






Many At Once: Capturing Intentions to Create And Use Many Views At Once In Large Display Environments

Jillian Aurisano^{1,2}  and Abhinav Kumar¹  and Abeer Alsaiani^{1,2}  and Barbara Di Eugenio¹  and Andrew Johnson^{1,2} 

¹ Department of Computer Science, University of Illinois at Chicago, Chicago, United States of America

² Electronic Visualization Laboratory



Figure 1: In three requests, a participant generates 15 views of her data. In the final request, the participant references a set of three views and poses a request to copy and pivot these views collectively to four new subsets of the data. The final result is a grid of views, which the participant used for an integrative data analysis task that spanned the views as a collection.

Abstract

This paper describes results from an observational, exploratory study of visual data exploration in a large, multi-view, flexible canvas environment. Participants were provided with a set of data exploration sub-tasks associated with a local crime dataset and were instructed to pose questions to a remote mediator who would respond by generating and organizing visualizations on the large display. We observed that participants frequently posed requests to **cast a net** around one or several subsets of the data or a set of data attributes. They accomplished this directly and by utilizing existing views in unique ways, including by requesting to **copy and pivot a group of views collectively** and posing a set of **parallel requests on target views** expressed in one command. These observed actions depart from multi-view flexible canvas environments that typically provide interfaces in support of generating one view at a time or actions that operate on one view at a time. We describe how participants used these ‘cast-a-net’ requests for tasks that spanned more than one view and describe design considerations for multi-view environments that would support the observed multi-view generation actions.

CCS Concepts

• **Human-centered computing** → **Empirical studies in visualization**;

1. Introduction

Large, multi-view environments present a variety of benefits in visual data exploration [AEYN11], particularly in contexts where users of the environment wish to juxtapose and arrange many views of data [AEN10, CCC*15], and generate integrative insights across these views [RFK*13]. However, interaction in these environments remains an area of active research [BSES17].

In this paper, we conducted an observational, exploratory study of visual data exploration in a large environment with a flexible canvas for displaying many views of data, using an approach similar to Grammel et al. [GTS10], where participants express their intentions to a remote mediator who responds on their behalf. This approach allowed us to study the intentions of the participants in-

dependent of any particular view generation paradigm or graphical interface, giving us access to what users of a large, multi-view environment would like to do when unconstrained by the design choices of existing tools.

Existing large and flexible canvas environments for displaying multiple views, which include both large displays and virtual canvases, (eg. [DHRL*12, GGL*14, JE13, BCC*05, ZZD14, YS19]) typically allow users to explore their data by producing one view at a time, either through drag-and-drop operations through a menu, through actions on elements within a single view, through trails of copied and pivoted single views, or through data-flow diagrams.

In contrast, we observed that **when participants expressed their intentions without constraint, they frequently posed re-**

quests for many views in one command, by asking for many subsets of the data and many data attributes at once. We term this ‘casting a net’. These requests were accomplished both through direct queries and by utilizing prior views that were persistently displayed on the canvas. When using existing views, participants frequently posed requests to copy and pivot these views, but they often did so in ways that ‘**scaled-up**’ their intentions, expressing **multiple, parallel copy+pivot actions to perform on a single view target**, or by **collectively copying and pivoting sets of views** in one command. These ‘cast-a-net’ requests enabled participants to efficiently produce sets of views with conserved features, which utilized the display space and allowed them to perform tasks that spanned many views.

In this paper, we contribute a detailed description of how participants efficiently expressed intentions to ‘cast a net’ to target many subsets of the data and data attributes. This includes collective and parallel actions on prior views on the display. We contribute a description of how these actions facilitated data exploration and discuss the design implications for large, multi-view environments.

2. Related Work

2.1. Large Displays

Recent research suggests a variety of benefits for information visualization in large display environments. When provided with “space to think”, analysts use large displays to organize analysis artifacts, encoding conceptual relationships by positioning related text documents together [AEN10]. Large displays also enable users to leverage movement and embodied cognition [JH15, EALN11] for improved memory in data intensive tasks. When perceptually scalable encodings are applied to data attributes, there is evidence to suggest that users can perform visual queries over large volumes of data, and over many related views of data [YN06, BI12, ROJL14]. Finally, given the ability to display more related views of data [RAF*], users appear to formulate integrative hypothesis that make use of these views [RJPL15]. In response to these findings, applications have been designed for large visualization environments targeting hybrid display of information [IDW*13], collaboration [MAN*14], presentation of large volumes of data [ARJ*15], and integration of 2d and 3d views [RFK*13].

Interaction with visualizations on large displays presents challenges and opportunities [AEYN11, BSES17]. Recent work has examined movement or proxemics as an input to visualization environments [BMG10, JHKH13], as well as multi-touch [JH14] and an ecology of devices [BFE15, MAN*14, HBED18].

Our work contributes to this body of work by examining use of a large display for a real visual data exploration scenario, but we capture intentions for views independent of a realized interface. Some of our findings echo Knudsen et al. [KJH12], where a whiteboard workshop captured interactions over many visual artifacts on large display surfaces. Our work complements this analysis by observing similar tasks that spanned many views and utilized large display areas. Our work expands upon Aurisano et al. [AKG*b, AKG*a], which presents preliminary analysis of this data, and Kumar et al. [KADE*16, KDEA*], which examines utterances and gestures from a natural language processing perspective.

2.2. View Construction and Multi-View Environments

A variety of flexible canvas environments have been created for information visualization, including virtual canvases with pan and zoom interaction. Broadly, these tools aim to enable users to generate many views of their data and position these views freely (e.g. [DHRL*12, GGL*14, JE13, BCC*05, ZZD14, YS19]).

View creation in these environments has been explored using a variety of interaction techniques. Initial views are often added to the flexible canvas through interaction with a menu, such as through drag and drop operations onto the canvas. Alternatively, natural language queries can create views, in systems such as FlowSense, which feature a natural language interface to the data flow model, where data is selected and transformed through a series of views [YS19]. Many of these systems also present ways to create new views through actions on existing views. One approach is to allow participants to copy and pivot a view target, or to create new views from selections within an existing view, to drill down into more focused portions of a dataset. These actions can facilitate the creation of visualization provenance trails, and aim to enable backtracking and revisions along the trail [DHRL*12, BCC*05]. Our work contributes to this line of research by providing empirical data for participant view creation intentions, in order to aid in the development of interfaces for view generation in large, flexible canvas environments.

Van den Elzen and Van Wijk focus on generation of multiple views from a single view, creating small multiple sets, in a trail that preserves context [vdEvW13]. Our work also finds value in generating many views at once, to efficiently explore the information space. We depart from this work, by focusing on a large, flexible canvas environment, where views can be freely positioned, and not in a layout that emphasizes visualization provenance. We also capture interaction types which might create many views at once.

View creation has been examined in a single view context in Grammel et al. This study used a remote mediator, who responded to requests for a view of the data. They found that view creation posed challenges for those not trained in visualization construction. Participants struggled to translate questions into appropriate visualizations. The take-away from this study is that even with robust graphical interfaces, which remove the need to learn coding or scripting for visualization creation, users may still make visualization construction errors [GTS10]. While we do not focus on view construction challenges in our study, and we examined view creation in a multi-view environment, we also observed under-specification of intended views as well as using the existing screen state as a shortcut in posing requests.

