



Workshop report: information visualization—human-centered issues in visual representation, interaction, and evaluation

Andreas Kerren¹
John T. Stasko²
Jean-Daniel Fekete³
Chris North⁴

¹Computer Science Department, School of Mathematics and Systems Engineering (MSI), Växjö University, Växjö, Sweden; ²School of Interactive Computing & GVU Center, Georgia Institute of Technology, Atlanta, GA, U.S.A.; ³INRIA, Unité de Recherche Futurs, Université Paris-Sud, Orsay Cedex, France; ⁴Department of Computer Science and Center for Human–Computer Interaction, Virginia Tech, Blacksburg, VA, U.S.A.

Correspondence:
Andreas Kerren, Computer Science Department, School of Mathematics and Systems Engineering (MSI), Växjö University, Vejdes Plats 7, Växjö SE-351 95, Sweden.
Tel: +46 470 767102;
fax: +46 470 84004.
E-mail: andreas.kerren@vxu.se

Abstract

From 28 May to 1 June 2007, a seminar on 'Information Visualization—Human-Centered Issues in Visual Representation, Interaction, and Evaluation' took place at the International Conference and Research Center for Computer Science, Dagstuhl Castle, Germany. One important aim of this seminar was to bring together researchers and practitioners from Information Visualization and related fields, as well as from application areas, for lively discussion and interaction. The seminar allowed critical reflection on actual research efforts, the state of field, evaluation challenges, and other important topics. This report summarizes the event.

Information Visualization (2007) 6, 189–196. doi:10.1057/palgrave.ivs.9500160

Keywords: Information visualization; visualization; human-centered aspects; evaluation; visual analytics; interaction; exploration; human–computer interaction

Introduction

Schloss Dagstuhl, the International Conference and Research Center for Computer Science, was initiated by the German government to promote informatics research at the international level. Dagstuhl seeks to foster dialogue among the computer science research community, advance academic education and professional development, and transfer knowledge between academia and industry. The Center is located in an 18th-century castle near Wadern in southwest Germany.

The primary feature of Dagstuhl is the week-long seminars on various topics in computer science. Dagstuhl Seminars are frequently described as being the most productive academic events that the participant researchers have ever experienced. The informal and friendly atmosphere fostered at the Center promotes personal interaction between the guests. Traditionally, there is no set program followed at Dagstuhl Seminars. Instead, the pace and the procedure are determined by the presentations offered during the seminar and the discussion results. In order to maintain the constant high standard, topics for the seminars and a list of participants are submitted in the form of a proposal and are reviewed by a scientific directorate. Consequently, participation in a seminar is by personal invitation only. Further general information about Dagstuhl Seminars can be found on the Dagstuhl Castle webpage.¹

One main goal of this seminar on Information Visualization was to bring together researchers and practitioners from the addressed research areas as well as from application areas, such as Bioinformatics, Finance, Geo Sciences, Software Engineering, and Telecommunications. Several international conferences include information visualization topics, each with a slightly different high-level objective. Another goal of the Dagstuhl seminar was to consolidate these diverse areas in one joint meeting.

The seminar allowed critical reflection on actual research efforts, the state of the field, evaluation challenges, and future directions. Participants were also encouraged to perform system demonstrations of prototypes and environments relevant to the seminar topics.

This report is organized as follows: In the next section, we review the composition of attendees and briefly explain how the programme was formed. The subsequent section discusses the seminar sessions and briefly highlights the main ideas that emerged from each. The last section discusses some of the products and next steps of the seminar.

Participation and program

About 40 researchers from 11 countries participated in this seminar. The majority of attendees were from the US and from Germany, but others came from Canada, Australia, and other European countries. Most attendees were affiliated with universities but others came from industrial research labs such as Microsoft and IBM, as well as other research institutes, such as INRIA, Fraunhofer, and NICTA.

