The Science of Interaction

Abstract

There is a growing recognition within the visual analytics community that interaction and inquiry are inextricable. It is through the interactive manipulation of a visual interface – the analytic discourse – that knowledge is constructed, tested, refined, and shared. This paper reflects on the interaction challenges raised in the visual analytics research and development agenda and further explores the relationship between interaction and cognition. It identifies recent exemplars of visual analytics research that have made substantive progress toward the goals of a true science of interaction, which must include theories and testable premises about the most appropriate mechanisms for human-information interaction. Six areas for further work are highlighted as those among the highest priorities for the next five years of visual analytics research: ubiquitous, embodied interaction; capturing user intentionality; knowledge-based interfaces; collaboration; principles of design and perception; and interoperability. Ultimately, the goal of a science of interaction is to support the visual analytics community through the recognition and implementation of best practices in the representation of and interaction with visual displays.

1. Introduction

A central precept of visual analytics is that it is through the interaction with a visual interface that human insight is created. As visual analytics is concerned with the relationship between visual displays and human cognition, merely developing novel visual metaphors is rarely sufficient to trigger this insight (where insight may be a new discovery or confirmation of a prior belief). These visual displays must be embedded in an interactive framework that scaffolds the human knowledge construction process with the right tools and methods to support the accumulation of evidence and observations into theories and beliefs.

The “science of interaction” is concerned with the study of methods by which humans create knowledge through the manipulation of an interface. As a science, it involves the development and testing of theories about the most effective means to support inquiry. Interaction, however, is an overloaded term. At one level, interaction typically refers to the set of controls provided to the user to manipulate an interface and the relationship between the user and that interface. At a more abstract level, there is the interaction between the user and the problem space. This higher-level interaction is a cognitive act that is enabled by computational tools, but it does not take place exclusively within them (nor, for that matter, through the use of any single tool). The science of interaction, therefore, is a much broader concept than just the principles for creating interface widgets. Some of one’s interaction with an information space might take place within the context of a software tool, but much of it occurs internally in one’s mind; traces of the process of inquiry will be found throughout all of the tools at one’s disposal.

To change fundamentally the nature of human interaction with information such that discovery is both natural and supported seamlessly by computational aids, new interaction research is needed. Users need to be connected to their data (or to analytic operations that provide insight into that data), not tethered to a device. To support ubiquitous analysis, where insights can be generated and tested wherever the mind is – not wherever the data and the tool happen to be – interaction with information spaces needs to be made available across devices, platforms, locations, use context, and collaborative settings. But rather than focus on point solutions for information analysis tools on individual platforms, effort must be devoted to understanding the relationship between interaction and inquiry such that coherent, consistent analysis capabilities are at the user's disposal wherever and whenever he or she is thinking about a problem space. As interaction research matures, transitions between tools should become more transparent, so we might not even be aware that we have moved from one system to another. Following from the mantra that good design just works, good interactive tools should not draw attention to the novelty of their operation. Instead, they should just seem natural and obvious, bolstering and never confounding human cognitive capacity. As best practices emerge, these practices need to be embodied in tools the community develops so users begin to see familiar interaction models in the new tools they learn. Integrated systems that encapsulate suites of capabilities need to be developed subsequently a consistent user experience is provided throughout the inquiry process.

As a checkpoint toward the development of this science of interaction, this paper identifies progress toward the major recommendations made in *Illuminating the Path: The Research and Development Agenda for Visual Analytics* [1] and articulates a series of research challenges for the future, centered around the problem of changing the dominant interaction paradigms to ones that support human knowledge construction as effectively as possible.

2. Review of Recommendations in *Illuminating the Path*

Chapter 3 of *Illuminating the Path*, "Visual Representations and Interactions", contains three high-level recommendations on research gaps related to the science of interaction.

- **Create a science of visual representations based on cognitive and perceptual principles that can be deployed through engineered, reusable components. Visual representation principles must address all types of data, address scale and information complexity, enable knowledge discovery through information synthesis, and facilitate analytical reasoning.**

  One of the most ambitious aims of the visual analytics R&D agenda is the construction of community standards for effective representations and interactions and the deployment of these standards through interoperable components. While the information visualization community has become more insistent on the recognition and implementation of best practices, only modest progress has been made toward making users' lives easier by enforcing these practices in the development of new tools. (For instance, suggestions have been made that the peer review process account for the degree to which new tools use accepted standards). While other papers in this issue focus on scale and reasoning, this paper examines the interplay between cognition and interaction as well as makes suggestions for how to support reusable, interoperable components.

- **Develop a new suite of visual paradigms that support the analytical reasoning process.**

  Visual metaphors appropriate for the range of data used to accomplish a task (spatial, temporal, textual, quantitative, and so on) must be created. While the visual analytics literature over the past five years is rife with examples of novel visual metaphors, few of these have reached the level of paradigm. Part of this difficulty comes from the lack of a common frame of reference for abstract data (unlike, for instance, geographic data for which maps provide a dominant visual paradigm). With time and sufficient evaluation studies, some of these techniques may become standard. However, treatment of uncertainty and the development of multi-format representations (such as those that might be suitable for displaying features in video, text, audio, and sensor feeds in a coherent display) have seen considerably less investigation. Given that most real-world analysis tasks involve data in multiple formats, this is a significant shortcoming.

- **Develop a new science of interactions that supports the analytical reasoning process. This interaction science must provide a taxonomy of interaction techniques ranging from the low-level interactions to more complex interaction techniques and must address the challenge to scale across different types of display environments and tasks.**

  As described in the next section, the development of interaction taxonomies has been a significant research theme in the past five years. The community's reflection on both interaction with a device or interface and interaction with a task or problem is an important step toward a coherent theory of human-information interaction. In particular, we have seen increased discussion of how visual analytics tools are used in the wild, from mobile devices in the field, to high-resolution workstation displays, to collaborative command and control environments. These real-world use cases provide exemplars from which principles of interface scalability can be abstracted.

3. Interaction as a Reasoning Aid

The interactive manipulation of computational resources is part of the reasoning process. Thus, interaction is always situated in the context of some problem or goal-directed activity. Acknowledging these situations – or better yet, incorporating direct support for them – can improve the ability of interactive interfaces to help humans reason. Interaction should not be an afterthought – a set of controls bolted on to a clever visual display to allow the user to modify it – but the first thing that is considered in the development of an analysis system. The interaction *is* the insight.