3. Methodology

In this section we present the design decisions in our evaluation and how these decisions allowed us to address our research goals. We also discuss limitations and how we address these limitations in our analysis. Our research goal was to observe visual data exploration in a large, wall-sized display environment to derive design goals for future systems that are grounded in how participants request new views, utilize and reference existing views on the display and utilize the display space in support of data exploration tasks.

Our research questions are :

1. In a multi-view environment how did participants request views?
2. How did participants use existing views to pose subsequent requests for new views?
3. How did the display space support analysis tasks that involve more than one view?

To address these questions, we had three broad goals in our study design: 1) **realism**: capture interactive intentions expressed in response to real visualizations of data within a realistic data exploration scenario in a large display environment; 2) **unrestricted expression of intentions**: capture interactive intentions independent of existing interfaces or interaction modalities, to capture what participants wanted to do when reasonably unconstrained; 3) **multiple rounds of view generation**: examine these intentions over an analysis session, with many rounds of visualization generation, in support of completing a data exploration task.

To meet these design goals we conducted an observational exploratory study in a laboratory setting, using a protocol that mirrored Grammel et al. [GTS10]. Recruited participants were given a data exploration task, and told to verbally express their intentions to a remote mediator, (a PhD student in data visualization), who was located in an adjacent room monitoring spoken and gestural communication from the participant over video and audio feeds. By locating the mediator in a different room, we distanced the participant from the interface used to generate new views. This allowed us to examine their behavior in an interface-agnostic setting. Like Grammel et al. participants were informed that the remote mediator was a person, and we do not simulate a system as in a Wizard-of-Oz study, as is used in other studies of interaction modalities in information visualization, [e.g. [WLJ*12, TS19]].

Unlike Grammel et al., we use a large, multi-view flexible canvas environment to persistently display prior responses, allowing us to look at how participants used these past views and the large display. In addition, we did not ask participants to specify an intended single view, but rather to ask anything that might aid in the exploration and analysis of the data. Study of abundant display space mirrors the work of Knudsen et al., but we use a digital environment with real views of data and many cycles of view construction [KJH12].

3.1. Piloting

To arrive at our final study design, we conducted pilots in two phases. In the first piloting phase, we performed an offline pilot with four remote subjects, who were presented with a document summarizing the data variables and a data analysis task, and had the opportunity to pose analysis or clarification questions over a two week period via email. This enabled us to refine the materials and add focused data exploration sub-tasks. Then, we conducted a pilot study with five participants in a laboratory environment. We refined our approach to responding to participant queries, particularly our approach to managing new visualizations as they were added to the display. We refined the experimental setup by shifting the cameras to ensure a better view of gestures.

Participant	Requests	Participant Data	
		Visualizations	Ratio
1	11	36	3.27
2	11	18	1.63
3	13	23	1.77
4	13	28	2.15
5	21	34	1.62
6	11	25	2.27
7	23	30	1.30
8	16	32	2.00
9	9	35	3.89
10	32	34	1.06
11	13	31	2.38
12	23	35	1.52
13	24	33	1.375
14	22	32	1.45

Table 1: Participant data: number, request count, number of views produced, view:request ratio.

3.2. Participants

14 participants (7 male and 7 female, ages 18 to 34), were recruited for the study, with an additional 4 participants in the first stage pilot and 5 in the second stage pilot. The participants were drawn from diverse fields including computer science, communications, business, speech-pathology education, biology and medicine. Participants had varied experience with visualization and data analysis, ranging from daily data analysis tasks (close to 50 percent of participants), to almost never conducting data analysis (20 percent of participants). All participants were familiar with common data visualization types and used computers daily. A few participants had used the large display environment for class or meetings, but they had not used it for data exploration.

Given this diversity, we do not draw conclusions that are specific to any particular background or level of expertise. Domain experts or novice analysis could be an area of focus in future work. However, since all participants were either students (10 participants) or professionals in data driven fields (4 participants), this group is appropriate to target for future realized systems. Participants completed their analysis in 45-90 minutes. Participant data (requests and views) is listed in Table 1.

3.3. Apparatus, Environment and Materials

We performed our study in a laboratory setting, allowing us to control the interface and environment, as well as manage the communication channels between participant and remote mediator. We opted to perform this study in a digital context, as opposed to a whiteboard, to capture multiple, continuous rounds of interaction, with visualization responses that were generated from real data.

The environment for the participants consisted in a large display wall (6.675 by 2.01 meters and 6144 by 2304 pixels) shown in Figure 2. Participants could refer to onscreen textual descriptions of 1) the data, including attributes and their values, and 2) the overall goal and sub-tasks. Paper copies of the task description and data




Figure 2: We conducted our study in a laboratory setting with the remote mediator in an adjacent room. The large display in the participant's environment presented a data and task description on the left side of the display, and a status bar and chat box at the top. The remote mediator had access to two video feeds, and two 4k displays mirroring the large display. The remote mediator generated views using Tableau, and dropped them to the user through Sage2.

description were removed, so the participant directed their attention to the display and gestured freely. On the top of the display we created a status bar, with an animation indicating when the remote mediator was working on responding to the request, and a chat box, in which the remote mediator could enter messages. This is depicted in Figure 2.

The remote mediator was isolated from the participant in a nearby room. As in Grammel et al. [GTS10], this allowed us to shield the participant from the interface used to generate new views, avoiding biasing effects and removing the influence of verbal or non-verbal feedback from the remote mediator. An in-room aide, a graduate student in computer science, was present in the room with the participant to explain the study protocol, address technical questions during the study and conduct the final interview. The remote mediator was not introduced to the participant until the study was complete. We chose to shield the participant from the mediator to encourage direct and honest feedback during the final interview. The remote mediator was provided with two video streams, showing the participant from the front and from behind, to capture pointing gestures and gaze, as well as facial expressions. The remote mediator viewed two 4k displays that mirrored the participant's large display, enabling the remote mediator to ensure optimal placement and sizing of the provided views. The remote mediator generated visualizations using Tableau on a laptop, and dropped exported static images of these views onto the large display using a collaborative large display software, Sage2 [MAN*14], which also supplied a laptop interface for re-sizing and positioning these views freely. The remote mediator used a chat box, to communicate with the participant, and a status bar, to show when the remote mediator was producing new views in response to participant queries. This is depicted in Figure 2.

3.4. Task

Their task was to explore crime data from a local, public data repository in order to decide how to deploy additional policing resources. We chose this task because we believed the data would be familiar to participants and that they would be motivated to explore this data out of personal interest. Each crime incident in the dataset included a **GPS coordinate**; a **neighborhood identifier** (from one of four

local neighborhoods); a date and time, which were used to infer the **time of day**, **day of the week**, **month of the year**, and **year** when the crime took place; a classification of the primary **crime type** (e.g. theft, burglary, assault...); as well the primary **location type** where the incident took place (e.g. street, residence, business). We supplied a list of general sub-tasks (e.g. examine changes over time, look for hotspots), to provide direction and a starting point to the participants. The addition of subtasks was based on feedback from participants in the first phase pilot.

3.5. Procedure

During the instruction phase, the in-room aide provided the data description and analysis tasks to the participants and instructed them to 'ask anything that would aid in your analysis', and were given no restrictions in the kinds of queries they could pose. The aide explained the interface, pointed out the location of cameras, and demonstrated through a short social exchange that the remote mediator could respond to spoken and gestural requests. Participants were encouraged to think aloud and describe their findings as they inspected provided views.

