ISSN 2186-7437

NII Shonan Meeting Report

No. 2018-2

Visual Analytics: Towards Human Machine Intelligence

Yingcai Wu
Nan Cao
David Gotz
Steffen Koch

January 29–February 1, 2018

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Visual Analytics: Towards Human Machine Intelligence

Organizers:
Yingcai Wu (Zhejiang University)
Nan Cao (Tongji University)
David Gotz (UNC-Chapel Hill)
Steffen Koch (University of Stuttgart)

January 29–February 1, 2018

Data mining and visualization have attracted considerable attention in recent years for exploring and understanding big data. Machine learning focuses on developing automatic algorithms to discover patterns in large data sets. Although huge successes have been achieved, existing approaches usually assume that a ground truth is readily available. In practice, this is often not the case. In some cases, manual data annotation is required, which is tedious, onerous at scale, and highly dependent on the judgment of the human annotators. A ground truth understanding of a dataset may simply not exist, such that it can be difficult if not impossible to mathematically model the unknown types of patterns we hope to find. Even when patterns can be modeled, the intuitive explanation and validation of the models can pose a major challenge. Moreover, the underspecified complex tasks with a very high dimensional space of input variables and parameters cannot be simply handled without the inclusion of human expertise and knowledge. Compared with data mining, visualization aims to produce intuitive visual representations of data. It allows people to quickly see and interact with the patterns in the data by making effective use of their high-bandwidth visual system. As an old saying goes a picture worth a thousand words, a good visualization will significantly improve the abilities of people to understand and interpret the data and analysis results. However, the enormous amount of complex data leads to the difficulty of creating concise, discernable, and intuitive visual representations. Visual summaries of big data can still easily overwhelm users.

To make best use of the advantages and bypass the disadvantages of data mining and visualization, visual analytics has recently been introduced to facilitate analytical reasoning by interactive visual interfaces. It presents data and analysis results in context, and thus, it can provide rich evidence that supports or contradicts the analysis results, and consequently, help with data interpretation and result validation. Analysts can annotate on (e.g., place labels on) or adjust results via interactive visualizations to supervise the underlying analysis procedure, for example, and thus, gradually produce increasingly precise analysis and correct results. Visual analytics have been used in many applications to tackle various important problems, such as tackling urban issues like traffic jam and air pollution, making better diagnostic and treatment decisions, preventing threats and fraud in business, optimizing rescue efforts, forecasting severe weather conditions, and achieving situational awareness during crisis.

We believe that visual analytics can enable human-centric computational intelligence by effectively integrating human knowledge and expertise into powerful computational algorithms through a high-bandwidth visual processing channel and user interactions. Despite recent impressive advances, designing developing effective visual analytics for big data still poses significant challenges for researchers and practitioners. There are still many research opportunities and open questions that we should address for creating visual analytics to enable effective human-machine intelligence.

In this Shonan Meeting we identified research opportunities as well as critical problems that can occur if applying visual analytics to enable human-machine intelligence. After hearing different incentive talks and a subsequent collection of important research questions. All participants of the workshop decided on topics to focus on during the course of the meeting. The three topics were discussed in the remaining days extensively in respective working groups:

- Guidance in the Human-Machine Analysis Process (working group 1)
- Steering Data Quality with Interactive Visualization (working group 2)
- Bias and Trust in Visual Analytics: Challenges and Opportunities for Effective Human Machine Intelligence (working group 3)

The outcome of the group discussions is described in more detail below. In addition, this report contains the titles and abstract of the given talks, a list of participants, and the workshop's schedule.

1 Guidance in the Human-Machine Analysis Process: Envisioning an intelligent computer assistant to facilitate an intelligent analyst

Authors

Christopher Collins, Tim Dwyer, Natalia Andrienko, Jaegul Choo, Tobias Schreck, Jing Yang, and Xiaoru Yuan

1.1 Motivation

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 tasks. Some of these tasks are relatively low-level and it is fairly 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. Thus, rather than replacing the human data analyst, we expect that sophisticated artificial intelligence will play a support, guidance or facilitation role in an interactive data analytics process for the foreseeable future.

In this paper we list the goals for and both the pros and cons of guidance, and we discuss the role that such machine facilitation can play in the key tasks from visualisation and also the more sophisticated model-generation tasks of visual analytics.

