

Science of Analytical Reasoning

Keywords: visual analytics, visualization, interaction, reasoning, cognition, sensemaking

Abstract

There has been progress in the science of analytical reasoning and in meeting the recommendations for future research that were laid out when the field of visual analytics was established. Researchers have also developed a group of visual analytics tools and methods that embody visual analytics principles and attack important and challenging real world problems. However, these efforts are only the beginning and much work remains to be done. This paper examines the state of the art in visual analytics methods and reasoning and gives examples of current tools and capabilities. It demonstrates that the science of visual analytics needs interdisciplinarity, indicates some of the disciplines that should be involved, and presents an approach to how they might work together. Finally, the paper describes some gaps, opportunities, and future directions in developing new theories and models, enacting them in methods and design principles, and applying the resulting tools to significant and complex practical problems and data.

Introduction

The science of analytical reasoning is an essential part of visual analytics (VA), as discussed in the introduction to this special issue and in *Illuminating the Path*.²⁷ To attack complex, exploratory, insight discovery, knowledge-building problems, visual analytics requires a foundation in the science and theory of analytical reasoning. Little in terms of controlled observations and practical results exists for higher level reasoning processes in perceptually rich environments, not even in the field of cognitive science. Thus it has been hard to build such a science. Visual analytics offers the promise of providing the practical basis where theories can be tested and (cognitive) analytics must be applied in order to have any hope of successfully solving very important, but quite challenging, real world problems. Although analytical reasoning is an essential part of VA, there are other essential components (computation, interactive visual representations, and analytic methods). All these must eventually be brought together to provide an overarching science of visual analytics. Approaches to doing some of this are discussed in other papers in the special issue.

Initial Recommendations and Where We Are

Chapter 2 of *Illuminating the Path* dealt with the science of analytical reasoning and laid out 8 recommendations for future research. We won't have space in this article to address all these recommendations individually. Instead we will present and discuss a couple of combined recommendations that cover the main points of the full set. (The reader is referred to Chapter 2 of *Illuminating the Path* for a discussion of the complete set of recommendations.)

Recommendation

Build upon theoretical foundations of reasoning, sense-making, cognition, and perception to create visually-enabled tools to support collaborative analytic reasoning about complex and dynamic problems.

Over a period of many years, there has been a large amount of research in reasoning, cognition, and perception that can be applied to the visual analytics framework and to the complex and dynamic problems that it must address. While the work has clear implications for visual analytics, it is essential that it must also be focused on areas that are critical for the design and evaluation of systems that aid human cognitive processing. The relative lack of this work so far in the VA community may be due to the lack of communication between visualization researchers and cognitive scientists whose methods could be used to shed light on analytic cognition in visually complex environments. Much of this work might be adapted from the existing literature; however progress in this area could be considerably improved by a higher level of communication. In addition, the focus of visual analytics is to make the tools visually-enabled, coupling visualizations and interactions with the human visual/understanding channel for maximum throughput integrated with human understanding and judgment. Interaction is the mechanism of coupling and thus interaction should be considered from the standpoint of coupling human reasoning and analytic processes with computer-based processes. Collaboration should be considered in two senses: collaboration between the human and the computer and collaboration among humans to provide enhanced perspective, more effective generation of new ideas and hypotheses, and diversity of viewpoints. Scalability in problem-solving must be supported in both types of collaboration.

Recommendation

Conduct research to address the challenges and seize the opportunities posed by the scale of the analytic problem.

The issues of scale are manifested in many ways, including the complexity and urgency of the analytical task, the massive volume of diverse and dynamic data involved in the analysis, and challenges of collaborating among groups of people involved in analysis, prevention, and response efforts.

A main realization among leaders in the visual analytics community is that the most challenging applications are not just large scale in terms of, say, the data that must be analyzed or even in terms of the dimensionality of the data, they are also complex in terms of the analysis task and this task involves reasoning, inference building, hypothesis creation and testing, and decision-making. The Sensemaking model of Pirolli and Card²¹ was the most comprehensive existing model addressing these issues from the standpoint of investigative analysis. It shows that the investigation must be iterative with data foraging, evidence building, hypothesis creation and testing, and decision-making overlapping one another. This is a general operational model that is relevant for all types of investigation.

However, it doesn't say much about what is inside each of the steps or how they should be connected.

Understanding these steps in terms of reasoning and argument building and connecting them via interactive interfaces is a main province of visual analytics. Beyond this, dynamic data, changing conditions, and new insights require an exploratory, adjustable approach to problem-solving. Such an approach is really not explicitly addressed in models such as Sensemaking. There are many aspects to dynamic problem-solving. Dynamic data require their own structure and analytical approach. In this regard, analyzing and organizing data in terms of temporal events is an important and, it seems, general approach, as tools such as EventRiver¹⁶, STAB¹, and GeoTime⁶ show. Beyond this, general models that explicitly consider dynamic behavior and temporal sequencing must be developed and integrated into visual analytics approaches. An illustration of this requirement is the need to incorporate flow models and explicit temporal sequence in routing for large scale emergencies¹⁴. Efficient evacuation of a very large building or evacuation routing in a city requires not just one shortest path route but a series of routes that can be used in alternating fashion so that no route becomes gridlocked. These must be available in a situation-aware context that is dynamic in order to be usable.