We opted to not include a learning phase in our study. We opted to allow analysts to discover system capabilities during the session, rather than through a learning phase to give as much time as possible to a single analysis. We opted to not provide views at the start because we did not wish to direct the analysis in a particular direction and we wished to capture initial queries as well as follow-up queries. We did not provide visualization templates, because we wanted participants to pose questions freely rather than specify views directly, which has already been investigated by Grammel et al. [GTS10]. Participants began with a blank canvas and a data and task description.

The remote mediator responded with visualizations where possible, including situations that could be answered with textual responses (e.g. "which crime occurred the most often?"). When participants posed a request, the mediator generated one or several views in response to their request using Tableau. If a request produced a multi-view response, all views were presented at once. Provided visualizations were numbered by request and given a title that communicated the subset of the data contained in the view (e.g. *Theft and Battery*), and the visualized data attributes (e.g. *Time of the Day*). All visualizations were saved and used in our analysis. In situations where the expected outcome to a request was unclear, the remote mediator made an appropriate guess rather than asking extensive follow-up questions. We opted for this response style because we did not want participants to feel that they needed to precisely specify views. We wanted them to pose requests freely. Participants were instructed to correct the mediator if the responses were not what they wanted. The mediator would select appropriate templates for any spoken data attributes, and would filter based on selected subsets of the data. Colors and scales were generally the defaults supplied by Tableau.

The visualizations presented to the participant were static images exported from Tableau, not interactive views. The benefit of this approach was that it allowed us to focus on view generation actions in a large display context, and bracket the challenge of view modification and multiple coordinated views, which would have introduced

a large range of design choices to our study [KPV*17]. The view coordination problem could be addressed in future work.

We made the decision to respond to all requests for new subsets of the data, new data attributes or new visual templates with new views, in new separate windows. In contrast, when participants wanted to modify the encodings, scales or layouts within a view (e.g. adding labels, changing color schemes), we treated these as view modification requests and replaced the old static image with a new static image reflecting the requested change. We opted to adopt this distinction because we wanted the participant to capitalize on an important feature of the large display space: the ability to externalize and spatialize their exploration of the data [AEN10]. By creating new views in response to requests to explore new subsets of the data and/or new data attributes, participants could relate and compare these new views to prior views, and access their prior findings by retrieving the associated visualization. A similar distinction can be found in Javed et al. [JE13], where new views are spawned in response to requests to explore a new portion of the data and attribute space. We also distinguished between new view and view modification request types in our analysis.

The remote mediator had control over view positioning. We attempted to make reasonable decisions in our protocol for view positioning. Related views were juxtaposed next to each other, small multiples were positioned in a line or a grid. If views displayed different temporal subsets of the data, we ordered these views from left to right from earliest to latest. Views that were active were generally centered and views that were inactive were moved to the side. We opted to position the views automatically for the user for several reasons. First, we learned during the second phase of piloting that participants struggled to interpret a set of views and make decisions about where to position them, and that verbal positioning instructions were time consuming. Second, views generated during the pilot study very quickly filled the display which made it challenging for participants to pose new requests. We could have provided a secondary device or interaction modality for view positioning, but this would have tethered the participants to a device and we wished to encourage interaction through the mediator. A limitation of this choice is that we captured fewer view layout requests, and this could be an interesting direction for future work. The layout protocol involved first deciding whether to move aside prior views and then arranging the new views in the central region of the display, which we call the ‘active’ region.

Participants decided when to end the session, based on when they felt they addressed the data exploration task. We did not require participants to complete all subtasks, because we wanted to encourage exploration driven both by the listed subtasks and by insights from visualizations. Participants generally used the subtasks as a starting point, and referred back to them when they had completed a thread of the analysis and wished to pivot in a new direction. As seen in table 1, the mediator provided an average of 30 views to an average of 17 requests. Following the session, the participant took a computerized survey and completed an interview with the aide. The remote mediator would visibly exit the session before the participant began the survey by deleting the chat box and status bar, in order to encourage candid responses. A complete description of our

procedure, along with visualizations produced and chat transcripts, is in our supplemental materials available online[†].

3.6. Analysis

We opted for qualitative analysis methods in order to capture rich behavior within a realistic scenario, and we use a grounded approach [CHCPM07]. The recorded video was transcribed in full. We also used the stored and numbered static visualizations and chat transcripts from the sessions. A team of three researchers reviewed a subset of the participants transcript and video. This team met several times to discuss high-level themes. We note that 1) participants expressed their intention to generate visualizations either directly or by utilizing existing views on the display and 2) participants frequently generated many views that could be arranged into coherent group with relatively few interactions through the mediator. These themes informed the adopted coding approach.

A primary coder created a visual record of each participant session. For each request, a visual ‘scene’ was created that depicted 1) snapshots from the video showing the display before the request, 2) the transcript of the request, 3) snapshots from the video showing the participant’s gestures to onscreen views, 4) the images of the views provided to the participant and 5) snapshots from the video showing the display following the request. We also created scenes for changes to the view layouts and think aloud. All scenes are available in our supplemental material online[‡]. We adopted this approach because we needed to rapidly review the transcript alongside the display state, the provided views and the participant’s movement and actions. Over 550 scenes were compiled in total, with 23-64 per participant. 215 scenes contained requests for new visualizations, which in our protocol consist in requests to explore a new subset of the data and/or a new data attribute. We did not focus in our analysis on think-aloud, on requests to modify a view (such as adding labels) or on requests to move or re-position views. This decision reflected the themes we identified during early review of the study. In addition, this focus allows us to contribute to an important step in data exploration- viewing a new portion of the data and attribute space- that is particularly relevant in a large display context, where users benefit from externalizing and spatializing their exploration process [AEN10]. In addition, since large displays with a flexible canvas do not pose a fixed limit on the number or composition of viewports onto the data, it relevant to focus on interactions to create new viewports that take advantage of this flexibility.

The primary coder used an open coding approach to refine a set of codes to apply to the visualization scenes. Codes were developed through an iterative, multi-pass process. These codes were discussed with two coding reviewers. The coding reviewers posed questions and flagged ambiguous cases. After discussion, the codes were modified through several passes. This review and discussion process was repeated several times, until the codes were relatively stable and addressed the themes.

The final codes capture both how views on the display were related to participant requests, and how many data attributes and subsets of the data were requested, an approximation of how many tasks participants performed that spanned more than one view. We noted unusual features within what we term referential requests,

and developed a set of codes specific to this request type. Codes are available in our supplemental material online[§].

4. Findings

In this section we describe our coded observations from participant visual data exploration sessions. Our coding scheme is divided into three parts. The first part identifies the ways in which participants utilized or did not utilize existing views to express their intentions. Essentially, these codes identify *how* participants expressed their intentions. We identified three strategies - direct (41 percent), referential (42 percent), and selection (17 percent). This primary division helped us to isolate different strategies participants used to express complex intentions efficiently.

