

Scalable Set Visualizations

Edited by

Yifan Hu¹, Luana Micallef², Martin Nöllenburg³, and
Peter Rodgers⁴

1 Yahoo! Research – New York, US, yifanhu@yahoo-inc.com

2 Aalto University, FI, luana.micallef@aalto.fi

3 TU Wien, AT, noellenburg@ac.tuwien.ac.at

4 University of Kent – Canterbury, GB, p.j.rodgers@kent.ac.uk

Abstract

This report documents the program and outcomes of Dagstuhl Seminar 17332 “Scalable Set Visualizations”, which took place August 14–18, 2017. The interdisciplinary seminar brought together 26 researchers from different areas in computer science and beyond such as information visualization, human-computer interaction, graph drawing, algorithms, machine learning, geography, and life sciences. During the seminar we had five invited overview talks on different aspects of set visualizations as well as a few ad-hoc presentations of ongoing work. The abstracts of these talks are contained in this report. Furthermore, we formed five working groups, each of them discussing intensively about a selected open research problem that was proposed by the seminar participants in an open problem session. The second part of this report contains summaries of the groups’ findings.

Seminar August 13–18, 2017 – <http://www.dagstuhl.de/17332>

1998 ACM Subject Classification Human-centered computing → Visualization, Theory of computation → Design and analysis of algorithms

Keywords and phrases algorithms, information visualization, scalability, set visualization, visual analytics

Digital Object Identifier 10.4230/DagRep.7.8.1

Edited in cooperation with Tamara Mchedlidze


1 Executive Summary

Yifan Hu

Luana Micallef

Martin Nöllenburg

Peter Rodgers

License  Creative Commons BY 3.0 Unported license
© Yifan Hu, Luana Micallef, Martin Nöllenburg, and Peter Rodgers

Sets are a fundamental way of organizing information. Visualizing set-based data is crucial in gaining understanding of it as the human perceptual system is an analytic system of enormous power. The number of different set visualization methods has increased rapidly in recent years, and they vary widely in the visual metaphors used and set-related tasks that they support. Large volume set-based data can now be found in diverse application areas such as social networks, biosciences and security analysis. At present, set visualization methods lack the facility to provide analysts with the visual tools needed to successfully interpret large-scale set data as the scalability of existing metaphors and methods is limited.



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 3.0 Unported license

Scalable Set Visualizations, *Dagstuhl Reports*, Vol. 7, Issue 8, pp. 1–22

Editors: Yifan Hu, Luana Micallef, Martin Nöllenburg, and Peter Rodgers



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

This seminar provided a forum for set visualization researchers and application users to discuss how this challenge could be addressed.

Existing set visualizations can be grouped into several families of techniques, including traditional Euler and Venn diagrams, but also node-link diagrams, map- and overlay-based representations, or matrix-based visualizations. Inevitably, the approaches taken to drawing these visualizations are diverse, for example node-link diagrams require graph drawing methods, whereas overlay techniques use algorithms from computational geometry. However, they are similar in a number of aspects. One aspect is the underlying set theory. For instance, theoretical results into the drawability of many of these set visualization techniques for different data characteristics is possible (as already done in example Venn and Euler diagram research). Another common aspect is that the visualizations are typically focused on an end-user, so perceptual, cognitive and evaluation considerations are an important concern.

A particularly pressing issue in set visualization is that of scaling representations. The number of data items can be large and many methods aggregate individual items. Yet, even using aggregation, the limit of the most scalable of these methods is considered to be in the region of 100 sets [1]. Typical application areas that make use of sets include, e.g., social networks, biosciences and security analysis. In these applications, there may be many millions of data items in thousands of sets. Other applications have high-dimensional data, where each item is associated with a large number of variables, which poses different scalability challenges for set visualizations.

A distinct feature of set visualization is that visualizations must support set-related, element-related, and attribute-related analysis tasks [1] that involve, e.g., visually evaluating containment relations, cardinalities, unions, intersections, or set differences. For example, bioscience microarray experiments classify large numbers of genes and multiple visualization tools have been developed to visualize this data. However, current efforts can only visualize small sections of the information at once [2]. Similar scalability challenges for set visualizations appear in many other applications as well. Hence, developing effective visualization methods for large set-based data would greatly facilitate analysis of such data in a number of important application areas.

Seminar Goals

The goal of this seminar was to bring together researchers with different backgrounds but a shared interest in set visualization. It involved computer scientists with expertise, e.g., in visualization, algorithms, and human-computer interaction, but also users of set visualizations from domains outside computer science. Despite the large number of set visualization techniques, for which there is often a considerable practical and theoretical understanding of their capabilities, there has only been limited success in scaling these methods. Thus the intended focus of this seminar was to discuss and study specific research challenges for scalable set visualizations concerning fundamental theory, algorithms, evaluation, applications, and users. We started with a few overview talks on the state of the art in set visualization, but then focused on small hands-on working groups during most of the seminar week. We aimed to accelerate the efforts to improve scalability of set visualizations by addressing open questions proposed by the seminar attendees, in order to produce concrete research outcomes, including new set visualization software and peer-reviewed research publications.

Seminar Program

1. On the first two days of the seminar we enjoyed five invited overview lectures on different aspects of set visualizations. The topics and speakers were chosen as to create a joint understanding of the state of the art of set visualization techniques, evaluations and applications. Silvia Miksch gave a systematic overview of set visualization techniques, grouped by types of visual representations and tasks, with a special focus on set visual analytics. Martin Krzywinski reported about his experiences on using visual analogies for showing set-based data in the area of genomics. Sara Fabrikant took a cartographer's view on visualizing sets and explained how successful cartographic maps work as information displays by taking not only the design but also the context and the user into account. Stephen Kobourov explained how large graph-based set data can be represented using a familiar map metaphor by showing several interesting data sets and their map representations. Finally, John Howse presented how set visualizations can be used as diagrammatic reasoning systems in logic.
2. In the open problem session on the first day of the seminar we collected a list of 13 open research problems that were contributed by the seminar participants. In a preference voting we determined the five topics that raised the most interest among the participants and formed small working groups around them. During the following days the groups worked by themselves, except for a few plenary reporting sessions, formalizing and solving their respective theoretical and practical challenges. Below is a list of the working group topics; more detailed group reports are found in Section 4.
 - a. **Mapifying the genome:** Can the axis of the entire genome be mapped on a 2-dimensional space based on gene function rather than a 1-dimensional line based on gene position?
 - b. **Area-proportional Euler diagrams with ellipses:** Can the use of ellipses extend the size of data that can be drawn with area-proportional Euler diagrams?
 - c. **Spatially informative set visualizations:** Can we improve spatial overlay-based set visualizations when allowing some limited displacement of the given set positions?
 - d. **Set visualization using the metro map metaphor:** How and under which conditions can the metro map metaphor be used to visualize set systems?
 - e. **Visual analytics of sets/set-typed data and time: challenges and opportunities:** What are the main research challenges and opportunities in the context of set visualizations that change over time and how can these be structured?
3. We had a flexible working schedule with a short plenary session every morning to accommodate group reports and impromptu presentations by participants. In two of those Wouter Meulemans and Nan Cao shared recent results of theirs related to set visualization.
4. During the week we encouraged participants to come up with suggestions for further strengthening this growing community of set visualization researchers. In a plenary session on Friday we collected and structured these ideas and made started planning future events related to set visualizations, see Section 1.