Relevant Visual Analytics Tools and Results

Several tools have been developed that address aspects of analytical reasoning. There are also relevant results involving evaluations of the tools. Some main examples are given next.

Financial Visual Analytics. The WireVis system² was created in cooperation with financial analysts at the Bank of America (BoA) to address fundamental problems they had in understanding their vast flow of financial transactions. Its initial purpose was to seek and discover fraud, especially wire transfer fraud. This is a very difficult investigative analysis problem in that tools must be exploratory, since new modes of deception are tried all the time, and must help human analysts see odd patterns over time in financial transactions. The need to involve highly trained human analysts to find meaning and discover new modes of deception means that these analyses are expensive and the human part of the system doesn't scale very well. WireVis directly addresses the challenge of scaling limited human resources by permitting the analyst to explore tens of thousands of transactions or more over extended time periods while still being able to dive in and look in detail at any transaction or group of transactions.

WireVis was established after extensive discussions with bank analysts and is built around the idea of *keywords*. (The keywords are a set of highly proprietary words developed by the analysts, through long investigative experience, for flagging transactions that may be of interest. The keywords can contain names of countries or cities, types of businesses, types of transactions, etc. and set the analytic context for this specialized analytic task.) WireVis has a 4 window interface where each window focuses on a key aspect of the transactional analysis (frequency of keywords in a transactional cluster, temporal trends of clusters by keyword, "search by example" for keyword and temporal patterns, and keyword relation within transactions). Each window has multiple interactions, and an interaction in one window produces an update in another, which shows relationships across views.

A fast, hierarchical clustering approach is applied to the transaction, and the time window shows daily activity for all clusters over a 13 month window. One can choose certain keyword patterns for any time period and do reclustering on the fly. Thus, even for repositories of BoA's size, the analyst can start from an overview and with a few clicks get to clusters where activity of individual accounts are discriminated, all the while keeping track of longer range temporal patterns and emerging relationships. The analysts have never had any of these capabilities

before. Although WireVis was initially focused on wire fraud, its capabilities are general and appropriately extended versions are now being considered for general financial analysis including risk and customer analysis.

An evaluation of high level reasoning processes was then undertaken with WireVis as the central application⁵. The analyst's operations were divided into strategies (overall approaches to solving a problem), methods (specific steps taken), and findings (conclusions drawn by the analyst after investigating a suspicious activity). In this study, expert analysts were subjected to multiple observational tools including video, "talk-aloud" recording, post-analysis reports, and interaction logging. However, a group of non-expert evaluators, using only the interaction logs represented in a visualization to make it easier to organize among strategies, methods, and findings, were able to recover 60% of the expert's strategies, 60% of their methods, and 79% of their findings. This is an important result because it has been notoriously difficult to infer user intent and especially higher level reasoning from interaction logs alone. The result also shows that considerable aspects of the reasoning process can be recovered without the elaborate and invasive processes of video recording, talk-aloud, and post-investigation reporting. There are many cases where it is hard to impose these observational processes or where they just haven't been imposed. In addition, there is evidence that imposing talk-aloud or other invasive activities during an analyst's investigations can affect, often negatively, reasoning processes such as discovery and spontaneous insight (a-ha moments)³. Finally, talk-aloud collection will be incomplete (analysts tend to *stop* talking when they become deeply immersed) as will post-investigation reporting (analysts forget detailed steps), so interaction log analysis provides a valuable addition permitting more complete understanding of the analyst's reasoning processes. In this study, it was also found that strategies of a few of the analysts were not clearly represented in the visual interface; if they had been (which would have been straightforward to do without loss of generality of WireVis), the strategies recovery would have been significantly higher. These results can be used to quickly and effectively train apprentices in best practices of analysis, to more fully understand the analyst's craft, and to uncover specific parts of the interactive interface that need improvement.

The above study raises the question of why the results from evaluation of interaction logs alone, even when undertaken by non-experts, were as good as they were. It appears that a significant part of the answer lies with the careful preparation of the visualization tools for the reasoning tasks at hand. However, the WireVis example also

shows that it is possible to do this while still keeping a strong, general aspect to the interface (though the keyword list may have to be swapped out or expanded). The WireVis outcomes are important because they suggest a generalizable approach to analytical reasoning problems. One should not just deal with cognitive tasks for a specific problem but should always attempt to generalize to whole classes of problems. In this case, a key outcome was the development of knowledge- and strategy-laden keyword representations that were intimately coupled to other carefully chosen representations of the data. The result was an approach that could be generalized to other types of financial analytics. Further the methods permitted deep analytical evaluation. These generalized approaches, their outcomes, and evaluations provide necessary building material for the science of analytical reasoning.

