




Contents lists available at ScienceDirect

Visual Informatics

journal homepage: www.elsevier.com/locate/visinf

Guidance in the human–machine analytics process

Christopher Collins^{a,*}, Natalia Andrienko^{b,i}, Tobias Schreck^c, Jing Yang^d, Jaegul Choo^e, Ulrich Engelke^f, Amit Jena^g, Tim Dwyer^h^a Faculty of Science, University of Ontario Institute of Technology, Ontario, Canada^b Fraunhofer Institute IAIS, Sankt-Augustin, Germany^c CGV Institute, Graz University of Technology, Graz, Austria^d Department of Computer Science, University of North Carolina at Charlotte, Charlotte, NC, USA^e Department of Computer Science and Engineering, Korea University, Seoul, South Korea^f CSIRO, Hobart, Tasmania, Australia^g IITB-Monash Research Academy, Mumbai, India^h Faculty of Information Technology, Monash University, Melbourne, Australiaⁱ City, University of London, London, UK

ARTICLE INFO

Article history:

Received 18 August 2018

Accepted 19 September 2018

Available online xxxx

Keywords:

Guidance

Visual analytics

Model evaluation

ABSTRACT

In this paper, we list the goals for and the pros and cons of guidance, and we discuss the role that it can play not only in key low-level visualization tasks but also the more sophisticated model-generation tasks of visual analytics. Recent advances in artificial intelligence, particularly in machine learning, have led to high hopes regarding the possibilities of using automatic techniques to perform some of the tasks that are currently done manually using visualization by data analysts. However, visual analytics remains a complex activity, combining many different subtasks. Some of these tasks are relatively low-level, and it is clear how automation could play a role—for example, classification and clustering of data. Other tasks are much more abstract and require significant human creativity, for example, linking insights gleaned from a variety of disparate and heterogeneous data artifacts to build support for decision making. In this paper, we outline the potential applications of guidance, as well as the inputs to guidance. We discuss challenges in implementing guidance, including the inputs to guidance systems and how to provide guidance to users. We propose potential methods for evaluating the quality of guidance at different phases in the analytic process and introduce the potential negative effects of guidance as a source of bias in analytic decision making.

© 2018 Zhejiang University and Zhejiang University Press. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license

[\(http://creativecommons.org/licenses/by-nc-nd/4.0/\)](http://creativecommons.org/licenses/by-nc-nd/4.0/).

1. Introduction

A research agenda by Thomas and Cook in 2005 proposed “Visual Analytics” as a more holistic approach to visual data understanding than was being explored at the time by visualization researchers and practitioners (Cook and Thomas, 2005). Specifically, it was intended to address problems of making visual analysis scale to very large quantities of data, but was also intended to broaden the research to consider human reasoning models and processes in order to create truly effective analysis tools. It also called for exploration of “human–machine interaction” and “mixed initiative supervisory control systems” (Cook and Thomas, 2005, p. 50). In the intervening years great strides have been made by machine

learning and data mining researchers to build automated data analysis tools with growing sophistication. We have begun to rely on algorithms in more areas of our day-to-day lives, a key example being the modern ubiquitous reliance by drivers on turn-by-turn navigation systems. In this scenario, the machine is truly the guide while the human performs the relatively menial operation of steering the car.

While in 2005 Thomas and Cook spoke frequently of “facilitation” to describe a visual analytics system’s role in supporting human data analysis, the notion of machine “guidance” has recently been introduced into visual analytics by Ceneda et al. (2017). The creep in language used to describe the collaboration between human and machine in the domain of visual analytics may imply a growing confidence that the future of visual analytics systems is one where analysts are guided to insights by the machine rather than achieving those insights through their own agency. However, we are not there yet—there is no practical visual analytics system

* Corresponding author.

E-mail address: christopher.collins@uoit.ca (C. Collins).

Peer review under responsibility of Zhejiang University and Zhejiang University Press.

<https://doi.org/10.1016/j.visinf.2018.09.003>2468–502X/© 2018 Zhejiang University and Zhejiang University Press. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

in which control is handed to the machine at the level to which humans allow themselves to be guided by in car-navigation systems. Such guidance systems in visual analytics remain largely theoretical, with practical systems incorporating guidance mostly only at the demonstration stage, as surveyed in Section 2. In Section 3 we broadly look at what the goals for a guidance system should be in an ideal visual analytics system. In Section 5 we look more closely at recent models of guidance in visual analytics and illicit their shortcomings and incompatibilities. In Section 6 we seek to unify and extend these models to suggest more concrete roles for guidance in visual analytics systems. The work characterizing guidance systems emerging from the visual analytics literature speaks of very high-level “analytical” and “sense-making tasks” that is significantly more abstract than the tasks generally considered in models of information visualization. In Section 7 we review these tasks and extend our discussion to include opportunities here also for guidance. In Section 8 we tie these together as building-blocks for a practical workflow for users of systems incorporating guidance. Finally, in Section 9 we look at how a guidance system in visual analytics can be evaluated for effectiveness.

2. Background

The topic of machine guidance for analytic activities has been of growing interest as the power of machine learning opens new opportunities. The recent paper by Ceneda et al. (2017) introduces a formal description of the opportunities for automated guidance in visual analytics, centered around the knowledge gaps, inputs and outputs, and guidance degree. We plan to expand on this model by broadening the concept of guidance to include just-in-time facilitation which may make analyses processes more efficient by presenting tools and templates at the appropriate moment. Ceneda et al. build their model of guidance atop van Wijk’s model of visualization (van Wijk, 2006), presenting opportunities for guidance at a high level in the process diagram. We investigate the potential role of facilitation across lower-level task taxonomies, e.g., Brehmer and Munzner (2013) and more sophisticated models of visual analytics (Andrienko et al., 2018). Specific instantiations of guidance have been reported, for example, helpful interventions when eye-tracking indicates an analyst is exhibiting signs of confusion (Conati et al., 2013; Toker and Conati, 2014; Panwar and Collins, 2018) or when a logging system detects sub-optimal search strategies (Brown et al., 2014). There have also been investigations into the role of machine intelligence in revealing a potential bias in an analysis process (Wall et al., 2017a), by exposing the differences between the data a user has seen and the overall characteristics of the full dataset. On the other hand, others have raised concerns about the potentially negative impacts of guidance, or machine learning in general, advocating instead for agency and freedom of the analyst (Dörk et al., 2013). Our exploration of the role of guidance will acknowledge the potential pros and cons of each form of guidance with examples from the literature where appropriate. We consider different levels of guidance and facilitation, from low-level operations on adjustment of visual displays to high-level analysis tasks such as model development and evaluation.

The term *guidance* refers to providing the user with help when the user experiences difficulties, e.g., does not know which tool to use or how to proceed in analysis. The term *facilitation* has a broader meaning. It includes guidance but also any possible ways to make the work of the analyst more efficient. There is also a subtle but important difference between the two terms where machine *guidance* may imply the machine has control or power over the user, while machine *facilitation* implies the machine plays a more neutral role in supporting the user’s agency.

The paper by Federico et al. (2017) introduced a conceptual model of Knowledge-Assisted visual analytics by including different knowledge types, namely domain and operational knowledge

as well as tacit and explicit knowledge of the user. Their model proposes to capture the explicit knowledge of the user either by externalization of the tacit knowledge or by a computer-simulated cognitive process. A limitation of the Federico model is that it considers only one channel of information flow from machine to human: the visualization itself. The Ceneda model include additional channels of information flow from machine guidance to human, via “cues” and “options” (which are presumably more transient notifications to the user than the visualization itself) and also “prescribing” changes to the “specification”.

Elements of user guidance can currently be seen in many state-of-the-art visual analytics (VA) systems: Advizor Solutions *Advizor*, Tableau, Qlik, TIBCO *Spotfire*, TIBCO *Jaspersoft SAS JMP*, SAS *Visual Analytics*, IBM *Cognos Analytics*, SAP *Lumira*, Microsoft *Power BI*, ESRI *ArcGIS*, *GeoTime*, and *Centrifuge*, each of which supports exploratory analysis, goal-driven analysis, or a mix of both. Some of these VA systems are targeted to a specific domain, such as ArcGIS and GeoTime focusing on location-based analytics, and Centrifuge targeting cybersecurity and unstructured network data. Other systems claim to provide a multi-domain interactive visualization and analytics solution.

Behrisch et al. (2018b) evaluated commercially available visual analytics systems for the following features: Data Handling and Management; Automatic Analysis; Complex Data Types, Visualization; *User-Guidance, Perception, Cognition*; and Infrastructure. State-of-the-art visual analytics systems are highly interactive and allow interactive exploration of the data for analysis. All of these VA systems support automatic predictive analysis for lower-level tasks such as clustering and outlier detection, classification, and regression, alert mechanisms, and auto-updating of data sources, but they fail to implement a logic-driven conditional analysis and comparative analysis of potential scenarios. *Advizor* is capable of calculating the intermediate visualization results via sampling-based calculations and predictive analysis methods whereas the SAS *Visual Analytics* system achieves the same using incremental updates from long-running tasks. The review of Behrisch et al. revealed that today’s VA systems are rich in features to reduce cognitive overload: recommending visualizations (e.g., Tableau’s *Show Me* feature), analyzing the data types, presenting a ranked list of visualizations, providing data previews, and offering guidance for analysis via some form of wizard. They also support building macros by recording and rerunning of procedures.

In general, the analysis workflow models implemented in the state-of-the-art VA systems are more interactive and are capable of dealing with large volume and variation of data, but guidance is still limited. Analysts are often left alone in an overwhelming and confusing visualization space too large to explore by themselves.

3. Goals and aspects of guidance

To discuss guidance in a systematic manner, we begin with defining the goals of guidance in visual analytics. We can start very broadly. What makes an ideal guide? The precise answer to this question will vary based on the domain, the type of data, the type of enquiry required. The goal of a user of a visual analytics tool most broadly is to gain accurate knowledge about a dataset in order to answer analysis questions. This is different from the example of in-car navigation, an existing and arguably successful guidance system, mentioned in our introduction, where the computer explicitly directs the user in a fairly menial task. The user’s cognitive load is minimized, but their knowledge of their environment grows little. Indeed they no-longer need to pay attention to geography, and their focus is reduced to keeping lanes and avoiding obstacles. It is a short step to replace the human-driver entirely with an AI for the remaining tasks. Some of what we currently describe as data analysis is also potentially or already completely automated. Share

trading (Cha, 2007) and tumor detection (te Brake et al., 2000; Babu Vallabhaneni and Rajesh, 2017) are examples where at least some of the tasks can be performed entirely autonomously.

Visualization's role in visual analytics is most compelling in data analysis tasks where the nature of the task is not easily specified ahead of time. The contents of the dataset are unknown, the opportunities for analyzing the data are not or only partially understood, and so on. In these situations, the abilities of humans to be aware of a broader context for the data and to action a growing awareness of the data into decisions that affect the real world is key and still well beyond the capabilities that we would entrust to an artificial intelligence. Thus, in what aspects can our currently limited machine processes complement human abilities? That is, what are appropriate and attainable *goals* for a guidance systems research agenda?