1.2 Relevant work

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. [5] 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 Wijks [16] model of visualization, presenting opportunities for guidance at a high-level in the process diagram. We plan to investigate the potential role of facilitation across lower-level task taxonomies, e.g. Brehmer and Munzner [3] and more sophisticated models of visual analysis [2]. Specific instantiations of guidance have been reported, for example, helpful interventions when eye-tracking indicates an analyst is exhibiting signs of confusion [6, 14] or when a logging system detects sub-optimal search strategies [4]. There have also been investigations into the role of machine intelligence in exposing potential bias in an analysis process [18], by exposing the differences between the data a user has seen and the overall character 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 [7]. Our exploration of the role of facilitation will acknowledge the potential pros and cons of each form of facilitation with examples from the literature where appropriate.

1.3 Background/ concepts/ terminology

The term *guidance* refers to providing help to the user when the user experiences difficulties, e.g., does not know what 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. 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.

1.4 Goals for machine facilitation

To discuss guidance and facilitation in a systematic way, we begin with defining the goals of guidance/facilitation. Examples include: to inform, to reduce cognitive load, to improve tool usability for novices, to avoid or reduce bias, to verify findings, and others. Then we define the aspects of guidance from the front end and back end perspectives. At the front end, these are the form, medium, ways of interaction and communication between the human and the computer, and the integration in the analytical process. At the back end, the aspects include the content, information sources, algorithmic implementation, inputs and outputs. An important aspect to consider is the evaluation of guidance: how to measure the improvements due to using guidance compared to non-guided analysis?

We posit the following as a manifesto of goals for machine facilitation:

- To avoid bias
- To inform

- To reduce cognitive load
- To make analysis more efficient
- To capture provenance
- To improve usability for novices
- For training
- Experience transfer, allow novices to perform expert level analysis
- Engagement
- Hypothesis testing
- Verification of findings
- Testing stability and sensitivity of findings
- To refine results and enhance discovery
- Should never be harmful
- Should suggest rather than prescribe

1.5 Background knowledge and capabilities of an intelligent guide/assistant

Envisioning an intelligent guide or assistant, we reason about the knowledge and capabilities that are required for fulfilling the expected functions. An intelligent assistant would need to have a knowledge base of (1) data types and structures, (2) possible relationships among data components, (3) possible patterns, such as trend and seasonality in time series, (4) existing visualisation and analysis methods, their applicability to data types and their capability to exhibit or detect patterns and relationships, (5) user actions, possible intentions, and possible difficulties. When the analysis starts, the analyst should understand the structure and properties of the loaded data, anticipate patterns and relationships that may exist, be able to find sources of additional related data. In the process of analysis, the assistant should be able to track the process, automate collecting provenance, understand current situation and anticipate further steps. The assistant needs to have an adaptive and growing understanding of the users intentions as they perform analysis.

1.6 Facilitation for the visual analytics process

The visual analytics process can be seen as the process of deriving a mental or formal model of some subject, which includes repeated evaluation and development of the currently existing model until it satisfies the criteria of correctness, coverage, generality, specificity, and fitness to the purpose. Intelligent support is to be provided for generation of an initial model, evaluation, improvement of the model, collection of provenance, and externalization of knowledge gained. The support includes informing about data properties, automated extraction of patterns and relationships, suggestion of methods and parameters, warning about data portions and components not covered, suggestion of annotation templates, and others.

1.7 Ways of getting feedback

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. Main types of such input are explicit and implicit. In the explicit case, the user proactively, or on request by the system, provides hints on the currently 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. Explicit input data also includes feedback collected from groups of users, e.g., collected in a distributed or crowd-based system. 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. In addition, new interaction modalities or user monitoring techniques including eye tracking, stress and cognitive load measures, recognition of speech or facial expressions, or brain-computer interfaces can be considered. 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 remains a challenge due to heterogeneity, size, and possible noise and uncertainties, especially in the implicit case.

1.8 Modalities for providing assistance

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.

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.

High-end immersive environments, such as large tiled display walls or CAVEs, can provide assistance to multiple users simultaneously. Recent advancement in low-cost augmented-, virtual- and mixed-reality devices, such as Microsoft Hololens, provides further opportunities for applying effective assistance with immersive environment in visual analytical applications. Further research in this area could also include how to coordinate multi-modalities in challenging real applications.

1.9 Possible ways to Implement

Given a users real-time usage and interaction logs of the system, such as mouse movements, click logs, eye tracking, algorithms for intelligent facilitation determines (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 (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).

