

VisTorch: Authoring Physical Visualization Dashboards using Handheld Projectors

Biswaksen Patnaik , Huaishu Peng , and Niklas Elmqvist 




Fig. 1: **The VisTorch.** (A) Our physical implementation of the VisTorch device. (B) Using the VisTorch to uncover charts embedded into a physical dashboard in the world by projecting them onto flat surfaces with fiducial markers.

Abstract—We present VISTORCH: a handheld device for projected immersive analytics consisting of a pico-projector, a multi-spectrum digital camera, and a touch input surface. VisTorch enables viewing charts in a *physical visualization dashboard*—a 3D arrangement of linked interactive visualizations in real space—by simply pointing the device at a nearby surface, thereby revealing any data placed in that physical location. Fiducial markers in the environment enables the device to anchor charts anywhere in a room. In our user study, we asked participants to first author a physical visualization dashboard by organizing charts in 3D space and then refer to these charts to answer questions. The participants were finally asked to present the data to a collaborator. The VisTorch device threads the needle in providing situated data—which a traditional handheld display is unable to do—while enabling common ground for multiple users without requiring multiple devices—which head-mounted AR goggles do require. Results from our study showcase the utility of physical arrangements of data for memory and recall as well as the benefit of projected data displays.

Index Terms—Ubiquitous analytics, immersive analytics, augmented reality, situated visualization.

1 INTRODUCTION

A fundamental premise of ubiquitous [20], situated [21], and immersive [33] analytics is the presentation of data *in situ*; i.e., integrating visual representations of data in the real world. There are many benefits to this approach [19]: (1) it increases the display space from a small set of monitors to potentially the entire area surrounding a user; (2) it supports *situated action* [50] as well as *distributed* [28] and *embodied cognition* [47] central to human reasoning; and (3) it facilitates multiple people working together in the same physical space [29]. Furthermore, in situations when the data has a connection to the user’s physical location, it also enables *embedding* the data in a location relevant to the data [56], such as temperature near a weather station, a time schedule at a bus stop, or electricity consumption on a refrigerator. However, such ubiquitous displays are a non-trivial technical challenge, with each solution having its own drawbacks. For example, fixed displays are static and mobile devices are typically limited in size, thus limiting the display area and the potential for data embedding. Augmented Reality (AR) using head-mounted displays is nearly ideal for the purpose, but

such devices are costly and not yet widely available. Even handheld AR [4], which is trivial using current mobile devices, is troublesome because the imagery is shown on a personal screen, making establishing deixis and common ground [7] between collaborators awkward.

We present VISTORCH (Fig. 1), a custom-built handheld device combining a laser pico-projector [17] with a camera and trackpad input (Figure 1). VisTorch enables a user to shine the projector at any surface in a room to reveal any physical visualization dashboard component located on that surface. The onboard camera tracks fiducial markers placed on the surface and calculates the projector’s orientation, allowing the projected content to be corrected to avoid distortion due to skewed perspective. Since VisTorch requires physical projection surfaces, visualization components cannot be placed in mid-air, which would not have been a restriction if the technology had been built using Augmented Reality. The technique also requires an explicit action to reveal content rather than merely looking, as is the case with a head-mounted AR device. However, in exchange the VisTorch method does not require wearing a bulky (and costly) HMD, and the physical action of pointing the projector to reveal data is akin to shining a flashlight to light up a darkened room, a peephole interaction [59] that is familiar to many. Furthermore, the projected image is visible by all participants co-located in the space, facilitating synchronous collaboration. We demonstrate techniques for placing and accessing charts in the world using the device.

To validate VisTorch, we conducted a user study involving 12 participants who used the device to author and then to access a physical visualization dashboard for reading dynamic data. The study was organized into three phases: (1) participants were first asked to organize a

• All authors are with University of Maryland, College Park, MD, United States. E-mail: {bpatnaik, huaishu, elm}@umd.edu

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

dashboard in 3D space; then (2) they were given a sequence of quick lookup tasks requiring them to refer to different parts of the dashboard; and finally (3) they were asked to give an informal presentation of the dataset to the experimenter, acting as a collaborator. While we do not claim that VisTorch is superior to head-mounted or handheld AR, we do think that the device may be a useful complement to AR technology in mobile and collaborative settings.

2 BACKGROUND

Data has been increasingly integrated into our surrounding world since the early days of ubiquitous computing [54]. Despite this, sensemaking has largely been restricted to offline settings, and it is only recently that analytics has become an anytime & anywhere activity [19]. Here we review the literature on such ubiquitous, immersive, and situated analytics and then discuss specific topics within visualization dashboards, ubiquitous display environments, and handheld projectors.

2.1 Ubiquitous, Immersive, and Situated Analytics

Already in 2013, Elmqvist and Irani proposed the notion of *ubiquitous analytics* (UA) [20] that would apply ideas from ubiquitous computing to data visualization. The original concept was primarily targeted for mobile devices, but the idea was rapidly extended to mixed and augmented reality [40]. The notion of *immersive analytics* (IA) [10, 33] is explicitly based on such immersive technologies, i.e., both VR and AR. *Situated analytics* (SA) [21, 52], in contrast, is a subset of both UA and IA that concerns data that has some physical referent to the real-world location where it is displayed. Willet et al. [57] take this concept a step further by distinguishing specific representations of situated data that are *embedded* into the real world.

Of course, the field of augmented reality [1, 45] has been visualizing data integrated into the physical world for a long time (e.g. [22, 27, 55]). However, most of these representations were mostly restricted to labels, navigational cues (e.g. arrows and distances), and visual highlighting (e.g., outlining a part to be replaced or a hatch to be opened). It is only recently, with the advent of higher-order analytical tasks, that people are starting to look at full-fledged situated [8] and ubiquitous visualization [53]. Several toolkits have been proposed for this purpose, including DXR [49], which uses a grammar-based specification language, and IATK [12], which provides several specialized Unity components for immersive and situated analytics applications. However, Unity—which is common for both of these toolkits—is proprietary software managed by a single vendor. Instead, VRIA [9] suggests the use of open web-based technologies for immersive analytics; this is also the approach taken in this paper. Furthermore, the VisTorch device in our work can be seen as having flavors of all three of these types of analytics; we provide access to ubiquitous visualizations (UA) in an immersive manner (IA), albeit restricted to surfaces in the user’s environment (i.e., not mid-air displays). Furthermore, some of the data that can be presented with the device is situated in nature (SA).

2.2 Dashboards and Memory

Interactive visualization dashboards have quickly become one of the most prolific forms of visualization for many practical applications [43]. However, given their widespread adoption and refined design principles in visualization practice—to the point where leading practitioners devote entire books to the topic [23]—you may be forgiven for not immediately associating the word *dashboard* with a physical control panel mounted in the front of a vehicle or aircraft.¹ Such physical dashboards consist of both gauges and displays for monitoring the state of the vehicle, as well as levers and wheels for controlling it.

A major benefit of dashboards, be they physical or virtual in nature, is that they enable the user to rely on spatial [48] and muscle memory [32] to refer to the dashboard’s constituent parts. Just like most drivers no doubt instinctively know to regularly check their speedometer while driving their car, so do visualization dashboards enable a financial

¹Going back further, a *dashboard* was originally a piece of wood or leather mounted in the front of a horse-drawn carriage or sleigh to protect the driver and passengers from mud and dirt thrown up by the horse’s hooves.

analyst to take in the state of their stock portfolio at a glance. This idea of building on spatial (and muscle) memory has been shown to be a powerful way to organize information; for example, Scarr et al. [44] discuss how spatial memory can become an organizing principle in computer interfaces. The idea is particularly powerful for visual representations, such as in the Data Mountain [41], where bookmarks are organized in a 3D (or 2.5D) terrain. It has been advocated as one of the strengths of immersive environments; however, Liu et al. [31] recently showed that a truly immersive wrap-around view organization is not beneficial for recall and user preference over flat organizations, and that a semi-circular layout may be the best compromise.