Future Plans

During the entire seminar, participants actively discussed ways how to disseminate, proliferate and promote scalable set visualization research in diverse specific areas, such as: set theory and diagrammatic reasoning; algorithms and graph theory; information visualization and

visual analytics; evaluation, users and application areas. This led to the concretization of the following future milestones, each of which is being coordinated by volunteered seminar participants:

- **Diagrams Workshop in 2018** on Set Visualization and Reasoning (SetVR)
The workshop aims at promoting set visualization to the Diagrams community, of which well-renowned mathematicians and logicians are members, thus proliferating relevant set theory and diagrammatic reasoning research;
- **IEEE VIS Workshop in 2019** on Set Visualization and Analytics (SetVA)
The workshop aims at promoting set visualization to the Information Visualization and Visual Analytics communities, at the premier forum for advances in information and scientific visualization, with the aim to generate new visualization and analytic techniques to handle large set-typed data;
- **Dagstuhl seminar in 2019** on Set Visualization and Analytics (SetVA) over Time and Space
This seminar has revealed, for the first time, the need for visualization and analytic techniques for the set-typed data that has an element of time and/or space; thus a follow-up Dagstuhl seminar will be organized to discuss this topic, once again among researchers with diverse set visualization backgrounds;
- **Set Visualization Workshop in 2020** in the Computational Geometry Week or collocated with Graph Drawing
This workshop aims at disseminating set visualization to a more algorithmic and computational geometry research community, to ensure the production of effective, yet efficient and scalable set visualization algorithms;
- **Set Visualization browser**, like <http://setviz.net>
This browser will collect and disseminate available set visualization techniques, making them easily accessible through various categorizations, such as the type of data analysis tasks or application areas they target;
- **Set Visualization book**
The book would serve as a guide for researchers who are new to set visualization and as a review of the current state-of-the-art of set-typed data in the various related domains.

We decided to have an annual set visualization workshop that each year focuses on one of (i) diagrammatic reasoning and logic, (ii) information visualization and visual analytics, and (iii) computational geometry and graph drawing, at premier venues of the respective research communities, to generate further research interest in all of these three diverse areas that are all important for scalable set visualizations.

Evaluation

According to the Dagstuhl survey conducted after the seminar, as well as informal feedback to the organizers, the seminar was highly appreciated. Particularly the small group size, group composition, and the seminar structure focusing on hands-on working groups was very well received. The seminar's goals to identify and initiate collaboration on new research challenges was very successful (also in comparison to other Dagstuhl seminars) as the participants rated the seminar highly for inspiring new research directions, joint projects and joint publications. We are looking forward to seeing the first scientific outcomes of the seminar in the near future and to continuing the efforts to support the growth of the set visualization community.

Acknowledgments

Schloss Dagstuhl was the perfect place for hosting a seminar like this. The unique scientific atmosphere and the historic building provided not only all the room we needed for our program and the working groups, but also plenty of opportunities for continued discussions and socializing outside the official program. On behalf of all participants the organizers want to express their deep gratitude to the entire Dagstuhl staff for their outstanding support and service accompanying this seminar. We further thank Tamara Mchedlidze for helping us collecting the contributions and preparing this report.

References

- 1 Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter Rodgers. The State-of-the-Art of Set Visualization. *Computer Graphics Forum*, 35(1):234–260, 2015.
- 2 Sebastian Behrens and Hans A Kestler. Using VennMaster to evaluate and analyse shRNA data. *Ulmer Informatik-Berichte*, page 8, 2013.

2 Contents

Executive Summary

Yifan Hu, Luana Micalef, Martin Nöllenburg, and Peter Rodgers 1

Overview of Talks

SetVA: Visual Analytics of Sets and Set-Typed Data: Challenges and Opportunities
Silvia Miksch 7

Visual analogies and explanations
Martin Krzywinski 7

The Visualization of Sets: A Cartographer’s View
Sara Irina Fabrikant 8

Set Visualization with Maps
Stephen G. Kobourov 8

Diagrams for Logic and Reasoning
John Howse 9

The Painters Problem
Wouter Meulemans 9

UnTangle Map: Visual Analysis of Probabilistic Multi-Label Data
Nan Cao 10

Working Groups

Mapifying the Genome
Radu Jianu, Martin Krzywinski, Luana Micalef, and Hsiang-Yun Wu 11

Euler Diagrams drawn with Ellipses Area Proportionally (EDEAP)
Fadi Dib, Peter Rodgers, Michael Wybrow 14

Spatially Informative Set Visualization
Thom Castermans, Mereke van Garderen, Wouter Meulemans, Martin Nöllenburg, and Xiaoru Yuan 15

Set Visualization Using the Metro Map Metaphor
Robert Baker, Nan Cao, Yifan Hu, Michael Kaufmann, Stephen Kobourov, Tamara Mchedlidze, Sergey Pupyrev, Alexander Wolff 18

Visual Analytics of Sets/Set-Typed Data and Time: Challenges and Opportunities
Daniel Archambault, Kerstin Bunte, Sara Irina Fabrikant, John Howse, Andreas Kerren, and Silvia Miksch 20

Participants 22

3 Overview of Talks

3.1 SetVA: Visual Analytics of Sets and Set-Typed Data: Challenges and Opportunities

Silvia Miksch (TU Wien, AT)

License © Creative Commons BY 3.0 Unported license
© Silvia Miksch

Joint work of Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, Peter Rodgers

Main reference Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, Peter J. Rodgers: “The State-of-the-Art of Set Visualization”, *Comput. Graph. Forum*, Vol. 35(1), pp. 234–260, 2016.

URL <http://dx.doi.org/10.1111/cgf.12722>

Sets comprise a generic data model that has been used in a variety of data analysis approaches. Such approaches involve analyzing and visualizing set relations between multiple sets defined over the same collection of elements. However, visualization / visual analytics of sets is a non-trivial problem due to the large number of possible relations between them. We provide a systematic overview of state-of-the-art techniques. We classify these techniques into six main categories according to the visual representations they use and the tasks they support. We compare the categories to provide guidance for choosing an appropriate technique for a given problem. Finally, we identify challenges and opportunities in this area and propose possible research directions. The most important challenge – in my point of view – is Visual Analytics of set systems over time, called “SetVA over Time”. Further resources on set visualization are available at <http://www.setviz.net>.

