

Middlesex University Research Repository

An open access repository of

Middlesex University research

<http://eprints.mdx.ac.uk>

Xu, Kai ORCID: <https://orcid.org/0000-0003-2242-5440>, Attfield, Simon ORCID: <https://orcid.org/0000-0001-9374-2481>, Jankun-Kelly, T. J., Wheat, Ashley, Nguyen, Phong H. and Selvaraj, Nallini (2015) Analytic provenance for sensemaking: a research agenda. IEEE Computer Graphics and Applications, 35 (3) . pp. 56-64. ISSN 0272-1716 [Article] (doi:10.1109/MCG.2015.50)

Final accepted version (with author's formatting)

This version is available at: <http://eprints.mdx.ac.uk/15814/>

Copyright:

Middlesex University Research Repository makes the University's research available electronically.

Copyright and moral rights to this work are retained by the author and/or other copyright owners unless otherwise stated. The work is supplied on the understanding that any use for commercial gain is strictly forbidden. A copy may be downloaded for personal, non-commercial, research or study without prior permission and without charge.

Works, including theses and research projects, may not be reproduced in any format or medium, or extensive quotations taken from them, or their content changed in any way, without first obtaining permission in writing from the copyright holder(s). They may not be sold or exploited commercially in any format or medium without the prior written permission of the copyright holder(s).

Full bibliographic details must be given when referring to, or quoting from full items including the author's name, the title of the work, publication details where relevant (place, publisher, date), pagination, and for theses or dissertations the awarding institution, the degree type awarded, and the date of the award.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Middlesex University via the following email address:

eprints@mdx.ac.uk

The item will be removed from the repository while any claim is being investigated.

See also repository copyright: re-use policy: <http://eprints.mdx.ac.uk/policies.html#copy>

Analytic Provenance for Sensemaking: A Research Agenda

Kai Xu Simon Attfield T.J. Jankun-Kelly
Ashley Wheat Phong H. Nguyen Nallini Selvaraj

Abstract

Sensemaking is a process of find meaning from information, and often involves activities such as information foraging and hypothesis generation. It can be valuable to maintain a history of the data and reasoning involved, commonly known as *provenance* information. Provenance information can be a resource for “reflection-in-action” during analysis, supporting collaboration between analysts, and help trace data quality and uncertainty through analysis process. Currently, there is limited work of utilizing *analytic provenance*, which captures the interactive data exploration and human reasoning process, to support sensemaking. In this article, we present and extend the research challenges discussed in a IEEE VIS 2014 workshop in order to provide an agenda for sensemaking analytic provenance.

Keywords. Provenance, Senesmaking, Visual Analytics, Collaboration, Data Quality.

Sensemaking is a process of finding meaning from information — a process of comprehension. It is the construction, elaboration and reconciliation of representations which account for and explain the information we receive about the world. Sensemaking often involves a variety of activities such as information foraging and triage, schematization, and hypothesis generation and validation. During complex sensemaking tasks, it can be valuable to maintain a history of the data and reasoning involved and the context within which sensemaking was performed – referred to as *provenance* information. Provenance information can be a resource for “reflection-in-action” during analysis, supporting collaboration between analysts, and help trace data quality and uncertainty through analysis process. It can also act as a resource after the event, supporting the interpretation of claims, audit, accountability, and training.

There has been considerable work on capturing and visualizing *data provenance*, which focuses on data collection and computation, and *analytic provenance*, which captures the interactive data exploration and human reasoning process. However, there is limited work of utilizing such provenance information to support sensemaking, in terms of improving efficacy and avoiding pitfalls such as uncertainty and human bias. A workshop was held during IEEE VIS 2014 with the aim of bringing together researchers involved in visual analytics

and various aspects of sensemaking to bridge this gap. The workshop participants considered emerging positions and findings related to the capture, processing, representation and use of provenance information to support complex sensemaking tasks. In this article, we present and extend the research challenges discussed in the workshop in order to provide an agenda for sensemaking analytic provenance.

The research challenges are organized in the *capture*, *visualize* and *utilize* order. The paper starts with a hierarchical provenance model that forms the basis of the following discussions on analytic provenance capture and visualization. While the focus is on using analytic provenance to support sensemaking (i.e., “utilize”), “capture” and “visualize” prescribe what analytic provenance information is available and the possible means to utilize it respectively. The two subsequent sections consider two facets of “utilization”: collaboration and uncertainty/trust. While there are many possible applications, these are two areas that analytic provenance can potentially make great impact and where many open research problems remain. Finally the Conclusion section summarizes the research challenges. While not a complete survey, the paper provides reference to publications that serve as examples to the research challenges discussed.