The program aimed to generate lively discussions. The presenters were asked not to give talks about their own research work. Instead, the organizers started the seminar with a discussion and the objective to collect interesting and important themes that should be discussed in later sessions. Attendees then could give talks within these different sessions and the talks could, for example, illuminate a specific aspect of the session theme, indicate open problems/difficulties, formulate concrete questions that should be discussed, and so on. From the initial discussion, 11 different session themes emerged, each on a sub-area of information visualization that fit into the overall seminar theme. A source for more detailed information, including a list of all participants, abstracts from each talk, and notes from each session can be found at Seminar 07221.²

Seminar sessions

The program included both general discussion and 11 topical sessions. Each of the topic sessions is reviewed briefly below.

Collaborative information visualization

Presenters: Jeff Heer, Frank van Ham.

The theme of this session was collaboration and how people could work together using information visualization systems. A number of examples of systems were presented as illustrations of how collaboration can play a role in InfoVis.

One example was the Vizster visualization of *Friendster* social networks. When the system was displayed on a large screen in a public area, people tended to spend more time in front of the visualization in groups rather than

as individuals. People told stories, encouraged each other to unearth facts and generally just engaged in social play using the visualization.

In another project from IBM, the system's developers built a way for people to annotate and comment on a visualization created by another person. For instance, a person can make marks on portions of a visualization or circle some key aspect of the visualization to draw attention to it. People can also comment on aspects of a visualization and their text comments are connected to the visualization through links. These annotations and comments also generate new unique URLs that can be shared and returned to.

The ManyEyes project also from IBM allows people to upload their own data sets, and then others can create visualizations of the data using the built-in views of the system. These views include a number of well-known representations such as bar charts, line charts, treemaps, bubblecharts, and stacked histograms. The InfoVis research group at IBM is studying how and why people create visualizations on the system and comment about the data sets stored there.

One key idea that emerged during the session was that groups view these kinds of visualizations more than individuals, thus we need to think about the ramifications of that fact. It certainly brings prior research in CSCW to our focus here. Also, people play different roles as they use collaborative InfoVis tools: some create visualizations, others comment or develop new insights communally, while others may even mark up or edit views.

Another key idea was that visual sensemaking is not only a cognitive exercise but it is also a social activity. Systems that promote collaboration have the opportunity to enhance this aspect.

The visualizations described in this session were massively collaborative and end-user driven, two characteristics not encountered much in InfoVis before. The speakers noted that comments about visualizations (and thus interaction and sensemaking) can occur not only in a visualization system but also on outside blogs, web postings, and so on.

Discussions after the talks pointed out that the type of collaboration discussed in the two talks was asynchronous. Relatively little work has focused on synchronous collaboration in InfoVis systems.

A question was asked whether one of the visualizations could be made partly private so only certain people could examine and comment on it. The speakers commented that this would run counter to the goals of the projects in that everything remain in the public sphere.

Two main themes resonated throughout this session. First, these talks were about examples of InfoVis tools being used to help people communicate, not just do analysis. This idea provides an interesting parallel to how computers have been thought of in general, first as tools for doing work but more now as aids for helping people to communicate with others. The second theme was that the projects discussed are all examples of the

'democratization of visualization,' that is, visualization being available to and intended for everyone.

Theoretical foundations of information visualization

Presenters: Matthew Ward, Robert Kosara, T.J. Jankun-Kelly, Natalia Andrienko.

The aim of this session was to discuss existing and missing theoretical foundations for InfoVis. The general consensus was that InfoVis currently lacks adequate foundational theory, resulting in *ad hoc* approaches. The speakers offered four different potential directions for new theories.

The first idea suggested exploiting existing concepts of Communication Theory to formalize the problem of predicting a visualization's effectiveness for communicating information. Since a data set may contain near infinite potential features it might actually be easier to measure entropy, or the loss of information along the visualization pipeline. For example, different data and visual transformations cause different types and quantities of entropy. Then, if analysts know where entropy has occurred in the pipeline, they can take corrective action. A serious challenge in this approach is the compromise between syntactic and semantic measures, the difference between data content and data implications.

The second idea was to establish a culture of Visualization Criticism, analogous to the practice done in art and literature. For this method to be effective, we must agree to a set of rules for criticism, develop a language for talking about visualizations, and promote its practice in publications and at conferences. For example, the paper review process could be organized as critiques, with credit given to the critiquers. EagerEyes.com provides several examples. There are significant academic social issues to overcome to make this approach successful.