To inform: This most broad goal echoes the overarching visual analytics described above: to grow knowledge of an unknown dataset. A guide could suggest useful starting points for human visual analysis. For example, the Voyager system is a mixed-initiative system that combines recommendations of potentially interesting views of particular dimensions in data, with a user interface that makes it easy for the user to browse and then drill into or recast results (Wongsuphasawat et al., 2017).

To mitigate bias: Human susceptibility and difficulty in avoiding different types of bias is well studied by psychologists. Wall et al. (2017b) discuss the danger of bias in visual analytics systems, but also the role it plays in analysis, where actively refuting bias by exploring the data, is a legitimate way to drive knowledge and consensus forward. The take-away for a guidance system therefore may not be to eliminate bias, but to keep the analyst aware of their own bias or the biases of others.

To reduce cognitive load: Guidance systems may keep track of analysis processes (capture provenance) and make suggestions to make analysis more efficient, for example extrapolating from current views or recent actions to suggest alternative next steps (Gotz and Zhou, 2009).

For training: Guidance systems may be used to improve usability for novices learning about a new visualization or visual analytics system. Approaches such as visualization by analogy (Ruchika-chorn and Mueller, 2015) and suggested interactivity (Boy et al., 2016) may be useful here, if they were made to dynamically appear on demand. Furthermore, guidance systems which draw on the analysis processes carried out by experts may be useful for experience transfer to raise the skill level of novices (Matejka et al., 2009).

For engagement: Mixed-initiative interaction systems may leverage bio-sensing and other tools of affective computing to provide personalized and just-in-time guidance to mitigate or prevent frustration and increase engagement in an analysis process (Conati et al., 2013; Panwar and Collins, 2018; McDuff et al., 2012).

To verify conclusions: Guidance systems may provide assistance in downstream tests related to visual analytics processes, such as running statistical tests on specific hypotheses, verifying findings, monitoring incoming data for changes in detected patterns, and testing the stability and sensitivity of findings.

In addition to supporting these goals, appropriate guidance should never be harmful, and should only suggest actions and views, rather than prescribe them.

Guidance may be classified broadly into low-level and high-level guidance. Low-level guidance deals with suggestions such as clicking a specific button or viewing particular data in a visualization. High-level guidance provides suggestions about the process of analytics. Strategies here include branching out ("showing something different"), which may reduce bias, reinforcing ("more like this") which could help to confirm hypotheses, and serendipity, or a guided random approach, which may increase discovery.

Furthermore, high-level guidance may also be associated with the tool itself rather than the data, to provide just-in-time help on the interface capabilities.

Guidance can be driven by a number of inputs, which will be discussed throughout this paper. As an introduction, some information that can be used to create guidance includes: interaction logs, view logs (what has been seen), data logs (what has changed in the data) as well as models of analyst knowledge, tasks, requirements, or data (attribute distributions, connections, outliers, etc.).

Guidance approaches can be described based on their front-end (user-facing) and back-end (system) characteristics. The front-end consists of the visual form, the interaction techniques, the style of communication between the system and users, and the integration of the guidance in the analytics process. The back-end consists of the content of the guidance, the algorithmic aspects, the inputs and outputs of guidance algorithms, and the specific information sources used to derive guidance suggestions.

4. Requirements of intelligent guidance

Guidance approaches supporting the goals given in Section 3 should meet certain desired requirements. From the literature, a number of requirements, which are often postulated, can be identified. We discuss a number of important abstract and more concrete requirements as follows.

Generally, a guidance system should be able to provide effective guidance, i.e., fulfill the set of guidance goals. In that, the guidance system should prove to be useful to users and allow the qualitative or quantitative measuring of its added value over non-guidance supported systems. The guidance functionality should be easily accessible by the user; the user feedback required for training (see Section 8.1) should be intuitive to provide and not disrupt the analysis process. The method of displaying suggestions or potentially helpful additional information needs to be aware of the environment available and avoid distraction, or obscuring the current visualization.

When providing guidance, the system should communicate why a suggestion or guidance is being given. Specifically, the guidance should resemble more of a *white box* instead of a *black box*, supporting user confidence in the appropriateness of the guidance. The system confidence about when and which guidance to provide may also be varying or low. The system should determine a suitable threshold when guidance is likely to be helpful; or at least, be clear about the level of confidence in the guidance when issuing a guidance step. It is also desirable that the provided guidance and user steps taken in response should be traced and added to the analysis provenance.

More generally, the provision of guidance should be adapted to the *context* of the user analysis process. Depending on the stage of analysis, e.g., exploration vs. confirmation, but also the task, the subset of data already seen, or the current and past views considered and their sequence, different kinds of guidance could be reasonable. Adaptation of the guidance system should include different guidance levels for different types of users, different levels of expertise, and different states of the user, e.g., being frustrated, confused, engaged, etc. Furthermore, different roles for guidance could be required in exploratory data analysis and in other aspects of visual analytics which are typically more goal-oriented. In goal-oriented tasks, the guidance or advice should be more specific than in exploratory analysis where guidance may run a whole suite of analytics and provide summary information or suggestions for exploration. Guidance should be adapted to the user but also be available to groups of users, e.g., during collaborative analysis settings if needed.

More specific requirements stipulate that the guidance should be provided at the right time and in the right mode. Regarding the

time, the system should monitor the analysis process and predict the appropriate times to guide. This could be, for example, times when a user might be confused or lost in the analysis stage. A heuristic to implement the latter could be to detect when users are revisiting previously seen states (Behrisch et al., 2014) or searching in a non-systematic manner (Brown et al., 2014). Regarding the mode, guidance may be provided in a *synchronous* manner, issuing notifications or interventions in the analysis process. Guidance could also be provided in an *asynchronous* manner, allowing the user to come back to the guidance at his or her discretion (Mehta et al., 2017). Gradual guiding is also possible, with the system providing some guidance at an early stage and more only on user request. The user should also be empowered to proactively set the level of guidance needed or ultimately, opt out of any guidance at all if so desired.

The requirements may also differ if specific visualization environments are considered, e.g., desktop-based vs. augmented or virtual reality analysis environments. Specifically, modes of interaction and visualization options may differ, e.g., requiring to use interaction paradigms such as voice or gestures.

The above requirements are among the most widely discussed, and there may exist more. Just from these examples, we observe the design and requirement space for guidance approaches is large. Existing solutions to date address specific goals and requirements.

5. Guidance in the visual analytics process

5.1. Limitations of the existing conceptual model of guidance

The conceptual model of guidance proposed by Ceneda et al. (2017) is based on van Wijk's generic model of visualization (van Wijk, 2006), in which visualization is applied to data using a specification (methods and parameters) and creates an image. A user perceives and interprets the image using his/her current knowledge, thereby increasing or modifying that knowledge. Based on their current knowledge, the user may perform interactive exploration, which affects the specification and thus modifies the image, which is further perceived and interpreted.

Ceneda et al. make a slight modification of this basic model by replacing the term 'Visualization' by 'Visual Analytics.' Like visualization in the original model, visual analytics transforms data into an image to be perceived by the user. In this view, visual analytics is considered basically as a combination of visual and analytic methods. Guidance can be applied to the specification of these methods. Hence, the guidance model from Ceneda et al. is not specific to visual analytics. Guidance is represented as helping the user create, perceive, and transform an image. A later paper (Federico et al., 2017) discussing the possible roles of explicit knowledge in a visual analysis process notes that inputs for guidance are explicit knowledge, data, and specification containing the full history of previous settings used in the exploration. These inputs are analyzed to generate specific suggestions.

The main limitation of this conceptual model is that it is too abstract to use practically in designing and implementing guidance tools. It shows *where* guidance can be provided and proposes a set of attributes to characterize guidance (*how*), but it does not propose or imply any approach to understanding *what* specific guidance can be provided.

For understanding this, it is necessary to consider more specifically what *knowledge* the user wants to derive from the data. Assuming that the user derives knowledge by perceiving images, the main task of guidance is to help the user create such images from which the required kind of knowledge can be effectively derived through perception. Hence, a designer of a guiding system needs to anticipate the kind(s) of knowledge that will or may be required. The guidance model from Ceneda et al. cannot help with this

because, similarly to 'visual analytics', it represents 'knowledge' at the highest possible level of abstraction, i.e., as a single atomic block in the overall scheme.

Furthermore, derivation of knowledge involves not only image perception and exploration but also verification of findings (Sacha et al., 2014), which is not considered by Ceneda et al., although this very important activity may also need to be guided. Again, to be able to provide specific guidance, it is necessary to anticipate the possible kinds of findings.

To extend the existing model of guidance beyond *where* and *how* to *what*, we build upon a recently proposed conceptual framework in which visual analytics is represented as a model building process (Andrienko et al., 2018). This framework will be further referred to as a model building framework (MBF).

5.2. Conceptualization of visual analytics as model building

The MBF is based on a definition of *knowledge* as a *model* of some piece of the world, which is called 'subject'. A model is any kind of representation, e.g., verbal, graphical, or mathematical. A model of a subject represents its aspects (i.e., components and their properties) and relationships between them.

The goal of analysis is to build an appropriate model of a subject using data (observations and measurements) partly reflecting this subject. Criteria for model appropriateness include correctness, comprehensiveness, fitness to purpose, generality, specificity, and so on.

The analysis begins with generating a tentative initial model. Throughout the analysis, the current model is repeatedly evaluated in regard to the appropriateness criteria and further developed if not yet appropriate. Besides these core activities, the analysis process may also include collecting provenance and externalizing the model obtained.

5.3. Required support to model building activities

The paper introducing the MBF (Andrienko et al., 2018) also discusses how these activities can be supported in visual analytics systems. Thus, the generation of an initial model requires the use of visual and/or computational methods promoting abstraction and generalization. Guidance may help with choosing such methods and using their results. For model evaluation, approaches adopted in statistics and machine learning can be applied to model components represented in a computer-readable form. Mental models or model components can be evaluated using visual and interactive means, which include re-application of previously used methods with different parameter settings, application of alternative methods, and taking different subsets of data.

The paper notes that, while there exist established practices of evaluating models in statistics and data mining, evaluation of mental models does not receive sufficient attention in visualization and visual analytics research. It is stated:

"Although many visual analytics systems and toolkits include interactive facilities for the operations mentioned above, it has to be a decision of the analyst to apply these operations. The existing software neither informs/reminds the analyst about the possible use of the available interactive techniques for mental model evaluation nor encourages the analyst to even concern about such an evaluation" (Andrienko et al., 2018, p. 289).

This reveals a clear need for user guidance. Similar considerations apply to model development, in which the analyst may need to rectify, expand, or simplify the current model based on results of an evaluation.

