

From Visualization to Visually Enabled Reasoning

Joerg Meyer¹, Jim Thomas², Stephan Diehl³, Brian Fisher⁴,
Daniel A. Keim⁵, David H. Laidlaw⁶, Silvia Miksch⁷,
Klaus Mueller⁸, William Ribarsky⁹, Bernhard Preim¹⁰, and
Anders Ynnerman¹¹

- 1 University of California, Irvine, USA – jmeyer@uci.edu
- 2 Pacific Northwest National Laboratory, USA – Jim.Thomas@pnl.gov
- 3 University of Trier, Germany – diehl@uni-trier.de
- 4 Simon Fraser University, Vancouver, Canada – bfisher@sfu.ca
- 5 University of Konstanz, Germany – keim@inf.uni-konstanz.de
- 6 Brown University, USA – dhl@cs.brown.edu
- 7 Vienna University of Technology, Austria – silvia@ifs.tuwien.ac.at
- 8 Stony Brook University, New York, USA – mueller@cs.sunysb.edu
- 9 University of North Carolina, Charlotte, USA – ribarsky@uncc.edu
- 10 University of Magdeburg, Germany – preim@isg.cs.uni-magdeburg.de
- 11 Linköpings Universitet, Sweden – Anders.Ynnerman@itn.liu.se

Abstract

Interactive Visualization has been used to study scientific phenomena, analyze data, visualize information, and to explore large amounts of multi-variate data. It enables the human mind to gain novel insights by empowering the human visual system, encompassing the brain and the eyes, to discover properties that were previously unknown. While it is believed that the process of creating interactive visualizations is reasonably well understood, the process of stimulating and enabling human reasoning with the aid of interactive visualization tools is still a highly unexplored field.

We hypothesize that visualizations make an impact if they successfully influence a thought process or a decision. Interacting with visualizations is part of this process. We present exemplary cases where visualization was successful in enabling human reasoning, and instances where the interaction with data helped in understanding the data and making a better informed decision.

We suggest metrics that help in understanding the evolution of a decision making process. Such a metric would measure the efficiency of the reasoning process, rather than the performance of the visualization system or the user. We claim that the methodology of interactive visualization, which has been studied to a great extent, is now sufficiently mature, and we would like to provide some guidance regarding the evaluation of knowledge gain through visually enabled reasoning. It is our ambition to encourage the reader to take on the next step and move from information visualization to visually enabled reasoning.

1998 ACM Subject Classification I.3.8 Applications, J.0 General

Keywords and phrases Interactive Visualization, Reasoning

Digital Object Identifier 10.4230/DFU.SciViz.2010.227

1 Introduction

Visualization of information alone does not provide new insights. It is the discourse between the human brain and the masses of information that enables reasoning and analytics. Information visualization and the science of interaction must not only focus on rendering

performance and methods of human-computer interaction (HCI), but also on successful reasoning. Traditionally, research on HCI uses results from cognitive sciences to enhance the user's experience and performance when interacting with a visualization system. A successful visual data analysis system is usually characterized by an improvement in these two user-related categories, measured by means of a user study.

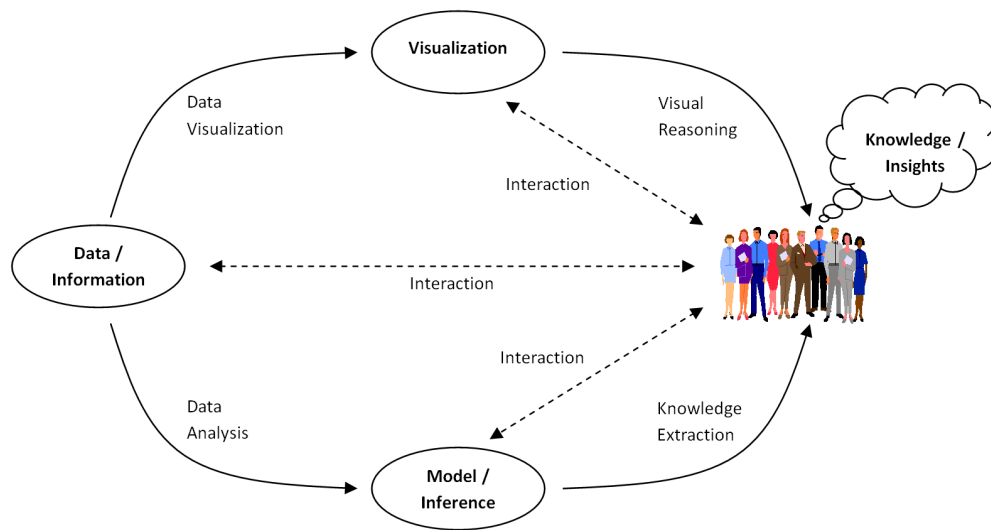
The next step in the evolution of visual analytics is the integration of interactive visualization and human reasoning. Interactive reasoning is a process that takes place when a dynamic visualization system responds to the user's input and aids the user in gaining *new insights*. A tight coupling between *cognition*, *interaction* and *visual analytics* is necessary to enable the user to make informed decisions. We suggest a new metric, which is required to evaluate the efficiency of visually enabled reasoning.

The next parts of this chapter provide an overview of some key definitions and define the new science of visually enabled reasoning. Chapter 2 gives a strong motivation for the shift from *visualization* to *visually enabled reasoning*. Chapter 3 outlines a roadmap to *visually enabled reasoning* and describes visual reasoning tools. A discussion is provided in chapter 4, which is intended to motivate the reader to take an integrated approach and to think about visual analytics as a tool that empowers humans to make better informed decisions. As a special feature, we present cases where *interactive visualization* and *visual analytics* were successful in *enabling human reasoning*.

1.1 Definition of Terms

Visual Analytics — Visual analytics is defined as the science of analytical reasoning facilitated by interactive visual interfaces [51]. This research area emerged from the fields of information visualization, scientific visualization, and data analysis. Thus, visualization as one method of analyzing data is combined with data analysis, e.g. statistics with respect to correlations, clusters and distributions of data. Visual analytics goes beyond visual data mining [28], which also emphasized a combination of data analysis and visualization by incorporating HCI. Visual analytics focuses on data analysis by means of interactive data visualization and human-computer interaction. It is understood that the human visual and cognitive system is the most powerful tool for understanding the complex relationships between data elements. Adapting a dynamic visualization system to match the abilities of the human cognitive system is an instrumental key to optimizing user experience and performance. A system that matches visual perception, with respect to resolution, focus, attention and detail without overloading human senses is most suitable for efficient interpretation of large data sets. Human reasoning is an important and indispensable element in this process. It is important to note that the cognitive process of reasoning does not end with a good understanding of the data. Conclusions must be drawn, leading to actions. These actions include planning, decision making, and going back to the visualization in order to change the simulation or planning scenario.

