

The science of interaction

William A. Pike^{a,*}
John Stasko^b
Remco Chang^c and
Theresa A. O'Connell^d

^aPacific Northwest National Laboratory,
PO Box 999, MSIN K7-28, Richland,
WA 99352, USA.

^bGeorgia Institute of Technology, Atlanta,
GA 30334, USA.

^cUniversity of North Carolina at Charlotte,
Charlotte, NC 28223, USA.

^dNational Institute of Standards and Technology,
Gaithersburg, MD 20899, USA.

*Corresponding author.
E-mail: william.pike@pnl.gov

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 article 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. Seven areas for further work are highlighted as those among the highest priorities for the next 5 years of visual analytics research: ubiquitous, embodied interaction; capturing user intentionality; knowledge-based interfaces; collaboration; principles of design and perception; interoperability; and interaction evaluation. Ultimately, the goal of a science of interaction is to support the visual analytics and human–computer interaction communities through the recognition and implementation of best practices in the representation and manipulation of visual displays.

Information Visualization (2009) **8**, 263–274. doi:10.1057/ivs.2009.22

Keywords: visual analytics; interaction theory; reasoning; collaboration

Introduction

A central precept of visual analytics is that the development of human insight is aided by interaction with a visual interface. 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 or negation 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

This article is a product of a workshop on the Future of Visual Analytics, held in Washington, DC on 4 March, 2009. Workshop attendees included representatives from the visual analytics research community across government, industry and academia. The goal of the workshop, and the resulting articles, was to reflect on the first 5 years of the visual analytics enterprise and propose research challenges for the next 5 years. The article incorporates input from workshop attendees as well as from its authors.

Received: 27 May 2009
Revised: 7 July 2009
Accepted: 7 July 2009

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. Emerging best practices need to be embodied in the 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 in such a way that a consistent user experience is provided throughout the inquiry process.

As a checkpoint toward the development of this science of interaction, this article identifies progress in interaction research since the publishing of *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.

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 inquiry.

In the process of inquiry, users' contexts help them identify relevant concepts and link them into appropriate structures. The inquiry that users of visual analytics systems engage in is often pragmatic, in that useful insight only emerges out of the experience of manipulating information.² This manipulative aspect of inquiry is crucial; the more ways a user can 'hold' their data (by changing their form or exploring them from different angles and via different transformations), the more insight will accumulate. These acts of conceptual manipulation have been described as *situation*,³ the bringing together of background contexts and current observations and analyses toward some goal. Situated cognition has been shown to be important to both formal and informal discovery, (for example, Clancey⁴ and Lave and Wenger⁵), and software interfaces can enable the enactment that is part of situation in that they help users select and reason with a set of concepts and contexts.

Lemke⁶ calls situation an 'ecology', evoking the dynamic interaction between concepts and thinkers in the process of knowledge construction. This dynamism is also a central premise of the theory of distributed cognition,^{7,8} which 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. In human-computer interaction research, distributed cognition has been used to describe the process by which people transform and propagate representations of information, thus facilitating analysis, learning and knowledge.⁹

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 goal-directedness or intentionality accounted for the reasoner's context and prior knowledge and was the mechanism by which order was brought to observations.¹⁰ The judgment, or perspective, of the individual inquirer is 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. 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.

A common motif in recent information visualization and 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,¹¹ breakdown occurs when expectations or prior mental models fail to adequately explain observations, often forcing the user to fundamentally reconsider their information space or their tools. 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 a situation that will lead to an emergent insight on that part of the user by challenging prior assumptions.

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 analysis tasks that are typically performed with visualizations. These primitive tasks – retrieve value, filter, compute derived value, find extremum, sort, determine range, characterize distribution, find anomalies, cluster and correlate – accommodate specific questions that might be asked of a visualization and can be composed into aggregate functions for more complex questions. The P-Set model¹³ offers an approach

for capturing a user's sequence of low-level 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 Aruvi¹⁴ integrates history tracking with diagrammatic knowledge capture.