1 Modelling

Analytic provenance information can be categorized using a four-layer hierarchical model based on its semantic richness [7]. Figure 1 shows this model using analyzing stock market as an example: the level of semantics increases from bottom to top. The bottom-level *events* consists of low-level user interactions such as mouse clicks and keystrokes, which have little semantic meaning. The next level up is *actions*, which are analytic steps such as querying the database or changing the zooming level of data visualization. The parameters such as data description and visualization settings are also part of the provenance. Further up are the *sub-tasks*, which are the analyses required to achieve the sensemaking goal. In the case of stock market analysis, examples are identifying top performing companies and determining long term trends. In the top-level is the *task*, i.e., the overall sensemaking undertaking, which is “analyzing stock market”.

Analytic provenance is closely linked both within and across layers. Within a layer, analytic provenance is linked temporally (i.e., one event happens after another) and logically (e.g., one action depends on the two previous actions). There are also connections across layers: a database query action consists of several mouse click and key stroke events, and it is part of a higher level sub-task level such as “comparing stock performance”.

2 Capture

Analytic provenance capture provides the data for its visualization. What provenance is available and its quality decides what provenance visualization




Figure 1: The hierarchical analytic provenance model shown with an example of analyzing stock market [7]. The semantic richness increases from bottom to top. The bottom layer are the *events* such as key presses and mouse clicks, which have little semantics. The next level up are *actions* such as the database query and visualization zooming. Further up are the *sub-tasks*, which usually are the analyses performed during the sensemaking. The top level *tasks* are the overall sensemaking undertaking.

is possible and its quality. Capturing lower level events and actions is relatively straightforward in a visual analytics system. However, such analytic provenance information alone is of limited use [7]. Tasks and sub-tasks provide important clues to the purpose and rationale underlies the sensemaking. However, they are largely part of users’ thinking, which a visual analytics system does not have direct access to. This is one of the biggest challenges in analytic provenance capture. There is a limited time window to capture such information; even the users themselves may forget what they were doing after a while, at which point it becomes very difficult to recover the analytic provenance information.

Existing approaches to capture high level analytic provenance can be broadly categorized into *manual* and *automatic* methods. The manual methods largely rely on users recording their analysis process and sensemaking tasks, whereas the automatic methods try to infer the higher level tasks and sub-tasks from lower level events and actions. While the manual approaches are usually more accurate, it can distract user from the actual analysis task, which may discourage users from recording analytic provenance. On the other hand, the automatic approaches do not introduce interruption to the sensemaking process, but their capability of inferring semantic-rich analytic provenance information is limited [7]. Personal differences introduces additional difficulty for automatically inferring higher-level analytic provenance. Users’ knowledge and experience have a considerable impact on the way they conduct analysis. As a result, the sensemaking process (i.e. the analytic provenance) can vary significantly from user to user, even with the same dataset and analysis task.

2.1 Manual

As previously discussed, manual capture mostly focuses on the task and sub-task level. Allowing user annotation is one of the most common forms: User creates *notes* or *annotations* that are associated with certain data, analysis result, and/or visualization. The content of a “note” is not limited to findings or discoveries; it can also include the thinking that leads to a finding or the relationships between findings. *Data-aware annotation* links the findings and associated visualization to the underlying data used to produce them, which makes it possible to apply new analysis and visual mapping at a later stage if further investigation is needed.

While individual note only represents a fraction of the analytic provenance, it is possible to provide a reasonably good overview of the sensemaking process if a number of notes and the connections between them are captured. However, this is only possible when users are willing to take notes, which can be perceived as distractions sometimes. There are two common strategies to alleviate this: minimizing interruption/cognitive effort and providing tangible benefits to the sensemaking task. Reducing interruption and cognitive effort can lower the likelihood that users are discouraged from recording analytic provenance. This can be achieved through integration with the analysis tools (so users do not need to switch between interfaces) or streamlining the recording process (e.g., with minimal mouse clicking and movement). Besides, it is likely to motivate user adoption if the analytic provenance captured can provide perceivable benefits to the analysis task, i.e., immediate support of sensemaking process. Examples include the ability to record discoveries during the analysis [18] and review/plan exploratory analysis for complex sensemaking task [9]. However, currently there is a lack of general design guidelines for how to achieve them, and there are few user studies evaluating how effective they are, in terms of both the benefits they bring and the potential cognitive cost they can introduce. Any progress related to these two challenges can have a considerable impact on the capture of analytic provenance and enable better support for the sensemaking.