Science of Interaction — The contemplated science of interaction through visual interfaces is broad. From the human perspective, it has both perceptual and cognitive components. Since we are focused on visually enabled reasoning, we will concentrate mostly on the cognitive component here. As shown in Figure 1, interactive visualization is the pervasive mediating component between the user and data, i.e., between information and insight. The science of interaction must formulate principles for all these interactions out of which both descriptive and predictive models must arise. From these principles and inference models, rules of design for interactive visualization systems and for the interactive reasoning will be formulated. There will be general rules that will apply across all visualizations, and



■ **Figure 1** Visually Enabled Reasoning.

there will be specific rules, based on the general rules but developed for domain-specific problems. In this set of principles and models, there must be a model for the human cognitive process engaged in the rhythm of interaction, that is, in simultaneously probing data and analyses, and in assessing visualizations. To emphasize, the human cognitive process must be modeled explicitly as an integrated part of any model of visually enabled reasoning. Once modeled, it can also be evaluated, and the efficiency of the reasoning process can be assessed (Figure 1).

Cognition — The Oxford Dictionary defines *cognition* as “the mental action or process of acquiring knowledge and understanding through thought, experience, and the senses ()”. In the context of visually enabled reasoning we extend theories of problem solving and reasoning to include ways in which experience and the senses integrate human knowledge and expertise with information from interactive visual analytic technologies. The ability of these systems to shift the burden of information processing to perceptual processes should enable cognitive operations to take place at a higher level over more complex cognitive “chunks”. In this view, cognition cross-cuts levels of the perception/decision/action hierarchy, and can best be examined within the framework of abilities and constraints that make up the human *cognitive architecture*.

Interactive Reasoning is the process of distinguishing between ideas in order to create new relations and insights based on collected evidence. A significant reasoning process is the building and testing of hypotheses (which may involve choosing between competing hypotheses). A hypothesis can either be an argument (e.g., a decision process that determines a course of action or point of view) or a model (e.g., a predictive representation of how something works). Evidence and also reasoning artifacts are derived from relevant information, data, analyses, or previous knowledge. Reasoning is intrinsically interactive. Here we are concerned with reasoning mediated by interactive visualization where the visualization is the means of presenting to humans the relevant information, analyses, and knowledge and also the interface through which the human manipulates the information, analyses, etc. to advance the reasoning process.

Visual Reasoning usually incorporates different people with considerably different qualifications and backgrounds. There are people developing tools for visual reasoning, and other people using these tools to analyze their data. In most scenarios, however, the people who acquire data (the authors of the data) are different from the actual customers. In medicine, for example radiology technicians and radiologists acquire data using a specific imaging protocol based on the request of a referring physician who is the actual end-user. A surgeon, for instance, might use these data for preoperative planning. The development of visual reasoning tools requires a close collaboration of all stakeholders from the onset.

The concept of participatory software development which is well-known in the HCI community since the 80ies [20][24] is essential to enabling end users to influence design decisions and ensure that the constraints of the particular domain are included in system design at an early stage. Visual reasoning and problem solving strategies are very complex issues, and it is very difficult to design them right from the beginning with reliable formal specifications directly for supporting scientists. The process of tool development based on a preliminary requirement analysis and later determining all the flaws of the design in a user study is just not effective. Instead, the authors and the users of the data should be considered as active co-developers and not as passive sources of information, who are supposed to answer questions such as "How often do you need this feature?" This active involvement of end users requires that tool developers communicate in a not too technical way with the domain expert. In particular, formal specifications, such as nested state transition diagrams or even use cases are not a good basis for a discussion with end users. The scenario-based approach by Carroll and Beth Rosson [43] is promising, because it yields specifications which can be discussed with end users, as well as guide the design, development and evaluation process.

Dynamics — Modeling human dynamics is a critical part of understanding and evaluating visually enabled reasoning. The study of human dynamics includes the adaptability and interaction between models and visual interfaces. Today, rapidly growing technologies such as Internet, mobile computing and sensor web have enabled new patterns of human interactions, from social networks to physiological functions [8]. Human dynamics has become more complex and more venerable. Unfortunately, our understanding of human dynamic behavior and machine interaction is very limited. Much is invisible. To make invisible visible is the goal of visual analytics, and to help model the complex, dynamic human machine interaction is the aim of this article.

Social networks represent a good model for studying complex human dynamics when many individuals and large computing networks are involved. The fundamental studies of social networks such as the *six degrees of separation* and the *power law of linked interactions* shed light on the scalability of human networking activities. Those remarkable models enrich our in-depth understanding of the dynamics in a very large network, which is a challenge to a visualization system.

Insight Gain — Gaining insights is a main goal of visual analytics or information visualization methods, however, as Yi et al. point out [57], although a few definitions of insight exist, no commonly accepted definition has emerged in the community. For example, Card et al. [9] declare that *the purpose of visualization is insight, not pictures* (p. 6). Saraiya et al. [44] define insight as an "individual observation about the data by the participant, a unit of discovery" (p. 444). North categorizes insight to be *complex, deep, qualitative, unexpected, and relevant* [34]. Yi et al. [57] identified four types of processes through which people gain insight: Provide Overview, Adjust, Detect Pattern, and Match Mental Model. According to Chang et al. [11] the scope of definitions of insight in the visualization community differs from that of the cognitive community: the definitions of insight in the

visualization community are generally broader but vaguer than those in cognitive science. They suggest defining insight in two parallel meanings: (1) a term equivalent to the cognitive science definition of insight as a moment of enlightenment, and (2) a broader term to mean an advance in knowledge or a piece of information.

1.2 A New Science

We propose a new science of human interaction with information. This new science is based on insight gain through visually enabled reasoning. Interactive visualization, which has matured into a set of well studied tools, plays an important role in this endeavor. Traditionally, user studies have been used to evaluate performance and utility of visualization systems. We propose a new metric which measures human reasoning instead of visualization or interaction performance. We will ask what kind of experiments can be conducted to measure human reasoning and whether user studies are an appropriate method for evaluating this process.

The paradigm of *visually enabled reasoning* is a whole new kind of evolution. How do we create a model to measure human reasoning? How do we measure insight gain? In order to answer these questions, cognitive scientists, visualization researchers and human-computer interaction specialists must collaborate in an interdisciplinary effort to define both system specifications and metrics for human reasoning. This collaboration has already begun, as evidenced by initial steps toward a human cognitive model undertaken by visualization and cognitive scientists working together [23].

1.3 Importance of Interdisciplinary Collaboration

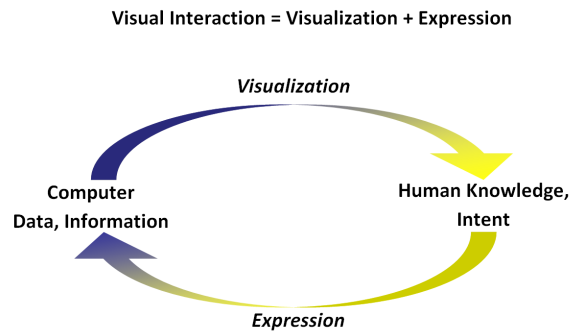
To enable scientists and domain experts to use advanced visual analytics system successfully requires a substantial knowledge of highly specialized scientific domains. This in-depth knowledge is usually not completely formalized. It is applied in an intuitive and implicit manner and thus cannot be easily extracted as a basis for computer support. As an example from engineering, the simulated flow in a designed engine model is analyzed with respect to complex flow patterns using a variety of analysis and visualization methods. In medicine, surgeons have to make treatment decisions based on the assessed risk associated with different interventions, often using interdisciplinary discussions, such as in a tumor board for their planning tasks. Intensive interdisciplinary collaborations are an essential basis for providing visual reasoning tools in such advanced application areas.

