowledge of user goals and actions can predict user performance on tasks with high accuracy. These predictions were used in evaluating the design of workstation upgrades for telephone operators. Surprisingly, this remains one of the canonical examples of such predictive models. At a lower level, Gluck developed a model (in ACT-R/PM) that predicted student performance on algebra problems based on the distribution of eye movements observed during eye-tracking [Glu99]. These findings support the notion that learning and complex reasoning may be decomposed into small scale, primitive actions, and that we should account for these in our own analysis of user interactions.

In this project, we aim to support hypothesis testing and generation for brain circuit research using cognitive modeling principles. By collecting user data in the form of captured video during usage, logged interaction histories, and eye-tracking, we aim to gain insight about – and build tools to support – the high-level effectiveness of our software. We expect that the processes underlying scientific reasoning and hypothesis formation are significantly more difficult to model than basic interface interactions, like mousing and button presses, to which Fitts's Law and other principles have been applied in predicting user performance. At the same time, if we are successful in modeling user reasoning and problem solving for scientific tasks, the results will be widely beneficial across the fields of scientific computing and visualization, visual analytics, and human computer interaction. This would be a core contribution by creating technology that supports scientific insight and knowledge discovery in a generalized way.

c.3 Circuitry and Network Exploration In computer science, networks of many sorts are abstracted to the concept of graphs. Many techniques [Ead84; DH96; FR91; FLM95; Tun94; LK05; BGHM07] and systems [BST03; Aub03; SMO+03] for displaying general graphs have been developed over the years. The value of this work has been demonstrated by its recent application to a wide range of domains. For example, researchers are able to understand dynamics of communities by visually exploring social networks [Fre00] and biologists can understand protein-interaction and gene regulation networks [SMO+03; HMW+05; ing; DBD+02; JYN+09].

However, research has shown that often general network visualization techniques need to be adjusted to match the particularities of the application domain. For instance, using generic graph-drawing techniques in the context of protein interaction networks yields visualizations that are not intuitive to proteomic researchers: their failure to incorporate protein cellular location and signaling pathway drawing conventions detracts from the visualizations familiarity [Bar; JM03; JYN+09]. In social network visualization exploiting the particular structure of social networks can lead to significantly better visualization [BM02; Fre00] and placing actors randomly on the display can cause misguided interpretations of an actor's role [BMK96].

Moreover, network visualizations have to be integrated into research contexts particular to each domain. For instance, the usefulness of a protein-interaction network is greatly enhanced if the user is allowed to map onto it experimental data [Bar; JYN+09], if the information is hyperlinked to external sources of meta-data, and if certain visual queries, such as highlighting connections pertaining to one protein, are supported.

Network visualization in the domain of brain connectivity is still in its infancy. A few of the online databases hosting connectivity information have developed their own visualization modules [K"ot04]. Others have developed personalized visualizations for specialized brain regions or organisms [wor]. Finally, in [BS09] the authors discuss the opportunity of approaching brain connections from a graph theoretic perspective. These approaches have several drawbacks: they fail to adequately merge findings from network visualization with neuroscientists' intuition; they are limited in scope to singular organisms, brain regions or connection databases; they don't allow users to integrate their own experimental data into the analysis; and they don't offer analysis features such as load/save capabilities, or hypothesis-formation support.

d A Decade of Developing Visual Analysis Tools for Science

Over the last decade, we have developed scientific visualization tools that have taught us a number of lessons about science and about developing tools to support it. Our experience provides us a solid foundation for attacking the proposed problem. In this section we describe a few examples of software we have created and lessons we have learned that we believe will be instrumental to our proposed work.

Brooks describes the principles of such work much better than we can in his paper "The Computer Scientist as Toolsmith II"[Bro96]. A central message is that, "hitching our research to someone else's driving problems, and solving those problems on the owners' terms, leads us to richer computer science research." While there are costs to this sort of approach, there are also benefits; our work has followed this philosophy with good results.

Our experiences with network visualization will help in developing the multi-scale circuit diagramming components of the proposed work. We have studied the workflow associated with protein-interaction networks analysis in the context of high-throughput protein activation experimental data. Proteomic researchers conduct experiments revealing the degree of activation of proteins and wish to relate this data to the body of documented protein interactions from the literature [JYN+09]. Linked views of the signaling network and experimental data, and quick access to the literature and to database information about the proteins and