The third idea focused on the problems of teaching InfoVis, and suggested the need for a formalism to describe different visualizations. In contrast to courses in physics or chemistry, InfoVis courses often consist of a series of visualization techniques without any underlying theory to guide students. Three potential models were proposed: an explorative model could formalize users' exploration path, a transform model formalizes the pipeline, and a design model emphasizes the consequences and benefits of design choices. Also needed are theory-based textbooks and greater discussion about InfoVis education. One problem is that humans play an essential role in InfoVis. Thus, psychology and related fields must be involved in such a theory.

The fourth and most provocative idea posited that we generate a taxonomy of all possible patterns in data, and identify visualization methods that will find each pattern. Such a theory could offer predictive power for what patterns can be discovered from a data set. It could also be used to create a visualization textbook intended for data analysts, in contrast to current textbooks intended for visualization researchers. However, there was significant

skepticism in whether such a taxonomy can be created, since visualization allows analysts to find previously unknown patterns that cannot be automatically mined.

The overall discussion centered on difficulties in generating a theory of InfoVis. For example, InfoVis is interdisciplinary, and theories must include user-centered aspects. Some participants pointed out that InfoVis is still in the early observational stages of a developing science, and perhaps not ready for predictive theory yet. Efforts to perform experiments should continue. But, we should also avoid becoming too comfortable with empiricism, forever building and testing tools, and should proceed to theory development.

The value of information visualization

Presenters: Chris North, John Stasko, Jean-Daniel Fekete.

In this session, the speakers engaged the interesting problem of how to better communicate the value of InfoVis and how InfoVis can be better disseminated. It was generally agreed upon that we, as InfoVis researchers, often have a difficult time explaining and showing how information visualization can benefit people. One reason for this challenge is that InfoVis systems often are designed to assist with the acquisition of *insight*, but, how does one define and quantify insight. Many insights are informal, that is, they involve more abstract notions and are challenging to precisely articulate.

Furthermore, InfoVis appears to be most useful in exploratory analytic scenarios, involving browsing and undirected information seeking, as opposed to other more concrete tasks like search. It is much easier to quantify and compare tools for search because one can more easily identify whether a search was successful or not and how long the search took. In exploratory information seeking, there is no well-defined goal state that can be examined and measured.

InfoVis systems can be most useful in scenarios involving browsing and exploration. Such scenarios typically occur when the person involved has a limited understanding of the domain being examined and cannot clearly articulate what is needed or sought. In fact, one view of InfoVis believes that it is not so useful for problem solving, that is, answering specific questions. Instead, it is useful for asking better questions or simply helping a person identify and articulate a question or problem.

InfoVis may be useful for both acquiring new insights and for simply gaining an understanding of a problem or situation more rapidly than would have been possible without a visual representation of the data. InfoVis systems appear to have much value in simply speeding up the time to gain situational awareness and knowledge about a problem at hand.

Speakers in the session also showed a number of example pictures and visualizations that researchers use to help make the case about the value of InfoVis. Frequently, these examples appear early in introductory talks about InfoVis or in an early lecture in an InfoVis course. The

examples also typically highlight findings that would not be discovered by automatic data mining.

Evaluation of information visualizations

Presenters: Sheelagh Carpendale, Keith Andrews, Helen Purchase.

The importance of evaluating InfoVis techniques and systems, and the challenges in doing so, were the main topics of this session. The speakers identified many different difficulties in performing InfoVis evaluations with one of the key reasons being that InfoVis systems are usually large, complex systems. For instance, one may want to evaluate an InfoVis technique on its own merits, but it is difficult to separate the evaluation of a technique from the specific system implementation of that technique.

Many different evaluation methodologies can and have been used for InfoVis systems, ranging from inspection-based techniques such as heuristic evaluation and cognitive walkthrough to more formal comparative experiments to more situated, observational studies such as ethnographies and case studies. Evaluations can be more formative to gain design insight and guidance or more summative to compare competing approaches and measure value. No matter what evaluation technique is chosen, however, a multitude of questions and issues emerge. What are the experimental hypotheses? Which implementation should be tested? Who should participate in the evaluation? Which data set should be used and what tasks are to be performed? What should be measured?