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 (does it allow users to efficiently and effectively perform each low-level operation?), these higher-level categories can be used to assess the kinds of goals to which an interface could be applied. Although 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.

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 discourse is formed around the feedback between user goals and the result of analyzing a particular representation or changing it through interaction. Representation techniques – the development of which is typically the province of information visualization – are low-level artifacts that support a user's intent to represent information in such a way as to make a particular determination. Acknowledging the distinction between techniques and intents is critical to the development of analysis support tools; the technique should never be considered an end in itself, but a means to support the user's information understanding.

Despite the development of initial taxonomies of interaction and representation components, an outstanding research challenge remains properly understanding the relationship between these components and the modes of inquiry – abduction, deduction and induction. In this order, these three modes of inquiry form a process flow for the construction of knowledge.¹⁶

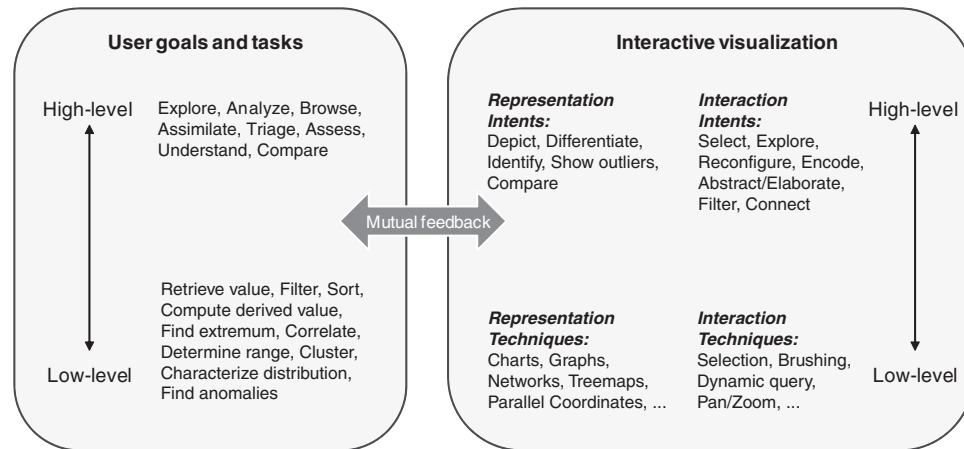


Figure 1: Analytic discourse involves the mutual feedback between user goals and tasks and the affordances of interactive visualization.

In abduction, observations derived from exploratory analysis stimulate possible hypotheses through an ‘appeal to instinct’. (What seems to make sense to the user?) The process of becoming familiar with a new data space can be thought of as an abductive process: the 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 reviewing the data, identifying gaps in it, 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 (that is, in what context the analyst’s answers will be used), and articulating existing assumptions. In constructing hypotheses or ‘abductions’, the analyst is engaged in exploration of the data space and the formation of mental models to explain observations.

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?) Deduction may result in refuting a previously formed hypothesis, as unlike abduction it is necessarily truth-preserving. In a confirmatory analysis, the user might skip the abductive step (instead starting with a set of alternative hypotheses as given) and engage in a ‘top-down’ deductive assessment. Here, the ability for the user to very quickly structure his or her information space to identify confirmatory features is crucial, in contrast to the exploratory need to reflect on multiple complementary displays while seeking out structure.

Lastly, inductive hypothesis testing selects the most likely explanation by looking for additional indicators and ranking alternative explanations. The verification phase of an analysis is typically inductive, where contra-indicators that might refute a hypothesis are sought,

alternative explanations are considered and biases are assessed, challenged and resolved. 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 and their answers over time.

Although the philosophy of science provides convenient models for understanding the process of inquiry, there are practical costs to implementing such models in interactive analysis tools. In assessing common interaction costs, Lam¹⁷ describes a ‘gulf of goal formation’ to account for the user’s cost of deciding on an intent. This intent is 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 actions that could be replaced with ‘shortcuts’ – 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.



