View Relations: An Exploratory Study on Between-View Meta-Visualizations

Søren Knudsen
University of Copenhagen
Copenhagen, Denmark
sknudsen@di.ku.dk

Sheelagh Carpendale
University of Calgary
Calgary, Alberta, Canada
sheelagh@ucalgary.ca

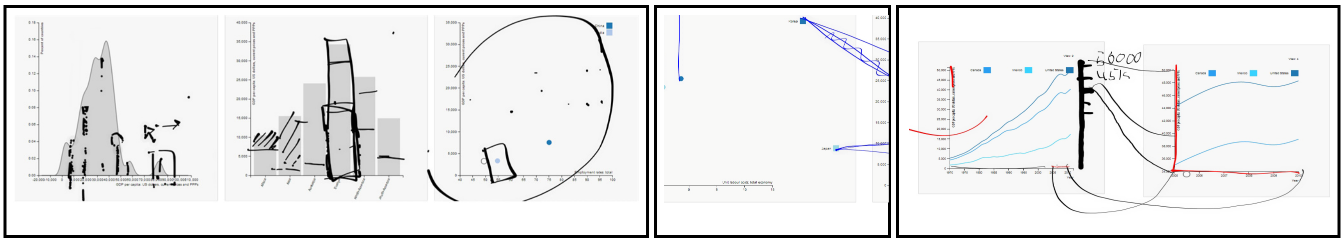

Figure 1. Examples of the sketches produced by the participants in our study. The participants sketched detailed ideas on top of visualization views we provided. For example, they considered brushing and linking (left), connected legend items to data bars (middle), and connected axes to indicate their scale relation (right).

ABSTRACT
In this paper, we explore the potential of using visual representations to support people in managing, organizing, and understanding relations between multiple visualization views. Multiple views can help people understand different facets of data and data processing, and are a crucial part of data analysis particularly when it is done collaboratively. Both the growing use of multiple views and the increasing display sizes have amplified the need to explore how to better help people to understand the relations between many views. To improve our understanding of how to visualize view relations, we invited visualization and interaction designers to critique and sketch representations of view relations. The participants provided design critiques, and sketched their own relation representations. Our findings expand the range and palette of ways of visually linking visualization views and suggest new directions for designing view relation representations to better support analysis with multiple views.

Author Keywords
Visualization; multiple and coordinated views; qualitative study; view relations; meta-visualization.

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

NordiCHI '16, October 23 - 27, 2016, Gothenburg, Sweden

ACM 978-1-4503-4763-1/16/10…$15.00
DOI: http://dx.doi.org/10.1145/2971485.2971566

INTRODUCTION
We explore the design space of representing relations between visualization views. Representing these visually and interactively has the potential to support people in managing, organizing, and understanding relations between multiple visualization views. As visualization research expands, and with the increasing demand from analysts for effective visualizations, multiple views are becoming more and more common. The prevalence of large displays that facilitate simultaneous display of several views has increased the need to provide effective support for understanding view relations. Research on coordinated and multiple views (CMV) has been a common theme in visualization research, including conferences devoted to the subject [28]. Many discussions have arisen about how CMVs and related techniques support people in analyzing data. These include discussions on linking common data between views [6,8,28], comparing data in different views [11], and preserving one's mental map from one view to the next [9]. This previous research focused either on introducing specific methods for showing view relations [8,9,12,32,38], or on creating multiple view systems to support analysis of a given dataset and its associated tasks [32].

Thus, previous work offers view relation methods and techniques that address specific problems. However, relatively little is known about the advantages and disadvantages of the different techniques, and only rarely have studies attempted to compare [14] or generalize them [11,20]. Despite evidence of the growing need for such techniques, we lack advice on how to use these techniques, their issues, and the potential factors to consider (e.g., the tasks to support, and potential interference of view relation techniques). Designing useful representations of view relations is challenging and complex. Currently, visualization designers have to piece together disparate knowledge of view coordination techniques. In response to this problem we study view relation techniques through an exploratory approach.

To gain a better understanding of this design space, we conducted a qualitative study, in which we asked nine visualization and interaction designers to critique existing ideas of representing view relations and to generate new ones. Our methodology draws on previous contributions in this area (e.g., [18,22,37]). During one-and-a-half hour sessions, the participants provided design critique, and sketched their own relation representations. From these activities, we analyzed the combined verbal, gestural, and sketched deliberations. From our work, we contribute:

1) A better understanding of the design space of techniques for representing view relations.
2) New, promising explicit ideas for representing view relations.
3) Suggestions as to how our findings might be useful in formulating new research questions about view relations.

TERMINOLOGY

For clarity we define the terminology we use in this paper:

Views as bounded areas that have their own use of spatial organization and display any part of a dataset using any visualization techniques. View boundaries may be represented visually using borders, backgrounds, or similar techniques (this definition draws upon [2,4,8,26]).

View relations as properties that people might use to associate views. This can include such factors as the data shown in views, or meta-data about views, such as the way visual variables are used in views to encode data, or the order in which people created views.

Representations of view relations as meta-visualizations that are intended to help people understand view relations. Representations of view relations cover all methods for showing these relations including: those that intersect the boundaries of views, matching encodings across views, or overlay visual marks on within-view visualizations to reveal between view relations.

RELATED WORK

Historically, cartographers have used insets similarly to the notion of views to show areas of maps at different scale, and represented their relations statically. In interfaces, such insets are often referred to as overview-plus-detail (see e.g., [7]). Ware & Lewis [38] introduced an interactive version of these separate magnified views using lines to indicate spatial relations between views, thus introducing a simple technique for showing scale relations.

