Scientific Visual Analytics: Advancing Theory and Practice through Cognitive Modeling

Brown University
DH Laidlaw†, BJ Raphael*, *CS*
S Sloman*, D Badre, *Cogn. Sci.*
D Rand, *Eco. & Evo. Bio.*

Howard Hughes Med. Inst. & Stanford U
M Schnitzer*, *Bio. & Appl. Physics*
Harvard Medical School
C Benoist* *Micro- & Immuno- Bio.*

Florida Int'l U.
R Jianu, *CS*
RI School of Design
F Drury, *Illustration*

† PI, * Co-PI

a Research Vision and Plan

We envision a theory of interactive visualization within the field of human-computer interaction (HCI). Such a theory would, perhaps fifty years in the future, let us accomplish most evaluations of an interactive visualization tool through predictive modeling and theoretical analysis. Such a theory is far in the future; we are still at an early stage where a significant part of the journey toward a theory is observational. In this expedition, we propose to move toward that vision. We will begin by observing individual users as they interact with examples of scientific visual analytics tools. As we move forward, however, we will experiment with using current models of human abilities to test our hypothesis that they can be used to predict user performance with our tools.

As Turing-award winner Brooks describes in his visionary "Computer Scientist as Toolsmith II" paper [12], driving applications can lead us to richer computer science research. We strongly subscribe to this concept, and so the work will be guided by scientific analysis applications in brain science and genomics. Understanding the human brain is a daunting enterprise. Brain researchers must interpret their voluminous data in the context of everyone else's voluminous data. Genomics researchers face an analogous challenge. These pressing needs, including specific examples from our labs, are further described in Sec. a.2. The software tools we propose will make these enterprises less daunting and more productive.

These driving applications will benefit the scientists that use them as well as providing examples that should benefit scientists in other disciplines. In recent years, advances in computing and experimental methods have resulted in the rapid growth of large-scale high-dimensional data in almost every area of scientific research. The growth of such complicated data holds enormous potential for scientific progress. At the same time, analyzing all this new information is a significant challenge. Humans are adept at finding patterns, and existing scientific visualization methods leverage that, but as the size and complexity of data increase, visualization alone is not enough. In this proposal, we aim to generate data analysis tools that enable human interaction with high-dimensional data at multiple levels of abstraction, allowing scientists to reason and make discoveries well beyond what is possible today.

In our specific example applications, we propose to apply knowledge about how people perceive and think to the development, evaluation, and optimization of visual analysis interfaces for scientific inquiry. Human perception and cognition are complex and sometimes counterintuitive; many perceptual illusions and cognitive biases are well known and their bases understood. Through that understanding they can sometimes even be turned from bugs into features.

We will include not only scientific models and principles of cognition and perception, but also those developed over centuries by artists and visual designers. Artists' phenomenological understanding of the response of viewers to stimuli captures behavior that some scientific models miss, and we propose to codify this understanding to influence interface designs for our driving applications. At Brown our visualization research has considered extensively the design and evaluation process of such tools. We will utilize visual design knowledge in the software design process as we have in the past [4, 40].

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 higher-level principles. Sloman and Badre, from Brown's Department of Cognitive, Linguistic, and Psychological Sciences, have a deep understanding of these areas of knowledge and will guide this aspect of the proposed work. They have knowledge about how scientists think that is developed and specific enough to provide guidance about what functionality to include in the tools and how to design them. This knowledge concerns behaviors including how scientists test hypotheses, how they mentally represent data, how they make decisions, and what kind of neural "processing systems" they can bring to bear.

The driving applications will be developed using a new software framework that will support our research agenda. This framework will capture observational data of users, codify perceptual principles and models that might explain observations, and integrate multiple models that might explain different aspects of observations. Within this framework, we will evaluate our tools as well as testing the modeling hypotheses we will study.

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

Our approach will diverge from standard software engineering in its primary goal of capturing and disseminating observations, positing generalizable models and theories to explain those observations, and testing those models and theories. This divergence can happen, however, only in the context of real tools and real users, necessitating the software engineering aspects.

To achieve our vision, we propose a number of specific activities with expected deliverables and outcomes. These activities are shown in Fig. 1.

The following sections outline our rationale and motivation for the elements of our proposed research. They also sketch some of the preliminary results that have helped us architect the research plan and provided us with some confidence about some of the riskier aspects. An overview of related work provides context and establishes significance. Finally, we describe the relevance of the work to the Expeditions in Computing program goals and characteristics.

This effort is clearly an ambitious one. But we have gathered a research group with the 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 compensate for that risk. Even if we are not completely successful, we believe that the tools will accelerate brain science and genomics, the knowledge gained will benefit other analytic applications, and we will come closer to a theory of scientific visual analytics.

a.1 Significance for Computing

Major computational challenges exist at the interface between human users and computers. In order for computing tools to augment cognitive analysis by human users – a common goal in visualization – they must provide functionality that is complementary to the human cognitive system. Traditionally, visualization research has focused on one direction of human-computer interaction: transforming data into computer-generated images that allow data features to be easily perceived and analyzed. This Expedition proposes a research agenda that will enable computers to interpret rich information about human users, to create predictions of user states and analysis outcomes from this information, and to optimize visualization design and output using these estimates. Recent developments in cognitive science and the acceleration of scientific data collection make this a timely venture that is well-aligned with the Grand Challenges and R&D agenda posed by Thomas and Cook at the Pacific Northwest National Laboratory [68]. Our research will advance core topics in Chen's list of top ten unsolved information visualization problems ("2. Understanding elementary perceptual-cognitive tasks" and "9. Causality, visual inference, and predictions") [15], as well as number of Johnson's top scientific visualizations problems ("4.

1. Develop an interactive analysis and reasoning software tool for brain scientists to explore and analyze connectivity in the brain via imaging data, published literature, and other experimental data.
2. Develop an analogous software tool for genomics researchers.
3. Support the linkage of analysis to external data by leveraging existing curated knowledge on the web, by providing web access to our own curated brain and genomics data, and by utilizing social and collaborative web interfaces to grow these resources.
4. Support the tools across a range of visualization hardware platforms, from laptops, tablets, and desktops to display walls and a new "retina cave" at Brown.
5. Instrument the tools to create a testbed for new interaction and visualization techniques.
6. Capture data of scientists using the tools, including interaction hardware events, video of body motion and body language, and tracking of what users are looking at.
7. Manually analyze captured data as training and evaluation data for subsequent automated model fitting.
8. Create a software framework for comparing captured data with models and principles from cognition, perception, and art.
9. Evaluate a series of cognition and perception models and principles for their ability to predict user performance measured in Item 6.
10. Using feedback from brain scientists and genomics researchers, define improved visualization, interaction, and analysis techniques.
11. Iteratively optimize these new techniques using predictions from the cognition and perception models and principles.
12. Validate the approach by confirming predictions empirically.
13. Integrate our research efforts with several Brown classes.
14. Make publicly available throughout the process our captured user-interaction data, our modeling testbed, the brain-connectivity and genomics visual-analysis software tools, and the interaction and visualization techniques developed.

Figure 1: Proposed activities and deliverables

Perceptual issues", "6. Human-computer interaction", and "15. Theory of visualization") [39]. This work will be applied to scientific visualizations; however, other human-computer interfaces that involve reasoning have the potential to be improved by results from this research.