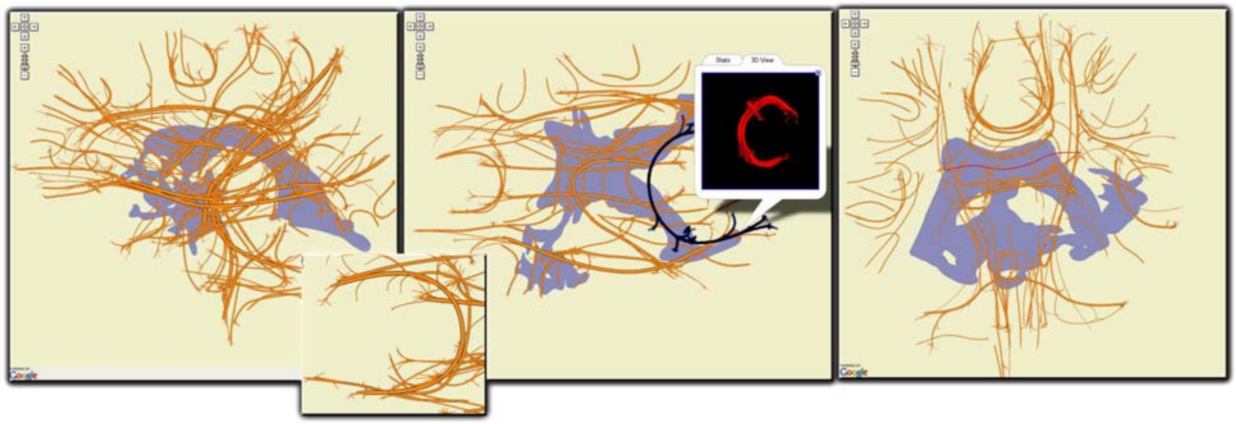

Figure 2: DTI tractography data projected onto the transverse, coronal and sagittal planes. Major tract bundles are represented schematically by their centroid tract; individual tracts in bundles are linked from the centroid bundle to their projected end-points. Bundles can be selected and pre-computed statistical data along with 3D poses of the tract bundle can be displayed.