3.2 Visual analogies and explanations

Martin Krzywinski (BC Cancer Research Center, Vancouver, CA)

License © Creative Commons BY 3.0 Unported license
© Martin Krzywinski

“The great tragedy of science—the slaying of a beautiful hypothesis by an ugly fact.” wrote Huxley, in a statement that is as much about how science works as about the irrepressible optimism required to practice it. But even greater is the tragedy of obfuscating facts with impenetrable figures and demoting their natural beauty by florid visuals. The issue isn’t one of pure aesthetics—lack of clarity, precision and conciseness in science communication slows our efforts to move forward. In the field of disease research, this has fateful impact on lives.

Through a series of examples from the field of genomics, I show examples of how visual analogies can fail and succeed. I point out that we must not only get the foundations right – choice of color, shape and encoding – but also carefully attend to organizing flow and continuity. The latter become particularly important for complex multi-panel figures.

Many of the kinds of data generated in genomics can be considered as sets. One example for immediate concern is the overlap between specific genomic mutations and cancer types. I argue that for large data sets, such as the recently reported inventory of mutations in 10,000 sequenced tumor genomes [1], there is value in using simple and familiar forms and resist the urge to create new and complex encodings when a series of standard ones will suffice.

References

- 1 A. Zehir et al. *Mutational landscape of metastatic cancer revealed from prospective clinical sequencing of 10,000 patients*. *Nature Medicine* 23, 703–713 (2017).

3.3 The Visualization of Sets: A Cartographer’s View

Sara Irina Fabrikant (*Universität Zürich, CH*)

License © Creative Commons BY 3.0 Unported license
© Sara Irina Fabrikant

Joint work of Members of the UZH Geographic Information Visualization Group (GIVA) and students. See URL below.

Main reference Fabrikant, S.I. (2016). Evidence-based Design. Invited talk at The Cartographic Summit. The Future of Mapping. Hosted by the International Cartographic Association (ICA) and Esri Inc., Redlands (CA), Feb. 9, 2016.

<http://videoembed.esri.com/iframe/4937/000000/width/960/0/00:00:00>

URL <http://www.geo.uzh.ch/en/units/giva/research.html>

A key research question for the cognitive cartographer is “How do (Geo)graphic Information Displays work?” Better yet: how do we need to design GID’s to support effective, efficient, and affective inference and decision making, and ultimately spatial behavior? It turns out that display DESIGN is not the most important factor. In an increasingly mobile information society, the use CONTEXT of the display is equally important. This, aside from the key factor: human USER. The design of cognitively supportive (geo)graphic information displays must include/consider the decision maker (i.e., individual and group differences including their spatial abilities, perception/cognitive capacities, background, training, etc.) as well as the use context of the display (i.e., uncertainty, time pressure, other users, environmental factors, etc.).

References

- 1 Kübler, I., Richter, K.-F., Fabrikant, S.I. *Visualization of risk and uncertainty to communicate avalanche hazard: An Empirical Study..* Proceedings, 28th International Cartographic Conference, International Cartographic Association, Jul. 2-7, 2017, Washington, D.C. USA. 2017.
- 2 Griffin, A.L., White, T. , Fish, C., Tomio, B., Huang, H., Robbi Sluter, C., Meza Bravo, J.V., Fabrikant, S.I., Bleisch, S., Yamada, M., Picanço, Jr., P., Jr. *Visualization of risk and uncertainty to communicate avalanche hazard.* International Journal of Cartography, DOI: 10.1080/23729333.2017.1315988.

3.4 Set Visualization with Maps

Stephen G. Kobourov (*University of Arizona – Tucson, US*)

License © Creative Commons BY 3.0 Unported license
© Stephen G. Kobourov

Joint work of Yifan Hu, Sergey Pupyrev

Main reference Yifan Hu, Emden R. Gansner, Stephen G. Kobourov: “Visualizing Graphs and Clusters as Maps”, IEEE Computer Graphics and Applications, Vol. 30(6), pp. 54–66, 2010.

URL <http://dx.doi.org/10.1109/MCG.2010.101>

Visualization techniques help us understand and analyze complicated datasets and allows us to perceive relationships, patterns, and trends. While statistical techniques may determine correlations among the data, visualization helps us frame what questions to ask. Many interesting data sets can be visualized as graphs: the vertices in the graph are the objects of interest (for example, researchers) and a link between two vertices in the graph indicates a relationship (for example, research collaboration). Graph visualization aims to present such data in an effective and aesthetically appealing way. We describe ways to visualize datasets with the help of conceptual maps as a representation metaphor. While graphs often require considerable effort to comprehend, a map representation is more intuitive, as most

people are familiar with maps and ways to interact with them via zooming and panning. We consider map representations of several interesting datasets: from movies on Netflix, to books on music on Last.fm, to maps of computer science.

3.5 Diagrams for Logic and Reasoning

John Howse (University of Brighton, GB)

License © Creative Commons BY 3.0 Unported license
© John Howse

Traditional diagrammatic representations of set-based systems include Venn diagrams and Euler Diagrams. Euler diagrams are well-matched to meaning and contain free rides where some derivations are readable from the diagram. Spider diagrams extend Euler diagrams to include the representation of elements. Sound and complete spider diagram reasoning systems have been developed, with expressiveness equivalent to first order monadic logic. More expressive notations have been developed including constraint diagrams, for software system specification, and concept diagrams, for ontology representation. These notations include the representation of relations. Finally, multi-diagrams can express information that would become cluttered when represented in a single diagram. An open question is when is it appropriate to represent information with multi-diagrams rather than single diagrams.

3.6 The Painters Problem

Wouter Meulemans (TU Eindhoven, NL)

License © Creative Commons BY 3.0 Unported license
© Wouter Meulemans

Joint work of Arthur van Goethem, Irina Kostitsyna, Marc J. van Kreveld, Wouter Meulemans, Max Sondag and Jules Wulms

Main reference Arthur van Goethem, Irina Kostitsyna, Marc J. van Kreveld, Wouter Meulemans, Max Sondag, Jules Wulms: “The Painter’s Problem: covering a grid with colored connected polygons”, CoRR, Vol. abs/1709.00001, 2017.


URL <http://arxiv.org/abs/1709.00001>

Small multiples are a powerful visualization technique exploiting juxtaposition for concurrently displaying data under different conditions. Set visualization is characterized often by its containers that link together the elements belong to the same set. The typical Euler diagram is well known and intuitive, but suffers from drawback such as color perception of blended semi-transparent colored overlays.