Furthermore, InfoVis systems are typically not used in a void. They are a component in a sense-making ecology. To adequately evaluate an InfoVis system, one must factor in and understand this larger environment.

We discussed empirical studies in graph drawing in more detail as an application example of evaluation in a related field. Work in this area could be a guide for evaluations in InfoVis. However, many concrete metrics in graph drawing (e.g. uniformity of edge length and minimizing edge crossings) make it easier to compare systems in that area. In InfoVis, one must also factor in and evaluate interaction provided by the system, which is often more difficult to do well.

Participants suggested that steps could be taken to help foster better InfoVis evaluations. By better documenting data sets and tasks used in studies, as well as providing actual source code and system settings used in the experiments, we promote and foster better replicability. Efforts like the IEEE InfoVis Conference Contest help by creating standardized data sets for use across systems. The BELIV (Beyond time and errors: novel evaluation methods for Information Visualization) Workshop held initially in conjunction with AVI '06 and CHI '08, is a meeting of researchers specifically focusing on evaluation issues in InfoVis.

Finally, a lively discussion ensued about whether *all* good InfoVis papers need to include evaluation.

Participants commented that an unwritten expectation of evaluation now appears to hold for conference and journal publications. There was general agreement that this view is too strong and should not be held. If a new technique is being introduced, it may be difficult to adequately evaluate it. Further, innovative and creative new ideas should be able to stand on their own.

Largely, it is a matter of how a research contribution is presented. If a paper makes a claim that it improves some process or problem, then the authors should show that this is done, typically through some form of evaluation. But if such claims are not made, then evaluation should not always be required. In fact, one might argue that it is an insult to evaluation to feel that it must be included in all papers. Evaluation is hard and very time-consuming. Throw-away evaluations that show little and are included just for the sake of doing so can 'pollute' a paper. It was also pointed out that for some papers a good analytic evaluation would be better than an empirical evaluation.

Interaction in information visualizations

Presenters: Helwig Hauser, Jonathan Roberts.

Interaction and exploration are two of the most important aspects of information visualization; some even argued that insight is formed through interaction. The session addressed this question with two angles: a compelling presentation showing a complex investigation that required many interactions to understand a large multi-dimensional dataset, and an entertaining poll asking participants to list what they think are the most important issues in InfoVis interaction and to rank them.

We began with a list of seven issues that were then refined into 15 categories:

- analytic query,
- management of the environment,
- coordination,
- meta-level,
- exploration state,
- interaction devices,
- data structures,
- algorithms,
- hardware,
- data,
- filters,
- flow,
- forms,
- transforms,
- navigation,
- selection.

The audience broke into small groups that discussed and listed what they thought were open problems. The category receiving the largest number of items was 'Exploration State' (history of interaction). However, the post-exercise discussion suggested that the initial



Figure 1 Results of participants' polls on important interaction issues.

taxonomy had many more interdependencies than were first apparent and came up with the following list of issues (Figure 1):

- uncertainty in data, tasks, representations, aggregations;
- recommendation systems;
- time-changing data;
- interactive sensitive dependence;
- gulf of execution (bridging the gap between goals and actions).

Visual analytics

Presenters: John Stasko, Carsten Görg, Stephan Diehl, Daniel Keim.

The aim of this session was to clarify what is meant by the term 'Visual Analytics' and to discuss how this area differs from InfoVis. A concise definition of Visual Analytics characterizes it as the science of reasoning facilitated by interactive visual interfaces (as in *Illuminating the Path*), but this definition likely does not explain its similarities and differences to InfoVis very well.

A more thorough description of Visual Analytics begins with the observation that InfoVis techniques and tools tend to not work as well when the data sets being examined become very large, for instance, billions or trillions of records. Conversely, automatic knowledge discovery approaches such as data mining can scale up more flexibly, but they only work well for well-defined and specific problems. Visual Analytics is then the combination of these two approaches, the tight integration of visual, human-directed analysis methods with automatic, database-supported methods. It is not simply applying one approach followed by the other, however. It must involve fundamental, deep integration of the two techniques. For instance, human-directed visual interfaces could help steer automatic mining algorithms, or

automated data filtering and selection tools could funnel data to visual interfaces.