Interaction Challenges for the Next 5 Years

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 interaction that need attention from the research community for a true ‘science of interaction’ to emerge. These categories of work – ubiquitous, embodied interaction; capturing user intentionality; knowledge-based interfaces; collaboration; principles of design and perception; interoperability; and interaction evaluation – 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.

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, interweaving with what is already known, and departing once outdated, 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?’ Although 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, while in other cases, automated transformations 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.¹⁸ Examples of applications customized for small devices (such as for emergency responders in Figure 2¹⁹) help elucidate design techniques, both for data management and interface construction, that accommodate the affordances of the platform. In Figure 2, techniques such as making user interface components transparent to avoid occluding other features and menus that can be moved or hidden help make information analysis on small screens possible.

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 accuracy and performance times for tasks such as finding attribute values or trends on large displays, Yost and North²⁰ 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,²¹ 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, as the training effect means that new tools, which use familiar metaphors, can be more easily adopted.

For visual analytics to be truly transformative, it is also vital that non-traditional interfaces be developed, not

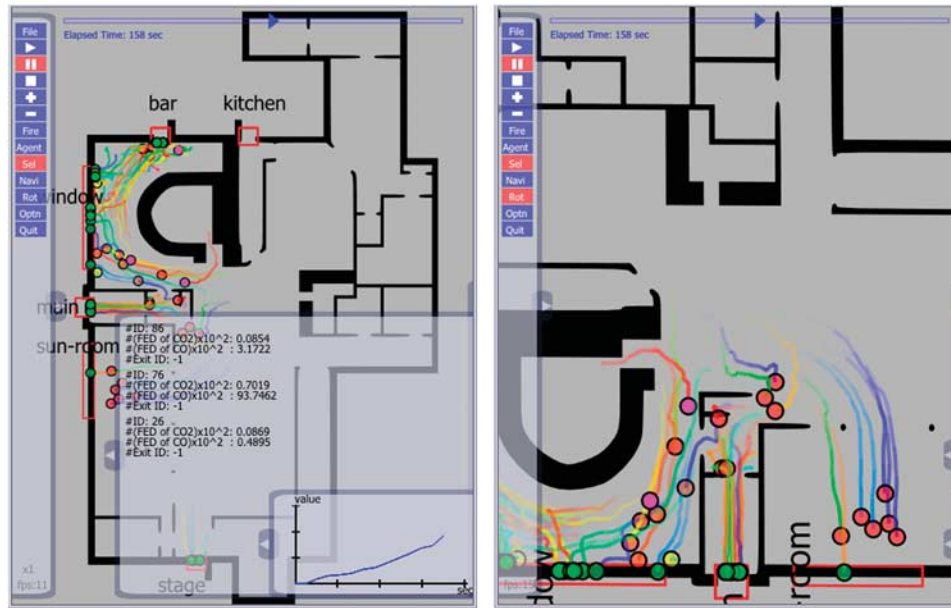


Figure 2: Mobile interface for building evacuation modeling uses techniques such as transparent overlays (left) to preserve context on small screens.¹⁹

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 interaction²² 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.⁹ 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. The SiteLens system, for instance, provides a prototype for overlaying visual depictions of historical sensor data atop live camera images on a handheld device, supporting geo-contextual exploration of relevant data.²³ In mobile environments, can information delivery be tailored not just to location, but to the task at hand? In collaborative environments, can the relationships between participants, their actions and discussions, be seamlessly integrated with the information displays they are using?

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



Figure 3: Action trails in HARVEST allow users to preserve their inquiry paths.²⁴

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.

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 Shrinivasan and van Wijk¹⁴, Dou *et al.*²⁵ and Heer *et al.*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 may 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, as 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. (An example of infrastructure to support provenance capture in visualization is VisTrails, which uses a workflow model to capture, share and reuse exploratory visualization processes.²⁷ 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 use 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.