With the eventual goal of a set-visualization technique that uses a disjoint representation and combines well with small multiples, we study the following problem:

Given a grid of square cells (see Figure 1, left), where each is either “red”, “blue”, “purple” (both red and blue) or “white” (neither), can we partition each purple cell into red and blue pieces, such that the union of all red cells and red pieces is connected, and likewise for blue? (refer to Figure 1, right)

We provide a characterization of instances that admit such a partition, allowing for an efficient algorithmic test. Moreover, if a partition is possible, we bound the maximal number of pieces needed for a single purple cell to five. If there are no white cells present, we can even improve this bound to two. That is, we can then partition each purple cell into a single red and a single blue piece, such that the result meets the connectivity constraint.



■ **Figure 1** Illustration of the solution

3.7 UnTangle Map: Visual Analysis of Probabilistic Multi-Label Data

Nan Cao (Tongji University, CN)


License  Creative Commons BY 3.0 Unported license
© Nan Cao

Data with multiple probabilistic labels are common in many situations. For example, a movie may be associated with multiple genres with different levels of confidence. Despite their ubiquity, the problem of visualizing probabilistic labels has not been adequately addressed. Existing approaches often either discard the probabilistic information, or map the data to a low-dimensional subspace where their associations with original labels are obscured. In this talk, we introduce a novel visual technique, UnTangle Map, for visualizing probabilistic multi-labels. In our proposed visualization, data items are placed inside a web of connected triangles, with labels assigned to the triangle vertices such that nearby labels are more relevant to each other. The positions of the data items are determined based on the probabilistic associations between items and labels. UnTangle Map provides both (a) an automatic label placement algorithm, and (b) adaptive interactions that allow users to control the label positioning for different information needs. Our work makes a unique contribution by providing an effective way to investigate the relationship between data items and their probabilistic labels, as well as the relationships among labels. Our user study suggests that the visualization effectively helps users discover emergent patterns and compare the nuances of probabilistic information in the data labels.

4 Working Groups

4.1 Mapifying the Genome

Radu Jianu (City, University of London, GB), Martin Krzywinski (BC Cancer Research Centre – Vancouver, CA), Luana Micallef (Aalto University, FI), and Hsiang-Yun Wu (TU Wien, AT)

License  Creative Commons BY 3.0 Unported license
© Radu Jianu, Martin Krzywinski, Luana Micallef, and Hsiang-Yun Wu

The human genome is a 1-dimensional structure of approximately 3 billion bases arranged across 23 pairs of chromosomes. Two challenges arise when attempting to visualize whole-genome data using the common and traditional genomic position axis, as in Figure 2.A. These issues are due to the fact that the data often annotate the small parts of the genome that code for proteins (exons), whose position on the genome is largely a byproduct of random shuffling during evolution and does not directly relate to the function of the genes (thus the reason why the position of the genes in a human genome is different from that of a mouse, as shown in Figure 2.B).

First, only about 2-3% codes for proteins, so the regions in which data appear are very sparse. This imposes a limit on the visibility of high-resolution elements in figures and visualizations, without smart down-sampling many of the data bins across regions of interest are smaller than a pixel or beyond visual acuity. Second, adjacency of genes along the genomic position axis does not generally relate to similar function or pathway. For example, the two genes that are involved the metabolism of tryptophan into melatonin are TPH, which is on chromosome 11 and ASMT, which is on chromosome X. The position of these genes is very difficult to guess on whole-genome plots and any correlation in values is essentially impossible to follow because the gene's positions are so far apart. In some cases the use of a position axis is helpful, such as whole-genome displays that show copy number variation, often shown in a format similar to Figure 2.C, which can quickly show large-scale structural changes (e.g. loss or gain of a chromosome arm) but cannot communicate the functional consequences of these changes.

We propose to address these two issues by remapping the axis of the genome from one based on position to one based on function. This would have the effect of creating a fixed coordinate system in which parts would be associated with function (e.g. cellular membrane, cell cycle, apoptosis, etc), thereby satisfying the Gestalt grouping principle of proximity, which states that elements that are positioned close to one another are perceived as related semantically. We argue for the need for a reordering that can be used in 1-dimensional data encodings (e.g. along a line) and 2-dimensional (similar to a map). This kind of remapping is commonly applied to gene-centric visualizations of the presence of mutations (or other structural or functional disturbance) (Figure 3), in which tens of genes are grouped by function. Our proposal extends this to the full genome.

Our approach is as follows. We would first define a similarity metric between genes based on their functional similarity, as defined in the Gene Ontology [2] or another resource. Multiple criteria for similarity can be used and possibly tuned to specialize applications, in which some groupings are more relevant than others. Second, we would apply clustering and layout algorithms (e.g., spring embedder, stochastic neighborhood embedding) to derive a 2-dimensional layout of the genes. Our expectation would be that genes that are close in this layout are functionally similar. This layout would be reshaped to fit into a rectangular area suitable as a canvas on which data can be drawn, in a similar way as already done for