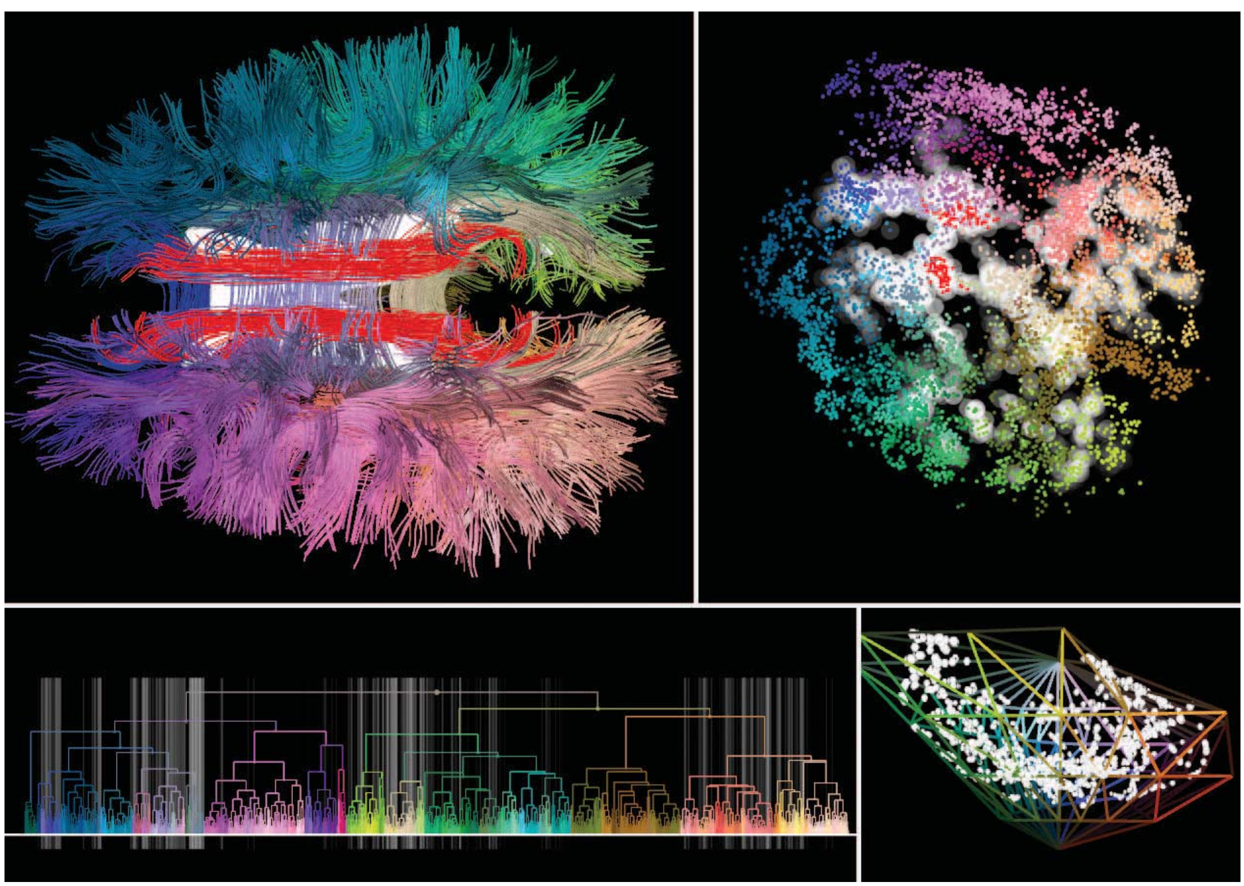

Figure 3: Coordinated DTI tractogram model exploration in lower dimensional visualizations: 2D embedding (upper-right), hierarchical clustering (lower-left), and L*a*b* color embedding (lower-right). A selection of a tract bundle (red) in the hierarchical clustering is mirrored in the other views.

signaling should apply to brain circuit analysis tools. We next describe some specific principles we expect to apply.

One principle is that the network or circuit should be presented in a frame that is relevant to the user. In [JYN+09] we found that placing proteins on a canvas in ways that contradict proteomic researchers' intuition and actual cellular location is counterproductive. Moreover, we found that integrating those networks within a frame that is familiar to the researcher helps the user relate to the network. Fig. 5 shows an example of a protein interaction network scaffolded onto a signaling pathway stylized drawing.

A second principle is that investigating a network at different levels of detail or granularity is beneficial. Fig. 4 illustrates some of our more recent work displaying a protein interaction network as a google-map. In this example, zooming out of the image has the effect of restricting the proteins and interactions that are displayed to the most significant ones, while zooming in allows one to drill into neighborhoods of more obscure proteins.

Third, we found that showing multiple views corresponding to different visual abstractions of the underlying data can help a user understand all of the data more effectively. In [JDL09b] we demonstrate that linking a 3D tractography model of diffusion MRI imagery to abstract low dimensional representations yields an improvement in interaction and data understanding, a technique illustrated in Fig. 3. In Fig. 2 we illustrate additional recent work in which brain tractography is abstracted to schematic projections onto the three principle projection planes. Preliminary results show that this mode of representation, when coupled with the original 3D model, is appreciated by neuroscientists. It also represents a first step in abstracting anatomical features to connectivity information and projecting it onto a meaningful space.

From our preliminary interactions as a group, we hypothesize that these findings will translate to effective brain circuitry analysis tools. The anatomical framework and the inclusion of familiar anatomical landmarks are likely to play an important role in how a user perceives a brain network. Users will want to investigate the network at different levels of granularity depending on their research interests. Coupling abstract network representations to real anatomical data or visualizations of supporting meta-data is likely to improve analysis.

Our experience with developing more imaging-oriented applications will help with the components of the proposed work that deal with microscopy and MRI data. Some representative examples include 3D flow visualization [RPKL06; PRLK05; PHW+05; FRS+03; CRH+10], diffusion MRI visualization and analysis [Zha06; JDL09b; ZDK+01], carpal kinematics analysis [MGL07; MCL09; MDAL04], and even archeology visual analysis [AVLJ00].

Finally, Brown's proximity to RISD provides a wealth of visual design expertise from not only the students but also the faculty, one of whom is an investigator on this proposal. Laidlaw and Drury have taught a "Virtual-Reality Design for Science" class several times, bringing together computer science students, art students, and scientists to explore the process of designing visual and interactive tools to accelerate science. We have documented some of what we have learned in the literature [KKVL05; VAJ+03; AJLD08; KAM+08].

Laidlaw's lab has been one of the leading visualization labs in formal evaluation of visualization methods [LKJ+05; FCL09; PFK+07; DJK+06]. Some of these evaluations were done in collaboration with perceptual psychologists from Sloman and Badre's department. They incorporated some of the kinds of experimental design issues that will be needed to execute our proposed work. Interestingly, these formal evaluations and the "critiques" of of our Brown/RISD class are related – in a sense, they are attempting to solve the same problem of understanding how humans and software interact effectively. This synergy is fascinating and enlightening, but unfortunately rare. We will leverage it in the proposed work.

e Cognitive Optimization: a Novel Approach to Human-Computer Interaction

Development of the software tool will involve an iterative process. An initial prototype will be developed based on software design and cognitive principles. That prototype will be tested first informally and then formally. The results of the testing will shape the next generation of the tool which will also undergo testing and refinement. More details of this process can be found in Sec. f.

