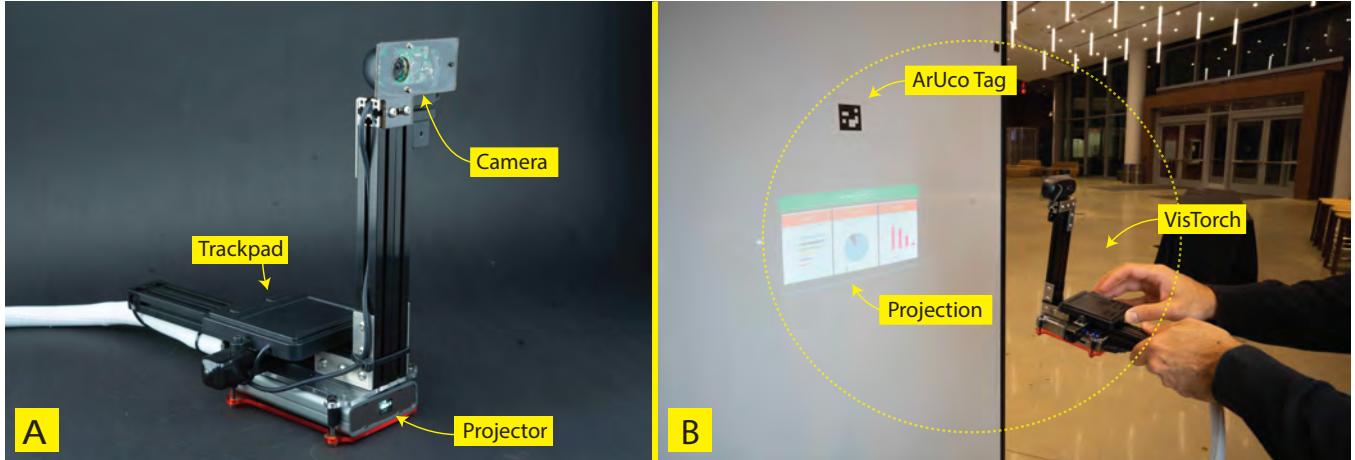


# VisTorch: Interacting with Situated Visualizations using Handheld Projectors

Biswaksen Patnaik  
University of Maryland  
College Park, Maryland, USA  
bpatnaik@umd.edu

Huaishu Peng  
University of Maryland  
College Park, Maryland, USA  
huaishu@umd.edu

Niklas Elmqvist  
Aarhus University  
Aarhus, Denmark  
elm@cs.au.dk



**Figure 1: THE VisTorch.** The device enables accessing situated visualizations by pointing a tracked projector at a physical surface, similar to shining a flashlight in a dark room. (A) Our implementation of the VisTorch device. (B) Using the VisTorch to uncover charts embedded into a situated dashboard in the world by projecting them onto surfaces tagged with fiducial markers.

## ABSTRACT

Spatial data is best analyzed *in situ*, but existing mixed reality technologies can be bulky, expensive, or unsuitable for collaboration. We present VisTorch: a handheld device for projected situated analytics consisting of a pico-projector, a multi-spectrum camera, and a touch surface. VisTorch enables viewing charts situated in physical space by simply pointing the device at a surface to reveal visualizations in that location. We evaluated the approach using both a user study and an expert review. In the former, we asked 20 participants to first organize charts in space and then refer to these charts to answer questions. We observed three spatial and one temporal pattern in participant analyses. In the latter, four experts—a museum designer, a statistical software developer, a theater stage designer, and an environmental educator—utilized VisTorch to derive practical usage scenarios. Results from our study showcase the utility of situated visualizations for memory and recall.

## CCS CONCEPTS

- Human-centered computing → Displays and imagers; Visualization systems and tools; Visualization application domains.

## KEYWORDS

Ubiquitous analytics, immersive analytics, augmented reality, situated visualization.

### ACM Reference Format:

Biswaksen Patnaik, Huaishu Peng, and Niklas Elmqvist. 2024. VisTorch: Interacting with Situated Visualizations using Handheld Projectors. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24), May 11–16, 2024, Honolulu, HI, USA*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3613904.3642857>

## 1 INTRODUCTION

A fundamental premise of ubiquitous [17], situated [18], and immersive [35] 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 [16]: (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* [52] as well as *distributed* [24] and *embodied cognition* [48] central to human reasoning; and (3) it facilitates multiple people working together in the same physical space [25]. Furthermore, in situations when the data has a connection to the user’s physical location, it also enables *embedding* the



This work is licensed under a Creative Commons Attribution International 4.0 License.

data in a location relevant to the data [58], such as temperature near a weather station, a time schedule at a bus stop, or electricity consumption on a refrigerator. However, implementing such ubiquitous displays are a non-trivial technical challenge, with each solution having its own drawbacks. For example, fixed displays are static, whereas 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, cumbersome, 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 between collaborators awkward.

