Journal Title
XX(X):1–13
©The Author(s) 2021
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/
SAGE

Effects of Screen-Responsive Visualization on Data Comprehension

Sriram Karthik Badam[1] and Niklas Elmqvist[2]

Abstract

Visualization interfaces designed for heterogeneous devices such as wall displays and mobile screens must be responsive to varying display dimensions, resolution, and interaction capabilities. In this paper, we report on two user studies of visual representations for large versus small displays. The goal of our experiments was to investigate differences between a large vertical display and a mobile hand-held display in terms of the data comprehension and the quality of resulting insights. To this end, we developed a visual interface with a coordinated multiple view layout for the large display and two alternative designs of the same interface—a space-saving boundary visualization layout and an overview layout—for the mobile condition. The first experiment was a controlled laboratory study designed to evaluate the effect of display size on the perception of changes in a visual representation, and yielded significant correctness differences even while completion time remained similar. The second evaluation was a qualitative study in a practical setting and showed that participants were able to easily associate and work with the responsive visualizations. Based on the results, we conclude the paper by providing new guidelines for screen-responsive visualization interfaces.

Keywords

Evaluation, responsive visualization, small displays, ubiquitous analytics.

Introduction

Visualization and visual analytics (VA) fields are increasingly taking advantage of novel device technologies such as multiple monitors[1,2], large displays,[3,4] tabletops,[5,6] and mobile devices[7–9] to go beyond traditional personal computers with mouse and keyboard interaction.[10,11] This movement towards *smart environments*[12] facilitates more complex and diverse sensemaking scenarios, where the analysts can ubiquitously interact with data through more than one device[13] and also efficiently collaborate with others during the sensemaking process.[14] However, a major challenge in developing visualization systems for these environments is the need for *responsive visualization*:[15] visual representations that can adapt to the display size,[15] similar to the idea of responsive web design for creating websites that can fit any display dimension. While there exist guidelines for designing visualizations with space-efficient layouts such as horizon graphs[16] and stacked charts,[17] the impact on sensemaking of using such representations on different displays—in terms of comprehension and insight generation—is largely unknown.

In this paper, we investigate the effect of display size on user comprehension and insight generation from responsive visualizations. We focus on coordinated multiple view (CMV) layouts[18,19] of visualizations because they offer a generic solution for dealing with multidimensional datasets and are popularly used in commercial visualization systems.[20,21] To this end, we developed a visualization system for interacting with large multidimensional datasets on a large 55-inch multi-touch display. This system consists of standard visual representations such as line charts, bar charts, scatterplots, and parallel coordinates to encode the data in the CMV layout. We also created two responsive versions of the system interface for small-screen devices such as smartphones and tablets. These small-screen interfaces overcome the lack of screen real estate for a CMV layout by (1) compressing the visualizations on the boundaries of the CMV layout and expanding them on demand (**boundary visualization**), and (2) adopting an overview+detail layout where the overview is a thumbnail version of the CMV layout and the detail is a visualization of interest from the multiple views (**overview visualization**).

To answer our research question, we conducted two complementary user studies using our visualization system to compare analyst performance—data comprehension and insight generation—on large vs. small (mobile) displays. Our primary contributions include:

- **Evaluation of Boundary vs. Overview:** Results from a controlled laboratory study comparing performance using the compressed boundary and overview interfaces for animated changes inside the visual representations. Our results show that visual changes on the boundary visualization take less time (or fewer replays) to interpret compared to the overview. The boundary visualization interface on the small display

[1]Apple Inc., Cupertino, CA, USA. Work done during doctoral studies at University of Maryland, College Park; [2]University of Maryland, College Park, MD, USA

Corresponding author:
Niklas Elmqvist, College of Information Studies, University of Maryland, College Park, MD 20742, USA.
Email: elm@umd.edu

Figure 1. Working with visualization systems on different displays. (A and B) Responsive visualization interfaces on a tablet and smartphone in front of the large display interface. (C and D) Close up of the two responsive visualization interfaces used in our studies—a overview+detail layout (C) and a compressed boundary visualization (D) interface.

also yields more accuracy than both the large display (surprisingly) and the overview on a small display.

- **Practical Usage:** Results from a qualitative study on the use of the three visualization methods (one large display and two small display interfaces) in a practical setting where the participants use them to verify complex insights about multidimensional data. We found that the space-efficient boundary representations reduced the number of interactions performed by the users compared to the overview representation without compromising accuracy.
- **Design Guidelines:** Practical guidelines for designing responsive visualizations, i.e. visual representations that scale to any display size, and responsive visual interfaces, that were derived from our findings from the two studies.

Background

Relevant to the idea of responsiveness, we first cover the research on scalable and space-efficient visualizations. We also discuss plastic user interfaces and visualization systems that go beyond a desktop to support devices of different display characteristics.

Space-Efficient Visualization

Increasing the amount of data that can be perceived from a visual representation is of crucial importance when dealing with large datasets. Much visualization research leverages space-efficient techniques based on space-filling (cf. treemaps [22] and sunburst diagrams [23]) or layout-optimized (cf. horizon charts [24]) visual encodings to visualize different forms of data. These techniques target a compact layout or representation that may trade readability of some data (at the cost of interaction) to increase the amount of information encoded within the display. For example, treemaps [22] encode hierarchical information structures by partitioning a given display space into rectangles.

Space-filling versions also exist for traditional visual representations. Keim et al. [25] extended bar charts to a space-filling representation, *pixel bar charts*, in which each pixel is used to represent detailed information about the data. Hao et al. [26] introduced pixel matrix displays to generalize the idea of capturing information at each pixel in a visualization. [27]

Heer et al. [16] investigated the advantages of using horizon graphs that are created by dividing and layering filled line charts, for time-series data. They found the layering approach of this space-efficient representation to be beneficial over line charts as display space decreases. Finally, other space-efficient encodings for line charts to increase the number of attributes captured by a single view—such as stacked area charts as well as composite visualizations [28]—can also effectively scale visualizations to display characteristics.

Beyond designing space-efficient visual encodings, the challenges of limited display space for a visualization interface can be mitigated through interaction such as zoom and pan or distortion (cf. Mélange [29]). However, zooming and panning causes loss of overview and distortion techniques [29] only work well for continuous visual spaces—not for a dashboard with multiple views. Alternatively, additional viewports can also be created: one for overview and one for detail to fit multiple views to smaller displays. [30] Overview+detail layouts are popular for efficiently utilizing a limited display space, and are shown to be more efficient than pan and zoom. [31] In geospatial visualization, offscreen representations [32,33] have also been explored to tackle the problem of showing a large visual space in a small screen real estate.

Plastic Interfaces and Responsive Visualization

The challenge of developing interactive systems that run on any physical device with varying display size and input was identified more than a decade ago. Thevenin and Coutaz formalized the idea of *plastic user interfaces* [34] that can adapt to a target device modality by means of adapting the rendering techniques and behavior of the system. Calvary et al. [35] created a reference framework for interfaces supporting multiple targets. To actually develop interfaces that can adapt to any device, it is important to understand the tradeoffs of using varying device modalities.

Tan et al. [36] studied a large projection display against a desktop monitor to quantify its benefits to an individual user. When using them at the same visual angle, the authors found that the users performed better on spatial orientation tasks on the large display, as it provided a greater sense of presence. Liu et al. [37] compared physical navigation on a ultra-high-resolution wall display against virtual navigation on a desktop for a data classification task (where items on the screen are organized into containers). They found that the wall display was more effective for such tasks with higher difficulty levels. Jakobsen and Hornbæk[38] studied the relation between display size and usability of visualizations. However, they did not find any significant benefit in using a large display for their scenarios. They focused on geographical maps and associated tasks in using them, and compared three different techniques (overview+detail, focus+context, and zooming) on three display sizes. These works implicitly offer important guidelines towards designing experiments that compare devices of different modalities.