Visual Analytics research in the United States grew primarily from homeland security applications and interests, but it is not fundamentally tied to that domain. Visual analytics research could apply to other domains such as personal information management, socio-demographics, healthcare, engineering, finance and business, just to name a few.

Visual Analytics also may have evolved, at least in part, from the view that too much InfoVis research focused on creating evocative visuals with relatively little regard to actual user needs and tasks. Visual Analytics appears to have a more application-driven focus.

Because the IEEE InfoVis Conference and the VAST Symposium are co-located and concern topics with much in common, potential confusion on where to submit research papers does exist. By clarifying the definition of this new area and illustrating how it differs from InfoVis, we can hopefully assist researchers in finding the best venue to present their work.

Examples of two systems and application areas in visual analytics were presented to further illustrate what this new area is about. The first involved investigative analysis of underlying plans and plots embedded across large document collections. The second illustrated visual mining in software archives to assist software developers with activities like debugging and refactoring.

Writing good InfoVis papers

Presenter: Tamara Munzner.

This session had a very pragmatic aim, to determine what makes a good InfoVis paper. Five different categories of papers were discussed: system, model, evaluation, technique, and design study papers. These five categories are described on the IEEE InfoVis Conference call-for-papers webpage. The audience was divided into groups, each discussing one of the five paper categories. At the end of the session, the groups presented their findings.

Of particular note was the discussion of Evaluation papers. Participants argued that the evaluation paper category should be explicitly broadened to encourage submissions on more types of evaluations, including requirements analyses and ethnographic studies as well as comparative studies. However, formative usability studies might better fit in the Design Study category. Distinctions could also be made between papers that attempt to prove the value of a visualization and those that attempt to inform future design. An important issue is that some reviewers might not have expertise in all evaluation methods, and so the InfoVis Conference reviewer database should have more subcategories for evaluation expertise. Another open question is how to handle papers that report null results.

It was agreed that evaluation papers should contain clear descriptions of the aim of the evaluation or

the research question, why the selected evaluation method was chosen to answer the question, and a complete description of the method, results, and lessons learned. While papers of other categories may also include evaluations, it was generally agreed that Evaluation papers should contain mostly evaluation related content.

For the Design Study paper category, a suggested prototypical paper should report on a system application where the developed visualizations are not innovative enough for a technical paper, and should offer a convincing argument that the needs of the application are really solved by the system. Another view is that Design Study papers should be about the design process undergone, not just the product. They should also enable generalization that motivates future designs or technical papers. An open question is whether strictly conceptual or mock-up designs would be acceptable.

In the discussion, one suggestion was to identify one or more excellent existing papers for each paper category, and to include those as exemplars on the InfoVis Conference webpage. Also, reviewers and readers should be made aware of the category of each paper. Thus, authors should state the category in their paper text, and the online submission system should require authors to tag their submission accordingly. It might also be helpful to provide a link to guidelines for good InfoVis papers from the InfoVis Conference webpage to mitigate the problem of diverse expectations.

Videos are very important for illustrating the benefits of highly interactive techniques and should be more encouraged in paper submissions. IEEE now requires that published supplemental videos for a conference paper must also be peer reviewed, and so authors cannot add videos to accepted papers after the review process is complete. One possible way to encourage video submission is to offer a 1 week later deadline for submitting a supplemental video for a paper. This would give authors time to make a video after finishing their paper, and would still be in time for reviewers.

The speaker presented a lively discussion of paper pitfalls including: too low-level with no big picture, least publishable units, too much content without any details, simultaneous submissions, no explicit statement of contributions, ignored previous work on similar problems or solutions, lack of initial motivation and overview, and too much jargon. Also presented was a list of questions to consider before beginning a new InfoVis related project: Are there accessible real users? Is there a real need for an InfoVis solution? Will the need persist? Is there a real analysis task? Is real data available? Is the available data the truly needed data, or merely the easy-to-collect data?

Integration of visualization domains

Presenters: Kwan-Liu Ma, Michael Schlemmer, Jason Dykes, Dirk Zeckzer.

The focus of this session was the integration of InfoVis with other closely related fields, such as Scientific Visualization (SciVis), Geographic Visualization (GeoVis), Bioinformatics Visualization (BioVis), etc. Several visualization examples were shown that demonstrated combinations of these fields.