In the process of inquiry, users' context helps them identify relevant concepts and link them into appropriate structures. These acts of conceptual manipulation have been described as *situation* [2], the bringing together of background contexts and current observations and analyses toward some goal. Software interfaces enable the enactment that is part of selecting and reasoning with a set of concepts and their associated contextual wrappers. Lemke [3] calls situation an "ecology," evoking the dynamic interaction between concepts and thinkers in the process of knowledge construction. Situated cognition has been shown to be important to both formal and informal learning and discovery [e.g., 4, 5]. Information analysis tools are essentially learning aids; they help their users learn about the patterns, trends, relationships or other features in their data.

Furthermore, the theory of distributed cognition [6, 7] posits that cognition is an emergent property of interactions between people and their environment through perception and action. Rather than viewing cognition as the mental processes occurring only within individual minds, distributed cognition recognizes the vital importance of people's interactions with artifacts. Interaction is the process by which people transform and propagate representations of information, thus facilitating analysis, learning, and knowledge [8].

Given the close coupling between interaction and cognition, the science of interaction must empirically validate theories about cognitive processes on its way toward producing knowledge-construction interfaces. For instance, there is a centuries-long thread of research into the interplay between interacting with information and human judgment. Kant, for example, introduced the notion of purposiveness in reasoning that encompasses action directed toward some end; this purposiveness accounted for the reasoner's context and prior knowledge and was the mechanism by which order was brought to observations [9]. The judgement, or perspective, of the individual inquirer was central to the inquiry. For visual analytics environments, this means that embodying the role, task, and worldview of the user in the analytic environment is the only way for user and environment to become collaborators in the discovery process.

Dialogical inquiry, or the interplay between human and tool (where each poses both questions and answers), is also vital to the future of interaction science. A user might approach a tool with a question in mind, or the tool might show patterns or features that prompt the user to form new questions rather than arrive at immediate answers. The inquiry that users of visual analytics systems engage in is often pragmatic, where through the manipulation of resources and the accumulation of experience in an information space, insights are generated. Pragmatism along the lines of that proposed by C.S. Peirce [10] holds that science is in its practice; inquirers replace faith in the *a priori* with their experiences in the world. In pragmatic inquiry, useful insight only emerges out of the *manipulation* of observations. This manipulative aspect of inquiry is crucial; the more ways a user can "hold" their data (by changing its form this way and that, exploring it from different angles, and via different transformation), the more insight will accumulate. Pragmatism has some similarities to Husserl's phenomenology [11], developed at about the same time. Both ground understanding in experience, but phenomenology adds the aspect of intentionality. Every cognitive action is directed at something; there is always an object of intention.

In dialogical inquiry, the role of the human-computer interface is to support the questioning process. The dialog can be between people, in the case of a collaborative system or between human and computer, or within an individual user and their "future self" – the one who is starting to understand the information space. Except in the case of some collaborative systems, this dialog is rarely made explicit, but it is always there. This dialog should also be playful, in that it is manifest as a free-flowing stream of ideas that are compared, evaluated, and tested using a variety of tools. Rather than a "game-based" interface that implies perhaps too strict a sense of goal-direction, the interface might be more like a "toy" – something that encourages open exploration without imposing the game-maker's rules on the exploratory process.

3.1 The Elements of Interaction

The information visualization community has begun to distinguish between low-level interactions (those between the user and the software interface) and high-level interactions (those between the user and the information space). Given the intentionality, or goal-directedness, implicit in both of these levels, it is useful to enumerate the various aims that a user might have in manipulating an interface.

In lower-level interaction, the user's goal is often to change the representation to uncover patterns, relationships, trends or other features. In an attempt to define the building blocks of a science of interaction, Amar et al. [12] define a set of low-level interaction primitives that generalize across visual analytic methods. These primitives – retrieve value, filter, compute derived value, find extremum, sort, determine range, characterize distribution, find anomalies, cluster, and correlate – that both accommodate specific questions which might be asked of a visualization and can be composed into aggregate functions for more complex questions. Such taxonomies can be used to evaluate the "completeness" of an interface: does it allow users to efficiently perform each of these operations?

The P-Set model [13] offers an approach for capturing a user's sequence of interactive steps in an application-agnostic fashion. Tracking the investigation process allows the user to see their current state in the context of prior exploration and can potentially inform future action. Tools such as Palantir (http://www.palantirtech.com) are now implementing history mechanisms that expose the sequence of interactive steps as a sensemaking aid, and Aruvi [14] integrates history tracking with diagrammatic knowledge capture. Given that the interactive manipulation of an interface is the outward manifestation of the user's reasoning process, working toward an ecosystem of visualization systems that capture these manipulative steps in a common fashion can serve to help users track their knowledge construction processes as they move between tools.

In higher-level interaction, the user's goal is to generate understanding. Here, understanding the intent of the interaction becomes critical. Yi et al. [15] define a taxonomy of interaction intents – select, explore, reconfigure, encode, abstract/elaborate, filter, and connect – that could constitute the components of a knowledge discovery or confirmation process. Just as lower-level interaction capabilities can be used to assess completeness of an interface, these higher-level categories can be used to assess the kinds of goals to which an interface could be applied. While no single visual analytics application might exhaustively support all possible user goals, collections of tools could be assembled that together meet the goals of a particular user task. To accomplish this feat, however, individual tools need to not only be mapped to the intents they afford, but interoperability challenges must be addressed to allow the seamless transfer of user data and findings across the multiple third-party components assembled in support of a task.

In assessing common interaction costs, Lam [16] introduces the "gulf of goal formation" to account for the user's cost of deciding on an intent. This intent is then translated into *execution* steps (for which goals are mapped to the tools and operations offered by a system and to the physical movements required to perform them) and *evaluation* steps (which involve perception, interpretation, and further manipulation of the resulting displays. Principles arising from this work – including the needs to reduce interface complexity, increase the predictability of interaction events, and identify long sequences of repetitive that could be replaced with "shortcuts" – further suggest that community standards are needed. A closer coupling between understanding the reasoning process inherent in the user's manipulation of the interface and the design of that interface can lead to visual analysis systems that better align with their user's goals. A downside of this approach is that it can lead to a profusion of problem-specific interfaces customized for the reasoning processes characteristic of a particular domain and no general principles that can apply across domains.

A common motif in recent information visualization interaction literature is that the inability of a user to fluidly manipulate an interface in furtherance of their goals is always a condition to be avoided. However, in many analysis tasks, goals are unstable, and a straightforward progression down a path of discovery is impossible. A breakdown in analytic discourse is not necessarily bad; in fact, it is often under conditions of breakdown that new discoveries are made. As a hermeneutic concept [17], breakdown occurs when expectations or prior mental models fail to adequately explain observations. Breakdowns might occur when users see something in a display that causes them to revise their goals or refine previously held assumptions about the information space. The fact that an interaction with a tool does not have the anticipated result may not always be indicative of a problem with the tool's design; it may be indicative of an emergent insight on that part of the user.