We present VisTorch (Fig. 1), a custom-built handheld device combining a laser pico-projector [13] with a camera and trackpad input (Figure 1). VisTorch lets a user to shine the projector at any surface in a room to reveal any situated visualization 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 VisTorch 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 in a darkened room, a peephole interaction [61] familiar to many. Furthermore, the projected image is visible by all participants co-located in the physical space, facilitating collaboration.

We validated VisTorch using both a user study and an expert review. The former user study involved 20 participants who used the device to author and then access situated visualizations for reading dynamic data. The study was organized into three phases: (1) participants were first asked to organize charts in 3D space; then (2) they were given a sequence of quick lookup tasks requiring them to refer to different charts; and finally (3) they were asked to give an informal presentation of the dataset to the experimenter, acting as a collaborator. The latter expert review engaged four experts in a think-aloud protocol while using VisTorch for an in-depth review involving their professional expertise: theater stage production, statistical data analysis, environmental education, and museum planning. Findings from both forms of validation showcase people's intelligent use of space [30] for data analysis and support our hypothesis that handheld projection for situated visualization can be a useful approach to data analysis in mobile settings.

## 2 BACKGROUND

Data is increasingly being integrated into our surrounding world since the early days of ubiquitous computing [56]. Despite this, it is only recently that data analytics and sensemaking has become an anytime and anywhere activity [16]. 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/Situated Analytics

In 2013, Elmqvist and Irani proposed the idea of *ubiquitous analytics* (UA) [17] that would apply ideas from ubiquitous computing to data visualization. The original concept was primarily targeted for mobile devices and physical displays and tabletops [29], but the idea was rapidly extended to mixed and augmented reality [41]. *Immersive analytics* (IA) [9, 35] is explicitly based on such immersive technologies. *Situated analytics* (SA) [18, 49, 54] is a subset of UA/IA that concerns data that has some physical referent [58] to the real-world location where it is displayed. Willett et al. [58] take this concept further by highlighting representations of situated data that are *embedded* into the real world.

Of course, the field of augmented reality has been visualizing data integrated into the physical world for a long time (e.g. [2, 19, 22, 57]). However, most of these representations were 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 that people are starting to integrate full-fledged situated [6] and ubiquitous visualization [55] as well as situated analytics [49] in AR. Several toolkits have been proposed for this purpose, including DXR [51], which uses a grammar-based specification language, and IATK [11], which provides several specialized Unity components for immersive and situated analytics applications. However, Unity—common for both of these toolkits—is proprietary software managed by a single vendor. Instead, VRIA [7] suggests the use of open web-based technologies for UA/IA; this is also our approach here. Furthermore, the VisTorch device in our work can be seen as having flavors of all three 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, the mobile form factor means that the device can be used to present situated data (SA).

### 2.2 Dashboards and Memory

Visualization dashboards have quickly become a prolific form of visualization in many disciplines [44]. A major benefit of dashboards is that they enable the user to rely on spatial [50] and muscle memory [33] to refer to the dashboard's constituent parts.

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. [46] 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 [42], where bookmarks are organized in a 3D (or 2.5D) terrain. Wright et al. [60] report on intelligence analysts using physical space to arrange data, and Andrews and North [1] famously showed how ample display space can be used to facilitate in-depth analysis. In particular, embodied human-data interaction [15] leverages spatial memory and embodiment for interacting with visualization. Physical data arrangements have been advocated as one of the strengths of immersive environments; however, Liu et al. [32] 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.

### 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 [40] 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 with a regular office to turn virtually any surface into a display. A more recent example with the same goal is the Microsoft RoomAlive project [28]. Similarly, the Everywhere Displays Projector [39] 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 [20].

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 [38] adapts the mapping from motor space to display space depending on the user’s location. In E-conic [37], 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 [36] 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 Displays and Interaction

Projectors have now been miniaturized to the point where they can be integrated into handheld devices. Dachselt et al. [13] examine this new generation of highly mobile *pico-projectors* and outline both existing work as well as a future research agenda. Similarly, Rukzio et al. [43] 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 [5] simulated a wrist-worn projector display and propose several interaction techniques for its use. Hotaru [53] (“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, enabling the projected view to change dynamically based on what part of the world is seen by the camera.