Knowledge-based interfaces

Despite 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, in which 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, Xiao *et al.*²⁸ 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 can be also 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,²⁹ 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. These models should also address work styles, user preferences and mental models of their tasks, workplace constraints, and even personality styles in collaborative settings. 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, Robinson³⁰ found that users' information organization strategies ranged from concept maps to timelines to piles. These organizational models reflected users' 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.¹⁷ 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.

Collaboration

Collaboration is characteristic of nearly all visual analytics work. Rarely does a single individual undertake the entire analysis process, from data acquisition through to reporting and decision making. Infrastructure for collaboration – emerging technologies for shared interactive displays – is a healthy research theme outside of visual analytics. However, greater research effort must be applied to developing collaborative systems that help map between user models and across synchronous and asynchronous collaboration modes. 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 *et al.*³¹ 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 Agrawala³² 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 (for example, Tomaszewski and MacEachren³³). 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 analytics systems, 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 tasks³⁴ using the brushing and linking interaction technique. 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 analytics systems will be important in advancing collaborative interaction as a science.

Principles of design and perception

Despite the growth of the visual analytics community and the development of successful technologies in the past 5 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, cognitive science and usability engineering 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. Recent efforts in the Visual Analytics Science and Technology (VAST) Symposium Challenge are beginning to use metrics-based evaluation to assess whether and to what degree visual analytics tools will empower users to be efficient, effective and satisfied.

An important 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 5 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 (for example, Ziemkiewicz and Kosara³⁵) is a step in the right direction.

When design practitioners and usability engineers 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. Design and evaluation roles must be kept separate, however, allowing user-centered evaluation to be run by independent usability engineers who are specifically trained in metrics-based evaluations, human factors and psychology.

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 because of 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 physiological 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. However, the lack of substantial involvement by the cognitive and perceptual science communities in visual analytics has hindered the development of principles for sound interfaces not just for visualization, but for analysis. Addressing this lack of research integration will be a primary challenge for the visual analytics community over the next 5 years. Funding programs that expressly support joint research are one way to achieve this, as are efforts by visual analytics researchers to include psychologists on their research teams.

Interoperability for integrated interaction

The advances that the visual analytics research community has made in the past 5 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 5 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 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. While intellectual property issues can complicate sharing analytic methods widely, funding agencies may wish to sponsor workshops that attempt to define interoperability specifications across the visual analytics research community. Such specifications will help ensure that analysis components adopted by those agencies will integrate well with existing platforms, and creating analysis components compatible with a community-wide substrate need not require the sharing of intellectual property.

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. These composable interfaces can be informed by the dataflow programming model, where operations in one analytic process trigger downstream changes in others. While visual programming languages and composable coordinated visualization environments are common in the research community, they are often most useful for developers or very skilled users, rather than for typical end users. Composable interfaces that can take multiple forms to reflect users' mental models of their tasks may meet with more success. If they capture the community's best practices and allow new advances to be rapidly plugged in, such environments have the potential to change the nature of information work.

Early examples of such environments (for example, Pike *et al.*³⁶) 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.

Evaluating the costs and benefits of interaction

Evaluation of visual analytics systems has been an active and important component of advancing the science of visual analytics. The paper on 'Visual Analytics Technology Transfer Progress'³⁷ in this special issue highlights the strides that have been made – and challenges that remain – in quantitatively and qualitatively measuring the benefits of visual analytics. There remains a particularly acute need to develop techniques for measuring the costs and benefits of interaction during the knowledge construction process. With few exceptions (for example, Lam¹⁷), interaction has not been isolated as an experimental variable. Usability studies that use the same visualization but vary the interaction techniques and methods can illuminate the specific benefits and costs of those interactions. Such studies can help create a science of interaction by contributing to agreed-upon methods and falsifiable hypotheses, making the role of interaction better defined and its effects more predictable.