Tools for Investigation and Reasoning. Other visual analytics tools have been developed that support reasoning and analytical processes. Jigsaw was developed to help investigators deal with large collections of documents, in particular to support investigative analysts in sensemaking where they must understand multiple entities and their relationships within the documents^{24, 23}. Jigsaw acts like a visual index into the documents highlighting connections between entities through multiple views, such as list, graph, and timeline-based representations. (See Figure 1.) There is also the capability to provide different views of the document text itself. Analysts can then follow a trail of entity connections in order to more fully understand the context of events and their detailed workings. Thus Jigsaw helps the analyst link seemingly unconnected events together to make more complete and coherent stories across the document collection. Jigsaw has been shared with multiple investigators and agencies, who have given feedback about its capabilities. There is interest in using it in some of these settings.

The Scalable Reasoning System (SRS)^{18, 19} is focused more on structured argumentation and provides a web-based diagrammatic reasoning interface that allows the users to record the structure of their arguments and hypotheses. Data brought in from other applications and from the Web are represented as “sticky notes”; the contents of these notes can be visualized using data clustering, timeline, and map-based displays. Using the idea of “reasoning artifacts” from structured argumentation, SRS tags these notes with reasoning roles, such as evidence, assumption, or hypothesis. Reasoning artifacts can be turned into edges that describe relationships between notes, and users can record confidence assessments for each artifact and the attached note. This provides an explicit knowledge structure for the argument. SRS uses Dempster-Shafer belief theory to compute likelihood scores for each artifact, allowing

users to see quickly the sum of likelihood and uncertainty for each component of their reasoning structures. By tracking the development of their hypotheses graphically, SRS can help users reflect on knowledge that might otherwise be kept tacit.

There are also tools that tackle the mixed-initiative aspects of the investigative process. Mixed-initiative systems are those where the human and computer work together, though sometimes independently, in intimate partnership, each doing what it does best. A mixed-initiative approach is necessary because the most challenging (and often most important) visual analytics problems require the insertion of the human ability to attach meaning or to create or extend hypotheses, yet the data are too large, the dimensions too high, the ramifications of a change or decision too many, for a human to handle unaided. These last areas are the province of the computer. In the Sensemaking approach, the investigative process is divided into two main, overlapping loops, the foraging loop and the sensemaking loop, the latter entailing higher order hypothesis creation, comparison, and evaluation. RESIN has been developed to support the first loop, except that it is extended to a foraging/analysis loop to make explicit the integrated analysis that occurs in visual analytics. RESIN provides an automated framework for the reasoning and analytical process, concentrating in particular on data selection and choosing the proper VA tools for analysis²⁸. (See Figure 2.) It relies on an AI blackboard –based⁴ software agent that employs interactive visualization, mixed-initiative problem solving, and time-series analysis to support predictive analysis from a particular viewpoint. (The blackboard architecture was developed to handle complex, ill-defined problems. In this case, a hierarchical system is developed with several agent-based knowledge sources for specific problem aspects. The blackboard is iteratively updated by the knowledge agent when its internal specifications match the current blackboard state.) The RESIN system enables analysts to explore large amounts of data in order to generate, track, and validate multiple hypotheses in an uncertain environment. This general approach has been applied to some specific problems. For example, in terrorism event analysis, RESIN uses information from the START Center’s Global Terrorism Database¹⁵ to match an event with a likely group or groups and to determine the threat level in the region in the near future based on past behavior of the group. RESIN is now being extended to also analyze socio-cultural factors that could be significant in predicting behavior trends.

To attack the sensemaking loop, STAB has been developed. STAB is based on the human proclivity to make sense of complex information or processes by constructing stories that causally connect events, ascribe intentions to actors, and make predictions about the world. Humans also use stories to communicate goals, hypotheses, explanations, and conclusions to one another. However, humans also have cognitive limitations and biases in constructing stories, including biases in collecting, interpreting, and using information. The aim of the STAB project is to develop an interactive approach that uses the structures and processes of story construction to support and guide information visualizations for complex problems^{1,9}. The STAB system provides an interactive story editor for entering specific stories as well as generic story plots. STAB uses its library of story plots to interpret input streams of events, generate multiple stories that provide causal explanations for sequences of events, and calculate confidence values for the generated stories. Current work on STAB investigates how the entities occurring in a story may focus information visualizations of input data streams. Also under investigation is how predictions made by a generated story may focus data foraging, analyses, and visualizations. This current work connects STAB closely with RESIN and activity is underway to combine the two to provide an overall mixed-initiative approach to Sensemaking. The combined system could be embedded within RESIN, which already has a structure for considering end-to-end reasoning and decision-making processes.