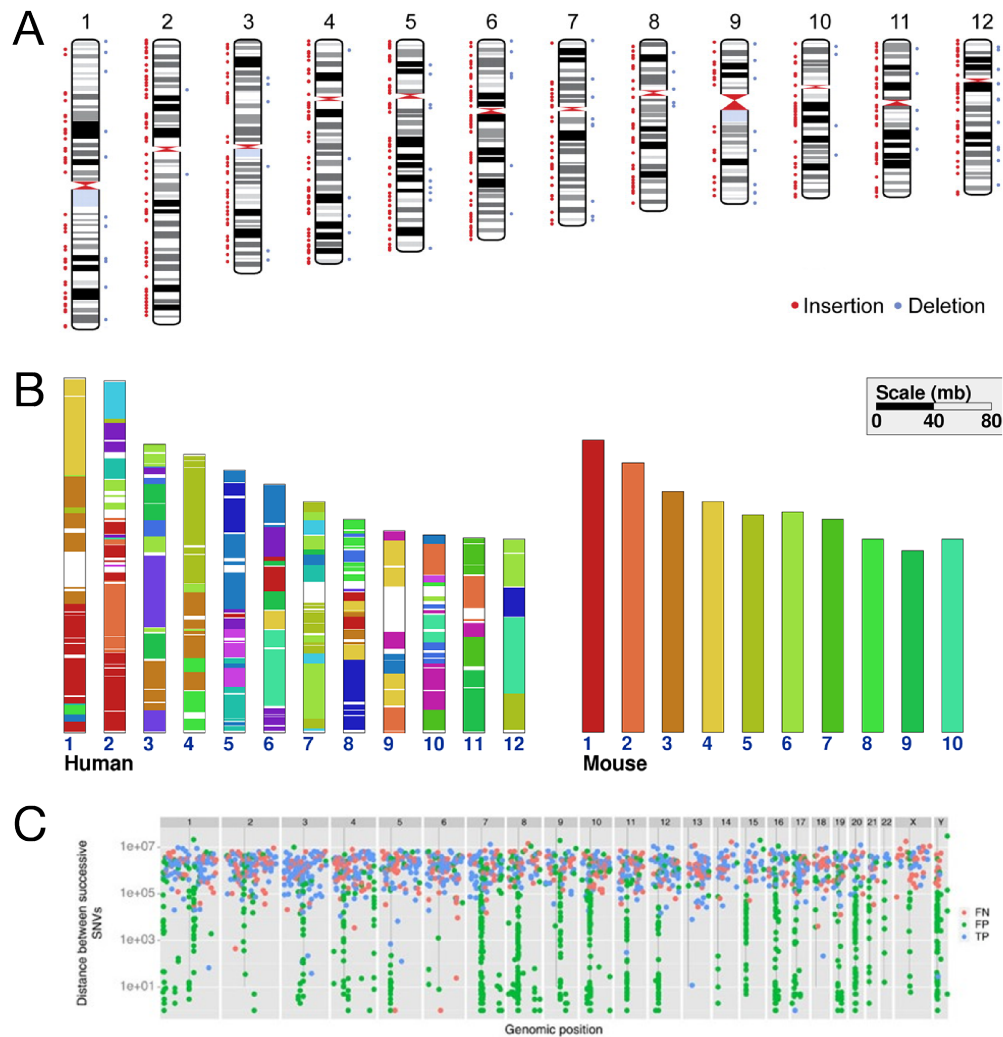


Figure 2 The human genome is 3 billion bases long. The longest chromosome (chr1) is about 250 million bases. Drawing data at single-base resolution is essentially impossible – the smallest visually discernible element in a figure covers about 250,000 bases. (A) The exact position of features such as mutations (insertions, deletions, SNPs) is impossible to assess. We cannot say which genes are affected by these mutations and, more importantly, what cellular function may be impacted by these changes. Figures like this help in understand the overall numbers and proportions of these changes but do not communicate the consequences of these changes. Adapted from Figure 3 in [6]. (B) The position of genes on the genome is driven by shuffling of genetic material within and among chromosomes during evolution. Shown here are regions of sequence similarity between human and mouse genomes. Source: [4]–Figure 2 (C) Features in data across an entire genome are very difficult to assess. Figures like this can be helpful to assess very large structural changes – when large contiguous parts of the genome are affected (e.g. deletion of a large region, or a chromosomal arm). Source: [1].




Figure 3 The idea of ordering genes by function is commonly applied when displaying information for small sets of genes that are recurrently mutated in disease. Here genes are grouped as “pancreatic cancer genes”, “chromatin remodelling”, “DNA damage report”, “axon guidance” and “known oncogenes”. Our proposal would extend this concept across the entire genome with gene order either globally fixed or based on specific applications. Source: [5].

example for graphs [3]. Third, using this 2-dimensional arrangement, we would derive a 1-dimensional order by choosing a path through the genes in 2-d (space filling, TSP, etc). This step would create an axis that would unfold the 2-d arrangement for linearized display.

This approach can be extended to generate a mapping between every base in the genome and a functionally correlated 2-dimensional position. Instead of laying out genes, their neighboring regulatory elements and inter-genetic regions would also be placed around the gene’s positions so that data that might relate to the gene’s activity would be drawn next to the gene. Regions of the genome very distant from genes (e.g. > 1 Mb away) that could not be unambiguously mapped to a function could be relegated into a separate part of the display. The output of this proposal would be coordinate transform file. Initially, it would remap intervals in the genome that correspond to genes to a point (or region) in a 2-dimensional rectangle and a 1-dimensional line. Subsequently, the file would include more parts of the genome, such as the regulatory regions and inter-genetic regions.

References

- 1 Tyler S Alioto, Ivo Buchhalter, Sophia Derdak, Barbara Hutter, Matthew D Eldridge, Eivind Hovig, Lawrence E Heisler, Timothy A Beck, Jared T Simpson, Laurie Tonon, et al. A comprehensive assessment of somatic mutation detection in cancer using whole-genome sequencing. *Nature communications*, 6:10001, 2015.
- 2 Michael Ashburner, Catherine A Ball, Judith A Blake, David Botstein, Heather Butler, J Michael Cherry, Allan P Davis, Kara Dolinski, Selina S Dwight, Janan T Eppig, et al. Gene ontology: tool for the unification of biology. *Nature genetics*, 25(1):25, 2000.
- 3 Emden R Gansner, Yifan Hu, and Stephen Kobourov. Gmap: Visualizing graphs and clusters as maps. In *Visualization Symposium (PacificVis), 2010 IEEE Pacific*, pages 201–208. IEEE, 2010.

- 4 Amit U Sinha and Jaroslaw Meller. Cinteny: flexible analysis and visualization of synteny and genome rearrangements in multiple organisms. *BMC bioinformatics*, 8(1):82, 2007.
- 5 Nicola Waddell, Marina Pajic, Ann-Marie Patch, David K Chang, Karin S Kassahn, Peter Bailey, Amber L Johns, David Miller, Katia Nones, Kelly Quek, et al. Whole genomes redefine the mutational landscape of pancreatic cancer. *Nature*, 518(7540):495, 2015.
- 6 Jinchuan Xing, Yuhua Zhang, Kyudong Han, Abdel Halim Salem, Shurjo K. Sen, Chad D. Huff, Qiong Zhou, Ewen F. Kirkness, Samuel Levy, Mark A. Batzer, and Lynn B. Jorde. Mobile elements create structural variation: Analysis of a complete human genome. *Genome Research*, 19(9):1516–1526, 2009.

4.2 Euler Diagrams drawn with Ellipses Area Proportionally (EDEAP)

Fadi Dib (Gulf University for Science&Technology – Mishreff, KW), Peter Rodgers (University of Kent – Canterbury, GB), Michael Wybrow (Monash University – Caulfield, AU)

License  Creative Commons BY 3.0 Unported license
© Fadi Dib, Peter Rodgers, Michael Wybrow


This working group addressed the research question “Can we draw general area-proportional Euler diagrams with ellipses?”. The goal was to increase the number of sets that can be drawn using Ellipses, from the current state of the art, which is limited to 3 [1]. We also aim to improve on the performance of techniques that use circles [3]. The target was to develop a web based software system that can accurately layout a larger number of diagrams compared to these previous techniques. This would enable those wishing to visualize such diagrams to access a more effective tool.

The motivation behind this work is the demand for area proportional diagrams to illustrate data in areas such as medicine, biosciences and numerous other disciplines. The goal is to ensure that the areas of overlapping regions are directly proportional to the cardinality in the data. The output of current tools is often inaccurate, in particular the circles used in current tools are known to be inaccurate for three sets forming Venn-3 [2]. Such a failure at a small set size is multiplied as the number of sets increase. Fortunately, it is known that far more accurate diagrams can be drawn when ellipses are used in place of circles, accuracy rates well over 90% can be achieved for Venn-3 diagrams [1].