The challenge with virtual dashboards is that many are large and therefore consume a significant part of the display needed for other tasks. In this paper, we propose the notion of a “physical” visualization dashboard spatially embedded into the real world, thereby distributing the views in the world (increasing display area) and associating the views with spatial features in the world (thus relying on spatial and muscle memory). A recent survey paper by Bach et al. [2] proposes design patterns for dashboard design; one of the categories discussed are “embedded” visualizations, but this refers to miniature panels integrated into existing apps, not in the physical world.

2.3 Ubiquitous Displays

If we could make every surface in every room a display, ubiquitous visualization would be trivial. Unfortunately, things are rarely this simple. The futuristic Office of the Future [39] from 1998 was ahead of its time. Using the notion of *spatially aware displays*, a “sea of cameras,” and ubiquitous projectors, the goal was essentially to meld a CAVE [16] with a regular office to turn virtually any surface into a display. Similarly, the Everywhere Displays Projector [38] combines a static projector with a computer-controlled rotating mirror to project imagery on any surface in a room. However, projectors have their own challenges—see below—and are not yet sufficiently ubiquitous to make every surface a display (and may never be). Several projector-based and screen-based approaches have been proposed since, but challenges such as coordination, interaction, performance, and interfacing remain [24].

One of the more obvious problems with large-scale multi-display environments is that the user’s view of different surfaces will depend on their physical position, which can affect the legibility of displays. Several approaches have been proposed to correct for the user’s dynamic perspective. The Perspective Cursor [37] adapts the mapping from motor space to display space depending on the user’s location. In E-conic [36], this idea is taken further to correct not just the cursor but the windows and graphical elements in a display environment based on the user’s dynamic position in the room. Finally, the Ubiquitous Cursor [35] uses a projector and a hemispheric mirror to project a low-resolution cursor anywhere in a physical space, correcting distortion based on room dimensions.

2.4 Projector-based Display and Interaction

Projectors have now been miniaturized to the point where they can be integrated into handheld devices. Dachselt et al. [17] examine this new generation of highly mobile *pico-projectors* and outline both existing work as well a future research agenda involving them. Similarly, Rukzio et al. [42] survey possible models for the use of pico-projectors to turn the world into a canvas for pervasive computer imagery.





Some of the early work in this space conducted design explorations before the technology even existed. Blaskó [6] simulated a wrist-worn projector display and propose several interaction techniques for its use. Hotaru [51] (“firefly”) discusses the use of a paired camera to enable touch interaction on the projected surface. Our approach in this paper couples a pico-projector with a camera to detect spatial features in the room, enabling the projected view to change dynamically based on what part of the world is within the camera’s field of view.



While many projector-based systems have a touch screen, there exist several alternate and more direct methods for interacting with the projected output. In the Wear Ur World project [34], Mistry et al. present a personal display device that hangs around the user’s neck and combines a miniature projection for augmenting the user’s surroundings

with computer-generated imagery and camera-based gesture tracking. ShadowPuppets [15] use hand shadows to allow multiple individuals to easily collaborate around a shared projected display. Willis et al. [58] show how multiple projectors can be used as pointing “beams” to support collaborators working together on the same physical surface. Similarly, Beardsley et al. [5] discuss how to use the handheld projector as a pointing device. The Skinput project [26] demonstrates the use of the human body as an input device by projecting a graphical interface on a person’s palm or forearm using a shoulder-mounted pico-projector. While our VisTorch system currently supports none of these advanced interaction modalities—the device uses a simple trackpad for input—they are intriguing possibilities for future investigation.

3 OVERVIEW: SITUATED VISUALIZATION DISPLAY PLATFORMS

A *situated visualization* [8, 55] is a visual data representation that is rendered in a physical location, often using data that is relevant to that location. Sometimes this is useful merely for the purpose of using the world as a canvas for non-situated tasks, such as writing email, editing a document, or checking social media; sometimes the tasks are contextual and location-dependent, such as navigating using a map, looking up reviews about a particular restaurant, or analyzing the traffic patterns in a busy intersection. There are several display platform that can help realize this kind of situated data representations, each with their own strengths (+) and weaknesses (–):

-  **External screens:** Fixed screens can be used to display visualizations in the world.
 - + High-resolution output and large size.
 - + Straightforward collaboration because display is shared.
 - Cannot easily be moved.
 - Current screens do not tend to support touch interaction.
 - Too costly to use everywhere beyond specialized spaces.
-  **Mobile devices:** Data can be shown on a mobile device such as a smartphone or tablet based on the device’s location [30].
 - + Handheld and flexible.
 - + Widespread adoption and (relatively) inexpensive.
 - + High-resolution screens that support touch interaction.
 - Displays are portholes into the digital rather than integrated.
 - Small screen size makes collaboration challenging.
-  **Virtual Reality (HMD):** The new generation of consumer-level Virtual Reality head-mounted displays (HMDs), such as Oculus Quest, Varjo Aero, and HP Reverb G2, enables visualizing data in immersive 3D space [3, 13, 14].
 - + Immersive 3D display environment.
 - + Data representations can be integrated into the world.
 - + Significant adoption and (relatively) inexpensive.
 - + High-resolution output and gestural interaction.
 - Not integrated with the real world.
 - Not mobile; physical navigation is limited.
 - Collaboration requires multiple devices, one per user.
-  **Augmented Reality (Handheld):** Computer-generated imagery embedded on camera view in a handheld mobile device [4].
 - + Immersive 3D display environment.
 - + Data representations can be integrated into the world.
 - + Significant adoption and (relatively) inexpensive.
 - + High resolution output and touch interaction.
 - Not directly integrated with the real world.
 - Collaboration requires multiple devices, one per user.

-  **Augmented Reality (HMD):** Augmented Reality HMDs embed digital imagery on top of the user’s view of the real world [11].
 - + Immersive 3D display environment.
 - + Data representations can be integrated into the world.
 - + High-resolution output and gestural interaction.
 - Poor adoption and (relatively) expensive.
 - Not integrated with the real world.
 - Not mobile; physical navigation is limited.
 - Collaboration requires multiple devices, one per user.
-  **Projectors (Handheld):** Mobile pico-projectors can be used to project computer-generated imagery on flat surfaces [17, 42, 46].
 - + 3D display environment.
 - + Data representations can be integrated into the world.
 - + Relatively inexpensive.
 - + Straightforward collaboration because display is shared.
 - + High-resolution output.
 - Displays limited to projection surfaces in the world.

Based on this review, we can classify these display platforms using a shared taxonomy based on the key characteristics uncovered above. Table 1 gives an overview of these device technologies.

We note that each technology has its strengths and weaknesses. In particular, HMD-based Augmented Reality is clearly the best platform for delivering situated visualization, but is still not widely available (at least partly due to high cost). This is exacerbated by the fact that collaborative data analysis, a key mechanism for many real-world data visualization tools [29], would require each analyst to have their own HMD in order to participate.

