Aalto University

Hall, Andreas; Ahonen-Rainio, Paula; Virrantaus, Kirsi

Insight provenance for spatiotemporal visual analytics

RESEARCH ARTICLE

Insight provenance for spatiotemporal visual analytics: Theory, review, and guidelines

Andreas Hall, Paula Ahonen-Rainio, and Kirsi Virrantaus

Department of Built Environment, School of Engineering, Aalto University, Espoo, Finland

Abstract: Research on provenance, which focuses on different ways to describe and record the history of changes and advances made throughout an analysis process, is an integral part of visual analytics. This paper focuses on providing the provenance of insight and rationale through visualizations while emphasizing, first, that this entails a profound understanding of human cognition and reasoning and that, second, the special nature of spatiotemporal data needs to be acknowledged in this process. A recently proposed human reasoning framework for spatiotemporal analysis, and four guidelines for the creation of visualizations that provide the provenance of insight and rationale published in relation to that framework, work as a starting point for this paper. While these guidelines are quite abstract, this paper set out to create a set of more concrete guidelines. On the basis of a review of available provenance solutions, this paper identifies a set of key features that are of relevance when providing the provenance of insight and rationale and, on the basis of these features, produces a new set of complementary guidelines that are more practically oriented than the original ones. Together, these two sets of guidelines provide both a theoretical and practical approach to the problem of providing the provenance of insight and rationale. Providing these kinds of guidelines represents a new approach in provenance research.

Keywords: visual analytics, provenance, spatiotemporal, cognition, reasoning, visualization, framework, guidelines, review

1 Introduction

Visual analytics (VA) seeks to combine the reasoning powers of humans and the automated analysis techniques of computers through interactive visual interfaces with the goal of

achieving effective analysis and decision making [45]. Research on provenance is an integral part of VA. It focuses on different ways to describe and record the history of changes and advances made throughout an analysis process [67]. Provenance information is important both for the analyst or analysts, e.g., to keep track of their progress, and for persons external to the analysis that might need to trace and verify the process, accuracy, and sources of an analysis. For example, in analysis connected to serious real-life problems it is important to be able to repeat the analysis process as well as to know the argumentation of decisions made during the analysis.

Provenance has traditionally been divided into analytical provenance, focusing on the analyst's reasoning processes [59,64,87], and data provenance, also known as data lineage, focusing on the computational workflow of the analysis [5,11,22,79]. This research belongs to the former, focusing on ways to communicate, through visualizations, the reasoning and the knowledge used by an analyst to persons external to the analysis process.

A recent organizational framework of provenance in VA further differentiates between five different types and six different purposes of provenance [67]. Of the types listed, this research deals with the provenance of insight, covering "cognitive outcomes and information derived from the analysis process," and the provenance of rationale, covering "the reasoning behind decisions, hypotheses, and interactions" [67]. Of the purposes listed, this research relates to those of recall, replication, collaborative communication, and, most strongly, presentation, which is described as the communication of "how an analysis was conducted, how the findings were determined, or how the data justifies a conclusion" [67].

To provide the provenance of insight and rationale two things need to be recognized. First, a profound understanding of human cognition and reasoning is needed; studying technical sciences alone is not sufficient. This might seem obvious but the truth is that the progress of understanding human reasoning in VA has been slow and misunderstandings still exist. To explain this several authors point to the lack of an accepted theoretical framework through which to investigate and understand this kind of reasoning [1, 52, 53, 60]. Second, VA usually deals with data sets that have both spatial and temporal components, i.e., spatiotemporal data. These kinds of data provide both technically and cognitively unique challenges and possibilities which have been studied extensively within the field of geographical information science [16]. This special nature of spatial and temporal data needs to be acknowledged.

In a recent dissertation Hall [31] proposed a human reasoning framework for spatiotemporal analysis and, on its basis, derived four guidelines for the creation of visualizations that communicate the reasoning of spatiotemporal analysis, i.e., visualizations that provide provenance. The framework is built on the prevailing theories on human cognition and reasoning in the domains of cognitive science and psychology. It was developed with the explicit goal of taking the special nature of spatial and temporal data into consideration. Hall [31] concluded that the guidelines are a good starting point for any project aiming at providing the provenance of insight or rationale in any domain addressing the visual analysis of spatiotemporal data but also points out that the guidelines are quite abstract and that further research is needed to make them more concrete.

This paper reviews available provenance solutions from the literature and discusses them in the light of the guidelines. The aim was to identify a set of key features that are of relevance when providing the provenance of insight and rationale and, on the basis of these features, to produce a new set of guidelines that are more technically oriented. The main research method is a literature review. The starting point was a set of projects that

Tools	References
Aruvi	Shrinivasan & van Wijk [78]
Click2Annotate, part of ManyInsights	Chen et al. [7]
CLIP	Mahyar & Tory [54]
CommentSpace	Willett et al. [85]
ExPlatesJS	Javed & Elmqvist [38]
GeoTime Stories	Eccles et al. [15]
GraphTrail	Dunne et al. [13]
HARVEST	Gotz & Zhou [29], Gotz et al. [28]
Human Terrain Visual Analytics Prototype (HTVAP)	Walker et al. [84]
Note Taking Environment (NTE)	Sacha et al. [71]
Sandbox, part of nSpace	Wright et al. [86]
SchemaLine	Nguyen et al. [57]
sense.us	Heer et al. [36]
SensePath	Nguyen et al. [58]
Tablet, part of Jigsaw	Liu et al. [51], Görg et al. [27]
The Provenance and Annotation Prototype[1] (P&A)	Groth & Streefkerk [30]
The Scalable Reasoning System (SRS)	Pike et al. [62, 63]
VisTrails	Bavoil et al. [3], Callahan et al. [6]

[1] The tool does not have a name in the original paper by Groth and Streefkerk (2006). To facilitate the readability of this paper the authors named it as they found best.

Table 1: The eighteen tools that made it to the paper, and their abbreviations when applicable, listed in alphabetical order.

emphasize the provenance of insight or rationale. These projects are listed by Ragan et al. in Table 2 on page 37 in [67]. The mode of the review can be described as sprawling: when interesting references were found in a paper these were added to the review. The focus was on identifying realized tools and software. Table 1 provides a list of the tools that made it to this paper and their references. Henceforth in this paper the tools are only referenced by their names. Note that this paper does not assess these tools in any way and that in many cases the tools were developed with a completely different goal than providing the provenance of insight or rationale. Rather, this paper discusses how specific features of these tools fit into the provenance context studied in this paper. For the same reason, this paper does not go into the evaluations and user studies done on these tools as the findings would not necessarily be transferrable to the context of this paper.

The rest of this paper is structured as follows. Section 2 gives a short theoretical foundation. The framework and guidelines in Hall [31] have not been previously published in the easily accessible format of a paper and are therefore presented here in Sections 3 and 4 respectively. A review of the available solutions can be found in Section 5. The paper ends by presenting a new and practical set of guidelines in Section 6. The new guidelines are discussed in Section 7, which also includes conclusions.

2 Theoretical foundations

This research takes the distributed cognition framework [1, 75] as a theoretical framework for the study of interaction with visualizations and computers, as suggested by many re-

searchers [14, 37, 52, 53]. Through reasoning we derive inferences or conclusions from a set of premises [74]. There are two main types of reasoning, deduction and induction [40, 42]. Inductive reasoning, which uses knowledge not given in the premises to go beyond the given information, is the cornerstone of human reasoning [2]. When reasoning inductively we use the meaning of premises and our knowledge to construct mental models that are iconic, i.e., based on the concept of resemblance [40–42]. They are, however, not images [26, 46], but rather abstract topological structures [26, 46, 68].

Spatial reasoning appears to have a primary role in our cognition [4, 20]. It is one of the most common and basic forms of human intelligence [17, 49, 65] and we use it to reason about other concepts, such as possession, circumstances, and time. Human reasoning about space and time is mainly qualitative in its nature [10, 17, 23, 48, 55, 69]. Qualitative categorical representations are also argued to be important when perceiving an object or scene visually [48].