Many visualization systems show multiple views that display different aspects of a dataset. North & Shneiderman [26] explained how CMV techniques allow people to explore data through a variety of interactions. Tweedie [35] argued the importance of considering meta-data visualizations. Weaver [39] added his notion of integrated meta-visualizations to this discussion. Baldonado et al. [2] provided design guidelines and North et al. [25] described techniques to implement CMV systems.

Brushing and linking is a common coordination technique, in which items selected in one view, are highlighted in other views. In addition, navigational coordination is often used to relate zoom and pan interactions in one view to other views. For many systems this remains sufficient, since there are only a few views, and they are positioned statically and visible throughout use of a system. However, most CMVs rely on interactions to let people discover view relations and have not used the design space we explore – persistent visual representations of relations between views. In 2007, Roberts [28] provided an excellent overview of the state of CMVs. He debated whether CMVs was a "solved problem" or if the visualization community had "barely scratched the surface of the subject". The amount of recent work (e.g., [6,8,9,17,21,32,36,43]) contributing new techniques for representing view relations suggests the latter.

Representing data flow has been studied both in visualization (e.g., [40]), and outside (e.g., [13]). Dunne et al. [9] combined data flow representations with visualizations of provenance. They connected individual views using lines, which indicate performed interactions (subset, pivot, and union), in a pan-and-zoom interface. The color of the lines matched visual data components (e.g., points, bars, etc.) in the views and these lines were drawn from border to border of the views. This technique represented interaction history, enabling people to reconstruct analysis trails. Other contributions have also aimed to support analysts' processes by visualizing analysis steps (e.g., [16,30]).

Hybrid visualizations, which combine existing visualization techniques [20], have also been the focus of contributions. Zhao et al. [43] combined tree-maps and node-link diagrams to reap the benefits of both techniques. Likewise, Henry et al. [17] presented a hybrid network visualization that combined the benefits of node-link diagrams and adjacency matrices to show both local and global structure. Aside sometimes lacking view boundaries, most hybrid visualizations fit our terminology. Collins & Carpendale [8] presented an alternative multi-view technique. They showed multiple 2D views arranged on planes in a virtual 3D environment and used line connections to show relations between views' data points. Navigating the 3D space allowed people to explore relations between datasets in different 2D representations such as scatterplots and tree-maps.

The visualization pipeline has also been considered as a focus of view relations. Tobiaz et al. [32] linked views through visualization of the visualization pipeline [5], in order to support co-located collaboration. Heer & Agrawala [15] discussed the idea of basing points of collaboration on this. In Lark, the pipeline was shown as a form of node-link tree. Nodes, which represented individual stages, were shown as circles. Views were connected to the tree as leaves. Thus similar views, were leaves on the same branch of the tree, giving a representation of views' relations.

While recent contributions (e.g., [11,20]) have aimed at systematizing some of these contributions, their relation to earlier taxonomies of CMV systems (e.g., [26]) remains unclear. Ragan et al. [27] provided a framework of provenance in visualization and data analysis, but did not focus on the visual means of showing provenance between views. Earlier work by Roberts [28] also discussed exploration processes and meta-information, with more focus on views. While visualization views are a familiar concept, it is quite disparate. Even according to our definition, views might be constructed by system designers (e.g., [29]), during interaction (e.g., [9]), or be a direct result of data (e.g., [17]). Since these aspects have received little attention in CMV research, we directly explore the potential of view meta-data and relations between views to reconcile these different concepts, within our notion of view relations. Thus, while the literature has suggested a great variety of view relations, we know little about their relative benefits, and lack an overview to understand, design, and evaluate them. Moreover, possibilities outside those suggested by the literature might exist. Therefore, we conducted a study to better understand the breadth and scope of showing view relations. We describe our methodology in the following.

STUDY METHODOLOGY

We conducted an exploratory study in which we invited visualization and interaction designers to critique and sketch representations of view relations. Our goal was to improve our understanding of this area and to expand the palette of possible representations of view relations. Before starting this, we considered our methodological options. Designing, implementing, and testing one single possible design would leave us with yet another point design, and limited insights into benefits and drawbacks of different techniques. Therefore, we considered approaches that would allow us to compare many ideas. Conducting a review of existing literature would provide us with an understanding of existing ideas and techniques. However, we were also interested in expanding the field. To go beyond simple interviewing, we chose to expand our understanding of the role of view relation representations by letting visualization and interaction designers discuss existing alternatives, and sketch their own new ideas.

To seed the discussion with existing alternatives, we implemented seven lo-fi prototypes. We chose diverse points from the existing range and limited the number to keep the study protocol manageable. This allowed us to get critiques from our participants, to gain more knowledge about the strengths and weaknesses of view relation representations, and to better understand their underlying properties. Previous work [33] has shown that providing participants with alternatives, results in more feedback. Therefore, we did not ask the participants to critique a single system, but to critique alternatives for representing view relations.

Additionally, we also asked the participants to sketch their own ideas, thus allowing them to extend the possibilities for design beyond the alternatives we provided. We argue that our choice of methodology offers a sweet spot between an expert review [34] and semi-structured interviews. In contrast to other evaluations of single representations of view relations (e.g., [9,32]), our approach opened several different opportunities such as understanding the breath of design possibilities, rather than discovering if a particular approach was to participants liking. The study took the form of a semi-structured inquiry-based interview, in which participants critiqued and sketched design alternatives. Our study design method allowed us to explore the advantages and disadvantages of different relation types and methods of representation. Thus, we obtained a broader understanding of view relations, and draw inspiration from the many varied designs created by our participants.