2 Motivation

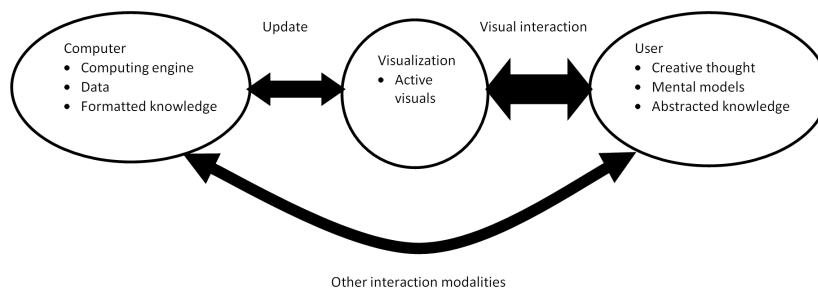
One of the key aspects of research on visually enabled reasoning will be focused investigation of the structure of human perception, cognition and action, i.e. human cognitive architecture. In this effort, we must understand how cognitive functions are distributed across perceptual and perceptually-guided motor processes.

For example, low-level perception of events is itself an inferential process. Irvin Rock's "logic of perception" [41][42] refers to the visual system's ability to compensate for inadequate sensory information (the so-called "poverty of the stimulus") and to reconcile conflicting sensory information through processes of unconscious inference. Perceptual inference is largely data-driven, and does not take into account the perceiver's conscious thoughts, beliefs, intentions, etc.

Pylyshyn's "cognitive impenetrability" test [37] distinguishes these two levels: while learning does train perception's inferential processes (and so individuals will differ one from



■ **Figure 2** Conveying visual information to the brain.



■ **Figure 3** Visualization-centric human-computer interaction.

another in their abilities, and experts will differ from novices), the perceiver's conscious thoughts, beliefs, intentions, etc. do not actively participate in perception. Thus input from end-users, e.g. verbal protocols of a "think aloud" session, can give only limited insight into their perceptual logic. The term "metacognitive gap" [17] describes this counterintuitive break between the ways in which these two logics must be understood, and hence the need for a new cognitive science of human interaction with visualization systems.

From the perspective of the cognitive architecture of sensemaking, when we understand a dataset we do so by attending to the conceptual implications of information that is itself constructed by the pre-conscious logic of perception. Visually enabled reasoning will be most successful when we fully understand how to design images and dialogs that enable the logic of perception to support the logic of conscious reasoning.

In Figure 2, visualization includes both the technical process of producing a visual display of information as well as the process of conveying this information into the human brain to form a mental image. Expression includes both the mental process of deciding what to express as well as the technical process of conveying this to the computer.

In Figure 3, the visualization is in the center and forms the medium that links the human and the computer. The widths of the links themselves illustrate the bandwidth of these channels. Note that the visualization-human bandwidth is perceived the largest, while that of the other modalities is much smaller. The visualization is not just an image, but an active and responsive process, equipped with autonomous methods, such as (semantic) zooming, graphics rendering, context+focus displays, etc. The visualization is updated by both user and computer, and it also stimulates or provokes new updates by both of these processes. In

this diagram, computer and human are equal processes both capitalizing on their individual strengths. The visualization forms a bilateral interaction medium.

The advance in computing power allows for the recomputation of parts of the visualization in real or near real time and thus enables modern interaction techniques like brushing and linking [16][53], panning and zooming [3], focus-context [33], magic lenses [4], as well as animated transitions. But not only the visualization, but also the underlying data and analyses can be partially recomputed.

Highly interactive and dynamic techniques are essential for supporting visual reasoning. With these techniques, changes between complex visualizations can now be directly observed instead of having to be interpreted which poses a much higher mental effort.

Visual or computational steering [54] is the concept that on the basis of a visualization of the current status parameters of a running simulation or analysis are changed and the results are visualized immediately. Thus, visual steering allows for the control of the simulation in real time, instead of running a simulation first and then doing a post-mortem analysis.

As a result, the visualization has turned from an end product of the analysis process to a user interface in an interactive and iterative analysis process.

3 Roadmap to Visually Enabled Reasoning

This chapter outlines the system components and conceptual considerations required for visually enabled reasoning. Most of the system components, such as interaction devices, dynamic and adaptive visualization, animation and sensory feedback, have already been developed and need to be combined with visual cognition techniques. We discuss tools for visually enabled reasoning and define metrics for evaluating these tools with respect to supporting the process of computer-aided human reasoning.

3.1 Science of Interaction

In terms of visually-enabled reasoning, interaction is the means by which the human and the computer work together (see also section 1.1 for a definition of terms). It makes the interface permeable and communication and collaboration possible. The visualization subsystem needs to be dynamic and adaptive, because the analysis and cognitive tasks can be significantly different depending on what direction the user takes in her investigation. Scientific research up to now usually employs user studies, task analyses, usability tests, and system performance evaluations for the evaluation, validation, and improvement of these systems. What has been missing to support visually-enabled reasoning is research that focuses on cognition, the reasoning and argument process itself, and the ability to make decisions. One aspect that must be focused on is the idea of a *mixed initiative system*, where the human and computer work together in intimate collaboration, each doing what it does best and then sharing results at the right time. New models that address the maintaining of cognitive flow for this mixed initiative system, through visual representation and especially through interaction, are starting to be developed [22][23], but much more work remains to be done.

Edward Tufte introduced the concepts of micro/macro reading, in which *detail cumulates into larger coherent structures* [52]. In this sense a figure or diagram can be graphically read at the level of larger contextual structure, at the detail level, or both view connected to (similar to Focus and Context techniques). A good example is the illustration of sleep and wake for newborn infants: each of the individual observations can be seen (micro reading), but collectively all observations reveal the larger 25-hour and then the 24-hour circadian cycles (macro reading). These concepts can be transferred to handle interactions. Interact

methods can be defined on the micro level (e.g., selecting data points) and on the macro level according to particular tasks.

The classical interaction methods proposed by Shneiderman [47] (overview, zoom, filter, details-on-demand, relate, history, and extract) or by Chuah and Roth [13] (basic visualization interaction [BVI] operations: graphical operations [encode data, set graphical value, manipulate objects], set operations [create set, delete set, summarize set, and others], and data operations [add, delete, derived attributes, and others]) can be seen on the micro level.