The subject matter of spatiotemporal analysis can be conceptualized as being composed of spatial, temporal, and thematic components of data [56, 61, 88, 89]. Spatial and temporal data identify locations in space and time in relation to some frame of reference. Any other data are thematic and typically a property that can be sensed, measured, and assigned a qualitative or quantitative value [56]. The task of the spatiotemporal analyst can be interpreted as finding the meaningful relations inside and between these three data components.

3   Human reasoning framework for spatiotemporal analysis

The human reasoning framework for spatiotemporal analysis identifies and relates the different cognitive mechanisms of human cognition to each other. The framework borrows its notation and builds to a large extent on the human cognition framework for information visualization by Patterson et al. [60]. It incorporates the framework of spatiotemporal concepts (the STC framework [31,32]), which classifies spatial and temporal relations into three basic types, distance, direction, and relations of topology, and has a three-level hierarchical structure of spatiotemporal concepts consisting of basic, compound, and labeled concepts. The framework is thus tightly tied to spatiotemporal analysis and its unique nature.

The framework is depicted in Figure 1 and details the flow of information in human reasoning from stimulus to response. It has a vertical structure to reflect the fact that reasoning is a dynamic interplay of bottom-up and top-down processes. The framework only covers the analytical reasoning system as postulated by the dual-systems theory [19,44,81,82], not the intuitive system.

External information enters human reasoning through encoding, which converts the stimuli into a neural representation in human memory. In spatiotemporal analysis space and time and the relations therein are always present in the encoded stimulus. If the stimulus is a map, innumerable spatial relations are implicitly conveyed to the analyst, regardless of what aspect of the data is under analysis. The framework emphasizes this by explicitly listing basic spatial and temporal relations as a major part of the encoded information. Reference information and thematic information are also given special attention as they play a major role in spatiotemporal analysis. The neural representation that results from encoding thematic and reference information needs, however, to be interpreted using, amongst other things, knowledge from long-term memory, before thematic information or reference infor-
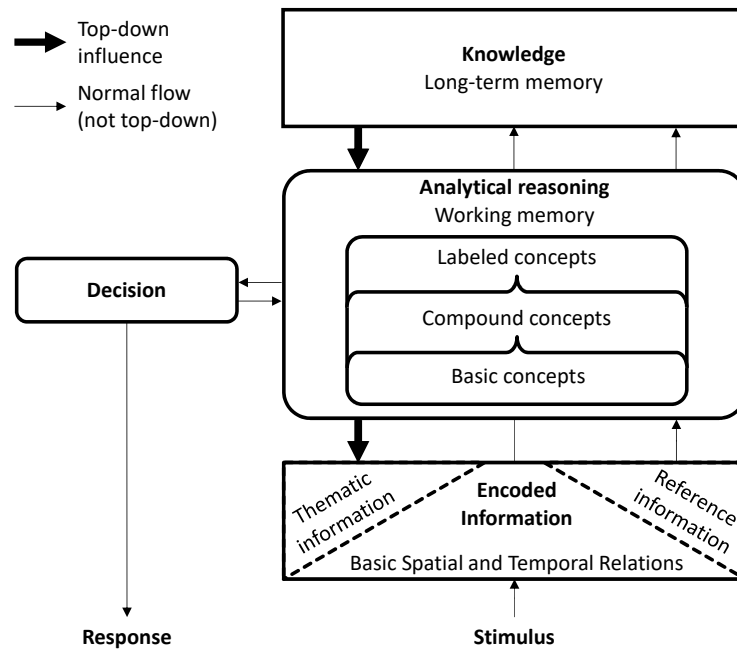
Figure 1: The human reasoning framework for spatiotemporal analysis. The flow of information starts with a stimulus and ends in a decision. Normal processing flow refers here to the traditional view of a one-directional flow of information in human perception. The components depicted in the diagram are explained in the text. The notation is inspired by Patterson et al. [60].

mation can be extracted. In this way thematic information differs from spatial and temporal relations, which are more directly conveyed. On the basis of the goals and demands of the present analysis task, the analyst may direct his attention to certain features of the stimulus, thus exposing the encoding of the top-down influence of attention and working memory. It is also important to remember the influence of exogenous, stimulus-driven attention on the encoding process [18]. These two kinds of attention are not explicitly depicted in Figure 1 but are included in the framework by Patterson et al. [60].

The actual analytical reasoning takes place in working memory. There, the encoded information is iteratively processed as mental models on several levels of abstraction, as illustrated by the STC framework. The basic concepts are distance, direction, and relations of topology and they are the direct equivalences to the basic types of spatiotemporal relations as defined by the STC framework. On the next level there are compound concepts. Contrary to the basic concepts, they need more than two points in space or time to apply and are in a way derived from and divisible into basic concepts. These are concepts such as intensity, size, shape, and distribution. The third level is called labeled concepts. These are basic and compound concepts enriched with thematic information or knowledge not inherent in the data. For example, the distribution of a set of points on a map is a compound concept, while the distribution of the subset of points that have a certain thematic value is a labeled concept. These concepts are processed as mental models or as parts of

mental models in our working memory. See Hall [31] and Hall and Ahonen-Rainio [32] for a more detailed discussion of these concepts.

Spatiotemporal concepts are usually given a qualitative categorical representation when processed in our working memory, which might be complex, or as simple as short or long. This categorization is based on the perceived information and on previous knowledge about this kind of information, depicted through the arrows indicating top-down influence from long-term memory and normal processing flow from encoding in Figure 1. This processing might result in an insight, understood as a changed mental model, i.e., new knowledge, which is then stored in long-term memory and which might result in a decision. The causally structured narratives that we use to make sense of and organize new insights and information [15,21,80] are created in working memory on the basis of previous knowledge, a top-down process, and the perceived information, a bottom-up process.

Reasoning in spatiotemporal analysis is inductive and based on knowledge. This knowledge is situated in long-term memory and is depicted at the top of Figure 1. The knowledge is, at least in part, stored in the form of iconic mental models that have an abstract topological structure that corresponds to the structure of what it represents. Causality is often central in these models. Once one of these mental models is activated and accessible to working memory it influences our encoding processes by guiding our attention. These mental models can also be activated directly by an encoded stimulus, depicted by the arrow from encoding to long-term memory in Figure 1.

The final stage of reasoning is a decision, which might result in a response. Decisions can also guide working memory, e.g., by changing certain goals or eliminating certain choices under consideration. A response might be directed towards the representation of data, resulting in a change of stimulus. This then restarts the cognitive processing, perhaps resulting in a changed mental model in working memory or the activation of a new mental model from long-term memory. This is the iterative manipulation of physical and mental representations of information that leads to insight and new knowledge, as described in the sensemaking process by Pirolli and Card [66] and in the knowledge generation model by Sacha et al. [73].

4 The original guidelines for visualizing reasoning

The guidelines for the creation of visualizations that communicate the reasoning of spatiotemporal analysis are organized under four headings. It is central that the visualizations communicate not only the cognitive outcomes, such as insights and other forms of analytical findings, but the whole cognitive process that led to these outcomes. They are directly derived from the human reasoning framework for spatiotemporal analysis depicted in Figure 1 and each guideline relates to one of the four components in the framework, as indicated.

(1) Decisive tangible pieces of information—relates to encoding

The visualizations need to communicate the decisive tangible pieces of information that were identified in the stimulus and on which the reasoning was based. This is the encoded information that becomes part of our mental models. This communication entails the inclusion of the original physical representation or some simplified representation of it. In this, there needs to be a special focus on the decisive spatial and temporal relations identified

in the data by the analyst. Important thematic information and the frame of reference also need to be taken into consideration at this stage.

(2) Decisive spatiotemporal concepts—relates to reasoning

The visualizations need to communicate the decisive spatiotemporal concepts and their qualitative categorical representations, i.e., the mental models iteratively processed in our working memory on several levels of abstraction. The qualitative categorical representations are needed to communicate how the analyst's differentiation of spatiotemporal relations into normal behavior and anomalies affected the reasoning. If relevant, the origin of this classification should also be communicated, i.e., whether it is based on the perceived information or on previous knowledge about this kind of information.