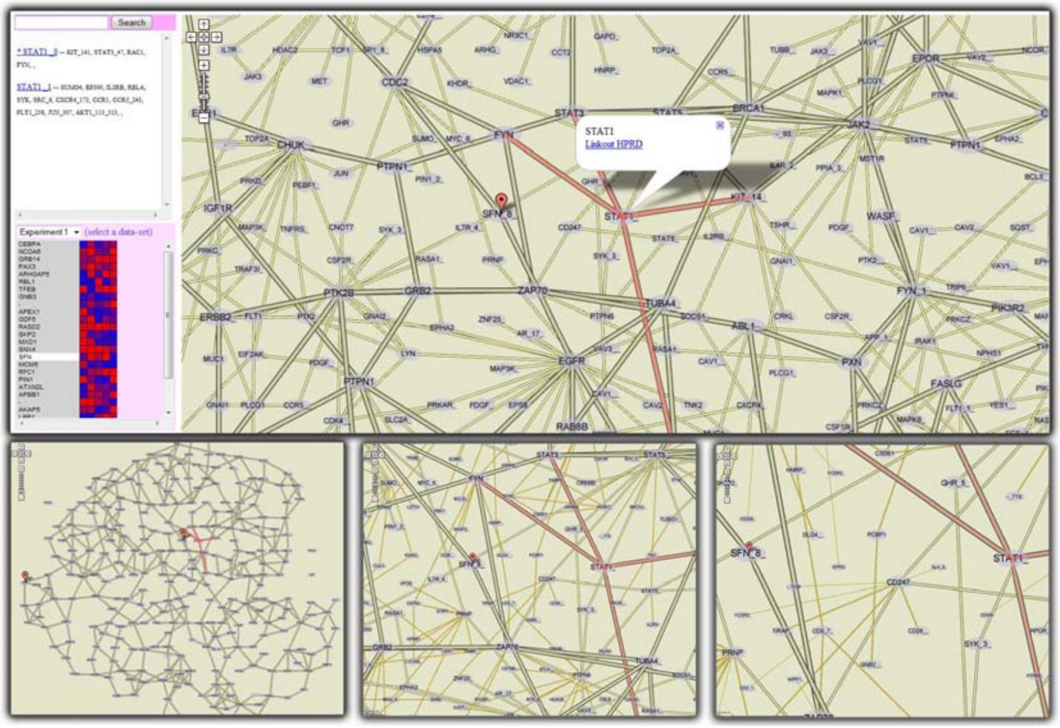

Figure 4: A google map of a protein interaction network. Outer zoom levels reveal only well known proteins and interactions while zooming in brings forth more obscure proteins.

e.1 Prototype Development The cognitive principles that will be used to develop the prototype come in three varieties. First, principles of perception and attention will determine the physical parameters of the display. Second, principles of goal selection will determine the number of tasks made available to the user and their accessibility at any point in time. Third, principles of judgment and reasoning will determine the design of the tool, the information that is on display at any given time, the actions that it makes easy, and the guidance and feedback that it provides.

e.2 The Principles of Perception The software will make available facilities for viewing data, like the network itself, and for performing other tasks (e.g., maintaining history, calculation facilities, note-taking facilities). The designers will have to make decisions about how to display the data (e.g., number of network graphs, network parameters, number of nodes per network) and the number of other functions to include in the display. Note that these are all dynamic entities. They change rapidly over time and they interact constantly.

Recent work in the study of attention has examined the number of objects that can be tracked at any given time [PS88] and the physical parameters that determine the identity of an object [SP99]. These studies present participants with a large number of objects moving in either random or systematically varying directions on screen. Participants' task is to track one or more of the objects; i.e., to maintain the identity of some number of objects as they move about. In some experiments, the objects disappear or become occluded or change in color or size. These experiments reveal the key physical parameters describing how people deploy and distribute attention to maintaining awareness of the existence and behavior of objects as they change over time. From this work, and from older work on the capacity of attention [Pas99], we can draw inferences about how networks should be drawn, the number of nodes and links that should be displayed, and how much changing information about the network can be maintained simultaneously.

e.3 Goal Maintenance Scientific problem solving entails multitasking. The scientist must decide what to look at, recognize patterns in the data, decide which hypotheses to test and actually test them, all