Yi et al. [56] studied different interaction methods and proposed a novel user intent-based categorization schema, which could be seen on the macro level as

1. Select: mark something as interesting (e.g., brushing),
2. Explore: show me something else (e.g., navigation),
3. Reconfigure: show me a different arrangement (e.g., swap x and y axis of a scatter plot),
4. Encode: show me a different representation (e.g., switching to a different visualization method),
5. Abstract/Elaborate: show me more or less detail (e.g., details on demand),
6. Filter: show me something conditionally (e.g., dynamic queries), and
7. Connect: show me related items (e.g., linking).

Yi et al. [56] also presented the different categorization schemata developed in the past, which makes quite clear that there is no established science of interaction available yet. However, according to the needs and purposes, different categorization schemata are used. As mentioned in Yi et al. [56], their proposed categorization schema does not cover all aspects, which were discovered in their literature and system research.

3.2 User Modeling

User modeling is a mature concept used in the design of intelligent user interfaces [29][30][31]. User modeling is concerned with deriving assumptions on the current user based on their previous activities with the system. These assumptions may be exploited to enhance the interactive use of a visual reasoning system. Thus, the system is adaptive with respect to its user.

The user model might be something as simple as a data structure with parameters and associated values. It might be more complex and represent also "why" the system has arrived at a certain assumption and describe the certainty of that assumption. In fact, sophisticated knowledge representations are exploited, e.g., in advanced hypertext systems [6][7].

Obviously, user modeling is a concept which is only applicable to systems that are used frequently; otherwise the user modeling effort does not pay off. Once this concept is implemented, experienced long-term users of a visualization reasoning system can be effectively represented and supported. At the lowest level, default values for visualization techniques, such as colors, parameters for diagram representations, isolines and so forth may be adapted to the preferences of the user. At a higher level, the selection of visualization techniques and the screen layout might be adapted. For instance, the system "classifies" the current task, looks for the visual reasoning strategies applied frequently in the past, and "suggests" a layout, where, for example, a 3-D visualization, a cross-sectional 2-D visualization and a 2-D histogram of data values and gradient magnitude are presented simultaneously. Similarly, the selection of data analysis techniques, such as cluster analysis, principal component analysis and correlation analysis might be adapted based on previous decisions of the user. Finally, the composition of data analysis results and visualization techniques is based on a huge parameter space and thus would benefit from narrowing this

space based on "intelligent" decisions. Therefore, a visually enabled reasoning system should include models of the following:

- User,
- Visualization Environment (Output Devices), and
- Data Analysis Methods.

These three components and their combination are essential for modeling visual reasoning.

User modeling, however, describes only one source for providing adaptive behavior. The selection and parameterization of visualization and data analysis techniques should also be adaptive with respect to the available computational output and device resources. Given the large variety of output devices, differing in spatial resolution, number of gray levels, physical size, and, for instance, the ability to render stereoscopic images, the suitability of visualization techniques strongly depends on the particular output device. Adaptive behavior should be guided primarily by existing knowledge of the properties of visual perception, including different abilities of user groups at varying age levels.

Finally, even the same user working in the same environment (computational hardware and output devices) often needs a variety of interactions and parameter adjustments for carrying out a "similar task". A large amount of these adjustments is due to the special characteristics of the data sets. Signal-to-noise ratio, frequency, the spatial distribution of unreliable data which should be removed, the amount and characteristics of inhomogeneity, and other imperfections of real-world measured data should be analyzed in order to "suggest" a meaningful sequence of visual reasoning steps.

The concept of exploiting the great potential of enhancing visual reasoning with adaptive components, however, is not straightforward. Successful examples for user models with the special requirements of visual reasoning in mind have not been developed yet. Existing strategies for collecting, structuring and representing user data have to be refined and evaluated. As a general issue of user modeling, privacy issues must be taken into account. Efficient user modeling inevitably requires collecting and analyzing large amounts of data of specific users. It must be ensured that this data is securely stored and that it cannot be accessed by unauthorized persons.

The experience of many users with adaptive behavior in computing software is often negative. In some cases, users could not predict the feedback of the system, resulting in a feeling of losing control. Moreover, the suggested solutions of the software often do not correspond with the choices that the user would make in a particular situation. Systems tend to over-generalize user interactions from the past. All of these problems have to be taken seriously.

Adaptivity has obvious limitations and therefore should be restricted to situations where a strong effect is likely to occur. Raskin pointed out that users are only effective if they form habits, if the system they use performs in a predictive way and if they are not overwhelmed by necessary, but semantically less relevant decisions [38]. If a system continuously changes its behavior — based on knowledge acquired by analyzing interaction patterns of the user — forming a habit becomes difficult, because the learning process of the user interferes with the adaptive system component [32]. Therefore, changes of the system's behavior should not occur frequently, they should be motivated and explained appropriately to the user, and they only come into effect after the user accepts them.

As stated above, a visual reasoning environment, in particular a collaborative one will potentially incorporate very different people with considerably different qualifications and backgrounds. These parameters will determine the inherent complexity and style of the visualizations used. A main obstacle in achieving true human-like artificial intelligence is the

fact that human consciousness is highly dependent on personal semantic models, knowledge, past experiences, skills, preferences, and the like, which are hard to capture and to encode in machines. We face the same obstacles when attempting the encoding of data and information into visual representations. These differences are expressed horizontally (same complexity, but different representation) as well as vertically (reduction of complexity). While the former is more a function of personal preferences, possibly motivated by professional or community background, the latter is a function of educational background, classification of the information visualized, and task-mandated (minimal) requirements. Thus, there will not be a single one visual reasoning environment that fits all participants, yet it must allow all participants to communicate with one another and with the computing engine as well. The key here is to develop a parameterized model of users and tasks, methodologies to acquire and test them, procedures to generate the user- and task-suitable visualizations, and finally appropriate means to translate one representation into another (also known as grounding [5][14]). For this, we need to capture personal preference vectors (in terms of visualization paradigms) and correlate them with other user information, such as background, education level, and others. These frameworks can then also be used to parameterize tasks and knowledge. A rich suite of user studies is needed to provide these models, and market research has developed statistical frameworks (such as conjoint analysis [19][26]) to efficiently acquire these, with a minimal set of users and user involvement. Such models will then yield adaptive user interfaces that can eventually predict the best visual representations of the information at hand. Finally, once such models are formulated, systems can automatically coach analyst users in the development of strategies or plans of attack for conducting more complex analyses. One will also be able to generate templates that cover the best-fitting strategies for most efficient analysis. Schneider-Hufschmidt et al. [45] provide a good overview on concepts of user modelling and using user models to enhance interaction, and most of the concepts described in this article are still applicable more than a decade later.

3.3 Visual Perception

Research in cognitive science aims at explaining how the human visual system creates perceptual experiences from visual stimuli. The results of this research have many implications for the design of visualizations. For example, due to the different densities and kinds of receptors in the human eye, color should be used for detail information in the user's focus, but not in the periphery (context), and motion can be used as a stimulus in the periphery. Many interesting insights have been gained by looking at visual illusions. Recently, Changizi et al. [12] proposed that many of these illusions are due to the predictive power of the human visual system that tries to compensate for neural delay. It is an open question how we can exploit this predictive power in the design of visualizations to convey information faster.