In this paper, we choose to focus on handheld projectors as an alternative technology for situated visualization. Just like other platforms, pico-projectors powered by mobile devices have their share of weaknesses: they require a projection surface, which means that mid-air immersive displays are impossible, and they also tend to rely on touch interaction on the mobile device itself. On the other hand, projectors are relatively inexpensive and they project a display that is shared by all participants unaided. Furthermore, the handheld “flashlight”-like peephole interaction [59] is familiar to many people.

4 VISTORCH: SITUATED DATA USING PICO-PROJECTORS

The VisTorch is a portable handheld projected immersive/ubiquitous analytics system that enables embedding data visualizations in physical space. It is a camera-projector system that reads fiducial markers placed in the environment and projects a perspective correct display on the surface. This provides a hand-controlled peephole interaction with data visualizations in physical space. To harness the embodied nature of the device, the interactions are based on deictic gestures.

4.1 Calibration, Tracking, and Rendering

We use ArUco markers [25] placed on flat surfaces in the environment to enable tracking the position and orientation of the device. While it is possible to use infrared ink to make markers invisible (e.g. [18]), our current implementation is based on markers visible to the naked eye.

Starting to use VisTorch in a physical marker requires a quick calibration phase, where the user pans the camera around to show all of the available markers. The device will then build an internal 3D representation of the physical space. Adding new markers to expand the available display surfaces is trivial and the calibration is fast.

Whenever the camera in the VisTorch sees a marker, it determines the position and orientation of the device in relation to the marker. If no marker is visible, the device emits a discrete blue pulsing pattern to signify that it is not currently tracking. Once the 3D position of the device is known, we calculate a perspective transform to apply to any imagery to be displayed. This makes it possible to render a distortion-free view of any data visualizations in that part of the space even if the device is held oblique to the surface.

PLATFORM	COST	AVAILABILITY	WORLD	COLLABORATION	MOBILITY	EMBEDDING
External screens	medium	widespread	porthole	shared	fixed	none
Mobile devices	medium	widespread	porthole	personal	mobile	none
Virtual Reality (HMD)	low	common	integrated	personal	room	none
Augmented Reality (handheld)	medium	widespread	porthole	personal	mobile	none
Augmented Reality (HMD)	high	rare	integrated	personal	mobile	embedded
Projectors (handheld)	medium	rare	integrated	shared	mobile	surfaces

Table 1: **Display platforms for situated visualization.** Each technology has its own strengths and weaknesses. Handheld projectors, which are the focus of this work, are still not widely used, but have many strengths for situated visualization: they facilitate collaboration and can yield embedding on real-world surfaces at a low cost.




Fig. 2: **The VisTorch.** (A) Close up view of the VisTorch. (B) Trackpad for interaction. (C) Camera for reading ArUCo markers. (D) Projector to display contents. (E) Using VisTorch to uncover data in the physical space. (F) VisTorch showing situated data about a Bluetooth speaker.

4.2 Interaction

The key VisTorch functionality is to enable placing, organizing, and viewing visualizations in space (Fig. 2, Fig. 3). A trackpad on the device facilitates interaction with displayed information.

- ✚ **Translation:** The projected contents undergo a translation transform such that they look pinned to a specific location in space. This is analogous to how things may become visible in a dark room when a flashlight is shined on them.
- ↔ **Scaling:** When the VisTorch is moved away from a projection surface while still pointed towards a marker, the projected display holds its size. Thus the scale of the displayed visualizations are held constant. This functionality is inactive when a chart is in placement view, allowing the user to intuitively control its size by moving the device closer or further away from the surface. When the chart is placed, its scale will be preserved. This embodied interaction makes it natural for a user to control not just the placement but also the size of the charts in the room.
- ☰ **User interface:** The device projects two views that can be toggled

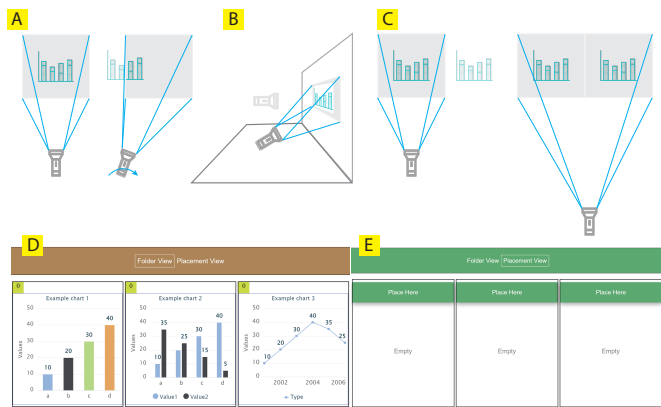


Fig. 3: **Interactions with VisTorch.** (A) Translation: Visualizations are pinned to surfaces. (B) Perspective: Oblique projection corrects perspective of display. (C) Scale: Multiple visualizations shown to obtain overview. (D) & (E) Folder view and Placement View screenshot.

between by a button: a *folder* view with a list of visualizations available to be placed in the environment and a *placement* view that holds all the visualizations that are placed around a certain marker. When pointed to a surface in the room that has a marker, the device opens up the folder view with the list of all the available visualizations available to be placed around. The placement view has empty placeholders to hold visualization that can be filled with any of the visualizations from the folder view. A visualization can be pinned to the surface by clicking the *place* button in the placement view, and can be deleted from the environment with a *delete* button. Multiple instances of a visualization can be placed in the environment; the number of instances of a visualization placed is reflected in the folder view.

- **Interacting with visualizations:** Since our VisTorch software is a browser-based solution (see below), the visualizations rendered in the space are standard HTML components. This makes it possible to use the trackpad on the VisTorch to interact directly with a visualization currently centered in the device as if using a mouse in a standard browser window.

4.3 Hardware Design

There are three primary components to the VisTorch hardware: a laser projector for displaying contents, a camera to read ArUCo markers, and a trackpad to facilitate interaction with the displayed content (Fig. 2).

- **Projector:** We use a Nebra AnyBeam laser pico-projector. The projector is focus-free as it uses MEMS based laser scanning technology to display images. We choose this projector as it is extremely portable measuring at 103mm × 60mm × 19mm, is fanless, and offers a plug in HDMI compatibility. Just weighing 133 grams and offering a 720p/60Hz resolution at 22 ANSI Lumens, the projector is perfect for handheld applications.

📷 **Camera:** For the camera, we use a Logitech C720 HD webcam.

🖱️ **Trackpad:** We use an Adafruit mini panel mount USB trackpad with trackpad surface dimensions of 60mm × 45mm.

We used T-slot aluminium extrusions from maker beam to design the frame of the VisTorch and custom-made laser-cut acrylic fixtures to attach the components together. The design has a physical separation between the camera and the projector to ensure the projected contents does not interfere with the recognition and tracking of the the ArUCo markers in the environment. The camera and the projector are vertically aligned so that the center of the camera and the projector are in line.

The current VisTorch device software runs on a Dell XPS15 Laptop that has an Intel Core i7-7700HQ CPU@2.80GHz processor and 16 GB RAM. The device is connected to the laptop through a 15ft extension cable, making it easy for the device to be moved around a room while having the laptop stationary in one corner of the room. We also experimented with an on-device computational unit, but opted for an external computer for our research prototype.