Human Cognitive Model. In order to organize the above and other work in a meaningful way and use it to establish design principles for visual analytics tool development, it is necessary to create a human cognitive model. Such a model is also necessary to treat the mixed initiative aspects of the visual analytics process. However, even in cognitive science, higher order reasoning processes associated with attaching meaning to results, developing and evaluating models or hypotheses, and making decisions, are in a sense black boxes, and there is no general, practical model. Visual analytics can give a fresh perspective to this problem and researchers have begun to develop a prototype human cognitive model (HCM)¹⁰. The model uses the cognitive science literature to clearly state and differentiate the strengths of the human and computer components of the mixed initiative system. Thus we have a basis for determining when a human should intercede (e.g., to attach meaning at a key point in the analysis and reasoning process) and what the computer should do to best prepare for this moment. A significant part of the computer's role is to keep track of pertinent information, including analysis processes and outcomes. We see one approach to keeping track of analyses and outcomes in the RESIN/STAB work above. These systems, plus tools

such as SRS, can also keep track of competing hypotheses and the evidence for each, helping to remove the natural human bias to prefer one hypothesis over another, even when, in many cases, substantial counter-evidence is available. In addition, the computer can use its vastly superior “working memory” to keep information available and accessible by highly interactive, exploratory visualization tools. Finally, of course, the computer has an ever more powerful ability to carry out computations and to do computation-based analysis. All this provides the opportunity for the computer to augment human discovery by computer-aided discovery. One approach to do this is through observation of what interests the human; the computer then suggests information that is semantically related, but not yet considered, or does computations to create new relevant information. The human is then free to explore or to ignore these suggestions.

The HCM leads to design principles for visual analytics tools¹⁰. For example, semantically rich and complex reasoning applications often require multiple windows, each focusing efficiently on one aspect of the problem. This need is established in WireVis, Jigsaw, SRS, and a number of other visual analytics tools. In order to minimize the human attentional overhead of dealing with multiple windows, each window must be efficiently focused on a task and obey the principal of “balanced interaction”. Balanced interaction goes beyond brushing and linking from information visualization in that it requires that an interaction in one window not only cause an update in another window but also be available, in a general way, in all windows. With balanced interaction, reasoning and interaction become merged for the user, and the attentional and cognitive switches required for handling different visualizations or different focuses among the windows tend to subside. This aspect is part of a general consideration of cognitive flow within the HCM. One goal of intuitive, exploratory visualization should be that the visualization should not hamper the rhythm of reasoning until the human chooses to refocus resources elsewhere. This sense of being “in the zone” allows the human collaborator to reason without encountering unnecessary attentional or cognitive impediments. In cases where task complexity exceeds the user’s ability to process information, or a cognitive impasse is reached for some other reason, the computer can provide a scaffolding of support by presenting the information within relevant context, suggesting what may have been overlooked, and keeping relevant information present. These considerations result in a series of design choices in terms of interaction. In particular, interaction should be, as much as possible, direct and intimate. Direct interaction insures that one deals directly with the artifacts of the reasoning/analysis process rather than with indirect representations that require cognitive shifts (e.g.,

Boolean inputs or pull-down menus where selections or queries must be entered). Intimate interaction insures that the interaction is so translucent to the human that it appears natural and obvious, thus maintaining the intimate collaboration between human and computer. A key method for maintaining direct, intimate interaction is through “search-by-example” where one merely indicates the pattern or relationship one is interested in (rather than, for example, having to construct an elaborate Boolean query). Prominent instances of search-by-example are selecting a keyword distribution or transaction pattern over time from an account of interest and then finding accounts with similar patterns (as in WireVis), searching for images similar to (or dissimilar from) a selected group of images, searching for video patterns similar to a selected one, or various text body search techniques. To fully achieve search-by-example requires a full-fledged visual analytics approach with a true marriage of visualization and analysis. Imposing the design principles described here does not just lead to improved exploration and knowledge-building for problem-solving, it may also lead to discoveries and insights that one wouldn’t find otherwise. There is evidence that spontaneous insights (a-ha moments) can actually be suppressed if a tool is not flexible enough to permit the user’s mind to roam freely³.

This treatment of interaction in terms of cognition and cognitive flow offers a broadened perspective on the development of an interaction theory discussed in another paper in this special issue²⁰. The interaction theory needs to take into account these issues of cognitive flow, higher level reasoning, and human cognitive modeling. The principles described here and others that may follow from further theoretical development should be part of the theory’s outcomes. In order to develop a more complete, working theory, much remains to be done with the HCM, as discussed in the next section.

Establishing a Science, Future Directions, and Opportunities

If we are to build a science of analytical reasoning, we must think about what this entails. A common definition of science is, “The intellectual and practical activity encompassing the systematic study of the structure and behavior of the physical and natural world through observation and experiment”. This definition extends to psychology and the cognitive sciences when we remember that mental processes are embodied in physical beings and must be studied that way. The science must have observations and undertake experiments that can be reproduced independently and that work to constrain or disprove theories. A usual way to build the science is to lay down basic, general principle

that define its scope and withstand testing and then build theories, models, or hypotheses that are themselves subject to observational and experimental confirmation. To pursue these issues, we must apply the scientific method, defined as “A method of procedure consisting in systematic observation, measurement, and experiment, and the formulation, testing, and modification of hypotheses”. As researchers in analytical reasoning build up a body of work that addresses all these aspects, we will have a science.

We must also identify who should be involved in this new science. Certainly visualization scientists, cognitive scientists, and psychologists should be and already are involved. Since analytics and reasoning can have a social component, other social scientists should be involved as well. Since this new science should be part of the broader science of visual analytics, developers of analysis methods (statistical, database, AI, etc.) may also take part. Though this intrinsic interdisciplinarity requires substantial effort, we should also keep in mind its benefits. As Bordon has pointed out, “Most advances in science come when a person for one reason or another is forced to change fields”.