Finally, most closely related to our work is HideOut [59], which uses handheld projectors to display digital content on real-world objects using infrared fiducial marker tracked by a camera. Compared to HideOut, our focus in this paper is exclusively on the use of such handheld projection for authoring and manipulating situated visualizations, and in our findings of how people arrange data in space to facilitate analysis [1, 30]. The AR Magic Lantern [23] is based on a similar handheld display unit combining SLAM-tracking using an Apple iPhone and a pico projector, but its focus is not specifically on visual data analysis.

## 3 OVERVIEW: SITUATED VISUALIZATION DISPLAY PLATFORMS

A *situated visualization* [6] is a visual data representation that is rendered in a physical 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 location-dependent, such as navigating, looking up reviews about a restaurant, or analyzing the traffic patterns in a busy intersection. There are several display platforms 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.
- **Mobile devices:** Data can be shown on a mobile device based on the device’s location [31].
- **Virtual Reality (HMD):** Virtual Reality head-mounted displays render data in immersive 3D space [12].
- **Augmented Reality (Handheld):** Data embedded in a camera view on a handheld mobile device [4].
- **Augmented Reality (HMD):** Augmented Reality HMDs embed digital imagery on top of the real world [10].
- **Projectors (Fixed):** Multiple fixed projectors turning a physical space into a display environment [28].
- **Projectors (Handheld):** Mobile pico-projectors rendering data displays on flat surfaces [13, 23, 43, 47, 59].

Table 1 presents a classification of these display platforms based on their characteristics. 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 [25], would require each analyst to have their own HMD in order to participate.

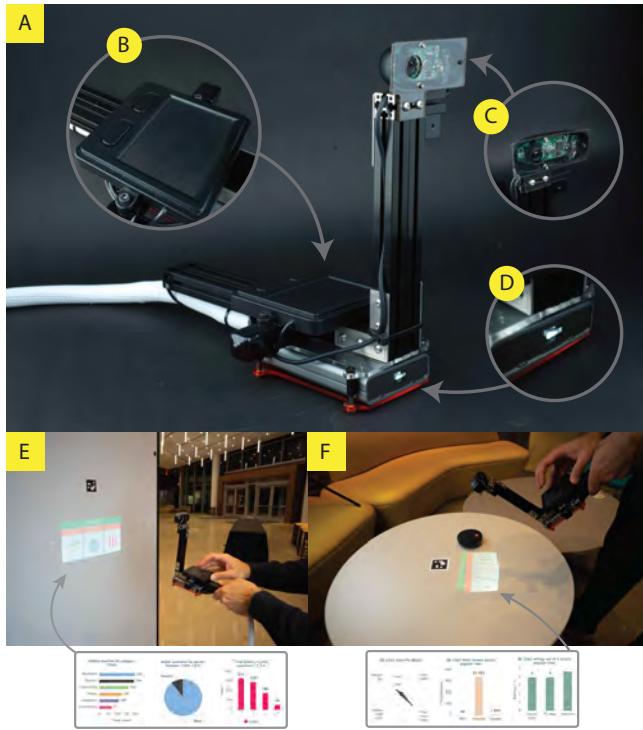
In this paper, we choose to focus on handheld projectors as an alternative technology for situated visualization. Of course, pico-projectors powered by mobile devices have their share of weaknesses: they require a projection surface, which means that mid-air immersive displays are impossible, and they 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 can be seen by all participants. Furthermore, a “flashlight”-like peephole interaction [61] is familiar to many people.

## 4 VISTORCH: SITUATED DATA USING PICO-PROJECTORS

The VisTorch is a portable handheld projected 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.

PLATFORM	COST	SUPPLY	WORLD	SHARING	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 (fixed)	medium	rare	integrated	shared	fixed	surfaces
Projectors (handheld)	medium	rare	integrated	shared	mobile	surfaces

**Table 1: DISPLAY PLATFORMS FOR SITUATED VISUALIZATION.** Handheld projectors, the focus of this work, are still not widely used, but have many strengths for situated visualization: they facilitate collaboration and can embed displays on physical surfaces at a low cost.