4.4 Software Architecture

The user interface of the VisTorch system is rendered in the browser and is built with standard web technologies—HTML, CSS, and JavaScript. The system uses a Python server built with Flask to perform image processing and tracking through OpenCV.²

⚙️ **Image processing:** Marker detection is done on the server. We use OpenCV’s ArUco library to detect the presence of markers in the environment. When a marker is detected on a surface, we use perspective transform to compute the transformation that helps make the projected image on the surface look perspective correct from the reference frame of the device. The actual transformation of the UI happens in the client (i.e., in the browser). We use socket communication to continuously exchange a stream of data between the server and the client. To remove any sudden changes to the projected display that may be caused by abrupt hand movements, we use a moving averages/sliding window smoothing algorithm in our image processing pipeline.

🖥️ **Display Rendering:** The UI is rendered in the browser. Based on the visible markers, the system renders selective elements through dynamic DOM manipulation. Once the rendering is complete for a certain marker, we use CSS to 3D transform the display pane correct for perspective. We can easily display information visualizations designed with HTML, CSS, and JavaScript with this pipeline. Our user study included visualizations created using Highcharts,³ but any web-based visualization library such as D3 or Vega-Lite can be used.

5 USER STUDY

We conducted an experiment to evaluate the usability of VisTorch. We see the device being used in scenarios of guided presentation where the expert guides an audience (one or more people) in obtaining an overview and/or answering questions about the data. In our study, we emulate this by having the test facilitator act as the audience and ask questions about the data. The ultimate purpose of our study is to determine if the affordance of placing visualizations and authoring ad-hoc dashboards in physical space helps create *deictic metaphors* that enhances *embodied* and *distributed* cognition.

5.1 Apparatus

We use the VisTorch system described in Section 4 to conduct our user studies. The study was conducted in a space that resembled an office setting. ArUco markers were placed in the space dividing it broadly into 5 surfaces of interaction that included 3 vertical walls and two table top surfaces (Fig. 4). VisTorch was tethered to the laptop by a 15ft extension chord and could easily be moved around in the space.

²<https://opencv.org/>

³<https://www.highcharts.com/>

#	AGE GROUP	DEGREE	EXPERTISE
♂ P1	25–30	PostDoc. (Ph.D.)	Good
♂ P2	25–30	PhD. student	Good
♂ P3	25–30	PhD. student	Good
♂ P4	30–35	PhD. student	Good
♂ P5	25–30	PhD. student	Good
♂ P6	25–30	Masters	Good
♂ P7	25–30	Masters student	Good
♂ P8	25–30	Masters student	Good
♂ P9	25–30	Masters student	Good
♂ P10	20–25	Masters student	Good
♂ P11	20–25	Masters student	Good
♂ P12	25–30	Masters student	Good

Table 2: **Participant demographics.** All participants reported having a good expertise of using data visualizations on a scale that ranged from no expertise, passing knowledge, good, and expert.

We conducted our experiment with the translation (where a visualization looks pinned to a specific point in physical space) and overview feature (moving away from the projection surface shows multiple visualizations placed across different markers) turned off. The study involved comparison tasks that needed multiple visualizations to be displayed together. Keeping the translation and overview induces cropping of images, which limits the size of the display area. This is a hardware limitation and can easily be solved with a higher resolution and a wider throw projector. Thus, in the experiment, the display was simply turned off when no marker was in the field of view of the camera.

5.2 Participants

We recruited 12 participants (6 identified as ♂ male, 6 as ♀ female) for our study (see Table 2). The age of the participants ranged from 21-35 years. Most participants were university students except one who was a working professional. We polled participants before the experiment about their expertise with with data visualizations. All participants reported their expertise as being *good* on a scale that ranged from *no experience*, *passing knowledge*, *good* and *expert*.

5.3 Experimental Factors

We involved the following factors in our experiment:

- 📄 **View Cardinality (V):** The views involved in completing a sense-making task.
 - ONE VIEW (1V) to complete a sensemaking task.
 - TWO VIEWS (2V) to complete a sensemaking task.
 - THREE VIEWS (3V) to complete a sensemaking task.
- 📍 **Data Type (T):** The type of data involved.
 - NON-SITUATED (NST): Visualization is on abstract data and its physical placement has no significance.
 - SITUATED (ST): Visualization is about a physical object or space and is placed near it.

5.4 Experimental Design

We used a within-participant factorial design where each participant participated in trials for all the view cardinalities V and data types T .

3	📄	View Cardinality V (1, 2, 3)
×	📍	Data Type T (<i>Non-situated</i> , <i>Situated</i>)
×	⋯	repetitions
18	♂♀	trials per participant

5.5 Tasks

Our experiment involved two sensemaking tasks: (i) 📄 **layout generation**, i.e. placing and arranging visualizations in physical space for an initial exploratory analysis, and (ii) 📍 **identifying data items** from

the created layout. While these were the overall tasks, we conducted the experiment in two parts by data type T : non-situated data and situated data. For each data type, a participant had to perform both the tasks: layout generation and identifying data items by number of views involved (V). The participant also had to **think out aloud** while performing the tasks so their rationale behind their actions could be noted. All the visualizations were generated with HIGHCHARTS.

Non-situated visualization (NST): For non-situated visualization, we used a Nobel Laureates multidimensional dataset. Here our goal was to determine how placing visualizations in physical space helps with embodiment.

LAYOUT GENERATION: The participants were asked to go through a list of visualizations (shown in the folder view) and organize them in space such that it made sense to them and would help them answer questions related to the visualizations. At this point there was no mention of the type of questions they would be answering. The organizational strategy for the layout was recorded.

IDENTIFYING DATA ITEMS: After the layout generation was complete, the participants were introduced to the type of questions they would be answering. The layout they had created earlier was now frozen and no changes to the layout was allowed. However, they could refer to the layout of visualizations as many times as needed to answer a question. The questions varied based on the view cardinality (V) to answer them. Three sets of questions involving one view ($1V$), two views ($2V$), and three views ($3V$) respectively were asked. Each set of questions had three repetitions. For each question, it was noted if the participant was able to retrieve the information from the layout they had created.

Situated visualization (ST): For situated visualizations, we created the scenario of a shopping experience where we placed three Bluetooth speakers on a desk. The participant played the role of the sales person and the test facilitator served as a customer who had questions about the product. Visualizations were designed with the data about the three speakers from the manufacturer and various sellers websites. The visualizations were clearly marked with the name of the speakers they showed the data about. For the situated data type, our primary goal was to have a qualitative feedback on how easy it was to create ad-hoc dashboards.

LAYOUT GENERATION: The participants were asked to go through the list of visualizations and arrange it next to the speakers. Unlike the previous case, here the participants were free to change and make new organizational layouts.

IDENTIFYING DATA ITEMS: The participants were asked to answer questions by the test facilitator that involved referring to one view ($1V$), two views ($2V$) and three views ($3V$). We recorded how new layouts of the visualizations were made to create ad-hoc dashboards to answer questions.

5.6 Procedure

Upon arriving for the study, participants were first screened based on their expertise with data visualization. All participants reported having good expertise. Then the participants were asked to give informed consent to participate in the study.

After giving their consent, the participant entered the study room where they were introduced to the study procedure. The test facilitator then demonstrated the VisTorch system, its components, and how to use it. The demonstration was made on a sample set of visualizations where the participant was shown 5 example visualizations and was asked to place them on various available display surfaces in the room. This helped them get familiar with all the basic functionalities of the system such as selecting visualizations, placing visualizations on a surface, and determining the number of instances of a visualization placed. The actual study began when the participant indicated that they were completely familiar with the functionality of the system.




Fig. 4: **Study space.** A 3D model of study space where the user studies were conducted. ArUco markers are shown in red. The number of visualizations (non-situated) placed by participants is indicated.