While the term “responsive visualization” is relatively new, the idea of adapting visualizations to displays has existed for a long time. Yost et al.[39] explored physically adaptive visualizations for taking advantage of the human perceptual abilities by say using light colors that blend with the background and can only be seen when close to a display. This form of responsiveness based on spatial attributes of the user was also explored by Isenberg et al.[4]. To apply this practice to the physical attributes of the devices, Leclaire and Tabard[15] created a web framework for responsive visualization called R3S.js. Badam et al.[40] explored the idea of plastic visual representation when transferring visualizations from a large to small display in multi-device environments by manipulating the visualization pipeline. In this paper, we focus on developing not just responsive visualizations but also visualization interfaces that are holistically responsive in nature.

Most recently, Hoffswell et al.[41] studied responsive visualization by reviewing 231 responsive news visualizations and interviewing five journalists about designing such visualizations. They use these findings to derive four design guidelines and implement them in a new responsive visualization design system that facilitates flexible, cross-device design workflows. In contrast, while we propose a novel design framework, our work makes no technical contribution beyond our user testbed, and our evaluation is more quantitative in nature. Furthermore, our focus is more on large displays, whereas Hoffswell et al. is geared more towards mobile and personal displays.

Visualization Beyond Desktop Displays

Large displays have received much attention in recent times. They can support thinking by the added space[1] and also support collaboration between analysts. To take advantage of the physical space in front of the large displays, recent works[42,43] proposed utilizing position, orientation, movement, distance, and identity of the users to interact with visualizations. Beyond this, cross-device visualization frameworks[7] for connecting large displays with portable devices have also been presented to support advanced collaboration scenarios. Langner et al.[9] investigated how multiple small-screen personal devices (tablets and smartphones) can be combined for visual exploration. Chen[44] supported visualization of large time series on a smartwatch through aggregated statistics placed on the borders along with a detail visualization in the center. Horak et al.[45] explored the combination of a large display and a smartwatch for visual data exploration and found benefits of this setup compared to a large display-only environment.

Among recent research on using large high-resolution displays for sensemaking, Reda et al.[46] studied two display modalities—a wall-sized display and a high resolution cylindrical display environment. They found that increasing the display size and resolution can improve insight quality and breadth.

Evaluation Overview

The related work highlights techniques to create visualizations and visual interfaces that convey data in a limited screen space. Based on this, there are two main ways to achieve responsiveness to a given screen for a visualization interface:

- *Space-efficient visual encoding*: to compactly convey given data by packing the visualization with as much information as possible (cf. horizon charts[16] or pixel bar charts[25]).
- *Space-efficient interface layout*: using multi-scale viewports—overview+detail[47]—or multi-focus techniques to place a complex interface in a small screen viewport.

In fact, time-series visualization interfaces developed by Chen[44] for smartwatches showcase these two methods in action—through border and overview+detail interfaces—to achieve responsiveness to very small displays.

We were interested in evaluating such responsive interfaces to understand their affordances in visual sensemaking. For this purpose, we considered a standard visualization interface for a large display and defined two small screen versions based on the above techniques. The large display interface represents a *classical visualization interface* for multidimensional data, with multiple visualizations coordinated through user interaction (CMV[18]) to support data exploration. Inspired by Chen,[44] we developed a *boundary visualization interface* for a small screen that compresses the visualizations on the boundary of the classical interface to fit to a smaller screen (using space-efficient visual encodings). We also created a *overview visualization interface* that uses a bird’s-eye view to show the entire interface with a detail representation to show the current visualization of interest (a space-efficient overview+detail layout). We evaluated the differences between the two responsive visualization interfaces on a small screen (Figure 1), compared to the original interface on a large display as baseline, in terms of data comprehension and insight generation through two user studies (motivated by the InfoVis evaluation pattern of complementary studies):[48]

1. *Controlled study:* First, a controlled study of differences between boundary and overview interfaces compared to the classical, in terms of quantitative measures—time and accuracy—for comprehending data in coordinated multiple views.
2. *Study of practical use:* The second evaluation was a qualitative study of the differences in a practical setting where insights needed to developed through interaction with the interfaces.

For the controlled study, the type of visualizations and complexity of the interface is controlled to extract reasonable findings from the quantitative measures. For the study on practical use, the visualization interfaces on both large and small displays were more supportive of open-ended visual exploration of the dataset through interaction.

In other words, the studies are complementary because Study 1 isolates low-level tasks in an internally valid and highly controlled study focused on quantitative metrics, whereas Study 2 engages participants in a more high-level and open-ended qualitative study. We then combine our findings in the Discussion section at the end of this paper.

User Study: Comprehension from CMV

This elaborate design space can be explored to design techniques for responsiveness by choosing alternative layouts, encodings, and content in visual interfaces based on the display characteristics. In fact, previous work on time-series visualization interfaces[44] showcases techniques grounded in this design space to adapt to a smartwatch, including: (1) border visualizations utilizing encodings that fit the visualizations into the border of the small display (adapting visual mapping), and (2) overview+detail layout that shows the entire interface in a small bird's-eye view with specific content of interest as a detail view (adapting layout).

We were interested in evaluating such responsive interfaces to understand their affordances in visual sensemaking. Therefore, the focus of this user study was to observe quantitative differences in time and accuracy between a classical visualization interface and two small-screen versions—boundary and overview interfaces.

Dataset and Interfaces

We chose to study three interfaces: (1) a classical visualization interface (**CV**) showing a multidimensional dataset, (2) a boundary visualization interface (**BV**) developed for a small screen where the boundary views of the classical interface are compressed, and (3) an overview interface (**OV**) for a small screen that uses an overview+detail layout. Figure 2 shows these three interfaces applied to a motion pictures (movies) dataset for this study. The movies dataset contains 3,201 movie records with 15 variables including information such as gross earnings, budget, genre, IMDB rating, Rotten Tomato rating, and date of release. We picked this dataset as it may be accessible and familiar to a general audience due to its real-world interest.

- **CV** *Classical Interface (Large Display):* The classical visualization interface consists of a CMV layout with five visualizations of a movies dataset. The center visualization is a scatterplot of movies organized by their budget (x-axis) and gross (y-axis), and colored based on their ratings. The visualization surrounding the center (focus) contain statistical information about these variables (average budget and gross) over the years and for different genres visualized with line charts (filled) and bar charts respectively. All visualizations take equal screen space. This interface represents a typical multi-view visual dashboard.
- **BV** *Boundary Interface (Small Display):* The boundary visualization interface is created for a small display by compressing the peripheral views in the CMV layout using space-efficient visual encodings. The center scatterplot visualization takes most of the space (80%). The rest of the screen is equally divided among the four surrounding visualizations of average budget and gross. The area chart turns into a horizon chart with three layers (cf. guidelines on sizing the horizon[16]) and the barchart adapts into a space-filling version with color intensity to capture value instead of bar size (cf. border views[44]).
- **OV** *Overview Interface (Small Display):* The overview visualization interface is developed using an overview+detail layout on the classical interface. An overview+detail layout helps show multiple scales of information on a given limited display space. In this case, it shows the entire classical interface as an overview by making the viewport smaller, with the detail view showing a particular visualization.