Provenance collection and model externalization involve explicit representation of findings, interpretations, judgments, inferences, and the final model in a form that can be transferred to

others. These activities require tools for annotating images, organizing and linking notes, and constructing more abstract representations such as knowledge graphs. The user may greatly benefit from intelligent guidance and facilitation. Thus, the system can propose structured templates to describe and organize findings or even construct draft annotations based on automatic detection of patterns in data.

5.4. Role of patterns in data analysis

An essential feature of a model is that it is a *generalized* representation (Andrienko et al., 2018, p. 283) in the sense that it refers to multiple observations taken together rather than consisting of specific representations of individual observations. Moreover, the generalization extends beyond available observations (i.e., represented in data or known to the analyst) to observations that could potentially be made in the future. This required feature of a model implies that users of visual analytics systems should be able to perceive multiple data items together and conceptualize them jointly as a meaningful whole. Such a whole is commonly called a ‘*pattern*’ in the visualization and visual analytics literature.

While being widely used, the term ‘*pattern*’ has not been explicitly defined. We propose the following working definition of a pattern:

A pattern is a representation of a collection of items of any kind as an integrated whole with specific properties that are not mere compositions of properties of the constituent items.

A pattern can be viewed as a possible component of the model being built. It can be included in the current model if it is relevant to the analysis goal. A relevant pattern needs to be evaluated, e.g., by checking if it can still be perceived or otherwise extracted after changing the visualization or analysis method, or by computing aggregate characteristics of the data items involved in the pattern and comparing them to corresponding aggregate characteristics of the remaining data.

Since generalization is the core and essence of model building, the analysis process is largely centered on patterns, which are extracted from data, evaluated, refined, organized, interlinked, annotated, and integrated into the final model. Consequently, the task of guidance is to help analysts extract and manage patterns.

5.5. Implications for designing and implementing guidance

5.5.1. Help in pattern extraction

Analysis starts with the generation of an initial model, which requires initial extraction of one or more patterns from data. It may also happen that some initial model already exists in the mind of the analyst (e.g., as a result of previous analyses of similar data). In such a case, the analyst needs to check if the pre-existing model complies with the current data, which requires the extraction of patterns from current data. Hence, the main task of guidance at the initial stage of analysis is to help the analyst extract patterns from data.

Pattern extraction takes place not only at the initial stage of analysis but also within the following loop of model evaluation and development. In model evaluation, the analyst modifies the data (e.g., takes a different sample), methods (e.g., uses another clustering algorithm or applies another visualization technique), or parameters, extracts patterns, and compares them with the ones extracted previously. Good correspondence gives evidence of pattern trustworthiness. In model development, the analyst may need to search for additional patterns that can refine or expand the current model, or for different patterns if previously extracted patterns have not been confirmed in the evaluation. Hence, guidance

in pattern extraction is relevant throughout the whole process of model building.

Patterns can be extracted from data mentally or computationally. Mental extraction requires an appropriate visual representation of the data enabling pattern perception. Computational extraction can be done using appropriate algorithms, and the results need to be presented to the analyst for interpretation, evaluation, and incorporation in the overall model. To support mental extraction of patterns, the guidance subsystem (further referred to as ‘*guide*’) can suggest suitable visualization techniques or automatically create effective visualizations. For computational extraction, the guide can suggest suitable algorithms. Both kinds of support can be possible if the guide knows, first, the structure and properties of the data, second, the analysis focus (which aspects are relevant), and, third, the analysis goal (e.g., description or prediction). Knowing the data structure and properties, the guide can anticipate what kinds of patterns can potentially exist. Knowing the analysis focus and goal, the guide can determine which kinds of patterns may be important and useful.

As an example, let us take the analysis task from the VAST Mini Challenge 1: “*investigate the circumstances of an epidemic outbreak in a city and forecast how it will develop further*” (Grinstein et al., 2011). The data consist of geographically referenced messages from social media, some of which mention disease symptoms. The data structure thus includes temporal, spatial, and textual components. Patterns can exist in the temporal evolution of the number and contents of the messages, the spatial distribution of the locations of message posting, and the joint spatiotemporal distribution of the messages and their contents. Possible patterns in the temporal evolution of the number of posts include temporal trend (increase or decrease), peak or pit, periodic repetition of some smaller pattern or pattern sequence, and random fluctuation. Possible patterns in the temporal evolution of the message contents include changes of keyword frequencies. Possible patterns in the spatial distribution of the posts include spatial uniformity, spatial clustering, spatial trends of the density (e.g., increase or decrease from north to south or from the center to the periphery), and spatial alignments. Possible patterns in the spatial distribution of the message contents include higher or lower frequencies of particular keywords in different parts of the territory. Possible patterns in the spatiotemporal distribution include concentration and dissipation, appearance, disappearance, growth, shrinkage, movement, merging, and splitting of spatial clusters, and changes of keyword frequencies in different parts of the territory. All these kinds of patterns are relevant to the task of describing the disease outbreak, whereas only temporal and spatiotemporal trends are important for the task of forecasting the further development.

Generally, to support pattern extraction at different stages of the analysis process (initial model generation, evaluation, and further development), the guide can (1) suggest or automatically choose visualization techniques showing data distributions (statistical, temporal, spatial, spatiotemporal) and correlations among data components and/or (2) suggest or automatically run appropriate pattern extraction algorithms, such as clustering, trend detection, or motif discovery. Since data may contain outliers, which may obstruct pattern extraction, the guide can also suggest methods for outlier detection and removal and help with the use of these methods.

5.5.2. Help in pattern and model evaluation

Evaluation of extracted patterns includes measuring their properties (how frequent, how high, how dense, how large, how fast, how regular, etc.), checking pattern stability or sensitivity with respect to changes of the methods by which they have been extracted and with respect to noise in data, and investigating the scope of the patterns (in what part of the data they exist). For computationally

extracted patterns, property measurements can be done automatically. To provide valuable help in evaluating mentally extracted patterns, the guide would need to know what specific patterns have been extracted, i.e., the analyst would need to mark the patterns in the visual display and annotate them in a structured form telling the guide what kinds of patterns they are. Regarding pattern stability/sensitivity and scope, the analyst may not only need help in choosing and applying suitable approaches to testing but may also need to be engaged in such testing (Andrienko et al., 2018). As a minimum, the guide may give general recommendations to look at the data in different ways, take different random samples, try other methods or parameters, introduce some noise in the data, etc. To provide more specific help, the guide needs to know the extracted patterns and the methods and parameters used for the extraction. Apart from that, the analyst should be given suitable tools, support, and guidance for comparing patterns obtained with different settings.

A model can be seen as a system composed of patterns that are linked by relationships, such as hierarchical (larger patterns include smaller ones), temporal, spatial, or causal. A model can also include patterns representing different parts of the data. Such patterns are linked by difference and partitioning relationships between the corresponding parts of the data. Model evaluation involves, apart from evaluating the constituent patterns, assessment of its overall correctness (consistency with the data), comprehensiveness (inclusion of all relevant aspects and relationships), coverage (inclusion of all available data), generality (applicability to data that were not used for model generation), specificity (representation of important distinctions and details), prediction capability, complexity (number of components and relationships), and resource efficiency (e.g., possibility to obtain similar results using less time and/or computer resources). At least some of these criteria can be addressed in guidance. Thus, regarding comprehensiveness and coverage, a guide can inform the analyst about data components or subsets that have not yet been considered, which can be done without having an explicit representation of the current model in a computer-readable form. Evaluation with regard to some other criteria can more easily be supported for computer-based models than for mental models. There are established techniques for testing computer-based models, which can be recommended to the analyst. Besides, the guide can show the distribution of good and bad model results, help in identifying and comparing data subsets for which model results are good and bad, and highlight uncertain and borderline cases.

For a mental model that is not externalized, the guide can only provide some general suggestions concerning potentially important evaluation criteria and common ways of assessment. Having a structured explicit representation of a mental model, the guide may be able to provide more specific help. Thus, having explicitly represented patterns, the guide may help analysts check: if there are similar patterns that refer to different data subsets and therefore could be joined; if there are different patterns that apply to the same data subset and thus require conflict resolution or redundancy removal; or if patterns extracted earlier can also represent data that have not been considered yet in the analysis. Let us illustrate this idea with an example.

An analyst studies social media posting activities in a touristic region containing mountains, valleys, lakes, forests, and towns. The data include posting times, location references (coordinates or place names), and texts (such as messages in Twitter or titles of photos). The analyst wants to see what places are popular in different seasons of the year and, based on the keywords occurring in the texts, what people do there. The analyst first takes a subset of data generated in winter and uses a density map to identify places with a high number of posts, which appear as “hot spots” on the map. The analyst outlines the boundaries of these places

and examines which keywords frequently occurred in the whole set of places. She finds two prevailing groups of keywords, one referring to winter sports and the other to Christmas and New Year celebration. The analyst separately selects the posts referring to the former and latter topics and explicitly records two spatial patterns: places used for winter sports and places where people celebrate Christmas and New Year.

The analyst then focuses on the data generated in the summer season, in which the keywords refer to winter and summer sports, hiking, and sightseeing. When the analyst takes the subset of posts referring to winter sports, an intelligent guide can find the previously recorded spatial pattern of winter sports places and superimpose it on the density map, helping the analyst to check if the pattern also applies to the summer season. The analyst notices that some places are used for winter sports both in winter and in summer. She refines the previous pattern by subdividing it into a pattern of places used for winter sports only in winter and those used for winter sports in any season. When the analyst creates a density map of sightseeing-related posts, the intelligent guide can detect that the previously recorded spatial pattern of Christmas and New Year celebration matches very well the density distribution. Indeed, the places having interesting sights to see are also popular as places for spending the winter holidays. When the guide puts the celebration pattern on top of the density map, the analyst may decide that the pattern can also represent sightseeing activities and extend the annotation of the pattern accordingly.

When the analyst investigates the spatial distribution of hiking-related keywords, the guide may check if these keywords also occur in the texts of the winter data subset. It may find that the frequency of these keywords in winter is sufficiently high to deserve attention and notify the analyst about this. The analyst might not have noticed these keywords among the most frequent terms occurring in the winter subset because they were dominated by the winter-specific keywords. After being notified by the guide, the analyst may look for the spatial distribution of the hiking-related keywords in the winter subset and record the corresponding spatial pattern. The guide may detect a high overlap of this pattern with the pattern of winter sports distribution and exhibit it to the analyst for resolving a possible conflict. In response, the analyst may select the places used both for winter sports and for hiking and compare the temporal distributions of the keywords related to winter sports and to hiking at the level of days or weeks. An absence of temporal overlap may mean that the place use depends on the weather, particularly, the presence of snow and/or ice.

In this imaginary scenario, the guide helped the analyst not only in evaluating extracted patterns but also in model development, which is discussed in the following subsection. Generally, the problem of support and guidance in model evaluation has not yet been sufficiently addressed in visual analytics research and thus provides a challenging but interesting and important research direction.