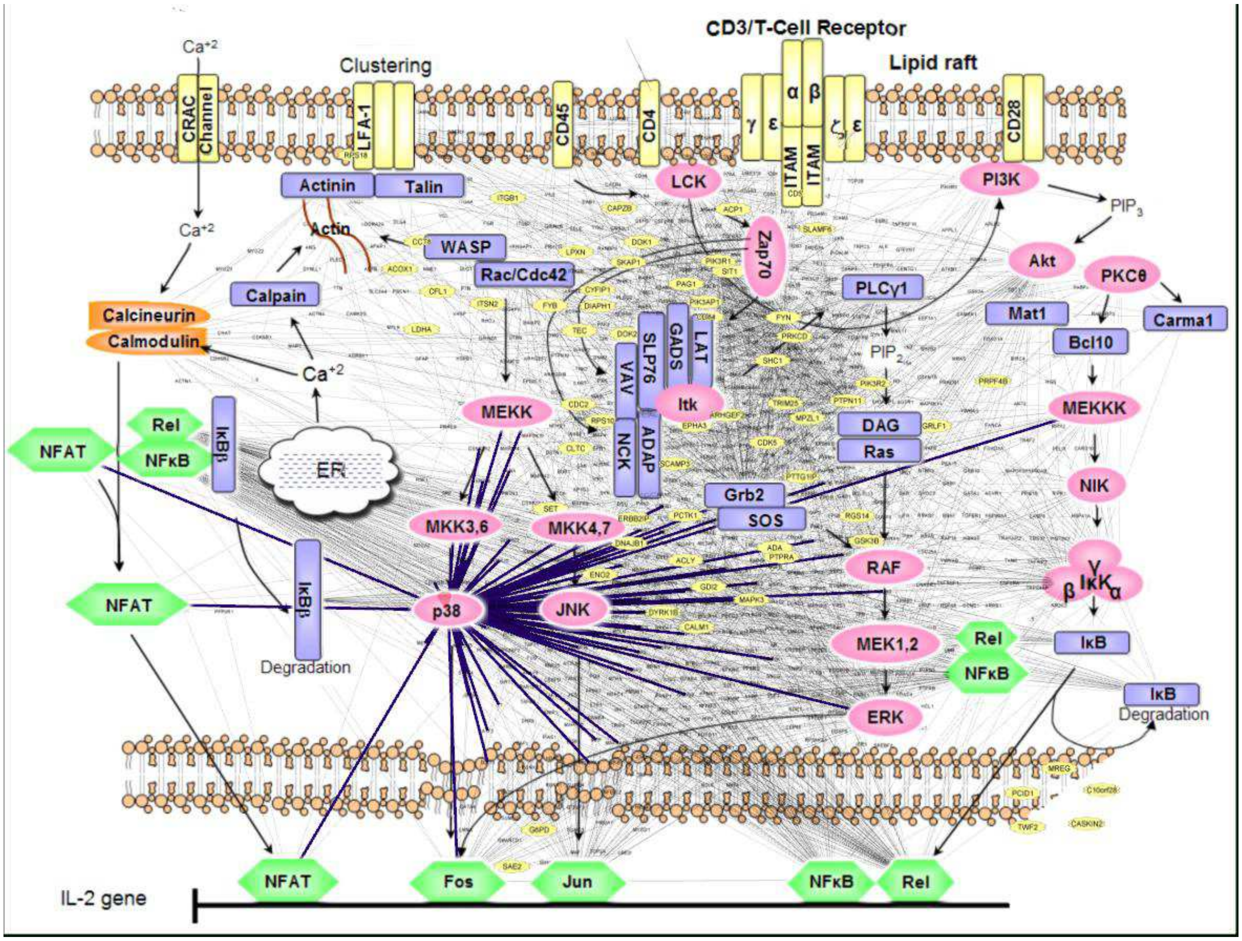

Figure 5: Proteins and interactions extracted from the Human Protein Reference Database are integrated into a stylized signaling pathway model.

the while keeping track of what has been investigated, what has not, any new ideas that come to mind while exploring, while also being mindful of their other everyday concerns. Research has studied the process of multitasking using a variety of paradigms including prospective memory [ME00], task switching [RM95] and goal priming [AD00; BC99].

Several principles have emerged from this work identifying the conditions under which different goals facilitate versus compete with one another. Goals facilitate when they share a larger purpose. For instance, sub-goals of a single task tend to facilitate one another. Goals also facilitate one another when the task performed to pursue them are similar. Otherwise, goals compete. Relations among goals also determine whether they should be pursued simultaneously or sequentially. We will use these principles to determine which tasks to make easily accessible to users and which to put in the background. Moreover, people differ in their ability to juggle multiple pursuits simultaneously. Our system will tune itself to individuals' work styles.

e.4 Hypothesis Testing and Problem Solving At the more abstract level, the software will be designed around three basic principles concerning how users think and solve problems.