Another cognitive resource that current visualization techniques do not fully exploit is the human visual memory [46][48] — our ability to store vast numbers of (sufficiently different) pictures.

3.4 Visual Reasoning Tools

Visual reasoning tools are essential for a wide range of tasks. In supervisory control tasks, often characterized by large displays, it is essential that abnormalities are presented (and eventually aurally added) in an attention-grabbing manner. In exploratory tasks, a wide flexibility is useful to enable browsing-like undirected exploration. Finally, many routine tasks involve visual reasoning, not the least important are software assistants for medical

diagnosis. For routine tasks, flexibility should be deliberately restricted or at least hidden behind some "Advance options..." sheet for expert users only. Not only reduced parameter sets but also guided (wizard like) interaction to predefined steps are appropriate to enable users to build habits. The incorporation of analytic capabilities is essential for these three kinds of basic tasks. In exploratory tasks, often involving huge and multi-dimensional data, analytic and aggregating capabilities are essential to cope with the large amounts of data at all. In supervisory tasks, analytic functions may enable not only the localization of a potential failure but also the classification of the problem, the analysis of its severity and potential sources. Finally, in routine tasks, users may be directed to suspicious features. The development of effective tools in any of these task categories requires a model of the task, a model of the user groups as well as a model of visualization options including possible combinations guided by knowledge on human perception.

4 Discussion

The state-of-the-art in visually enabled reasoning has evolved from baby steps (beginning of visualization) to teen years (introduction of interactive visualization) to adulthood (maturing of visualization and interaction research and device technology). We are now approaching the senior years, and we are exploring new ways to use these established technologies for the next evolutionary step, which is visually enabled reasoning.

This chapter discusses challenges and shortcomings of current systems, proposals for improved interactive systems that take human cognition into account, and the potential need for a new science of interaction with information.

4.1 Current Shortcomings in Interactive Visualization Systems

The word *visualization* refers to the process of creating a mental understanding and notion of an object or phenomenon of which information is conveyed to the mind through our sensory channels of perception. Sometimes this is simply referred to as *insight*. Vision is just one of the sensory channels that can be used, but as it is in many aspects the dominating one it has lent its name to the creation of the mental notion of perceived subjects. Visualization is therefore in many cases facilitated by computer rendered images and relies on the power of the human visual sense to analyze the content of images. It should be noted that the mental process of visualization does not only rely on vision, but also makes use of our other senses. The reason for the dominance of the visual channel is found in the high information bandwidth that it creates to the human brain, as we are fundamentally visual beings. This is also reflected in the semantic metaphors we use to describe understanding — "I see", "To Visualize" or indeed "Insight". The inclusion of other senses in the visualization process is sometimes referred to as perceptualization, which emphasizes the collaboration of all human senses. This is a somewhat confusing use of the word and could also be misleading as it points more to the creation of the impression rather than the insight gained.

The apparent success of visualization based on the visual sense is only a part of the reason why other senses have been less explored in visualization. It can be argued that in comparison with vision other senses are more qualitative, less precise and harder to render input for. Also the equipment is expensive and the bandwidth low compared to vision. Despite this there have been many efforts to develop multi-sensory visualization. Most of these efforts have, however, not delivered the added value needed to compensate for the effort involved in generating the associated stimuli, and multi-sensory visualization has not been given high priority in the visualization community. Furthermore, immature implementations

and demonstrators have deterred many users from exploiting the potential of multi-sensory visualization.

There is, however, renewed interest in multi-sensory visualization. Rendering is now reaching a very high level of quality and to be able to bring visualization to the next level, addressing larger and more complex data, more information channels need to be utilized. It has also been shown to that, for instance, haptic interfaces can improve the process of visualization and shorten time to insight. To fully explore the use of multi-sensory visualization a rigorous understanding of what additional sensory channels can contribute, which methods are really effective, and for which applications the additional information channels can provide key information is needed. Used in an effective way it is deemed probable that the multi-sensory visualization can improve the visualization process as much or more than further improvement of rendering of images.

People tend to use both hands if they manipulate 3-D objects [25][40]. In medicine, two-handed interaction has been successfully applied, e.g., to pre-operative planning in neurosurgery. Hinckley et al. [27] argue that for the interaction tasks involved, e.g., exploration of a brain with free orientation of head and cutting plane, the most intuitive handling can be achieved with two-handed 3-D interaction where the dominant hand does fine-positioning relative to the non-dominant hand. In an empirical evaluation they demonstrated that physicians use these interaction techniques efficiently after only a short learning period.

4.2 Improved System Designs

The new science of interaction enabling visual reasoning as experience within visual analytics technologies and systems requires interfaces and interaction something specific of the data type and sometimes specific to the application. For example IN-SPIRE is a suite of technologies designed for unstructured text analytics [1]. The interlinking multiple visual representations, lists, multiple query types, temporal, affect, and other visual representation were deigned to effective text analytics. It applies to a wide range of applications from news, reports, blogs, planning documents, science articles, and many more. Some of the foundational interactions techniques are used with other data types but the implementations of the interactions are quite different such as those for video analysis of news [18].

Some data types and applications will quire unique interactions to ease the cognitive burden between the users, often non computer specialists, and their information. Such an interface can be seen in recent financial analytics system built specifically for fraud detection [10]. The many cyber applications based on billions of transactions are another call of application specific interfaces. Our thinking for a new science of interaction with visual analytics systems must include both the foundational interactions and some application and data type specific applications.

4.3 Cognitive Aspects of Reasoning

Alan Newell posited "bands" of mental activity ranging from biological-level neural firings over 10 ms and below, cognitive operations that take place on the order of seconds, rational activities that take place on the order of minutes and so on.

What Newell did not explicitly consider were that many temporal constraints on cognitive processing are due to the cost of acquiring information from the environment through motor activity. At lowest level eye and head movements strategically (albeit unconsciously) sample the visual world so as to support processes of perceptual inference discussed above. Newer cognitive science research suggests, the time required to execute an eye movement to acquire

needed information constrains the speed of cognitive processing [2]. In reading, for example, the processing time of a fixated word is slowed so as to enable the eye to have the time to make a saccade to the next word in the sentence. Given the hard constraint of the time required to make an eye movement, this "just-in-time" cognitive processing reduces the load on short-term memory in reading.

The correspondence of eye movement and processing times is characteristic of many perceptuomotor "interactive routines". These routines comprise epistemic actions that reveal information, externalizing actions that modify the perceptual world to reflect conceptual understanding, and coordinating actions that bind concepts to content. In the case of expert performers (e.g. skilled musicians or very experienced computer users) these interactive routines are effectively "compiled", taking place automatically under supervisory control of conscious problem solving processes.

This line of investigation has significant implications for visualization applications that support human reasoning. Wayne Gray's "soft constraints" cognitive cost accounting hypothesis [21][55] posits that small changes in the time required for the user to acquire information from a visual display can impact information comprehension and discourse and cause significant shifts in task performance and strategy. Work by Po et al. [36] demonstrated that presence of a cursor and delay in its response had profound impacts on users' ability to target display items via voice and pointing. The effects of these small changes in display response was attributed a shift between dorsal and ventral visual pathways.