1.10 Validation of Facilitation

Evaluating the impact of facilitation is an open research problem. Past work on evaluating recommendation systems will inform our exploration of methods to validate the appropriateness of guidance, understand the impact of guidance on insights, and the potentially (de)biassing effects of facilitation. For example, recommendation systems are traditionally evaluated on the accuracy of the recommendation: does the provided suggestion suit the needs of the user at the given moment (i.e. is it accepted by the user). Newer models of validation argue that recommendations can also be valuable if they increase the coverage of the users knowledge of the data space, or if they increase the occurrence of serendipitous discovery [8, 9]. We will explore the parallels from this work to the concepts of guidance in visual analytics.

2 Steering Data Quality with Interactive Visualization

Participants

Gennady Andrienko, Nan Cao, Seokhee Hong, Shixia Liu, Conglei Shi, Yu-Shuen Wang, Yingcai Wu

2.1 Motivation

Data-centric approaches have been increasingly prevalent to variety of problems in different domains. During the data collecting stage, some inaccurate or error data can be included, which may affect the further usage of data. Thus, one key issue on preparing and processing data is how to ensure the data quality. Data cleansing techniques has played a critical role to achieve this.

Data cleansing has been researched for a long time in database and knowledge discovery area [1, 15]. Recently, there are several works in human computer interaction on data cleansing in tabular data [12, 11]. However, due to the increasing complexity of the data collected through variety of ways (e.g., GPS data, text, video), it is more and more challenging to effectively and accurately cleanse the data. In most of the cases, domain knowledge from experts is important to guide for better performance of data cleansing algorithms. Therefore, it is of great interest to study how to better combine user-guided methods with system-guide methods during the analysis, and information visualization and visual analytics are the important part to achieve this goal.

In the paper, we first report the related work from different research areas. Based on that, we propose a general visual analytics framework of data cleansing. We then discuss the challenges and opportunities from three aspects: human related, data related, and task related.

2.2 Related Work

For the past two decades, researchers have been extensively studying on data cleansing. Abedjan et. al., summarized the data error detection algorithms and well established data cleansing tools [1]. Broeck et. al., proposed a general three-stage framework on data cleansing, including screen stage, diagnosis stage,

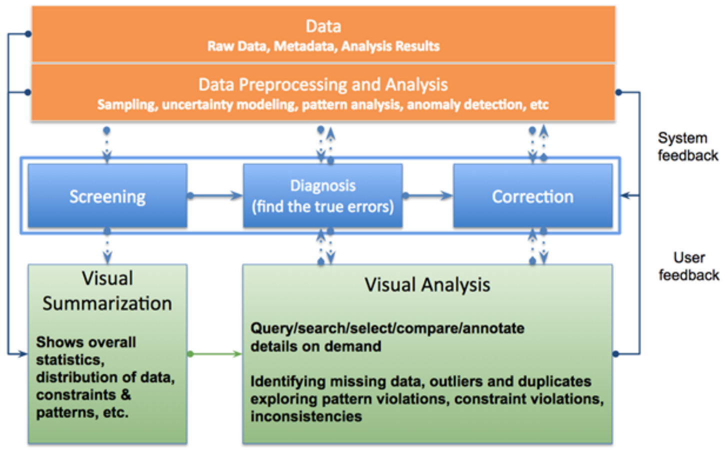

Figure 1: General Data Cleansing Pipeline

and correction stage [15], which is a great summary of user workflow on data cleansing. There are also several works focusing on interactive data cleansing tool design. Profiler [12] was designed to interactively detect and visually summarize the outliers from data. Wrangler [11] targeted on interactively creating data transformation scripts. One major drawback for both tools is that only tabular data is supported. Beyond these work, Sean et. al., further summarized the research direction on how visualizations and interaction techniques can help data wrangling [10]. However, there is a lack of high level abstraction on the generic framework on designing a visual analytic system for data cleansing.

In our work, we adapt the three-stage framework [15], and propose a visual analytics framework, focusing on how interactive visualizations can help on each stage. Inspired by [10], we propose and align the challenges and research opportunities with the our visual analytics framework, which we hope can better guide the visual analytics research on data cleansing.