Figure 1 summarizes the relationship between high- and low-level interaction. The interactive controls provided by the individual display device provide access to a set of low-level representation and interaction techniques that support higher-level intents. Analytic discourse is the relationship between these techniques and the user's goals and tasks, which involve low-level choices about manipulating interactive controls and higher level goals surrounding the problem being investigated. The higher-level tasks are those that are actually part of the inquiry process. It can be useful to think of these higher-level tasks as one or more of the three modes of inquiry – abduction, deduction, and induction – which Peirce [18] linked into a process flow for the construction of knowledge. In abduction, observations derived from exploratory analysis stimulate possible hypotheses through "an appeal to instinct" (What seems to make sense to the user?). Deduction follows, in which the consequences of those hypotheses are examined (If the emergent hypothesis is true, can the question to which it is an answer be reframed to assess the validity of the claims which the hypothesis would imply?). Lastly, inductive hypothesis testing selects the most likely explanation by looking for confirmatory indicators and ranking alternative explanations. Induction is not truth-preserving, as future observations may alter or contradict a hypothesis, but in interaction design, this is an advantage. The user's goals are similarly mutable, and frequently the intent of a visual analysis task is to identify the *best* explanation for an observation, acknowledging that there are potentially multiple explanations and that no analysis tool is likely to provide access to all possible data as well as all possible ways of exploring that data. Visual analysis tools simply give users the means to ask questions and must support the evolution of those questions over time.

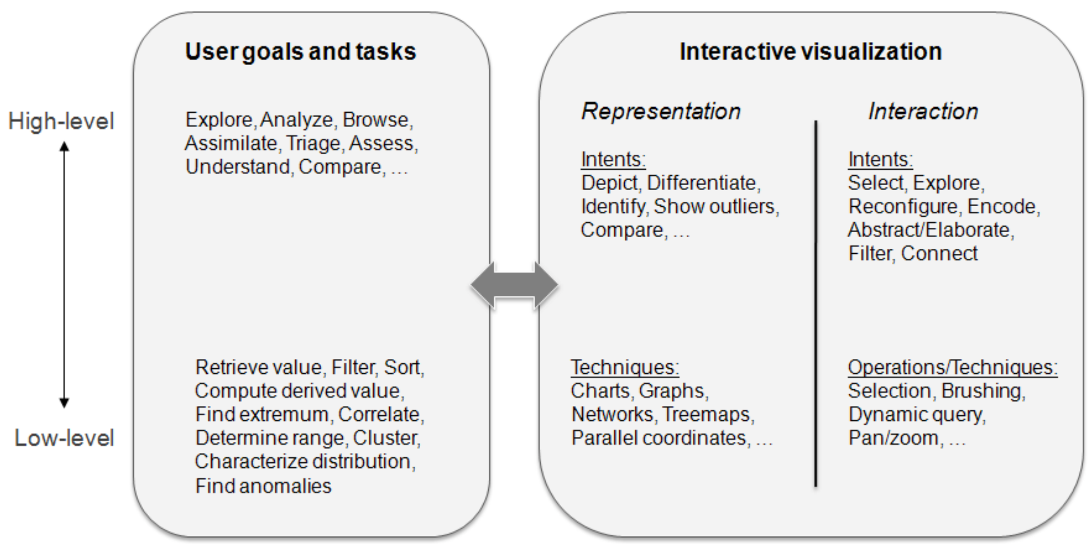

Figure 1. Analytic discourse involves the relationship between user goals and tasks and the affordances of interactive visualization.

Peirce's model is adequate for cases where the user is attempting to generate understanding *ab initio*, although more straightforward confirmatory tasks are also possible. In a confirmatory analysis, the user might skip the abductive step (having performed it already) and engage in a "top-down" deductive assessment. Here, the ability for the user to very quickly structure their information space to identify confirmatory features is crucial, in contrast to the exploratory need to reflect on multiple complementary displays while seeking out structure.

The interaction process can be considered as the recursive engagement with an information space at three levels – familiarization, hypothesis construction, and verification – that align with the three modes of inquiry. In *familiarization* (an abductive process), a user is beginning to understand the problem and form goals as well as identify sources of data if they are not already given. The familiarization process involves identifying gaps in available data, determining what tools and methods can apply to it, transforming data into formats usable by those methods, identifying changes in the data if it has been examined previously, understanding what the customer needs (i.e., in what context the analyst's answers will be used), and articulating existing assumptions. In *hypothesis construction* (which may be both abductive and deductive), the analyst is engaged in exploration of the data space and the formation of mental models to explain observations. In the *verification* phase (typically inductive), alternative explanations are considered and biases are assessed and challenged.

4. Interaction Challenges for the Next Five Years

The science of interaction has ambitious goals. In light of the studies that have been accomplished since the publication of the visual analytics R&D agenda, this section attempts to clarify some specific research needs and enumerate the components of this science that need attention from the research community. These categories of work – ubiquitous, embodied interaction; capturing user intentionality; knowledge-based interfaces; collaboration; principles of design and perception; and interoperability – are essential if visual analytics is to move from a mode of producing single-purpose proofs of concept to universally impactful systems that encompass the best solutions from across the research and development community.

4.1 Ubiquitous, Embodied Interaction

It is important that the research community's focus on better understanding the relationship between inquiry and interaction not lose sight of the fact that analytic interaction is embedded in a user's experience in the world. While software developers might focus on *tools*, users focus on *problems*. These problems live in users' minds and follow them from tool to tool, from place to place, and from one day to the next. Rarely is a user's problem solved in the confines of a single software environment and during just the period of time in which that user is directly interacting with the environment. Since an information worker's life is generally composed not of singular analysis tasks but of continuous engagement with information, which is constantly arriving and interweaving with what is already known, how can we provide these users with coherent interaction experiences across the range of interfaces they might use in the course of their work life? Since the manipulation of ideas that leads to insight may occur anywhere – and even subconsciously – how can we provide access to all relevant information whenever and wherever a user is engaged with a problem?