If these theories are correct, interactions of temporal patterns in human-information dialog can interact with the intrinsic time course of cognitive processes to support or impede cognitive processing. This suggests that temporal rhythms in solo and collaborative use of technology can both detect and support "flow" [15] of effective cognitive processing and fluency of interaction. Addressing the sequential nature of human-information dialog will require new empirical methods that integrate mathematical modeling of sequences of interaction and human-mediated qualitative research methods.

Many of us are now working within research projects that involve user evaluations of applications and visualization tools at different points throughout the development process. Unfortunately we see a tendency in our community to rely almost exclusively on quantitative methods of evaluation, often associated directly with theories of cognition and sometimes taken without question or modification from the fields of HCI and cognitive science. While these fields are naturally close to our work and collaborations with HCI and cognitive science are and will continue to be very fruitful in our research and in evaluations, we would also like to point out that there are other approaches to user evaluation that can shed light on elements of visualization which will also be very useful for future work. Inspiration to this more varied and qualitative research on usability draws its inspiration from the groundbreaking study of human-machine interaction 'Plans and Situated Action' by Suchman [49][50]. Examples of qualitative methods that we suggest could be useful additions to our projects include:

Interviews with users — both interviews which follow a preset interview guide and open interviews (which allow the user to speak more freely about her/his experience of the tool) can be helpful not only in analyzing how well the user has succeeded in interpreting the data we are presenting, but also in discovering otherwise unknown issues with the tools. More importantly, this method can sometimes uncover questions to and about the data that are important to users but which were not specified in the research project's original remit. Often this is a result of the fact that 'the user' is generally a more heterogeneous category than we imagine it to be. Interviews with several different users can show how this heterogeneity impacts and is impacted by our visualization tools.

Discourse analysis — by this we suggest that it is sometimes useful for us to analyze the discourses which surround that data that we are attempting to visualize. Rather than relying solely on official documentation to describe the data which is provided to or included in our research projects, we suggest that analyzing more unofficial discourses (sourced through searches of academic, trade and news reports, for example) that surround a data set can also shed light both on alternative interpretations of the data and on other possible users, whose needs and impressions should also be considered when we are constructing tools with which to visualize data.

Observation of the tools in use — while this type of method can be used to make quantitative analyses of a tool (by counting eye movements or timing task completion, for example), we suggest that qualitative observations, based on the ethnographic methods like those employed within the field of Science, Technology and Society could also lead us to understand the way people actually use and, perhaps more importantly, 'misuse' our tools, forcing us to see our work in entirely different contexts. Much of this observation could and should be done outside of the laboratory, in 'real world' settings, to give us a better feel for how our research behaves and is experienced once it leaves our hands.

The emphasis on visual display and human visual processing in the literature reflects our understanding that the visual modality is central to human experience and the most likely candidate modality for effective technological enhancement of human reasoning. Behind this assumption lies a concern that adding finely-textured information from other senses — hearing and touch — will distract from the information better presented in the dominant sensory domain. We find this assumption unwarranted. Cognitive neuroscientists have long known that vision itself comprises multiple sensory channels with integration of those channels occurring in higher visual areas that must draw from multiple neural maps of smaller sensory channels such as color, shape, orientation, etc. Many of these areas also draw from sensory channels from hearing and touch.

4.4 Do We Need a New Science of Human Interaction with Information?

It is clear that the study of interaction must be significantly expanded and deepened from what has been done so far. How interaction works with visual display really hasn't been studied in any depth. Yet, it is clear with the new highly interactive exploratory tools that are now being developed that interaction is a very important part of the total system, even if it is not well understood. It is with this and the needs for visually-enabled reasoning in mind that the developers of the visual analytics research agenda called for a new "science of interaction" [51]. This call has been reiterated and progress so far on the science of interaction detailed recently [35][39].

4.5 Unexpected Discoveries

In the discussion of visual reasoning tasks, three categories were defined: exploratory, supervisory, and routine. Of these three, exploration is the most demanding in terms of reasoning, because it involves a process where the investigator does not know what she is looking for (at least not in detail). Hence, the process is one where discovery is emphasized. Once the investigator makes a discovery, she must assess what it means and how it fits into the context of what she already knows. Often a model or argument, which we call here the hypothesis, must be formed that weaves together known facts, the newly discovered evidence, the task at hand, and other relevant knowledge. The hypothesis has predictive

capability and is testable; new evidence must be collected to validate it. There are often multiple competing hypotheses, so one must gather further evidence that lends support to one or the other. Finally, the hypotheses may need to be modified based on new evidence collected or if there is a dynamic situation where circumstances change over time. All these processes are reasoning processes and can be complex and iterative. If, further, one has both complex reasoning *and* large scale, dynamic data with perhaps many related variables, then this is a visual reasoning task requiring visual analyses.

In the visual reasoning process, interaction is key for two reasons. First, exploration implies probing the data in overview, in relations within the data, and in detail. It is only by doing this that discoveries can be made. Second, exploration and discovery must intimately involve the investigator since only she can determine the context, meaning, and relations of the discoveries made. In particular, there cannot be an automated analysis that will extract meaning and relations because the nature of the discovery is not known beforehand and thus cannot be planned for. Therefore, interaction is key, and it must furthermore be through a visual interface.

We will give examples for the class of problems, involving both complex reasoning and large scale data, for which visual reasoning through an interactive visual interface is required. Certainly problems in bioinformatics are of this type. Researchers often have to compare and contrast annotated genomic data for several species, which includes large amounts of information such as relevant publications, various statistical analyses including of the sequences themselves, microarray results, protein expression results, and other information. Often a gene function can only be fully understood as part of a gene network, such as regulatory networks that suppress or enhance functions associated with specific diseases. These are complex reasoning processes involving exploratory analysis for which visualizations are required.

Another example is integrated computer experiments involving weather and environmental effects. Weather drives all these computer experiments (e.g., wind patterns, clouds, rain patterns, etc.), and to get high resolution results, which are necessary for detailed environmental impact studies, one must start with high resolution weather. The weather inputs then drive high resolution air quality models with hundreds of time-dependent output fields involving interacting chemical constituents and items such as various types of particulates. The results are truly stupendous in size and enormously complex in their relations and interactions. Two- and three-dimensional visualization techniques have been used for some time, but to study and understand the interacting 4D fields in the context of factors such as changing pollution sources, population and traffic patterns, and other factors requires visually-enabled reasoning. Without visually-enabled reasoning, there is no hope of understanding these complex processes. General, complex problems that have many applications include exploratory analysis and understanding of large scale collections of text or multimedia. Large text collections, of course, appear in many contexts from the above bioinformatics problems to business intelligence and legal evidence-gathering for large scale civil or criminal investigations. An example of the latter is the ENRON fraud investigation, which involved the sifting of billions of documents of all types, from email exchanges to internal memos and reports, by teams of lawyers. Visual analysis tools such as INSPIRE, developed by PNNL, have proved quite successful for these types of analyses, which cannot be done in such detail by any other method. In its latest versions, INSPIRE is also an example of a tool that is being enhanced to more fully support visually-enabled reasoning. Large scale multimedia collections (containing related images, text, video, closed captions, etc.) are also notoriously difficult to analyze. Unannotated image collections, for example, must

be sorted by hand, and even this categorization will not provide a view of all the relations one might need for an exploratory analysis. Likewise, video collections must be watched, even if annotated at some level, in order to understand, analyze in detail, and relate their contents. For collections containing millions or billions of images (such as collections that can be gleaned from the Web) or tens of thousands of hours of video (such as produced by broadcast news in a period of days), the task is unmanageable. Now exploratory visual analysis tools have been developed that effectively attack both these problems and point the way towards full visually-enabled reasoning capabilities.