The second part of our coding scheme looked at whether a participant's request targeted a single data attribute and a single subset of the data, which we term a **targeted** request that could be presented in one view, or whether the request **cast a net** around several subsets of the data and/or several data attributes within the information space. We divide cast-a-net requests into browse, compare and complex multifaceted. This coding gives us access to *what* multi-view, multi-subset, multi-data-attribute intentions participants requested through the mediator. For the purpose of this discussion, we define a subset of the data as a set of rows from a tabular data set, where the rows are selected based on one or several filters. By data attribute we mean the columns of a tabular data set, such as the day of the week the crime occurred on, the crime type, or the neighborhood of the crime. These two code parts and their frequencies are summarized in Figure 3.

The third part is applied specifically to referential requests. In this part, we examined the number of views that are targeted in a referential request and the number of actions that are specified on the target(s). This allowed us to concretize our observation that, when unconstrained, participants 'scaled-up' their intentions to create or operate on many views at once, to extend their exploration to data subsets and data attributes, and to perform tasks that spanned more than one view.

4.1. Direct, Referential and Selection Requests

We observed that participants efficiently posed complex requests through the mediator. A significant way that they accomplished this was by using existing views, either as templates or for selection/drill-down. To isolate these requests, we looked at how participants referenced or utilized 'active views', or views in the center of the display, in formulating their request, and any verbal or gestural indication toward these views. In cases where participants referenced view targets in posing their requests, we labeled these as **dependent requests**, because they relied in some way on existing views. These constituted 59 percent of visualization requests. The remainder we termed **direct requests**. In these requests the participant specified the intended view with no reference to existing views and they represented 41 percent of view requests.

Of dependent view requests, most were labeled as referential.

Referential view requests came in the form of "Can I see this, but...", where participants indicated a view target or targets using speech and/or gestures, expressed an action or actions to perform on the target(s) which resulted in the outcome, a new view or views on the display. In effect, participants specified a desire to copy and pivot the target or targets to a new portion of the "information space". Of the 215 visualization requests, 92 were referential requests, around 42 percent. It should be noted that referential requests did not include requests to correct an error, such as indicating that a particular view did not address the expressed query. These only included requests that referenced an existing view and used it as a shortcut in expressing a request to the mediator, where this request would pivot the targets to a new portion of the dataset. A request to correct an error would be classified based on the content of the request, and could be targeted, selection or referential.

In the third major category, selection, participants requested a new view that focused on a region selected from a target view. For instance, a participant looking at a breakdown of crimes by day of the week may then ask to see Friday crimes, selecting the 'Friday' subset of the data, within a new view with respect to a new attribute. These represented 17 percent of all visualization requests. The resulting trail of views would have a hierarchical relationship which includes a parent view that is selected from, and a child view which displays the selection.

4.2. Target vs Cast-a-Net

In the second part of our coded observations, we distinguish requests that were **targeted** to a single data attribute (e.g. day of the week) and a single subset of the data (e.g. Thefts in 2014), from requests that **cast a net** over a number of data attributes and subsets of the data. These codes capture whether participant explored by expressing a focused question to which a singular response could be provided, or whether they wished to look across many portions of the dataset at once. This distinction allowed us to analyze requests where participants efficiently arrived at sets of views that took advantage of the display space and facilitated multi-view analysis tasks, which helped us address our research goal of understanding data exploration and view creation in a large display environment. Of cast-a-net requests, we observed three major categories- browse, compare and complex multifaceted- based on the particular data subsets and attributes enumerated in the request. Most requests were labeled as cast-a-net requests, and we explore the implications of this in the discussion.

4.2.1. Targeted

A targeted request is one where the participant specified a single subset of the data and a single attribute of interest, which could be responded to within a single view. A targeted request might include asking for a map of thefts or frequencies of thefts by day of the week. Alternate views could be provided, such as using alternate visual templates, colors, scalings and other encodings, but the portion of the data shown and major divisions and aggregations would remain consistent across alternate views. Overall, 35 percent of requests were labeled as targeted.

[§] <https://github.com/uic-evl/many-at-once-paper>

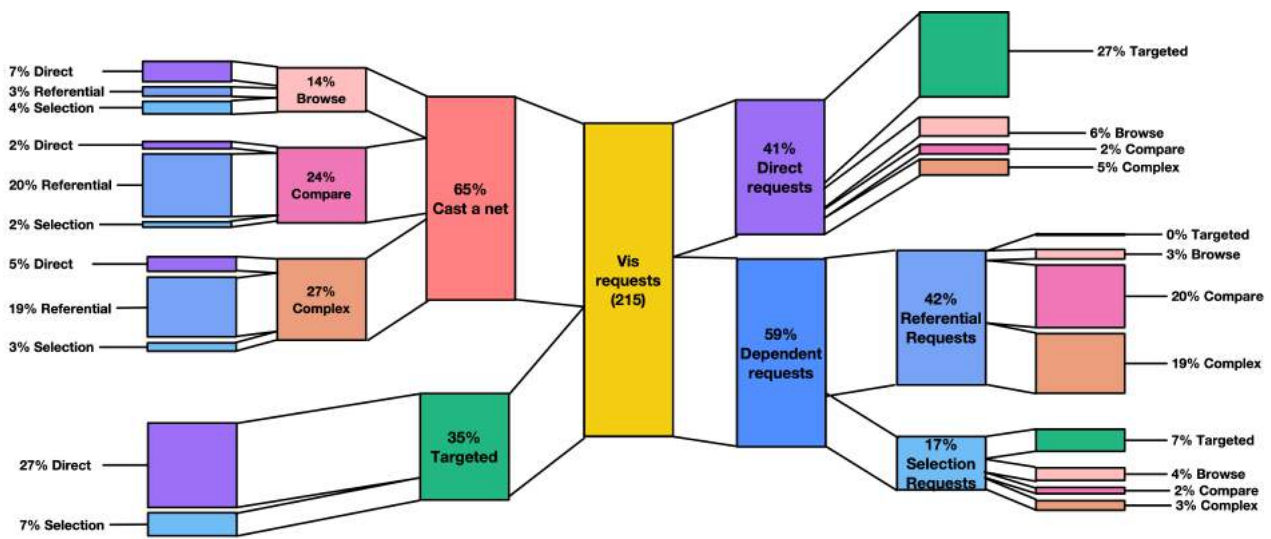


Figure 3: This figure presents an overview of the coded observations of how participants utilized existing views to pose their requests, and how many data attributes and subsets of the data they requested in one command.

4.2.2. Cast a net: Overview

The remaining 65 percent of requests spanned more than one subset of the data and/or more than one data attribute. These requests typically elicited either a set of views or one or several large multifaceted views, where several aggregations of data are presented within one window, such as a divided bar chart or a multi-line chart. We classified ‘cast-a-net’ requests into 3 categories: compare, browse and complex multifaceted requests, depicted in Figure 4

All referential requests were coded as cast-a-net, even if one new view is produced from the request. We did this because the new view possessed a relationship to the target and the participant generally used the resulting pair or set of views to perform tasks that spanned more than one view. This decision also allowed us to capture how the referential request arrived at multi-view states in supporting of browsing, comparing and faceting, similar to the states achieved through direct cast-a-net requests.

4.2.3. Cast-a-Net: Browse

In browse cast-a-net requests, participants expressed a single subset of the data that they wished to focus on (e.g. Thefts on Saturday), but requested several data attributes within that subset (eg. ‘by year and by month and by hour’). In effect, the participant expressed the intention to browse several attributes and views within a focused area. The resulting views would allow the participant to browse for trends, features and patterns within the subset of the data of interest. Cast-a-net browse requests constituted 14 percent of all visualization requests.