In order to achieve the goals, a JavaScript software system was started during the seminar, running a basic hill climber, and limited number of criteria merged into a weighted sum to form a fitness function. This result from the seminar demonstrated promise that, with refinement, the tool could meet the required goals.

Since the seminar, the group has regularly met remotely to refine the system. Much effort has gone into improving the area measurements, the calculation of fitness components, the testing framework and speed of optimization. However, before release (and write up as a research paper) the following activities need to be completed:

- Identification and Implementation of additional criteria. Local optima occur when ellipses that should intersect are entirely contained in on another, or are entirely separate. This leads to situations where there are no incremental moves to ensure the ellipses overlap as required.
- Testing of various optimizers. The current hill climber is viewed as overly simplistic and a more effective optimizing algorithm will be implemented, perhaps based on simulated annealing.



■ **Figure 4** Output from an early version of the EDEAP software.

- The testing framework needs improvement. The weights for the criteria must be set (in line with both objective data and subjective view of what a good diagram might look like). Of particular note is that the result must be compared against the current state-of-the-art, the software of [3].

References


- 1 Luana Micallef and Peter Rodgers. eulerape: Drawing area-proportional 3-venn diagrams using ellipses. *PLOS ONE*, 9(7):1–18, 07 2014. URL: <http://www.eulerdiagrams.org/eulerAPE>, doi:10.1371/journal.pone.0101717.
- 2 P. Rodgers, G. Stapleton, J. Flower, and J. Howse. Drawing area-proportional euler diagrams representing up to three sets. *IEEE Transactions on Visualization and Computer Graphics*, 20(1):1–1, Jan 2014. doi:10.1109/TVCG.2013.104.
- 3 L. Wilkinson. Exact and approximate area-proportional circular venn and euler diagrams. *IEEE Transactions on Visualization and Computer Graphics*, 18(2):321–331, Feb 2012. doi:10.1109/TVCG.2011.56.

4.3 Spatially Informative Set Visualization


Thom Castermans (TU Eindhoven, NL), Mereke van Garderen (Universität Konstanz, DE), Wouter Meulemans (TU Eindhoven, NL), Martin Nöllenburg (TU Wien, AT), and Xiaoru Yuan (Peking University, CN)

License © Creative Commons BY 3.0 Unported license
© Thom Castermans, Mereke van Garderen, Wouter Meulemans, Martin Nöllenburg, and Xiaoru Yuan


Set visualization for elements with (geo)spatial locations is often dominated by the elements being shown in exact spatial locations. We set out to investigate methods where we do not keep elements at their exact location, but rather let them move away from it in order to further clarify the set visualization. However, elements are not allowed to move arbitrarily: rather, they must remain near their spatial location, to allow high-level spatial patterns to remain visible. Since spatial accuracy is not perfect, but not completely lost either, we refer to this as *spatially informative set visualization*.



■ **Figure 5** Point set of 15 points, each specifying one or two colors.



■ **Figure 6** The shortest plane support for the point set in Figure 5.



■ **Figure 7** A Kelp-style rendering of the support shown in Figure 6.

Definitions. The first problem we faced was to define what it means to “clarify” a set visualization. We formalize set visualization using the existing concept of a *support graph*: a graph G on the elements is called a support of a set system, if each set induces a connected subgraph in G . Using such a support, we can now define what a good set visualization is, by setting requirements on the support graph. In particular, we set the following goals for a good support G : (a) G should be plane, that is, have straight-line edges without intersections; (b) G should be short, that is, the sum of edge lengths is small.

We would like to study how much the support improves if we are allowed to move points. However, little is known about finding good supports even with fixed elements. We thus first defined and studied the problem below (refer to Figures 5 and 6 for an illustration). Here we use color to indicate set memberships, as is common in research involving hypergraph supports.

SHORTESTPLANE SUPPORT: Consider a set P of points in \mathbb{R}^2 and a set C of colors. A function $\chi: P \rightarrow \mathcal{P}(C) \setminus \{\emptyset\}$ maps each point in P to a nonempty subset of the colors in C . Find a plane graph G with vertex set P such that G has minimal total length and G is a support: for each color $c \in C$, the subgraph induced by $\{p \in P \mid c \in \chi(p)\}$ is connected.

We use **PLANE SUPPORT** to refer to the variant where the length of G is not considered and we just want to decide whether such a graph G exists at all. We call a point p a k -colored point if $|\chi(p)| = k$. We observe that plane supports lead to Kelp-style rendering [5, 7] where sets intersect only at common members; see also Figure 7. A short support graph minimizes ink, following Tufte’s principles [8].

A sample of related work. Planar supports without fixed elements [4] or short supports with fixed elements [1, 6] have been studied, but not their combination. The case where each element belong exactly to one set is also studied in a Steiner setting [2], where additional points may be placed. If all points are 1-colored, **PLANE SUPPORT** reduces to finding nonintersecting spanning trees. For $|C| = 2$, this was already solved by Bereg et al. [3]. To the best of our knowledge, the problem is still open for $|C| > 2$.

Results. We studied the above mentioned problems and arrived at a number of interesting results. First, we observe that there always exists a plane support for a given point set, as long as there is at least one point which is mapped to all $|C|$ colors by χ .

► **Observation 1.** **PLANE SUPPORT** is always TRUE, if there exists at least one $|C|$ -colored point.

However, as soon as we insist on having a short length, the problem turns out to be NP-hard, even under some seemingly simplifying restrictions.

► **Theorem 2.** *SHORTESTPLANE SUPPORT is NP-hard, even if one or more of the following conditions hold: (a) $|C| = 2$, (b) either there is no $|C|$ -colored point or all 1-colored points are of the same color; (c) the resulting support graph G must be a tree.*

The last restriction in the theorem above insists that our resulting support graph G should be a tree. If we look more closely at this condition, this implies that the $|C|$ -colored points must induce a connected subtree of G . In other words, if $|C| = 2$, there is a backbone connecting all 2-colored points; the remaining 1-colored points are connected to this backbone via single-colored trees. One may wonder whether starting with the minimum spanning tree as such a backbone can lead to an approximation algorithm. Unfortunately, we must answer this negatively, even if we add the requirement that G must be a tree.

► **Lemma 3.** *Assume $|C| = 2$. There exists a family \mathcal{F} of point sets such that, for each point set $P \in \mathcal{F}$, the Euclidean minimum spanning tree of its 2-colored points is not a subgraph of any plane support tree whose length is within a constant factor of the shortest plane support tree of P .*

Finally, we developed an ILP that solves SHORTESTPLANE SUPPORT. It can be customized to allow a number of crossings, weigh crossings into the optimization, and/or require that G is a tree. The ILP also allows points to have a set of candidate positions that can be selected, such that we can start developing spatially informative set visualizations.