First, people have a tendency to be overly myopic in the hypotheses they choose to test [KH87]. This manifests in several ways: we tend to consider only a limited number of hypotheses, failing to consider alternatives; we focus on data that is consistent with the hypothesis we have in mind, neglecting other data; we interpret data in a way that supports our prior hypothesis, failing to appreciate when data are disconfirming.

The most simple and effective way that psychologists have found to reduce this bias is to ask people to "consider the opposite" [LLP84]. That is, simply cuing reasoners to consider an opposing hypothesis when problem solving or interpreting data is often enough to get them to do so. In other words, the psychological roadblock is not always an inability to consider alternative hypotheses, but a meta-cognitive failure to consider doing so. A second remediation method is to ask people to think diagnostically rather than causally [FDS10]. Instead of thinking about a problem in terms of how causes lead to effects (e.g., how does my model of cortical processing lead to this network diagram?), more alternatives are considered when the converse problem is considered (e.g., what is the probability that my model of cortical processing is correct given this network diagram?). Both of these principles can easily be implemented to nudge scientists along.

The second cognitive principle is that people think in terms of mechanisms that operate over time [Slo09]. That is, they break up tasks into functional units. But these functional units do not merely produce specific outputs for given inputs, they do so with known or expected time courses. In that sense, they are causal and not merely computational. Also, these mechanisms come at different scales: some describe operations at microscopic levels and others describe more abstract functional units. In each case, the mechanisms come in structured bundles that can be described as networks or circuits. Mechanisms relate variables of 5 types: causes, disablers, enablers, preventers, and effects. The general form of a mechanism follows that of a structured equation: The effect is a joint function of the other 4 variables. This is not merely a mathematical relation because it obeys a temporal constraint: the effect cannot occur prior to the causes. Moreover, people have a rough idea what the delay should be in the appearance of the effect when the activities of the other variables are known.

This way of thinking is particular helpful as a way for scientists to structure their knowledge of the brain. The brain is studied and described at multiple levels, from the cellular to the systems level. Each level consists of networks of functional units and each functional unit has an expected time course. When a scientist looks at imaging data, he or she does not merely see a visual pattern but rather a dynamic entity that has a purpose, where each pattern of activity exists in the service of some effect. In the canonical case, the activity represents the response to a stimulus that follows a causal chain to arrive at an ultimate effect, a motor response.

The third cognitive principle is that people think in two ways, intuitively and deliberatively [Eva03; Slo96; SW01]. Intuition involves recognizing patterns and coming to conclusions that relies on memory and associative principles like similarity and contiguity. People are not conscious of the process of intuitive reasoning, only of the result. On the other hand, people think deliberatively. They transform symbols over time using executive control mechanisms housed in working memory. People are conscious of both the process and the conclusion of such reasoning.

When solving problems, the two systems of thought have two forms of interaction. For simple problems that can be solved in habitual ways, the normal operation of the system is for the intuitive response to dominate. But if a competing deliberative response comes to mind, then that is what determines performance. A scientist who sees a familiar network configuration will intuitively identify it. But if the context of the observation rules that possibility out, then normally the classification will be abandoned. The second form of interaction concerns problems that require focused thinking over time, like puzzle solving. Such problems are generally solved by a cyclical process in which deliberation offers representations to the intuitive system [NS72]. The intuitive system then tries to find a meaningful pattern in the representation. If it succeeds, then the response is validated by the deliberative system. If it fails, the deliberative system tries to construct a more successful representation.

The distinction between the two forms of thinking will guide the design of the software tool by implementing two sorts of modules. One type will be specialized for intuitive reasoning. It will be much less constrained in the amount of information that it displays and will be designed to afford pattern recognition. The other type will be specialized for deliberation. It will show only a small amount of information at a given time and will facilitate small, analytical reasoning steps.

e.5 Formal Testing After informal testing of the prototype, and at each stage of the tool's development, it will undergo formal user testing. This will involve picking a set of simple, solvable problems concerning brain circuitry and asking a set of novice (probably undergraduate) brain researchers to solve them. Some of the researchers will use the software tool and others will use the software tools that are normally used by researchers in the relevant field. We will make sure that levels of training with the two sets of tools is approximately equivalent. We will choose problems that are challenging but not impossible. Some will be largely analytical in the sense that standard methods can be applied while others will require thinking "outside the box" (insight problems).