Our aim is to pursue a research agenda for scientific visual analytics, which differs from traditional analytics applications, like disaster management and social-network analysis, in ways that present unique challenges to computer scientists and their collaborators. First, the end users of science tools are primarily scientists who are generating hypotheses about experimental data; the result of refining hypotheses could be experimentation and collection of new data. Second, scientists cannot be expected to become "power users" of analytics software like intelligence professionals. Their tools must be learnable and intuitive enough to preserve the cognitive capabilities needed for hypothesis generation and experimental design and to allow the time between uses for experiments. To address these unique challenges, we propose a framework that couples visualizations with cognitive models. It will be used to analyze captured user data visually and to hypothesize about cognitive models. Subsequently, models will be used to predict user states and analysis outcomes. We believe the framework will enable the discovery and evaluation of design principles and optimizations.

Visualization has been successful at accelerating science and analytics despite having very few theories of why or how it works. Our proposed Expedition will establish an agenda of developing descriptive and predictive models that relate analysis outcomes to information about human users, goals, and computer interfaces. To the best of our knowledge, we will be among the first to systematize the way cognitive theory is integrated with scientific visualization tools. This agenda has both immediate and long-term potential to change the way visualization science is conducted and the way applications are developed and used in real settings.

A Novel Framework for Human-Centered Visualization Our three-part framework will:

1. let visualization developers use instrumentation and monitoring devices to collect and encode user interactions and observations, like galvanic skin response and eye tracking, over time; it will enable visual analysis of user data and goals for scientific visual analytics;
2. feed interaction and user data to a cognitive prediction manager application, which will run software modules that implement predictive models of user performance and user cognitive states;
3. provide reports and a dashboard of model predictions over time that are used to understand and improve the visualization user experience.

The framework is described in more detail in Sec. a.5. Extensibility will be an important aim in our work. New visualizations will have access to an application programming interface (API) for cognitive tools. Standard interfaces for both visualization development and cognitive-model development will be created and advertised on the Web and in academic conferences and publications. Software that conforms to these standards will be interchangeable within the framework. This will give visualization developers opportunities to incorporate current findings in cognitive science into novel interface features, then share their software.

a.2 Need for Visual Analysis Software in Scientific Applications

We describe these needs in terms of our two driving application areas.

Brain Connectivity The first target scientist community for our analysis software is brain scientists studying the architecture and function of the rodent and human brains at multiple scales. Three labs involved in this expedition are a part of this community.

The human brain is one of the most complex organs 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 and cognition to learning and memory. However, our understanding of the neural basis underlying these brain functions are limited because the networks of connections are so complex and because recent advances in molecular and imaging techniques have expanded the data about those networks to an overwhelming level.

We illustrate the needs with an example from the Schnitzer laboratory at Stanford. A research direction there is the study of neural circuits that underlie sensorimotor learning in rodents. They study the prefrontal cortex of the rat, 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 hypotheses 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, they 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 [13] and is shown in Fig. 2. In preliminary work, Schnitzer's group developed a prototype "MindMap" to begin addressing these issues (Fig. 3). In this diagram of brain regions and connections, which must be reproduced impractically large to be legible, the challenges of displaying and understanding the data are clear.

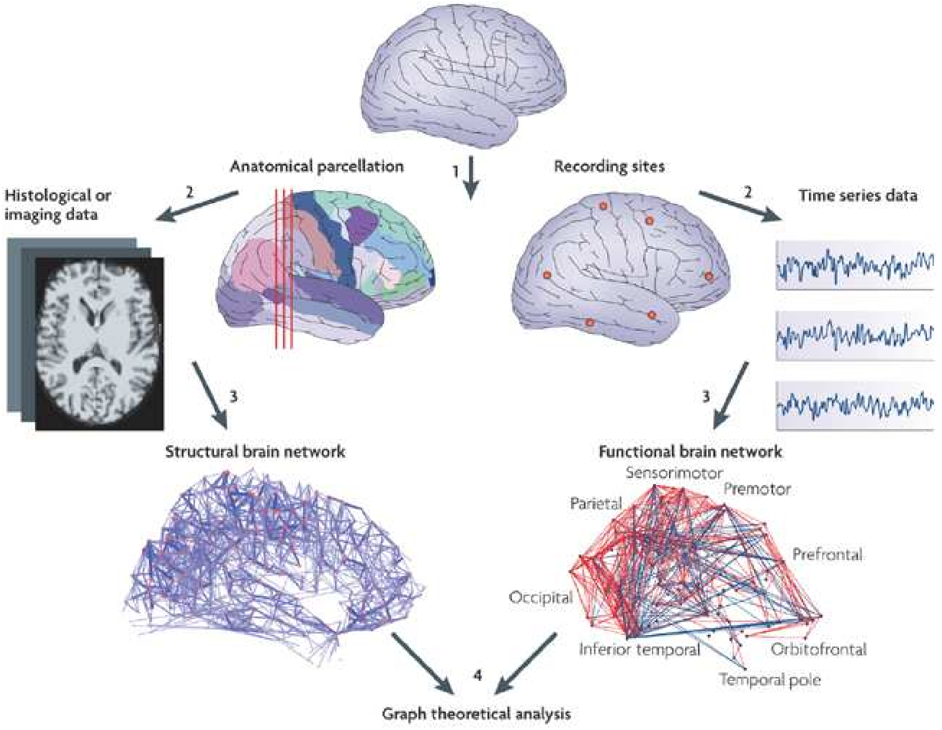

Nature Reviews | Neuroscience

Figure 2: An example of the multi-level workflow of brain network analysis (from [13]).

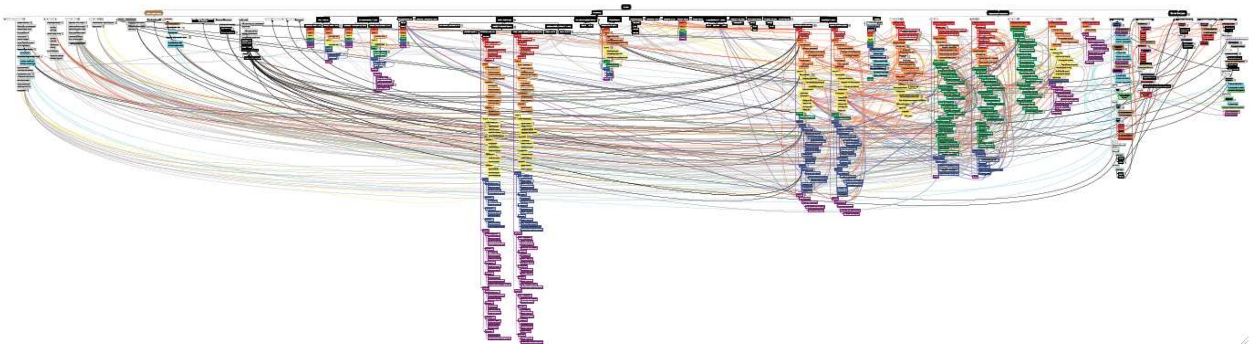

Figure 3: Mindmap displaying connections across brain regions derived from brain science literature. Brain regions are arranged horizontally with hierarchical decompositions below the top level. Arcs show connections. The challenges of displaying and understanding the data are substantial.

Badre's laboratory explores problems analogous to Schnitzer's but studies cognitive control in the human brain and at a larger anatomical scale. Cognitive control is our ability to plan and guide our behavior based on internally maintained goals. Understanding cognitive control function requires studying how the prefrontal cortex (PFC) operates dynamically within systems-level networks. One way to infer how prefrontal cortex communicates and interacts with other parts of the brain is to assess co-activation in MRI signal between different voxels in the brain. Badre's laboratory currently looks at fMRI data obtained from two different sources: one source is publicly available repositories of neuroimaging data sets or online platforms that synthesize a large amount of fMRI data extracted from published articles (e.g., NeuroSynth [74], BrainMap [26]), and the other source is data collected by his lab and other collaborators [69].