Ubiquitous computing offers the potential to make data and computational resources accessible anytime and anywhere, but research into ubiquitous interaction is needed to make sure that these resources are provided in a useful, usable state. Crucially, interaction research needs to ensure that transaction costs for performing analytic operations are minimized at every step, while not forcing users into fixed processes that fail to allow for breakdown conditions. For instance, in confirmatory assessments, the user might simply want the answer to a question: "what's the cheapest price to fly from New York to Los Angeles next month?" While we can provide all the data to perform this assessment visually, in certain user contexts a question answering system, rather than an exploratory interface, may be more appropriate. If the price is below a certain threshold, or the user lacks confidence in the answer provided, *then* an exploratory visual interface might be offered to enable further investigation. If an analyst has a spark of insight during his or her morning commute, what interface can be provided to best afford assessment of that insight with whatever computational resources (such as a mobile device) are accessible at that place and time?

Ubiquitous interaction means that rather than build point solutions, which work for one of each possible user context (the field, office, collaborative environment, and so on), we need to be concerned with creating core analytic capabilities that can be transformed for each of these contexts but that remain consistent across them.

Recent research into the extremes of interaction context – the use of very small and very large displays – suggests that there may be some interaction principles that can remain common across displays and in other cases, automated transformations that may need to occur depending on the use context. For instance, studies of mobile device interfaces have determined that fisheye techniques can be more successful than zoomable interfaces because they better preserve a user's orientation in an information space [19]. Examples of applications customized for small devices (such as for emergency responders in [20]; Figure 2) help elucidate design techniques, both for data management and interface construction, that accommodate the affordances of the platform.

At the other end of the spectrum, large high-resolution displays allow focus and context to be managed in fundamentally new ways. In an empirical study of the accuracy and performance times for tasks such as finding attribute values or trends on large displays, Yost and North [21] found that through both physical navigation and human perceptual abilities, users were able to perform some tasks more quickly as interface size grew. In large displays, users preferred embedded visualizations presented in context rather than small multiples. Given the same task but the reality of having to perform it in different contexts, lessons learned about the effectiveness of interfaces at these extremes can be translated into principles for automated presentation.

There has also been research into interface metaphors that remain consistent across platforms. FacetMap [22], for instance, is a faceted browsing technique that can provide different levels of aggregation on differently sized displays. FacetMap visualizations use an identical interface across platforms, although the level of detail that can be rendered in a single view changes. Increased research emphasis needs to be placed on metaphors that provide consistency for users, since the training effect means that new tools, which use familiar metaphors, can be more easily adopted.

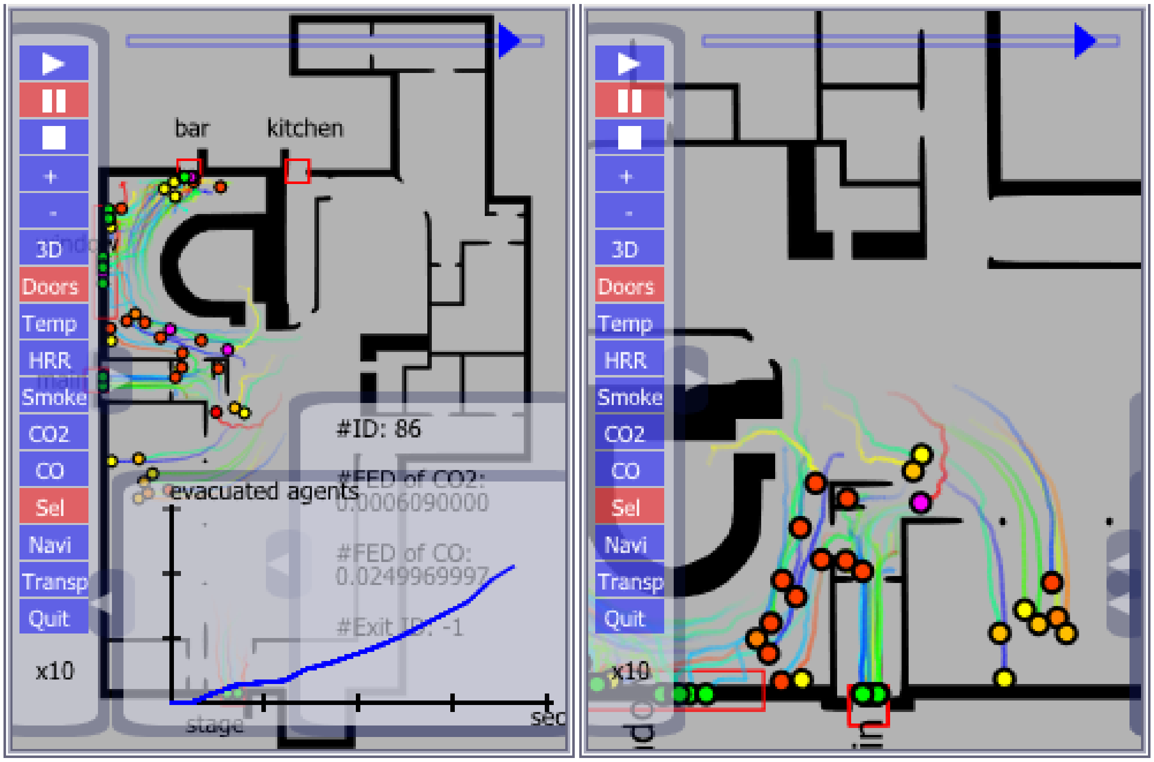

Figure 2: Mobile interface for building evacuation modeling uses techniques such as transparent overlays (left) to preserve context on small screens [20].

For visual analytics to truly be transformative, it is also vital that non-traditional interfaces be developed, not just in the common "off the desktop" realms of mobile and large displays but also in the realm of mixed reality and context-aware computing. The notion of embodied interaction [23] suggests that information artifacts, such as components of a visual display, take on meaning through their use. This use is nearly always embedded in a physical, social, and cultural context that should not be ignored; cognition involves the interplay among a range of distributed artifacts that may be physical, mental, or digital [8]. A law enforcement officer using a visual tool on a mobile device during a field interview is not interacting with a piece of software; he or she is interacting with an incident or a suspect, and the software tool is supporting that interaction. Research into interfaces for augmented reality, taking into account all of the physical and task-related context that surrounds an interaction, can lead to new principles for transparent design whereby information systems automatically recognize context and support their users' information work with just the right set of tools for the task at hand. New ways of combining physical and virtual information as well as real-time and historical data are needed. In mobile environments, can information delivery be tailored to the task at hand? In collaborative environments, can the physical relationships between participants, their actions and discussions, be seamlessly integrated with the information displays they are using?