Participants

We recruited 13 paid participants (5 female, 8 male) from the general student population of our university. The participants were between 18 and 45 years of age. All participants self-reported as proficient computer users. Furthermore, 8 participants had previously used visualization as a means of data analysis; however, it was mostly limited to charting in Excel, MATLAB, Mathematica, R, and SPSS. Only two participants had experience working with interactive visualization tools such as NodeXL, Gephi, and Tableau.

Apparatus (Devices)

The participants used a 55-inch display—Microsoft Perceptive Pixel—as the large display, and a 8.9-inch Google Nexus Tablet as the small display. The large display has a resolution of 1920×1080 pixels, while the tablet is 2048×1536 pixels (but with an effective CSS resolution of 1024×768 pixels). The interfaces were developed using web technologies—HTML, CSS, and JS—and the D3 framework (https://d3js.org/).

Tasks

To measure the low-level costs of responsiveness through boundary and overview transformations in visual sensemaking, we chose to study typical visual analysis tasks such as value and trend identification, and comparison. The major feature of a CMV interface is the **coordination** across views: user interaction on one visualization leading to visual changes in others. Therefore, it is an important aspect whose efficiency should be preserved across interfaces as they are made responsive to smaller screens. For this purpose, we created tasks for the controlled study in which the participants had to verify statements about the data values and trends in the visualizations in each interface.

We chose to simulate data selections on a focus view to see changes on the surrounding context of the CMV layout. The animations used to convey the changes in the CMV interfaces had a one-second duration (suggested by Robertson et al.[49]).

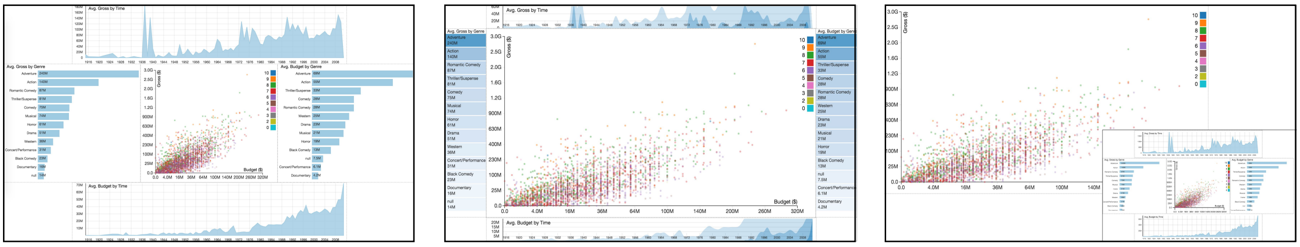

(a) Classical Visualization Interface (used on a large display)

(b) Boundary Visualization Interface (used on a small display)

(c) Overview Visualization Interface (used on a small display)

Figure 2. (Left to right) Classical, boundary, and overview visualization interfaces used in the controlled study with the movies dataset. Boundary and overview interfaces are developed for a small screen (a tablet).

The animations were also staged—the axes transitioned first followed by changes to the visual elements (or marks) in the visualizations. Here, marks correspond to rectangles in the bar chart, circles in the scatterplot, and areas in the filled line chart and the horizon chart. Staged animations are recommended by past research.[50] The statements provided to the participants are of two types: (1) perceiving changes in values in a chart in the interface (whether they increased or decreased over the animation created through simulated selections), and (2) perceiving retainment or reversal of a trend in a chart (whether the trend stayed the same or reversed at some location).

To come up with the statements, we created a list of observations with a goal of answering high-level questions about the data such as "how are the top rated movies different than others?," "are directors good at only specific genres?," and "are there differences in the typical gross from genres over time?" After this, we extracted the selection interactions that led to these observations, to simulate them in the actual experiment. The observations from answering the high-level questions were modified to create the statements for the study that can be either true or false. These statements were tested in two pilot studies with (1) a visualization expert to verify the correctness and complexity of the observations behind the statements, and (2) a novice student to verify if the statements are comprehensible. The statements were revised based on their feedback.

For controlling the study, statements had similar complexity—requiring the user to look at animated visual changes for two or more items within a visualization. Examples statements used in the participant tasks include,

1. The average budget values for thriller/suspense and musical genres are higher than default.
2. The average gross for 1972 and 2009 are lower than default.
3. The trend in average budget values between 2000 and 2008 is similar to default.
4. The trend in average gross values for Romantic Comedy and Horror genres is opposite to default.

Here, "default" refers to the state of the dataset before animation, and "trend" refers to the change trajectory. Note that these definitions were explained to the participants and verified during training.

Experiment Design and Procedure

We used a within-subjects design with the participants using all visual interfaces to verify statements about the data. Each interface was assigned a random set of statements, and the interface order was counterbalanced. This ensured that there was no effect of statement and interface order. There were eight statements (task repetitions) for each interface: four about time (line charts) and four about genres (bar charts). The statement type (S)—value comparison or trend comparison—and the data type (D)—about time (line chart) or genre (bar)—are also factors in our experiment.

Each study session started with the participant reading and signing a consent form, as well as completing a demographic survey. Following this, they went through a training procedure on an assigned interface, including how to interpret the line charts, bar charts, and the statements. Their accuracy was tested during the training to make sure they understand the statements and the interface. Each task required participants to, (1) understand a statement and identify which visualization to look at, (2) click a *play* button to simulate an interaction leading to animated changes in the visualizations (after a 3-sec countdown), and (3) determine if the statement is true or not and submit the answer. The animations were shown after a 3-sec countdown to ensure attention from the participants. They were also asked to verbally explain their reasoning for the answers before moving on to the next task. The participants were allowed to replay the simulated interactions for each statement any number of times. Following this, they click a *next* button to move to the next statement. Following the tasks on one interface, they completed a Likert-scale survey rating the efficiency, ease of use, and enjoyability of the interface. They then moved on to other conditions to follow the same process. Each session lasted for 50 minutes or less.

Measures

We recorded the accuracy of each participant's assessments (true/false) as well as the number of interaction replays performed to reach the assessment. We also recorded the time taken for completing each task.

Hypothesis

Based on our design, we formulate the following hypothesis:

H1 The classical visualization interface will be more accurate than boundary and overview visualization interfaces as the large display contains more visual space, which can help track the visual changes in charts.

H2 Classical interface will be faster than boundary and overview interfaces for the same reason as above.

H3 Boundary will be faster and more accurate than overview as it uses the display space more efficiently with space-efficient encodings.

Analysis and Results.

Here, we discuss accuracy/correctness and completion time (through number of replays and time taken) for the tasks by reporting on the results from statistical analyses. Figure 3 visualizes these results by calculating point estimates and 95% confidence intervals (CI) based on 1,000 percentile bootstrap replicates. Considering recent concerns with null-hypotheses testing[51] and APA recommendations[52] regarding p-value statistics, our analysis combines the best of both worlds by reporting p-values as well as confidence intervals from bootstrapping (see Figure 3).

Accuracy. Table 1 summarizes the main effects and interactions on accuracy using logistic regression (all assumptions valid).

Table 1. Effects of factors on accuracy (logistic regression).

Tasks	Factors	df, den	F	p
All	Display Interface (I)	2, 293.6	4.71	.009**
	Data Type (D)	1, 288.5	.09	.76
	Statement Type (S)	1, 299.3	1.67	.19
	I * D	2, 290.6	.51	.59
	I * S	2, 295.4	1.05	.35
	D * S	1, 289.1	1.39	.23
	I * D * S	2, 297.8	.53	.58

* = $p \leq 0.05$, ** = $p \leq 0.01$.