(3) Cognitive outcomes—relates to decision

The visualizations need to communicate the cognitive outcomes, such as insights and other forms of analytical findings, as well as the narrative that guided the organization of these outcomes into existing knowledge structures. Regarding a specific insight, these visualizations will ideally communicate both the experience of having it as well as the product of that experience. Decisions made also need to be communicated if they are relevant.

(4) Knowledge—relates to knowledge

The visualizations need to communicate the knowledge used. Here knowledge does not only refer to representations of knowledge in long-term memory, partly in the form of mental models, but also to the mental models processed in working memory during the process of reasoning. As these mental models have an abstract topological structure, the preference, referring to the congruence principle of Tversky et al. [83], is that the visualizations should also have an abstract topological structure. When these abstract topological structures are being created, the possible role of causality in these mental models, or chains of mental models, needs to be considered.

The visualizations will take both the bottom-up and top-down aspects of the reasoning into consideration if they are constructed according to these guidelines. In this way they will capture the dynamic interplay between bottom-up and top-down processing that constitutes human reasoning. Please refer to Hall [31] to get an idea of how these guidelines might be applied and in what kind of visualizations that they might result.

5 Review of available solutions

This review is divided into two parts. The first part looks at the solutions related most closely to Guideline 1. This entails ways of communicating the original physical representation used by the analyst in the analysis process, i.e., the stimuli, as well as ways of differentiating in more detail the information identified therein. There is a need for both an overview and detail, so to speak [77], with special attention being paid to decisive spatial and temporal relations. The second part deals with solutions more closely related to the ideas of Guidelines 2, 3, and 4, i.e., ways of communicating the reasoning and knowledge of the analyst as well as the cognitive outcomes of the analysis. This is more a question of explicating the bits and pieces of knowledge and information that were important for a certain analysis outcome and their relation to each other, with special attention being paid to spatiotemporal concepts.

5.1   How can the stimuli be captured?

To communicate the original physical representations and the decisive details therein, there needs to be a way to record or link back to the visualizations used. In general there are two ways of recording these kinds of data: designed systematic capture and manual user capture [35,87]. The former usually aims at capturing all the visualization states, while the latter is about letting the user capture the relevant states. In some cases the delimitation between the visualization states to be automatically captured is clear-cut. Such an example is VisTrails, developed within the domain of scientific visualization, which was one of the first tools to formally propose a designed systematic capture. In other domains, such as VA, which include a lot of interaction with the visualizations, e.g., brushing, filtering, and zooming, the delimitation between different visualization states is fuzzier. In these cases designed systematic capture can be realized at different semantic levels.

To semantically model user analytical behavior many researchers apply the model by Gotz and Zhou [29], shown here in Figure 2. They characterized user behavior into a hierarchical system consisting of tasks, sub-tasks, actions, and events. These four categories describe levels of decreasing semantic richness where activities in one category are combined into sequences that form a single activity on the following level. Events such as mouse clicks and keystrokes and actions such as filtering and sorting are easy to capture but low in semantic content. Sub-tasks and tasks represent more concrete analytical goals. They have a higher semantic content but are difficult to capture without user input. One system that applies this model is the visual analytic system HARVEST. It performs systematic capture on the action level by capturing the user's actions, e.g., query, filter, and sort activities, by bundling events such as mouse clicks and drags into action-level expressions of user activity.


Figure 2: An illustration of the hierarchical model of user analytical behavior by Gotz and Zhou [29].

Tools that apply designed systematic capture usually provide a visual overview of the stored states. The simplest structure is probably a linear trail, such as the "action trail" of HARVEST or the "breadcrumb track" of the Scalable Reasoning System (SRS). More common is a tree structure, where nodes usually represent visualization states and edges the changes applied to a certain state (the parent node) to achieve the new state (the child node). Typical examples of this are VisTrails, Aruvi, and the Provenance and Annotation Prototype (P&A). Figure 3 gives typical examples of a linear trail and a tree structure. In Aruvi [78] the history tree can be ordered by time so that it resembles a timeline and shows the temporal order of the visualization states. Some tools also use a pure timeline view in-

stead of a tree structure, e.g., SensePath. A variation on the tree structure is to use the actual visualization states as nodes. Examples of this can be found in ExPlatesJS and GraphTrail.

There is also the option of not applying any specific structure. In the Human Terrain Visual Analytics Prototype (HTVAP) snapshots representing visualization states are stored without any explicit structure. The user can browse and reorder them on the basis of a set of quantitative measures. To further help analysts handle the potentially large number of snapshots each of them is enriched with a graphical summary that succinctly summarizes the important aspects of the visualization state, including predefined data source uncertainty, the spatial and temporal extent of the data, and the visual variables used.


Figure 3: Two illustrations of typical visual overviews of stored states. The linear structure in the upper row is inspired by the action trail of HARVEST and the tree structure in the bottom row is inspired by the history tree of Aruvi.

These visual overviews of the stored states, especially those applying a tree or timeline structure, also often work as an overview of the interactions made with the system. A way to allow a more detailed understanding of the interaction with the visualization states is to record a screen capture video, a feature introduced by SensePath to support human-computer interaction researchers in their analysis of subjects' browser-based sensemaking process. When a recorded action is selected on a timeline view, there is a browser view that displays the web page where the action was performed and in addition to this a video view that automatically jumps to the corresponding part of the screen capture video. The video can provide additional contextual information about subjects' interaction, such as scrolling and mouse movements. Lipford et al. [50] discuss the use of video for a similar purpose and they recommend that video should not be used without providing additional information such as recorded actions.

Many tools provide some kind of annotation functionality that enables users to record, organize, and communicate findings. These kinds of functionalities can be used to communicate the decisive details of the original physical representations. Maybe the easiest to implement is textual annotation of stored states. This kind of functionality is offered in many tools, such as P&A, Sense.us, and HTVAP. Although expressive, written text needs to be read and can easily remain disconnected from the details of the visualization. Overlaid graphical annotations, on the other hand, are implicitly conveyed and provide a way to easily explicate details in a visualization, such as pointing to or circling around certain details. They are also good for identifying spatial and temporal relations [33], as can be seen in Figure 4. Sense.us [36] allows traditional textual annotations to be enhanced with graphical annotations such as free-form ink, lines, arrows, shapes, and text that can be drawn

over the visualization view. ExPlatesJS [38] provides a similar graphical annotation feature but instead of just single views being annotatable the whole "exploration canvas," where each visualization and data operation that is created is laid out in a node-link structure, is annotatable.


Figure 4: Overlaid graphical annotations (here in red) are good for identifying spatial (and temporal) relations. Map from MapAnt (www.mapant.fi), data from National Land Survey of Finland (www.nls.fi), both provided under the Creative Commons license.

Graphical annotations are of two general types: overlay and data-aware. They are discussed by Heer et al. [36] and Heer and Shneiderman [35]. Overlay annotations are a direct and unconstrained way of interacting with a visualization but they lack an explicit tie to the underlying data. By making annotations data-aware they can persist over operations such as filtering, aggregation, and rotation and thus not be tied to only one particular view of the data. One way to make annotations data-aware is to realize them as selections [34]. GeoTime Stories, a space-time cube tool for the analysis of events, allows data-aware graphical annotation of events by means of curves, arrows, callouts, outlines, and enlargements. These annotations can be re-rendered when the view is updated (e.g., when the space-time cube is rotated).

Textual annotations can also be made data-aware. The ManyInsights tool implements an approach called Click2Annotate that allows the user to annotate interesting aspects of a view by brushing the relevant data (e.g., a set of points), storing it as a finding, and choosing its type from a pre-defined set of types, such as cluster or outlier, corresponding to the sub-task level in the model by Gotz and Zhou [29]. The system then automatically generates an annotation that describes the data and that the user can manually modify or add to. As these annotations are data-aware they can be flagged in all views where the relevant data items can be viewed.