4.2 Capturing User Intentionality

Necessary to the tailoring of information displays to users in support of embodied interaction is the need to recognize what the user is trying to achieve through the interaction. Some visual analytics systems have added annotation capabilities that allow users to manually record their thought processes as they work. These annotations serve as a textual representation of goals and strategies that are useful for reconstructing the steps one took to reach a finding or, when shared with others, can allow them to better evaluate the finding, but they are not usable by the software tools themselves. While techniques like P-Set can be used to represent the “how” of an analysis process, typically lost is the “why.” Knowing *why* a user is pursuing a particular path is crucial to a visual analytics tool’s ability to modify its presentation, suggest alternatives, or identify additional information for the user.

Representing analytic provenance means we do not just capture the source of data and the transformations that were applied to it, nor only the sequence of interaction steps that occurred, but we develop mappings between application events and cognitive events. Can the temporal relationships between application events (such as what events tend to occur in proximity to each other and whether they occur in rapid succession or in drawn out periods of reflection) be used as indicators of insight?

Recently, increased attention has been devoted to the problem of capturing higher-level thought processes *in situ*. The HARVEST system uses a hierarchical model of events, actions, subtasks, and tasks to categorize a user’s activity [24]. In HARVEST, a visualization state can be saved together with the “action trail” (Figure 3) that constitutes its provenance. These actions are composed of exploration actions (those involved in accessing and exploring data), insight actions (those involved in marking visual objects or taking notes) and meta actions (those that operate over the action trail itself, such as undo and redo). The authors of HARVEST discuss examples of third-party visualization tools being mapped to this provenance schema, suggesting that it may be possible to distill interaction with any visualization into a set of user-activity building blocks.

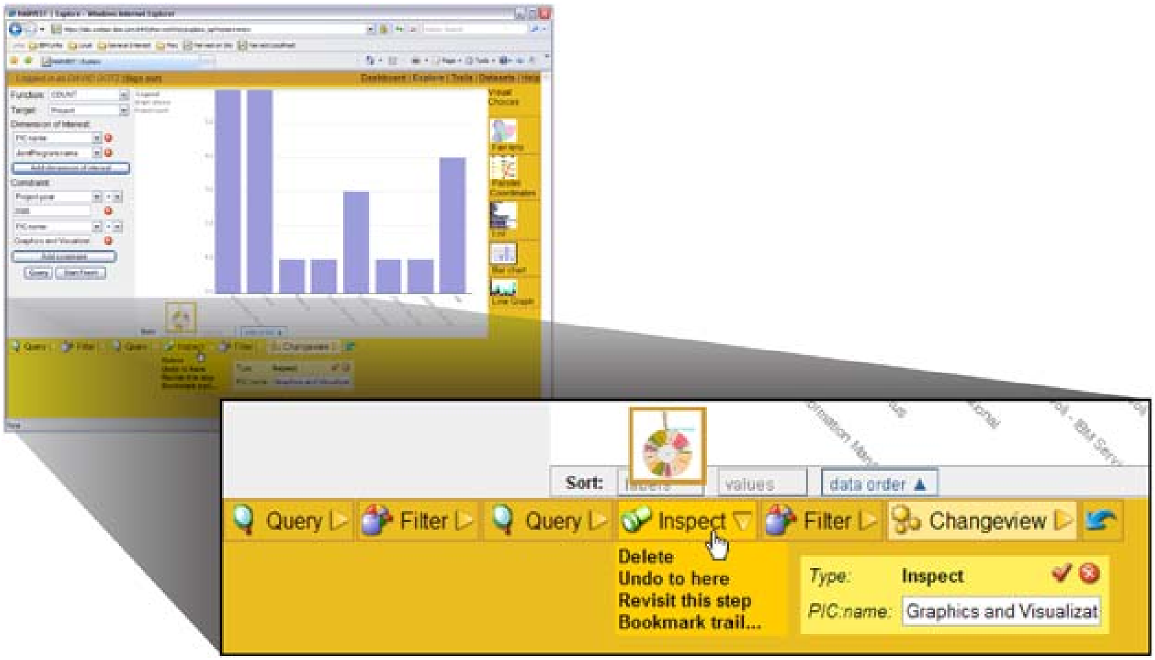

Figure 3: Action trails in HARVEST allows users to preserve their inquiry paths [24].

Dou et al. [25] take an alternative approach, using human coders to explore the logs created during other users' interactions with a financial analysis tool. The results of this coding were compared to "ground truth" collected by think-aloud narration from the original analysts. The more successful of the techniques employed by these human coders in recognizing insight could be used as models for automatically identifying the findings and methods of visual analytics tool users.

Meta visualization tools such as those used in [14, 25, 26] are useful as history-preserving tools in support of an individual user's exploration, but additional work is needed to extend visual analytics systems to be able to represent exactly what the insight is and why it is important. The actual insight is generally expressed either as a textual narration or a bookmark to a view. If, as we discuss in more detail in the next section on knowledge-based interfaces, the insight can be represented in a form that is useful both to the human and the software system, it will be possible to automate knowledge collection during the interaction process and customize the interface to align with a user's particular strategy (for instance, by recognizing interaction strategies that are typically more successful in leading to insight and suggesting interaction paths that follow them).

Additionally, these initial studies in capturing analytic provenance are necessarily developed as proofs of concept within a single visual analysis application. However, since analysts rarely complete all of their work within the confines of a single tool, it will become necessary to abstract these taxonomies into community standards to which developers will map the particular operations supported by their applications. Without such standards, analysts typically resort to capturing their thinking and evidence manually because there is no other way to collect it seamlessly across tools. With a common provenance schema and a growing body of tools that support that schema, the analyst's toolbox will come to contain applications that can feed each other's results; findings from one tool will be passed to another, built upon, and passed on, preserving audit trails throughout the entire analysis process. Analytic provenance also needs to be captured over longer time periods than just a single analysis session; it ought to be possible to capture, and allow users to reflect upon, the long-term learning that occurs as a user grows familiar with an information space. Recognizing these longer-term knowledge construction processes can prompt a system to help the user re-evaluate older findings in light of new knowledge.

4.3 Knowledge-based Interfaces

Given the widely acknowledged intersection between interaction and knowledge construction, the ability of visual analysis tools to represent and reason with human knowledge is underdeveloped. One strategy for representing analytic insight as more than a view or an annotation is to begin incorporating computational representations of human knowledge into visual systems. Frequently, knowledge representation formats like description logics or ontologies are used for information schema mapping or entity-relation search applications, where formal semantics are necessary for machine reasoning. Knowledge representations can also be used to mark up human expressions of insight in machine-readable formats and in a more consistent fashion than narrative text annotations allow.