5.5.3. Help in model development

Model development means an improvement of the current model with regard to issues identified by evaluation. For a computer model that is insufficiently correct or insufficiently specific, the guide can suggest modification of modeling method parameters, trying alternative methods, or dividing heterogeneous data into more homogeneous parts and replacing an overall model by a combination of more specific models. Data partitioning followed by extraction of more specific and accurate patterns can also be suggested for improving mental models. For increasing model comprehensiveness and coverage, the guide can show data subsets and aspects that have not yet been covered and support pattern extraction from them. The guide can also inform the analyst about existence of alternative analysis and modeling methods or additional data sources. Like initial model generation, model development involves pattern extraction, which was discussed earlier (Section 5.5.1).

5.5.4. Help in provenance collection

To support provenance collection, it is necessary to keep track of all operations and methods applied to data (Xu et al., 2015). The history can be represented visually, e.g., as a graph (Shrinivasan and van Wijk, 2008). Gotz and Zhou (2009) describe how a taxonomy of the user's actions can be used for automatic capture of semantically meaningful and logically organized provenance. This requires the system to have a "semantic" user interface organized according to the action taxonomy. Besides exploration history tracking, it is necessary to enable and facilitate annotating. Annotations can be valuable not only for tracking the provenance but also for providing help in model evaluation and development (see Section 5.5.2). To encourage and facilitate the creation of annotations, the guide can propose structured templates or even draft annotations that can be completed by the analyst with small effort. The basis for such help is the knowledge of possible patterns, methods used, and, for computational methods, properties of results obtained. Particularly, when a system is oriented to specific types of data and analysis tasks, the patterns, methods, and result properties are known to the system designer, who can infer the required contents and structure of annotations and thus prepare suitable templates (Eccles et al., 2008; Chen et al., 2010). Thus, the system Click2Annotate provides pre-defined templates for popular types of patterns, or facts: cluster, outlier, rank, difference, correlation, and compound fact consisting of two or more smaller patterns. Templates or draft annotations for mentally extracted patterns can be generated in response to analyst's interactive marking of observed patterns in a visual display.

Besides creating annotations, the guide should help the analyst to organize and manage them. For example, in Aruvi (Shrinivasan and van Wijk, 2008), users can attach notes to nodes of a graph representing exploration history and link them to states of the visual display, and similar functions exist in some other systems (Gratzl et al., 2016; Walker et al., 2013), which, however, do not provide guidance for annotation creation and management. A kind of guidance exists in the system Sandbox (Wright et al., 2006), which provides automatic layout for collected information pieces and notes using templates of analytical frameworks, such as process models. The analyst needs to choose a suitable template explicitly.

Annotations created in the process of analysis represent pieces of the analyst's knowledge (i.e., mental model of the analysis subject) and thus can be used for externalization of the final model.

5.5.5. Help in model externalization

As noted in the MBF paper (Andrienko et al., 2018, p. 290), fragmentary notes attached to some states of visual displays do not form an adequate representation of the model the analyst has in mind. For a more complete and systematic external representation of the model, the analyst needs to organize the notes and, possibly, provide additional comments. When the externalized model is only meant to be viewed by humans, the analyst may represent the model, together with the provenance, in the form of a "story" (Eccles et al., 2008; Gratzl et al., 2016; Walker et al., 2015). To construct a good story, the analyst may need guidance, which is not available in current systems. In Sandbox (Wright et al., 2006), the user constructs a graph from pieces of evidence and notes, which represents the user's mental model. Besides providing automated layouts based on workflow templates, Sandbox can also automatically derive a concept map from the text notes.

A concept map is a commonly used form of knowledge representation (Dwyer, 2016). It can be used for transferring knowledge between human analysts (Zhao et al., 2018), and this format is also suitable for computer processing. The system HARVEST (Gotz et al., 2006) enables the user to create and manage concepts and instances, maintains a base of concepts, instances, and annotations, represents the collected knowledge pieces visually, and allows the

user to link them to corresponding evidence in a data display. The user can also analyze and further develop the knowledge synthesized. The system facilitates the user's activities by automated knowledge management and visualization but does not provide explicit guidance.

The MBF paper notes that further research is needed on supporting both provenance collection and knowledge externalization, and this also applies to guidance. Particularly, the externalization of computer models in a human-readable form needs to be addressed. So far, this has been done for specific kinds of models, such as causal networks (Wang and Mueller, 2017), which can be represented in a graph form.

6. Knowledge of an intelligent guide

In discussing the visual analytics process in Section 5 and, particularly, possible support to it in Section 5.5, we have already mentioned different kinds of knowledge and capabilities that are required for fulfilling the expected functions. Here we summarize and, where necessary, extend these requirements. Fig. 1 schematically represents the visual analysis process and the possible types of guidance and help, and Fig. 2 shows what kinds of knowledge can be used for providing them.

Prior to analysis, an intelligent guide would need to have a general knowledge base of (1) data types and structures, (2) possible relationships among data components, (3) possible errors and uncertainties in data, (4) existing approaches to detecting and correcting data errors and to dealing with uncertainties, (5) possible patterns, such as trend and seasonality in time series, (6) possible analysis tasks and types of patterns that may be relevant to them, (7) existing visualization and analysis methods, their applicability to data types and their capability to exhibit or detect patterns and relationships, (8) possible user actions and possible purposes for doing them. It would also be good if the guide could have a model of the user, including the user's expertise and background, goals and questions; however, user modeling has not yet been addressed in visual analytics research.

When the analysis starts, the guide should (1) understand the structure and properties of the loaded data, (2) anticipate patterns and relationships that may exist, and (3) be able to find sources of additional related data. The guide should also help the user with (4) detecting and correcting errors in the data and (5) getting aware of data uncertainties and their possible effects on the analysis.

In the process of analysis, the guide should be able to (1) track the process, (2) facilitate collecting provenance, (3) understand the current situation and anticipate further steps. For the latter, the guide needs to have an adaptive and growing understanding of the users' current knowledge and further intentions as they perform an analysis. Such understanding can be derived from generated annotations and concepts and established relationships among them, and from continual feedback, either explicit or implicit, from the user.

When the analysis finishes, the guide should help the user to externalize the mental model built based on the previously tracked analysis process and collected annotations and concepts. The guide should propose appropriate arrangements for the collected material and, when possible, automatically construct draft representations, such as a knowledge graph or a story, which can be edited and completed by the analyst with small effort.

7. Guidance for visualization tasks

The previous sections investigated the role of guidance in visual analytics processes, and the knowledge a guidance system would require to provide recommendations. In this section we dive into the low-level tasks that analysts carry out when working with




Fig. 1. Schematic representation of the visual analysis process and the possible types of guidance and help. The ovals represent the analyst's activities: IM—generate initial model, EvM—evaluate model, DM—develop model, CP—collect provenance, ExM—externalize model. The double lines represent the process flow; the directed lines show the sequence while the undirected lines link activities performed in parallel. Thin arrows represent data and information flows. The blocks with grey background represent possible kinds of support and guidance, and the glowing lines connect them to corresponding activities in the analysis.




Fig. 2. Types of knowledge that can be used for supporting and guiding the visual analysis process. The representation of the analysis process is the same as in Fig. 1. The grey boxes represent the knowledge of an intelligent guide, and the glowing directed lines represent derivation of new knowledge.

visualization components within visual analytic processes, and explore some design ideas of how guidance may fit into each task. Wherever relevant, we relate these low-level tasks to the high-level view of the visual analytics process discussed in Section 5. To

structure the discussion, we enumerate guidance for the *Why?* and *How?* levels of the typology by Brehmer and Munzner (2013, Fig. 1). For each task, our investigation explores guidance ideas, including inputs to guidance and the potential role in bias.

7.1. Guidance opportunities for why tasks

The *why* typology of Brehmer and Munzner (2013) focuses on the analyst intent about why a task is performed, and breaks down into four major categories: consume, produce, search, and query. Some of the task types are closely related to the earlier presented model-building view of the visual analysis process (Section 5.2).

Consume relates to the use of a visualization to consume information in a variety of domain contexts. There are three specific tasks within this category: present, discover, and enjoy. *Presenting* visualizations, in a collaborative setting or meeting, could be aided by guidance before and during the presentation. Before presenting, guidance could provide suggestions on views to include, and order these views in a suggested storyboard. This guidance could be informed by logs of the analysis process and the views used by the analyst leading up to the presentation, as in HARVEST (Gotz et al., 2006). In a system supporting knowledge externalization (Section 5.3), the guidance can be informed by explicit representations of discovered patterns and externalized components of the analyst's mental model that has been built in the process of analysis (Sections 5.2, 5.5.5).

When analysis provenance and/or externalized knowledge constructs are not available, or no detailed analysis has yet been conducted, *discover* tasks can still be supported. According to the earlier presented ideas for guidance in the visual analytics process (Section 5), this corresponds to support in pattern extraction (Sections 5.4, 5.5.1). The data types and pattern analysis algorithms could be used to suggest visualization views on the data, and highlight the detected patterns, similar to the view recommendations of Wongsuphasawat et al. (2016, 2017). Discover task guidance algorithms could also be informed by search heuristics to detect patterns and data correlations, by user navigation history, and by accumulated interaction logs from other analysis sessions and other users. Guidance in discover tasks may decrease the impact of bias the user brings to the task, by suggesting potential visualization views that do not support preconceived notions about the data. However, similarly, discover assistance may also restrict the analyst's understanding to only those subsets of data which are suggested by the system. Therefore, the guiding subsystem should be designed so as to promote comprehensive exploration of the entire dataset and viewing the data from diverse perspectives. It should also encourage and facilitate the validation of findings (Section 5.5.2). Guidance in discover tasks could be implemented using interesting point detection, subspace analysis, quality measures (Tatu et al., 2011), and through learning user relevance (Behrisch et al., 2014; Healey and Dennis, 2012).

Using visualization for *enjoyment* could be supported through guidance to views on data which are aesthetically pleasing, surprising, or related to the user ('egocentric' views). Guidance for enjoyment could be further driven also by affective metrics such as bio-sensing measures of emotion.

The *produce* use of visualization relates to a using visualization to create new artefacts, such as views of data, derived attributes, or groupings of data items, as well as annotations and externalized concepts (see the block "User-generated materials" in Figs. 1, 2). Guidance possibilities under this goal include suggesting views of data based on statistical analysis of data attributes and suggesting settings of variable parameters in a visualization system to create views with visible trends or clusters. In addition, guidance could assist the creation of visual summaries and overviews of large datasets, and provenance tracking systems (such as HARVEST) could produce explorable histories of the analysis process. Annotations on a view may be the product of a produce task—guidance could be provided to indicate data items related to those items involved in annotation actions. The discussion of possible assistance

for analysis provenance collection and knowledge externalization (Sections 5.5.4, 5.5.5) is also relevant to the produce tasks.

The *search* uses of visualization are often precursors to other tasks. In order to discover information or produce new views or data, an analyst must first find the data of interest. Brehmer and Munzner break search down into four tasks based on whether the target is known and whether the location is known. When a specific target and its location are known, the task is simply to *lookup* the information. Guidance can be used here in the form of autocomplete and predictive search. If the target is known and the location is unknown, the task is to *locate* the information, and a guidance system could track which parts of a visualization have been explored and guide the user toward unseen views. A locate guidance may also triage likely locations based on known characteristics of the target to narrow the search space.