In our study, we focus on abstract visualization tasks, such as comparing sets of data, recording provenance, and navigating views, and less on the domain of the visualized data or tasks related to the data. Thus, based on Munzners' nested model for visualization design and validation [24], we believe that it is most sensible to invite visualization and interaction designers to participate.

Participants

For participants we selected people who were actively working on the design, implementation and evaluation of visualizations and visual interfaces. Our participants had all published in top venues in visualization, interaction design, or human computer interaction. They were all current active researchers. There were five men and four women, most either had a related MSc or were currently PhD students, and one had a related PhD.

Apparatus

During the study, the participants worked with seven alternative design scenes. Full resolution images and interactive versions are available from supplemental material online[1]. The scenes consisted of visualization views and view relation representations. The views could be dragged around using touch, which resulted in updates to the relation representations (e.g., lines connecting data points between views). Some scenes captured ideas from related work, while others were novel. Below we describe the intended focus of each of the scenes. We designed the scenes to work as conversation catalysts, and not as faithful reproductions of the work that inspired them.

[1] http://view-relations-study.sorenknudsen.com

Scene 1 (S1) consisted of an area chart, two bar charts and two scatterplots with legends. We used lines as relation representations to connect individual views to indicate performed filtration interactions, and enable people to understand analysis history, thus drawing on inspiration from Dunne et al. [9]. However, unlike their work, we showed links directly from data points and bars. **Scene 2** (S2) consisted of three scatterplots and a bar chart. The scatterplots had legends that used a common spatial layout. We used lines to connect individual data points to data points and representations of aggregate data in other views (e.g., data bars). Linking individual data points is a common MCV technique. However, work by Collins & Carpendale [8], inspired us to connect individual data to representations of aggregate data. **Scene 3** (S3) consisted of 5 line charts with legends. We used lines to connect views that showed data at different scale thus drawing inspiration from Ware & Lewis [38]. Their technique formed the basis for the variations of visual representations of scale relations in this scene. **Scene 4** (S4) consisted of four bar charts and four scatterplots with legends. We used lines to connect views based on a visual representation of the visualization pipeline. This was inspired by work by Tobiaz et al. [32]. **Scene 5** (S5) consisted of four bar charts and a line chart with a legend. We used lines to connect multiple views. Particularly, data from multiple bar charts were integrated data into a single line chart. Like S1, the work by Dunne et al. [9] inspired the design of S5. However, this scene used the union operation and pivot actions of their work. **Scene 6** (S6) consisted of two line charts and two scatterplots, that all had legends. We used lines to show relations to and from the legends. **Scene 7** (S7) consisted of a dendrogam and two scatterplots with legends. We used color to show relations between elements of the dendrogram and the scatterplots' colored axis marks.

The scenes were implemented in D3 [3] and ran in the Chrome browser. All scenes visualized data obtained from OECD (http://stats.oecd.org). The scenes were shown on an 84", 4k display at 30Hz supporting touch and pen interaction. The large display provided participants freedom to layout view arrangements and ample space to think [1].

Procedure

Each session lasted approximately one-and-a-half hours, and consisted of three phases. In the first phase, participants were briefed about the study; signed a consent form; and answered a short questionnaire about demographics and experience with data analysis, visualizations and the data and technologies used in the study. We then introduced the dataset used in the seven scenes.

The middle phase consisted of two parts. In part A, participants looked at, interacted with, and critiqued the seven scenes. In part B, participants sketched their own relation representations between scene views based on a description of the views. They used a digital pen to sketch on top of the same design scenes stripped from showing view relations. We counterbalanced the order of Part A and B. We found that participants in both groups contributed many novel designs, and observed almost all findings across the two groups. Similarly, we randomized the scene order to not favor specific scenes, although each participant went through the scenes in part A and B in the same order.

We probed participants with questions. During part A, we asked factual questions, descriptive questions, and evaluative questions. For instance, we asked participants to tell us the GDP in 2010 for Canada, to describe the relations between two views, and to state their preferences for the relations shown. After experiencing and discussing a scene, participants were invited to ask questions about the representations, letting us understand their interpretation. We continued to the next scene when participants had answered our questions and we theirs. During part B, we asked participants questions similar to those in part A, this time relating to what they sketched. For example, we asked participants to describe their sketches and choices of representations.

To conclude, we debriefed participants in a short semi-structured interview. We asked participants; (a) which advantages and disadvantages they observed from seeing the relations represented visually; (b) which relations seemed most important and why; and (c) which methods of representing relations seemed most useful and why. Finally, we asked participants about the study methodology, and how they thought the tools they used during the session had influenced their ideas and sketches (e.g., if they had been limited by the pen thickness or choice of colors).

Data Collection

We recorded the participant sessions using a video camera pointing towards the display from an approximate 30-degree angle and 4-meter distance, which showed the 84" display in the image center. We used an audio recorder installed above the display to record an intelligible signal, and captured display state during the sessions.