Post-hoc analysis with Tukey HSD revealed significant differences between boundary visualization and overview visualization interfaces ($p = .009$), and the classical interface and boundary visualization on tablet ($p = .017$). Boundary visualization interface (total correct answers = 91/104) on the tablet was more accurate than the overview interface (total correct answers = 74/104) and the classical interface (total correct answers = 76/104). Figure 3-left showcases these effects with the Boundary interface outperforming the classical and overview interfaces—with the effect being stronger between Overview and Boundary. There were no significant differences between the classical and overview interfaces. This rejects our first hypothesis (**H1 rejected**) since boundary visualization interface was significantly more accurate.

Number of Replays. As described above, time measurement is influenced by the time taken to interpret the statement in each task rather than just comprehension of the visualization and animation. In fact, the participants took similar amount of time to answer the statements in each condition as seen Figure 3-right. We therefore focus on the number of interaction replays (an integer) and analyzed it using a generalized linear regression model (reported in detail in Table 2). Results from an RM-ANOVA analysis were similar.

Table 2. Effects of factors on replays (gen. linear regression).

Tasks	Factors	df, den	F	p
All	Display Interface (I)	2, 24	3.50	.046*
	Data Type (D)	1, 12	.04	.84
	Statement Type (S)	1, 12	.00	.97
	I * D	2, 24	.14	.86
	I * S	2, 24	.35	.70
	D * S	1, 12	.26	.61
	I * D * S	2, 24	2.27	.12

* = $p \leq 0.05$, ** = $p \leq 0.01$.

Post-hoc analysis with Tukey HSD revealed significant differences between boundary and overview visualization interfaces ($p = .038$). There were no significant differences for the other two combinations. Overall, boundary visualization interface (mean = 1.39, s.d. = 0.86) had significantly less interaction replays than overview visualization interface (mean = 1.68, s.d. = 0.95). This can also be confirmed from Figure 3-middle where the size of non-overlapping region between Boundary and Overview interfaces is more than 80% of the bands (showcasing a large effect)[53]).

Based on this, our second hypothesis is rejected (**H2 rejected**). However, the third hypothesis is confirmed since boundary visualization (BV) was faster based on interaction replays and more accurate than overview visualization on tablet (**H3 confirmed**).

Subjective Preferences

After each session, the participants rated the techniques on three interfaces: efficiency, ease of use, and enjoyability, on a Likert scale ranging from 1 (e.g., strongly disagree) to 5 (e.g., strongly agree). Figure 4 showcases the differences between the three conditions as perceived by the participants. Across all the scales, both classical and boundary interface are perceived to be better than the overview interface. This especially reflects in the number of participants strongly disagreeing on all three scales for the Overview condition on the tablet device. On the other hand, the differences between classical and boundary interfaces are little. Therefore, to better explain our results, we report the participant feedback in the following section.

An ANOVA conducted for each of these scales showed that that the effect of interface was significant for all three at the .05 level. Post-hoc analyses revealed that, (1) boundary (mean = 4.076, s.d. = 0.759) was more efficient ($p = .043$) than the overview interface (mean = 3.23, s.d. = 1.363), (2) classical interface (mean = 4.153, s.d. = 1.068) was easier ($p = .011$) to use than overview interface (mean = 2.923, s.d. = 1.255), and (3) classical interface (mean = 4.076, s.d. = 1.037) was more enjoyable ($p = .029$) than overview (mean = 3.153, s.d. = 1.405).

Participant Feedback

The differences in performance can be explained through participant comments during the tasks as well as their preferences.

The boundary interface was better than classical (for some). Participants (8/13) expressed that the boundary

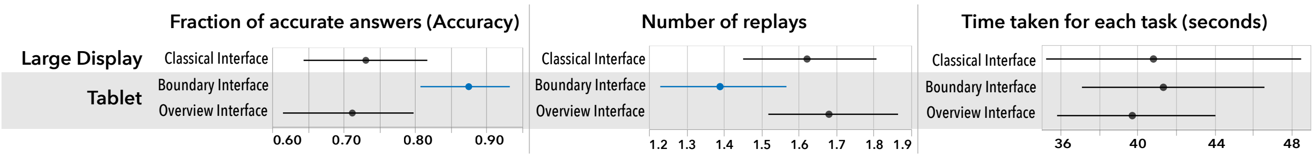

Figure 3. (Left to right) Point estimates and 95% confidence intervals based on 1,000 bootstrap replicates for the measures from the user study. As observed, the boundary visualization interface (in blue) on the tablet outperformed other interfaces in terms of accuracy and number of replays (fewer in Boundary). However, the effect is smaller for the latter measure and absent for the time measure.

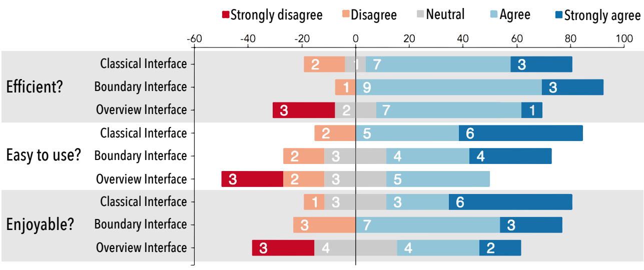

Figure 4. Likert scale ratings provided by the participants for efficiency, ease of use, and enjoyability of the three conditions.

interface suited the form factor of the small display used in this user study. This is mainly for reasons, (1) the space-efficient encodings, although require some training, convey information effectively in the small space, and (2) the convenience of sitting and viewing the tablet device from a reasonable viewing distance (compared to the large display). Participants (4/13) commented that the classical interface on the large display required them to track a larger physical space on the screen to verify the statements in the tasks. This could have also contributed lower accuracy on the classical interface compared to the boundary (BV), which is a surprising result. For instance, P3 commented, "*it is hard to look for the information that I want [on the classical interface] on large display. On boundary interface it is much easier to locate*." Boundary interface was also seen to be more suited to the small screen than the overview interface. P5 said "the visualizations on the boundary are much easier to follow than the overview." In the post questionnaire, majority of the participants (12/13) therefore preferred boundary transformation—i.e., adapting the visual encodings to be more space efficient to achieve responsiveness—compared to the Overview. This is irrespective of the visualization (line or bar chart), and this aspect also reflects in the lack of an effect of the visualization type in our results. Four participants preferred the boundary interface on the tablet over the classical interface on the large display, citing the efficient nature of the boundary interface on a familiar personal device that they use everyday.

The overview interface was familiar, but not effective. Participants felt that the overview interface provided a more familiar set of visualizations, since some of them (7/13) were not familiar with horizon charts. However, the added cost in tracking changes closely in a small overview overcame the convenience of reading familiar visual representations (line and bar charts). Five participants explicitly cited the small size of the overview visualization for giving it a poor rating on all scales. Therefore, overview transformation—a transformation of the layout of the interface to achieve responsiveness without changing the visual encodings—is not ideal when it comes to low-level sensemaking tasks.

Qualitative Evaluation: Practical Usage

The focus of the controlled study was on observation of controlled visual changes in a visualization interface. However, the controlled study was focused on low-level tasks in a specific setting. To explore these phenomena in a more ecologically valid setting, we also conducted a study for a more open-ended use of the interfaces where participants were asked to complete high-level tasks. This would yield more organic and realistic usage of our techniques. Since our focus here is more on understanding rather than comparative performance, we opted to make this a qualitative study.