When one does not have a specific target in mind, the task is to either *browse* through a known location or *explore* when the location is not known. Guidance for these tasks encompasses many of the approaches suggested for *discover*, *produce*, and *locate*. Guidance could be driven by user behavior and data characteristics, to guide exploration into parts of the visualization space which are similar to (or different from) data which has already been explored. Support for bookmarking significant views for revisitation would assist the browsing and exploration process. Visualization types and specific views could be suggested to reveal patterns detected in the data, such as attributes with strong dependencies, or sections of graphs that are densely connected. To avoid potential bias in the exploration process, a guidance system should be transparent about why it is making recommendations, and the scope of the recommendation space. That is, with exploration guidance, users should be aware what sorts of guidance are not possible, so that they can consider how much to rely on the system. Care should be taken to design guidance for exploration to assist users but not to fully prescribe exploration pathways, which could lead to serious issues of algorithmically-driven bias.

The final group of *why* tasks in the typology are the *query* tasks, which act on targets once they are found. After a lookup or locate task, *identify* returns target item characteristics (often referred to as "details on demand"). Here, guidance could, for example, highlight target attributes that are statistically unusual or rank attributes to surface those predicted to be more important in the current context. In a browse or explore task, identify returns item references, such as a reference to the item with the highest value on a particular attribute. In this case, guidance may provide contextual information in the identification process, such as the identity of similar data items.

The *compare* query task can be facilitated through guidance that may suggest which items may be interesting to compare to the selected target, or which attributes exhibit patterns which differentiate items from one another. Finally, guidance in the *summarize* task may suggest views on the data which are appropriate to summarize the selected items or attributes. For example, a summary of a choropleth map may be a histogram of the distribution of values, while a suggested summary of a graph may be an abstracted view of the graph nodes and their connections, grouped into communities. Automatically generated natural language captions which can be edited may also be useful as guided summaries.

7.2. Guidance opportunities for how tasks

The *how* aspect of the Brehmer and Munzner typology aligns well with the user-centered view of low-level tasks by Yi et al. (2007). These are the tasks which people are most likely familiar with when thinking of working with visualizations. Brehmer and Munzner distinguish between three classes of methods, and we will introduce guidance opportunities for each in this section.

Methods that *encode* data in a visualization include choosing the visual variables to use to represent data and designing the layout of marks on the page with the aim to facilitate pattern extraction (Section 5.5.1). The space of guidance for encoding is already rich, with recommendation engines such as Google (2018) and Wongsuphasawat et al. (2017, 2016) already deployed for use. Systems can suggest appropriate visual encodings based on the cardinality and distribution of values. For example, a hue encoding would not be appropriate for high cardinality categorical data. Furthermore, guidance systems have been designed to assist in the specific encoding choices, such the perceptual-modeling driven recommendations of Colorgorical (Gramazio et al., 2016). Guidance systems for encoding data into visualizations risk introducing bias through the selection of encodings and pre-designed visualizations which are possible, and the data transformations applied to place data in the view. For example, if data are automatically clustered, scaled, or outliers removed, that may change the interpretation of the resulting visualization.

When we think of interaction with an information visualization, the specific interactions that come to mind fall under the *manipulate* class in the typology. *Navigate* tasks can be supported through guidance which suggests new views on data, which can be helpful, in particular, for model evaluation and development (Sections 5.5.2, 5.5.3). This guidance can be discrete, suggesting a gallery of new viewpoints for the visualization, or continuous, suggesting directions for operations such as pan and zoom. Guidance systems could be driven by back end data modeling and tracking of user views in order to guiding to areas which have not been seen, or parts of the visualization which contain data which is similar to, or different from the current view. This approach could reduce potential bias or misinterpretation of the data by revealing counterfactual views and assisting an analyst to ensure coverage of the visualization. The Voder system exemplifies this type of guidance, extracting facts from a dataset and suggesting views into the data which will reveal the evidence supporting the facts (Srinivasan et al., 2018).

The *arrange* method refers to adjusting layout of the data items. For example, a guidance system may suggest a layout change to reorder parallel coordinates axes to make correlations apparent, or may suggest a log scale transformation of an axis to make trends visible, thereby facilitating pattern extraction (Section 5.5.1). Guidance may also be incremental and in conjunction with user input, for example responding to user adjustments to, e.g. graph layout, by learning the layout strategy and propagating it to other parts of the graph. Guidance may also be offered to arrange items in a visualization or views in a coordinated workspace according to a similarity or distance metric. Such guidance could highlight similarities and clusters, however incorrect or coincidental associations could introduce accidental patterns and encourage misreading.

The *change* task adjusts a visual encoding, such as changing the styling of items in a scatter plot or changing the representation of a chart. Guidance systems such as the system of Srinivasan et al. (2018) present suggestions of alternative charts, for example, suggesting a box plot as an alternative to a bar chart for data containing outliers. Systems which support easy transition between chart types, such as Tableau and Excel, could be augmented with suggestion systems to reveal alternative designs. Furthermore, guidance systems could explain the reason behind a suggestion—box plot for outliers, stacked bar for two attributes, etc. The Voder system uses extracted facts to suggest chart types which more readily reflect particular aspects of the data (Srinivasan et al., 2018). Guidance could also be provided to adjust mappings, such as applying a scale transform on data (Heer and Agrawala, 2006), or a new color scale (Gramazio et al., 2016), for example, to increase discernability of differences or reduce visual clutter. Change

guidance may facilitate pattern evaluation (Section 5.5.2) and reduce bias in analytic processes by encouraging users to view data from different perspectives.

Filter methods add or remove elements from view, either temporarily (hide/reveal) or permanently (add/delete). Guidance in filter operations can be used to declutter views, such as the MDL Treecut algorithm for automatically hiding branches of tree diagrams which carry little information (Veras and Collins, 2017). Guidance can also be provided to detect noise and suggest appropriate thresholds to filter data, such suggesting parameters for a band-pass filter on sensor data. In multivariate data, guidance may suggest attributes to investigate first, for example, attributes with unusual distributions or attributes which are highly correlated. This could speed the process of investigating high dimensional data. Guidance may also be used to suggest values for parameter settings, range sliders, and other filter widgets which may reveal views which are interesting (e.g. have clearly detectable patterns, outliers, etc.). For example, Scented Widgets provide cues to users about potentially interesting parameter values (Willett et al., 2007). Suggestions of this type may be informed by back-end correlation, clustering, and data analysis, as well as explicit specifications of user interest, or models of interest learned from user behavior.

Finally, *aggregate* methods change the granularity of the display of items in a visualization, such as the steps in a time scale or clustering level of nodes in a graph. Algorithms which are aware of both on-screen clutter and information content, such as the minimum description length-informed treecuts of Veras and Collins, can guide the selection of aggregation levels which balance information density and usability (Veras and Collins, 2017). Guidance could also suggest methods of aggregation, appropriate parameter settings for cluster thresholds, or highlight data items which may be aggregated by the user. While guided aggregation approaches may make for more scalable visualizations, it is likely that suboptimal guidance could lead to missing details or overly cluttered views. For aggregation tasks, it is also desirable to provide guidance for testing whether patterns that can be observed are sensitive to aggregation parameters, e.g., whether the overall shape of a histogram significantly changes after modifying the bin size.

The final methods of Brehmer and Munzner's *how* typology are methods which *introduce* new elements to a visualization, for example, to represent derived or imported data. We have covered guidance in these more visual analytic tasks in Section 5.

8. Building blocks for implementing guidance

The previous sections discussed rationale, requirements and examples of guidance-based systems. We observed that guidance can help at different stages in the visual data analysis process, and that there are different approaches to implementing guidance, depending on the type of system. We propose to abstract the implementation of guidance approaches by a typical *input* ↔ *compute* ↔ *output* workflow. We next discuss building blocks for implementing guidance based on these steps, namely, the collection of input from the user (Section 8.1), the computation of guidance steps to give by the system (Section 8.2), and the actual presentation of the guidance to the analyst (Section 8.3). Note that we here give non-exclusive examples of commonly used techniques. We also note that the above input ↔ compute ↔ output workflow typically does not end with one output step, but may be highly interactive with iterations between the steps, allowing the system to learn about the user and application context to guide.

8.1. Obtaining user input for guidance

For a facilitation system to decide when, what kind, and how to provide guidance to the user, input data about the user and context of the analysis process is required. Such input is obtained typically either *explicitly* and *implicitly*. We can also distinguish if the input is obtained exclusively from a single user, or from some group of users facing similar analysis tasks.

In the explicit case, the user proactively, or on request by the system, provides hints on the current analysis phase, information need, perceived relevance of data or views, etc. This is typically provided by interaction dialogues. An example is relevance feedback, where the relevance of selected views is rated by the user, which in turn may trigger a search for similar or dissimilar views to facilitate exploration. For example, the system of Behrisch et al. (2014) requests explicit user feedback on the perceived relevance of different scatter plots, and uses this to inform guidance for exploration. Explicit input data also include feedback collected from groups of users, e.g., collected in a distributed or crowd-based system. An interesting option, particularly for recurring analysis problems, is to build a feedback database from analysis sessions by different users, hence re-using the analysis information.

In the implicit case, the system relies on observations of the interactive analysis process and decides on the facilitation actions to take. Such observations can stem from usage logs taken from mouse and keyboard interaction, as the user operates the analysis system. For example, Brown et al. (2014) use low-level features of interaction (mouse movement statistics) to infer the user analysis performance. In systems supporting provenance collection and knowledge externalization (Section 5.3), the guiding subsystem may base its guidance decisions on analyzing user-generated materials (Fig. 2).

Besides classic interaction channels, a system may also use new interaction modalities and user sensing techniques, including eye tracking, stress and cognitive load measures, recognition of speech or facial expressions, gestures, or brain-computer interfaces. For example, Shao et al. (2017) used eye-tracking records of which areas of a scatter plot matrix have been explored so far, to inform guidance. Panwar and Collins (2018) use GSR sensing and eye tracking to detect user frustration to provide just-in-time guidance.

While all of these provide rich sources of input data for the system to decide on guidance, selecting and preprocessing appropriate feedback data for use with guidance algorithms is a challenge due to heterogeneity, size, and possible noise and uncertainties. Particularly in the explicit case, one also must consider the effort involved for the user to provide feedback. For an efficient analysis process, the cost of providing feedback must be traded off with the improvement brought by guidance.

8.2. Computing guidance steps

Given a user's real-time usage and interaction logs of the system, such as mouse movements, click logs, eye tracking, algorithms for intelligent facilitation determine (1) what to recommend (e.g., potentially useful data items to look at, new visualization views to provide, interactions to perform), (2) when to recommend (e.g., by identifying when a user is lost), and (3) what forms to take to recommend (e.g., passive non-intrusive suggestion or active replacement with a new view). These tasks involve the use of prior, session-specific, and situation knowledge, according to Fig. 2.