Analysis

We analyzed the recorded material based on grounded theory [19,31]. We started with concepts from the literature [6,8,9,11,17,20,21,28,29,32,36,38,43], and looked for new ideas and concepts while analyzing the data. Participants told us about their design ideas, criticisms and suggestions, verbally, via gestures, or with sketches and annotations. In analyzing the observations, we gave equal weight to all the different ways in which they indicated their ideas. We read the transcripts, we watched the videos, and we studied the captured screenshots. As coding concepts emerged, they often would include a range of possible participant indications. Likewise, a particular design for showing view relations would often make use of several of the concepts for which we coded. In the first analysis pass, the first author went through all the material, keeping notes of interesting moments while obtaining an overview. During this pass, a range of concepts were identified and discussed by the co-authors

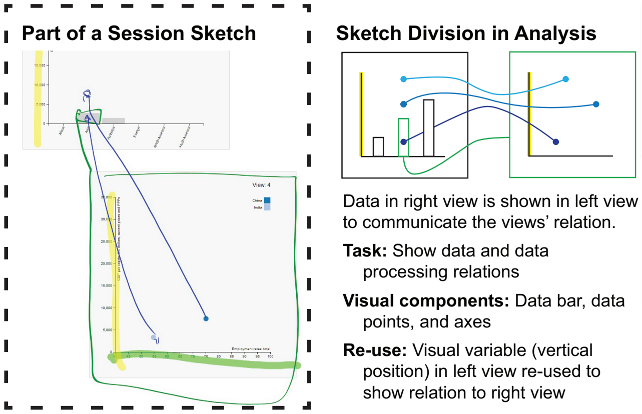

Figure 2. Process of analyzing screenshots from experiments. The figure shows a quarter of a scene 1 session sketch. The 63 sketches were analyzed by identifying view relation concepts, which are described in the results section.

	Concepts	Scenes								
		1	2	3	4	5	6	7	D	T
1	Task	1	0	0	1	1	1	0	1	2
2	Interaction	2	2	1	3	1	1	4	3	6
3	Brushing and Linking	1	2	0	2	1	1	1	0	3
4	Axis Relations	3	1	5	2	2	5	0	1	5
5	Legend Relations	0	3	0	0	0	6	1	1	7
6	Visual Components	8	6	7	6	7	7	4	0	9
7	Grouping Views	4	3	2	4	2	1	0	1	8
8	Show more information	4	5	0	0	4	3	1	1	8
9	Direction, Flow, and Order	4	1	1	0	0	1	0	0	5
10	Line Arrows	5	3	3	2	4	5	2	0	7
11	Strength	1	2	0	1	1	0	1	0	4
12	Clutter & Scalability	1	3	0	1	2	1	2	4	5
13	Interference with Views	5	1	2	0	1	0	1	1	6

Table 1. Overview of the concepts in the empirical data. The columns, scene 1 to 7 (see Figure 1), and D (debriefing), show number of participants for which the concept emerged. Column T shows the total number of participants who mentioned this concept. Note that T does not add columns.

in meetings held for approximately every 60 minutes of observed video. We identified 13 concepts to develop. Table 1 provides an overview of the countable results according to the concepts. In the second pass, notes taken during the first pass were used to revisit the video material in context of these concepts. During this pass, we used screenshots of the session sketches for each of the seven scenes in part A and B, and marked up relevant concepts in the sketches (Figure 2).

The concepts provided a basis in concrete empirical data, and were useful for identifying participants' statements and sketches. Thus, they allowed us to analyze the range of ideas that participants expressed. However, they are less useful from the perspectives of understanding existing and designing new representations of view relations. Thus, we examined each concept from these perspectives. From this, we mapped them to four themes which we describe next.

FINDINGS

In this section, we describe our findings that emerged from analyzing the empirical data. In all our descriptions, we use the empirical data to discuss the themes, and to exemplify how this might be considered in design. We refer to analysis concepts as ***[concept (number)]*** where number indicates the number of sessions we observed it in. We present our findings in terms of four themes: tasks, representations, interactions, and challenges.

Tasks

Participants directly considered the tasks that view relation representations might support ***[task (2)]***. For example, a participant stated that the importance of seeing particular types of relations depends on tasks and goals: *"It depends on what you want. If you want to follow a specific country, then this relation, in that case is more important. It completely depends on the context"*. Similarly, a participant stated: *"I don't imagine I would always want to see all the different things"* in relation to highlighted axes.

At the end of the sessions, participants were given the possibility to prioritize the different things they had considered. Here, participants either emphasized the possibility of seeing connections between data points (similar to brushing and linking), understanding filters and aggregation between views, or understanding axis scale relations. In addition to the direct considerations about tasks, other statements related to tasks. We describe these next, and indicate their relation to other concepts.

Construct overview [axis relations (5); legend relations (7); grouping views (8)]: A common consideration was to help people construct a mental overview of the views. For example, participants sketched relations between axes by highlighting them (Figure 2) to help construct an overview of what the different views showed and connected axes' ranges by lines. Likewise, four participants considered showing fewer views with increased information density by integrating two or more views, and four and three participants respectively connected and encircled views to show that they were similar in some way such as showing the same data, or using the same encoding.

Compare [axis relations (5); interaction (6)]: All nine participants considered the need to compare data between views. For example, five participants considered representations of axis scale relations (Figure 1, right). Additionally, one of these considered showing such relations when arranging views to line up (Figure 3, left). Additionally, one participant considered that moving and rotating bar charts would facilitate comparison (Figure 3, right).

Provenance and story-telling [task (2); direction, flow, & order (5); line arrows (7)]: Constructing an overview and comparing views are well-known visualization tasks. However, two participants also emphasized the potential of using representations of view relations to convey stories. For example, a participant used a hand gesture to indicate direction of views' connection, suggesting that the visualization was

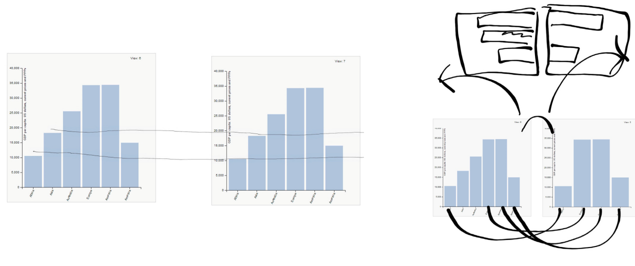

Figure 3. Participants considered different methods of interacting with views to show their relations.