The study began with the non-situated data condition where the participant first explored the list of visualizations and created a layout by arranging them in physical space. The arrangement had to be such that it made it easy for them to answer questions about the visualizations efficiently. However, at this point the type of questions they would be answering was not yet explained. The test facilitator then recorded the layout generated by the participant. The next phase consisted of three sets of questions that could be answered by referring to one visualization, two visualizations, and three visualizations, respectively. Before each set of question was asked, participants were told how many visualizations they would have to access to answer the questions. Participants were given the option to change the layout before they answered each set of questions but could not make any changes while answering questions. However, they could refer to their layouts as many times as needed.

The second part of the study consisted of situated visualizations where the scenario of a shopping experience was enacted. Three Bluetooth speakers were placed and the participant was given a list of visualizations about the speakers. The participant started by organizing the visualizations about each speaker next to it. This was followed by a session where the test facilitator enacting as a potential customer asked question about the speakers. Three sets of questions were asked similar to the previous condition, questions involving one view, two views, and three views. Here the participant was free to make changes to the existing layout or make new ones to answer questions.

The participants were asked to think aloud during the entire study, and the conversations were audio and video recorded. The test facilitator also made notes on how the system was being used in-situ. The sessions were also video recorded. At the end of both the non-situated and situated phase, participants were asked to fill a NASA TLX (task load index) assessment. The average time for completion of each of these study conditions was recorded.

Once the study was completed, the participant was compensated

with a \$10 gift card. A typical session lasted between 50-60 minutes.

6 RESULTS




Fig. 5: **Participant (N=12) authored layout.** The green dots show the placement of 1/2/3 visualizations (non-situated) by markers by each participant overlaid on physical space.

Here we discuss the results from our study. We start with an overview of the results and then dive into the details by data type. For both the conditions, we collected open ended qualitative feedback on the usability of the device. We noted the average time taken to complete the tasks with each data type as well collected NASA-TLX assessment. We also report on the strategies used by the participants as they used a think aloud method to complete tasks. We collected organizational layouts in space for the non situated data. This however wasn't important in case of situated data as the data was tied to physical objects.

6.1 Non-situated Data

For the non-situated data condition, we do a detailed analysis of how the visualizations were organized in space by the participants. Then we report on the time taken to complete the task and show the results of the NASA TLX assessment.

■ **Layout generation:** Participants authored different layouts in the physical space to get an overview of the data. A custom designed heatmap of all the visualizations across all participants overlaid on the physical space is shown in Fig. 4. The total number of visualizations placed at each marker across all participants is shown in yellow circles. The green pie represents the % of all visualizations placed by surface. We see *Wall A* and *Wall B* being used the most across all participants and *Desk B* the least. When given an option to change the layout that the participants had generated, none of the participants opted to make any. All the participants chose to use the layout that they had generated in the initial exploration phase.

The layout generated by individual participants is shown in Fig. 5. Here, we observe various spatial distributions varied by participant. All participants P1-12 generated clusters of visualization based on similarities, attributes, or features that they felt to be of importance. For instance, P1 stated “*I arranged the visualizations by similarity to one another.*” Said P2, “*I made the themes and I remember where the themes were,*” “*This [pointing to wall A] was the introductory panel and this [pointing to desk B] was a geography thing.*” P3 said “*This [pointing to wall A] was categorical data and personal information.*” And P7 stated “*I had mapped the structure in my mind to the physical world,*” while P11 said “*The way I arranged it made it easy [to complete tasks].*” However, P12 mentioned they did not follow any specific strategy to arrange all the visualizations and went through all of the markers one after another to answer questions.

⌚ **Time of completion:** The time of completion of the tasks is shown in Fig. 6a. On an average participants took 21.8 minutes to complete the task. P12 took the longest as they had no specific task. P12 took the longest as they had no specific task. P12 took the longest as they had no specific task. P12 took the longest as they had no specific task.

⚖ **NASA TLX Assessment:** The NASA TLX assessment for the participants is shown in Fig. 6c. We observe that the *mental demand and effort* are rated high. From the feedback, we inferred the while the sensemaking task required mental effort, the device made it easier to accomplish it. Participants noted that the physical demand was rated high because of the weight of the device. We discuss about the weight of the device in the following sections.

6.2 Situated Data

For the situated data condition, we did not record any layout strategies as the visualizations were tied to physical artefacts. However, this condition did have comparative tasks where visualization (say device-features of a speaker) had to be compared to get an overview. In such cases, the participants created ad-hoc dashboards by putting the (device-features) visualizations together in the vicinity of the physical artefacts (speakers). For instance, all the speakers were placed on Desk A and ad-hoc dashboards were created on either on Wall A or Wall B.

💬 **Qualitative Feedback:** We asked the participants about their experience in enacting the sales person and answering questions about the products (speakers in this case). P1 mentioned “*[it] Helps build mental connection with the product*”. P2 mentioned that that was an interesting scenario and quoted “*you can see the product and the images[visualizations] together*”. P4 and P7 both mentioned that that would be helpful in a shopping experience when the customer shares the same view/display as the sales person as opposed to the customer looking into a personal tablet screen.

⌚ **Time of completion:** The time of completion of the tasks is shown in Fig. 6b. On an average participants took 6.6 mins to complete the task.

⚖ **NASA TLX Assessment:** The NASA TLX assessment for the participants is shown in Fig. 6d. Participants mentioned that physical demand was rated higher because of the weight of the system which over time got a little tiring.




Fig. 6: **Quantitative metrics from the user study.** (a) and (b) show the time taken to complete the tasks by task type. (c) and (d) show the NASA TLX assessment by task type. The legend shows the 21 point scale and the labels on the bars show the number of responses per point by color. The scale ranges from 1 (very low, yellow) to 21 (very high, pink). For *Performance* 1 (yellow) represents *perfect* and 21 (pink) represents *failure*.

6.3 Experiences with VisTorch: A Qualitative Review

Here we summarize all the experiences and open-ended feedback of all the participants based on the features of the the VisTorch device. Participants also mentioned how they would potentially use the VisTorch device and we have included those thoughts too.

- **Deixis in Guided Presentation:** P3 mentioned how the deictic affordance of the device would help them create guided presentations, quoting *"I would like to place stuff spatially on the walls and window and we can see it when I point to that"* and *"I don't have to be there physically, they[audience] could interact with the flow of information that I have thought of [by following a pre-defined path in space to discover placed contents]"*. P5 liked the affordance of pointing to reveal and mentioned *"It was easy that I point to a direction and easily see a visualization"*.
- **Spatial Arrangement and Mental Models:** Almost all participants emphasized the importance of self authoring layouts and how it helped them memorize information by creating a mental model of the layout. P1 mentioned *"I like that you can place [visualizations] in your surrounding"* and added *"[Situated visualizations] Help build mental connection with the product"*. P2 mentioned *"I made the themes [of organizing visualizations in the layout] and I also remember where the themes were"*. They also added how the experience of overlaying digital information on physical objects helps enhance the experience of making decisions about the physical products. P7 mentioned *"As long as I spend time to organize [visualizations], it's easier for me to quickly grab the information"*.
- **Authoring Ad-hoc Dashboards:** P2 mentioned *"I did like the make your own dashboard on the go"*. P3 mentioned *"I like that I can place the same visualization multiple times"* while talking about making dashboards to do comparison tasks. Both P9 and P11 mentioned that they liked the ability to make dashboards on the fly to answer questions and P11 added that such a feature helps make comparative tasks easier.
- **Direct Manipulation in Physical Space:** P3 mentioned *"The feel of this is very user friendly, I feel like I am holding something [visualization] in my hand, taking it and placing it"* and added *"I can interact with it as if it were a physical thing and change it on the fly"*. They also added that the interaction was seamless.
- **Room for Improvement:** We also collected feedback from the participants on how the device could be improved.