4.2.4. Cast-a-Net: Compare

In compare cast-a-net requests, participants requested different subsets of the data with respect to a common data attribute. For example, a request to examine two neighborhoods by crime type, would

be classified as a comparison request, because the participant specified one data attribute and several subsets of the data. Sensible responses to these requests include multiple views in separate windows or in multifaceted views within the same window, such as multi-line charts or grouped bar charts. These views allowed participants to compare distributions, trends, or spatial hotspots across multiple subsets of the data. Cast-a-net compare requests constituted 24 percent of all visualization requests.

4.2.5. Cast-a-Net: Complex Multifaceted

In complex multifaceted requests, participants would express interest in several subsets of the data and several data variables. Responses to these requests would include multiple views with permutations of the subsets and variables of interest. At times, participants might request views that allowed them to simultaneously browse within several subsets of the data, and compare these subsets against a set of common data attributes, with each dimension presented in a grid. At other times, complex requests might warrant combinations of multifaceted views, to enable participants to facet the data in different ways. Cast-a-net complex multi-faceted requests constituted 27 percent of all visualization requests. In cases where participants requested a subset of the data with respect to two attributes, it could be argued that a single multifaceted view could be offered in response. However, we found that these requests could elicit several kinds of responses. For example, a request to see thefts by year and neighborhood, could be responded to with a multi-line chart or with a grouped bar chart or with small multiples. Due to this complexity, felt that requests for two attributes belonged in this category.

4.3. Creating Many Views with Referential Requests

Within referential requests, we captured the number of targets, actions and outcomes of the referential request. Of the 92 referential

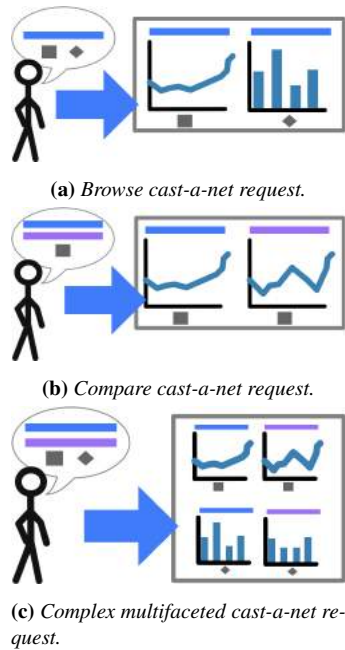


Figure 4: *a. Browse cast-a-net requests:* Participants indicate one subset of the data, in this graphic represented by the long blue bar, with respect to more than one data attribute, in the graphic represented by the gray square and diamond. Participants used these requests to focus on a subset of interest, and explore several data attributes. *b. Compare cast-a-net requests:* Participants indicate one data attribute of interest, in this graphic represented by the gray square, with respect to more than one subset of the data, in the graphic represented by the long blue and purple bars. Participants used these requests to compare subsets with respect to a single data attribute. *c. Complex cast-a-net requests:* Participant indicate more than one data attribute and/or subset of the data. In this graphic they requested two subsets of the data with respect to two attributes, and the responses were depicted in a grid. Participants used these requests to explore complex combinations of attributes and subsets.

requests, the majority targeted a single view, expressed a single operation to perform on that view, producing a single outcome (37 requests). The remaining referential requests were coded as one-to-many (25 requests), many-to-one (6 requests), and many-to-many (21 requests). These requests enabled participants to efficiently express desires for complex sets of views.

4.3.1. One-to-One

In one-to-one referential requests, the participant indicated a single view and specified a single operation to perform on this view, which would produce a single outcome, in our case a new view on the display. For example, one participant pointed to a multi-line chart, displaying data by day of the week, and asked "Can I have a look at this (pointing to the target) by month of the year?". By referencing this complex template, the participant was able to shift from one data attribute (day of the week) to another (month of the

year), while still retaining the other components of the view. The conserved features serve to link these views together.

4.3.2. One-to-Many: Parallelized Copy+Pivot

One-to-many referential requests occurred where participants referenced a single view target, but expressed an intention to perform **multiple operations in parallel on this view**, producing a set of views unified by preserved features from the original template. For instance, in Figure 5a, the participant pointed to a view showing thefts by time of the day, asked "Give me the same of this (pointing) with battery, deceptive practices, criminal damage and assault, please". The mediator took the specified view, preserved the template and x-axis, and repeatedly changed the filter from theft to the enumerated crime types, producing a new view for each of the specified types. The participant then scanned the set of views and identified differences in the hourly distribute of battery crimes, when compared to the other crime types. Parallel actions of this kind are highly efficient. **Rather than request each new view one at a time participants opted to bundle the actions together within a single request.**

4.3.3. Many-to-Many: Collective Copy+Pivot

In 21 cases, we observed participants making referential visualization requests by indicating many view targets, through pointing or speech, and then expressing one or several operations to perform **collectively** on the indicated targets. We term these actions 'many-to-many', because many views were targeted by the participant, with the intention to produce many views and extend the reach of their exploration. In one case, pictured in Figure 5b, the participant points to two views, one of which shows theft by month and the other theft by day of the week. The participant asks "Can I get these same charts but just for battery." To respond to this request, the moderator pivoted the two views, producing two new views. When all the views were juxtaposed on the display, the participant then compared the number of battery and theft crimes by day of the week and month, and identified several differences.

In many of these cases, participants collectively operated on sets of views with conserved features. For instance, a set of views with a common set of filters could be pivoted to a new set of filters. A set of views with a common visual template could be pivoted to a new set of data attributes. We speculate that these **commonalities across views served to signal to the participants that sets of views could be referred to collectively and acted on as unit.**

4.4. How, What and How Many

Examining coded observations in combination, several interesting features emerge. Of direct requests, one third were labeled as cast-a-net. Participants would pose direct browse requests either by asking for general information about their area of interest or by bundling several data attributes together, often using language that applied to several data attributes (e.g. 'where crimes occur' included several data attributes) or by wanting to know when crimes occurred, and failing to specify temporal aggregation. Direct browse requests could be seen as related to under-specification of intentions or the high cognitive load of interaction in the absence

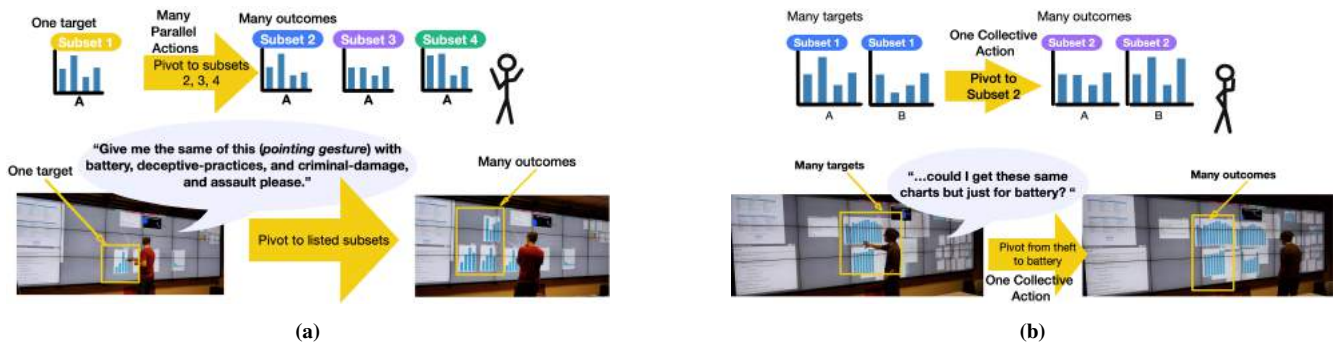


Figure 5: *a.* In this one to many referential request, a participant indicates a view target and expresses multiple operations to perform in parallel on the target, enabling a comparison task across the resulting views. *b.* In this many to many referential request, a participants collectively duplicates a pair of views, and performs tasks on the resulting view grid.

of a visual interface, which came up during interviews with participants.