Annotations can be marked with descriptive terms, called keywords or tags, to organize them for future navigation, filtering, or search. Tagging systems are non-hierarchical and inclusive and are often contrasted with hierarchical and exclusive taxonomies [25]. Examples of tools that provide a tagging functionality are ManyInsights, which allows the user

to create their own tags, and CommentSpace, which uses a small, fixed vocabulary of tags (question, hypothesis, to-do, evidence-for, and evidence-against).

5.2   How can the reasoning be captured?

Moving on to the subject matter of Guidelines 2–4, things get more abstract. There is no designed systematic capture that can record the reasoning of the analyst. Analysts need to record this information themselves. What is key here is to provide the possibility of linking the bits and pieces of information and knowledge that the reasoning was based upon together to show how they led the analyst to a specific insight or decision. A popular solution is to use a node-link structure for this purpose, often called reasoning graphs. Typical examples of this are CLIP and Jigsaw's Tablet. Figure 5 provides a typical example of a reasoning graph. These are described as visual thinking spaces that support sense-making through allowing users to record, organize, and share their findings in the form of node-link graphs and notes. Interesting entities, such as people, locations, or events, are added as nodes and relationships between them as links. Aruvi [78] is another example. Compared to CLIP, it provides a larger set of different nodes and connectors that allow the user to create more expressive graphs. A third and more refined example is SRS [62, 63], which provides the ability to record and link evidence, assumptions, hypotheses, and other artifacts using a node-link structure.


Figure 5: A typical and fictional example of a reasoning graph inspired by the one found in Jigsaw's Tablet.

There are also other approaches than the node-link structure. The analytical software environment nSpace provides the user with a flexible and expressive workspace called the Sandbox where users can "write anywhere, group anything, [and] place anything anywhere" [86]. It supports node-link structures but can also handle text snippets and images and puts the focus on the ease of spatial arrangement of these. Another solution is to use a textual narrative, or story, to communicate the reasoning. This more formal and linear approach is used by GeoTime Stories. Instead of a graph or workspace, the main output is an authored interactive report.

Important distinctions for these kinds of systems are if they are stand-alone or integrated systems and if the evidence that the reasoning is based upon is left out of, linked

to, or included in the recorded reasoning. Stand-alone systems are mainly for organizing thought and communicating findings, not for analyzing information. CLIP is a pure standalone tool, designed for recording findings made in other contexts. It does not support either the inclusion of evidence, or linking to it. Many tools, however, integrate functions for analyzing data and recording and organizing findings. This makes it easier to include or link to evidence, such as the visualization states used. In Aruvi the analysis of data and recording of reasoning literally take place side by side. There, a node in the reasoning graph can be linked to the original visualization state that the piece of information in the node is based upon. At the same time Aruvi also displays the location of this visualization state in a history tree of stored states. SRS also applies an integrated approach, with the additional option of including visualization states in the recorded reasoning. The idea is that any feature of a view, or an entire view, can be saved and added to the node-link structure as a node that contains a snapshot of that view. These nodes can be expanded to a full view that consists of an interactive visualization of what the analyst saw.

There are also tools that are not pure stand-alone or integrated systems. Sandbox, which is a part of nSpace, is integrated with a tool for information retrieval and analysis called TRIST [43]. The analyst can add any relevant information from TRIST by dragging it into the Sandbox workspace and saving it as an object. These objects then work as links to the original interactive views in TRIST. It is also possible to include the original visualization states in the Sandbox workspace. There is also drag–and-drop support for information from other sources, such as text editors and web explorers, but these objects do not link back to the original sources. In relation to these sources, Sandbox is a standalone system. The future might be standalone systems that can be plugged into any VA system. In a recent publication Sacha et al. [71] present their current progress in developing a notetaking environment (NTE) that supports importing visualization states from, and recording interactions with, any VA system.

There is an evident difference when it comes to the structures for recording reasoning between tools that are mainly designed to support analysts in organizing their thoughts, and tools that are designed to support the analyst in presenting results. The former emphasize divergent and expanding structures, which are better suited for exploratory work, while the latter focus on more linear structures, which are natural to use when it is already known how things are connected and the causal relations between input and output. Typical examples of tools that are designed to support analysts in organizing their thought are CLIP, Jigsaw's Tablet, and Sandbox.

Two interesting and dissimilar examples of tools more focused on enabling the presentation of results are GeoTime Stories and HTVAP. As already mentioned, GeoTime Stories provides the analyst with the ability to author narrative reports, or stories, that link to decisive views of events in time and space. The story is a top-down approach that promotes the sharing of observations and understanding of complex phenomena. They can be presented to high-level decision makers without additional explanations. A narrative approach is also used in HTVAP. Its narrative is, instead of a written report, a graph-node structure consisting of a set of ordered and annotated snapshots of visualization states that communicate "facts, assumptions, assertions and interpretations" [84]. Figure 6 illustrates the HTVAP narrative approach. This function was developed with a special focus on providing support for analysts giving short verbal briefings to superiors. In these tools the story is used as a vessel for communication and understanding in a similar way as in narrative visualization [24, 47, 76].

Figure 6: An illustration of the graph-node structured narrative of HTVAP. It consists of a set of ordered and annotated snapshots of visualization states. This fictional example uses lorem ipsum instead of real text.

Trust is an important aspect of analysis. Sacha et al. [72] describe analysts' trust as depending on the extent of their "awareness of the underlying uncertainties generated on the system side." Xu et al. [87] list three contexts of trust in analysis: "trust in the data, trust in the process, and trust in the result." Several tools provide support for analysts to record their trust in findings. In SRS analysts can describe their trust in evidence and the degree to which each node in the reasoning graph supports their hypotheses or conclusions. For this purpose each node and link in the graph is equipped with an adjustable bar graph that represents its current belief state. SRS propagates the trust and support scores across the graph and produces likelihood estimates for hypotheses and conclusions automatically. NTE provides similar functions by enabling the analyst to apply trust ratings to imported visualization states and notes. Figure 7 illustrates how ratings are aggregated into an evidence bar in NTE that describes if a hypothesis is rejected or accepted. Sandbox also provides these kinds of features.

Whether the tool supports designed systematic capture of the visualization states used or only relies on manual user capture has a considerable effect on the available design solutions when it comes to how visualization states are added to a reasoning graph. With tools that support designed systematic capture it is possible to let the user choose from the set of saved visualization states. This kind of design enables reasoning graphs to be created post-analysis because the decisive visualization states can still be accessed. That analysts explicitly need to save visualization states to use them in a reasoning graph might be problematic as it requires the analysts to be aware of their importance immediately, although this might only become clear later in the analysis. HTVAP presents a clever solution for how visualization states can be added to a reasoning graph from a set of stored states.
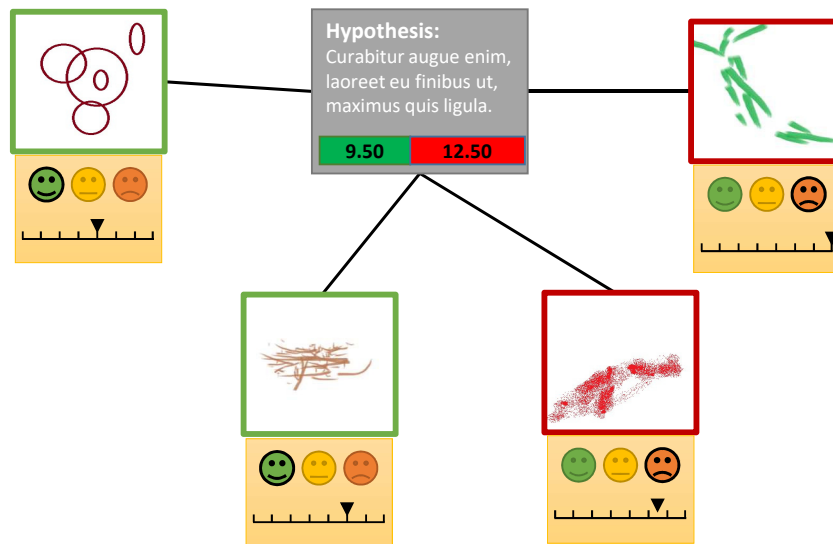
Figure 7: Simplified illustration that describes how ratings are aggregated into an evidence bar in NTE that describes if a hypothesis is rejected or accepted. This fictional example uses lorem ipsum instead of real text.