When human perceptual abilities and machine reasoning combine, new mixed-initiative interfaces become possible. For instance, [27] introduces a technique for turning the features and patterns that a user identifies visually into rules that can be used to automatically recognize additional cases of those features in large data sets. Approaches such as this are not only useful for offloading the burden of search to an automated tool once exemplars have been identified, but it can also be used as formal records of the structure in data that constituted the discovery. Without this formal markup, consumers of an analysis product may have difficulty in understanding exactly what it was in a particular view that caught the analyst's attention.

Techniques that examine the format and structure of the data the user is analyzing, such as the "Show Me" feature in Tableau [28] make use of another form of knowledge representation. These techniques rely on formal models of the relationship between data types and visualization affordances to recommend displays that are likely to result in useful insight.

One component of future knowledge-based interfaces should be user models that account for the role, context, prior knowledge, and aims of the individual or group using a visual analytics tool. Even users with similar backgrounds solving similar tasks will use idiosyncratic reasoning strategies and therefore require tools that accommodate these strategies. In a study of analysts performing a typical exploration and hypothesis construction task, Robinson [29] found that users' information organization strategies ranged from concept maps to timelines to piles. These organizational models reflected their internal cognitive representation of the problem, and the design of future analysis environments should allow users to choose the interface model that best aligns with their perspective.

User models should also extend to formal descriptions of the concepts and relationships the user cares about most. These descriptions can be used to emphasize those concepts in an interactive display, reducing the amount of effort needed to find the items of greatest interest in a complex information space. The insight-detection techniques described earlier are one mechanism for inductively generating these models; visual tools can bootstrap representations of a user's interests and characteristic strategies over time. Such user models can address the need for "steerability" in mixed-initiative interfaces identified by Lam [16]. While automated search and discovery systems will be vital for helping users deal with ever-growing information spaces, only if these interfaces accommodate users' viewpoints will they be widely adopted.

4.4 Collaboration

Collaboration is characteristic of nearly all visual analytics work. Rarely does the entire analysis process, from data acquisition through to reporting and decision-making, take place with a single person. While infrastructure for collaboration – emerging technologies for shared interactive displays – is a healthy research theme outside of visual analytics, requiring greater attention is the need for collaborative systems to help map between user models. Cooperative knowledge construction and communication requires this mapping to facilitate efficient and appropriate re-use of knowledge resources as well as sound decision making by consumers of another's analysis.

In proposing a framework for multi-analyst work, Brennan [30] uses the notion of "private perspectives" to describe the user models we discuss above. Within each perspective, facts and confidence can be formally represented as logical rules, allowing private perspectives to be fused into a shared view. However, the technique assumes that analysts will use shared voice channels to communicate their reasoning behind the perspective, suggesting that the technique is most appropriate for synchronous collaboration within small groups. Support for this sort of explicit collaboration must be matched by systems for implicit collaboration, where a sufficiently robust representation of a user's reasoning is preserved for later consideration by others.

Heer and Agrawala [31] enumerate many of the requirements that must be managed in designing collaborative visual analysis systems. Extant systems address many of these challenges – such as indicating reference (to what is one user attempting to draw others' attention) and managing sideband discussions for clarification and decision making (e.g., [32]). However, preserving sufficiently rich descriptions of an analyst's activities in asynchronous group work such that others can effectively ask questions of that analyst in the absence of their physical or virtual presence is important.

Furthermore, when designing collaborative visual analytical systems, it is important to note that interaction techniques developed for single-user systems do not always have the same effects in a collaborative system. For example, Isenberg and Fisher present a collaborative system that enables multiple users to perform co-located document analysis tasks [33] using the interaction technique *brushing and linking*. This technique has been used extensively in (single-user) visualizations, especially in systems that utilize multiple coordinated views [15]. However, brushing and linking presents new benefits as well as challenges when applied to a collaborative environment. For example, while the technique allows users to maintain common ground and awareness, it can also blur the boundary between individual and group work. This example suggests that reexamining existing interaction techniques and developing new ones specifically for collaborative visual analytical systems will be important in advancing collaborative interaction as a science.

4.5 Principles of Design and Perception

Despite the growth of the visual analytics community and the development of successful technologies in the past five years, the community has not seen interdisciplinary participation to the extent necessary to make more significant progress on the challenges of analytic interaction. The design and cognitive science communities, in particular, need to be more deeply engaged in visual analytics research. It is the responsibility of the visual analytics enterprise to form substantive collaboration with these communities, bringing experts from those fields into our research teams.

An important first step in involving the design community in visual analytics occurred during the kickoff meeting for the Canadian Network of Visual Analytics Centres, at which design panels critiqued existing analysis tools from a user-centered perspective. The culture of substantive design critique is not yet part of the visual analytics enterprise, yet such critiques are vital if research-grade systems (which most of the products of the visual analytics community in the past five years represent) are to be transitioned into operational use. In many evaluation studies of analytic interfaces, the design of the interface is often being evaluated more than the underlying analytic algorithm, even if the intent of the study is otherwise. Recent research in identifying appropriate visual metaphors for particular cognitive tasks (e.g., [34]) is a step in the right direction.

When design practitioners are involved in visual analytics tool development and evaluation from the start, good design practices and aesthetics in visualization design will begin to permeate the community. Just as joint research funding programs are beginning to support better cooperation between visualization and data sciences research communities, joint programs that involve design activities (often funded and performed under humanities programs) must be started.

Likewise, although cognitive science has long been identified as a pillar of visual analytics, there has to date been relatively modest involvement of the cognition and perception communities in visual analytics research programs. During analysis, information is constantly represented in new ways: information elements gain and lose prominence; give birth to new information; or disprove and thereby eliminate other information elements. Understanding the intersection of cognition and the dynamic nature of information is integral to understanding interaction. However, the limits of human cognitive abilities have largely gone unexplored. There is evidence of biological changes to brains due to interaction with technology, but brains are not evolving as fast as information is increasing. While preventing cognitive overload is frequently raised as an aim of visual analytics tools, the conditions that constitute overload in exploratory tasks are not well understood.

There is also a biological dimension to cognition and perception. The science of interaction requires understanding the constraints imposed by the biology of the human eye, and information must be presented in a way that accommodates physical limitations. Users may even end up with eye fatigue or strain because the presentation has pulled their eyes constantly to the periphery of the display when they are trying to work with data points in the middle.