Comparison cast-a-net requests were most frequently accomplished via-reference to existing views. In these cases, participants might be examining a view and then would ask to pivot to a new filter. Essentially participants wanted to know if their observations extended to other subsets of the data. Complex referential requests frequently involved either faceting the target view, such as by subdividing the view by an additional data attribute, or by a collective many-to-many copy and pivot operation, which would produce a grid of views with conserved features in each dimension, allowing for smooth movement between two different multi-view tasks, one accomplished by scanning horizontally across views within the grid and the other by scanning vertically.

While many selection requests were targeted, requesting the selected subset of the data with respect to a single data attribute, other cases were more complex. Half of selection requests 'cast a net' around the selection, either requesting several data attributes for that selection (browse, 25 percent of selection requests), or faceting this selection with respect to several data attributes (complex, 17 percent of selection requests). In other cases, participants made several selections from the target view and compared these selections across a conserved attribute (12 percent of selection requests).

4.5. Cases

We observed that the cast-a-net requests and referential requests produced coherent sets of views and enabled a variety of analysis tasks that integrated information across many views. Several participants used **repeated requests of these types to efficiently create many views in a few interactions with the mediator**. The views could then be **positioning in grids and clusters**, to perform **simultaneous multi-view analysis tasks**, browsing, comparing, trend identification, and faceted exploration.

The first case we wished to highlight involved a participant who used four queries to produce 29 visualizations. She began with a direct browsing request focusing on the neighborhood around the university, which resulted in 7 views focused on the university. This many-to-many referential operation was repeated, for two more

neighborhoods resulting in a screen state with 28 views, seven for each neighborhood and four for each attribute, in just 3 requests. These views were presented in a grid that permitted her to perform a between-neighborhood comparison task across pairs of views and a within-neighborhood browsing task within several views showing different data attributes.

In the second case study that we wished to highlight, the participant made a series of referential requests, resulting in three multifaceted views that showed a common subset of the data (crimes in 2014) and a common aggregation in a multifaceted bar chart by the four neighborhoods. Then in her final request, she targeted 3 views for a collective and parallel copy+pivot operation to produced 15 views, which collectively covered 4 data attributes within 5 years. Walking from left to right, and scanning vertically, she could smoothly move between trend analysis within a neighborhood, comparing trends across neighborhoods, as well as browsing for interesting patterns within a neighborhood and year. This case is pictured in Figure 1.

4.6. Interviews and Surveys

Our interviews and surveys with participants enabled us to examine how participants experienced the data exploration sessions, their impressions of the quality of the visualizations and their reaction data exploration with a large, multi-view, flexible canvas.

4.6.1. Overall impressions

Many participants directly commented that they they liked the experience. One participant stated *"The experience was amazing. Most of my queries were satisfied through the visual presentations. The data provided enough data for my understanding of the visualizations. The data analysis expert understood all my questions and I got a prompt visual response"*. In the survey the participants responded to a set of questions on a five-point Likert scale. All of the participants felt that the mediator always or usually understood their requests (50 percent within each score). 66 percent felt that the responses always helped them analyze the data and the remaining 33 percent felt that the responses usually helped them analyze the data. 93 percent of participants felt the responses met their expectations all (53 percent) or most (40 percent) of the time.

Participants noted that responses with multiple views were valuable. For example, one participant stated *“It was impressive to see the data and be able to compare contrast it in many different ways. Each visual makes you consider a new aspect and/or want to inquire about new data to find new patterns.”* . Another stated, *“The multiple responses were very helpful. Sometimes the additional responses helped answer a complex question, and could be used to compare more detailed responses to more general ones.”* In survey responses, 60 percent of participants preferred getting multiple responses, with the remaining either holding no opinion (33 percent) or preferring one response (7 percent).

4.6.2. Blank canvas challenges

Several participants described challenges related to knowing where to start, formulating requests verbally, and facing a blank canvas. One participant noted *“The cognitive load of like thinking about what I want to visualize and translating that is just more steps vs like, I want to look at that, click click click, doing it myself.”* In contrast, some participants appreciated the ability to offload tasks onto the mediator. One participant stated *“It is much more convenient to just say and get things done, rather than implementing your own. It lets you, at least in my case, I could completely focus on what I wanted to do, instead of ‘do I click here, should I draft that?’.* What am I trying to solve, that is all I focused on. I loved that part.” Another noted that the process of verbalizing their intentions may have helped them with planning and decision making stating *“...sometimes the act of describing a chart helps you figure out exactly what you want. Or, in some cases, you realize that what you’re asking for doesn’t make sense and you change your mind.”*

4.6.3. View organization challenges

Participants who commented on the window positioning approach adopted in the study, where the mediator automatically positioned views for the participant, tended to have more negative impressions. These challenges are noteworthy, because even though the mediator had extensive experience positioning and displaying views for the participants, and managing large numbers of views, doing so manually posed challenges and was imperfectly executed at times. A few participants wanted control over view positioning and suggested a touch screen to enable this. While this was not possible in our study, it would be sensible in a realized system to provide some direct control to the user in managing the views on screen. However, from the study pilots, we knew that some automatic decisions in managing the views was needed from the mediator, otherwise visual clutter was a significant barrier.

5. Discussion

In this section we integrate our coded observations from the data exploration sessions with participant comments from the interview and survey in order to consider the design implications of our findings.

5.1. Arriving at Many, Not Just One

We found that cast-a-net view generation was a common request style. Multiple view responses were appreciated by our partici-

pants, and we describe many cases where the groups of views produced from cast-a-net requests enabled tasks that spanned more than one view. Essentially, **this request style allowed participants to rapidly create groups of views that were useful together as a collection.** This also allowed participants to **rapidly utilize the available display space, at times filling the whole display in a few actions.**

In contrast, many visualization environments aim to help participants arrive at a single view, or a series of single views, that address their questions. Some systems do this through intuitive interface design, such as Tableau and its precursor Polaris [STH02]. In systems where alternate views are presented [GDA*15], often these are framed as alternate options to help users find useful single views, or as ways to accommodate ambiguity in the expression of intentions. In other cases, many views are presented to users in order to help guide a faceted exploration of visualization recommendations, as in Voyager [WMA*15]. But, since Voyager uses a bookmarking mechanism, allowing users to mark useful single views, the end goal is still framed as helping users create a set of useful single views of their data.