**Figure 2: ANATOMY OF 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.

#### 4.1 Calibration, Tracking, and Rendering

We use ArUco markers [21] 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. [14]), 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

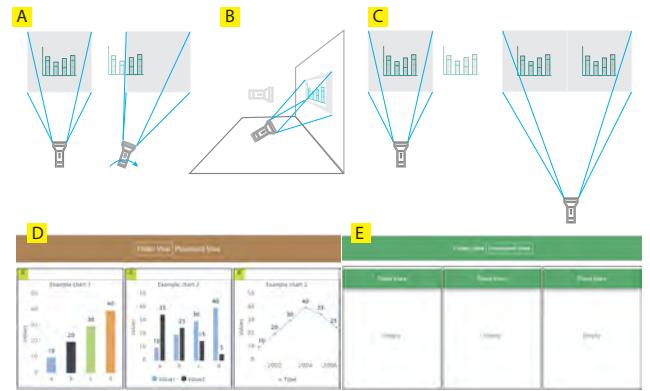
3D representation of the physical space. Adding new markers to expand the display space is trivial and calibration is fast.

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 displayed imagery. 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.

Small movements induced due to unstable hand movement causes fast transformations in the image, making displays jittery and difficult to read. To avoid this, we employ a moving averages based smoothing technique to stabilize the image. Although this makes the display readable, it does yield some latency in responsiveness.

#### 4.2 Interaction

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



**Figure 3: INTERACTIONS WITH VISTORCH.** (A) Translation: Visualizations are pinned to surfaces. (B) Perspective: Oblique projection corrects display perspective. (C) Scale: Multiple visualizations show overview. (D) & (E) Folder view and Placement View screenshots.

❖ **Translation:** The projected contents undergo a translation transform such that they look pinned to a specific location in space. This is analogous to how objects in darkness become visible when illuminated by a flashlight.

↔ **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 is 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 is saved. This embodied interaction makes it natural to control not just the placement but also the size of a chart.

☰ **User interface:** The device projects two views that can be toggled between using 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 visualizations available for placement. The placement view has empty placeholders for holding visualizations drawn 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 placed is reflected in the folder view.

► **Interacting with visualizations:** Since our VisTorch software is browser-based, the visualizations 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 visual contents, a camera to read ArUco markers, and a trackpad to facilitate interaction (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 chose this projector because it is extremely portable (133g and measuring 103mm × 60mm × 19mm), is fanless, and offers a plug in HDMI compatibility with 720p/60Hz resolution at 22 ANSI Lumens.

○ **Camera:** Logitech C720 HD webcam.

► **Trackpad:** Adafruit mini panel mount USB trackpad with trackpad surface dimensions of 60mm × 45mm.

We used off-the-shelf T-slot aluminium extrusions from MakerBeam 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 do not interfere with the recognition and tracking of the ArUco markers in the environment. This avoids any errors in reading the fiducial markers by the camera due to an overlap of the projected contents on the marker when projected from about 3 feet from the surface. An alternate design

would combine camera and projector and use a frame sync. 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 an extension cable (about 4.5m long), 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.

The overall cost of the VisTorch—not counting the laptop—is less than \$400, with the projector being the largest expense. Adding on-board computing would still yield a cost well below \$500, putting the device at significantly lower cost than current-generation AR HMDs (\$3,500 for a Microsoft HoloLens 2 and an Apple Vision Pro).

### 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 with Flask to perform image processing through OpenCV.<sup>1</sup>

❖ **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 compute a reverse perspective transform to make the projected image on the surface look perspective correct from the reference frame of the device. 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 caused by abrupt hand movements, we use a moving averages smoothing algorithm in our image processing pipeline.

● **Display Rendering:** The UI is rendered in the browser. The system renders selected elements based on visible markers 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. This also allows using the VisTorch for any existing web-based visualization. Our user study included visualizations created using Highcharts.<sup>2</sup>

### 5 USER STUDY

We conducted an experiment to evaluate the user experience of VisTorch. In our study, we emulate collaborative data analysis and presentation by having the test facilitator act as the audience and ask questions about the data. The purpose of our study was to determine if the affordance of placing visualizations in physical space helps create *deictic metaphors* that enhance *embodied* [48] and *distributed cognition* [24].

<sup>1</sup><https://opencv.org/>

<sup>2</sup><https://www.highcharts.com/>

(a) Participants P1-P10.				(b) Participants P11-P20.			
#	AGE GROUP	DEGREE	EXPERTISE	#	AGE GROUP	DEGREE	EXPERTISE
P1	25–30	Postdoc	Good	P11	20–25	Masters student	Good
P2	25–30	PhD. student	Good	P12	25–30	Masters student	Good
P3	25–30	PhD. student	Good	P13	25–30	PhD. student	Good
P4	30–35	PhD. student	Good	P14	25–30	Masters student	Good
P5	25–30	PhD. student	Good	P15	30–35	Doctorate	Passing
P6	25–30	Masters	Good	P16	25–30	PhD. student	Good
P7	25–30	Masters student	Good	P17	20–25	Masters student	Passing
P8	25–30	Masters student	Good	P18	30–35	PhD. student	Expert
P9	25–30	Masters student	Good	P19	25–30	Masters student	Good
P10	20–25	Masters student	Good	P20	25–30	PhD. student	Expert

**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.

## 5.1 Apparatus

We use the VisTorch system to conduct our user studies. The study was conducted in a space that resembled an office setting. ArUco markers were placed in the space to divide it into 5 surfaces of interaction that included 3 vertical walls and 2 tabletop surfaces (Fig. 4). VisTorch was tethered to the laptop by a 15ft extension chord and could easily be moved around in the space.