- **Device Ergonomics:** All participants mentioned that the device was heavy and it got a little tiring while doing tasks. Some participants opted to put the device down for a few seconds in between tasks for a break.

P2 mentioned it relatively difficult to project on horizontal surfaces as compared to vertical surfaces.

- **Resolution:** Participants P3, P7, P8 noted that while the tasks were doable with the current system, it would be more comfortable to have a projection with a higher resolution.
- **Display Stabilization:** The display took a few seconds to stabilize after it was pointed to a surface. Participants P3, P6, P8, P12 mentioned that it would improve their experience if the stabilization could be done faster.
- **User Interface:** The current system enables selecting and placing visualizations one at a time. P6 mentioned that it would be a better experience to have multiple selections enabled. When the device was held at an oblique angle, as the perspective correction kicked in, occasionally edges of the *folder view* would get cropped. Participants P1, P7, P8, P9, P11 mentioned that this would occasionally interfere with their interaction with the displayed contents.

7 DISCUSSION

We studied how the VisTorch is used in context of guided presentations for situated and non-situated data. We evaluated the affordances of the device in sensemaking and reported on the qualitative feedback from the participants. Here, we provide our interpretation of the results and describe how VisTorch can be generalized across various scenarios.

7.1 Explaining the Results

Deixis in Guided Presentation. We observed that the participants used deictic metaphors in presenting an overview of their layouts to the test facilitator. They not only used point and reveal gestures with the VisTorch to bring attention but also used deictic words such as *there & here, these & those* to refer to visualizations. Understanding such deictic metaphors was easy for anyone co-located in the space thus making it easy to narrate an overview of their layouts.

Spatial Arrangement and Mental Models. We observed almost all participants mention the importance of self-authored layouts of visualizations which helped them memorize it. Although the layouts generated by the participants differed from one another, it was interesting to see that the participants being accurate in finding the visualizations while completing tasks. In general participants preferred to organize the layouts on vertical surfaces as opposed to horizontal surfaces. This is probably because it's relatively inconvenient to aim the VisTorch at an angle onto horizontal surfaces compared to vertical surfaces. Participants also preferred to use corners in the space as we see in Fig. 5 where *Wall A* and *Wall B* were heavily used. A reasoning behind this strategy is that corners provide multiple surfaces to interact with. We also observed that the participants preferred to use the surfaces that were away from bright light sources. We acknowledge that the low brightness of the projector is a limitation of the device.

Authoring Ad-hoc Dashboards & Direct Manipulation in Physical Space. Participants mentioned that the feature of self-authoring visualization dashboards helped them custom design and arrange them in space and sometimes, make them on the fly. This paired with the metaphor of manipulating visualizations by pointing and interacting made their experience seamless as if it were a physical artefact. We believe that the features of the VisTorch device help provide this experience. VisTorch is portable, handheld, and affords overlaying digital information on physical space. It enables interacting with information by direct manipulation of the projected display. This helps create a virtual infinite dashboard that can be overlaid on the physical surrounding where objects of everyday life act as anchors to hold information.

7.2 Generalizing the Results

We have evaluated the affordances of VisTorch and its performance for sensemaking. Here we generalize the results of our study in terms of *users*, *context* of use, *environment* of use and the displayed *contents*.

Users: Though most of the participants of our study were university students, we can see the device being used by anyone who has a basic understanding of visualizations and computers. The participants have described the device as P3: "user friendly" and P6: "easy to use and straightforward".

Context: We evaluated VisTorch in the context of a guided presentation where an expert authors physical visualization dashboards and guides an audience. We observed that for both situated and non situated data, the affordance of anchoring data into the physical space helps create a custom infinite physical dashboard of visualizations. We see this being used in scenarios of collaborative sensemaking. We have described potential scenarios of use in the following section.

Environment: In our study, we evaluated our prototype in a laboratory setting where VisTorch seemed to work best. We acknowledge that the device has limited brightness for its use in outdoor settings or rooms that are brightly lit. However the brightness of the projected display can be improved by moving the device closer to the projection area. Given these conditions, we see the device easily being used in indoor spaces such as offices, classrooms, retail stores, warehouses etc. or outdoors at night with limited lighting. We used fiducial markers printed off an office printer to tag our study space. Thus any of the spaces we discussed can easily be augmented by a pack of these printed tags by tapping them to various surfaces. The current system can entirely be powered by the USB ports of the laptop it's tethered to thus making it portable and easy to use.

Contents: The VisTorch uses web technologies - HTML, CSS and JavaScript to render contents. This pipeline makes it extremely convenient to custom design visualizations for the system as any web-based visualization library such as D3, Vega-Lite, HighCharts etc. can be used. Any interactivity programmed in the visualizations such as panning, zooming, brushing, tooltips, transitions etc. are supported by the system and the on-board trackpad makes it easy to interact with visualizations in-situ.

7.3 Potential Usage Scenarios

We see VisTorch used in scenarios of guided presentation of visualizations where an expert guides an audience.

Let's imagine the experience of purchasing a product, for instance, a car. Our decisions of buying a car are shaped by various factors - the research we do online, the countless reviews we watch, the personal preferences we have, (sometimes) the recommendations of our friends or family and, the convincing information that the salesperson presents. In such a scenario, we envision VisTorch to knit this fabric of scattered information and present it in front of us in context of the product, here, the car. It will help overlay all the infographics and features next to product, the car, and this display is shared by the salesperson and the customers. Quick comparisons across products can be made in-situ for better decision making.

In a different scenario of a classroom, let's imagine how the audience refers to the visualizations in the presentation when asking a question. Often times it's the slide number or the title of the visualization or a vague description of it. We envision a potential use of VisTorch here where a copy of the visualizations from the presented slides is overlaid in the physical space in clear sight of the audience. In such a scenario, if one were to refer to a visualization they could easily refer to it by the physical artifact in the surrounding it's tagged to, for instance "Could you please pull up the visualization that you put next to the window and show me how it compares to the one you placed next to the door?" Here we see the device being used to augment traditional display systems in collaborative sensemaking.

7.4 Limitations and Future Work

While VisTorch has all the affordances needed to author physical visualization dashboards in the real world, the implementation still has room for improvement.

Participants from our study felt that the device is heavy. This is due to the fact that we used T-slot aluminium frames for the device framework, which are maker-friendly for quick prototypes but have the downside of being heavy. Also, the prototype is not weight-balanced for its context of use. We can overcome this by designing an enclosure with a material of suitable weight class and balancing the weight of the device by carefully placing the internal components.

A few participants mentioned improving the resolution of the device as well as the occasional cropping of the user interface when the projector is held at oblique angles to a surface. This is doable by switching to a higher resolution and a wider throw-angle projector. This will not only improve the resolution of the projected display, but also reduce cropping. Although the resolution of the current projector is not the best, it does offer other advantages such as being focus free, low power, lightweight, and extremely portable. At the time of the design the device, this was the best option available on the market.

We stabilize the image of the display and smooth out any small movements by employing a moving averages based stabilization technique. This induces a small latency where it takes a few seconds to stabilize the image after a sudden movement is made. This latency can be improved by parallelizing the stabilization in image processing.