Ongoing research. We are aiming to use our ILP to investigate how various conditions affect the length of the support graph, such as the difference between forcing G to be a tree or not, allowing a number of intersections, etc. Also, we plan to use it to compare how well (existing) heuristic and approximation algorithms work. Finally, we aim to use it to investigate the effect of allowing elements to move from their original position. Care needs to be taken to measure structural differences, rather than shortening effects obtained simply by moving points closer to one another.

References

- 1 H. A. Akitaya, M. Löffler, and C. D. Tóth. Multi-colored spanning graphs. In Y. Hu, M. Nöllenburg (eds), *International Symposium on Graph Drawing and Network Visualization*, LNCS 9801, pp. 81–93, 2016.
- 2 S. Bereg, K. Fleszar, P. Kindermann, S. Pupyrev, J. Spoerhase, and A. Wolff. Colored Non-Crossing Euclidean Steiner Forest. In K. Elbassioni, K. Makino (eds), *International Symposium on Algorithms and Computation*, LNCS 9472, pp. 429–441, 2015.
- 3 S. Bereg, M. Jiang, B. Yang, and B. Zhu. On the red/blue spanning tree problem. *Theoretical Computer Science*, 412(23):2459–2467, 2011.
- 4 K. Buchin, M. van Kreveld, H. Meijer, B. Speckmann, and K. Verbeek. On Planar Supports for Hypergraphs. *Journal of Graph Algorithms and Applications*, 14(4):533–549, 2011.
- 5 K. Dinkla, M. van Kreveld, B. Speckmann, and M. A. Westenberg. Kelp Diagrams: Point Set Membership Visualization. *Computer Graphics Forum*, 31(3pt1):875–884, 2012.
- 6 F. Hurtado, M. Korman, M. van Kreveld, M. Löffler, V. Sacristán, A. Shioura, R. I. Silveira, B. Speckmann, T. Tokuyama. Colored Spanning Graphs for Set Visualization. *Computational Geometry: Theory and Applications*, to appear, 2017.
- 7 W. Meulemans, N. H. Riche, B. Speckmann, B. Alper, and T. Dwyer. KelpFusion: A hybrid set visualization technique. *IEEE Transactions on Visualization and Computer Graphics*, 19(11):1846–1858, 2013.
- 8 E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press (Cheshire, CT), 1983.

4.4 Set Visualization Using the Metro Map Metaphor

Robert Baker (University of Kent, UK), Nan Cao (Tongji University, CN), Yifan Hu (Yahoo! Research, US), Michael Kaufmann (University of Tübingen, DE), Stephen Kobourov (University of Arizona, DE), Tamara Mchedlidze (Karlsruhe Institute of Technology, DE), Sergey Pupyrev (Facebook, US), Alexander Wolff (Universität Würzburg, DE)

License © Creative Commons BY 3.0 Unported license

© Robert Baker, Nan Cao, Yifan Hu, Michael Kaufmann, Stephen Kobourov, Tamara Mchedlidze, Sergey Pupyrev, Alexander Wolff

We consider a visualization style for hypergraphs that is inspired by schematic metro maps. Such maps are common for urban citizens, who all know that the stations traversed by the same colored curve belong to the same metro line. This intuitive understanding of grouping have been employed to visualize other abstract data forming hypergraphs. For example, Foo [3] turns personal memories into a metro map, Nesbitt [4] and Stott et al. [10] use the metro map metaphor to visualize relationships between PhD theses and items of a business plan, Sandvad et al. [7] for building Web-based guided tour systems, and Seskevce [8] uses it for visualizing historical events. One of the most popular applications is the visualization of the movies and movie genres by the creators of the website Vodkaster.


We formalize the problem of constructing such a visualization for given hypergraph as follows. Let $H = (V, \mathcal{E})$ be a hypergraph with vertex set V and edge set $\mathcal{E} \subseteq 2^V$. A *metro-map drawing* of H is a graphical representation where each node in V is depicted by a point in the plane and each hyperedge $e \in \mathcal{E}$ by an open continuous curve that passes through the points corresponding to the vertices in e . In case two hyperedges contain the same vertex, their curves both pass through the point representing this vertex and may either *touch* or *cross* at this point. We call the latter situation a *vertex crossing*, and we call a crossing of hyperedge curves that is not a vertex crossing an *edge crossing*. A metro-map drawing of a hypergraph is called *monotone* if all hyperedge curves are monotone with respect to the x-axis.

Some Simple Observations


Since both vertex and edge crossing may impair the readability of the metro-map drawing of a hypergraph, we want to characterize the hypergraphs that can be represented without, or with a few, vertex and edge crossings. We observe that each hypergraph with at most four hyperedges can be represented without vertex and edge crossings. We call a hypergraph $H = (V, \mathcal{E})$ *k-vertex-complete* for some $k \leq |\mathcal{E}|$ if any subset $E \subset \mathcal{E}$ of at most k hyperedges has a distinct vertex in common, that is, there is an injective function $f: E \rightarrow V$ such that $f(E) \in \bigcap_{e \in E} e \neq \emptyset$. We call $|\mathcal{E}|$ -vertex-complete hypergraphs simply *vertex-complete*. We observe that a 2-vertex-complete hypergraph with five hyperedges does not have a metro-map drawing without vertex and edge crossings. This follows simply from the fact that K_5 is not planar. Next, we consider drawings with vertex crossings but without edge crossings. We can show that every vertex-complete hypergraph admits a metro-map drawing without edge crossing. The idea behind the proof is to exploit the vertices to realize the intersections among the hyperedge curves. For an example, see Figure 8.

A Heuristic

For practical applications, we propose a heuristic that constructs a metro-map drawing of a given hypergraph. The heuristic consists of four steps.



■ **Figure 8** A metro-map drawing is a visualization of a hypergraph where the metro lines represent hyperedges and the stations represent hypervertices. (a) A metro-map drawing of a vertex-complete hypergraph with three hyperedges (violet, red, yellow), (b) a mirrored copy of (a), and (c) a metro-map drawing of a vertex-complete hypergraph with four hyperedges. The drawing is constructed recursively by routing the new (blue) hyperedge through a vertex that is only contained the new hyperedge and then through the concatenation of (a) and (b). The new hyperedge does not contain any vertex of (a), and it does contain every vertex of (b).



■ **Figure 9** A preliminary drawing of the support graph of a dataset

In the first step, we simplify the hypergraph, by ignoring all vertices that belong to a single hyperedge, and contracting all vertices that belong to the same set of hyperedges.