There are two more design solutions related to the topic of Guideline 2–3 that deserve to be mentioned. First, many tools provide some kind of tagging system or taxonomy to describe the role of nodes in a reasoning graph. In SRS reasoning artifacts can be tagged as evidence, assumption, spatial or temporal pattern, causal relationship, or hypothesis. In NTE verifying, falsifying, and neutral tags for pieces of evidence are used. In Figure 7 these tags are represented by the happy, sad, and neutral faces. Second, some tools also provide timeline visualizations alongside the reasoning graphs. For example, CLIP allows nodes in the reasoning graph to be linked to and visualized on a timeline. Related to this, SchemaLine presents a solution for how entities on a timeline can be grouped together visually to facilitate analysis.

5.3   Synthesis: Key features

This section presents a synthesis of the review of available provenance solutions that was performed. A set of key features for an ideal provenance tool is identified. This is done on the basis of the guidelines and with the intention to provide the provenance of insight and rationale, with the main purpose being communication, according to the framework by Ragan et al. [67].

One of the first questions when designing a tool for recording and communicating reasoning is what kind of format to use. The recommendation of this paper is to use a graph with a node-link structure, i.e., a reasoning graph. This is an abstract topological structure as recommended in Guideline 4 and it is also one of the most popular solutions, on the evidence of the review. It furthermore enables the user to easily record causal links between pieces of information and knowledge. Optionally, support for other kinds of or-

ganizational metaphors, such as those listed by Robinson [70], could be provided. What is of the uttermost importance is that these reasoning graphs include or link to the original physical representations that the reasoning was based upon, the bottom-up aspect of reasoning. This can, e.g., be achieved through including the actual visualization states in the graph (as in NTE, Figure 7), making them available in an adjacent view (as in Aruvi), or by including some simplified representations of them that link to the original visualization (as in HTVAP, Figure 6). If the original visualizations were interactive, then the visualizations accessed through the reasoning graph should also be interactive.

These reasoning graphs should be created post-analysis in order not to interrupt the actual analysis. This is not to say that graphing capabilities do not need to be provided to analysts during analysis, only that for reasoning graphs to be effective for presenting the results and outcomes of analysis they will necessarily be different from graphs for organizing thought. In the kind of deterministic situation that prevails once the outcomes of analysis and the goal of a presentation are clear it is easier to create effective graphs that only include the relevant information. These graphs will thus be linear rather than divergent and clearly communicate the narrative of the analysis (see Guideline 3 and Figure 6).

A prerequisite for a smooth post-analysis creation of graphs is that there is a designed systematic capture of the visualization states used. Manual user capture at this stage would simply be too laborious and difficult. The saved visualization states need to be easily accessible and browsable by the analyst. This can be facilitated, e.g., by presenting the stored states as snapshots that include graphical summaries and allow the analyst to order them by qualitative measures (as in HTVAP) or by providing a tree-like visualization (illustrated in Figure 3) of the interaction process that created the states (as in VisTrails), possibly orderable by time (as in Aruvi).

Providing annotation abilities is of equal importance to including the original visualizations. Through textual annotation analysts are able to describe their reasoning and communicate their narrative. It also enables them to pinpoint the role of thematic information, the frame of reference, spatiotemporal concepts, their own knowledge, and other aspects mentioned in the guidelines. Providing graphical annotation is equally important as it enables the analysts to specify the decisive tangible pieces of information identified in visualizations, especially spatial and temporal relations (see Guideline 1 and Figure 4). Data-aware annotations have many strengths, especially in interactive interfaces where they can stay visible although the view changes, and are to be preferred if realizable. Overlay graphical annotations and data-loose textual annotations are, however, good alternatives in most situations. Annotation capabilities should be available both during analysis and post-analysis when creating reasoning graphs. Annotations created during analysis should be automatically captured and linked to the corresponding visualization view. Annotated visualization states should be easily discernible when browsing saved states to facilitate access to these supposedly important states.

Tagging systems or taxonomies can be used to enhance the information value of both annotated visualization states (as in CommentSpace) and entities in reasoning graphs (as in SRS). As for stored states and annotations, tags can also facilitate browsing and reasoning (the faces in Figure 7 are one example of the use of tags). There are many taxonomies and similar systems available that, if not usable as such, could work as a starting point. Some examples are the fact taxonomy by Chen et al. [8], the classification of visualization insights by Choe et al. [9], the classification of movement patterns by Dodge et al. [12], and the STC framework [31, 32]. A tagging system based on the STC framework could be useful to

conform to Guideline 2. In a similar fashion entities in reasoning graphs and annotations can be tagged with the degree of trust that the analyst has in them.

A big challenge for analysts creating reasoning graphs post-analysis is how to remember how the reasoning progressed and which visualization states were actually decisive. A solution is to use cued recall, which can enable analysts to "remember more of their rationale and decision points" [50]. Cues can be provided to the analyst in the form of interaction logs [39], screen capture video (as in SensePath), eye-tracking recordings [30, 50, 67], and think-aloud recordings (used for annotating visualizations in P&A).

6 New practical guidelines

As a summary four new guidelines are provided.

(1) Use node-link reasoning graphs as a communication format
The graphs should be created post-analysis so as not to infer with the analysis. They should furthermore have a clear communication objective, i.e., a certain decision or outcome that they aim at communicating, and a linear structure that clearly communicates the narrative and important causal links. They can be enhanced with tags indicating the level of trust and other relevant information.

(2) Include the original physical representations
The reasoning graph needs to include or provide access to the decisive visualization states and the other physical representations that the reasoning was based upon. The smoother the access, the better. This requires there to be a designed systematic capture during analysis and the captured states to be easily browsable and accessible post-analysis. In this way the spatial and temporal relations identified in the data by the analyst are also included. This is especially important when analyses have included maps, timelines and similar spatiotemporal representations that include substantial amounts of possibly important relations.

(3) Provide annotation capabilities
Textual and graphical annotation capabilities should be provided both during analysis and post-analysis. Annotations should be automatically captured and linked to the corresponding visualization states. Tags, possibly identifying spatiotemporal concepts and relations, can be used as one type of annotation. Graphical annotations are, especially in maps, timelines and similar spatiotemporal representations, an effective and important way to pinpoint the decisive spatial and temporal relations identified in the data by the analyst.

(4) Offer cued recall
When reasoning graphs are created post-analysis analysts need help to remember their rationale and decision points. This kind of help can be achieved through cued recall.

7 Discussion and conclusions

While the original four guidelines were built around theories from cognitive science and psychology, and as such are quite abstract, this study created a set of complementary guidelines that are instead centered on tool features. The new set of guidelines is thus more concrete, practical, and comprehensible than the original ones. They do not, however, act as

substitutes for the original guidelines but work as a good complement to them. Together they provide both a theoretical and practical approach to the problem of providing the provenance of insight and rationale.

The idea of providing these kinds of guidelines represents a new approach in provenance research. Up to now, research on provenance has mostly been practically oriented and focused around a certain analysis tool or problem. These two sets of guidelines combine theoretically oriented findings from cognitive science with technical solutions from VA and look at provenance on a general level. These guidelines can work as a starting point for any project aiming at providing the provenance of insight or rationale in any domain addressing the visual analysis of spatiotemporal data.

Although these guidelines have been developed for a spatiotemporal application domain they are probably transferable to other domains with no, or only minor, changes. The main contribution of the spatiotemporal focus is the advice to include the original physical representation and the advice to enable graphical annotations. The original four guidelines are more centered around the spatiotemporal application domain and thus complement these new guidelines from that point of view as well.