2.3 Visual Analytics Framework

We adapt the three stages of data cleansing aforementioned and propose the visual analytics framework on data cleansing, shown in the Figure 1. The raw data is first processed by some data mining algorithms, such as sampling, pattern analysis, and anomaly detection. The output of the analysis result together with the raw data are the input of the three data cleansing stages: screening, diagnosis, and correction. In the screening stage, some visual summarization techniques are applied for users to quickly gain insight into the data and find potential errors. In diagnosis stage, users can then focus on analyzing them by interactively exploring related data in more depth to figure out the actual errors. In correction stage, users can apply changes on data to correct the errors. They can modify the raw data, analysis methods, and / or the analysis result. This framework can apply to different types of data, and we plan to use several concrete examples (multimedia data, textual data, trajectory data, and network data) to show how the framework works.

2.4 Challenges

Based on the proposed visual analytics framework, we identify the challenges brought by the complexity of human, data, and, tasks. In this section, we summarize all the challenges below.

2.4.1 Human complexity

1. **Lack of domain knowledge**
   Better integrating domain knowledge is key to success of the system design. However, in some cases, analysts are lack of sufficient knowledge or expertise regarding new types of data or new data sets).

2. **Limitations of perception/cognition**
   There have been thorough studies on humans limitation of visual perception and cognition. For complicated types of data, it is challenging to design a visual analytic system while keeping the complexity of the system within the perception limitation.

2.4.2 Data complexity

1. **Veracity**
   Veracity refers to the trustworthiness of the data, and it is the key goal for the data cleansing. The data errors (e.g., incorrect values and missing fields) introduced by data collecting are needed to be corrected.

2. **Volume**
   When the data volume continues increasing, it poses challenges of the stability on both computing methods and visual representations. The interaction design will be also affected.

3. **Velocity**
   In many cases, new data are generated constantly. For real-time analysis, it is challenging to handle especially high-frequency and high-volume data stream.

4. **Variety**
   Different types of data (e.g., table, text, image and network) and data fusion methods further introduce the complexity of the system design.

2.4.3 Task complexity

1. **Generalizability**
   Different analysis tasks and methods have different requirements to data quality. A data set may be suitable for method A but not suitable for method B (including Domain Complexity).

2. **Uncertainty**
   Uncertainty exists in data cleansing results. This should be taken into account by analysis methods.

3. **Task dependency**
   The analysis tasks are from each step of the framework, and these tasks

have dependencies among them. This fact increases the difficulty in choosing the suitable methods for analysis tasks.

2.5 Research Opportunities

We summarize the research opportunities also from the same three aspects for creating a visualization-driven framework for visually steering data quality.

2.5.1 Human-related opportunities

1. How to tackle knowledge integration and calibration?
2. How to scale up through collaborative visual interfaces / progressive methods?
3. How to conduct mixed initiative guidance (system initiative guidance and user initiative guidance)?
4. How to evaluate the VA driven data cleansing system (quality metrics, human experiment based on tasks: measuring time, error, and intuitiveness?)?

2.5.2 Data-related opportunities

1. How to correct data efficiently?
2. How to transform data for facilitating problem detection and correction?
3. How to derive quality metrics for different types of data (e.g., text, image, video, and network)?
4. How to handle the data at all scales? (e.g., terabyte and petabyte?)
5. How to visualize different types of data in an integrated interface?

2.5.3 Task-related opportunities

1. How to relate data quality to requirements of different analysis tasks and methods?
2. How to design visualization and visual analysis approaches for screening, diagnosis, visualization/visual analytic, and correction?
   (a) How to visualize summarization of screening results effectively and efficiently?
   (b) How to make the data processing algorithm transparent to users?
   (c) How to visualize diagnosis results intuitively to ensure trustworthy analysis and enable more effective human decision making?
   (d) How to design interactive, easy-to-use system to better support correction step?

3 Bias and Trust in Visual Analytics: Challenges and Opportunities for Effective Human Machine Intelligence

Participants

David Gotz, Steffen Koch, Zhicheng "Leo" Liu, Ross Maciejewski, Benjamin Renoust, Guodao Sun, Jing Yang, Ye Zhao

3.1 Introduction

User trust in a system is a key condition for a successful relationship between humans and machines. Building this trust is especially critical in Visual Analytics which places user interaction at the center of the system (Human-in-the-loop). The rise in Machine Learning performances makes it unavoidable for the design of visual analytics system, although it often remains a black box of which the output cannot always be explained.

The construction of trustworthy relationship requires a delicate balance while both ends could show some type of bias. Any such bias would pose a threat to this construction of trust, reducing the effectiveness of the human-machine intelligence process. The outcome of this group discussion is a reflection on how trust is positively or negatively impacted during the whole visual analytics process, while further identifying sources of bias.