Algorithms that can perform these tasks can be developed in two directions: pre-defined *rule-based* and *learning-based*. The former, which relies on the expert knowledge, can be implemented as follows. As one of the simplest form, if a user explicitly asks for help

via a dedicated user interfaces for that purpose, the algorithm can provide active facilitation. If users do not do anything for a certain amount of time, it would be an indicator that he or she needs some facilitation. On the other hand, if the system detects that a user repeatedly performs the same thing (e.g., repeatedly invoking the same views), it may be another sign of facilitation needs.

As the second type of algorithms, one can cast this problem as a supervised learning problem based on a collection of past user logs, or direct/indirect feedback data. That is, user logs work as input features, while target variables to predict correspond to their implicit/explicit user feedback about whether they need facilitation at a particular moment in time and/or whether the provided facilitation was helpful or not. Once a sufficient amount of user logs are collected, such a formulation can open up a possibility to apply various state-of-the-art machine learning approaches. For example, a recently popular sequence-based prediction model called recurrent neural networks, e.g., long short-term memory, can be a good candidate to tackle our facilitation tasks. In the work of Behrisch et al. (2014), a decision tree classifier was trained from user relevance feedback and image features of candidate views to determine previously unseen but relevant views. Shao et al. (2017) used a similarity function defined for scatter plot views to recommend previously unseen scatter plot patterns from a scatter plot matrix, implementing guidance based on serendipity search (see also Section 3).

Recent work on *quality metrics* (or measures) for visualization can be used to inform both rule-based and learning-based guidance systems. A visual quality metric quantifies the expected usefulness of a data view, based on certain view properties. For example, if a view is highly cluttered (Ellis and Dix, 2007), it is unlikely that a user will be able to identify or compare any patterns. Hence, a measure of clutter can be seen as a quality measure. In recent studies, different quality measures have been proposed for many important visualization techniques (Behrisch et al., 2018a; Bertini, 2011). Quality measures can be used by a system in a rule-oriented manner by suggesting views of high quality measure. This can help to reduce time-consuming interactive searches for view parameters or data selections to create relevant views. Quality measures can also be used in a learning-oriented way, for example, learning to adapt which quality measures to use for a given analysis scenario.

In general, a drawback of machine learning is that it requires a significant amount of training data for a competitive performance, which can easily annoy or frustrate users. As ways to alleviate this problem, online and incremental learning, as well as transfer learning in this facilitation tasks, would be a promising research direction. Regarding quality measures, we note these are often heuristically defined and involve different parameters to set. Also, the kind of quality measure to use is a decision problem. A promising direction can also be information-theoretic approaches to assess the quality of views (Chen and Jänicke, 2010); a number of existing view quality measures implement information-theoretic concepts.

8.3. Modalities for guidance

Major modalities of providing assistance include textual or visual channels. Visual channels, such as color, highlighting, and animation, can provide different levels of attendance depending on which type of visual signal is applied. Textual information can provide more details, while a high attention cost may be required.

High-end immersive environments, such as large tiled display walls or CAVEs, can provide assistance to multiple users simultaneously (Klapperstueck et al., 2018). Recent advances in low-cost augmented-, virtual- and mixed-reality devices, such as Microsoft

Hololens, provide further opportunities for applying effective assistance with immersive environments in visual analytics applications (Stuerzlinger et al., 2018). In addition to the traditional channels for providing assistance in visualization and visual analytics, sound/voice (hearing), touch/motion (haptic) or other nontraditional sensory channels can provide effective assistance if properly used (McCormack et al., 2018). Natural language generation also holds promise for providing interpretative guidance for complex visualizations (Srinivasan et al., 2018).

Further research in this area could also include how to coordinate multi-modalities in challenging real applications, for example, analytics tasks situated in a difficult environment: operators in the field, on the factory floor, the hospital ward or operating theatre and so on (Thomas et al., 2018).

9. Validation of guidance

As with any visual analytics approaches, an artificial intelligence guided/facilitated visualization needs to be validated starting from its problem definition stage. Munzner's nested processing model for visualization design and validation (Munzner, 2009) and the seven evaluation scenarios of Lam et al. (2012) provide excellent guidelines for the design and validation of visualization. van Wijk (2013) and Carpendale (2008) provide an informative introduction to available evaluation techniques such as lab experiments, insight-based studies (North, 2006; North et al., 2011), and field studies (Shneiderman and Plaisant, 2006). Lam et al. (2012) provide valuable guidance on when to use which evaluation techniques. In this section, we focus on several unique challenges and potential solutions facing validation of guidance and readers can refer to the aforementioned articles for general approaches to visualization evaluation.

9.1. Objectives for guidance

Past work on evaluating recommender systems will inform our exploration of methods to validate the appropriateness of guidance, understand the impact of guidance on insights, and the potentially (de)biasing effects of guidance. For example, recommender systems are traditionally evaluated on the accuracy of the recommendation and whether it suits the needs of the user at the given moment and in a given context (i.e., is it accepted by the user). Newer models shifted away from solely providing accurate recommendations but rather a ranking of a set of recommendations taking into account not only their accuracy but also other objectives such as diversity and novelty. Specifically, Ge et al. (2010) propose the evaluation of recommender systems with regard to serendipity and coverage while Kaminskas and Bridge (2016) extend this evaluation design space with diversity and novelty. We will explore the parallels between these works and the concepts of guidance in visual analytics.

Coverage: In recommender systems, coverage refers to the degree to which recommendations cover the set of available items and the degree to which recommendations can be generated to all potential users (Ge et al., 2010). Ge et al. (2010) further define coverage as (1) prediction coverage, namely the percentage of the items for which the system is able to generate a recommendation, and (2) catalogue coverage, namely the percentage of the available items which could ever be recommended to a user. In the context of guidance, we can define prediction coverage and catalogue coverage as the percentage of insights relevant to a task that can be discovered from the views/interactions guidance can lead to, and the percentage of the relevant insights that can be discovered from the views/interactions that guidance can bring to a user. Having good coverage is essential to break the information bubble (Resnick et al., 2013) and potential bias caused by automated guidance. A

high catalogue coverage can help users find useful initial patterns and diverse alternative patterns in model construction even if they have an ill-defined starting point. A high predicted coverage reduces the risk of leaving useful information unexplored in the reasoning process. However, measuring coverage is rather challenging since the ground truth (e.g., the total number of relevant insights) is often unknown and its difficult to conduct exhaustive experiments to learn how many insights can be discovered in a guided approach. To address this challenge, heuristic approaches can be used, such as conducting an insight-based lab experiment (North, 2006; North et al., 2011) to prove that a guided approach has better coverage than a current practice, conducting multi-dimensional in-depth long-term case studies (Shneiderman and Plaisant, 2006) to collect expert opinions, or having algorithmic experiments to test how effectively and efficiently the guided system can access information from varying starting exploration points.

Diversity: A recommender system should propose a diverse range of different suggestions. In information retrieval systems, for example, offering a short list of only the closest matches to the search terms may lead to a very homogeneous list. A diverse list should include in the results list a variety of options even if they are different from the search terms (Carbonell and Goldstein, 1998), which allows for disambiguation of user queries (such as Jaguar being an animal or a car) and may lead to a higher chance of satisfying a user's expectations and needs. A broad catalogue coverage would support a diversity of recommendations, which should be extended across all suggestion mechanisms of a guidance system in order to encourage *serendipitous discovery*.

Serendipity: Serendipitous insights are those that arise unexpectedly yet prove valuable (Niu et al., 2018). Besides, serendipitous discoveries are often encountered in the natural human information-seeking process (Niu et al., 2018). North highlights the value of visualization as an analysis technique that provides serendipitous insight (North, 2006). Such "good surprises" are often a driving power for users to continue their exploration and should be encouraged by guidance. Niu et al. (2018) present an inspiring user study for serendipity evaluation of a computational module-facilitated health information system. Modeling serendipity as surprise and value (further decomposed as being useful and being interesting), they calculated the serendipity rating of artifacts as the aggregate of the three ratings of surprise, usefulness, and interestingness. In follow-up interviews, they asked whether the users encountered surprise, liked the surprise, and thought surprising results were interesting or useful, as well as analyzed whether the computational measures of surprise were correlated with user-perceived surprise.

Novelty: Kaminskas and Bridge (2016) define novelty as an objective of a recommender system in addition to serendipity. While closely related, novelty refers more generally to an item that was previously unknown to the user but does not necessarily carry the notion of being surprising. As Niu et al. (2018) describe it, novel items are discovered during a process of actively looking for new information whereas serendipitous items were discovered without actively looking for them. Serendipitous items are argued by Niu to be a subset of novel items in the discovery process. Diversity is important in the process of serendipitous discovery. In the context of information retrieval, novelty is a measure of difference and newness of a result compared to other results in the set. Analogously, in a visual analytics guidance system, novel suggestions would be different to previous suggestions and may extend the user's capabilities, broaden their knowledge base, or guide them in new analysis directions that were previously unexplored.

9.2. Evaluating guidance in the analytical reasoning process

Guidance should be an integrated part of the analytical reasoning process, which is difficult to be evaluated with simple tasks and quality assessing metrics. Complex tasks, such as information foraging tasks, often need to be used. The performance metrics are also multidimensional, which may include not only quality of the resulting artifacts, but also metrics about the exploration process itself, such as smoothness, transparency, and cognitive load of the exploration process. Here we highlight several practices that may inspire the design of guidance evaluation in the analytical reasoning process.

Employing complex information seeking tasks: For guidance aiming at goals such as training, engagement, and bias reduction, complex information seeking tasks, such as information foraging and browsing, may often provide deeper insights than simpler tasks such as factor finding. For example, [Willett et al. \(2007\)](#) employ an information foraging task for comparing prototypes with different levels of guidance. In this controlled experiment, subjects were asked to explore a dataset and collect evidence relevant to a task hypothesis assigned to them. Exploration history, discoveries, and user preference arising from this task provided rich insights about the prototypes evaluated.

Making use of exploration history: History tracking, which automatically records the exploration process of users, may provide very useful information for validation of guidance. In the problem definition stage, analyzing exploration history of existing practices may help designers accurately locate steps in the reasoning process where guidance is needed most. History can also help researchers evaluate alternative guidance approaches or figure out which, when, and how guidance features work. For example, [Hijikata \(2004\)](#) presents a user study that captured different types of mouse operations and compared the keywords checked by the subjects and the keywords extracted according to the mouse operations, to identify the most useful operations that can be used for guidance. [Shrinivasan and van Wijk \(2008\)](#) extract usage characteristics from exploration histories, and link them with free thoughts users recorded and comments in the follow-up interview, to discern the usefulness of different components in a guided visualization system.