Dataset and Interfaces

Just like Study 1, the qualitative evaluation also involved three interfaces (Figure 5). To expand the scope of visual exploration, we used an airline flights dataset for this study. This dataset contains information of passengers, seats, and flights between cities in the US from 1990 to 2009, as well as population and distance between cities.*

CV *Classical Interface:* This interface consists of a dashboard-style CMV layout to view the attributes in the Flights dataset using multiple views: (1) the aggregated attribute values (e.g., total number of flights) by cities through bar charts on the corners, (2) temporal statistics (e.g., passengers over years) using filled line charts on top and bottom, and (3) all cities in United States on a geographical map at the center of the interface with links between them signifying flight connections. The interface supports touch-based querying of the visualizations (Figure 6) including object selection by tap on bar charts, range selection through drag and drop on line charts and parallel coordinates[54], and freeform shapes on the scatterplots and maps. Upon interaction, the system retrieves the selected data and updates the rest of the views.

BV *Boundary Interface:* In this interface, the views on the boundary of the classical interface are transformed to be space-efficient: (1) space-filling rectangles with color saturation capturing value instead of bar chart, (2) using horizon charts instead of line charts, and

*Flights dataset: http://www.transtats.bts.gov/

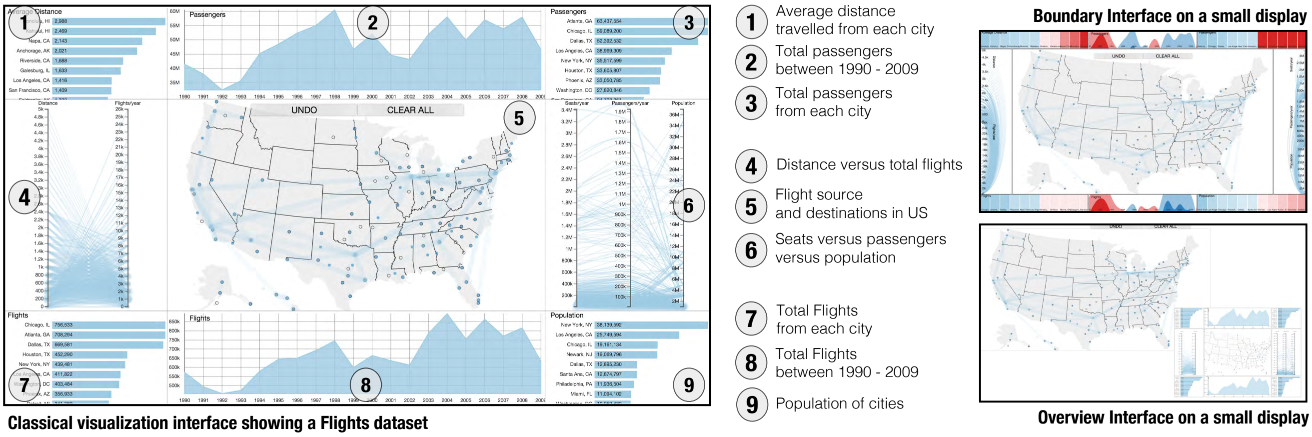

Figure 5. (Left) Classical interface with multiple coordinated views for a Flights dataset. (Right) Responsive versions of the classical interface interface developed for small displays—the boundary visualization interface (BV) expands a focus visualization by compressing the views on the boundaries and the overview visualization interface (OV) creates a overview+detail layout to work with the multiple views. The focus (or detail) visualization can be changed by tapping a view on the boundary or the overview map.

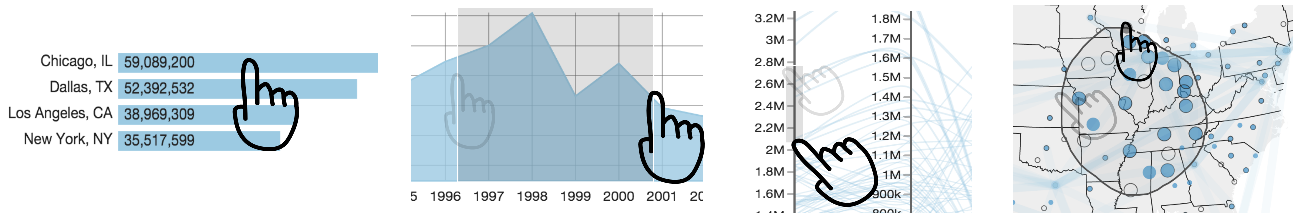

Figure 6. (Left to right) Touch interactions for selection on different visualizations to support visual exploration.

(3) folding parallel coordinates into arc diagrams (all dimensions are aligned). The boundary visualization interface follows these transformations and can be seen in Figure 5. From an interaction perspective, the main view in the center can now be changed by clicking on one of the boundaries to open an expanded visualization (*focus switching*), such as transitioning between a line chart and a horizon chart (Figure 7).

OV *Overview Interface:* Similar to the controlled study, this interface was designed for a small display and uses an overview+detail layout to explore multiple views. The items in the overview can be selected to switch the detail view (focus switching).

Participants

We recruited eight paid graduate students (2 female, 6 male) experienced in visual data analysis from HCI and related labs in our university. Their experience corresponds to: (1) reporting results in their research papers, (2) data analysis as part of their graduate research, or (3) developing visualization tools.

Method and Tasks

We wanted participants to interactively explore different parts of the dataset. More specifically, we wanted them to drill down by more than one dataset attribute to explore the data. This can provide an understanding of the complexity of observations during visual sensemaking made with these interfaces. Our definition of observation complexity is the corresponding number of attributes interacted with (cf. Reda et al.[46]).

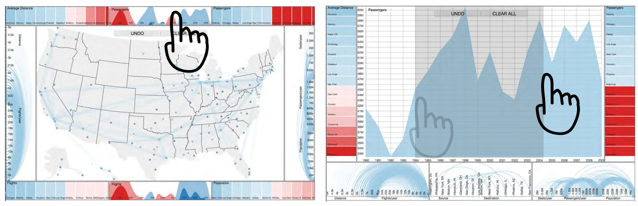

Figure 7. (Left) To interact with a visualization on the boundary (or in the overview) interface on a small display, the users change the focus by tapping the boundary view (or the overview). (Right) Then they can perform the interactions listed in Figure 6.

To guide the participants' visual exploration process, we provided them with a list of statements of three complexity levels (in contrast to the controlled study where the statements had the same complexity). Examples of these statements include:

- **Level 1:** 2008 recorded the highest number of passengers from cities on the West Coast of the United States.
- **Level 2:** Since 2006, San Francisco had more flights serving long distances (> 2500 miles) than New York.
- **Level 3:** Among East Coast cities, New York had the most number of short distance flights (< 1000 miles) during 2001–2003.

While we asked participants to collect insights regarding these statements throughout the study, we were less interested in the veracity of their responses (some of the statements had no single correct answer) and more concerned with their workflows using the interfaces. Think-aloud utterances, which included insights collected in response to the above statements, were recorded for later analysis.

Procedure

The procedure for this study was similar to the controlled study. The participants trained on an assigned interface by completing a tutorial, experimenting with the interactions, and verifying the sample statements. Following this they did six tasks on the assigned interface (two for each statement level). They then went through other conditions following the same procedure. They were asked to think aloud while interacting with the interfaces and verifying the statements. Their interactions were recorded and they were requested to provide feedback on their experience with the interfaces after the session. For boundary and overview interfaces, participants used a 5.7 inch Samsung Galaxy Note 4 smartphone with a 2560×1440 pixel resolution. The order of the interfaces was counterbalanced across participants. Each session lasted for about 60 minutes.