3.2 Trust

Trust has been widely defined by sociology and psychology, but we will restrain its definition to the one defined for systems [17]. Trust is a relationship between two components, in our case, a system and a user. A trusted system should display a set of properties that could be relied on, upon which a user can in turn execute tasks correctly. Any violation of these properties would impede the user's performance.

While systems are built upon models that may not always exhaustively capture the required properties for users to achieve tasks at full performance, users can still rely and build trust on partial information. An everyday example would be a guidance system not having all pedestrian paths registered: a user having the knowledge of a specific path would be able to adjust a proposed route walking this specific path, while still relying on the system after updating the rest of the route.

Trust is then a context dependent phenomena. Different users aim at achieving different goals relative to the tasks a visual analytics is designed for. We will focus on the elements that can break trust or help build/recover trust in our context of visual analytics.

We identified a few elements that would erode trust. Visual analytics is dynamic by nature, forcing users to interact and often change representation, and test different hypotheses that may negatively impact its relationship with the system. This would be even worse if the proposed system lacks of stability, or if experiments could not be exactly repeated. Putting the "Human in the loop" brings the knowledge extracted and presented by the system in contrast to the Human's knowledge and expertise. A conflict emerges when there is a mismatch between what a user knows and what a system represents, regardless of whether the system is right or wrong. For example, statistical significance and importance to users do not always match: it could be the representation of topic models that did not match user expectations, or even medical codes that are not clinically relevant. Although statistically correct, these examples erode user trust in the significance of the rest of the results. When training a model, one should be careful of the distance between the training data and the application data, which could be far apart. This points ourtthe need for intuitive interpretability of the results, while too much transparency can have also drawbacks. On one hand, knowledge of failures (even explainable) can erode trust in future results. On the other hand it helps users adjust the interpretation of results relatively to the system boundaries. There certainly is a good balance in between.

This brings us to the drives a system can implement to help building trust. As we mention about interpretability, it helps users better construct predictable expectations. Although it could lead to confirmation bias as we will discuss in the next section, small confirmation loops help users better rely on the system while expanding exploration. The same elements that erode trust when missing help assess reliability when they are tracked. This would include stability, repeatability, accuracy, transferability, and explainability (but in the language of the target users). Provided a feedback loop, especially when the systems helps building machine learning models, users can actually see positive impacts and improve interaction with the system over time. When confirming information, users often need provenance, and go to the most detailed view to check the origin of the data and confirm their insight. Because systems sometime fail, it is important to be transparent and display the root cause of system error. All the aforementioned cues help users distinguish the algorithms limits and understand what a system is good and bad at. The ultimate goal would be to directly transfer this knowledge by physically displaying those limits and give an accurate understanding of a system uncertainty and bias.

Of course, the properties of a system required to be trustworthy heavily depend on the different types of tasks its users need to achieve. We can however distinguish between three different groups of users: model builders, domain experts, and everyday users. In the context of Visual Analytics we often focus on model builders, to which explaining errors is essential for improving the models, and on domain experts who need to be aware of errors to assess models before decision making. Being aware of errors would be even more critical for everyday users who may not possess enough domain knowledge to identify system bias and may take wrong decisions in the end.

As we believe Visual Analytics design could include elements to positively influence trust, here would be a list of key questions. Transparency is a critical aspect, but when to be transparent? And how transparent? What would be the trade-off between interpretability and accuracy in name of trust? Since information visualization convey large amounts of information to users, how does it impact trust in general? More precisely, what aspects in visualization would tend to increase or decrease trust? Would data literacy and visual literacy help better improve trust? What aspects of trust are dependent to culture? Is there a way to deceive users and induce trust the same way consumers are influenced by advertisement and deceived by magicians? We will probably need to connect to psychology literature to answer some of those questions.

3.3 Bias

Building trust requires a good communication between user and system. Like any situation of communication, it can fail at different points: what is intended to be communicated, what is actually sent, what is received, and what is understood. In each of those steps, there may be a mismatch – even small – induced by the different process of Visual Analytics or the perception and understanding of user. We may refer to such mismatch as a bias. In this section, we study different sources of bias that would impede communication, and ultimately may erode trust. We investigate where in Visual Analytics design is bias introduced.