We will compare problem solving performance on both types of problems using the different tool sets using a variety of measures: likelihood of solving the problem, reaction time to solve the problem, and quality of solution. We will also obtain qualitative information about performance by querying participants about what they are thinking at regular intervals and tracking their eyes. This will reveal if our software tool is changing basic problem-solving strategy.

f Project Plan and Schedule

We plan to execute our multidisciplinary research agenda over the course of 5 years. The primary output of this project will be a visualization system for the analysis of brain circuitry, with a target audience of brain scientists seeking to understand neural pathways and connectivity. This software will be developed by the Brown group with continual and closely integrated collaboration with the brain research labs. The goal of this tool is to allow researchers to view brain circuits at multiple scales and to support sophisticated analysis of research hypotheses.

In order to maintain a community-driven approach, software development will follow a spiral engineering pattern, with system components and prototypes of increasing functionality deployed for each iteration of the software. These deployments will be to the brain research labs for evaluation and feedback. In addition, interviews and studies of users will be directly used to shape the cognitive models of user reasoning that the Brown group will use to design, modify, and evaluate the system interface. At all points in the development process, collaboration among the research groups will be used to ground design in what is important to users and to incorporate an understanding of the scientific reasoning process into a visual analysis tool.

Timeline

Our plan is divided into five phases. We anticipate that they will correspond to the five years of the proposed work, however there will clearly be some overlap among the phases and some variability in their size. Each phase/year will be driven by multiple evaluatable milestones, e.g., requirements documents, design sketches, design reviews, releases at various scopes, and experiments.

Phase/Year 1: Knowledge sharing, modeling, setup, and interface sketching. In this phase the labs will work together to gather requirements for the software and to develop an initial set of rule sets and user models for the cognitive modeling aspect of the project. This phase will include early testing of the cognitive models to determine whether they accurately predict user behavior and to what degree they are relevant to the analysis needs of the users. During this phase we will iterate on the requirements via sketched interfaces and visualizations; user feedback and cognitive analysis will guide the refinement of the sketches[KAM$^+$08]. A system for capturing user event history, video, and tracking will be developed and combined with a quick prototyping application development system used for several visualization applications [JDL09a]. We will also set up the SourceForge repository, including its bug tracker and forums for developers and users.

Phase/Year 2: Proof of concept design and prototype. An initial software design will be developed based on the requirements gathered in Phase/Year 1. This design will incorporate the highest priority requirements and will be prototyped to demonstrate to the brain scientists that their data is being accurately and understandably displayed, that the interface is usable, and that the system has the potential to show them information or support analysis that was not possible before. The prototype will be made available on SourceForge but without significant publicity. It will be used in our own laboratories to test the interface and the testable cognitive components.

Phase/Year 3: Primary system development cycle. During this main phase of deployment, testing, and design, we will refine our system design through small-scale evaluations within the Laidlaw lab and across the other three labs. This phase will include an iterative series of increasingly complete prototype systems, each of which will be deployed to the brain science labs and evaluated with respect to basic usability, task performance, and how well they support reasoning. Interface evaluation metrics will include completion time and accuracy with simple information-extraction tasks, insight generation, and qualitative feedback from interviews or focus groups. In addition to evaluating the interface, we will also evaluate the cognitive models used to refine and guide interface design. To do so, we will be evaluating the models' predictions of user behaviors and states against actual user data, such as computer interaction logging, video logging, and physiological measures such as eye-tracking data (for attention) and skin conductance response (for affective states). The results of these evaluations will be used to improve our cognitive model, which will subsequently be used to improve system design and responsiveness to the user. Throughout this phase, significant evaluation results and novel aspects of the system will be published in appropriate conferences and journals.

Phase/Year 4: Initial Public and Publicized Release. Iterations on the implementation will be continued into Phase/Year 4 leading up to a public release approximately halfway through the year. This version will include all the intended functionality of the system, including the ability to incorporate and display multiple types of brain connectivity data (fMRI, diffusion fMRI, optical microscopy, and other experimental data); the ability to compare and understand circuits at multiple levels of detail, from cell types to behavior; and finally, should be cognitively optimized to support reasoning and analysis at a high level, rather than just displaying data. During the second half of the year, feedback will be gathered and analyzed more broadly, and community-building efforts utilizing SourceForge mechanisms, for potential developers, and possibly something more user oriented for users who are not developers. User data capture will continue as will internal improvements.