One challenge arises when researchers from Badre's laboratory analyze the fMRI data is that they often want to ask questions about indirect connections, that is, to generate hypotheses about if and how the interactions between two brain regions are mediated by some other brain regions. This mode of inquiry is much more difficult than inquiry about direct connectivity given the potentially large amount of connectivity data and the currently limited methods for visualizing the data, especially if the researcher is interested in more than one level of indirection. Another challenge stems from the fact that there is no widely agreed-upon brain ontology. Brain regions can be organized based on either their cognitive functions or the brain's anatomical structure, and there are many variations for each of these two organizing principles. While searching for connectivity information about a particular brain region, it is easy to miss data that are relevant but have used different naming or are on a different level of granularity than the researcher has in mind. This necessitates translating or establishing mappings among different ontologies if a researcher wants to gather as much relevant information as possible when generating hypotheses.

The task of studying PFC within systems-level networks becomes even more challenging when researchers combine data and evidences derived using different techniques or measured on different species to reason about brain connectivity. For example, while researchers in Badre's laboratory mainly work with MRI data that reflect functional connectivities, they also sometimes refer to data obtained from tracer experiments on macaque, as their anatomical brain structures are more extensively studied than those of the human brains. It is very beneficial, and even necessary, to combine heterogeneous brain connectivity data, as no single technique can capture every aspect of the brain. This is very challenging, as the network derived from one data source is determined by the technique, the species, and the ontology defined by the experimenters or curators. As a result, researchers need to mentally merge all the different networks, discover how they complement each other, and reason about the source of conflicts.

Adding to the challenge, several independent neuroimaging projects, such as the 1000 Functional Connectomes Project [10] and the Human Connectome Project [70], aim to provide an abundant quantity of large scale human brain network data to the community. These projects provide an extensive resource to study brain networks in large populations; however, the exploration and understanding of brain connectivity across such populations remains a daunting challenge.

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 that might underlie different aspects of brain function. They also provide a convenient framework for comparing neural connectivity in normal and diseased brains, and thus could be used to diagnose neurological disorders 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 fraction of them study networks within the brain directly or would benefit from better understanding connectivity in the context of their research problems.

Connectivity Analysis in Genomics Interactions, relationships, networks, and hierarchies are ubiquitous in models that explain gene sequence and expression data. Such networks are studied by three of the research labs involved in the proposed work.

Understanding genomic activity and variation is critical to a broad array of high-impact scientific applications. High-throughput genomic approaches and rapid declines in the cost of DNA sequencing are producing great advances in the amount of data available to researchers in these areas. It is more important than ever to be able to analyze the complex interactions of these genes and protein signaling pathways.

As an example, the immune system, which the Benoist lab studies, has been compared to the brain in terms of its ability to process information from the environment, adapt quickly in response to challenges, and store information in long-term memory. Perhaps because the immune system is more accessible to experimental dissection than the brain, there is a good knowledge of the bewildering complexity of cell-types it includes (>250 different cell types can be distinguished). A full description of the variation in the genome's activity in these cell-types, during their differentiation or their responses to pathogens or antigens, are fundamental to understanding how the immune system works. The Immunological Genome Project [29] (ImmGen), a consortium of laboratories led by Benoist, has created a microarray compendium of all immunological cell types. Computational biologists that were part of ImmGen have started using this vast collection of data to create regulation models, grouping genes into discrete functional modules and determining likely regulators for each module, thus deriving network structure of regulation in immune cells.

In another example IGERT students working with Rand and Raphael are pursuing metagenomics, a new discipline of large sequencing of DNA/RNA from an environment. Metagenomics is enabled by the rapidly declining costs of DNA sequencing technologies. Brown University hosts an IGERT in Reverse Ecology (PI: Rand), that studies genomic variation across individuals and how genomic variation in a population is impacted by changes in the environment. A metagenomics study generates terabytes to petabytes DNA/RNA sequence data, with this data being billions of short sequences, rather than full-length genomes.

Algorithms are used to automate the assembly and alignment of this sequencing data, but these algorithms do not resolve the data unambiguously. Moreover, the DNA/RNA sequences alone do not directly answer the underlying biological questions. Just as one example, one driving scientific question in the IGERT program is how aquatic microorganisms adapt to varying amounts of salinity in marshlands. Answering this question requires analysis that integrates ecological measurements, organismal diversity, gene sequences, and interactions defined by metabolic networks within and between species.

Biologists must be empowered with tools that allow them to generate and test hypotheses in high-dimensional genomic and systems biology data. Many visualization programs have been introduced to address parts of this challenge: e.g. a multitude of genome browsers, including UCSC, Ensembl, Integrated Genome Viewer (IGV), and Galaxy allow interactive display of genomic data according to a reference genome sequence; Cytoscape is widely used for network data. However, the ability to interact across multiple biological scales from the individual genome through gene interactions within a species, and to interactions across species requires new visualization and interaction paradigms.

One final genomics example involves analysis of full genomes from many humans. As tens of thousands of human and cancer genomes are presently being sequenced (e.g. through NIH projects like the 1000 Genomes Project [2] and The Cancer Genome Atlas) there is increasing demand for tools to interpret this data. Interactive environments that combine automated interpretation from algorithms with human expertise will leverage the strengths of both types of analysis. While ImmGen is a powerful resource for the immunology community, the expression data and the derived module/network structures are currently distributed using simple web based querying tools with minimal visualization components. For example, regulatory modules are simply listed in the order they were identified by an algorithm along with a few pre-computed heatmaps indicating the microarray expression patterns of genes and regulators contained in each module. While useful for browsing the modules and understanding their potential function, more advanced visualization tools are now needed to allow genomicists to think in terms of "multiple hop" regulation mechanisms.

Improved tools for analyzing these networks and combining multiple sources of data will make it possible for researchers to form new hypothesis and improve the speed and accuracy of their experimental workflows and their science.

a.3 Related Work in Models and Principles of Cognition and Perception In this section, we put our work in context with other HCI advances that have been influenced by cognitive theory. We describe analysis methods and principles that we will build on in developing scientific visual analytics.

Task and Workflow Analysis A major consideration when undertaking this predictive modeling task is that human actions must be analyzed across many temporal orders of magnitude. Here, 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 [6]. Illustrating this idea, Gluck developed a model using the ACT-R/PM architecture that predicted student performance on algebra problems based on eye movements [22]. Models have also been used to decompose arbitrary tasks and predict completion times, and to compare systems based on these predictions. In Project Ernestine [24], 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, even decades later this remains one of the canonical examples of such predictive models. Significant opportunities remain in combining and cascading predictions about human cognition and behavior into actionable theories.

Analyzing user behavior is a reasonable step in generating hypotheses about higher-level cognitive activities. Recent work in visual analytics has provided a rich set of eye-tracking analysis methods we plan to build on in our Expedition [7]. These methods can be used to derive predictions and hypotheses about attention and visual analysis strategies. Laidlaw's lab recently performed an observational study of visualization use by immunobiologists that found divergent strategies for scatter-plot analysis [17]; preferences for these strategies might be related to cognitive traits, like locus of control, that we can evaluate. A recent tool that enables similar types of user analysis is CrowdScape, which visualizes the interaction histories of crowd workers on MTurk in a way that allows easy identification of different types of behavior [59]. In our work, we will build on analysis techniques like these that have been applied for manual assessment of interface effectiveness and user behavior. We believe data like these, in addition to other variables like user traits and environmental factors (e.g., display settings), can be fit to models of cognitive states and activities that will inform visualization and interaction design.