Observations

We analyzed the participant interactions, their think-aloud utterances, their collected insights, and their general feedback using informal single-pass thematic analysis. Below we report on the main themes in this data.

The overview interface required more interactions for similar outcomes. We observed differences in number of focus switches between the two small screen interfaces (Table 3). Fewer focus switches were observed for the boundary interface than the overview interface. This means that the participants directly interpreted the boundary visualizations to develop insights rather than expanding the views (e.g., Figure 7). In an ideal scenario, the number of focus switches on the small display interfaces should be the same as the complexity level of the observation being made (i.e., number of selections required). However, all participants found it hard to recall the location of their desired visualization on the overview interface. This led to more interactions to develop insights of similar complexity on Overview. In case of Boundary, only two participants (P3 and P6) switched the focus multiple times to figure the visualization they want, instead of reading the labels. All participants (except P7) felt that boundary visualization interface reduced the interactions compared to overview. Participant P5 expressed the reason as, "*Being able to read [the space-efficient encodings] and determine responses from the boundary tiles was a strong benefit.*"

Table 3. Number of focus switches on small display interfaces.

Level	Boundary		Overview	
	mean	s.d.	mean	s.d.
1	1.1	0.3	2.4	2.8
2	2.3	0.8	3.8	1.9
3	3.2	1.0	6.5	3.5

Large display required additional physical effort (for some participants). When developing observations that require more than one interaction (with more than one view in the interface), participants (5/8) felt the large display interface required them to focus their attention on a large screen space in their field of view to track changes when interacting with a view. Some (4/8) would even physically navigate in front of the display naturally when interacting with the classical interface. This gave them a feeling that they needed to perform more physical effort to achieve the same results and many participants commented on this aspect. Note that we did not control the distance of the participant from the large display, it was up to them to find a comfortable location. Participant P3 (used large display after other devices), for instance, mentioned that, "standing close to the large display to touch a visualization and moving to reach other views makes me work more than sitting in front of a table to use the boundary tool [on the smartphone]."

Space-efficient encodings did not compromise insights. Overall, the boundary representations in BV created through space-efficient encodings did not deter the participants from developing insights. Even though the participants performed fewer focus switches on the boundary interface, they still answered the tasks requiring more interactions with similar accuracy levels compared to overview interface on smartphone and the classical interface on the large display. Three participants performed worse on the overview interface, while the accuracy levels in the observations remained similar across the three interfaces for other participants. Mistakes made by these three participants were due to interpreting the wrong chart or failing to interact on the right attributes. Few participants (3/8) felt discomfort reading the space-efficient versions of the parallel coordinates (i.e., arc diagrams) when there are more than three coordinates in the diagram. However, this was not strongly reflected in the accuracy of their observations.

Discussion

In this section, we describe the rationale behind our study along with the limitations. We also present guidelines to explore the space of responsiveness in visual interfaces.

Study Design Rationale

Our study required making several design decisions since this was our first step in solidifying the space of responsiveness. We formulated the design rationale by considering the past research on perception and visualization design.[16,17,44] Here we discuss these rationale in detail.

- *Display sizes:* We used the 55-inch device as a large display due to its natural ability to show multiple views effectively, whereas the tablet acts as a small screen. These device types are also common in modern sensemaking processes.[11]
- *Focus on specific transformations:* We focus on two transformations—boundary and overview—to achieve responsiveness. This is because they characterize two popular methods (cf. Chen[44]) to maintain space efficiency in a visualization, by compressing the visual

encodings (boundary visualization) or by adapting the interface to have multiple viewports.

- *Labels, axes, and animations on visualizations:* We chose to show the labels and axes on each visual representation, since we believe that it is fair to assume that axis and labels would be present in a responsive visualization on any device and they would also be scaled appropriately to the display. The animations caused by the user interactions in the CMV interface had a one-second duration, during which visual elements transitioned in a cascading/staged fashion—first axes and then the visual marks. Staged transitions are recommended by past research on data animations from Heer and Robertson.[50]

- *Classical interface on the small display:* Scaling down the classical interface to the small display is not naturally responsive as it has severe limitations based on the small display and number of visualizations. But we still considered this condition in our pilot for the controlled study. We conducted two pilot studies with a visualization expert and a novice student to test the three conditions, along with the classical dashboard interface on the small screen. However, we found that the latter created issues for the participants when interpreting the features in the visualizations. Therefore, we did not include that condition.

- *Replays:* The controlled experiment (Study 1) provided the ability to replay the animated changes multiple times. Participants therefore replayed the animations to provide their best answers (see Figure 3 for number of replays). This choice ensures that the accuracy measured from the experiment translates to practical settings, where users could choose to repeat their interactions to better understand the data. However, this choice can impact the measurement of task time as individual participant confidence levels and preferences can influence the number of replays. We used best practices from past research to design the animations and reduce this effect.[49,50]

Limitations and Future Work

Our evaluation is just an initial step towards responsive visualization. For one thing, our work focuses solely on coordinated multiple view (CMV) layouts consisting of many smaller views. It could be argued that a high-resolution large display would be best suited for a single, highly-detailed visual representation that supports both digital as well as physical navigation (i.e., walking around the display). We chose to disregard such singleton visual representations, but this could be studied in future works.

Another limitation of our study is the use of specific responsive mappings in the boundary visualization interface. We chose horizon graphs[16] and a space-filling barchart[44] representations to convey the content of line and bar charts, since they have been explored by previous work. Our work here is therefore by no means exhaustive in the space of responsive visualization and there most certainly exist alternative techniques, some of them possibly more efficient than the two we propose and evaluate here. A significant future work is to apply our techniques to other visual representations. Towards this, the visualization-rich interface introduced in our study of practical use offers a starting point.

Finally, our participant population was both limited in size—13 in our controlled study and 8 in our qualitative study—as well as limited to graduate students in human-computer interaction (HCI) and related topics. Given that our results are clear and our sample size is consistent with comparable prior art, we still believe in the robustness of our findings. Second, while our HCI student participants represent a sample from a convenience population, we believe that their task performance or familiarity with charts is not significantly superior to information workers who are the primary target audience for our work. Third, we explicitly wanted to engage participants with significant computing experience to best reflect expert use in our qualitative study. Finally, our findings are also consistent with the general intuition: that a visual representation—the boundary visualization—specifically designed for space-saving on small displays performed better to the alternatives who are not, whereas this effect was not present on large displays where space is plentiful. In other words, we believe our findings are more or less accurate, representative, and generalizable.

Design Guidelines

Based on the results from our user studies, we provide the following preliminary guidelines for visualization designers.

- *Target responsive encodings and combine with layout transformations.* The boundary visualization proved to be an effective design as it used responsive visual mappings. Our participants were able to effectively work with the space-efficient encodings compared to other interfaces and preferred them over the overview+detail layout. To an extent, our boundary interface also uses a responsive layout. It compresses the boundary visualizations while adapting the overall layout of the interface. Therefore, the responsive visual encodings can be combined with layout transformations to achieve further responsiveness, since there may be a limit to how much a visual encoding can be compressed.

- *Highlight visual objects for tracking changes.* Some participants placed their fingers on objects they wanted to track on the interface on both small and large displays. Explicit highlights can support this practice.

- *Provide change indicators on the large display.* Users found it difficult to track changes on the large display due to its size. Visualizing change over time, such as time-lapse representations, can help improve user performance for large displays and can achieve better responsiveness.