2.2 Automatic

One of the main disadvantages of manual capture is the requirement of direct input from users. Automatic approaches try to address this by inferring higher level analytic provenance from what can be automatically captured. As discussed earlier, it is easier to capture analytic provenance at the event and action level. Therefore, most automatic approaches try to infer sub-task and task-level information from event and action provenance.

This turns out to be a difficult challenge. An experiment studied how much of a user’s reasoning process can be recovered from user action information [3]. A domain-specific sensemaking task was used and experts were recruited to analyze the user action log. Higher-level analytic provenance manually inferred from the interaction logs were compared with the ground truth obtained through interview. The results showed that 79 percent of the findings, 60 percent of the

methods, and 60 percent of the strategies were correctly recovered. The accuracy is not high even in such a constrained setting with domain experts doing the inference. Given the diversity of data and analysis involved in the sensemaking and the difficult of replicating expert knowledge/thinking in a computer system, the chance of having a generic technique that can accurately infer semantic-rich analytic provenance information for a variety of analysis tasks is not high.

Instead, existing methods either constrain the problem/analysis domain or Aim for less semantically rich analytic provenance. By limiting the choice of data and analysis/visualization, an inference algorithm has better chance to make the right guess. However, even within a specific domain (such as finance), the types of data and analyses involved are still of very large amount. Also, being limiting on the data and analysis can constrain the system capability, having a negative impact on the sensemaking task.

Given the difficulty of inferring task/sub-task information, a few methods target less semantic-rich provenance. One such example is “action chunking”, i.e., identify a group of actions that are likely to part of the same sub-task, without knowing what the sub-task is. Such approaches apply heuristics to infer patterns from action logs based on repeated occurrence and proximity in data/visualization space or analysis time [7]. Such chunking information can be useful in several ways. For example, the system can prompt user to take a note if such an action usually occurs within a specific sequence. Also, the grouping information can be used for aggregation when large amount of provenance information is to be visualized. This method is later extended to monitor user behavior for implicit signals of user intent and uses the information to suggest alternative visualization [6]. It is an open research problem to explore similar analytic provenance that can be effectively inferred and provides semantic information that can be used for supporting sensemaking.

For future research, a promising direction is the development of “hybrid” or “semi-auto” approaches, i.e., mixing the manual and automatic capture to combine their strength. For example, the previously mentioned “action chunking” can be further improved with user feedback: the algorithm can “learn” or improve itself using the user input that whether a group of actions form a sub-task. The improved algorithm can in turn help improve the manual capture by prompting users to take notes if such an action is expected within certain “action chunks” from previous experience. This type of approach is not limited to “chunking”. For example, an algorithm that predicts sub-tasks can ask for user feedback (i.e., whether the prediction is correct or not) and use the information to improve itself. Similar approaches can be used to uncover user intention or analysis strategies.

3 Visualization

Most existing provenance visualization methods focus on the action layer, which can be automatically captured and still offers certain level of semantic information. Often included are a series of user actions and notes, together with the

information of the data, computational analysis, and visualization that are associated with them. In most cases it is difficult to show all these information at once. Instead, existing methods often display selected provenance based on their design goals, with details on demand.

Node-link diagrams are a popular choice among methods that aim to show an overview of the sensemaking process [1, 4, 12, 15]. They usually follow the temporal order or the casual relationship among actions. In such methods, nodes represent a summary of system state and the edges represent actions that transit system from one state to another. While providing an overview of the sensemaking structure, in many cases node-link diagrams do not have sufficient details for understanding the semantics of user action. To provide more context, the most common approach is multiple-coordinated views that show the note and system state only for a selected step [12, 15]. This usually works well with many visual analytics systems, which already have view for each type of information: showing the sensemaking context essentially restores the system to a previous state. However, such setup still requires users to go through a process step by step, sometimes back and forth, to understand an analysis sequence, which places heavy cognitive work load on the user’s memory. Methods such as GraphTrail [4] show multiple system states and the links between them at the same time. By allowing zoom and pan, users can choose between overview of the analysis structure and details of individual system state. However, analysts can easily generate dozens or more system states within a short period, and this starts to reach the limit of such methods. To further improve scalability will require filtering or aggregation, and the research challenge is to guide the users such that interesting patterns in the sensemaking process are not lost. This will depend on the understanding of the provenance semantics, so for example unimportant actions can be filtered or a sub-task can be used as an aggregation of a series actions. This is closely related to the provenance capture discussed in the previous section.