We plan to draw on theories of information discovery and ideation to put user observations and stimuli into the context of the user's analytical process. For instance, Kerne and Smith created a theoretical framework that describes the coupling between digital representations and a human user's cognitive processes during the information discovery process [43]. The framework decomposes the information-based reasoning and problem-solving processes into some smaller, more tractable stages, and provides a basis for designing experiments that study the dynamics between the computational representations and individual cognitive stages. For example, it was used to compare how two different representations for information collections might affect the emergence of new ideas [42]. The framework could be a valuable reference as we investigate aspects of the user's cognitive process to model, and when we define the scope of individual models.

Improving Interfaces with Predictions and Principles While Project Ernestine validated the use of cognitive, perceptual, and motor modeling to improve operator tools, other models of user state and principles of perception have been applied to visualization and interface tools. Recently, Solovey, et al., created Brainput, which uses near-infrared spectroscopy (fNIRS) to predict users' states, like high or low cognitive load, in real time. Measurements from the fNIRS sensor were fit to a model of cognitive load during a robot navigation task. These predictions were used to adapt the control interface to support multitasking [65]. In general, sensors that help measure states like stress or excitement are becoming increasingly more accessible and less intrusive, e.g., Affectiva's wrist-worn Q Sensor (`http://www.affectiva.com/q-sensor/`), and inform models of higher-level cognitive states. We believe these will be viable, critical inputs in future visual analytics tools that enhance analysts' and scientists' cognitive abilities. Horvitz's twelve principles for mixed-initiative interfaces provide a basis for designing interfaces that respond to predictions about user states and goals [30].

Forming perceptual hypotheses has also been proposed as a foundation for visualization design and evaluation. Ware advocates formulating perceptual theories for visual mappings, then evaluating and iteratively refining the theory [73]. These perceptual theories can be useful both as guidelines for designers and for automatically optimizing visualizations. Recently, Bateman, et al., found that stylized marks in infographics can sometimes improve viewers' recall of the data, without loss of accuracy [9]. This called into question older design wisdom about minimizing 'chart junk'. Our proposed work aims to evaluate similar design heuristics through scientific experimentation. These perceptual principles can also help inform and simplify models of visual analysis. Pineo and Ware used a simple model of the human visual system to simulate users' abilities to follow a particle in a flow field visualization [55]. While challenging, simulating human visual-search behaviors has the potential to transform some visualization design problems into optimization processes that can be done automatically or with little human involvement, like the human-in-the-loop algorithm of House, et al., for visual parameter optimization [31].

Finally, three more types of principles from the cognitive science community will be incorporated in this work when designing tools for scientists. First, principles of perception and attention will determine the physical parameters of the display. For example, recent work in the study of attention has examined the number of objects that can be tracked at any given time [57] and the physical parameters that determine the identity of an object [62]. Second, principles of goal selection will determine the number of tasks made available to the user and their accessibility at any point in time. Research has studied the process of multitasking using a variety of paradigms including prospective memory [50], task switching [58], and goal priming [3, 8]. Third, principles of problem-solving and reasoning will determine the design of visualization tools; the information that is on display at any given time; the actions that are given priority status; and the guidance and feedback that it provides. This principle will guide the design in ways that avoid common cognitive errors such as "confirmation bias," the tendency to seek out evidence to support a conclusion and ignore evidence that disconfirms it. The most simple and effective way that psychologists have found to reduce this bias is to ask people to "consider the opposite" [47]. A second remediation method is to ask scientists to think diagnostically when evaluating visualization features (e.g., "What is the likelihood that this design will help my analysis?") rather than causally [19] in order to consider a wider field of possible designs. We will connect our findings back to these cognitive principles and theory in order to generalize our results to other visualization domains and broader topics in HCI.

Expanding Evaluation Metrics Our proposed work will be used to evaluate interactive visual tools using cognitive analysis. Recently the HCI community has acknowledged that to effectively evaluate software in more realistic scenarios, we need more evaluation metrics that go beyond measurements like task response time and accuracy [25]. In particular, improving evaluation has become a focus of several calls to action in the field of visualization [56, 15, 39]. Some researchers have begun to describe new evaluation techniques that target other attributes of user state. Below, we describe a selection of metrics that have demonstrated promise as effective components of evaluations, and that could become pieces of our scientific visual-analytics framework.

One example of an evaluation method that goes beyond traditional metrics in visualization is *insight-based evaluation*, which tests how well a tool generates "aha!" moments [52]. Visualization is frequently used in applications aimed at enabling insight discovery through analysis. Therefore, measuring insight can sometimes assess the utility of techniques more directly than metrics like response time, especially in exploratory visualizations. This method has been used to evaluate scientific visualizations [60, 61], including the GREMLIN tool by Laidlaw's lab [53]. A significant challenge with the method is using experts to code insights manually; we plan to use our framework, which will collect and help analyze multiple streams of user data, to find data-driven ways of identifying insights that will make the method more practical.

Measuring and predicting *affect*, or emotional state, has also shown promise in better tailoring software and hardware to users. Lewis, et al., studied how user affect can be primed with digital stimuli, and linked changes in affective state to changes in ideation [45]. Recently, McDuff, et al., built an interface for recording and visualizing users' affective states to help the users learn more about their daily habits and personally reflect [51]. This work establishes feasibility for our plans to infer user states from lower-level user data, and are also aligned with our goal to support users' abilities to do scientific analysis. In fact, research has shown that affect is intimately related to the decision-making process [46].

Other factors will affect how we compare and use metrics like these in our work. Individual differences exist for software use [16] and must be considered when evaluating visualization tools. Crowdsourcing is becoming more popular in visualization user studies, but individual differences between participants are not well understood, despite earlier work that replicated classic graphical perception tests using MTurk [27]. To advance this area, we plan to evaluate cognitive models with both domain science experts and non-expert crowd participants when possible. We also plan to collect information about social and work environments for application users, based on previous findings that these environments can influence visualization efficacy [32]. Understanding the relationships between cognitive-perceptual models, individual differences, and external factors (e.g., display size and settings) will advance visualization theory and let us comprehensively evaluate and design scientific visual-analytic tools.

a.4 Related Work in Scientific Network Analysis

Network Visualization Network visualization in biological domains, like brain circuit analysis and genomics, remains an incompletely solved problem. Some online databases hosting connectivity information have developed their own visualization modules [67]. Others have developed personalized visualizations for specialized brain regions or organisms [1]. Still other systems include visualization features that link to online gene or protein interaction databases [11, 64]. We may be able to layer genomic data atop such systems.

However, these approaches each have one or more drawbacks: they fail to adequately merge findings from network visualization with scientists' intuition; they are limited in scope to singular organisms, atlases, or databases; they don't allow users to integrate their own experimental data into the analysis; or they don't offer analysis features such as load/save capabilities, or hypothesis-formation support [54]. We propose to address all of these limitations.