Flexible-canvas environments often adopt this single view framing, by focusing on how the environment accommodates the display of a trail of single views [KC17]. But, if we return to the original insight from ‘space to think’ [AEN10], the value of a flexible canvas for sensemaking was in arranging and grouping analysis artifacts around conceptual relationships and these groups of artifacts were useful when considered together. An example of an approach which does foster the creation of many views in one action can be found in Van den Elzen and Van Wijk, where small multiples are created from a target reference view in a visualization provenance trail [vdEvW13], which provides an example of the value of this kind of approach which could be applied to large displays.

We suggest that large, flexible canvas environments should target view creation interactions that create sets of views with conceptual relationships that are valuable when considered together, not just a chain of useful single views. Developing interactive interfaces that enable cast-a-net view generation would help users accomplish this and take full advantage of a large display area.

5.2. Between-View Relations and Collective and Parallel Actions

We observed that participants used collective many-to-many copy+pivot referential requests to create many views of data at once. It is possible that **when groups of views have common features, users might be inclined to act on this set collectively, rather than one at a time.** For instance, when a set of views has a common filter, it may seem intuitive to pivot this set to a new filter collectively, in one command. These relationships are often described as ‘between-view relationships’ [KC16, KC17]. Displaying these relationships is often framed as a tool for showing an analysis process or for enabling users to synthesize information across views. But, depicting these relationships may also encourage efficient many-to-many referential interactions. Further exploration of how to enable these efficient interactions, and the contexts in which users would like to act on sets of views collectively is an interesting area for future research.

5.3. View organization: Not yet realized

The mediator used flexible positioning of views to create custom groupings and arrangements that reflected the content of the views. Based on piloting this study, positioning views for the participants served to help with visual clutter and to communicate complex multi-view responses using spatial positioning. However, freely positioning and re-positioning the many views of data generated during each study posed significant challenges and was imperfectly realized. While it is clear that users of large display environments want to have control over view positioning, it is also clear that free positioning is time consuming and potentially difficult. Algorithmic positioning approaches, such as tiling, are fast but do not take into account the content of the windows being arranged, which limits their utility. It is generally assumed that the benefits of flexible, manual positioning in large, multi-view environments outweigh the costs, in time and visual clutter. Although this study does not directly challenge this assumption, we suggest that additional algorithmic view positioning tools that take into account between-view relations might make it easier to manage many views on data.

In addition, our layout approach focused on grouping and ordering views based on their content, rather than displaying views in a temporal order or by visualization provenance. We can't comment directly on the impact this had on how participants posed requests for new views, but it is plausible that seeing views grouped together may have encouraged collective actions on groups of views. Whereas organizing views more clearly around visualization provenance might encourage different requests, such as selection requests. Further research is needed to understand the impact of different view organization strategies on participant behavior.

5.4. Large displays and cast-a-net requests

Large display environments allow users to externalize their exploration process, and juxtapose many views of their data. A smaller display, with a smaller viewport, can show fewer views at high detail at once. The cast-a-net requests, which form multi-view groupings, take advantage of the space to show many views at once, and allow users to rapidly arrive at those states on the display.

In addition, in our study there were many instances where participants moved physically in the space, stepping back to view many visualizations at once, stepping forward to take note of details. When views were positioned in a grid, they could move left to right, or look up and down, to toggle between different between-view analyses. This physical navigation has been found to be beneficial for data exploration in large display environments. Since cast-a-net view creation actions allowed participants to arrive at these compositions of views on the display, it is possible that this interaction style complements large display spaces.

Finally, there is significant interest in exploring non-mouse and keyboard interaction modalities in large display environments, particularly ones that allow users to interact without being tethered to one place. This includes touch, mobile devices, smart watches, speech, mid-air gestures, and proxemics. It is possible that the cast-a-net view generation approach, through open queries around points of interest as well as referential requests, may be effectively realized with one or several of these interaction modalities. This

ought to be considered in future research on large displays and multi-modal interaction.

6. Conclusions

We contribute observations of intentions to create many views at once to accommodate tasks that span more than one view. Using a methodology in which participants explore data on a large display by directly expressing their data exploration intentions to a remote mediator, we were able to examine how participants want to explore data independent of realized visual interfaces, tools or interaction approaches. We noted that participants posed requests in ways that cast a net around sets of data attributes and subsets of the data. They accomplished this both using direct requests, and by referencing and selecting from existing views. We describe how participants used these actions to create sets of views which accommodated tasks that spanned more than a single view. The takeaways from this study are that flexible canvas systems should consider techniques to facilitate creation of many views at once for multi-view analysis tasks.

Acknowledgements

This work was supported initially by NSF award IIS 1445751 and currently by NSF award CNS 1625941. Thanks to Khairi Reda, Alberto Gonzalez and Jason Leigh for contributing in early phases of the project.

References

- [AEN10] ANDREWS C., ENDERT A., NORTH C.: Space to think: large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2010), ACM, pp. 55–64. [1](#), [2](#), [5](#), [10](#)
- [AEYN11] ANDREWS C., ENDERT A., YOST B., NORTH C.: Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *Information Visualization* 10, 4 (2011), 341–355. [1](#), [2](#)
- [AKG*a] AURISANO J., KUMAR A., GONZALES A., LEIGH J., DI EUGENIO B., JOHNSON A.: Articulate2: Toward a conversational interface for visual data exploration. [2](#)
- [AKG*b] AURISANO J., KUMAR A., GONZALES A., REDA K., LEIGH J., DI EUGENIO B., JOHNSON A.: Show me data: Observational study of a conversational interface in visual data exploration. [2](#)
- [ARJ*15] AURISANO J., REDA K., JOHNSON A., MARAI E. G., LEIGH J.: Bactogenic: a large-scale comparative genome visualization for big displays. *BMC bioinformatics* 16, 11 (2015), S6. [2](#)
- [BCC*05] BAVOIL L., CALLAHAN S. P., CROSSNO P. J., FREIRE J., SCHEIDEGGER C. E., SILVA C. T., VO H. T.: Vistrails: Enabling interactive multiple-view visualizations. In *VIS 05. IEEE Visualization, 2005.* (2005), IEEE, pp. 135–142. [1](#), [2](#)
- [BFE15] BADAM S. K., FISHER E., ELMQVIST N.: Munin: A peer-to-peer middleware for ubiquitous analytics and visualization spaces. *IEEE Transactions on Visualization and Computer Graphics* 21, 2 (2015), 215–228. [2](#)
- [BI12] BEZERIANOS A., ISENBERG P.: Perception of visual variables on tiled wall-sized displays for information visualization applications. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2516–2525. [2](#)
- [BMG10] BALLENDAT T., MARQUARDT N., GREENBERG S.: Proxemic interaction: designing for a proximity and orientation-aware environment. In *ACM International Conference on Interactive Tabletops and Surfaces* (2010), ACM, pp. 121–130. [2](#)