Hall [31] did not address how the actual recording of provenance should happen, only what aspects of the reasoning need to be communicated. This study addresses this deficit by explicitly discussing what kinds of features are needed in a tool that implements the original four guidelines. By recommending that the graphs be created post-analysis this paper suggests that the recording of analysts' reasoning cannot be achieved during analysis without fragmenting the actual analysis process too much. To make up for the inconvenience and impracticality of creating reasoning graphs post-analysis this paper suggests the use of cued recall, perhaps in the form of interaction logs and screen capture video.

The next research step would be to create a provenance tool that implements these guidelines and perform user testing with it to study how it works in practice. Testing and refining the ease of use of the tool for recording reasoning would be as important as testing if it provides provenance. Here a successful implementation of the post-analysis graph creation and browsing of visualization states is key. User testing is also needed to figure out how to best provide cued recall. Implementing a tagging system based on the STC framework is also a direction for future research.

For a discussion on the human reasoning framework for spatiotemporal analysis and the original four guidelines, see Hall [31].

Acknowledgments

References

[1] ARIAS-HERNANDEZ, R., GREEN, T. M., AND FISHER, B. From cognitive amplifiers to cognitive prostheses: Understandings of the material basis of cognition in visual analytics. *Interdisciplinary Science Reviews 37*, 1 (March 2012), 4–18. doi:10.1179/0308018812z.0000000001.

[2] ARTHUR, W. B. Inductive reasoning and bounded rationality. *The American Economic Review 84*, 2 (1994), 406–411.

[3] BAVOIL, L., CALLAHAN, S. P., CROSSNO, P. J., FREIRE, J., SCHEIDEGGER, C. E., SILVA, C. T., AND VO, H. T. VisTrails: Enabling interactive multiple-view visualizations. In *VIS 05. IEEE Visualization, 2005.* (October 2005), IEEE, pp. 135–142. doi:10.1109/visual.2005.1532788.

[4] BORODITSKY, L. Metaphoric structuring: Understanding time through spatial metaphors. *Cognition 75*, 1 (April 2000), 1–28. doi:10.1016/s0010-0277(99)00073-6.

[5] BOSE, R., AND FREW, J. Lineage retrieval for scientific data processing: A survey. *ACM Computing Surveys 37*, 1 (March 2005), 1–28. doi:10.1145/1057977.1057978.

[6] CALLAHAN, S. P., FREIRE, J., SANTOS, E., SCHEIDEGGER, C. E., SILVA, C. T., AND VO, H. T. VisTrails: Visualization meets data management. In *Proc. 2006 ACM SIGMOD International Conference on Management of Data* (June 2006), S. Chaudhuri, V. Hristidis, and N. Polyzotis, Eds., ACM, pp. 745–747. doi:10.1145/1142473.1142574.

[7] CHEN, Y., BARLOWE, S., AND YANG, J. Click2annotate: Automated insight externalization with rich semantics. In *2010 IEEE Symposium on Visual Analytics Science and Technology* (October 2010), IEEE, pp. 155–162. doi:10.1109/vast.2010.5652885.

[8] CHEN, Y., YANG, J., AND RIBARSKY, W. Toward effective insight management in visual analytics systems. In *2009 IEEE Pacific Visualization Symposium* (Piscataway, NJ, April 2009), IEEE, pp. 49–56. doi:10.1109/pacificvis.2009.4906837.

[9] CHOE, E. K., LEE, B., AND SCHRAEFEL, M. C. Characterizing visualization insights from quantified selfers' personal data presentations. *IEEE Computer Graphics and Applications 35*, 4 (July 2015), 28–37. doi:10.1109/mcg.2015.51.

[10] COHN, A. G., AND HAZARIKA, S. M. Qualitative spatial representation and reasoning: An overview. *Fundamenta Informaticae 46*, 1 (2001), 1–29.

[11] DAVIDSON, S. B., AND FREIRE, J. Provenance and scientific workflows: Challenges and opportunities. In *Proc. 2008 ACM SIGMOD International Conference on Management of Data* (June 2008), L. Lakshmanan, R. Ng, and D. Shasha, Eds., ACM Press, pp. 1345–1350. doi:10.1145/1376616.1376772.

[12] DODGE, S., WEIBEL, R., AND LAUTENSCHÜTZ, A.-K. Towards a taxonomy of movement patterns. *Information Visualization 7*, 3–4 (September 2008), 240–252. doi:10.1057/palgrave.ivs.9500182.

[13] DUNNE, C., RICHE, N. H., LEE, B., METOYER, R., AND ROBERTSON, G. GraphTrail: Analyzing large multivariate, heterogeneous networks while supporting exploration history. In *Proc. SIGCHI Conference on Human Factors in Computing Systems* (2012), ACM Press, pp. 1663–1672. doi:10.1145/2207676.2208293.

[14] DYKES, J., MACEACHREN, A. M., AND KRAAK, M.-J. Advancing geovisualization. In *Exploring Geovisualization*, J. Dykes, A. M. MacEachren, and M.-J. Kraak, Eds. Elsevier, 2005, pp. 691–703. doi:10.1016/b978-008044531-1/50454-1.

[15] ECCLES, R., KAPLER, T., HARPER, R., AND WRIGHT, W. Stories in GeoTime. *Information Visualization 7*, 1 (March 2008), 3–17. doi:10.1057/palgrave.ivs.9500173.

[16] EGENHOFER, M. J., CLARKE, K. C., GAO, S., QUESNOT, T., FRANKLIN, W. R., YUAN, M., AND COLEMAN, D. Contributions of GIScience over the past twenty years. In *Advancing Geographic Information Science*, H. Onsrud and W. Kuhn, Eds. GSDI Association Press, 2016, pp. 9–34.

[17] EGENHOFER, M. J., AND MARK, D. M. Naive geography. In *Spatial Information Theory: A Theoretical Basis for GIS*, A. U. Frank and W. Kuhn, Eds., vol. 988 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 1995, pp. 1–15. doi:10.1007/3-540-60392-1_1.

[18] EGETH, H. E., AND YANTIS, S. Visual attention: Control, representation, and time course. *Annual Review of Psychology 48*, 1 (February 1997), 269–297. doi:10.1146/annurev.psych.48.1.269.

[19] EVANS, J. Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology 59*, 1 (January 2008), 255–278. doi:10.1146/annurev.psych.59.103006.093629.

[20] FILIPOVIC, L., AND JASZCZOLT, K. *Space and time in languages and cultures: Language, culture, and cognition*. John Benjamins Pub. Co, Amsterdam, June 2012. doi:10.1075/hcp.37.

[21] FISHER, W. R. Narration as a human communication paradigm: The case of public moral argument. *Communications Monographs 51*, 1 (March 1984), 1–22. doi:10.1080/03637758409390180.

[22] FREIRE, J., KOOP, D., SANTOS, E., AND SILVA, C. T. Provenance for computational tasks: A survey. *Computing in Science & Engineering 10*, 3 (May 2008), 11–21. doi:10.1109/mcse.2008.79.

[23] GALTON, A. Spatial and temporal knowledge representation. *Earth Science Informatics 2*, 3 (May 2009), 169–187. doi:10.1007/s12145-009-0027-6.

[24] GERSHON, N., AND PAGE, W. What storytelling can do for information visualization. *Communications of the ACM 44*, 8 (August 2001), 31–37. doi:10.1145/381641.381653.

[25] GOLDER, S. A., AND HUBERMAN, B. A. Usage patterns of collaborative tagging systems. *Journal of Information Science 32*, 2 (April 2006), 198–208. doi:10.1177/0165551506062337.

[26] GOODWIN, G. P., AND JOHNSON-LAIRD, P. N. Reasoning about relations. *Psychological review 112*, 2 (2005), 468–493. doi:10.1037/0033-295x.112.2.468.