The work that comes closest to our proposed applications, Cytoscape [63], provides a graph visualization platform targeted at systems biology. Cytoscape lets users choose from over a dozen layouts for node-link diagrams and lets them manually assign many visual properties, ranging from the border color of node glyphs to the font face and opacity of labels. However, the platform does not help users know when and how to choose from the design parameters and algorithms it offers. In fact, this typifies the problem we will address in our proposed work by enabling the discovery and practice of principled visualization optimizations that improve network analysis by end users.

Knowledge Integration Some of the recent work in eScience and digital libraries have enabled transparent access to multiple information sources united under a global domain ontology. In [66], for example, a system was built that allows researchers to access multiple bioinformatics resources using declarative queries. The existing work provides valuable experience that we could build upon when integrating existing knowledge sources for brain and genomics connectivity. Our work will differ from the existing work though, in that we will focus on how to design visual interfaces on top of the integrated sources to facilitate scientific discovery.

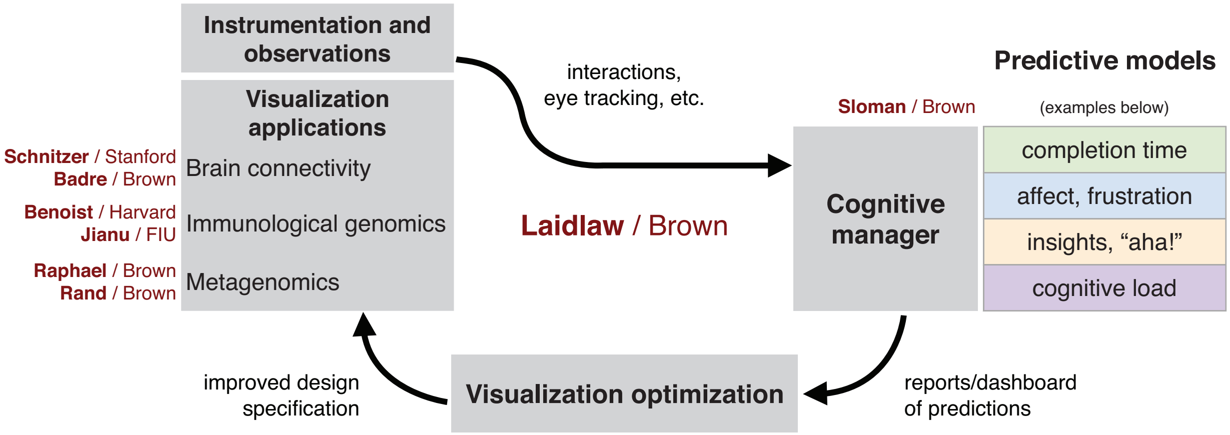

Figure 4: Proposed information flow for cognitively-optimized visualization. Labs and institutions are shown next to the applications and modules to which they will significantly contribute. Laidlaw's lab at Brown will oversee all parts and develop software for each. The predictive models shown on the right side are examples of models that can plug into the manager application (described in Sec. a.5) that generates data used to improve and optimize visualizations.

a.5 Proposed Framework for Human-Centered Visualization

In this section, we describe the design and evaluation plan for our human-centered visualization framework. The parts of the framework are illustrated in Fig. 4.

Application Development A spiral development cycle for applications will be used to implement an initial set of features based on requirements from domain experts in brain science and genomics. Sec. a.6 details proposed novel features informed by cognitive principles that will support scientists' analyses. We will prepare tools and practices for monitoring, recording, and encoding user actions and observations synchronized with application behavior. Applications will let us test new tools and practices that output encodings of data including: interview responses, user actions, video of participants, screen capture, galvanic skin response (GSR), and eye tracking.

Both system data and user data will be collected. System data include what the application is showing and environmental details, like display size and resolution. User data can be further divided into five categories: user intentions, user traits, user interaction histories, user states, and user analysis outcomes. Below is a list of examples of each of the five types:

- user intentions: goals, analytical tasks
- user traits: psychological traits, domain expertise, visual analysis expertise
- user interaction histories: sequences of input events, body motion, facial expressions, skin response, eye-tracking data
- user states: attention flows, cognitive load, affect
- user analysis outcomes: task accuracy, completion time, insight measures, user-experience measures

Among the five types of user data, intentions and traits will be largely descriptive. We will design standards for gathering and encoding these two types of information. In many cases, they will be used primarily for qualitative analysis. The other three types of information will be encoded and used primarily as inputs to the model manager application.

Evaluation of visualizations will include gathering qualitative and quantitative feedback from expert users in a laboratory setting. These evaluations will occur after iteratively implementing new features in the visualizations; they will ensure that we are producing tools that are enabling domain scientists to better analyze their data. Evaluation of new instrumentation and monitoring libraries will include quantitative profiling of data-collection speed, storage, and cost.

Cognitive Manager Application and Predictive Models The second phase in developing our framework will be designing a model manager application that receives and parses the system data and user data listed in the previous section. We will provide an API to access user data from the manager. An experimental aim is to use this manager application to fit user data to predictive models of user states or analysis outcomes over time. This aim is risky but could have high impact for visual analytics. If successful, this part of the framework will let us evaluate cognitive models relevant to scientific analysis as well as design principles for supporting visualization users based on models. We will focus on the following types of predictions:

- insights during the visual analysis process
- cognitive load
- frustration or engagement
- task accuracy and completion time

We will consider combinations of models as well. For instance, a user's frustration and his or her ability to have insights might be negatively correlated. In this way, predictions could cascade into one another to model higher-level behaviors. By incorporating models into a system, we will make hypotheses like this easier to test than is currently possible.

The manager system will use an open format like JSON to process user data and output predictions. This will enable future software to plug into the framework. Development of predictive models will be performed by Laidlaw's and Sloman's labs. Visualization tasks that will produce data to evaluate models will be designed with input from collaborators in the science domains. User studies with applications will be performed to collect data and compare predictions with reported user data, e.g., from interviews. In cases where specialized equipment like an eye tracker is not needed to validate a predictive model, we will perform quantitative studies with crowdsourced participants on MTurk in addition to domain scientists. Crowdsourcing some experiments with MTurk could give greater statistical power for testing hypotheses than using only traditional user studies.

Improving Visualization Applications with User Predictions As we develop predictive models that plug into the framework, we will manually review reports generated by trials in the user studies mentioned above. Laidlaw's and Sloman's labs will work together with input from Drury to codify a practice for reviewing timelines of user predictions alongside screen-captures. This review process will enable a human-in-the-loop solution for identifying high-level features or workflows that affect user states.

This process will be repeated as new models of user states and analysis outcomes are hypothesized and implemented. During this process, we will form hypotheses about design guidelines that let visualization designers in the future side-step negative user states (e.g., high probability of frustration, low probability of "aha" moments) and design for positive user experiences, like Jianu and Laidlaw's interface nudges toward effective visualization analysis [36]. We will refine and validate guidelines by applying them to applications and iteratively evaluating, i.e., cycling through the phases shown in Fig. 4. We will test interactions between new design principles, which could contradict one another. For instance, a new interface design might decrease user response time, but increase cognitive load or frustration. We will explore effective ways of communicating conflicts or synergies like this to the application designer. A long-term goal is building infrastructure for automatic optimizations that can be applied to aspects of visualization software based on predictions and principles.

a.6 Proposed Cognitive Support for Scientific Visual Analytics In addition to the framework described in Sec. a.5, the Expedition will contribute novel scientific visualization applications with features supporting the cognitive processes involved in analyzing connectivity and genome data. These features will be targeted at the needs illustrated in Sec. a.2. Our aim is to implement applications and use them as testbeds for optimizing visualization designs, features, and workflows for improved analysis, using predictive cognitive models and principles.