Visual Analytics heavily uses machine learning and data analytics, hence sharing the same risk for bias. For example, we previously mentioned a risk for bias in classification, when there is a distance between training data and test set – *e.g.* using features extracted from a classifier trained on ImageNet on a medical images. More specifically, a system may be sensitive to statistical and data bias, from the original sampling or during transformation. We may think of sampling bias during preprocessing data or while interacting with a system, by selecting wrong population or wrong variables – *e.g.* as a result of drastic dimension reduction. Undersampling data may result in a false interpretation of results. Sampling could be sensitive to data missingness and imputation may also include new bias – this may be particularly true while linking different datasets.

Beyond data themselves, modeling may also induce bias. The choice of the right algorithm with its right parameters is critical. One of the most common bias for a model is to be overfitted (or sometimes underfitted), leading to irrelevant predictions. Selection bias in training may lead to a model that cannot properly generalize to its application. Feature selection and engineering may disregard characteristics of the data important to users.

On the other end of Visual Analytics are users. They also face bias while interpreting outcomes. Bias may source from user psychology in a family referred as cognitive bias[1]. Visual Analytics may be particularly sensitive to confirmation bias, in which users pay specific attention to what confirm their own hypotheses (albeit a process somewhat required to build trust). Intention bias slightly varies from the confirmation bias as users are more likely to value insights found on purpose rather than those discovered by chance (this is sensitive for exploratory Visual Analytics often supporting serendipity). Framing effect lead to different interpretation of a same result depending on how it is presented. Anchoring bias can rise in an interactive system from the different paths leading to a same result while inducing different interpretations.

While interacting with the machine, we may face other bias issued from the experimental setting or the workflow. For example, we may face resistant users that would systematically try to reject outputs of a system. In other cases, the system and the data may be used for a task completely unrelated to what it was first designed and gathered for. While it would be nearly impossible to build trust in the first case, users overly trust the system or the data in the other case. The produced insight may not be reliable in the latter case.

[1] There even is a European project RECOBIA dedicated to investigate this phenomenon `https://www.recobia.eu`, as well as a VIS workshop `http://decisive-workshop.dbvis.de/`

The visualization part of a system is also subject to many different types of bias, often use to intentionally trick viewers of visualization in interpreting a visualization in a specific manner[2]. Perception bias may sometimes result in optical illusions but they can also induce misinterpretation, such as grouping visual elements together that do not actually group together in the data. Visual encoding can be particularly subject to bias, *e.g.* using a linear color mapping on a power-law distributed data would not help represent the distribution. Visual representations may also be ambiguous, such as with force-directed graph layouts, non-connected nodes appearing in the same location may be interpreted as close neighbors. Interpretation is of course influenced by user level of visualization literacy, knowledge of graph layout would prevent misunderstanding the previous example.

Towards addressing the question of bias in Visual Analytics, we also identified a few key questions: Can intelligent visualization avoid these pitfalls, while maintaining benefits? What would be the correct visual analysis to avoid bias x, bias y, bias z? While in the "User in the Loop" approach, which types of bias may occur at each step? How may we connect with "ad hoc model building" in Machine Learning community? How do we validate absence of bias in a system? How do we connect bias and task taxonomies?

Although Visual Analytics is exposed to many sources of bias, it has the potential to make things faster, easier, more understandable/comprehensible. By knowing and understanding these sources of bias, we see open challenges for Visual Analytics to tackle and materialize bias for consumers to avoid, while perceptual bias may also be used to the advantage of the system [13]. One should also be careful that the discovery of more ways to improve modeling can introduce risk of "time wasting" for chasing dead ends (such as chasing noise, and ending with overfitting models).

3.4 Conclusion

Building trust is crucial for user involvement, system adoption and efficient human-machine intelligence. Very much like visualization theory itself, it takes elements of cognitive psychology, perception and communication to address this issue. We attempted here to highlight important points for tackling the issue of designing visual analytics systems while optimizing trust. One key problem is to address bias at the different stages it occurs. This document also attempts to shed light on the different shapes bias can take, while underlining the unique opportunity we believe the field of Visual Analytics has to contribute to the solution.

[2] `http://www.vislies.org/`

Overview of Talks

The following invited talks very given by participants of the meeting to provide different perspectives on the topic and to inspire the discussions in the workshop.