As the goal of interaction is to build knowledge, generate insight and perform analysis, the effects of interaction should be measured according to their support for these goals.^{38–40} Although the evaluation methods proposed so far in the visual analytics community have focused on assessing the benefits of visual analytics as a whole, the key role played by interaction needs to be studied independently. While it is clear that visual representations can be informative without interactions (for example, in the form of static information graphics), and interaction cannot function alone without visual representations, exactly what kind and degree of benefit is realized by allowing a user to interact with visual representations is still undetermined. For example, Green *et al.*⁴⁰ 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'. What quantitative measures can be developed to assess the degree to which interaction enables the state of flow?

Assessing the usability of interactive visual interfaces means not just measuring user preference, but understanding what users need to do their jobs effectively. The International Standards Organization defines usability as having three parts – efficiency, effectiveness and satisfaction.⁴¹ Although measures for efficiency and



satisfaction are straightforward, effectiveness is complex. How did the interface enable the user to perform his or her job better? What interaction techniques improved the soundness of the resulting analysis? Quantifications of insight, knowledge and cognitive flow are candidate measures, but they are also difficult to assess.⁴² To enable better comparison between studies, standardized metrics for effectiveness are needed. And with new metrics, new evaluation paradigms might be necessary. New approaches have been proposed to overcome the challenges of performing formal evaluations in real-world settings (for example, Plaisant⁴³, Isenberg *et al.*⁴⁴); enhancing these approaches with methods for testing interaction independently from visual representation is crucial to verifying and validating interaction techniques in visual analytic systems.

Conclusion

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 seven key research areas – ubiquitous, embodied interaction; capturing user intentionality; knowledge-based interfaces; principles of design and perception; collaboration; interoperability; and interaction evaluation – 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. This body of work will inform the design of interactive, visually rich environments such that system developers can be more confident that their design choices will work, and users will find an interaction experience that is both consistent and supportive of sound analysis processes. 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.

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 US Department of Homeland Security Science and Technology Division. The Pacific Northwest National Laboratory is managed for the US Department of Energy by Battelle Memorial Institute under Contract DE-AC05-76RL01830. Certain commercial equipment, instruments, materials, services or companies are identified in this article. This in no way implies endorsement or recommendation by the National Institute of Standards and Technology.