We describe the main features for the proposed brain connectivity and genomics analysis tools below. This is not an exhaustive list. We will follow an iterative development process, so features will evolve and emerge. Furthermore, we expect to generalize components of these features for both brain and genomics applications, and for future visual-analytics applications.

Brain Circuitry Analysis Application

1. *Seamless integration of multiple data sources.* The brain circuitry tool will provide unified access to major brain databases, including the Brain Architecture Management System, CoCoMac, BrainMaps and the Allen Brain Atlas, all of which are resources used by researchers from the Schnitzer and Badre labs currently. We will also build in pointers to publications on PubMed and integrate imaging data repositories. In addition to these curated data sources, the system will let users import their own experimental data and analyze them in the broader research context. Finally, the biggest challenge towards achieving seamless integration of multiple data sources is the lack of a uniform ontology as described in Sec. a.2. We plan to address this problem by creating a translation service for ontologies that operates in the backend and lets users query resources across multiple schemas.
2. *Integration of multiple 2D network visualizations and 3D structural visualizations.* We will provide an appropriate set of visualization techniques to users based on initial requirements for our collaborators' analytical tasks. This will make it easier to choose useful representations from the wide variety of existing network-visualization techniques and the analyses they best support [71]. Furthermore, we have observed a coupling effect between visual representations and users' mental data models (e.g., brain anatomy, signaling pathways) [37] that we will support when possible.
3. *Visualization of uncertainty throughout data flow.* Uncertainty can arise in multiple stages, including data collection, data transformation and modeling, and data aggregation both within and across different sources. We will monitor, estimate, and visually show users where uncertainty occurs as well as the degree of the uncertainty. The development of this feature will be tightly coupled with some of the cognitive modeling work, where we will extend existing research on uncertainty visualization techniques to study the relationships among the user's perception of uncertainty, the visualization, and the analytical task, and uses the research results to inform the design of uncertainty visualization for the tool.
4. *Automatic suggestions of visualization methods based on user tasks.* We plan to build a suggestion engine that generates suggestions for visualization methods based on the user's current analysis goal. This feature will be enabled by constructing and integrating a predictive model of the user's intended goal based on query and browsing history.
5. *Reusable user annotations.* We will let users input annotations about publications or data, including uncertainty for experimental results, compatibility of experimental techniques, relevancy of different resources, and unstructured notes. These annotations will be persistent between users and sessions to support long-term and collaborative analysis tasks.
6. *Hypothesis simulation.* As an experimental feature, we plan to design a hypothesis simulation engine that lets users input a hypotheses about brain connections using a simple, specialized grammar. We plan to display disconfirming and supporting evidence for these hypotheses from the existing knowledge base. Later, we will explore more complex interactions between hypotheses.

Brain-Circuitry Tool Prototype We have designed and made available online a prototype of the brain circuitry tool. The tool is web-based to allow easy access and future extension into a centralized portal to the existing online resources. The tool currently visualizes rat brain connectivity data from Brain Architecture Management System (BAMS) using two different network visualizations, one with circular layout and one as a force-directed graph. We have instrumented the interface to collect both high level user tasks such as "search for connections between Striatum and Thalamus" and low level user events such as mouse clicks and idle time. and save notes about connections and publications.

Preliminary user feedback from brain scientists in Schnitzer lab and Badre lab have inspired the proposed features for the brain circuitry application listed above. They gave positive comments on the prototype's ability to show them the existing publications in BAMS in a compact and easy to navigate manner, and were excited about the tool's potential to transform the current literature search paradigm into a more powerful, flexible exploration paradigm. However, they expressed varying degrees of satisfaction about the current ontology which is static and determined solely by the curation scheme used by BAMS, depending on their research focus. Also, limitations of the current available visualizations were observed when the users attempted to perform analytical tasks unanticipated before the prototype. They also expressed concerns about uncertainty in the data and suggested that annotation and hypothesis simulation would be valuable features.

Schnitzer and Laidlaw are now working together to gather a group of expert users to try out the prototype during their research. We expect to be able to collect preliminary usage data and user feedback which will be valuable for generating a more concrete plan for both the cognitive modeling and application development aspects of the project.

Genomics Connectivity Analysis Application

1. *Integration of multiple immunological data sources.* The proposed genomics analysis application will unify data-sources critical to analyzing immunological data in the context of regulation, pathways, and multi-species comparisons. Example data sources include NCBI, Ensemble, GTeX, and Kegg. It will also incorporate knowledge from publication repositories like PubMed, as well as experimental data from users.
2. *Integration of visualization and analysis tools.* We will integrate the tools that our collaborators' analyses rely on. To do this, we will draw on design paradigms from the information visualization and visual-analytics communities that support combining analysis techniques. For instance, while testing ImmGen, we used the Focus+Context paradigm [14] to structure information by schemas that are familiar to researchers: analytic questions, level of depth (e.g., organ, cell, genome), etc.
3. *Personalization of visualization complexity based on user expertise.* Our experience with ImmGen taught us that the level of complexity in visualizations needs to be tailored to researchers' preferences in order for these tools to have high adoption rates. After personalizing these applications, we will try gradually increasing the visual complexity and features in order to unobtrusively steer users toward richer analysis activities.
4. *Online collaboration platform.* We will deploy a significant part of our visualization framework online and augment it with collaborative research tools that support asynchronous analysis [28] at multiple levels of scales and integration: between researchers in a single group, between different groups working on a single project, or publicly, within the entire immunological community. We will instrument our website to capture annotation and analysis data along with interaction data.
5. *Visual methods for analyzing genomic data at multiple scales and scopes.* We will develop novel methods for decomposing genomic data that is typically large, diverse, complex, and interconnected. Initially, we plan to adapt the Focus+Context paradigm for metagenomics data to show small subsets of the sample, like observed genomic phenomena within a species, in addition to large-scale relationships, like how those phenomena are impacted by the increased prevalence of a particular virus.

Generalizability We believe that the functionality described in this section can be generalized to many other domains. With the drastically increasing amount of digitized information, more and more research areas are facing the challenge of effectively integrating and building upon digital information from heterogeneous sources. Data from many disciplines can be abstracted as networks, and deeper understanding of the coupling between the user's background, expertise and research questions, and the visualization types will likely to benefit any domain that uses network visualizations for analysis purpose. The experience and guidelines derived from designing a system that allows a user to ask "what if" questions and flexibly express their problems to the system will also likely be applicable to any application that is designed to support complex reasoning.

a.7 Preliminary Results in Capturing and Cognitively Analyzing User-Interaction Data

Laidlaw's lab has been one of the leading visualization labs in experimental evaluation of visualization methods [44, 20, 18]. Some of these evaluations were done in collaboration with perceptual psychologists from Sloman and Badre's department. They involved 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 visual design 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.

Working with researchers from Benoist's lab, Laidlaw's lab recently performed a task analysis of visual exploration of gene expression data in the course of an immunobiology research project. This analysis uncovered patterns of interaction and important differences in behavior across users [17]. In a related project, researchers in Laidlaw's lab are working on designing a model to capture visualization user tasks based on user interaction history. Such a model will help bridge the gap between observable user interactions and higher-level cognitive tasks, potentially easing the application of relevant cognitive principles. These findings and ongoing work in this area will inform our understanding of these types of cognitive scientific research tasks.