Distinguishing different types of cognitive load: A plausible goal of guidance is to reduce the cognitive load. However, it would be overly simplistic to say that this should be a goal overriding all other concerns, since there are different types of cognitive load ([Sweller, 1988](#)), and not all of them are undesirable. Yet, extraneous cognitive load is an overhead that interferes with understanding. It is induced by system designs without sufficient consideration of the structure of information and the cognitive process ([Sweller, 1988](#)). Reducing extraneous cognitive load is a goal of facilitation. Meanwhile, germane cognitive load is desired since it represents users' efforts to process and comprehend the materials. It is devoted to schema acquisition and thus enhances learning ([Sweller, 1988](#)), which is an important goal of facilitation. [Liu et al. \(2013\)](#) argue that the indications of an effective exploratory visualization system are low extraneous cognitive load and high germane cognitive load. This statement is also true for guided visualization systems. In many practices ([Kang and Stasko, 2008](#)), cognitive load is measured by questionnaires such as the NASA TLX survey or through physiological measurements such as electroencephalography or pupil dilation measurements. With these approaches alone, it is difficult to identify extraneous cognitive load and germane cognitive load from the perceived cognitive load. To address this challenge, performance measures and user comments must be considered in cognitive load analysis. [Liu et al. \(2013\)](#) propose a practical approach that encodes user comments into three categories, namely *Engagement*, *Neutral*, and *Frustration*,

and uses their counts as indicators of which type of cognitive load dominating the perceived cognitive load.

Developing metrics for automated monitoring: While user studies are useful for the design and validation of guidance systems it is instrumental to quantify the effectiveness of guidance non-intrusively during real-world analytics tasks. Measurable attributes need to be defined and integrated into a model that quantifies the overall effectiveness of guidance. Such attributes may relate to system related parameters (e.g., item coverage and visualization design space) as well as user-related parameters, including user profiles, interaction histories, and cognitive load.

9.3. Avoiding harmful guidance

Counter to the benefits of guidance discussed throughout this paper, there is a risk that guidance systems may also be harmful. For example, mixed-initiative systems in the past have been found to be overly interruptive and distracting if they try to provide help when it is not wanted. In addition, provided assistance may be inappropriate ([Krisch, 2016](#)). If suggestions do not suit the user and context, this can, at best, be frustrating for the user. Worse, it could lead the analyst down an unhelpful, unimportant, or misleading path. This could be overly biasing, for example, guidance algorithms following a “more like this” method could keep an analyst stuck in the early conclusions, never revealing new data.

On a higher level, prevalent facilitation in data analysis could lead to *analytic atrophy*: if the analyst begins to rely too heavily on the suggestions of the system, they may fail to think critically in important situations. Due to these risks, we advocate for a human-in-the-loop process, whereby guidance is just one tool in the analyst's toolbox. Guidance systems should reveal the *reason* behind particular suggestions and offer clarity on the level of confidence the system has in the suggestion. The locus of control for the entire process must remain with the human analyst, allowing for guidance to be easily dismissed.

10. Conclusion

Data analysis is a non-trivial process, which often requires the use of multiple diverse tools, looking at data from different perspectives, and applying various transformations to the data. It may not always be obvious to an analyst what should be done at a given moment and what tool to use for that. Guidance is seen as a primary means to resolve this problem, which is called the “knowledge gap” ([Ceneda et al., 2017](#)). However, guidance can be beneficial to analysts even when they have no knowledge gap. It may be a means to make the analysis process more efficient, insights better validated, analysts more confident, bias avoided, and results clearly presented. In this paper, we have reviewed in a systematic way different aspects of guidance in the visual analysis process. While the earlier proposed conceptual model of guidance ([Ceneda et al., 2017](#)) is purely descriptive, we have tried to make our review more practically oriented by reasoning not only where in the analysis process and what kind of guidance can be provided but also how this can be implemented. In particular, what kinds of knowledge and information are needed for this.

To summarize, we have investigated the role of guidance in visualization and visual analytics, including the goals of guidance, the role in analytic processes and visualization tasks, implementation, and validation strategies. As interactive machine learning techniques and new interaction modalities become more integrated into visual analytics processes, we see many opportunities for deeper insights, easier workflows, and better accuracy from the human-machine analytic complex.

Acknowledgments

This paper is the direct result of an NII Shonan Meeting at the Shonan Village Center in Japan. We acknowledge the hospitality of the Center in making this research possible. This work was partly supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), [grant RGPIN-2015-03916], the Fraunhofer Cluster of Excellence on “Cognitive Internet Technologies” and by the EU through project Track&Know (grant agreement 780754).

References

- Andrienko, N., Lammarsch, T., Andrienko, G., Fuchs, G., Keim, D., Miksch, S., Rind, A., 2018. Viewing visual analytics as model building. *Comput. Graph. Forum* 37 (6), 275–299.
- Babu Vallabhaneni, R., Rajesh, V., 2017. Brain tumour detection using mean shift clustering and GLCM features with edge adaptive total variation denoising technique. *Alexandria Eng. J.* 12, 666–671.
- Behrisch, M., Blumenschein, M., Kim, N., Shao, L., El-Assady, M., Fuchs, J., Seebacher, D., Diehl, A., Brandes, U., Pfister, H., Schreck, T., Weiskopf, D., Keim, D., 2018a. Quality metrics for information visualization. *Comput. Graph. Forum (EuroVis State of the Art Report)* 37 (3), 652–662.
- Behrisch, M., Korkmaz, F., Shao, L., Schreck, T., 2014. Feedback-driven interactive exploration of large multidimensional data supported by visual classifier. In: *IEEE Conf. on Visual Analytics Science and Technology*. pp. 43–52.
- Behrisch, M., Streeb, D., Stoffel, F., Seebacher, D., Matejek, B., Weber, S.H., Mittelstaedt, S., Pfister, H., Keim, D., 2018b. Commercial visual analytics systems—advances in the big data analytics field. *IEEE Trans. Vis. Comput. Graphics*.
- Bertini, E., 2011. Quality metrics in high-dimensional data visualization: An overview and systematization. *IEEE Trans. Vis. Comput. Graphics* 17 (12), 2203–2212. <http://dx.doi.org/10.1109/TVCG.2011.229>.
- Boy, J., Eveillard, L., Detienne, F., Fekete, J.D., 2016. Suggested interactivity: Seeking perceived affordances for information visualization. *IEEE Trans. Vis. Comput. Graphics* 22 (1). <http://dx.doi.org/10.1109/TVCG.2015.2467201>.
- Brehmer, M., Munzner, T., 2013. A multi-level typology of abstract visualization tasks. *IEEE Trans. Vis. Comput. Graphics* 19 (12), 2376–2385.
- Brown, E.T., Ottley, A., Zhao, H., Lin, Q., Souvenir, R., Ender, A., Chang, R., 2014. Finding Waldo: Learning about Users from their Interactions. *IEEE Trans. Vis. Comput. Graphics* 20 (12), 1663–1672.
- Carbonell, J., Goldstein, J., 1998. The use of MMR, diversity-based reranking for reordering documents and producing summaries. In: *Proc. of the Int. ACM SIGIR Conf. on Research and Development in Information Retrieval*. pp. 335–336.
- Carpendale, S., 2008. Evaluating information visualizations. In: *Information Visualization*. Springer, pp. 19–45.
- Ceneda, D., Gschwandtner, T., May, T., Miksch, S., Schulz, H.-J., Streit, M., Tominski, C., 2017. Characterizing guidance in visual analytics. *IEEE Trans. Vis. Comput. Graphics* 23 (1), 111–120.
- Cha, M.-H., 2007. Automatic ordering method and system for trading of stock, bond, item, future index, option, index, current and so on, US Patent 7,181,425.
- Chen, Y., Barlowe, S., Yang, J., 2010. Click2Annotate: Automated Insight Externalization with rich semantics. In: *2010 IEEE Symp. on Visual Analytics Science and Technology*. pp. 155–162.
- Chen, M., Jänicke, H., 2010. An information-theoretic framework for visualization. *IEEE Trans. Vis. Comput. Graphics* 16 (6), 1206–1215. <http://dx.doi.org/10.1109/TVCG.2010.132>.
- Conati, C., Hoque, E., Toker, D., Steichen, B., 2013. When to adapt: Detecting user’s confusion during visualization processing. In: *Proc. of the Int. Workshop on User-Adaptive Visualization*.
- Cook, K.A., Thomas, J.J., 2005. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*, Tech. Rep. Pacific Northwest National Lab.
- Dörk, M., Feng, P., Collins, C., Carpendale, S., 2013. Critical InfoVis: Exploring the politics of visualization. In: *Extended Abstracts on Human Factors in Computing Systems (Alt.CHI)*.
- Dwyer, T., 2016. Network visualization as a higher-order visual analysis tool. *IEEE Comput. Graph. Appl.* 36 (6), 78–85. <http://dx.doi.org/10.1109/MCG.2016.117>.
- Eccles, R., Kapler, T., Harper, R., Wright, W., 2008. Stories in GeoTime. *Inf. Vis.* 7 (1), 3–17.
- Ellis, G.P., Dix, A.J., 2007. A taxonomy of clutter reduction for information visualisation. *IEEE Trans. Vis. Comput. Graphics* 13 (6), 1216–1223. <http://dx.doi.org/10.1109/TVCG.2007.70535>.
- Federico, P., Wagner, M., Rind, A., Amor-Amorós, A., Miksch, S., Aigner, W., 2017. The role of explicit knowledge: A conceptual model of knowledge-assisted visual analytics. In: *Proc. of the IEEE Conf. on Visual Analytics Science and Technology*.
- Ge, M., Delgado-Battenfeld, C., Jannach, D., 2010. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In: *Proc. ACM Conf. on Recommender Systems*. ACM Press.
- Google, 2018. See and use suggested and analysis in a spreadsheet. <https://support.google.com/docs/answer/6280499>.
- Gotz, D., Zhou, M.X., 2009. Characterizing users’ visual analytic activity for insight provenance. *Inf. Vis.* 8 (1), 42–55.
- Gotz, D., Zhou, M.X., Aggarwal, V., 2006. Interactive visual synthesis of analytic knowledge. In: *IEEE Symp. on Visual Analytics Science and Technology*. pp. 51–58.
- Gramazio, C.C., Laidlaw, D.H., Schloss, K.B., 2016. Colorgical: creating discriminable and preferable color palettes for information visualization. *IEEE Trans. Vis. Comput. Graphics* 23 (1), 521–530.
- Gratzl, S., Lex, A., Gehlenborg, N., Cosgrove, N., Streit, M., 2016. From visual exploration to storytelling and back again. *Comput. Graph. Forum* 35 (3), 491–500. <http://dx.doi.org/10.1111/cgf.12925>.
- Grinstein, G., Whiting, M., Liggett, K., Nebesh, D., 2011. IEEE VAST Challenge 2011, 2011. <http://hcil.cs.umd.edu/localphp/hcil/vast11/>. (Last Accessed 29 January 2014).
- Healey, C.G., Dennis, B.M., 2012. Interest driven navigation in visualization. *IEEE Trans. Vis. Comput. Graphics* 18 (10), 1744–1756.
- Heer, J., Agrawala, M., 2006. Multi-Scale banking to 45 degrees. *IEEE Trans. Vis. Comput. Graphics* 12 (5), 701–708.
- Hijikata, Y., 2004. Implicit user profiling for on demand relevance feedback. In: *Proc. of the Int. Conf. on Intelligent User Interfaces*. pp. 198–205.
- Kaminskas, M., Bridge, D., 2016. Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Trans. Interact. Intell. Syst.* 7 (1), 1–42. <http://dx.doi.org/10.1145/2926720>.
- Kang, Y.-a., Stasko, J., 2008. Lightweight task/application performance using single versus multiple monitors: A comparative study. In: *Proc. of Graphics Interface 2008*. pp. 17–24.
- Klapperstueck, M., Czauderna, T., Goncu, C., Glowacki, J., Dwyer, T., Schreiber, F., Marriott, K., 2018. ContextuWall: Multi-site collaboration using display walls. *J. Vis. Lang. Comput.* 46, 35–42.
- Krisch, J.A., 2016. Why You Hated Clippy, That Annoying Microsoft Paperclip, 2016. <https://www.vocativ.com/319986/clippy-digital-assistant/index.html>, (last checked 08 June 2018).
- Lam, H., Bertini, E., Isenberg, P., Pleasant, C., Carpendale, S., 2012. Empirical studies in information visualization: Seven scenarios. *IEEE Trans. Vis. Comput. Graphics* 18 (9), 1520–1536.
- Liu, Y., Barlowe, S., Feng, Y., Yang, J., Jiang, M., 2013. Evaluating exploratory visualization systems: A user study on how clustering-based visualization systems support information seeking from large document collections. *Inf. Vis.* 12 (1), 25–43.
- Matejka, J., Li, W., Grossman, T., Fitzmaurice, G., 2009. CommunityCommands: Command recommendations for software applications. In: *Proc. ACM Symp. on User Interface Software and Technology*. pp. 193–202.
- McCormack, J., Roberts, J.C., Bach, B., Freitas, C.D.S., Itoh, T., Hurter, C., Marriott, K., 2018. Multisensory immersive analytics. In: Marriott, K., Schreiber, F., Dwyer, T., Klein, K., Riche, N.H., Itoh, T., Stuerzlinger, W., Thomas, B.H. (Eds.), *Immersive Analytics*. Springer, ISBN: 978-3-030-01388-2 (in press).
- McDuff, D., Karlson, A., Kapoor, A., Roseway, A., Czerwinski, M., 2012. AffectAura: An intelligent system for emotional memory. In: *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. pp. 849–858.
- Mehta, H., Bradley, A.J., Hancock, M., Collins, C., 2017. Metatation: Annotation as implicit interaction to bridge close and distant reading. *ACM Trans. Comput.-Hum. Interact.* 35 1–41.
- Munzner, T., 2009. A nested process model for visualization design and validation. *IEEE Trans. Vis. Comput. Graphics* 15 (6), 921–928.
- Niu, X., Abbas, F., Maher, M.L., Grace, K., 2018. Surprise me if you can: Serendipity in health information. In: *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. pp. 23:1–23:12. <http://dx.doi.org/10.1145/3173574.3173597>.
- North, C., 2006. Toward measuring visualization insight. *IEEE Comput. Graph. Appl.* 26 (3), 6–9.
- North, C., Saraiya, P., Duca, K., 2011. A comparison of benchmark task and insight evaluation methods for information visualization. *Inf. Vis.* 10 (3), 162–181.
- Panwar, P., Collins, C.M., 2018. Detecting negative emotion for mixed initiative visual analytics. In: *Extended Abstracts of the 2018 CHI Conf. on Human Factors in Computing Systems*.
- Resnick, P., Garrett, R.K., Kriplean, T., Munson, S.A., Stroud, N.J., 2013. Bursting your (filter) bubble: Strategies for promoting diverse exposure. In: *Proc. of the Conf. on Computer Supported Cooperative Work Companion*. pp. 95–100. <http://dx.doi.org/10.1145/2441955.2441981>.
- Ruchikachorn, P., Mueller, K., 2015. Learning visualizations by analogy: Promoting visual literacy through visualization morphing. *IEEE Trans. Vis. Comput. Graphics* 21 (9), 1028–1044. <http://dx.doi.org/10.1109/TVCG.2015.2413786>. <http://ieeexplore.ieee.org/document/7061477>.
- Sacha, D., Stoffel, A., Stoffel, F., Kwon, K., Ellis, G., Keim, D.A., 2014. Knowledge generation model for visual analytics. *IEEE Trans. Vis. Comput. Graphics* 20 (12), 1604–1613.
- Shao, L., Silva, N., Eggeling, E., Schreck, T., 2017. Visual exploration of large scatter plot matrices by pattern recommendation based on eye tracking. In: *Proc. ACM IUI Workshop on Exploratory Search and Interactive Data Analytics*.