Design Space: Achieving Responsiveness

We conclude the paper by introducing a design space for responsive visual interfaces derived from our findings. Some of the techniques described here were used in the user studies, while others remain to be explored in future work.

A *responsive visualization* automatically adapts its visual representation (and interactions) depending on the physical and computational capabilities of the device it is being rendered on[15,40,41]. Achieving responsiveness is traditionally viewed from a small display's perspective—how to design the interface primarily for the small screen (this is sometimes called "mobile first" in responsive web design)[55]. This is because an interface designed for a small display can easily (if not optimally) scale to a large display, but a large display interface when scaled to a small display faces issues with readability of the interface content. This is something we also observed with our classical interface on a small display.

Modern smartphones may have higher display resolution than even some large displays, but the effective screen viewport still remains smaller than commercial large displays (also known as the CSS viewport. Solving the challenge of responsiveness for small displays can naturally provide solutions for the large displays, and vice versa. This can further enable the creation of immersive analytics spaces, which contain many large and small displays for sensemaking. Therefore, here we discuss the design choices for transforming a visualization interface on a physically large display to a small display. We take an holistic view towards characterizing the design choices for responsive visualization by describing ways to adapt an entire visualization interface—including the layout, visual mappings of individual visualizations, and also the content shown within the visualizations.

Adapting Layout

A coordinated multiple view layout in a visualization interface consists of an $m \times n$ grid of coordinated views. When working on a small display, the layout of the grid, including the width, height, and positions of the individual elements, can be modified to fit the display. The design choices in layout manipulation include:

- **Stack:** Similar to responsive web design, the layout can be adapted by stacking the views vertically or horizontally. However, this can lead to a loss of positional information and spatial relationships within the CMV layout. For example, views that capture the same dimensions may be rearranged to no longer be in spatial proximity.

- **Distort:** Distortion techniques (e.g., fisheyes[56]) can be used to magnify areas of interest within a layout while compressing the rest. These techniques can be applied to a CMV layout to focus on specific views of interest while compressing the rest to save space.

- **Proxy:** When there are specific views of interest (focus region) in the user interface, the rest of the views can be replaced with a proxy widget(s) that can range from markers, icons, or even simpler visualizations. This proxy can also be an overview panel as used in our evaluations.

Adapting Visual Mapping

The visual representation may also need to be adapted to the available display space. Based on the existing literature described in the related work, we identify three choices:

- **Fill:** Space-filling techniques can take full advantage of the available space to restructure the visualization. These techniques allow the graphical primitives to cover the entire space (e.g., treemaps[22]). Similarly, pixel matrices[26] can capture the trends in line charts and bar charts by capturing information at each pixel in the visualization view.

- **Layer/Fold:** For charts with graphical primitives (e.g., paths) spanning either the X or Y dimension, layering (or folding a dimension) saves space by splitting the dimension into parts that are overlaid and distinguished using other visual variables such as color or opacity if needed. For example, horizon charts[24] perform layering to significantly save space over traditional line charts.

- **Merge:** Alternatively, merging visualizations within the interface into composite representations can save space. As Javed and Elmqvist[28] discuss, there are four design choices for composite visualization: juxtaposition, superimposition, overloading, and nesting.

Adapting Underlying Data

Finally, the data content embedded within the visualization can be changed to make it more comprehensible on a smaller display. The dimension corresponds to avoiding sharp and unreadable features of a visualization on a small display by intelligently adapting these features at a data level. However, this is a double-edged sword since lower resolution data may not promise the same insights.

- **Aggregation and sampling:** These techniques are often used for managing the amount of information rendered through grouping, discretization, and sampling the data. They can be repurposed in responsive visualization to manage the content based on the physical display size.

- **Identify points of interest:** Another approach for managing content is to switch to only showing perceptually important points. These points capture important features within a visual representation in terms of visual perception[57].

Adapting Interaction

Interaction mechanisms in a visualization are connected to the input modality (e.g., mouse or touch) and the elements within the visual representation. As identified in existing literature[15], differences in input modalities can be bridged by abstracting input events to work across device platforms (e.g., click becomes tap on touch displays). However, such adaptation may not be enough as the input is also affected by the display size. For example, its hard to precisely touch a point on a smartphone due to the so-called "fat finger" problem.

Furthermore, when the visual representation changes (e.g., based on the above design choices), adapting the interaction model becomes more complex and dependent on the representation itself. Therefore, adapting interaction requires further consideration from a user experience standpoint beyond the layout, visual, and data transformations introduced earlier.

Conclusion

In this paper, we evaluated two responsive designs of a visualization interface for varying display size. Our two complementary studies—a controlled study and a qualitative study—provided insights into how the small screen interfaces perform against each other as well as the classical interface on a large display. We found that the space-saving boundary visualizations on the small display led to higher accuracy than both the large display interface and an overview interface on small display for comprehending visual changes. We attribute this to the efficient use of the space, compared to the large display. Similarly, in practical use, participants were able to directly use the boundary visualizations to answer questions; thus, saving the number of operations (focus switches) compared to the overview interface on the small display.

Beyond the studies, we provided a holistic picture through a design space for achieving responsiveness in visualization interfaces. Our future work includes exploring more responsive designs as well as taking an inverse perspective to this problem: expanding a small display interface to a large screen. We also call upon the visualization community to consider our design space and evaluate techniques for achieving responsiveness, along with us, to solidify the concept of responsive visualization.

Acknowledgments

This work was supported by the U.S. National Science Foundation award IIS-1539534. Any opinions, findings, and conclusions or recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the funding agency.

References

1. Andrews C, Endert A and North C. Space to think: large high-resolution displays for sensemaking. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 55–64.
2. Andrews C and North C. Analyst's workspace: An embodied sensemaking environment for large, high-resolution displays. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*. pp. 123–131.
3. Amatriain X, Kuchera-Morin J, Hollerer T et al. The AlloSphere: Immersive multimedia for scientific discovery and artistic exploration. *IEEE Multimedia* 2009; 16(2): 64–75.
4. Isenberg P, Dragicevic P, Willett W et al. Hybrid-image visualization for large viewing environments. *IEEE Transactions on Visualization and Computer Graphics* 2013; 19(12): 2346–2355.
5. Isenberg P and Carpendale S. Interactive tree comparison for co-located collaborative information visualization. *IEEE Transactions on Visualization and Computer Graphics* 2007; 13(6): 1232–1239.
6. McGrath W, Bowman B, McCallum D et al. Branch-explore-merge: Facilitating real-time revision control in collaborative visual exploration. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. pp. 235–244.
7. Badam SK and Elmqvist N. PolyChrome: A cross-device framework for collaborative web visualization. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. pp. 109–118.
8. Sadana R and Stasko J. Designing and implementing an interactive scatterplot visualization for a tablet computer. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*. pp. 265–272.
9. Langner R, Horak T and Dachselt R. VisTiles: Coordinating and combining co-located mobile devices for visual data exploration. *IEEE Transactions on Visualization and Computer Graphics* 2017; 24(1): 626–636.
10. Lee B, Isenberg P, Riche NH et al. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Transactions on Visualization and Computer Graphics* 2012; 18(12): 2689–2698.
11. Roberts J, Ritsos P, Badam SK et al. Visualization beyond the desktop—the next big thing. *IEEE Computer Graphics and Applications* 2014; 34(6): 26–34.
12. Weiser M. The computer for the 21st Century. *Scientific American* 1991; 265(3): 94–104.
13. Elmqvist N and Irani P. Ubiquitous analytics: Interacting with big data anywhere, anytime. *IEEE Computer* 2013; 46(4): 86–89.
14. Isenberg P, Elmqvist N, Scholtz J et al. Collaborative visualization: definition, challenges, and research agenda. *Information Visualization* 2011; 10(4): 310–326.
15. Leclaire J and Tabard A. R3s.js–towards responsive visualizations. In *Proceedings of the Workshop on Data Exploration for Interactive Surfaces*. pp. 16–19.
16. Heer J, Kong N and Agrawala M. Sizing the horizon: the effects of chart size and layering on the graphical perception of time series visualizations. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 1303–1312.
17. Javed W, McDonnel B and Elmqvist N. Graphical perception of multiple time series. *IEEE Transactions on Visualization and Computer Graphics* 2010; 16(6): 927–934.
18. Roberts JC. State of the art: Coordinated & multiple views in exploratory visualization. In *Proceedings of the IEEE International Conference on Coordinated and Multiple Views in Exploratory Visualization*. pp. 61–71.
19. Wang Baldonado MQ, Woodruff A and Kuchinsky A. Guidelines for using multiple views in information visualization. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*. pp. 110–119.
20. Godinho PIA, Meiguins BS, Meiguins ASG et al. PRISMA – a multidimensional information visualization tool using multiple coordinated views. In *Proceedings of the International Conference on Information Visualization*. pp. 23–32.
21. Sarikaya A, Correll M, Bartram L et al. What do we talk about when we talk about dashboards? *IEEE Transactions on Visualization and Computer Graphics* 2018; 25(1): 682–692.
22. Johnson B and Shneiderman B. Tree-maps: A space-filling approach to the visualization of hierarchical information structures. In *Proceedings of the IEEE Conference on Information Visualization*. pp. 284–291.
23. Stasko J and Zhang E. Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy

visualizations. In *Proceedings of the IEEE Symposium on Information Visualization*. pp. 57–65.

24. Saito T, Miyamura HN, Yamamoto M et al. Two-tone pseudo coloring: Compact visualization for one-dimensional data. In *Proceedings of the IEEE Symposium on Information Visualization*. pp. 173–180.
25. Keim D, Hao MC, Ladisch J et al. Pixel bar charts: a new technique for visualizing large multi-attribute data sets without aggregation. In *Proceedings of the IEEE Symposium on Information Visualization*. pp. 113–120.
26. Hao MC, Dayal U, Keim D et al. A visual analysis of multi-attribute data using pixel matrix displays. In *Proceedings of SPIE Visualization and Data Analysis*. pp. 1–9.
27. Keim DA. Designing pixel-oriented visualization techniques: Theory and applications. *IEEE Transactions on Visualization and Computer Graphics* 2000; 6(1): 59–78.
28. Javed W and Elmqvist N. Exploring the design space of composite visualization. In *Proceedings of the IEEE Pacific Symposium on Visualization*. pp. 1–8.
29. Elmqvist N, Henry N, Riche Y et al. Mélange: Space folding for multi-focus interaction. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 1333–1342.
30. Plaisant C, Carr D and Shneiderman B. Image browsers: Taxonomy and guidelines for developers. *IEEE Software* 1995; 12(2): 21–32.
31. Hornbæk K, Bederson BB and Plaisant C. Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM Transactions on Computer-Human Interaction* 2002; 9(4): 362–389.
32. Bezerianos A and Balakrishnan R. The vacuum: facilitating the manipulation of distant objects. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 361–370.
33. Ghani S, Riche NH and Elmqvist N. Dynamic insets for context-aware graph navigation. *Computer Graphics Forum* 2011; 30(3): 861–870.
34. Thevenin D and Coutaz J. Plasticity of user interfaces: Framework and research agenda. In *Proceedings of INTERACT*. pp. 110–117.
35. Calvary G, Coutaz J, Thevenin D et al. A unifying reference framework for multi-target user interfaces. *Interacting with Computers* 2003; 15(3): 289–308.
36. Tan DS, Gergle D, Scupelli P et al. With similar visual angles, larger displays improve spatial performance. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 217–224.
37. Liu C, Chapuis O, Beaudouin-Lafon M et al. Effects of display size and navigation type on a classification task. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 4147–4156.
38. Jakobsen MR and Hornbæk K. Interactive visualizations on large and small displays: The interrelation of display size, information space, and scale. *IEEE Transactions on Visualization and Computer Graphics* 2013; 19(12): 2336–2345.
39. Yost B, Haciahmetoglu Y and North C. Beyond visual acuity: the perceptual scalability of information visualizations for large displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 101–110.
40. Badam SK and Elmqvist N. Visfer: Camera-based visual data transfer for cross-device visualization. *Information Visualization* 2017; DOI:10.1177/1473871617725907.
41. Hoffswell J, Li W and Liu Z. Techniques for flexible responsive visualization design. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, pp. 1–13. DOI:10.1145/3313831.3376777.
42. Badam SK, Amini F, Elmqvist N et al. Supporting visual exploration for multiple users in large display environments. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST)*. p. to appear.
43. Kister U, Reipschläger P, Matulic F et al. Bodylenses: Embodied magic lenses and personal territories for wall displays. In *Proceedings of the 2015 ACM International Conference on Interactive Tabletops & Surfaces*. ACM, pp. 117–126. DOI:10.1145/2817721.2817726.
44. Chen Y. Visualizing large time-series data on very small screens. In *Short Paper Proceedings of the IEEE VGTC/Eurographics Conference on Visualization*. The Eurographics Association, pp. 37–41. DOI:10.2312/eurovisshort.20171130.
45. Horak T, Badam SK, Elmqvist N et al. When David meets Goliath: Combining smartwatches with a large vertical display for visual data exploration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, pp. 19:1–19:13.
46. Reda K, Johnson AE, Papka ME et al. Effects of display size and resolution on user behavior and insight acquisition in visual exploration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 2759–2768.
47. Cockburn A, Karlson A and Bederson BB. A review of overview + detail, zooming, and focus + context interfaces. *ACM Computing Surveys* 2008; 41(1): 2.
48. Elmqvist N and Yi JS. Patterns for visualization evaluation. *Information Visualization* 2015; 14(3): 250–269.
49. Robertson G, Cameron K, Czerwinski M et al. Animated visualization of multiple intersecting hierarchies. *Information Visualization* 2002; 1(1): 50–65.
50. Heer J and Robertson G. Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics* 2007; 13(6): 1240–1247.
51. Cumming G. The new statistics: Why and how. *Psychological science* 2014; 25(1): 7–29.
52. VandenBos GR. Publication manual of the american psychological association. http://www.apastyle.org/manual/, 2016.
53. Sullivan GM and Feinn R. Using effect size—or why the P value is not enough. *Journal of Graduate Medical Education* 2012; 4(3): 279–282. DOI:10.4300/JGME-D-12-00156.1.
54. Inselberg A and Dimsdale B. Parallel coordinates. In *Human-Machine Interactive Systems*. Springer, 1991. pp. 199–233.
55. Wroblewski L. Mobile first. http://www.lukew.com/ff/entry.asp?933, 2009.
56. Furnas GW. Generalized fisheye views. In *Proc. of the ACM Conference on Human Factors in Computing Systems*. pp. 16–23.
57. Chung FL, Fu TC, Luk R et al. Flexible time series pattern matching based on perceptually important points. In *Proceedings of the International Joint Conference on Artificial Intelligence Workshop on Learning from Temporal and Spatial Data*. pp. 1–7.