Additional recent work has focused on simplifying quantitative evaluation methods for visualization tools. Laidlaw's lab created a semi-automated system, Tome, that analyzes discrete user interactions with visualization GUIs and predicts task completion times [23]. These predictions are useful in evaluating task performance for a crowd of end users and can be modified to evaluate how unimplemented GUI features affect task speed. To do this, Tome interacts with the modeling tool CogTool [38] that simulates these tasks inside a cognitive modeling architecture; as explained in Sec. a.5, this coupling of models with software interfaces is a specific aim of the proposed work.

a.8 Preliminary Results in Developing Visual Analysis Tools for Science at Brown

Over the last decade, we have developed numerous scientific visualization tools that have taught us a number of lessons about science and about developing tools to support it. While some of these tools have been directed at our two driving application areas of brain science and genomics, our experience with other application domains is also relevant to the proposed work. Our experience provides us a solid foundation of knowledge, laboratory culture, and infrastructure 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.

Our experiences with network visualization will help in developing the multi-scale circuit and network 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 [37]. 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. Linked views of the signaling network and experimental data, and quick access to the literature and to database information about the proteins and signaling should apply to brain circuit analysis tools.

Our experience developing dissemination tools for ImmGen's genomic data will help us generalize the network analysis principles to all three genomic sub-areas. In the context of the ImmGen project, a vast effort aimed at documenting the immunological genome of the mouse, we developed methods of distributing the projects data in readily analyzable forms using the familiar Google Maps interface. Visualizations included a genome viewer, planar multidimensional data maps, heatmaps, and a network representation. In the context of cancer research we have developed a novel, interactive visualization model for comparative analysis of structural variants and rearrangements in human genomes, with emphasis on data integration and uncertainty visualization [53].

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 [21], diffusion MRI visualization and analysis [75, 34], carpal kinematics analysis [48, 49], and even archeology visual analysis [5]. Two publications in particular capture the brain science visualization work we intend to build on. In the first [35], 2D schematic diagrams of complex fiber tracts are displayed in a readily accessible stylized manner; in the second [34], multiple visualization approaches for viewing human brain connectivity are linked to facilitate interoperation and understanding.

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, Drury, is an investigator on this proposal. Laidlaw and Drury have collaboratively 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. Some of what we have learned is documented in examples cited here [41, 72, 4, 40].

a.9 Knowledge Transfer Public access has been from the outset an important aspect of this project. ImmGen data are already made publicly available prior to publication, following NIH guidelines for large community resource projects. Modeled on this, a new community oriented web portal will offer access to all scientific data we have curated: HCI, software, brain science, and genomics. Extended community features will be designed to make the sites into platforms for collaborative discovery so that users can not only consume information but also produce it in the form of annotations, feedback, discussion threads, and wiki entries. The portal will include an API and documentation for building on and sharing software in our proposed framework. Versions of the software 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, VRG3D, that we are developing with other NSF support (OCI-0923393). Releases will be advertised through blog postings, publications, direct contact, and during conference presentations. We will also propose to teach tutorials and classes at conferences. Additionally, our brain science and genomics participants will be encouraged to use the tools in classes and other educational settings in order to help students understand connectivity in an accessible visual form. We expect this data sharing plan to encourage both broad dissemination and the building of a lasting, long-term research community. Finally, as with any scientific endeavor, publication of results will be an important mechanism to transfer knowledge.

a.10 Timeline While our overall plan is enumerated at the beginning of the project description, some additional information about staging is included here.

Our plan is divided into three phases to take place over five years. Each phase will be driven by multiple evaluatable milestones, e.g., requirements documents, design sketches, design reviews, releases at various scopes, and experiments. The timing is approximate; we expect the stages to overlap.

Years 1-2: Knowledge sharing, data gathering, modeling, setup, and prototyping. In the first year 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. We will also evaluate and gather the external data that will be available for analysts. As a part of that data gathering, community web tools will be put in place. 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 [40]. 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 [33]. This will provide preliminary data to help ground our cognitive modeling efforts and provide feedback for software development. We will also set up the SourceForge repository, including its bug tracker and forums for developers and users. The initial design will incorporate the highest priority requirements and will be prototyped to demonstrate to the brain scientists and genomics researchers 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. Once a first prototype is in place, we will instrument the tool to gather usage data and anecdotal feedback from neuroscientists.

Years 2-4: 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 labs. This phase will include an iterative series of increasingly complete prototype systems, each of which will be deployed to the brain science and genomics 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. 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. During the second year, feedback will be gathered and analyzed more broadly, and we will employ community-building efforts utilizing SourceForge mechanisms and our web sites. User data capture will continue, as will internal improvements.

Years 4-5: Validation and dissemination. The final applications will first be deployed to the brain science and genomics labs, with the goal of incorporating their use into ongoing research. Researchers in these labs will empirically evaluate 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 final versions of the systems to the public in accordance with our data sharing plan.

a.11 Contributions to Expedition Program Goals The proposed research contributes to the goals of the Expedition Program as follows.

Goal: To catalyze research into deep scientific questions, hard problems, and compelling applications in CS-related areas. Our work reaches far into four different research areas: human-computer interaction, cognitive modeling, the study of connectivity in the brain, and the study of genomics. In human-computer interaction, a framework for predictive modeling of human interactions with computers has the potential to catalyze significant follow-on research in such models. The models we experiment with ourselves have the potential to accelerate others' research in user interface design not only in two scientific domains, but across many and likely outside of science as well. Additionally, there are significant open challenges in computing related to visualizing data in the proposed systems biology applications. These applications will load large-scale datasets – eventually, at the scale of the human genome or human brain complexity, for entire populations of individuals – and must support storage, navigation, transformation, and sharing of this data. Beyond computing, our proposed research addresses fundamental questions about how the brain works, from the level of cellular connections up to analytical reasoning. It also addresses fundamental questions of genomic variation across individuals and environments.

Goal: To inspire pursuit of careers in CS and engineering. The students and postdocs in this unusually interdisciplinary research will likely continue on in research careers that reflect some of the novel aspects of the work. The PI's experience with earlier interdisciplinary research projects is that they tend to attract students who are not attracted by more traditional computer science. Students interested in biology or visual design tend to be more diverse than those interested in computer science and engineering. Brown University also devotes considerable attention to creating an environment of diversity in its student population. The proposed work will be able to draw on that diversity.

The proposed work will also add significantly to three computer science courses at Brown that link education and research. One, "Interdisciplinary Scientific Visualization" centers around designing and executing research projects by emulating the US model of research design and funding. Students identify a research problem with a collaborator from another discipline, explore potential solutions, write a "funding" proposal, peer review the set of proposals, do the research, write it up, and present it all during one 13 week semester. They get a taste of the excitement, challenge, and risk inherent to interdisciplinary research in a context where the real risk is minimal. This class will serve as a first line of outreach for our proposed work, broadening exposure from the handful of students directly involved as research assistants to a dozen or more each time it is taught. From past experience, we expect that some of these students will go on to participate actively in the proposed work or other research projects. A number of past students in this class are now in successful research careers.

A second course, "Cognition, Human-Computer Interaction, and Visual Analysis" is a research seminar focused on how humans and computers can interact effectively when performing scientific analysis. Students in the course complete semester-long research contributions, which in the past have included user studies on visual search using touch screens; a design for EEG visualization based on a task analysis; and analyzing the effectiveness of using MTurk to perform quality control on MRI scans. These students will benefit by having access to a cutting-edge investigation in human-centered visualization happening on-site at Brown with interdisciplinary, complementary collaborations. We expect that students will have extraordinary opportunities to participate in larger, longer-term experiments than in the past by building off the Expedition agenda, and collaborating with the PIs and research teams.