Visual Analytics Approaches and Future Research Areas

Tobias Schreck, Graz University of Technology

Advances in data acquisition and storage technology lead to the creation of increasingly large, complex data sets across application domains as diverse as science, engineering, business, social media, or team sports analysis. Important user tasks for leveraging large, complex data sets include finding relevant information, exploring for patterns and insights, and re-using of data for authoring purposes. Novel methods in visual-interactive data analysis allow to tightly integrate knowledge of domain analysts with automatic data analysis methods, offering solutions for complex analysis problems. We discuss visual-interactive data analysis techniques from our work that support search and analysis in a variety of different data types and enabling novel application scenarios. Specifically, we discuss approaches for visual exploration of patterns in team sports data, example- and sketch-based search in multidimensional data sets, and interactive regression modeling. We conclude the talk by discussing future research challenges in the area.

Interactive Model Analysis with Interactive Visualization

Shixia Liu, Tsinghua University

In most AI applications, machine learning models are often treated as a black box. Users usually refine and improve the models according to performance metrics such as accuracy. Because of lacking of a comprehensive understanding of the working mechanism of these models, it is hard to build an effective two-communication between a human and a computer, which limits the further adoption of the models. To solve this problem, we have developed a set of visual analytics approaches to help users understand, diagnose, and refine a machine learning model. This talk presents the major challenges of interactive machine learning and exemplifies the solutions with several visual analytics techniques and examples. In particular, we mainly focus on introducing the following three aspects: 1) create a suite of machine learning techniques that produce more explainable models, while maintaining a high level of learning performance (prediction accuracy); 2) develop a set of visual analytics techniques that enable human users to understand and diagnose machine learning models; 3) a semi-supervised model refinement mechanism. Based on these, we develop an interactive model analysis framework, which is exemplified by deep learning, ensemble learning, and the topic model.

Visualization, Machine Learning and Algorithmic Aversion: Considerations for Human-Machine Intelligence

Ross Maciejewski, Arizona State University

As the amount of data available for analysis has increased, leaps in machine learning and data mining techniques have occurred, enabling large-scale modeling of all sorts of phenomena. Such modeling is often performed offline in a relatively black-box manner where results are presented to be used (or ignored) by the domain experts. Here, the visual analytics community postulates that the integration of domain knowledge into an interactive sense-making loop will improve modeling results from machine learning claiming that experts have some inherent knowledge that cannot be easily encapsulated by the machine learning. Anecdotal evidence from the visualization community has suggested that the direct integration of domain knowledge does improve the overall model efficacy. However, research from the management science community has found mixed results of human-in-the-loop. As such, how much (if any) human should be included as part of machine learning? In this talk I will cover issues of biases, trust, and future challenges for human-machine intelligence in relation to visual analytics.

Wrestling with Temporal Event Sequences

Leo Zhicheng Liu, Adobe Systems Inc.

Making sense of temporal event sequences is challenging because of the scale and complexity of such data. In this talk, I share our experiences in developing visual analytics approaches for analyzing event sequence data. The projects discussed here represent three different perspectives: designing novel visual representations to reduce visual cluttering, applying data mining algorithms to reduce the data and visualize the mined patterns, and devising novel mining techniques that scale to large datasets and produce more interpretable results. I conclude this talk with some reflections on these projects and future directions.

Immersive analytics: Interactive data analysis using the surfaces and spaces around us

Tim Dwyer, Monash Univeristy

Humans struggle to understand the masses of complex data they now accumulate. Visual data analytics offers a solution, and we are exploring the potential for new immersive display and interaction technologies to greatly enhance this potential. Immersive Analytics is a new research field developing the first practical and theoretical frameworks for immersive data analysis. Our work is informed by controlled studies and systematic design exploration; and user-centred design of practical tools for immersive data analytics. Findings that lead to more effective, engaging and collaborative systems for data analytics will ultimately allow people to make more informed decisions from data.

List of Participants

- Prof. Yingcai Wu, Zhejiang University
- Prof. Nan Cao, Tongji University
- Dr. Steffen Koch, University of Stuttgart
- Prof. David Gotz, UNC-Chapel Hill
- Prof. Xiaoru Yuan, Peking University
- Prof. Guodao Sun, Zhejiang University of Technology
- Prof. Jaegul Choo, Korea University
- Dr. Benjamin Renoust, Osaka University
- Prof. Yu-Shuen Wang, National Chiao Tung University
- Dr. Zhicheng "Leo"Liu, Adobe Systems Inc.
- Prof. Ross Maciejewski, Arizona State University
- Prof. Jing Yang, UNC Charlotte
- Prof. Ye Zhao, Kent State University
- Prof. Tim Dwyer, Monash University
- Prof. Tobias Schreck, Graz University of Technology
- Prof. Shixia Liu, Tsinghua University
- Prof. Christopher Collins, University of Ontario Institute of Technology
- Dr. Conglei Shi, Airbnb
- Prof. Seokhee Hong, University of Sydney
- Prof. Gennady Andrienko, "Fraunhofer Institute IAIS City University London"
- Prof. Natalia Andrienko, "Fraunhofer Institute IAIS City University London"