Besides providing a deeper understanding of the sensemaking process, analytic provenance can directly support to some sensemaking tasks. One example is (visual) narrative construction, during which user composes findings into a coherent “story”. A narrative can include raw data, analysis results, visualization, and user notes. Narratives describe the final conclusions in the context of the sensemaking process that leads to them, a useful feature for reporting and team collaboration. The DIVA system [18] allows interactive construction of narratives from user annotations and associated visualization states (Figure 2). The SchemaLine [10] allows users to create hypotheses or narratives by grouping notes along the timeline (Figure 3).

Analytic provenance has also been used to help users review their sensemaking process and guide further exploration, which is particularly useful for analysis of complex dataset such as those with high dimensionality. Such methods [9, 14] visualize the sensemaking space so user can easily see which part has been explored, e.g., which data dimension and which values within that dimension have been analyzed. Users can use this information to plan their further analysis and system can also use this information to suggest related but




Figure 2: A narrative created in DIVA: each bookmark (box) is a saved visualization state (including the uncertainty information), together with the note (text at the bottom). Related bookmarks are linked together to form a narrative [18].

unexamined data.

The research of how the provenance visualization can support sensemaking is still in its early stage and many research challenges remain. For example, neither the DIVA nor the SchemaLine provides any support for narrative construction beyond connecting saved states/notes. They entirely rely on the users to identify the relevant findings and identify the relationships among them. In terms of sensemaking guidance, support for analyses involve high-dimensional data and/or long investigation process is almost non-existent.

4 Collaborative Sensemaking

So far the application of analytic provenance to support sensemaking is mostly for individuals. Provenance can remind people how to interpret their own findings, direct them to areas where their analysis is lacking, and even help them to conceptualise (or make sense of) what it is that they are trying to do. However, as visual analytic systems move from the research lab to real world applications, collaboration becomes an increasingly significant issue. It is an issue ripe for research, where provenance information can play an important role. Hence we focus some attention on this.




Figure 3: In SchemaLine, each piece of text is an analyst note, positioned along the time axis according to the temporal information in the associated data. Users can link related notes to form a “schema”, which can be either a hypothesis during the early stage of analysis or a narrative to present the final findings. There are three schemata in this example represented as differently colored rectilinear paths [10].

The need for visibility is demonstrated to some extent through the use of “coordination artefacts” that collaborators sometimes use, such as written plans, procedures, timetables, schedules, checklists and other mechanisms which offer cues about intentions and action. By providing a trace of activities of collaborators who may be acting at a distance and asynchronously, provenance information has the potential to play an important role in providing cues for collaboration. Seeing the record of the actions of others allows the inference of their intent that may not be present in their results alone. As such, one of our research issues is the coordination—or handoff—of provenance between collaborators.

In addition to supporting collaboration around common goals, provenance information can also provide a basis for sharing best practice. What counts as best practice may not be immediately evident and may need to be identified over time and in relation to pre-defined success measures or indicators. Nevertheless, capturing the way that tasks have been tackled through provenance provides an opportunity for reconstructing successful approaches and identifying their significant features after the fact. This could then provide a basis of training and processes. Conversely, provenance may also provide case studies for failed processes for training. However, using failed processes without damaging the reputation of the involved users is an open research problem.

4.1 Privacy

A key issue to be considered when designing any system which involves the recording and retrieval of people’s actions is that of privacy. It may be unethical or even illegal to record all of a user’s actions performed on a system without their prior permission to do so. In the design of systems which record provenance, designers need to consider exactly what data is recorded, what that data will be used for, and by whom. Depending on context, it may or may not be appropriate to design a system which is sufficiently ‘socially translucent’ such that people can be held accountable for their actions [5]; privacy may require

that there are contexts in which they should not.