- Shneiderman, B., Plaisant, C., 2006. Strategies for evaluating information visualization tools: Multi-dimensional in-depth long-term case studies. In: Proc. of the 2006 AVI Workshop on beyond Time and Errors: Novel Evaluation Methods for Information Visualization. ACM, pp. 1–7.
- Shrinivasan, Y., van Wijk, J., 2008. Supporting the analytical reasoning process in information visualization. In: Proc. of the SIGCHI Conference on Human Factors in Computing Systems. ACM New York, NY, USA, pp. 1237–1246.
- Srinivasan, A., Drucker, S.M., Endert, A., Stasko, J., 2018. Augmenting visualizations with interactive data facts to facilitate interpretation and communication. In: Proc. IEEE InfoVis, IEEE Trans. Vis. Comput. Graphics. <http://dx.doi.org/10.1109/TVCG.2018.2865145> (in press).
- Stuerzlinger, W., Dwyer, T., Drucker, S., Görg, C., North, C., Scheuermann, G., 2018. Immersive human-centered computational analytics. In: Marriott, K., Schreiber, F., Dwyer, T., Klein, K., Riche, N.H., Itoh, T., Stuerzlinger, W., Thomas, B.H. (Eds.), Immersive Analytics. Springer, ISBN: 978-3-030-01388-2 (in press).
- Sweller, J., 1988. Cognitive load during problem solving: Effects on learning. *Cognitive Science* 12 (2), 257–285.
- Tatu, A., Albuquerque, G., Eisemann, M., Bak, P., Theisel, H., Magnor, M., Keim, D., 2011. Automated analytical methods to support visual exploration of high-dimensional data. *IEEE Trans. Vis. Comput. Graph.* 17 (5), 584–597.
- te Brake, G.M., Karssemeijer, N., Hendriks, J.H., 2000. An automatic method to discriminate malignant masses from normal tissue in digital mammograms. *Phys. Med. Biol.* 45 (10), 2843.
- Thomas, B.H., Welch, G.F., Dragicevic, P., Elmqvist, N., Irani, P., Jansen, Y., Schmalstieg, D., Tabard, A., ElSayed, N.A.M., Smith, R.T., Willett, W., 2018. Situated analytics. In: Marriott, K., Schreiber, F., Dwyer, T., Klein, K., Riche, N.H., Itoh, T., Stuerzlinger, W., Thomas, B.H. (Eds.), Immersive Analytics. Springer, ISBN: 978-3-030-01388-2 (in press).
- Toker, D., Conati, C., 2014. Eye tracking to understand user differences in visualization processing with highlighting interventions. In: Proc. User Modeling, Adaptation, and Personalization. Springer, pp. 219–230.
- Veras, R., Collins, C., 2017. Optimizing hierarchical visualizations with the minimum description length principle. *IEEE Trans. Vis. Comput. Graph.* 23 (1), 631–640.
- Walker, R., ap Cenyyd, L., Pop, S., Miles, H.C., Hughes, C.J., Teahan, W.J., Roberts, J.C., 2015. Storyboarding for visual analytics. *Inf. Vis.* 14 (1), 27–50. <http://dx.doi.org/10.1177/1473871613487089>.
- Walker, R., Slingsby, A., Dykes, J., Xu, K., Wood, J., Nguyen, P.H., Stephens, D., Wong, B.W., Zheng, Y., 2013. An extensible framework for provenance in human terrain visual analytics. *IEEE Trans. Vis. Comput. Graphics* 19 (12), 2139–2148. <http://dx.doi.org/10.1109/TVCG.2013.132>.
- Wall, E., Blaha, L.M., Franklin, L., Endert, A., 2017a. Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics. In: Proc. IEEE Visual Analytics Science and Technology.
- Wall, E., Blaha, L., Paul, C.L., Cook, K., Endert, A., 2017b. Four perspectives on human bias in visual analytics. In: DECISIVE: Workshop on Dealing with Cognitive Biases in Visualizations.
- Wang, J., Mueller, K., 2017. Visual causality analysis made practical. In: Proc. IEEE Conf. on Visual Analytics Science and Technology (VAST).
- van Wijk, J.J., 2006. Views on visualization. *IEEE Trans. Vis. Comput. Graphics* 12 (4), 421–432.
- van Wijk, J.J., 2013. Evaluation: A challenge for visual analytics. *Computer* 46 (7), 56–60.
- Willett, W., Heer, J., Agrawala, M., 2007. Scented widgets: improving navigation cues with embedded visualizations. *IEEE Trans. Vis. Comput. Graphics* 13 (6), 1129–1136.
- Wongsuphasawat, K., Moritz, D., Anand, A., Mackinlay, J., Howe, B., Heer, J., 2016. Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE Trans. Vis. Comput. Graphics* 22 (1), 649–658.
- Wongsuphasawat, K., Qu, Z., Moritz, D., Chang, R., Ouk, F., Anand, A., Mackinlay, J., Howe, B., Heer, J., 2017. Voyager 2: Augmenting visual analysis with partial view specifications. In: Proc. of the SIGCHI Conf. on Human Factors in Computing Systems. ACM, pp. 2648–2659.
- Wright, W., Schroh, D., Proulx, P., Skaburskis, A., Cort, B., 2006. The sandbox for analysis: Concepts and methods. In: Proc. of the SIGCHI Conf. on Human Factors in Computing Systems. ACM, pp. 801–810.
- Xu, K., Attfield, S., Jankun-Kelly, T., Wheat, A., Nguyen, P.H., Selvaraj, N., 2015. Analytic provenance for sensemaking: A research agenda. *IEEE Comput. Graph. Appl.* 35 (3), 56–64. <http://dx.doi.org/10.1109/MCG.2015.50>.
- Yi, J.S., Kang, Y.a., Stasko, J., Jacko, J., 2007. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Trans. Vis. Comput. Graphics* 13 (6), 1224–1231.
- Zhao, J., Glueck, M., Isenberg, P., Chevalier, F., Khan, A., 2018. Supporting handoff in asynchronous collaborative sensemaking using knowledge-transfer graphs. *IEEE Trans. Vis. Comput. Graphics* 24 (1), 340–350. <http://dx.doi.org/10.1109/TVCG.2017.2745279>.