Critical to setting forth an interdisciplinary science are the ways in which disparities in the research questions, methods, data, and arguments can be reconciled between fields. One approach, taken by cognitive science, attempts to bridge component disciplines of psychology, artificial intelligence, philosophy, neuroscience, linguistics, and anthropology using the concept of a “trading zone”^{26, 8} formed in part by methods that cross-cut the various disciplines. Thagard points out that computational models of cognition (e.g. SOAR, ACT-R) serve to “draw out the unforeseen empirical consequences of cognitive theories and display their limitations”. One can argue a similar role could be played for computational models of cognitive systems in visual analytics. One can further argue that a trading zone should be set up for the partners in developing the science of analytical reasoning.

A parallel approach to interdisciplinarity, Translational Science³⁰ avoids the “pure” versus “applied” science distinction, focusing instead on building application-driven basic research paradigms²⁵. The idea of a scientific discipline that spans exploration of underlying phenomena and complexities of real-world practice originated in the health sciences. Their need for effective “evidence-based medicine” required them to more closely coordinate clinical studies and research in underlying physiological, biochemical, and biophysical phenomena. For a translational science of VA, we emphasize the reciprocal flow of knowledge between studies of real- world

practices, which generates laboratory research directions, and the results of laboratory studies, which are directed by field work to more effectively address technology designers' "need to know" about analytical cognitive processing. As in the health sciences, translational research methods will require a co-evolution of field and lab approaches⁷. This blurring of the somewhat arbitrary distinction between pure and applied science is especially appropriate and useful for VA. It is quite evident that the challenges and problems that VA was established to undertake are deep and thus require new, deep basic research and methods, even in the core areas of visualization science, data analysis and knowledge acquisition, and other areas.

Beyond these general considerations, there is much specific work that must be done to build the science of analytical reasoning. Pottenger et al.¹⁷ propose an approach based on the core idea that an interactive analysis system can use data from user interactions to infer high-level knowledge of the analyst's state within a Sensemaking process, and then use this high-level knowledge to provide feedback that encourages a maximally effective route through the Sensemaking process²¹. In other words, the software should know enough about what the user is doing to be able to support a Sensemaking profile that provides high efficiency and minimizes errors due to known human cognitive limits.

One can see how this higher-level knowledge might be obtained through machine learning, say, by having a Sensemaking expert together with an analyst go over the user interaction sequence and manually label Sensemaking states and then train an HMM or Markov Net to predict where transitions occur. However, experience in machine learning shows that this would produce too large a number of parameters to learn without a huge amount of training data, and the model would likely be too sensitive to variations in users' investigative styles. With such a setup we would be unlikely to learn much useful knowledge about the Sensemaking process.

Another, perhaps better, way to gain high-level knowledge is to encode more detailed knowledge of the elements of the Sensemaking, from the level of raw user interactions up to the higher-level abstractions of the model. This is in line with the general observation above about the need to "get inside" the sensemaking steps; the work already done developing and evaluating visual analytics tools for reasoning tasks could also be used to develop this detailed knowledge^{2, 5, 18, 19}. In representing the Sensemaking model there are three primary classes of objects: Stages,

Artifacts, and Data Tasks. The dependency structure of these classes can be encoded in a description language. The Sensemaking stages and artifacts are represented as classes, and the relationships between them can be encoded in the class properties. This will produce a set of hierarchical relationships that shows the Sensemaking tasks at the highest layer of abstraction, the artifacts at the next lower layer, and the data tasks at the lowest layer. The data tasks can be recorded directly from the user interface, where recording them corresponds to instantiation of objects of the classes of the Sensemaking ontology.

The structure of this ontology immediately suggests how the Sensemaking tasks can be inferred from data tasks. In addition, this ontological model of the process can be analyzed to find entailments that provide limits to the set of Sensemaking stages that are logically possible given a history of data tasks performed and access to artifacts, thus limiting the scope and complexity of the learning problem. However, if we make the relationships in the ontology such that the current Sensemaking state is always a matter of strict necessity, we would have a model that required a much too restrictive user interface and workflow. Thus we see the need of learning from actual analysts rather than specifying the entailments of the model too strictly. In addition, the ontological structure itself may have to be loosened so that, for instance, categories of general knowledge can be included rather than just task-specific knowledge and so that the ontology can be dynamic.