Furthermore, the way in which provenance is captured, and the level of accountability that can be embedded in data captured by a system, may affect the way people use it. For instance, when data is aggregated and individual users are unable to be identified, they may be less reluctant to explore or experiment, as a level of plausible deniability preserved. Conversely, when a system is able to record individual user identities and their actions, people may be less inclined to perform certain actions in fear of for example recrimination from their superiors.

One potential solution to this is the preservation of anonymity. But this poses a further issue: when users collaborate, it can be important for them to be able to identify other members of the team and their contribution. In order to remedy this, a system may incorporate some level of internal/external split, meaning that some information can only be held internally within a group. Additionally, there may be different levels of privilege within a system according to a hierarchical structure, with different levels of access according to a user's level within that structure. However, the acceptance of this from a cultural perspective must be taken into consideration. Cultures where there is a higher power distance and more well defined hierarchy will find such solutions more acceptable than those with a more equal distribution of power.

4.2 Handoff of Provenance

Essential to asynchronous collaborative work in many disciplines is the process of handoff. Handoff is the transfer of responsibility for a task from one person or team to another, which by necessity is often accompanied by an exchange of information. This may include information about state of a domain of interest at a given point, work that has been carried out or work that is planned. A good deal of research has been done to understand handoff in the medical domain. In nursing there are well established practices and protocols which have been studied in some depth [11]. And handoff has also been studied in air traffic control, crowdsourcing and robotic system. We see value in examining the existing literature on handoff, particularly where researchers have studied the discourse involved. We also see the value of performing further cross-domain studies in areas such as intelligence analysis or software development. The overall aim would be to identify and abstract common principles of handoff with implications for the use of provenance information. Questions here would revolve around the way in which experienced domain practitioners have learned to abstract and communicate the essentials of complex episodes, outcomes and future possibilities. Handoff discourse, however, is likely to be very context dependent. So in abstracting away from that context, attention would need to be paid to the role of factors such as common ground and material artefacts play in allowing assumptions to be made and details skipped. Rather than such phenomena making handoff practices difficult to generalise, however, we see them as potentially indicating the kinds of common ground and artefacts which would make provenance information usable by others — something that it would be useful to understand.

5 Uncertainty and Trust

Underlying the challenges faced by visualizing provenance and understanding its use in collaboration are questions about the validity of the process and its record. Original data may be of low quality, depictions and interactions with the data may exacerbate uncertainty, leading to a lack of trust (or over-trust) of the result [13]. These two issues—uncertainty and trust—present significant challenges in the successful use of analytic and data provenance in sensemaking.

5.1 Uncertainty

Uncertainty, in our context, are variations from the stated value introduced to our data before or during its analysis. Before analysis, uncertainty stems from lack of precision in measurement, inconsistencies in recorded results, or missing values; these are issues from data provenance. During the analysis, tools in the workflow can modify or introduce uncertainty—sampling and aggregation, such as that done to ensure privacy, transforms even certain values into a representational result with some variance from the population; these are issues from the analytic provenance. The result of any individual visualization is thus uncertain, even if it is presented in a manner that hides this variation.

From our workshop discussion, there are three main challenges driven by the uncertainty in the analytics process. First, it is unclear from a general standpoint how to characterize uncertainty—what are the appropriate metrics for different types/sources of uncertainty, and how do they appropriately propagate through workflows [13]? Error analysis is well studied for arithmetic operations, but how do they combine under sampling, aggregation, or other transformation? How do we quantify and propagate uncertainty due to multiple witness statements, intelligence reports, or other non-quantitative measures? There is a research opportunity to characterize a reusable typology of uncertainty factors with known propagation methods. This typology will likely be built from domain specific examples of uncertainty first before a more general model is known.

The second, connected challenge relates to using the uncertainty to guide insight discovery. Even if the uncertainty in the process is understood, it is unclear how to model what the user currently knows about the data (and its certainty) or the extent of the analysis space covered. Metrics about the exploration process can assist [8], and the methods alluded to in the previous sections can partially address this challenge. There is a research challenge to integrate a model of uncertainty into these recommendation and insight modeling systems. Especially challenging is modeling and highlighting the unknowns—what is the uncertainty hiding in the data, or what are the range of valid results.