A third course that will benefit from this research is "Virtual Reality Design for Science." This course, jointly listed and taught at Brown and the Rhode Island School of Design, teaches design students enough science so that they can author new interactive tools for scientists. We plan to automate the process of evaluating these interfaces without going through the months-to-years implementation process, providing a demonstration of the acceleration our research will make possible.

Goal: To stimulate research, education, and knowledge transfer that promise scientific, economic and other societal benefits. Our plan to make the modeling testbed, the captured interaction data, the interactive brain connectivity analysis software, and the genomic analysis software all widely available should effectively transfer knowledge into several scientific communities. Both binary and source distributions will be provided. The research and development that ultimately result using these resources should provide broad scientific benefits, as already described. In particular, we expect undergraduate students and junior graduate students who have just started doing research in the two driving application areas will benefit significantly from the software that will be developed. Junior researchers are often well-trained in their knowledge about the brain regions or genes which are central to their research focus, yet might not have the system-level knowledge necessary to tackle some of the harder research problems. Traditionally, the only way to acquire the necessary connectivity knowledge is to comb through existing publications, which can be very time-consuming and even discouraging. The software we are going to develop will help bring down the barrier with the centralized, easily accessible system-level knowledge, and allow junior researchers to start early on thinking about deeper research problems. Anecdotal feedback from several junior researchers on the prototype for the brain circuitry software has already shown the tool's potential in helping them to gain more insight and ask deeper questions. Broader societal benefits that may result are difficult to predict specifically, but more effective ways for people to use computers has great potential.

a.12 Demonstration of Expedition Characteristics In addition to matching the Expedition program goals well, the project we propose has characteristics expected of expedition projects.

Characteristic: Foster creativity, informed risk-taking, and synergy. Our project is creative in bringing together several disciplines and in advancing all of them. These coordinated advances, each dependent on elements from other disciplines, are indicative of an Expedition that is greater than the sum of its parts.

Our research will continue a tradition of research creativity within and across the investigators' research groups. Our groups have already demonstrated that they value complementary and collaborative research, and Brown has a long tradition of blending education and research as effectively as any institution we know. If anything, we expect that the risk inherent to the work we propose and the environment we may build is beyond what reviewers will view enthusiastically. But we feel the risk is worth the potential rewards.

Characteristic: Draw upon well-integrated, diverse teams. Our proposed work is structured around a vision that integrates all of the investigators – each is essential to the success of the project. The investigators' work is diverse; six different departments are represented by our group. Each member brings unique abilities. Although only three of us are formally in a computer science department, our vision is centered around a core computer science problem: understanding human computer interaction and applying that understanding for scientist users.

Characteristic: Stimulate effective knowledge transfer. We propose to spend significant effort to transfer knowledge across disciplines and outside of the immediate research group, as outlined in Sec. a.9.

Characteristic: Demonstrate elements that enable discovery. Our experimental data capture, cognitive modeling testbed, and brain connectivity analysis tools are directly targeted at enabling discovery.

a.13 Summary of Synergy

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. The synergy of the group is outstanding. Though relatively small, it covers a surprisingly large range of disciplines. Although the group includes computer scientists, artists, psychologists, genomics researchers, and neuroscientists, the shared goals of understanding the brain and genomics brings us all together. Research in each of the disciplines is advanced by the proposed work, so motivation to succeed is shared by all. No discipline can be removed without fundamentally changing the nature and reducing the scope of the overall research – the sum is greater than what any subset could accomplish. The group and the project seem ideal for an Expedition in Computing.

b Leadership and Collaboration

We have assembled an excellent team of recognized experts in cognitive science, neuroscience, computer science, genomics, computational biology, and visual design. Their complementary expertise covers a breadth unusual in such a small group. The scientific problem areas of genomics and brain science and the research thrust of cognitive optimization of 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 years. Laidlaw has collaborated and published with all of the other participants save Rand and Badre. Badre and Sloman are in the same department and Sloman has attended one of Badre's classes. Laidlaw and Sloman participated together in a working group studying a dual-system model of cognition. Laidlaw and Sloman have taught each other's students and co-advised as well. Schnitzer provided the inspiration for attacking the brain-connectivity problem by contacting Laidlaw two years ago, and their groups have been interacting since on the neuroscience needs, a prototype "MindMap" that Schnitzer and his group developed, and an interactive prototype developed by Laidlaw and his group as detailed in Sec. a.6. 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. Drury has also participated in a number of perception oriented studies of visualization and interaction tools. Laidlaw, Benoist, and Jianu have worked together closely on the ImmGen project over the last 5 years and are co-authors on the major publications from that project. And Raphael and Laidlaw have collaborated on a genome-rearrangement visualization tool applied to cancer genomes [53]. Raphael and Rand work together in the Center for Computational Molecular Biology, where an IGERT they both participate on will leverage the tools we propose. Rand also collaborates with researchers at the Marine Biological Laboratory (MBL) at Woods Hole, and an institutional relationship between the MBL and Brown will facilitate collaboration on metagenomics analysis.

Brown provides a supportive environment for multidisciplinary work such as we propose. The diverse student body is 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. This institute also provides numerous brain scientists who may be interested in using and testing the tools that we will develop. Several other centers at Brown also provide important capabilities that we will leverage, as outlined in the Facilities section.

Because of the relatively small size of the group, the existing relationships among the members, and the natural incentives of the project, the management structure will be relatively simple. The PI, Laidlaw, will be responsible for monitoring the program and interacting with the other research lab leaders to track progress. Because the work is so interdependent, shared deliverables and dependencies will help to keep the research labs synchronized. A full-time postdoctoral scholar will help Dr. Laidlaw with this coordination. While this leadership structure is simple, Laidlaw has had excellent results with similar organizational structures in past collaborations of similar scope. One of the reasons this works well is that almost all of the elements are naturally advantageous for the the participants doing them. None of the research groups is acting as a service to the others. The research goals are linked so that the overall project is a win for all of the participants, providing natural incentives that will lead to success.

c Intellectual Property and Ethics

We intend to provide open access to the tools we produce, the data we acquire, and our modeling efforts. We do not intend to commercialize the work we propose. Data access will be consistent with the protection of human subjects as overseen by our IRBs.

Brown provides research ethics training to all students and postdocs who receive funding, and all will be required to participate. Brown also provides a reporting system for anyone with research ethics concerns and an ombudsperson available to faculty and postdocs. Each faculty member in our research team regularly meets with the researchers they manage, and part of that mentoring includes the opportunity to engage in discussion of ethical issues that may arise. The PI will also be available to all personnel at all sites as needed to discuss and help resolve any ethics questions or concerns.

d Experimental Systems and Shared Experimental Facilities

Much of the proposed work requires facilities typical to each individual investigator, e.g., computers and laboratory space. We propose to acquire a few resources specific to this project including a pupil tracking system. We will also leverage a number of interaction and display devices already in place in Brown's Computer Science Department and its Center for Computation and Visualization (CCV). These include tiled display walls, stereo-enabled desktop displays, an ultra-high-resolution Wheatstone stereoscope, haptic devices, and a virtual-reality cave expected to produce first light in mid 2013. This NSF-funded cave will feature a 360 degree field of view, resolution at the limit of human perception, and head tracked stereo imagery. It will be able to emulate displays of many different form factors because of its display properties.