References

- 1 Thomas, J.J. and Cook, K.A (eds.) (2005) *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. Los Alamitos, CA: IEEE Computer Society Press.
- 2 Peirce, C. (1905) What pragmatism is. *The Monist* 15(2): 161–181.
- 3 Solomon, K., Medin, D. and Lynch, E. (1999) Concepts do more than categories. *Cognitive Science* 3(3): 99–104.
- 4 Clancey, W. (1994) Situated cognition: How representations are created and given meaning. In: R. Lewis and P. Mendelsohn (eds.), *Lessons from Learning*. Amsterdam, The Netherlands: North-Holland, pp. 231–242.
- 5 Lave, J. and Wenger, E. (1991) *Situated Learning: Legitimate Peripheral Participation*. New York: Cambridge University Press.
- 6 Lemke, J. (1997) Cognition, context, and learning: A social semiotic perspective. In: D. Kirshner, and J. Whitson (eds.), *Situated Cognition: Social, Semiotic, and Psychological Perspectives*. Mahwah, NJ: Erlbaum, pp. 37–55.
- 7 Hutchins, E. (1995) *Cognition in the Wild*. Cambridge, MA: Massachusetts Institute of Technology Press.
- 8 Kirsch, D. (2006) Distributed cognition: A methodological note. *Pragmatics and Cognition* 14(2): 249–262.
- 9 Liu, Z., Nersessian, N. and Stasko, J. (2008) Distribution cognition as a theoretical framework for information visualization. *IEEE Transactions on Visualization and Computer Graphics* 14(6): 1173–1180.
- 10 Kant, I. and Pluhar, W. S. (1987) *Critique of Judgment*. Indianapolis, IN: Hackett Publishing.
- 11 Gadamer, H.-G. (1975) *Truth and Method*. New York: Continuum.
- 12 Amar, R., Eagan, J. and Stasko, J. (2005) Low-level components of analytic activity in information visualization. In: J. Stasko, and M. Ward (eds.), 2005 IEEE Symposium on Information Visualization. Minneapolis, MN. Piscataway, NJ: IEEE Computer Society Press, pp. 111–117.
- 13 Jankun-Kelly, T., Ma, K. and Gertz, M. (2007) A model and framework for visualization exploration. *IEEE Transactions on Visualization and Computer Graphics* 13(2): 357–369.
- 14 Shrinivasan, Y. and van Wijk, J. (2008) Supporting the analytical reasoning process in information visualization. In: M. Czerwinski and A. Lund (eds.) CHI 2008, Proceeding of the 26th Annual SIGCHI Conference on Human Factors in Computing Systems; 5–10 April. Florence, Italy. NY: Association for Computing Machinery Press, pp. 1237–1246.
- 15 Yi, J., Kang, Y., Stasko, J. and Jacko, J. (2007) Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 13(6): 1224–1231.
- 16 Peirce, C. (1877) Fixation of belief. *Popular Science Monthly* 12(November): 1–15.
- 17 Lam, H. (2008) A framework of interaction costs in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 14(6): 1149–1156.
- 18 Buering, T., Gerken, J. and Reiterer, H. (2006) User interaction with scatterplots on small screens – A comparative evaluation of geometric-semantic zoom and fisheye distortion. *IEEE Transactions on Visualization and Computer Graphics* 12(5): 829–836.
- 19 Kim, S., Jang, Y., Mellema, A., Ebert, D. and Collins, T. (2007) Visual analytics on mobile devices for emergency response. In: W. Ribarsky and J. Dill (eds.) IEEE Symposium on Visual Analytics Science and Technology: VAST '07; 30 October–1 November. Sacramento, CA. Los Alamitos: IEEE Computer Society Press, pp. 35–42.
- 20 Yost, B. and North, C. (2006) The perceptual scalability of visualization. *IEEE Transactions on Visualization and Computer Graphics* 12(5): 837–844.
- 21 Smith, G. *et al.* (2006) FacetMap: A scalable search and browse visualization. *IEEE Transactions on Visualization and Computer Graphics* 12(5): 797–804.
- 22 Dourish P. (2001) *Where the Action Is: The Foundations of Embodied Interaction*. Cambridge, MA: Massachusetts Institute of Technology Press.