Some of the speakers talked about the need for InfoVis in SciVis. They showed three levels of integration:

1. InfoVis next to SciVis to visualize *data attributes too numerous or abstract to show* using SciVis;
2. InfoVis coordinated with SciVis to visualize and manipulate *complex parameters used to control* SciVis;
3. InfoVis embedded in SciVis where both were merged and it was difficult to decide what part was InfoVis and what part was SciVis.

The boundaries were not clear, but the two first levels are easy to achieve and not unique to SciVis (InfoVis can be used to display complex abstract attributes associated with any database in Biology, Medical Imaging, etc.) The last level of integration was difficult to discuss in general but some interesting examples were shown.

There seemed to be an issue with integrating the respective rendering pipelines. SciVis uses 3D geometrical projections in which visual features essentially follow the projection, whereas InfoVis uses 2D geometrical projections where sizes should not be distorted to remain comparable and preattentive. This was considered a problem to merge the two in current systems.

The case of GeoVis seemed much closer to InfoVis. Some common properties of GeoVis that are perhaps familiar to InfoVis are that errors in data are common, scale and form are very important, and interpretations are typically subjective, imperfect, and incomplete. In GeoVis, maps help by providing a common structure for visual synthesis and an artifact on which to project tacit knowledge. GeoVis grew out of the previously established informal ethnographic methods of geography research, and hence embraces informal approaches and its tools do not need to be experimentally proven to be accepted in the geography research community. In fact, in geography, ground truth is a contentious issue. It was pointed out that the field of geography can greatly benefit from improved visual literacy by introducing visualization into general education. In terms of integration, it was pointed out that Cartography has good theories that InfoVis could better exploit.

A lively discussion produced disagreement about what the differences are, if any, between the fields. However, it was generally agreed that integration is natural, and that InfoVis is clearly needed and utilized in many other domains. But when InfoVis is applied to another domain it is generally then given the name of that domain (e.g. InfoVis in SciVis is still SciVis). More important is the question of how to choose the best representation. Methods from any of the Visualization fields can be chosen, but

the question is which methods work best for any given problem?

Human-centered information visualization and broader use

Presenters: Jarke van Wijk, Chris Weaver, Achim Ebert, Gennady Andrienko

The aim of this session was to discuss human-centered aspects in InfoVis and the question of how to promote the use of InfoVis to a much broader audience. Several successes and failures were considered. To get started, approaches to visualization research were grouped into four categories: User-centered, toolsmith, computer science, and curiosity-driven.

The novel 'curiosity-driven approach' suggests that the user is ourselves, the visualization researchers. In this case, the visualization researcher creates a visualization to satisfy his or her own curiosity about a particular problem. As an example, one of the speakers was curious about myriahedral projections of the globe onto a 2D plane, and created a series of fascinating animations of various projections as they 'unwrapped' the earth. This discussion of curiosity-driven visualization led to an idea for broader use in which the general public is trained to do this kind of work for themselves, by integrating visual literacy into primary education. Indeed, we are becoming a visual society, so citizens should be able to speak the visual language.

The question 'what is between expert and casual visualization?' generated much discussion and debate. Perhaps there is a region between experts and casual users that consists of curiosity-driven hobbyists or semi-serious amateurs. Examples domains include teaching, science, journalism, social interaction, and personal expression. Tools developed for this category of visualization might offer users the ability to customize visualization designs or to create mash-ups with tools such as Google Earth. As a live example, the speaker created a customized visualization of his personal MP3 music collection.

For an example of the Toolsmith approach to visualization, a novel 3D document management system was presented. This experience highlighted the difficulty of designing general purpose tools for a broad class of users, especially when attempting to introduce new technologies.

For an example of the User-Centered approach, one of the speakers presented a geovisualization solution that was designed to support a specific group of scientists who study deforestation. However, the targeted users were apparently never satisfied. A hypothesis was that the visual analysis was inherently complex, and the users preferred simple specific statistical answers over the more informal and complex visual insight. Some potential solutions were suggested: ostensible simplicity (offer some powerful operations, but through simple interactions), user guidance, and incremental intelligence. It may be necessary to introduce visualization capabilities

gradually to allow the user groups to learn new methods of work.