4.6 Interoperability for Integrated Interaction

The advances that the visual analytics research community has made in the past five years have largely been embodied in point solutions – individual tools or methods that demonstrate a new algorithm, a novel visual metaphor, or a new set of design principles. What the community must work toward in the next five years are mechanisms to turn these singular advances into components of integrated suites that support the end-to-end process of analysis. New platforms upon which individual solutions will reside are needed. In many user communities, deploying new tools is difficult politically, technically, and culturally. These problems can be mitigated in part by a recognition that new methods must often fit within existing workflows; demonstrating how a tool or technique integrates with the intended user's existing activities and goals as well as with the information systems he or she already uses is crucial for adoption.

Interoperability is vital to the science of interaction because analysis occurs in a workflow. Each component in that workflow will be a party in the analytic discourse, so each must acknowledge and respond to the contributions of other components. How can information best be passed among tools, and how can each tool build upon the discoveries made in others? It will not be a wise use of effort for research teams to implement complete analysis packages that, simply for the sake of completeness, replicate functionality available elsewhere. Instead, focus should be placed on creating the analytic substrate to which new capabilities will connect. This way, "gold standard" implementations can be made accessible to all members of the community, and valuable research funding can be devoted to novel development rather than redundant implementations.

Conceiving of visual analytics techniques as components in a larger, interoperable ecosystem can also lead to new kinds of composable interfaces that make analytic discourse more flexible than it can possibly be within the bounds of a single tool. Systems that can be re-wired by the analyst to meet changing goals – or ideally, that re-wire themselves – allow the diversity of an analyst's work to take place within an integrated environment. If it captures the community's best practices and allows new advances to be rapidly plugged in, such an environment has the potential to change the nature of information work.

Early examples of such environments (e.g., [35]) have explored the development of service-based analytic systems, where atomic components for data preparation, transformation, and display can be linked into mashups. Service-based analysis allows interfaces for interactive discourse to be constructed in a platform- and place-agnostic fashion. However, community standards for how to move meaning, not just data, between components are needed. Such standards will allow each component in a workflow to describe the knowledge structures that emerge from it.

4.7 Evaluating the Costs and Benefits of Interaction

Evaluating visual analytical systems and tools has been an active and important aspect of advancing the science of visual analytics. In the past five years, the community has made great strides in quantitatively and qualitatively measuring the benefits of visual analytical tools as well as discovering new methodologies for performing such evaluations in varying environments (see the paper "Technology Transfer Progress" in this special issue). However, in this great body of work, there has been limited effort in measuring the costs and benefits of interaction specifically. With few exceptions (e.g., [16]), interaction has not been studied in isolation to determine how it facilitates reasoning and knowledge building. As a science, interaction needs to be understood through the use of scientific methods such that its role is better defined and its effects more predictable.

Since the goal of interaction is to build knowledge, generate insight, and perform analysis, the effects of interaction should be measured accordingly [36, 37, 38]. While the evaluation methods proposed so far in the visual analytics community have been focused on understanding visual analytics as a whole, there is increasing awareness that interaction plays a key role to the success of visual analytical systems and therefore needs to be studied independently. For example, Green et al. [38] proposed a model for visual analytics based on human cognition and considered interaction to be responsible for engaging the user and keeping the user in a continuous, uninterrupted "cognitive flow". Under this model, the costs and benefits of interaction can be measured in a more quantitative manner.

The challenge for evaluating interaction in visual analytics starts with understanding the relationship between interaction and visual representation. In recent years, research in visualization and visual analytics has mostly focused on discovering new visual representations while using interaction as a supporting tool [15]. While it is clear that visual representations can be informative without interactions (e.g., in the form of static bar graphs or a pie charts), and interaction cannot function alone without visual representations, exactly what and how much a user can benefit from having the ability to interact with visual representations is still undetermined. In the case of complex problem solving that involves the use of sophisticated visual analytical system, this question is even more difficult to answer but also much more important.

Once the role of interaction is better defined, successful evaluation of interaction requires knowing what to measure. Insight, knowledge, and cognitive flow are some potential candidates, but they are also difficult to assess [39]. On the other hand, measuring task performance speed and accuracy are quantitative and reliable, but the results do not always adequately reflect the benefits of interaction (i.e., knowledge building or insight generation). In order to have a consistent basis of comparison, the community can benefit from identifying new metrics that are shared, acknowledged, and used by all researchers in this domain.

Finally, with new metrics, new evaluation paradigms might be necessary. For evaluating visual analytics systems, new approaches have been proposed to overcome the challenges of performing formal evaluations in real-world settings (e.g., [40, 41]). When evaluating interaction, similar challenges would likely occur such that the use of existing evaluation methods might not be appropriate. Identifying and inventing appropriate methods that support testing interaction independently from visual representation and utilize new evaluation metrics will be an important aspect of verifying and validating interaction and interaction techniques in visual analytic systems.

5. Conclusions

Since its inception, the field of visual analytics has emphasized the centrality of interaction with visual environments to the knowledge construction process. Indeed, it is now widely recognized that the interaction *is* the inquiry. Interaction is not just the manipulation of interface controls in a software environment but the discourse the user has with his or her information, prior knowledge, colleagues, and environment.

Through work in six key research areas – ubiquitous, embodied interaction; capturing user intentionality; knowledge-based interfaces; principles of design and perception; collaboration; and interoperability – the science of interaction can be advanced into a body of theory and practice that guides how humans will engage with information spaces in the future. Understanding the relationship between actions performed during use of an analysis tool and modes of inquiry can lead to systems that are able to recognize, reflect, and support the generation of insight by their users.

Acknowledgements

This work has been supported by the National Visualization and Analytics Center (NVAC) located at the Pacific Northwest National Laboratory in Richland, WA. NVAC is sponsored by the U.S. Department of Homeland Security Science and Technology Division. The Pacific Northwest National Laboratory is managed for the U.S. Department of Energy by Battelle Memorial Institute under Contract DE-AC06-76RL01830.

References

1 Thomas J and Cook K, eds. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Press: Los Alamitos, CA, 2005; 200pp.

2 Solomon K, Medin D, and Lynch E. Concepts do more than categories*.* *Cognitive Science* 1999; **3**(3): 99-104.

3 Lemke J. Cognition, context, and learning: A social semiotic perspective, In: Kirshner D and Whitson J (Eds). *Situated Cognition: Social, Semiotic, and Psychological Perspectives*. Erlbaum: Mahwah, NJ. 1997. 37-55.

4 Clancey W. Situated cognition: How representations are created and given meaning, In: Lewis R and Mendelsohn P (Eds). *Lessons from Learning*. North-Holland: Amsterdam. 1994. 231-242.