Meeting Schedule

Check-in Day: January 28th (Sun)

- Welcome Banquet

Day1: January 29th (Mon)

- Welcome session and introduction of participants
- Group Photo Shooting
- Talk by Tobias Schreck
- Talk by Shixia Liu
- Talk by Ross Maciejewski
- Talk by Leo Zhicheng Liu
- Talk by Tim Dwyer
- Suggestion and selection of three concrete research topics to focus on
- Building of three discussion groups

Day2: January 30th (Tue)

- Meeting of all participants
- Work in discussion groups
- Synchronization of discussion groups

Day3: January 31th (Wed)

- Meeting of all participants, synchronizing on topics of groups
- Work in discussion groups
- Excursion and Main Banquet

Day4: February 1st (Thu)

- Meeting of all particpants and final discussions
- Wrap up

References

[1] Z. Abedjan, X. Chu, D. Deng, R. C. Fernandez, I. F. Ilyas, M. Ouzzani, P. Papotti, M. Stonebraker, and N. Tang. Detecting data errors: Where are we and what needs to be done? *Proceedings of the VLDB Endowment*, 9(12):993–1004, 2016.

[2] N. Andrienko, T. Lammarsch, G. Andrienko, G. Fuchs, D. Keim, S. Miksch, and A. Rind. Viewing visual analytics as model building. In *Computer Graphics Forum*. Wiley Online Library, 2018.

[3] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, 2013.

[4] E. T. Brown, A. Ottley, H. Zhao, Q. Lin, R. Souvenir, A. Endert, and R. Chang. Finding waldo: Learning about users from their interactions. *IEEE Transactions on visualization and computer graphics*, 20(12):1663–1672, 2014.

[5] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski. Characterizing guidance in visual analytics. *IEEE transactions on visualization and computer graphics*, 23(1):111–120, 2017.

[6] C. Conati, E. Hoque, D. Toker, and B. Steichen. When to adapt: Detecting user's confusion during visualization processing. In *UMAP Workshops*. Citeseer, 2013.

[7] M. Dörk, P. Feng, C. Collins, and S. Carpendale. Critical infovis: exploring the politics of visualization. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, pages 2189–2198. ACM, 2013.

[8] M. Ge, C. Delgado-Battenfeld, and D. Jannach. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 257–260. ACM, 2010.

[9] M. Kaminskas and D. Bridge. Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 7(1):2, 2017.

[10] S. Kandel, J. Heer, C. Plaisant, J. Kennedy, F. van Ham, N. H. Riche, C. Weaver, B. Lee, D. Brodbeck, and P. Buono. Research directions in data wrangling: Visualizations and transformations for usable and credible data. *Information Visualization*, 10(4):271–288, 2011.

[11] S. Kandel, A. Paepcke, J. Hellerstein, and J. Heer. Wrangler: Interactive visual specification of data transformation scripts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3363–3372. ACM, 2011.

[12] S. Kandel, R. Parikh, A. Paepcke, J. M. Hellerstein, and J. Heer. Profiler: Integrated statistical analysis and visualization for data quality assessment. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, pages 547–554. ACM, 2012.

[13] B. E. Rogowitz, L. A. Treinish, and S. Bryson. How not to lie with visualization. *Computers in Physics*, 10(3):268–273, 1996.

[14] D. Toker and C. Conati. Eye tracking to understand user differences in visualization processing with highlighting interventions. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 219–230. Springer, 2014.

[15] J. Van den Broeck, S. A. Cunningham, R. Eeckels, and K. Herbst. Data cleaning: detecting, diagnosing, and editing data abnormalities. *PLoS medicine*, 2(10):e267, 2005.

[16] J. J. Van Wijk. Views on visualization. *IEEE transactions on visualization and computer graphics*, 12(4):421–432, 2006.

[17] P. Verissimo, M. Correia, N. F. Neves, and P. Sousa. Intrusion-resilient middleware design and validation. *Information Assurance, Security and Privacy Services*, 4:615–678, 2009.

[18] E. Wall, L. M. Blaha, L. Franklin, and A. Endert. Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, 2017.