"*trying to tell a story*". Likewise, in S6, a participant said that the lines that connected legends and data point between views made it possible to *"follow the data"*. Seven participants used arrows at least once. For example, participants sketched arrows between data bars and legends, to show relations between data in two views, which is less direct than linking data points to data points.

Within-view focus [clutter & scalability (5); interference with views (6)]: Finally, eight participants considered that sometimes, it is most important to be able to read a view without being distracted from relation representations. For example, in S3, a participant was concerned that highlighting the border and axes of view 4 to indicate an overview-plus-detail relation took focus from the data in the view.

Representations

Similarly to asking about task, we asked participants at the end of the session to consider the different representations they had sketched themselves, and those that we had used in the design scenes. Eight participants considered line connections to be useful, as long as they had limited interference with within-view visualizations. Likewise, five participants considered colors to be useful to indicate relations. They collectively remarked that colors should be used consistently within and between views. Two participants emphasized axis relations.

Visual Components [visual components (9); clutter & scalability (5); interference with views (6)]: All nine participants considered which components of a view to use in representing relations during the course of their session. They used all the different visual components of the view, the data and the meta-data (e.g., axes and legends). They connected data points and bars to other views' borders and sketched rectangles to group parts of data in one view to data in other views. For example, a participant considered showing an overview-plus-detail relation between two line charts, by connecting the line start and end points in the detail view (Figure 4), in contrast to using axes.

Reducing clutter and interference was a common reason for not connecting lines to data points. Six participants drew lines to legend items to reduce clutter. One participant stated: *"I want to connect this to the legend to reduce clutter*". Another participant said: "*I connect to the legend to not interfere with the lines*" (Figure 5).

Legends [legend relations (7)]: Relations between either entire legends or parts of legends were considered by seven participants. Five participants framed the legend in a circle or rectangle. A participant also considered connections between individual legend items and between individual and multiple items. Finally, in S6, where three views used identical legends, two participants considered separating the legend into a separate view. In contrast, a participant that didn't consider legend relations, suggested three interaction designs based partly on legends.

In designing S2, we used a spatial layout of legend items as a way of indicating legend relations. A participant, who noted this, stated that *"it is nice that the spatial position of legend items [match across views]"*. Contrarily, another participant initially thought that a legend item that was close to a data point was another data point.

Axes [axis relations (5)]: Five participants considered the usefulness of seeing relations between axes. For example, they all highlighted axes and used lines to connect axis ranges. In contrast to connections to legends, participants primarily argued for including axes in relation representations to convey something related to axes. Three participants commented that the different aspect ratios of views in S3 made it difficult to see that one showed a scaled version of the other. They argued that representing the scale relation resolved the issue. Likewise, a participant noted difficulties in spotting whether views that look alike are similar without visual support: *"The fact that you have three visualizations that all look very similar, while one of them has different axis labels. I feel like that should be highlighted in some way, right. 'Cause otherwise you're playing this game of spot the differences. And I don't think [people] are very good at that"*. The participant later noted about S7's relation representation *"at a glance you know that none of the axes are the same"*.

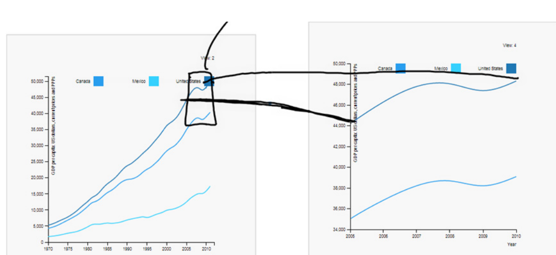

Figure 4. A participant connected line start and end points in a detail view with the same points in an overview in S3.

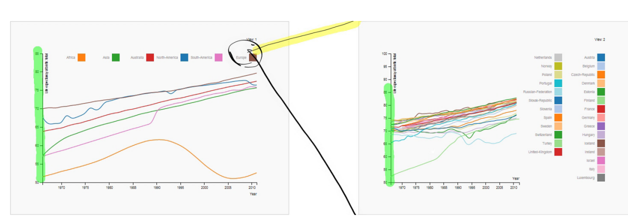

Figure 5. A participant sketched a relation representation between a legend and a view in S6.

Spatial position [direction, flow, and order (5); interaction (6); task (2)]: Position is an important visual variable, and can be used as a way to represent view relations. Participants considered using position to convey the direction of view relations, the flow of data between views, or the reading order of views. For example, a participant stated that some marks "*make you read the visualization in a specific order*". While arranging views, another stated that he was "*reading it left to right, top to bottom*". Likewise, a participant stated "*so this takes that data over there [pointing with both hands]*", and showed with hand gestures how views connected, suggesting that the visualization was "*trying to tell a story*".

Re-using encoding [show more information (8)]: Eight participants provided sketches that re-used a given views' encoding to show more information on either the relation representations, or directly on the representation already in the given view or another view. For example, six participants re-used the color of linked data points for the color of the link. This is a simple and well-known technique. Two participants used lines to connect data points in a scatterplot to data bars in a bar chart. Here, the lines' end points also encoded the specific data values on the vertical axis in the bar chart (Figure 1 and 2), thus re-using the bar chart's visual variable (vertical position), in the relation representation. Additionally, two other participants sketched variations of these. The participants both considered these as a way of summarizing the relevant factors from one view in another, and as a way to provide additional information.