[27] GÖRG, C., LIU, Z., AND STASKO, J. Reflections on the evolution of the Jigsaw visual analytics system. *Information Visualization 13*, 4 (October 2014), 336–345. doi:10.1177/1473871613495674.

[28] GOTZ, D., WHEN, Z., LU, J., KISSA, P., CAO, N., QIAN, W. H., LIU, S. X., AND ZHOU, M. X. HARVEST: An intelligent visual analytic tool for the masses. In *Proc. First International Workshop on Intelligent Visual Interfaces for Text Analysis* (February 2010), S. Liu, M. X. Zhou, G. Carenini, and H. Qu, Eds., ACM Press, pp. 1–4. doi:10.1145/2002353.2002355.

[29] GOTZ, D., AND ZHOU, M. X. Characterizing users' visual analytic activity for insight provenance. *Information Visualization 8*, 1 (January 2009), 42–55. doi:10.1057/ivs.2008.31.

[30] GROTH, D. P., AND STREEFKERK, K. Provenance and annotation for visual exploration systems. *IEEE Transactions on Visualization and Computer Graphics 12*, 6 (November 2006), 1500–1510. doi:10.1109/tvcg.2006.101.

[31] HALL, A. *Reasoning in Spatio-Temporal Analysis – Theory, Provenance, and Applications*. PhD thesis, Aalto University, 2016. http://urn.fi/URN:ISBN:978-952-60-7049-0.

[32] HALL, A., AND AHONEN-RAINIO, P. Analysis of basic relations within insights of spatio-temporal analysis. In *Modern Trends in Cartography, Selected Papers of CARTO-CON 2014*, A. Vondrakova, J. Brus, and V. Vozenilek, Eds., Lecture Notes in Geoinformation and Cartography. Springer International Publishing, September 2014, pp. 409–423. doi:10.1007/978-3-319-07926-4_31.

[33] HALL, A., AHONEN-RAINIO, P., AND VIRRANTAUS, K. Knowledge and reasoning in spatial analysis. *Transactions in GIS 18*, 3 (August 2013), 464–476. doi:10.1111/tgis.12049.

[34] HEER, J., AGRAWALA, M., AND WILLETT, W. Generalized selection via interactive query relaxation. In *Proc. SIGCHI Conference on Human Factors in Computing Systems* (April 2008), ACM Press, pp. 959–968. doi:10.1145/1357054.1357203.

[35] HEER, J., AND SHNEIDERMAN, B. Interactive dynamics for visual analysis. *Queue 10*, 2 (February 2012), 30:30–30:55. doi:10.1145/2133416.2146416.

[36] HEER, J., VIÉGAS, F. B., AND WATTENBERG, M. Voyagers and voyeurs: Supporting asynchronous collaborative visualization. *Communications of the ACM 52*, 1 (January 2009), 87–97. doi:10.1145/1435417.1435439.

[37] HOLLAN, J., HUTCHINS, E., AND KIRSH, D. Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI) 7*, 2 (June 2000), 174–196. doi:10.1145/353485.353487.

[38] JAVED, W., AND ELMQVIST, N. ExPlates: Spatializing interactive analysis to scaffold visual exploration. *Computer Graphics Forum 32*, 3pt4 (June 2013), 441–450. doi:10.1111/cgf.12131.

[39] JEONG, D. H., DOU, W., LIPFORD, H. R., STUKES, F., CHANG, R., AND RIBARSKY, W. Evaluating the relationship between user interaction and financial visual analysis. In *2008 IEEE Symposium on Visual Analytics Science and Technology* (October 2008), IEEE, pp. 83–90. doi:10.1109/vast.2008.4677360.

[40] JOHNSON-LAIRD, P. N. *How We Reason*. Oxford University Press, October 2006. doi:10.1093/acprof:oso/9780199551330.001.0001.

[41] JOHNSON-LAIRD, P. N. Mental models and human reasoning. *Proc. National Academy of Sciences 107*, 43 (October 2010), 18243–18250. doi:10.1073/pnas.1012933107.

[42] JOHNSON-LAIRD, P. N., AND KHEMLANI, S. S. Toward a unified theory of reasoning. *Psychology of Learning and Motivation 59* (2013), 1–42. doi:10.1016/b978-0-12-407187-2.00001-0.

[43] JONKER, D., WRIGHT, W., SCHROH, D., PROULX, P., AND CORT, B. Information triage with TRIST. In *2005 International Conference on Intelligence Analysis* (May 2005), pp. 2–4.

[44] KAHNEMAN, D. *Thinking, fast and slow*. Macmillan, 2011.

[45] KEIM, D., ANDRIENKO, G., FEKETE, J.-D., GÖRG, C., KOHLHAMMER, J., AND MELANÇON, G. Visual analytics: Definition, process, and challenges. In *Information Visualization: Human-Centered Issues and Perspectives*, A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, Eds., Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2008, pp. 154–175. doi:10.1007/978-3-540-70956-5_7.

[46] KNAUFF, M. A neuro-cognitive theory of deductive relational reasoning with mental models and visual images. *Spatial Cognition & Computation 9*, 2 (May 2009), 109–137. doi:10.1080/13875860902887605.

[47] KOSARA, R., AND MACKINLAY, J. Storytelling: The next step for visualization. *Computer 46*, 5 (May 2013), 44–50. doi:10.1109/mc.2013.36.

[48] KOSSLYN, S. M., CHABRIS, C. F., MARSOLEK, C. J., AND KOENIG, O. Categorical versus coordinate spatial relations: Computational analyses and computer simulations. *Journal of Experimental Psychology: Human Perception and Performance 18*, 2 (1992), 562–577. doi:10.1037/0096-1523.18.2.562.

[49] LEVINSON, S. C. *Space in language and cognition: Explorations in cognitive diversity*, vol. 5 of *Language Culture and Cognition*. Cambridge University Press, 2003. doi:10.1017/cbo9780511613609.

[50] LIPFORD, H. R., STUKES, F., DOU, W., HAWKINS, M. E., AND CHANG, R. Helping users recall their reasoning process. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on* (October 2010), A. MacEachren and S. Miksch, Eds., IEEE, pp. 187–194. doi:10.1109/vast.2010.5653598.

[51] LIU, Z., GÖRG, C., KIHM, J., LEE, H., CHOO, J., PARK, H., AND STASKO, J. Data ingestion and evidence marshalling in Jigsaw. In *2010 IEEE Symposium on Visual Analytics Science and Technology* (October 2010), IEEE, pp. 271–272. doi:10.1109/vast.2010.5653042.

[52] LIU, Z., NERSESSIAN, N., AND STASKO, J. Distributed cognition as a theoretical framework for information visualization. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (November 2008), 1173–1180. doi:10.1109/tvcg.2008.121.

[53] MacEachren, A. M. Distributed cognition: A conceptual framework for understanding map-based reasoning. In *Proceeding of the 27th International Cartographic Conference* (August 2015), International Cartographic Association. http://icaci.org/files/documents/ICC_proceedings/ICC2015/papers/2/940.html.

[54] Mahyar, N., and Tory, M. Supporting communication and coordination in collaborative sensemaking. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (December 2014), 1633–1642. doi:10.1109/tvcg.2014.2346573.

[55] Mark, D. M. Spatial representation: A cognitive view. In *Geographical Information Systems: Principles, Techniques, Management and Applications*, P. A. Longley, M. F. Goodchild, D. J. Maguire, and D. W. Rhind, Eds., 2 ed., vol. 1. Wiley, 2005, pp. 81–89.

[56] Mennis, J. L., Peuquet, D. J., and Qian, L. A conceptual framework for incorporating cognitive principles into geographical database representation. *International Journal of Geographical Information Science 14*, 6 (September 2000), 501–520. doi:10.1080/136588100415710.

[57] Nguyen, P. H., Xu, K., Walker, R., and Wong, B. L. W. SchemaLine: Timeline visualization for sensemaking. In *2014 18th International Conference on Information Visualisation* (July 2014), IEEE, pp. 225–233. doi:10.1109/iv.2014.14.