- 23 White, S. and Feiner, S. (2009) SiteLens: Situated visualization techniques for urban site visits. In: D.R. Olson *et al.* (ed.), Proceedings of the Twentyseventh Annual SIBCHI Conference on Human Factors in Computing Systems. Boston, MA. NY: Association for Computing Machinery Press, pp. 1117–1120.
- 24 Gotz, D. and Zhou, M. (2009) Characterizing users' visual analytic activity for insight provenance. *Information Visualization* 8(1): 42–55.
- 25 Dou, W. *et al.* (2009) Recovering reasoning processes from user interactions. *Computer Graphics and Applications* 29(3): 52–61.
- 26 Heer, J., Mackinlay, J., Stolte, C. and Agrawala, M. (2008) Graphical histories for visualization: Supporting analysis, communication, and evaluation. *IEEE Transactions on Visualization and Computer Graphics* 14(6): 1189–1196.
- 27 Silva, C., Freire, J. and Callahan, S. (2007) Provenance for visualizations: Reproducibility and beyond. *Computing in Science and Engineering* 9(5): 82–89.
- 28 Xiao, L., Gerth, J. and Hanrahan, P. (2006) Enhancing visual analysis of network traffic using a knowledge representation. In: P. C. Wong and D. Keim (eds.) 2006 IEEE Symposium on Visual Analytics Science and Technology: VAST '06; 31 October–2 November. Baltimore, MD, NY: IEEE Computer Society Press, pp. 107–114.
- 29 Mackinlay, J., Hanrahan, P. and Stolte, C. (2007) Show me: Automatic presentation for visual analysis. *IEEE Transactions on Visualization and Computer Graphics* 13(6): 1137–1144.
- 30 Robinson, A. (2008) Collaborative synthesis of visual analytic results. In: D. Ebert and T. Ertl (eds.) Symposium on Visual Analytics Science and Technology: VAST '08; 21–23 October. Columbus, OH, Los Alamitos, CA: IEEE Computer Society Press, pp. 19–24.
- 31 Brennan, S. *et al.* (2006) Toward a multi-analyst, collaborative framework for visual analytics. In: P.C. Wong, and D. Keim (eds.), IEEE Symposium on Visual Analytics Science and Technology: VAST '06; 31 October–2 November, Baltimore, MD. NY: IEEE Computer Society Press, pp. 129–136.
- 32 Heer, J. and Agrawala, M. (2008) Design considerations for collaborative visual analytics. *Information Visualization* 7(1): 49–62.
- 33 Tomaszewski, B. and MacEachren, A. (2006) A distributed spatiotemporal cognition approach to visualization in support of coordinated group activity. In: J. Landgren and S. Jul, (eds.) Third International Conference on Information Systems for Crisis Response and Management (ISCRAM 2006); 10–13 May, Gothenburg, Sweden. Gothenburg: University of Gothenburg. <http://www.iscram.org/ISCRAM2009/papers/>, accessed 13 August 2009.
- 34 Isenberg, P. and Fisher, D. (2009) Collaborative brushing and linking for co-located visual analysis of document collections. *Computer Graphics Forum* 28(3): 1031–1038.
- 35 Ziemkiewicz, C. and Kosara, R. (2008) The shaping of information by visual metaphors. *IEEE Transactions on Visualization and Computer Graphics* 14(6): 1269–1276.
- 36 Pike, W. *et al.* (2009) The Scalable reasoning system: Lightweight visualization for distributed analytics. *Information Visualization* 8(1): 71–84.
- 37 Scholtz, J., Cook, K., Whiting, M., Lemon, D. and Greenblatt, H. (2009) Visual analytics technology transition progress. *Information Visualization* 8(4): 294–301.
- 38 Plaisant, C., Fekete, J. and Grinstein, G. (2008) Promoting insight-based evaluation of visualizations: From contest to benchmark repository. *IEEE Transactions on Visualization and Computer Graphics* 14(1): 120–134.
- 39 Amar, R. and Stasko, J. (2004) A knowledge task-based framework from design and evaluation of information visualizations. In: M. Ward, and T. Munzner (eds.) IEEE Symposium on Information Visualization (InfoVis 2004); 10–12 October. Austin, TX. Piscataway, NJ, pp. 143–150.
- 40 Green, T., Ribarsky, W. and Fisher, B. (2008) Visual analytics for complex concepts using a human cognition model. In: D. Ebert and T. Ertl (eds.) IEEE Symposium on Visual Analytics Science and Technology: VAST '08; 21–23 October. Columbus, OH, Los Alamitos, CA: IEEE Computer Society Press, pp. 91–98.
- 41 International Standards Organization (ISO) (1998) *Ergonomic requirements for office work with visual display terminals (VDTs)*. In: *Part 11: Guidance on Usability*. Geneva, Switzerland: ISO. ISO Standard Number: ISO 9241-11:1998.
- 42 Chang, R., Ziemkiewicz, C., Green, T. and Ribarsky, W. (2009) Defining insight for visual analytics. *Computer Graphics and Applications* 29(2): 14–17.
- 43 Plaisant, C. (2004) The challenge of information visualization evaluation. In: M.F. Costabile (ed.), Proceedings of the Working Conference on Advanced Visual Interfaces (AVI 2004); 25–28 May, Gallipoli, Italy. NY: Association for Computing Machinery Press, pp. 109–116.
- 44 Isenberg, P., Zuk, T., Collins, C. and Carpendale, S. (2008) Grounded evaluation of information visualizations. In: E. Bertini, A. Perer, C. Plaisant, and G. Santucci (eds.), Proceedings of the 2008 Conference on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization (BELIV 2008); 5 April, Florence, Italy. NY: Association for Computing Machinery Press, pp. 56–63.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.