Finally, although the current system is tethered, we envision a future implementation to be wireless with an onboard computer. We also envision the implementation of the device to incorporate an IR camera where the fiducial markers in the environment are painted with IR ink, making them invisible to the human eye but visible to the camera.

8 CONCLUSION

We have presented VisTorch, a handheld device for projected immersive analytics in physical space. Our work started with an analysis of the design space of multiple situated visualization display platforms by their strengths and weaknesses. We showed the features of the system and how it affords flashlight-like interaction overlaying digital information on physical space. We evaluated these features in context of analytical sensemaking in an informal user study. A generalization of our results shows the utility of the device as a projected data display.

ACKNOWLEDGMENTS

This work was partly supported by grant IIS-1901485 from the U.S. National Science Foundation. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the funding agency.

REFERENCES

- [1] R. T. Azuma. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, 6(4):355–385, 1997. doi: 10.1162/pres.1997.6.4.355 2
- [2] B. Bach, E. Freeman, A. Abdul-Rahman, C. Turkay, S. Khan, Y. Fan, and M. Chen. Dashboard design patterns. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):342–352, 2023. doi: 10.1109/TVCG.2022.3209448 2

- [3] A. Batch, A. Cunningham, M. Cordeil, N. Elmqvist, T. Dwyer, B. H. Thomas, and K. Marriott. There is no spoon: Evaluating performance, space use, and presence with expert domain users in immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):536–546, 2020. doi: 10.1109/TVCG.2019.2934803 3
- [4] A. Batch, S. Shin, J. Liu, P. W. S. Butcher, P. D. Ritsos, and N. Elmqvist. Evaluating view management for situated visualization in web-based handheld AR. *Computer Graphics Forum*, 2023. to appear. 1, 3
- [5] P. A. Beardsley, J. van Baar, R. Raskar, and C. Forlines. Interaction using a handheld projector. *IEEE Computer Graphics and Applications*, 25(1):39–43, 2005. doi: 10.1109/MCG.2005.12 3
- [6] G. Blaskó, S. Feiner, and F. Coriand. Exploring interaction with a simulated wrist-worn projection display. In *Proceedings of the IEEE International Symposium on Wearable Computers*, pp. 2–9. IEEE, Piscataway, NJ, USA, 2005. doi: 10.1109/ISWC.2005.21 2
- [7] S. E. Brennan, K. Mueller, G. Zelinsky, I. Ramakrishnan, D. S. Warren, and A. Kaufman. Toward a multi-analyst, collaborative framework for visual analytics. In *Proceedings of the IEEE Symposium on Visual Analytics Science & Technology*, pp. 129–136, 2006. 1
- [8] N. Bressa, H. Korsgaard, A. Tabard, S. Houben, and J. Vermeulen. What’s the situation with situated visualization? A survey and perspectives on situatedness. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):107–117, 2022. doi: 10.1109/TVCG.2021.3114835 2, 3
- [9] P. W. S. Butcher, N. W. John, and P. D. Ritsos. VRIA: A web-based framework for creating immersive analytics experiences. *IEEE Transactions on Visualization and Computer Graphics*, 27(7):3213–3225, 2021. doi: 10.1109/TVCG.2020.2965109 2
- [10] T. Chandler, M. Cordeil, T. Czauderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, F. Schreiber, and E. Wilson. Immersive analytics. In *Proceedings of the International Symposium on Big Data Visual Analytics*, pp. 1–8. IEEE, Piscataway, NJ, USA, 2015. doi: 10.1109/BDVA.2015.7314296 2
- [11] Z. Chen, Y. Su, Y. Wang, Q. Wang, H. Qu, and Y. Wu. MARVisT: authoring glyph-based visualization in mobile augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 26(8):2645–2658, 2019. doi: 10.1109/TVCG.2019.2892415 3
- [12] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer. IATK: An immersive analytics toolkit. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*, pp. 200–209. IEEE, Piscataway, NJ, USA, 2019. doi: 10.1109/VR.2019.8797978 2
- [13] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott. ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 71–83. ACM, New York, NY, USA, 2017. doi: 10.1145/3126594.3126613 3
- [14] M. Cordeil, T. Dwyer, K. Klein, B. Laha, K. Marriott, and B. H. Thomas. Immersive collaborative analysis of network connectivity: CAVE-style or head-mounted display? *IEEE Transactions on Visualization and Computer Graphics*, 23(1):441–450, 2017. doi: 10.1109/TVCG.2016.2599107 3
- [15] L. G. Cowan and K. A. Li. ShadowPuppets: supporting collocated interaction with mobile projector phones using hand shadows. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 2707–2716. ACM, New York, NY, USA, 2011. doi: 10.1145/1978942.1979340 3
- [16] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: the design and implementation of the CAVE. In *Proceedings of the ACM Conference on Computer Graphics and Interactive Techniques*, pp. 135–142. ACM, New York, NY, USA, 1993. doi: 10.1145/166117.166134 2
- [17] R. Dachsel, J. Häkkinen, M. Jones, M. Löchtefeld, M. Rohs, and E. Rukzio. Pico projectors: firefly or bright future? *Interactions*, 19(2):24–29, 2012. doi: 10.1145/2090150.2090158 1, 2, 3
- [18] M. D. Dogan, A. Taka, M. Lu, Y. Zhu, A. Kumar, A. Gupta, and S. Müller. InfraredTags: embedding invisible AR markers and barcodes using low-cost, infrared-based 3d printing and imaging tools. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 269:1–269:12. ACM, New York, NY, USA, 2022. doi: 10.1145/3491102.3501951 3
- [19] N. Elmqvist. Anywhere & everywhere: A mobile, immersive, and ubiquitous vision for data analytics. *Communications of the ACM*, 2023. to appear. 1, 2
- [20] N. Elmqvist and P. Irani. Ubiquitous analytics: Interacting with big data anywhere, anytime. *IEEE Computer*, 46(4):86–89, 2013. doi: 10.1109/mc.2013.147 1, 2
- [21] N. A. M. ElSayed, B. H. Thomas, K. Marriott, J. Piantadosi, and R. T. Smith. Situated analytics: Demonstrating immersive analytical tools with augmented reality. *Journal of Visual Languages & Computing*, 36:13–23, 2016. doi: 10.1016/j.jvlc.2016.07.006 1, 2
- [22] S. Feiner, B. MacIntyre, T. Höllerer, and A. Webster. A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment. *Personal Technologies*, 1(4):208–217, 1997. 2
- [23] S. Few. *Information Dashboard Design: The Effective Visual Communication of Data*. O’Reilly Press, Sebastopol, CA, USA, 2006. 2
- [24] F. Garcia-Sanjuan, J. Jaen, and V. Nacher. Toward a general conceptualization of multi-display environments. *Frontiers in ICT*, 3:20, 2016. doi: 10.3389/fitc.2016.00020 2
- [25] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and M. J. Marín-Jiménez. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, 47(6):2280–2292, 2014. doi: 10.1016/j.patcog.2014.01.005 3
- [26] C. Harrison, D. S. Tan, and D. Morris. Skinput: appropriating the body as an input surface. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 453–462. ACM, New York, NY, USA, 2010. doi: 10.1145/1753326.1753394 3
- [27] S. Henderson and S. Feiner. Exploring the benefits of augmented reality documentation for maintenance and repair. *IEEE Transactions on Visualization and Computer Graphics*, 17(10):1355–1368, 2011. doi: 10.1109/TVCG.2010.245 2
- [28] E. Hutchins. *Cognition in the Wild*. MIT Press, Cambridge, MA, USA, 1995. 1
- [29] P. Isenberg, N. Elmqvist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Hagen. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization*, 10(4):310–326, 2011. doi: 10.1177/1473871611412817 1, 3
- [30] B. Lee, R. Dachsel, P. Isenberg, and E. K. Choe. *Mobile Data Visualization*. AK Peters Visualization Series. Chapman and Hall/CRC Press, Boca Raton, FL, USA, 2022. 3
- [31] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer. Effects of display layout on spatial memory for immersive environments. *Proceedings of the ACM on Human-Computer Interaction*, 6(ISS):468–488, 2022. doi: 10.1145/3567729 2
- [32] Q. Liu and E. Jorgensen. Muscle memory. *Journal of Physiology*, 89(Pt 4):775–776, 2011. doi: 10.1113/jphysiol.2011.205088 2
- [33] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, eds. *Immersive Analytics*, vol. 11190 of *Lecture Notes in Computer Science*. Springer International Publishing, New York, NY, USA, 2018. doi: 10.1007/978-3-030-01388-2 1, 2
- [34] P. Mistry, P. Maes, and L. Chang. WUW - Wear Ur World: a wearable gestural interface. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, pp. 4111–4116. ACM, New York, NY, USA, 2009. doi: 10.1145/1520340.1520626 2
- [35] M. A. Nacenta, R. L. Mandryk, and C. Gutwin. Targeting across display-less space. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 777–786, 2008. 2
- [36] M. A. Nacenta, S. Sakurai, T. Yamaguchi, Y. Miki, Y. Itoh, Y. Kitamura, S. Subramanian, and C. Gutwin. E-conic: a perspective-aware interface for multi-display environments. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 279–288. ACM, New York, NY, USA, 2007. doi: 10.1145/1294211.1294260 2
- [37] M. A. Nacenta, S. Sallam, B. Champoux, S. Subramanian, and C. Gutwin. Perspective cursor: perspective-based interaction for multi-display environments. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 289–298. ACM, New York, NY, USA, 2006. doi: 10.1145/1124772.1124817 2
- [38] C. S. Pinhanez. The everywhere displays projector: A device to create ubiquitous graphical interfaces. In *Proceedings of the International Conference on Ubiquitous Computing*, vol. 2201 of *Lecture Notes in Computer Science*, pp. 315–331. Springer, 2001. doi: 10.1007/3-540-45427-6_27 2
- [39] R. Raskar, G. Welch, M. Cutts, A. T. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image-based modeling and spatially immersive displays. In *Proceedings of the ACM Conference on Computer Graphics and Interactive Techniques*, pp. 179–188. ACM, New York, NY, USA, 1998. doi: 10.1145/280814.280861 2
- [40] J. Roberts, P. Ritsos, S. K. Badam, D. Brodbeck, J. Kennedy, and N. Elmqvist. Visualization beyond the desktop—the next big thing.