[58] Nguyen, P. H., Xu, K., Wheat, A., Wong, B. L. W., Attfield, S., and Fields, B. SensePath: Understanding the sensemaking process through analytic provenance. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (January 2016), 41–50. doi:10.1109/tvcg.2015.2467611.

[59] North, C., Chang, R., Endert, A., Dou, W., May, R., Pike, B., and Fink, G. Analytic provenance: Process interaction insight. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems* (2011), ACM Press, pp. 33–36. doi:10.1145/1979742.1979570.

[60] Patterson, R. E., Blaha, L. M., Grinstein, G. G., Liggett, K. K., Kaveney, D. E., Sheldon, K. C., Havig, P. R., and Moore, J. A. A human cognition framework for information visualization. *Computers & Graphics 42* (August 2014), 42–58. doi:10.1016/j.cag.2014.03.002.

[61] Peuquet, D. J. It's about time: A conceptual framework for the representation of temporal dynamics in geographic information systems. *Annals of the Association of American Geographers 84*, 3 (September 1994), 441–461. doi:10.1111/j.1467-8306.1994.tb01869.x.

[62] Pike, W. A., Bruce, J., Baddeley, B., Best, D., Franklin, L., May, R., Rice, D. M., Riensche, R., and Younkin, K. The scalable reasoning system: Lightweight visualization for distributed analytics. *Information Visualization 8*, 1 (October 2008), 71–84. doi:10.1109/vast.2008.4677366.

[63] Pike, W. A., May, R., Baddeley, B., Riensche, R., Bruce, J., and Younkin, K. Scalable visual reasoning: Supporting collaboration through distributed analysis. In *2007 International Symposium on Collaborative Technologies and Systems* (May 2007), IEEE, pp. 24–32. doi:10.1109/cts.2007.4621734.

[64] PIKE, W. A., STASKO, J., CHANG, R., AND O'CONNELL, T. A. The science of interaction. *Information Visualization 8*, 4 (January 2009), 263–274. doi:10.1057/ivs.2009.22.

[65] PINKER, S. *How the Mind Works*. W. W. Norton, New York, 2009.

[66] PIROLLI, P., AND CARD, S. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proc. International Conference on Intelligence Analysis* (2005), V. A. MacLean, Ed., Mitre McLean, VA.

[67] RAGAN, E. D., ENDERT, A., SANYAL, J., AND CHEN, J. Characterizing provenance in visualization and data analysis: An organizational framework of provenance types and purposes. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (January 2016), 31–40. doi:10.1109/tvcg.2015.2467551.

[68] REISBERG, D., AND HEUER, F. Visuospatial images. In *The Cambridge Handbook of Visuospatial Thinking*, P. Shah and A. Miyake, Eds. Cambridge University Press, 2005, pp. 35–80. doi:10.1017/cbo9780511610448.003.

[69] RENZ, J., AND NEBEL, B. Qualitative spatial reasoning using constraint calculi. In *Handbook of Spatial Logics*, A. Marco, I. E. Pratt-Hartmann, and J. F. A. K. van Benthem, Eds. Springer Netherlands, 2007, pp. 161–215. doi:10.1007/978-1-4020-5587-4_4.

[70] ROBINSON, A. C. Supporting synthesis in geovisualization. *International Journal of Geographical Information Science 25*, 2 (March 2011), 211–227. doi:10.1080/13658810903430916.

[71] SACHA, D., BOESECKE, I., FUCHS, J., AND KEIM, D. A. Analytic behavior and trust building in visual analytics. In *EuroVis 2016 - Short Papers* (2016), E. Bertini, N. Elmqvist, and T. Wischgoll, Eds., The Eurographics Association, pp. 143–147. doi:10.2312/eurovisshort.2016117.

[72] SACHA, D., SENARATNE, H., KWON, B. C., ELLIS, G., AND KEIM, D. A. The role of uncertainty, awareness, and trust in visual analytics. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (January 2016), 240–249. doi:10.1109/tvcg.2015.2467591.

[73] SACHA, D., STOFFEL, A., STOFFEL, F., KWON, B. C., ELLIS, G., AND KEIM, D. Knowledge generation model for visual analytics. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (December 2014), 1604–1613. doi:10.1109/tvcg.2014.2346481.

[74] SAMARAPUNGAVAN, A. Reasoning. https://web.archive.org/web/20161214194356/http://www.education.com/reference/article/reasoning, 2009. Last Accessed December 14, 2016.

[75] SCAIFE, M., AND ROGERS, Y. External cognition: How do graphical representations work? *International Journal of Human-Computer Studies 45*, 2 (August 1996), 185–213. doi:10.1006/ijhc.1996.0048.

[76] SEGEL, E., AND HEER, J. Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (November 2010), 1139–1148. doi:10.1109/tvcg.2010.179.

[77] SHNEIDERMAN, B. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. 1996 IEEE Symposium on Visual Languages* (1996), IEEE, pp. 336–343. doi:10.1109/vl.1996.545307.

[78] SHRINIVASAN, Y. B., AND VAN WIJK, J. J. Supporting the analytical reasoning process in information visualization. In *Proc. SIGCHI Conference on Human Factors in Computing Systems* (2008), ACM Press, pp. 1237–1246. doi:10.1145/1357054.1357247.

[79] SIMMHAN, Y. L., PLALE, B., AND GANNON, D. A survey of data provenance in e-science. *ACM SIGMOD Record 34*, 3 (September 2005), 31–36. doi:10.1145/1084805.1084812.

[80] SLOMAN, S. *Causal models: How people think about the world and its alternatives*. Oxford University Press, 2009.

[81] SLOMAN, S. A. The empirical case for two systems of reasoning. *Psychological Bulletin 119*, 1 (1996), 3–22. doi:10.1037/0033-2909.119.1.3.

[82] STANOVICH, K. E. *The robot's rebellion: Finding meaning in the age of Darwin*. University of Chicago Press, 2004. doi:10.7208/chicago/9780226771199.001.0001.

[83] TVERSKY, B., MORRISON, J. B., AND BETRANCOURT, M. Animation: Can it facilitate? *International Journal of Human-Computer Studies 57*, 4 (October 2002), 247–262. doi:10.1006/ijhc.2002.1017.

[84] WALKER, R., SLINGSBY, A., DYKES, J., XU, K., WOOD, J., NGUYEN, P. H., STEPHENS, D., WONG, B. L. W., AND ZHENG, Y. An extensible framework for provenance in human terrain visual analytics. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (December 2013), 2139–2148. doi:10.1109/tvcg.2013.132.

[85] WILLETT, W., HEER, J., HELLERSTEIN, J., AND AGRAWALA, M. CommentSpace: Structured support for collaborative visual analysis. In *Proc. SIGCHI Conference on Human Factors in Computing Systems* (April 2011), W. E. Mackay, S. Brewster, and S. Bødker, Eds., ACM Press, pp. 3131–3140. doi:10.1145/1978942.1979407.

[86] WRIGHT, W., SCHROH, D., PROULX, P., SKABURSKIS, A., AND CORT, B. The Sandbox for analysis: Concepts and methods. In *Proc. SIGCHI Conference on Human Factors in computing systems* (April 2006), ACM Press, pp. 801–810. doi:10.1145/1124772.1124890.

[87] XU, K., ATTFIELD, S., JANKUN-KELLY, T. J., WHEAT, A., NGUYEN, P. H., AND SELVARAJ, N. Analytic provenance for sensemaking: A research agenda. *IEEE Computer Graphics and Applications 35*, 3 (May 2015), 56–64. doi:10.1109/mcg.2015.50.

[88] YUAN, M. Use of a three-domain repesentation to enhance GIS support for complex spatiotemporal queries. *Transactions in GIS 3*, 2 (March 1999), 137–159. doi:10.1111/1467-9671.00012.

[89] YUAN, M. Representing complex geographic phenomena in GIS. *Cartography and Geographic Information Science 28*, 2 (January 2001), 83–96. doi:10.1559/152304001782173718.