In the second step, we construct a so-called support graph. A *support* of a hypergraph $H = (V, \mathcal{E})$ is a graph $G = (V, E)$ with the property that, for each hyperedge e of H , the graph $G[e]$ induced by e is connected. A support is *path-based* if $G[e]$ is Hamiltonian. It is NP-complete to compute a path-based support with the minimum number of edges [2]. We propose a heuristic algorithm for finding a path-based support for a given hypergraph.

In the third step, we lay out the support graph in the plane. In doing so, we try to ensure that all vertices belonging to the same hyperedges lie close by and that the paths representing the hyperedges have simple shapes. A preliminary drawing at this step is shown in Figure 9.

Finally in the fourth step, we feed the above drawing into a mixed-integer program (or some other existing algorithm) that generates a metro map layout [5].

The above steps can be implemented in many possible ways. The performance of our approach needs to be compared with existing similar approaches [1, 6, 9] experimentally.

References

- 1 Basak Alper, Nathalie Henry Riche, Gonzalo Ramos, and Mary Czerwinski. Design study of LineSets, a novel set visualization technique. *IEEE Trans. Vis. Comput. Graph.*, 17(12):2259–2267, 2011. doi:10.1109/TVCG.2011.186.
- 2 Ulrik Brandes, Sabine Cornelsen, Barbara Pampel, and Arnaud Sallaberry. Path-based supports for hypergraphs. *J. Discrete Algorithms*, 14:248–261, 2012. doi:10.1016/j.jda.2011.12.009.
- 3 Brian Foo. The memory underground. <http://memoryunderground.com>.
- 4 Keith V. Nesbitt. Getting to more abstract places using the metro map metaphor. In *Proc. 8th Int. Conf. Inform. Vis. (IV'04)*, pages 488–493. IEEE, 2004. doi:10.1109/IV.2004.1320189.
- 5 Martin Nöllenburg and Alexander Wolff. Drawing and labeling high-quality metro maps by mixed-integer programming. *IEEE Trans. Vis. Comput. Graph.*, 17(5):626–641, 2011. doi:10.1109/TVCG.2010.81.
- 6 Francesco Paduano and Angus Graeme Forbes. Extended LineSets: a visualization technique for the interactive inspection of biological pathways. *BMC Proceedings*, 9(Suppl. 6)(S4):13 pages, 2015. doi:10.1186/1753-6561-9-S6-S4.
- 7 Elmer Sandvad, Kaj Grønbaek, Lennert Sloth, and Jørgen Lindskov Knudsen. A metro map metaphor for guided tours on the Web: the Webwise guided tour system. In Vincent Y. Shen, Nobuo Saito, Michael R. Lyu, and Mary Ellen Zurko, editors, *Proc. 10th Int. World Wide Web Conf. (WWW'01)*, pages 326–333. ACM, 2001. doi:10.1145/371920.372079.
- 8 Jure Seskevce. SNAP metromaps. <http://metromaps.stanford.edu>.
- 9 Dafna Shahaf, Carlos Guestrin, and Eric Horvitz. Trains of thought: Generating information maps. In *Proc. 21st Int. Conf. World Wide Web (WWW'12)*, pages 899–908. ACM, 2012. doi:10.1145/2187836.2187957.
- 10 Jonathan M. Stott, Peter Rodgers, Remo Aslak Burkhard, Michael Meier, and Matthias Thomas Jelle Smis. Automatic layout of project plans using a metro map metaphor. In *Proc. 9th Int. Conf. Inform. Vis. (IV'05)*, pages 203–206. IEEE, 2005. doi:10.1109/IV.2005.26.

4.5 Visual Analytics of Sets/Set-Typed Data and Time: Challenges and Opportunities

Daniel Archambault (Swansea University, GB), Kerstin Bunte (University of Groningen, NL), Sara Irina Fabrikant (Universität Zürich, CH), John Howse (University of Brighton, GB), Andreas Kerren (Linnaeus University – Växjö, SE), and Silvia Miksch (TU Wien, AT)

License © Creative Commons BY 3.0 Unported license


© Daniel Archambault, Kerstin Bunte, Sara Irina Fabrikant, John Howse, Andreas Kerren, and Silvia Miksch

This breakout group has been composed of six people in total (see Figure 10) and represented the interdisciplinary characteristics of the overall Dagstuhl seminar consisting of experts in geographic information systems/geovisual analytics, visualization/visual analytics, formal diagrammatic reasoning, visual modelling of logic-based systems, machine learning, and data discovery in databases.

The aim of this group was to discuss, elaborate, and structure possible challenges and opportunities of set visualizations that change over time. As a first step, we discussed the communalities and specifics of sets, set-typed data, and time, to arrive at a common understanding of the various terminologies used in various cognate fields (see Figure 11). We collected state-of-the-art articles and research papers about set visualizations, spatio-temporal visualizations, visualization task taxonomies and classifications, etc. We examined various



■ **Figure 10** (front line) Andreas Kerren, Silvia Miksch, John Howse, Kerstin Bunte; (back line) Sara Fabrikant, Daniel Archambault



■ **Figure 11** Structuring the various terminologies used in the different – however still computer science-related – communities.

application domains, ranging from social networks to the digital humanities, and to various kinds of flows. Finally, we defined twelve current challenges and future opportunities, which the group aims to elaborate in more detail in the next steps.

Participants

- Daniel Archambault
Swansea University, GB
- Robert Baker
University of Kent –
Canterbury, GB
- Kerstin Bunte
University of Groningen, NL
- Nan Cao
Tongji University – Shanghai, CN
- Thom Castermans
TU Eindhoven, NL
- Fadi Dib
Gulf University for
Science&Technology –
Mishreff, KW
- Sara Irina Fabrikant
Universität Zürich, CH
- John Howse
University of Brighton, GB
- Yifan Hu
Yahoo! Research – New York, US
- Radu Jianu
City, University of London, GB
- Michael Kaufmann
Universität Tübingen, DE
- Andreas Kerren
Linnaeus University – Växjö, SE
- Stephen G. Kobourov
University of Arizona –
Tucson, US
- Martin Krzywinski
BC Cancer Research Centre –
Vancouver, CA
- Tamara Mchedlidze
KIT – Karlsruher Institut für
Technologie, DE
- Wouter Meulemans
TU Eindhoven, NL
- Luana Micallef
Aalto University, FI
- Silvia Miksch
TU Wien, AT
- Martin Nöllenburg
TU Wien, AT
- Sergey Pupyrev
Facebook – Menlo Park, US
- Peter Rodgers
University of Kent –
Canterbury, GB
- Mereke van Garderen
Universität Konstanz, DE
- Alexander Wolff
Universität Würzburg, DE
- Hsiang-Yun Wu
TU Wien, AT
- Michael Wybrow
Monash University –
Caulfield, AU
- Xiaoru Yuan
Peking University, CN