The development so far of these tools for exploration and discovery indicates that they will also be quite useful and effective for more modes problems in both size and complexity. It is our position that visually-enabled reasoning tools with high interactivity will be of great use whenever one is faced with an open-ended problem involving the meaning of data or information.

These issues of exploration, discovery, and the role of interaction are, of course, also central issues in visual analytics, which is most succinctly described as "the science of analytical reasoning facilitated by interactive visual interfaces". But whether one approaches visually-enabled reasoning and the science of interaction from the viewpoint of visual analytics or from another direction (e.g., scientific visualization and computational science), the basic needs and the science that must result are the same.

5 Success Story

MSU ERC Space Shuttle Study — In 1990, the National Science Foundation established an Engineering Research Center (ERC) for Computational Field Simulation at Mississippi State University (MSU). The fulfillment of the center's mission is illustrated by the John Glenn space shuttle flight. The center has significantly contributed to the art and practice of "unstructured grid generation", yielding high quality grids in significantly less time. The center focused a team on coupling its structured grid CFD algorithm knowledge within a portable, scalable computational architecture onto unstructured grid solver technology. This required substantial research in both boundary layer gridding and solution algorithms. As it turned out, the parallel solver (research) code had just been assembled for the first time when the Space Shuttle mission STS-95 was launched. NASA Johnson Space Center called seeking simulated analysis of the Space Shuttle Orbiter during the return flight after the Orbiter drag chute door was lost during main engine startup. The NASA engineers wanted to know the dynamic pressure in the region of the missing chute door in order to estimate the aerodynamic loadings during reentry. The ERC group read a previously supplied Space Shuttle Orbiter geometry into the ERC's integrated simulation environment (SOLSTICE) and created the grids within hours. Initial simulation results were computed on a high performance computer within two days. The significance of this endeavor was not that NASA actually needed the results for successful reentry, but rather that the ERC had been able to take a tough real world problem and compute the solutions in two to three days after receiving the geometry description. This demonstrated an achievement that was a direct result of the researchers' ability to simulate very complex real world problems with complex geometries in relative motion. These accomplishments have come from directed cross-disciplinary efforts involving various technologies: grid generation, field solution algorithms, and scientific visualization, coupling human reasoning with computer and computational engineering. The task could not have been accomplished without combining all of the various talents and technologies. *(Source: NSF Engineering Research Center at Mississippi State University)*

Acknowledgements

The authors would like to thank Lucille Nowell, who is with the United States Department of Energy, for her valuable suggestions and encouragement. The authors would also like to thank David S. Ebert, Hans Hagen, Kenneth I. Joy, and Daniel A. Keim, the organizers of Dagstuhl (Leibniz Center for Computer Science, Germany) Seminar 07291 on *Scientific Visualization*, Wolfgang Lorenz, Technical Administrative Director, and Reinhard Wilhelm, Scientific Director.

References

- 1 IN-SPIRE. <http://in-spire.pnl.gov> (accessed March 2, 2010).
- 2 D. H. Ballard, M. M. Hayhoe, P. K. Pook, and R. P. Rao. Deictic codes for the embodiment of cognition. *The Behavioral and Brain Sciences*, 20(4):723–742, 1997.
- 3 Benjamin B. Bederson, James D. Hollan, Ken Perlin, Jonathan Meyer, David Bacon, and George Furnas. Pad++: A zoomable graphical sketchpad for exploring alternate interface physics. *Journal of Visual Languages and Computing*, 7(1):3–31, 1996.
- 4 Eric A. Bier, Maureen C. Stone, Ken Pier, William Buxton, and Tony D. DeRose. Toolglass and magic lenses: the see-through interface. In *SIGGRAPH Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*, pages 73–80, New York, NY, 1993. ACM. ISBN 0-89791-601-8.
- 5 S. Brennan, K. Mueller, G. Zelinsky, Ramakrishnan IV, D. Warren, and A. Kaufman. Toward a multi-analyst, collaborative framework for visual analytics. In *IEEE Symposium on Visual Analytics Science and Technology*, pages 129–136, Baltimore, MD, October 2006.
- 6 P. Brusilovsky, A. Kobsa, and J. Vassileva. *Adaptive Hypertext and Hypermedia*. Kluwer Academic Publishers, Dordrecht, Netherlands, 1998.
- 7 Peter Brusilovsky. Methods and techniques of adaptive hypermedia. *User Modeling and User-Adapted Interaction: Special Issue on Adaptive Hypertext and Hypermedia*, 6(2–3):87–129, 1996.
- 8 Yang Cai and Judith D. Terrill. Visual analysis of human dynamics: An introduction to the special issue. *Information Visualization Journal*, 5:235–236, 2006.
- 9 S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, San Diego, CA, 1999.
- 10 R. Chang, M. Ghoniem, R. Kosara, W. Ribarsky, J. Yang, E. Suman, C. Ziemkiewicz, and A. Sudijianto. WireVis: Visualization of categorical, time-varying data from financial transactions. In *IEEE Symp. on Visual Analytics Science and Technology (VAST)*, Sacramento, CA, Oct. 30–Nov. 1 2007.
- 11 R. Chang, C. Ziemkiewicz, T. Green, and W. Ribarsky. Defining insight for visual analytics. *IEEE Computer Graphics and Applications*, 29(2):14–17, March/April 2009.
- 12 Mark A. Changizi, Andrew Hsieh, Romi Nijhawan, Ryota Kanai, and Shinsuke Shimojo. Perceiving the present and a systematization of illusions. *Cognitive Science: A Multidisciplinary Journal*, 32(3):459–503, 2008.
- 13 M. C. Chuah and S. F. Roth. On the semantics of interactive visualizations. In *IEEE Symposium on Information Visualization (InfoVis)*, pages 29–36, San Francisco, CA, 1996.
- 14 H. H. Clark and S. E. Brennan. Grounding in communication. In L. B. Resnick, J. Levine, and S. D. Teasley, editors, *Perspectives on socially shared cognition*, pages 127–149. Washington, DC, 1991. Reprinted in *Groupware and computer-supported cooperative work: Assisting human-human collaboration*, R. M. Baecker (ed.), Morgan Kaufmann Publishers, Inc., San Mateo, CA, pp. 222–233, (1992).
- 15 M. Csikszentmihalyi. *Flow*. Simon & Schuster, 1994.