Teaching information visualization

Presenter: Keith Andrews

In advance of this session, 18 participants completed a short questionnaire about their InfoVis teaching activities: course length, attendees, materials used, examinations, assignments, practical exercises, etc.

Results were discussed, for example:

1. Most InfoVis courses were taught at the graduate level,
2. Many did not use a fixed textbook, but the most popular books were Colin Ware's and Robert Spence's,
3. Students typically had to prepare a short presentation about a research paper,
4. All InfoVis courses used practical exercises,
5. About half of the courses had exams,
6. One of the most common practical exercises is to have students use InfoVis tools and then critique them.

A common problem reported by InfoVis teachers is how to logically organize the topics in the course. For example, Robert Spence's 2nd edition textbook on Information Visualization follows the pipeline model of representation, presentation, and interaction. One suggestion is to consider four cross-cutting dimensions: data types, domains, techniques, and methodologies. Another important dimension is cognitive and perceptual issues.

Participants also noted that a large public collection of images and videos would be very helpful for instructors. However, this could cause copyright problems.

The state of the field

Presenter: Martin Theus

In this session, participants discussed the current state of the field and solved problems. A deliberately provocative claim was made that many problems in the area of statistical data visualization (DataVis) were already solved. A demonstration of the Mondrian system (rosuda.org/Mondrian/) was used to show that data tables, even with missing values, could be analyzed using several classical statistical visual representations and explored through selection, filtering, and linked views.

It was pointed-out, though, that InfoVis is much broader than DataVis. DataVis focuses on well-structured data tables, with limited data types and scales, that map to graphics in a relatively straightforward manner using established views. There are thus fewer degrees of freedom available to designers and little left to invent.

However, InfoVis can learn from and expand on successes in DataVis. DataVis has a cookbook of established visual methods for different types of analyses. DataVis exploits formal statistical techniques to strengthen data graphics (e.g. confidence intervals in boxplots), enabling a more formal approach to

graphical analysis and testing. DataVis uses linked views as an interactive graphical form of testing conditional distributions. DataVis exploits data summarization and analytical building blocks.

It was claimed that since statistical visualization and analysis are well understood and clear, the users of the Mondrian system are satisfied with it. Whereas, InfoVis systems were said to often suffer from serious usability issues and difficulties for users. This is due to the greater complexity of data structures, and the critical role of interaction in InfoVis. Hence, usability is a top unsolved priority, and we need to look to our customers to ensure that we are solving their problems. We need to look inside and beyond the current community. Fair, constructive criticism is necessary and important.

Outcome and final comments

The organizers and participants decided to publish a book that should document and *extend* the findings and discussions of this Dagstuhl Seminar. Beforehand, the organizers gained the agreement of Springer Press to publish an LNCS State-of-the-Art issue on the seminar theme. The book will cover the problems discussed in the various sessions in detail.

Understanding of human-centered issues in the area of Information Visualization is of growing importance. This seminar was one of the first events that allowed participants to discuss such topics in an informal environment

together with researchers from different related fields. All participants discussed together without reservation in a friendly and cooperative way. Of course, we hope that this event was the basis for common future research cooperations by some means. A final evaluation performed by the Dagstuhl Center at the end of the seminar showed that the most participants were very pleased with the scientific content and quality of the seminar.

Acknowledgements

We thank all participants at the seminar for the lively discussions as well as the scientific directorate of Dagstuhl Castle for giving us the opportunity to organize this event. The abstracts and talks can be found on the Dagstuhl website for this seminar.² In addition, many attendees agreed to take notes during the seminar sessions. These notes were the basis for writing this report and are also available for download on the seminar website. Last but not least, the seminar would not have been possible without the great help of the staff of Dagstuhl Castle. We would like to acknowledge all of them for their assistance.

References

- 1 Dagstuhl Castle—International Conference and Research Center for Computer Science. <http://www.dagstuhl.de/en/about-dagstuhl/>.
- 2 Seminar 07221 ‘Information Visualization—Human-Centered Issues in Visual Representation, Interaction, and Evaluation’. <http://www.dagstuhl.de/07221>.