Additional data in the relation representation [show more information (8)]: Two participants provided sketches that went beyond relating components shown in views. They put additional data into the relation representation, and thus used this to provide new information. For example, a participant merged lines from multiple legend items and connected these to data bars in other views. Doing so, line thickness represented the fraction of the data bar that was used in the relation representation to *"encode more information"* (Figure 6).

[Strength (4)]: Line thickness is a well-known example of conveying relative differences. However, aside from the example in Figure 6, no participants sketched relations that used line thickness. In contrast, two participants used a double line to signify equality and one participant sketched curvy lines to signify weak relations.

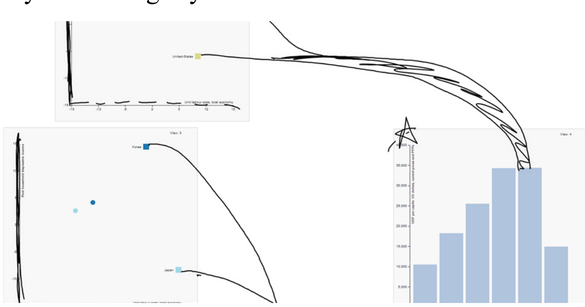

Figure 6. An example of participants' re-use of within-view representations in S2.

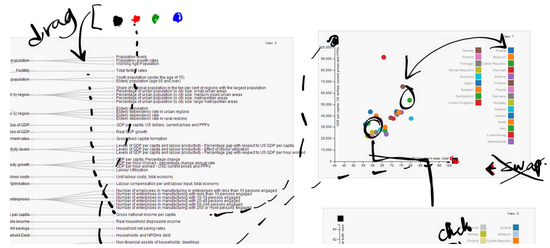

Figure 7. A participant imagined dragging colors across data variables in the dendrogram (left) and dropping the colors on the scatterplot axes (right).

While representing strength of relations was relatively rare, three participants talked about relation strength. For example, a participant stated *"This connection is not strong"* regarding a relation between two bar charts in S5 that showed the same data aggregated differently.

Interactions

Six participants considered potential interactions in regards to view relation representations ***[interaction (6)]***. Examples include rotating and aligning views spatially to compare them. Participants also considered other means of interaction. For example, in S1, a participant suggested that agency in interaction could make the relation between the bar charts in S1 easier to understand, and suggested that dragging a link between a data bar and a scatterplot that showed more details could update the scatterplot.

[brushing and linking (3)]: Three participants considered brushing and linking. For example, a participant imagined integrating the information shown in the five views in S1 into three views as a condensed representation, and then use brushing and linking to work with the data (Figure 1, left).

At the end of the sessions, we asked participants if there was anything they had neglected. One participant said that *"it was nice to be able to lay things out.* Other participants considered the emphasis they had given to interaction versus representation. One participant suggested that s/he hadn't *"thought too much about interaction"*. In contrast, another participant noted that they were primarily focusing on interaction.

Reconfiguring axes [interaction (6); axis relations (5)]: Two participants considered interactions to facilitate axis configuration in S7, where a dendrogram represented all data fields. One participant imagined dragging colors across data variables in the dendrogram and dropping the colors on the axes in the scatterplot views (Figure 7). The other participant imagined that the scatterplot axes could be reconfigured by first tapping the colored axis marks in the scatterplots. This would highlight the axis data variable in the dendrogram view. Afterwards, the highlighted data variable could be dragged to reconfigure the scatterplot axes.

Challenges

Eight participants considered challenges of relation representations, particularly during debriefing at the end of the sessions. Five participants specifically considered the main drawback of view relation representations in terms of scalability. Seven participants mentioned color scalability and risk of inconsistent use of color, while four participants considered overdrawing and clutter. Likewise, six participants also considered interference with views during the sessions.

Solutions to clutter [visual components (9); clutter & scalability (5); interference with views (6); interaction (6)]: Interestingly, participants also came up with suggestions for alleviating clutter such as: connecting to legends, connecting to view boundaries, connecting to axes, and other types of connections between non-data marks. Additionally, they suggested that connections might be shown based on interactions.

DISCUSSION AND FUTURE WORK

We described four themes that related to the thirteen concepts identified during coding. As described, many participants considered how to represent meta-visualization relations such as axis and legend relations. Such meta-view relation representations, which are rare in other visualization research, are more common in our study. More concretely, participants considered showing axis and legend relations using visual components. Additionally, some participants thoughtfully weighed the use of visual components in relation to reducing clutter. Some participants also considered interaction with regards to highlighting axis relations when arranging views spatially.

While previous work has looked at showing relations between visual *data* components (e.g., [8,17,29,36,42]), comparing techniques (e.g., [14,23]), and showing relations between visual *view* components (e.g., [9,21,32,38,41]), focus on visualizing relations between visual *meta-data* components seems to be missing. We see this as an interesting area for future work, which can build on the combination of our findings and existing work in this area (e.g., [12,17,29,36,38]). Related to the use of visual components, the concept of re-using within-view representations seems interesting. While this has not been the focus of previous work, examples of the approach exist in the literature. For example, colors within views have been used to color lines between views (e.g., [8,10]), and the position of data points have been used to anchor relation lines to axes and chart edges [36]. There has been less focus on meta-data components, in contrast to the focus on data and views in the initial descriptions of CMV systems (e.g., [26]). We think this concept deserves more focus and that it presents a promising direction for future research.