- [BSES17] BADAM S. K., SRINIVASAN A., ELMQVIST N., STASKO J.: Affordances of input modalities for visual data exploration in immersive environments. **1, 2**
- [CCC*15] CHANDLER T., CORDEIL M., CZAUDERNA T., DWYER T., GLOWACKI J., GONCU C., KLAPPERSTUECK M., KLEIN K., MARRIOTT K., SCHREIBER F., ET AL.: Immersive analytics. In *2015 Big Data Visual Analytics (BDVA)* (2015), IEEE, pp. 1–8. **1**
- [CHCPM07] CRESWELL J. W., HANSON W. E., CLARK PLANO V. L., MORALES A.: Qualitative research designs: Selection and implementation. *The counseling psychologist* **35**, 2 (2007), 236–264. **5**
- [DHRL*12] DUNNE C., HENRY RICHE N., LEE B., METOYER R., ROBERTSON G.: Graphtrail: Analyzing large multivariate, heterogeneous networks while supporting exploration history. In *Proceedings of the SIGCHI conference on human factors in computing systems* (2012), ACM, pp. 1663–1672. **1, 2**
- [EALN11] ENDERT A., ANDREWS C., LEE Y. H., NORTH C.: Visual encodings that support physical navigation on large displays. In *Proceedings of Graphics Interface 2011* (2011), Canadian Human-Computer Communications Society, pp. 103–110. **2**
- [GDA*15] GAO T., DONTCHEVA M., ADAR E., LIU Z., KARAHALIOS K. G.: Datatone: Managing ambiguity in natural language interfaces for data visualization. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology* (2015), ACM, pp. 489–500. **10**
- [GGL*14] GRATZL S., GEHLENBORG N., LEX A., PFISTER H., STREIT M.: Domino: Extracting, comparing, and manipulating subsets across multiple tabular datasets. *IEEE transactions on visualization and computer graphics* **20**, 12 (2014), 2023–2032. **1, 2**
- [GTS10] GRAMMEL L., TORY M., STOREY M. A.: How information visualization novices construct visualizations. *Visualization and Computer Graphics, IEEE Transactions on* **16**, 6 (2010), 943–952. **1, 2, 4**
- [HBED18] HORAK T., BADAM S. K., ELMQVIST N., DACHSELT R.: When david meets goliath: Combining smartwatches with a large vertical display for visual data exploration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018), ACM, p. 19.
- [IDW*13] ISENBERG P., DRAGICEVIC P., WILLETT W., BEZERIANOS A., FEKETE J.-D.: Hybrid-image visualization for large viewing environments. *IEEE transactions on visualization and computer graphics* **19**, 12 (2013), 2346–2355. **2**
- [JE13] JAVED W., ELMQVIST N.: Explates: spatializing interactive analysis to scaffold visual exploration. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 441–450. **1, 2, 5**
- [JH14] JAKOBSEN M. R., HORNBAEK K.: Up close and personal: Collaborative work on a high-resolution multitouch wall display. *ACM Transactions on Computer-Human Interaction (TOCHI)* **21**, 2 (2014), 11. **2**
- [JH15] JAKOBSEN M. R., HORNBAEK K.: Is moving improving?: Some effects of locomotion in wall-display interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), ACM, pp. 4169–4178. **2**
- [JHKH13] JAKOBSEN M. R., HAILE Y. S., KNUDSEN S., HORNBAEK K.: Information visualization and proxemics: Design opportunities and empirical findings. *IEEE transactions on visualization and computer graphics* **19**, 12 (2013), 2386–2395. **2**
- [KADE*16] KUMAR A., AURISANO J., DI EUGENIO B., JOHNSON A., GONZALEZ A., LEIGH J.: Towards a dialogue system that supports rich visualizations of data. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue* (2016), pp. 304–309. **2**
- [KC16] KNUDSEN S., CARPENDALE S.: View relations: An exploratory study on between-view meta-visualizations. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction* (2016), ACM, p. 15. **10**
- [KC17] KNUDSEN S., CARPENDALE S.: Multiple views in immersive analytics. **10**
- [KDEA*] KUMAR A., DI EUGENIO B., AURISANO J., JOHNSON A., ALSAIARI A., FLOWERS N., GONZALEZ A., LEIGH J.: Towards multimodal coreference resolution for exploratory data visualization dialogue: Context-based annotation and gesture identification. *SEMDIAL 2017 SaarDial*, 48. **2**
- [KJH12] KNUDSEN S., JAKOBSEN M. R., HORNBAEK K.: An exploratory study of how abundant display space may support data analysis. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design* (2012), ACM, pp. 558–567. **2, 3**
- [KPV*17] KOYTEK P., PERIN C., VERMEULEN J., ANDRÉ E., CARPENDALE S.: Mybrush: Brushing and linking with personal agency. *IEEE transactions on visualization and computer graphics* **24**, 1 (2017), 605–615. **5**
- [MAN*14] MARRINAN T., AURISANO J., NISHIMOTO A., BHARADWAJ K., MATEEVITSI V., RENAMBOT L., LONG L., JOHNSON A., LEIGH J.: Sage2: A new approach for data intensive collaboration using scalable resolution shared displays. In *Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2014 International Conference on* (2014), IEEE, pp. 177–186. **2, 4**
- [RAF*] REDA K., AURISANO J., FEBRETTI A., LEIGH J., JOHNSON A.: Visualization design patterns for ultra-resolution display environments. **2**
- [RFK*13] REDA K., FEBRETTI A., KNOLL A., AURISANO J., LEIGH J., JOHNSON A., PAPKA M. E., HERELD M.: Visualizing large, heterogeneous data in hybrid-reality environments. *IEEE Computer Graphics and Applications* **33**, 4 (2013), 38–48. **1, 2**
- [RJPL15] REDA K., JOHNSON A. E., PAPKA M. E., LEIGH J.: Effects of display size and resolution on user behavior and insight acquisition in visual exploration. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), ACM, pp. 2759–2768. **2**
- [ROJL14] REDA K., OFFORD C., JOHNSON A. E., LEIGH J.: Expanding the porthole: leveraging large, high-resolution displays in exploratory visual analysis. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems* (2014), ACM, pp. 2047–2052. **2**
- [STH02] STOLTE C., TANG D., HANRAHAN P.: Polaris: A system for query, analysis, and visualization of multidimensional relational databases. *IEEE Transactions on Visualization and Computer Graphics* **8**, 1 (2002), 52–65. **10**
- [TS19] TORY M., SETLUR V.: Do what i mean, not what i say! design considerations for supporting intent and context in analytical conversation. In *2019 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2019), IEEE, pp. 93–103. **3**
- [vdEvW13] VAN DEN ELZEN S., VAN WIJK J. J.: Small multiples, large singles: A new approach for visual data exploration. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 191–200. **2, 10**
- [WLJ*12] WALNY J., LEE B., JOHNS P., RICHE N. H., CARPENDALE S.: Understanding pen and touch interaction for data exploration on interactive whiteboards. *IEEE Transactions on Visualization and Computer Graphics* **18**, 12 (2012), 2779–2788. **3**
- [WMA*15] WONGSUPHASAWAT K., MORITZ D., ANAND A., MACKINLAY J., HOWE B., HEER J.: Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE transactions on visualization and computer graphics* **22**, 1 (2015), 649–658. **10**
- [YN06] YOST B., NORTH C.: The perceptual scalability of visualization. *IEEE Transactions on Visualization and Computer Graphics* **12**, 5 (2006), 837–844. **2**
- [YS19] YU B., SILVA C. T.: Flowsense: A natural language interface for visual data exploration within a dataflow system. *IEEE transactions on visualization and computer graphics* (2019). **1, 2**
- [ZZD14] ZGRAGGEN E., ZELEZNIK R., DRUCKER S. M.: Panoramic-data: Data analysis through pen & touch. *IEEE transactions on visualization and computer graphics* **20**, 12 (2014), 2112–2121. **1, 2**