5 Lave J and Wenger E, *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press: New York, 1991; 138pp.

6 Hutchins E, *Cognition in the wild*. MIT Press: Cambridge, MA, 1995pp.

7 Kirsch D. Distributed cognition: A methodological note*. Pragmatics and cognition* 2006; **14**(2): 249-262.

8 Liu Z, Nersessian N, and Stasko J. Distribution cognition as a theoretical framework for information visualization*. IEEE Transactions on Visualization and Computer Graphics* 2008; **14**(6): 1173-1180.

9 Kant I, *Critique of Judgment*. Hackett: Indianapolis, 1987; 576pp.

10 Peirce C. What pragmatism is*. The Monist* 1905; **15**(2): 161-181.

11 Husserl E, *Logical Investigations*. Routledge: London, 1970; 877pp.

12 Amar R, Eagan J, and Stasko J, Low-level components of analytic activity in information visualization, in *2005 IEEE Symposium on Information Visualization*. 2005. p. 111-117.

13 Jankun-Kelly T, Ma K, and Gertz M. A model and framework for visualization exploration*. IEEE Transactions on Visualization and Computer Graphics* 2007; **13**(2): 357-369.

14 Shrinivasan Y and van Wijk J, Supporting the analytical reasoning process in information visualization, in *CHI 2008, Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*. 2008. p. 1237-1246.

15 Yi J, Kang Y, Stasko J, and Jacko J. Toward a deeper understanding of the role of interaction in information visualization*. IEEE Transactions on Visualization and Computer Graphics* 2007; **13**(6): 1224-1231.

16 Lam H. A framework of interaction costs in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 2008; **14**(6): 1149-1156.

17 Gadamer H-G, *Truth and Method*. Continuum: New York, 1975; 551pp.

18 Peirce C. Fixation of belief. *Popular Science Monthly* 1877; **12**(November): 1-15.

19 Buering T, Gerken J, and Reiterer H. User interaction with scatterplots on small screens - A comparative evaluation of geometric-semantic zoom and fisheye distortion. *IEEE Transactions on Visualization and Computer Graphics* 2006; **12**(5): 829-836.

20 Kim S, Jang Y, Mellema A, Ebert D, and Collins T. Visual analytics on mobile devices for emergency response. *2007 IEEE Symposium on Visual Analytics Science and Technology* 2007 (Sacramento, CA); 35-42.

21 Yost B and North C. The perceptual scalability of visualization. *IEEE Transactions on Visualization and Computer Graphics* 2006; **12**(5): 837-844.

22 Smith G, Czerwinkski M, Meyers B, Robbins D, Robertson G, and Tan D. FacetMap: A scalable search and browse visualization. *IEEE Transactions on Visualization and Computer Graphics* 2006; **12**(5): 797-804.

23 Dourish P, *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press: Cambridge, MA, 2001pp.

24 Gotz D and Zhou M. Characterizing users' visual analytic activity for insight provenance. *Information Visualization* 2009; **8**: 42-55.

25 Dou W, Jeong C, Hyun D, Stukes F, Ribarsky W, Lipford H, and Chang R. Recovering reasoning processes from user interactions. *Computer Graphics and Applications* 2009; **29**(3): 52-61.

26 Heer J, Mackinlay J, Stolte C, and Agrawala M. Graphical histories for visualization: Supporting analysis, communication, and evaluation. *IEEE Transactions on Visualization and Computer Graphics* 2008; **14**(6): 1189-1196.

27 Xiao L, Gerth J, and Hanrahan P. Enhancing visual analysis of network traffic using a knowledge representation. *2006 IEEE Symposium on Visual Analytics Science and Technology* 2006 (Baltimore, MD); 107-114.

28 Mackinlay J, Hanrahan P, and Stolte C. Show me: Automatic presentation for visual analysis. *IEEE Transactions on Visualization and Computer Graphics* 2007; **13**(6): 1137-1144.

29 Robinson A, Collaborative synthesis of visual analytic results, in *2008 IEEE Symposium on Visual Analytics Science and Technology*. 2008: Columbus, OH. p. 19-24.

30 Brennan S, Mueller K, Zelinsky G, Ramakrishnan I, Warren D, and Kaufman A. Toward a multi-analyst, collaborative framework for visual analytics. *IEEE Symposium on Visual Analytics Science and Technology 2006* 2006 (Baltimore, MD), IEEE; 129-136.

31 Heer J and Agrawala M. Design considerations for collaborative visual analytics. *INformation Visualization* 2008; **7**: 49-62.

32 Tomaszewski B and MacEachren A. A distributed spatiotemporal cognition approach to visualization in support of coordinated group activity. *Third International ISCRAM Conference* 2006 (Newark, NJ)

33 Isenberg P and Fisher D. Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections. *Computer Graphics Forum* (Proceedings of EuroVis), 2009. To appear.

34 Ziemkiewicz C and Kosara R. The shaping of information by visual metaphors*. IEEE Transactions on Visualization and Computer Graphics* 2008; **14**(6): 1269-1276.

35 Pike W, Bruce J, Baddeley B, Best D, Franklin L, May R, Rice D, Riensche R, and Younkin k. The Scalable Reasoning System: Lightweight visualization for distributed analytics*. Information Visualization* 2009; **8**: 71-84.

36 Plaisant C, Fekete J, and Grinstein G. Promoting Insight-Based Evaluation of Visualizations: From Contest to Benchmark Repository. *IEEE Transactions on Visualization and Computer Graphics* 14, 1 (Jan. 2008), 120-134.

37 Amar R and Stasko J. A Knowledge Task-Based Framework for Design and Evaluation of Information Visualizations. *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on*, pp.143-150.

38 Green T.M., Ribarsky W, and Fisher B. Visual analytics for complex concepts using a human cognition model. *Visual Analytics Science and Technology, 2008. VAST '08. IEEE Symposium on*, pp.91-98.

39 Chang R, Ziemkiewicz C, Green T.M., and Ribarsky W. Defining Insight for Visual Analytics. *IEEE Computer Graphics and Applications*, vol. 29, no. 2, pp. 14-17, March/April, 2009.

40 Plaisant C. The challenge of information visualization evaluation. *Proceedings of the Working Conference on Advanced Visual interfaces*. AVI '04, 109-116.

41 Isenberg P, Zuk T, Collins C, and Carpendale S. Grounded evaluation of information visualizations. In *Proceedings of the 2008 Conference on Beyond Time and Errors: Novel Evaluation Methods For information Visualization*. BELIV '08, 1-8.