- 16 Stephen G. Eick and Graham J. Wills. High interaction graphics. *European Journal of Operations Research*, 81(3):445–459, March 1995.
- 17 B. Fisher, S. Fels, K. MacLean, T. Munzner, and R. Rensink. Seeing, hearing, and touching: putting it all together. In *ACM SIGGRAPH 2004 Course Notes*, page 8, Los Angeles, CA, August 08–12 2004. SIGGRAPH, ACM Press, New York, NY.
- 18 Mohammad Ghoniem, Dongning Luo, Jing Yang, and William Ribarsky. Newslab: Exploratory Broadcast News Video Analysis. In *IEEE Symp. on Visual Analytics Science and Technology (VAST)*, pages 123–130, Sacramento, CA, Oct. 30–Nov. 1 2007.
- 19 J. Giesen, K. Mueller, E. Schubert, L. Wang, and P. Zolliker. Conjoint analysis to measure the perceived quality in volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1664–1671, 2007.
- 20 M. Good. Participatory design of a portable torque-feedback device. In *Proceedings of the ACM Conference on Human Factors, Computing Systems*, pages 439–446, New York, NY, 1992. ACM Press.
- 21 W. D. Gray and W. Fu. Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, 28(3):359–382, 2004.
- 22 Tera Green and William Ribarsky. Using a human cognition model in the creation of collaborative knowledge visualizations. In *Proceedings of SPIE Defense & Security Conference, vol. 6983*, pages C1–C10, 2008.
- 23 Tera Green, William Ribarsky, and Brian Fisher. Building and applying a human cognition model for visual analytics. *Information Visualization Journal*, 8(1):1–13, 2009.
- 24 Jonathan Grudin. Why groupware applications fail: Problems in design and evaluation, Office: Technology and People. 4(3):245–264, 1989.
- 25 Y. Guiard. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of Motor Behavior*, 19(4):486–517, 1987.
- 26 A. Gustafsson, A. Herrmann, and F. Huber. *Conjoint analysis as an instrument of market research practice*, pages 5–45. Conjoint Measurement. Methods and Applications. 2000.
- 27 K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proceedings of ACM CHI Conference on Human Factors in Computing Systems*, pages 452–458, Boston, MA, 1994.
- 28 Daniel Keim. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 8(1):1–8, January–March 2002.
- 29 A. Kobsa and J. Fink. An LDAP-based user modeling server and its evaluation. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research*, 16(2):129–169, 2006.
- 30 A. Kobsa and W. Wahlster. Computational linguistics. *Special Issue on User Modeling*, 14(3), 1988.
- 31 A. Kobsa and W. Wahlster. *User Models in Dialog Systems*. Symbolic Computation. Springer, New York, etc., 1989.
- 32 Bill Kules. User modeling for adaptive and adaptable software systems, <http://www.otal.umd.edu/UUGuide/wmk/> (accessed May 18, 2009), University of Maryland, College Park, MD, April 19 2000.
- 33 Y. K. Leung and M. D. Apperley. A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction*, 1(2):126–160, 1994.
- 34 C. North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, 2006.
- 35 W. A. Pike. Science of interaction. *Information Visualization Journal, Special Issue on The Future of Visual Analytics*, 8(4), 2009.

- 36 B. Po, B. Fisher, and Booth K. S. Pointing and visual feedback for spatial interaction in large-screen display environments. In A. Butz, A. Krüger, and P. Olivier, editors, *Smart Graphics 2003*, Springer Lecture Notes in Computer Science, Heidelberg, Germany, 2003. Lexicle Ltd, York, UK.
- 37 Z. Pylyshyn. Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *The Behavioral and Brain Sciences*, 22(3):341–365, discussion 366–423, 1999.
- 38 Jef Raskin. *The Intelligent User Interface*. Addison-Wesley, 2001.
- 39 William Ribarsky, Brian Fisher, A. E. Turner, and William M. Pottenger. The science of analytic methods and reasoning. *Information Visualization Journal*, 2009. To be published.
- 40 Felix Ritter, Bernhard Preim, Oliver Deussen, and Thomas Strothotte. Using a 3d puzzle as a metaphor for learning spatial relations. In *Proceedings of Graphics Interface*, pages 171–178, Montreal, Canada, May 15–17 2000. Morgan Kaufmann Publishers.
- 41 I. Rock. *The logic of perception*. MIT Press, Cambridge, MA, 1983. ISBN 0-2621-8109-6.
- 42 I. Rock. *Indirect perception*. MIT Press, Cambridge, MA, 1997. ISBN 0-2621-8177-0.
- 43 Mary Beth Rosson and John Millar Carroll. *Usability Engineering*. Morgan Kaufmann, 2002.
- 44 P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. 11(4):443–456, 2005.
- 45 M. Schneider-Hufschmidt, T. Kühme, and U. Malinowski. *Adaptive user interfaces: Principles and Practice*. Amsterdam: North-Holland, 1993.
- 46 Roger N. Shepard. Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Behavior*, 6:156–163, 1967.
- 47 B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *IEEE Symposium on Visual Languages*, pages 336–343, Boulder, CO, 1996.
- 48 Lionel Standing. Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology*, 25(2):207–222, 1973.
- 49 Lucy Suchman. *Plans and Situated Action. The Problem of Human-Machine Communication*. Cambridge University Press, NY, 1987.
- 50 Lucy Suchman. *Human-Machine Reconfigurations*. Cambridge University Press, NY, 2007.
- 51 J. Thomas and K. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Center. IEEE CS Press, 2005.
- 52 E. R. Tufte. *Envisioning information*. Graphics Press, 1990.
- 53 L. A. Tweedie, R. Spence, D. Williams, and R. Bhogal. The attribute explorer. In *Proceedings of the Video Track of the ACM Conference on Human Factors in Computing Systems*, pages 435–436, Boston, MA, April 1994.
- 54 R. van Liere, Mulder J. D., and J. J. van Wijk R. Computational steering. *Future Generation Computer Systems*, 12(5):41–450, 1997.
- 55 B. Z. Veksler, W. D. Gray, and M. J. Schoelles. From 1000 ms to 650 ms: Why interleaving, soft constraints, and milliseconds matter. In *Proceedings of the 8th International Conference on Cognitive Modeling*, Ann Arbor, MN, 2007.
- 56 J. S. Yi, Y. ah Kang, J. T. Stasko, and J. A. Jacko. Toward a deeper understanding of the role of interaction in information visualization. 13(6):1224–1231, November/December 2007.
- 57 J. S. Yi, Y. ah Kang, J. T. Stasko, and J. A. Jacko. Understanding and characterizing insights: How do people gain insights using information visualization? In *BELIV: Proceedings of the 2008 Conference on Beyond Time and Errors*, pages 1–6, New York, NY, 2008. ACM.