Much work remains to be done with dynamic data and the temporal structure of data, especially with streaming data and with data that accumulates into long histories. This is not just a problem of data representation and transformation, though certainly the approaches described in the paper in this issue by Kasik et al.¹³ are essential. Dynamic data must be also considered from the standpoint of reasoning and analytical processes. Here the event-based approaches mentioned above^{16,6} are a start in the right direction, but much more remains to be done. Beyond this, modeling approaches that fit, in a mixed-initiative sense, into a human/computer system must be developed. For large scale emergencies, fixed plans are often quickly invalidated by events. (Large-scale emergencies are just one example where these approaches to dynamic data and temporal processes are needed.) One needs, rather, a *dynamic* model-based plan, that can be updated and responds to unexpected events in the current situation. Otherwise even expert decision-makers can be overwhelmed. The responses to hurricanes Katrina and Rita give ample evidence of this issue. Although incompetence and failure to respond quickly were major factors in the

Katrina disaster, it is questionable whether the best decisions would have been made even if responders and officials had been much better at doing their jobs, once the levees started failing and the situation changed completely. A general model that could take the new situation into account running within a visual analytics system so that it was readily available to decision-makers would have been a great help. In the case of Rita, previously established evacuation routing plans failed completely because people did not act as expected, trying to leave all at once due to Katrina being fresh in their minds. One aspect that must immediately be faced if evacuation, critical infrastructure, and other models are to be integrated into visual analytics systems is to make the models fast and interactive; otherwise they won't be compatible with the VA interactive interfaces nor with the need to respond to a disaster as it unfolds. Trade-offs must be made between accuracy versus speed and, more generally, with not having complete resources (lots of data could be missing, for example). Systems such as RESIN³¹ are looking at this but much more remains to be done. In addition, the models must be placed into the analytical reasoning environment in such a way that they fit the human reasoning and decision-making process. The HCM¹⁰ and cognitive analyses are starting to provide some direction here, but, again, much more must be done.

There is much work to do with respect to mixed-initiative approaches and human cognitive modeling. Although there are some promising initial results, the surface of this broad and deep area has just been scratched. With respect to HCM, many experiments and evaluations should be done to validate the precepts of the model and to see how they apply to specific tools and tasks. Although the initial modeling is founded on principles from cognitive science and related areas that already have some validation and evaluation, the cognitive analyses that we must attack with VA tools are more complex, larger, and deeper than what have been considered before. If nothing else, the required HCM must be more comprehensive, and it will be necessary to see how precepts that may have some independent validity work together.

Some VA tools have been used collaboratively^{12, 23} and some work has been done on design considerations¹¹, but not much work has been done that studies collaboration in a more fundamental way with respect to analysis and reasoning. Certainly, tools such as Jigsaw, SRS, and others have capabilities that lend themselves to collaboration because they enact knowledge structures and track reasoning processes that can then be shared. From the standpoint of the scope of this paper, sharing knowledge through a knowledge structure and sharing aspects of the

reasoning/argument-building processes are essential to meaningful collaboration. However, a fundamental approach is needed that investigates the form that collaboration should take, what artifacts should be shared, in what form, and at what stage of the reasoning process. This approach would help us decide the design and functionality of the interactive VA tools. At least three types of collaboration might come into play. The first is collaboration within a group of equal status though maybe with different tasks, say a group of analysts. The second is collaboration across groups, say between analysts and professionals in charge of emergency planning. The third is vertical collaboration within a hierarchy between, say, analysts and their managers.

The HCM must also be made predictive and practical. To be practical, the model must have several diverse instances of application so that one can clearly see how its precepts can be applied and can get a concrete sense of the success and effectiveness of the application. To be truly powerful, the HCM must be predictive so that one has an operational approach for not only designing but *optimizing* VA tools for complex reasoning applications. At the core of the predictive model is the goal that actions taken by the computer—presentation of new data, modification of existing data, computation and analysis of patterns of data etc. should be done in such a way that they do not interfere with the human's train of thought or flow of reasoning. Indeed, the temporal constraints of human memory, perception, and cognitive processing are such that optimizing the sequence of computational operation to match those constraints are likely to enhance the depth of analysis that the user is capable of. "In the zone" is the term often used to indicate the state of heightened cognitive efficiency and insightful thinking that can be achieved. In the predictive HCM, one would want to predict both qualitatively and quantitatively what being in the zone means, how it can be achieved, and such things as what certain types of distractions cost, how cognitive efficiency can be measured, and how (especially interaction) techniques can be ranked. One could then imagine having a cost/benefit model such as van Wijk's²⁸, but much more detailed and predictive, with which to determine the value of a particular VA approach. This work is complicated, however, by the fact that research in neuroscience²⁹ and other areas indicates that the focus required, for example, to complete certain types of cognitive tasks more quickly may significantly impede the human's ability to have a flash of insight or an important new idea³. An environment that supports more free association and less focus is superior for this. Thus there is the tension of trying to balance contradictory characteristics, both of which may be needed.

Acknowledgments

The authors acknowledge the contributions of Jason Perry and Christopher Janneck to the discussion of the proposed Sensemaking model. This work is supported in part by the U.S. Department of Homeland Security Science and Technology Division.