Likewise, previous literature also suggests creating relation representations that reduce interference. For example, by fixing lines that connect data points to axes and view borders [36], or by allowing people to control link visibility [29]. In our study, participants often expressed concern for clutter, and many suggested designs that used meta-data components as indirect ways to show relations to data. While one participant considered line thickness as a technique to represent the value of data, a few participants considered it as a way to indicate the strength of the relationship between views itself. We believe that relation strength might be an interesting aspect to consider in relation representations. Participants rarely considered showing the existence of brushing and linking relations. In fact, when we asked, participants found these relations to be less important and suggested they might simply be experienced through interacting with a visualization.

Additionally, interacting with relation representations could provide opportunities for working efficiently with data, but this is under-explored. It is necessary to understand the tradeoffs between interactive (e.g., brushing & linking) and static representations of relations. Also, while many participants were concerned with clutter, further studies will be needed to discover successful ways of dealing with clutter in combination with meta-visualizations. These seem to be under-explored challenges that deserve more attention and point to new and promising directions for future research.

Design scenes versus findings

Many of the identified concepts (1-6, and 9) consider the intention behind the design of meta-visualizations. Other concepts (7, 8, and 10-13) looked more at the specifics of the design. The five participants who critiqued the relation designs before sketching produced many novel designs. Rather than limiting their imagination; seeing possible methods seemed to lead to more ideas. A participant stated: *"Given these representations [within the views (e.g., scatterplots, bar charts, dendrogram trees)], I will only think based on the representations I have available"*. This means that our findings relate closely to the visualizations used in the design scenes, and clearly points to the potential for further research in this area, providing additional design ideas. However, many of our findings can be applied broadly. For example, the specific visualization types had little impact when participants considered how to show relations between data points in two views. Considering these options with other view representations might inspire other representations of view relations, and could be an interesting future research direction.

Our findings suggest that several tasks that relation representations may support are meta-level tasks (e.g., constructing overview and telling stories). Although participants considered relations that could help understand data, they focused on relations between views that help meta-analysis tasks such as navigation. There might be several reasons for this: First, participants that work with visualization and interaction design might focus less on data than domain experts. Second, we provided limited interaction possibilities for working directly with the data.

Additionally, relation representations might support more varied tasks. For example: Collins & Carpendale [8] intended to *"provide for a visualization space in which multiple data-related visualizations [could] be analyzed"*; Shneiderman & Aris [29] intended to *"enable users to gain meaningful high-level information from an overview as well as to ascertain the details of each node and link"*; and Tobiaz et al. [32] intended to *"support co-located collaboration among information workers who are making use of information visualizations in their analysis process"*. Ragan et al. [27] conducted a literature study of provenance types and purposes. While their focus on views was limited, they described a range of underlying potential view relation types, which might be represented. The results of our study might be interestingly applied to these types of potential view relations.

Implications for Design

Our study revealed the following important insights:

- **Axes and legends**: Representing meta-visualization relations through using axes and legends offers many new research opportunities.
- **Within-view encoding for relation representations**: Participants re-used a views' encoding to show more information using the relation representations, as part or on top of, the representation already in the view. While re-using color is common, participants came up with novel ideas based on position.
- **Additional data in the relation representation**: Previously, meta-visualization merely indicated existence of a relation. Participants enriched these with data.

These insights provide fertile grounds for new and promising view relations representations, which designers might consider in designing novel visualizations.

CONCLUSION

Considering the importance of being able to track complex, multi-person, multi-view analysis processes, we conducted an exploratory study on representing view relations that will help illuminate this research direction. In our study we asked visualization and interaction designers to critique and sketch possibilities for representing view relations, and identified many concepts relevant for studying and designing representations of view relations. Our approach generated rich data, providing a better understanding of the design space of techniques for representing view relations, and pointing to novel and promising ideas for representing view relations. Further, our findings can be useful in formulating new research questions about view relations.

ACKNOWLEDGMENTS

This research was supported in part by the Danish Council for Strategic Research grant 10-092316, NSERC (Natural Science and Engineering Research Council, Canada), AITF (Alberta Innovates Technology Futures) and SMART Technologies.

REFERENCES

1. Andrews, C., Endert, A., and North, C. Space to Think: Large High-resolution Displays for Sensemaking. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press (2010), 55–64.
2. Baldonado, M.Q.W., Woodruff, A., and Kuchinsky, A. Guidelines for Using Multiple Views in Information Visualization. *Proc. AVI*, ACM Press (2000), 110–119.
3. Bostock, M., Ogievetsky, V., and Heer, J. D^3 Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011) 2301–2309.
4. Card, S.K., Mackinlay, J., and Shneiderman, B. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann.
5. Chi, E.H. and Riedl, J.T. An operator interaction framework for visualization systems. *Proc INFOVIS*, IEEE (1998), 63–70.
6. Claessen, J.H. and Van Wijk, J.J. Flexible linked axes for multivariate data visualization. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011) 2310–2316.
7. Cockburn, A., Karlson, A., and Bederson, B.B. A Review of Overview+Detail, Zooming, and Focus+Context Interfaces. *ACM Comput. Surv.* 41, 1 (2009) 2:1–2:31.
8. Collins, C.M. and Carpendale, S. VisLink: Revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007) 1192–1199.
9. Dunne, C., Henry Riche, N., Lee, B., Metoyer, R., and Robertson, G. GraphTrail: Analyzing large multivariate, heterogeneous networks while supporting exploration history. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press (2012), 1663–1672.
10. van den Elzen, S. and van Wijk, J.J. Multivariate Network Exploration and Presentation: From Detail to Overview via Selections and Aggregations. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014) 2310–2319.
11. Gleicher, M., Albers, D., Walker, R., Jusufi, I., Hansen, C.D., and Roberts, J.C. Visual comparison for information visualization. *Information Visualization* 10, 4 (2011) 289–309.
12. Gratzl, S., Gehlenborg, N., Lex, A., Pfister, H., and Streit, M. Domino: Extracting, Comparing, and Manipulating Subsets Across Multiple Tabular Datasets. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014) 2023–2032.
13. Green, T.R.G. and Petre, M. Usability analysis of visual programming environments: a "cognitive dimensions" framework. *Journal of Visual Languages & Computing* 7, 2 (1996) 131–174.
14. Griffin, A.L. and Robinson, A.C. Comparing Color and Leader Line Highlighting Strategies in Coordinated View Geovisualizations. *IEEE Transactions on Visualization and Computer Graphics* 21, 3 (2015) 339–349.
15. Heer, J. and Agrawala, M. Design considerations for collaborative visual analytics. *Information Visualization* 7, 1 (2008) 49–62.