Phase/Year 5: Dissemination. The final system will first be deployed to the brain science labs, with the goal of incorporating its use into ongoing research. Researchers in these labs will observe how the system is used in practice and whether it successfully improves their ability to analyze data and test hypotheses in a real-world setting. At this point we will also release a final version of the system to the public as an open-source distribution, with the source code available via SourceForge and binary installers available through SourceForge and our labs' websites. We will continue to support this release and to log usage data in order to improve the system and better understand how it is used in practice. During this final phase we will also make a choice about the mechanism to be used for sustainability (see next section).

g Outreach, Education, and Sustainability Plan

As described in the timeline, several versions of the system will be deployed to the public as an open-source project via SourceForge, using their public license derived from the GNU General Public License. We have been using this approach for a virtual reality library we are developing with other NSF support (OCI-09-23393). We will also provide binary installers for Windows and Mac so that scientist users can download and use our software. This release will be advertised through blog postings, publications, direct contact, and during conference presentations in order to raise its visibility among the larger visualization, human-computer interaction, and brain science communities. We will also propose to teach tutorials and classes at conferences if there is interest. Additionally, researchers from the brain science labs will be encouraged to use the tool in classes and other educational settings in order to help students understand connectivity in an accessible visual form.

The system will include instructions and easy mechanisms for reporting bugs and providing other feedback. We will address serious maintenance issues as needed. Usage data will be gathered from users if they opt in upon downloading the tool; observation of the tool's use will be undertaken in the brain science labs involved in the proposal. These observations will be used to inform future iterations of the software.

Dr. Laidlaw will be responsible for overseeing the sustainability of the software. Our hope is to interest a company in pursuing SBIR funding to provide continued maintenance, support, and development past the end of the proposed project. An alternative model would be a mechanism like the one that supports ITK through the National Library of Medicine at NIH. If neither of these is possible, we will work to develop open-source community support through the SourceForge repository and its developer and user forums.

h Potential for Broader Use

The success of our proposed research has the potential to be broadly applicable to scientists in many other domains, as well as to software developers and users outside of science. Advances in creating effective visualization interfaces should be applicable to almost any scientific user. Advances in using cognitive optimization as a design and development principle should apply to almost any interactive software tool. Support for reasoning with sketches and networks should be applicable beyond brain networks – not only to disciplines that study networks, but to any analysis oriented problem area.

Applications already under development in Laidlaw's lab illustrate several other scientific domains where similar tools should be valuable, including the study of protein signaling, gene regulation, and possibly their interaction. Other examples of problem domains that might benefit include the study of social networks, epidemiology, and the analysis of large bodies of text.

Our novel "cognitive optimization" approach to supporting analysis systems has the potential to provide valuable input to almost any interactive tool for visual analysis. Because the code implementing it will be available as part of our distribution, it can be extended or applied by other researchers.

This project will provide broad access to a state-of-the-art visualization tool necessary for researching current problems in brain and cognitive science. Our software build and deployment process will allow users to get the most current stable version of our tool on demand or automatically.

Finally, much of the interaction data we collect will be anonymized and made available to the research community. This will be valuable to the research community; our goal is to share the insights we gain from studying user interaction patterns and the effectiveness of visualization techniques. These data will range from low-level eye-tracking and mouse-click logs to high-level information on decision making and concept modeling. This data repository has the potential to help numerous scientific areas within cognitive science and human computer interaction, including decision making, dual-systems research, perception and vision, and user interface optimization.

i A Cohesive Multi-disciplinary Team

We have assembled an excellent team for accomplishing the proposed research. We have recognized experts in cognitive science, neuroscience, computer science, and visual design. Their complementary expertise covers a breadth unusual in such a small group. The problem area of brain science and the research thrust of cognitive optimizing user interfaces are quite synergistic, which helps to reduce the number of disciplines necessary to attack the problem.

Importantly, the faculty investigators have been intellectually engaged with one another for some time. Badre and Sloman are in the same department and Sloman has attended one of Badre's classes. Laidlaw and Sloman are participants in an ongoing working group studying a dual-system model of cognition. Two of Sloman's Ph.D. students attended Laidlaw's "Cognition, Human-Computer Interaction and Visual Analysis" class last year, and one of Laidlaw's students attended two of Sloman's classes this year. Schnitzer provided the inspiration for attacking this problem by contacting Laidlaw several months ago, and they and their groups have been interacting since on what the neuroscience needs are. Laidlaw and Drury have been teaching a class "Virtual-Reality Design for Science" over the last seven years and have learned much about working collaboratively between the design of software and the design of visual and interactive artifacts.

Brown provides a supportive environment for multidisciplinary work such as we propose. The diverse student body is very creative, and we plan to leverage that by involving undergraduates in this research. Brown has porous disciplinary boundaries which permit easy collaboration. Brown's Brain Sciences Institute is an example of a multidisciplinary organization that leverages this easy collaboration; it also provides supportive infrastructure, including biological imaging, centralized talk announcements, small seed funding for new collaborations, undergraduate research opportunities, and multiple examples of successful collaborations.

While admittedly risky, we believe that the proposed work would provide broadly valuable benefits to all the disciplines directly involved, to many other scientific disciplines, to software development and design, and to knowledge workers.