References

1. Adams, Summer and Goel, Ashok (2007) "Making Sense of VAST Data," *Proc. IEEE Intelligence and Security Informatics*, New Brunswick, New Jersey, pp. 270-273.
2. Chang, Remco, Ghoniem, Mohammad, Kosara, Robert, Ribarsky, William, Yang, Jing, Suma, Evan, Ziemkiewicz, Caroline, Kern, Daniel, and Sudjianto, Agus (2007) "WireVis: Visualization of Categorical, Time-Varying Data from Financial Transactions," *Proc. IEEE VAST 2007*, .pp. 155-162.
3. Chang, Remco, Ziemkiewicz, Caroline, Green, Tera Marie, and Ribarsky, William (2009) "Defining Insight for Visual Analytics," *IEEE Computer Graphics and Applications*, vol. 29, no. 2, pp. 14-17.
4. Corkill, D. (1991) "Blackboard Systems," *AI Expert*, 6(9):pp.40-47.
5. Dou, Wenwen, Jeong, Dong, Stukes, Felesia, Ribarsky, William, Lipford, Heather, and Chang, Remco (2009) "Recovering the Reasoning Process from User Interactions," *IEEE Computer Graphics & Applications*. To be published.
6. Eccles, Ryan, Kapler, Thomas, Harper, Robert, and Wright, William (2008) "Stories in Geo Time," *Information Visualization*, 7(1), pp. 3-17.
7. Fisher, B. (2009) "Science and Smart Graphics," *Information Technology* (in press).
8. Gallison, P. (1997) "Image & logic: A material culture of microphysics," *Chicago: University of Chicago Press*.
9. Goel, Ashok K., Adams, Summer, Cutshaw, Neil, and Sugandh, Neha (2009) "Playing Detective: Using AI for Sensemaking in Investigative Analysis," *Georgia Tech GVU Technical Report (# GIT-GVU-09-03)*.
10. Green, Tera, Ribarsky, William, and Fisher, Brian (2008) "Visual Analytics for Complex Concepts Using a Human Cognition Model," *Proc. IEEE VAST 2008*, pp. 91-98.
11. Heer, J, Agrawala, M (2008) "Design Considerations for Collaborative Visual Analytics," *IEEE Information Visualization 2008*, pp. 49-62.

12. Isenberg, P, and Fisher, D (2008) "Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections," *Computer Graphics Forum*, to appear.
13. Kasik, David et al (2009) "Data Transformations and Representations for Information Generation," *Information Visualization*, Special Issue on the Future of Visual Analytics, Vol. 8(4).
14. Kim, Sangho, Shekhar, Shashi, and Min, Manki (2008) "Contraflow Transportation Network Reconfiguration for Evacuation Route Planning," *Transactions on Knowledge and Data Engineering (TKDE)*, 20(8).
15. LaFree, G. and Dugan, L. (2007) "Global Terrorism Database, 1970 – 1997," [Computer file]. ICPSR04586-v1, College Park, MD: University of Maryland [producer]. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2007-04-04.
16. Luo, Dongning, Yang, Jing, Fan, Jianping, Ribarsky, William, and Luo, Hangzai (2009) "EventRiver: Interactive Visual Explorations of Streaming Text," *VisCenter Technical Report CVC-UNCC-09-05*.
17. Perry, J.M., Janneck, C.D., and Pottenger, W.M. (2009) "Supporting Cognitive Models of Sensemaking in Analytics Software Systems," *DIMACS Research Report 2009-12*, Rutgers University.
18. Pike, W.A., May, R., and Turner, A. (2007) "Supporting Knowledge Transfer through Decomposable Reasoning Artifacts," *HICSS*, doi.ieeecomputersociety.org/10.1109/HICSS.2007.508.
19. Pike, W.A., Baddeley, B., Best, D., Franklin, L., May, R., Rice, D.M., Rienische, R., and Youngking, K. (2008) "The Scalable Reasoning System: Lightweight Visualization for Distributed Analytics," *IEEE VAST 2008*, pp. 131-138.
20. Pike, WA, et al. (2009) "Science of Interaction," *Information Visualization* Special Issue on the Future of Visual Analytics, Vol. 8(4).
21. Pirolli, P. and Card, S. (2005) "The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis," *Proceedings of International Conference on Intelligence Analysis* 2005: 2-4.
22. Plaisant, Catherine, et al, (2008) "Evaluating Visual Analytics at the 2007 VAST Symposium Contest," *IEEE Computer Graphics and Applications*, Vol. 28, No. 2, pp. 12-21.
23. Robinson, A (2008) "Collaborative Synthesis of Visualization Analytics Results," *IEEE VAST 2008*, pp. 19-24.

24. Stasko, John, Gorg, Carsten, and Liu, Zhicheng, (2008) "Jigsaw: Supporting Investigative Analysis through Interactive Visualization", *Information Visualization*, Vol. 7, No. 2, pp. 118-132.
25. Stokes, D. E. (1997) "Pasteur's quadrant : Basic science and technological innovation," *Brookings Institution Press*, Washington D.C.
26. Thagard, Paul (2005) "Being Interdisciplinary: Trading Zones in Cognitive Science," *Interdisciplinary Collaboration: An Emerging Cognitive Science*, S.J. Derry, C.D. Schunn, M.A. Gernsbacher (Eds) LEA, Mahwah NJ.
27. Thomas, J. and Cook, K. (2005) "Illuminating the Path: The Research and Development Agenda for Visual Analytics," *National Visualization and Analytics Center*.
28. van Wijk, JJ (2005) "The value of visualization," *IEEE Visualization, Proceedings of Vis 2005*, pp. 79 – 86.
29. Wilson, T.D., and Schooler, J.W. (1991) "Thinking too much: introspection can reduce the quality of preferences and decisions," *Journal of Personality and Social Psychology*, vol. 60(2), pp. 181-192.
30. Woolf, Steven H. (2008) "The Meaning of Translational Research and Why It Matters," *JAMA*, 299(2):211-213.
31. Yue, Jia, Raja, Anita, Liu, Dingxiang, Wang, Xiaoyu, and Ribarsky, William (2009) "A Blackboard-based Approach towards Predictive Analytics," *Proceedings of AAAI Spring Symposium on Technosocial Predictive Analytics*, pp 154-161, Stanford University, CA.