Both of the previous challenges require an understanding of how the uncertainty affects the user’s understanding and their sensemaking. Thus, how to synthesis understandable uncertainty that fits the user’s model of uncertainty is our final research challenge. Sensemaking under uncertainty needs to be studied to characterize and mitigate misunderstandings that occur due to this inherit

lack of information. While every sensemaking tasks begins with lack of knowledge of the final result(s), this “unknowledge” is different from sensemaking under uncertainty, where the process itself cannot be given full trust. Understanding the best practices for mitigating uncertainty in the process will assist users make decisions under uncertainty.

5.2 Trust

Even in an certain sensemaking process, levels of trust in the results may vary; uncertainty makes trusting decisions more fraught. In sensemaking analytics, trust appears in three contexts: Trust in the data, trust in the process, and trust in the result. Each presents challenges for research.

Trust in the data is an issue of data quality. Uncertainties introduced in the measurement, storage, and access of data all affect the trust in its validity. Provenance of the the data sourcing and workflow can be used as part of trust decisions regarding that data; as a computational artifact, data quality provenance can also be part of synthetic trust models [2,17]. The research problems here are both on the representation of the data’s quality (what are the appropriate metrics? How do these interact with our uncertainty propagation models as part of the workflow?) and on its communication to the user (how to indicate when a user is making risky inferences from data under low quality conditions? How do we depict the consequence of different quality representations in terms of workflow computational usage or result fidelity?).

A user’s trust in the analytical process, while related to uncertainty, incorporates other measures—the user’s trust in the data, their believed expertise of the material, and the cognitive biases they bring to the analysis. How can a computer synthesize a model of trust built from these factors? Venters’ et al. [17] and Sacha [13] suggest tracking the provenance of the data and the analytical process to measure a user’s trust in the process—tighter exploration loops suggest confidence whereas scattered exploration suggest distrust of the process. An open research challenge is to formally measure and quantify a trust inference model from given user explorations. While examples have been gathered, such as classroom visualization usage [16], more work is needed to generalize. It is also an open question of how to detect and communicate biases in the analytic process; inferring when a user is not exploring potentially fruitful avenues due to unconscious inattention is vital in robust recommendation systems.

Trust in the result is tied to their confidence in the process and the original data. Previously, we spoke of a model for the user’s inference and thus confidence of the process and result; in concert with that model would be one for measuring the risk associated in using the result. If uncertainty cannot be eliminated, it could be mitigated if appropriate measures of risk could be devised [17]. Determining appropriate risk models is an open research problem, and tying the risk to the uncertainty/trust is also an open challenge.

6 Conclusions & Research Agenda

Visual analytics can be improved via a better understanding of the behavior during the analytic process in support of sensemaking—provenance can be used for self reflection and exploration guidance, can facilitate collaboration, and help us understand what we can trust from possibly uncertain data. These separate aspects share the theme of lack of understanding—as a community, we do not know how best to utilize what we know about our processes to assist making decisions about what we know. We have presented several challenges raised from these topics during our IEEE VIS 2014 workshop as part of a research agenda for the community. Taken together, they form a four part research agenda:

- **Enhance provenance capture** to better support more accurate and higher level inference from analytic provenance. These may be manual, automatic, or hybrid, but such inference can assist in understanding the provenance process for better prediction, process correction, and decision making.
- **Develop and validate provenance visualizations for sensemaking.** Current research has only scratched the surface of the semantically rich space of information present in the provenance; to support the enhanced provenance capture recommended above, additional visual presentations are needed. Also, visualization techniques need to be scaled up to support long and complex sensemaking process.
- **Investigate privacy-aware methods to utilize collaborative provenance** that provide the appropriate level of detail depending on the sensemaking task and the role of the user. Proper collaboration will also require deeper understanding and generalization of the handoff provenance between collaborators in different domains.
- **Extend error propagation through provenance pipelines to wider types of uncertainty** via better typologies and studies of sensemaking risks under uncertainty. This agenda is synergistic with enhance provenance capture—better intent inference can be used to build model of trust in the sensemaking, whereas improved uncertainty models can correct over trusting inference models.

Systems and practices for supporting sensemaking are a vital part of the larger visual analytics context. We see that in the future, as visual analytics broadens its reach, better support for sensemaking will require solving these and other analytic provenance challenges.