16. Heer, J., Mackinlay, J., Stolte, C., and Agrawala, M. Graphical histories for visualization: supporting analysis, communication, and evaluation. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008) 1189–96.
17. Henry, N., Fekete, J.-D., and McGuffin, M.J. NodeTrix: a hybrid visualization of social networks. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007) 1302–9.
18. Isenberg, P., Tang, A., and Carpendale, S. An exploratory study of visual information analysis. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press (2008), 1217–1226.
19. Isenberg, P., Zuk, T., Collins, C., and Carpendale, S. Grounded evaluation of information visualizations. *Proc. BELIV*, ACM Press (2008), 1.
20. Javed, W. and Elmqvist, N. Exploring the design space of composite visualization. *Proc. PacificVis*, IEEE (2012), 1–8.
21. Javed, W. and Elmqvist, N. ExPlates: Spatializing Interactive Analysis to Scaffold Visual Exploration. *Computer Graphics Forum* 32, 3pt4 (2013) 441–450.
22. Knudsen, S., Jakobsen, M.R., and Hornbæk, K. An exploratory study of how abundant display space may support data analysis. *Proc. NordiCHI*, ACM Press (2012), 558–567.
23. Liu, X. and Shen, H.-W. The Effects of Representation and Juxtaposition on Graphical Perception of Matrix Visualization. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press (2015), 269–278.
24. Munzner, T. A nested model for visualization design and validation. *IEEE transactions on visualization and computer graphics* 15, 6 (2009) 921–928.
25. North, C., Conklin, N., and Saini, V. Visualization schemas for flexible information visualization. *Proc. INFOVIS*, IEEE (2002), 15–22.
26. North, C. and Shneiderman, B. *A Taxonomy of Multiple Window Coordinations*. University of Maryland, College Park, MD.
27. Ragan, E.D., Endert, A., Sanyal, J., and Chen, J. Characterizing Provenance in Visualization and Data Analysis: An Organizational Framework of Provenance Types and Purposes. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016) 31–40.
28. Roberts, J.C. State of the art: Coordinated & multiple views in exploratory visualization. *Proc. CMV*, IEEE (2007), 61–71.
29. Shneiderman, B. and Aris, A. Network visualization by semantic substrates. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006) 733–740.
30. Shrinivasan, Y.B. and van Wijk, J.J. Supporting the analytical reasoning process in information visualization. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press (2008), 1237–46.
31. Strauss, A. and Corbin, J. *Basics of qualitative research*. SAGE Publications, Thousand Oaks, CA.
32. Tobiasz, M., Isenberg, P., and Carpendale, S. Lark: coordinating co-located collaboration with information visualization. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009) 1065–72.
33. Tohidi, M., Buxton, W., Baecker, R., and Sellen, A. Getting the Right Design and the Design Right: Testing Many Is Better Than One. *Proceedings of the 24th annual ACM conference on Human factors in computing systems (CHI '06)*, (2006), 1243–1252.
34. Tory, M. and Möller, T. Evaluating visualizations: do expert reviews work? *IEEE Computer Graphics and Applications* 25, 5 (2005) 8–11.
35. Tweedie, L. Characterizing interactive externalizations. *Proc. CHI*, ACM Press (1997), 375–382.
36. Viau, C. and McGuffin, M.J. ConnectedCharts: explicit visualization of relationships between data graphics. *Computer Graphics Forum* 31 (2012) 1285–1294.
37. Walny, J., Carpendale, S., Riche, N.H., Venolia, G., and Fawcett, P. Visual thinking in action: Visualizations as used on whiteboards. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011) 2508–2517.
38. Ware, C. and Lewis, M. The DragMag Image Magnifier. *Proc. CHI*, ACM Press (1995), 407–408.
39. Weaver, C. Visualizing coordination in situ. *Proc. INFOVIS*, IEEE (2005), 165–172.
40. Young, D. and Shneiderman, B. A Graphical Filter/Flow Representation of Boolean Queries: A Prototype Implementation and Evaluation. *Journal of the American Society for Information Science* 44, 6 (1993) 327.
41. Zgraggen, E., Drucker, S.M., Fisher, D., and DeLine, R. (s|qu)eries: Visual Regular Expressions for Querying and Exploring Event Sequences. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press (2015), 2683–2692.
42. Zhao, J., Collins, C.M., Chevalier, F., and Balakrishnan, R. Interactive exploration of implicit and explicit relations in faceted datasets. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013) 2080–2089.
43. Zhao, S., McGuffin, M.J., and Chignell, M.H. Elastic hierarchies: combining treemaps and node-link diagrams. *Proc. INFOVIS*, IEEE (2005), 57–64.