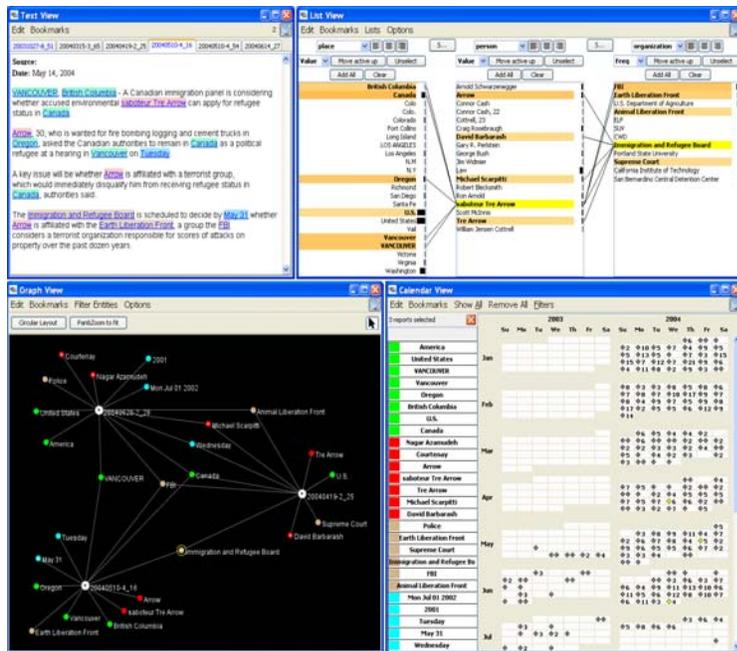


Figure 1. Jigsaw multiwindow interface²⁴. The document view is in the upper left, the list view is in the upper right, the calendar view is to the lower left, and the graph view is to the lower right.

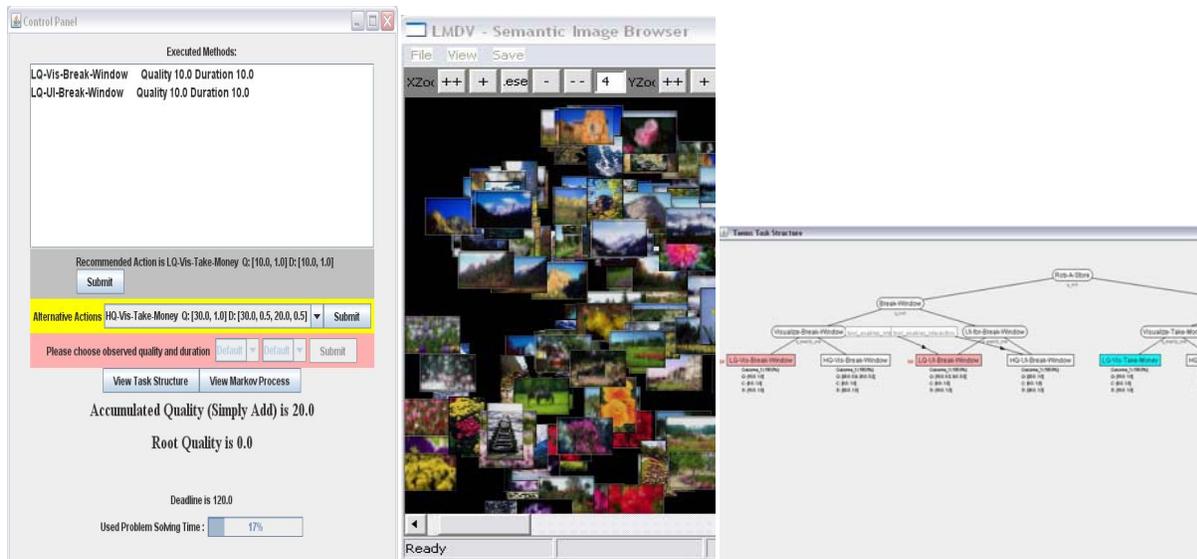


Figure 2. RESIN³ interface for image browsing and analysis application. The control panel on the left is set for a task with a tight deadline. It sets the scale of imagery to be browsed in the center tool. The panel on the right updates information on real-time execution of sub-tasks so that the user can allocate time appropriately for completing the end-to-end task.