References

- [1] L. Bavoil, S. Callahan, P. Crossno, J. Freire, C. Scheidegger, C. Silva, and H. Vo. VisTrails: Enabling Interactive Multiple-View Visualizations. In *IEEE Conference on Visualization*, pages 135–142. IEEE, 2005.

- [2] C. Bors, T. Gschwandtner, S. Miksch, and J. Gärtner. Qualitytrails: Data quality provenance as a basis for sensemaking. In *Proceedings of the International Workshop on Analytic Provenance for Sensemaking*, 2014.
- [3] W. Dou, D. H. Jeong, F. Stukes, W. Ribarsky, H. R. Lipford, and R. Chang. Recovering Reasoning Processes from User Interactions. *IEEE Computer Graphics and Applications*, 29(3):52–61, May 2009.
- [4] C. Dunne, N. H. Riche, B. Lee, R. A. Metoyer, and G. G. Robertson. GraphTrail: Analyzing Large Multivariate and Heterogeneous Networks while Supporting Exploration History. In *ACM Conference on Human Factors in Computing Systems*, pages 1663–1672, 2012.
- [5] T. Erickson and W. A. Kellogg. Social translucence: An approach to designing systems that support social processes. *ACM Transaction on Computer-Human Interaction*, 7(1):59–83, Mar. 2000.
- [6] D. Gotz and Z. Wen. Behavior-driven visualization recommendation. In *Proceedings of the 14th International Conference on Intelligent User Interfaces*, IUI '09, pages 315–324, New York, NY, USA, 2009. ACM.
- [7] D. Gotz and M. X. Zhou. Characterizing users' visual analytic activity for insight provenance. *Information Visualization*, 8(1):42–55, 2009.
- [8] T. J. Jankun-Kelly. Using Visualization Process Graphs to Improve Visualization Exploration. In J. Freire, D. Koop, and L. Moreau, editors, *Provenance and Annotation of Data and Processes*, volume 5272 of *Lecture Notes in Computer Science*, pages 78–91. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [9] A. Lunzer. Lightweight provenance-driven exploration. In *Proceedings of the International Workshop on Analytic Provenance for Sensemaking*, 2014.
- [10] P. H. Nguyen, K. Xu, R. Walker, and B. W. Wong. SchemaLine: Timeline visualization for sensemaking. In *2014 18th International Conference on Information Visualisation (IV)*, pages 225–233, July 2014.
- [11] R. P. Variations and commonalities in processes of collaboration: the need for studies of multiple settings. *Computer Supported Cooperative Work: CSCW: An International Journal*, 20(1-2):37–59, 2011.
- [12] W. A. Pike, R. May, B. Baddeley, R. Riensche, J. Bruce, and K. Younkin. Scalable visual reasoning: Supporting collaboration through distributed analysis. In *International Symposium on Collaborative Technologies and Systems*, pages 24–32. IEEE, May 2007.
- [13] D. Sacha, H. Senaratne, B. C. Kwon, and D. A. Keim. Uncertainty propagation and trust building in visual analytics. In *Proceedings of the International Workshop on Analytic Provenance for Sensemaking*, 2014.

- [14] A. Sarvghad and M. Tory. Exploiting history to reduce interaction costs in collaborative analysis. In *Proceedings of the International Workshop on Analytic Provenance for Sensemaking*, 2014.
- [15] Y. B. Shrinivasan and J. J. van Wijk. Supporting the analytical reasoning process in information visualization. *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*, page 1237, 2008.
- [16] C. T. Silva, E. Anderson, E. Santos, and J. Freire. Using vistrails and provenance for teaching scientific visualization. *Computer Graphics Forum*, pages 75–84, 2011.
- [17] C. C. Venters, J. Austin, C. E. Dibsedale, V. Dimitrova, K. Djemame, M. Fletcher, S. Fores, S. Hobson, L. Lau, J. McAvoy, A. Marshall, N. T. Townend, V. Viduto, D. E. Webster, and J. Xu. To trust or not to trust? developign trusted digital spaces through timely reliable and personalized provenance. In *Proceedings of the International Workshop on Analytic Provenance for Sensemaking*, 2014.
- [18] R. Walker, A. Slingsby, J. Dykes, K. Xu, J. Wood, P. H. Nguyen, D. Stephens, B. L. W. Wong, and Y. Zheng. An extensible framework for provenance in human terrain visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2139–2148, Dec. 2013.