34(6):26–34, Nov 2014. doi: 10.1109/MCG.2014.82 2

- [41] G. Robertson, M. Czerwinski, K. Larson, D. C. Robbins, D. Thiel, and M. van Dantzich. Data mountain: Using spatial memory for document management. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 153–162, 1998. doi: 10.1145/288392.288596 2
- [42] E. Rukzio, P. Holleis, and H. Gellersen. Personal projectors for pervasive computing. *IEEE Pervasive Computing*, 11(2):30–37, 2012. doi: 10.1109/MPRV.2011.17 2, 3
- [43] A. Sarikaya, M. Correll, L. Bartram, M. Tory, and D. Fisher. What do we talk about when we talk about dashboards? *IEEE Transactions on Visualization and Computer Graphics*, 25(1):682–692, 2019. doi: 10.1109/TVCG.2018.2864903 2
- [44] J. Scarr, A. Cockburn, and C. Gutwin. Supporting and exploiting spatial memory in user interfaces. *Foundations and Trends in Human-Computer Interaction*, 6(1):1–84, 2013. doi: 10.1561/1100000046 2
- [45] D. Schmalstieg and T. Höllerer. *Augmented Reality: Principles and Practice*. Addison-Wesley, Boston, MA, USA, 2016. 2
- [46] J. Schönig, M. Löchtefeld, M. Rohs, and A. Krüger. Projector phones: A new class of interfaces for augmented reality. *International Journal of Mobile Human-Computer Interaction*, 2(3):1–14, 2010. doi: 10.4018/jmhci.2010070101 3
- [47] L. A. Shapiro. *Embodied Cognition*. Routledge, New York, NY, USA, 2011. 1
- [48] Y. Shrager, P. J. Bayley, B. Bontempi, R. O. Hopkins, , and L. R. Squire. Spatial memory and the human hippocampus. *Proceedings of the National Academy of Sciences*, 104(8):2961–2966, 2007. doi: 10.1073/pnas.0611233104 2
- [49] R. Sicat, J. Li, J. Choi, M. Cordeil, W.-K. Jeong, B. Bach, and H. Pfister. DXR: A toolkit for building immersive data visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):715–725, 2019. doi: 10.1109/TVCG.2018.2865152 2
- [50] L. A. Suchman. *Plans and Situated Actions — The Problem of Human-Machine Communication*. Cambridge University Press, Cambridge, UK, 1987. 1
- [51] M. Sugimoto, K. Miyahara, H. Inoue, and Y. Tsunesada. *Hotaru*: Intuitive manipulation techniques for projected displays of mobile devices. In *Proceedings of the IFIP TC13 Conference on Human-Computer Interaction*, vol. 3585 of *Lecture Notes in Computer Science*, pp. 57–68. Springer, 2005. doi: 10.1007/11555261_8 2
- [52] B. H. Thomas, G. F. Welch, P. Dragicevic, N. Elmqvist, P. Irani, Y. Jansen, D. Schmalstieg, A. Tabard, N. A. M. ElSayed, R. T. Smith, and W. Willett. Situated analytics. In *Immersive Analytics*, vol. 11190 of *Lecture Notes in Computer Science*, pp. 185–220. Springer International Publishing, New York, NY, USA, 2018. doi: 10.1007/978-3-030-01388-2_7 2
- [53] J. Vermeulen, C. Collins, R. Dachselt, P. Irani, and A. Joshi. Reflections on ubiquitous visualization. In *Mobile Data Visualization*. Chapman and Hall/CRC Press, Boca Raton, FL, USA, 2021. 2
- [54] M. Weiser. The computer for the 21st Century. *Scientific American*, 265(3):94–104, 1991. doi: 10.1145/329124.329126 2
- [55] S. White and S. Feiner. SiteLens: Situated visualization techniques for urban site visits. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 1117–1120. ACM, New York, NY, USA, 2009. doi: 10.1145/1518701.1518871 2, 3
- [56] W. Willett, Y. Jansen, and P. Dragicevic. Embedded data representations. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):461–470, 2017. doi: 10.1109/TVCG.2016.2598608 1
- [57] W. Willett, Y. Jansen, and P. Dragicevic. Embedded Data Representations. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):461–470, Jan. 2017. doi: 10.1109/TVCG.2016.2598608 2
- [58] K. D. D. Willis, I. Poupyrev, S. E. Hudson, and M. Mahler. SideBySide: ad-hoc multi-user interaction with handheld projectors. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 431–440. ACM, New York, NY, USA, 2011. doi: 10.1145/2047196.2047254 3
- [59] K. Yee. Peephole displays: pen interaction on spatially aware handheld computers. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 1–8. ACM, New York, NY, USA, 2003. doi: 10.1145/642611.642613 1, 3