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) disabled to minimize cropping artifacts due to our specific hardware implementation. 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 providing enough room for contents to move around without cropping. Thus, in the experiment, the display was simply turned off when no marker was visible.

## 5.2 Participants

We recruited 20 paid participants (13 identified as male, 7 as female) for our study (see Table 2). The age of the participants ranged from 21–35 years. Most participants were university students, except two who were working professionals. We polled participants before the experiment about their expertise 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 number of visualizations required to complete a specific task: one (1V), two (2V), or three (3V). For example, a 2V task would require combining findings from two separate visualizations to answer, such as looking up the production year of a specific Bluetooth

speaker model, and then using that information to find comparable speakers made that same year.

**• Data Type (T):** The type of data involved: Non-SITUATED (NST): abstract data where physical location has no significance; and SITUATED (ST): context-specific data placed near the referent physical object.

## 5.4 Experimental Design

We used a within-participant factorial design where each participant participated in trials for all conditions:

3	■	View Cardinality V (1, 2, 3)
×	•	Data Type T (Non-situated, Situated)
×	•••	repetitions
18	■	trials per participant.

The order of conditions was fixed for all participants. Our rationale for a fixed data type *T* order was to first elicit spatial organization strategies from participants for abstract data. We were concerned that starting with situated data would bias these strategies even for non-situated data. This was also why we restricted participants from changing the spatial organization in the second phase for non-situated data. For the view cardinality *V*, we chose a fixed order to have participants work up from simple tasks involving only a single view to increasingly more challenging ones.

## 5.5 Metrics and Analysis

The sessions were audio and video recorded. We recorded organization layouts in both time and space. The participants were asked to think aloud during the entire study. The test facilitator also made notes on how the system was being used in-situ. One researcher analyzed the collected data using an inductive approach. This analysis was primarily based on the video and audio recordings while referring to the session notes to study specific highlights. We manually analyzed the entire audio transcript and summarized the findings.

## 5.6 Tasks

Our experiment involved two sensemaking tasks: (i) **layout generation**, i.e. placing and arranging visualizations in physical

space for an initial exploration, and (ii) **Q identifying data items** from the created layout. We conducted the experiment in two parts by data type  $T$ : situated and non-situated data. For each data type, a participant had to perform both tasks: layout generation and identifying data items by number of views involved ( $V$ ). We ensured that the questions required referring to all visualizations at least once. All the visualizations were generated with Highcharts.

**Non-situated visualization (NST):** Here we used a Nobel Laureates multidimensional dataset consisting of 10 visualizations. This data was abstract and had no natural mapping to physical space.

**LAYOUT GENERATION:** The participants were asked to go through a list of visualizations (shown in the folder view) and organize them in space as desired. The organizational strategy for the layout was recorded.

**Q 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 were allowed. However, they could refer to the layout of visualizations any number of times to answer a question. The questions varied based on the number of views ( $V$ ) needed to answer them. Three sets of questions involving one view ( $1V$ ), two views ( $2V$ ), and three views ( $3V$ ), respectively, were asked. Each question was repeated three times.

**Situated visualization (ST):** Here we created the scenario of a shopping experience with three Bluetooth speakers on a desk. Nine visualizations were designed with the data about the three speakers (3 per speaker) from the manufacturer and various retail websites. The visualizations were clearly marked with the name of the speakers they showed the data about. Our primary goal was to generate qualitative feedback on the user experience of creating and referring to ad-hoc physical dashboards.

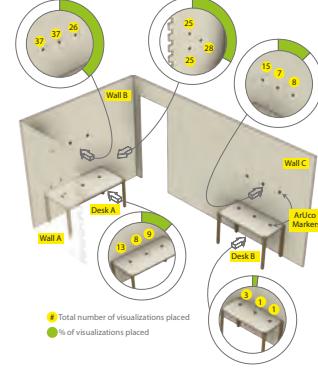
**LAYOUT GENERATION:** The participants were asked to go through the list of visualizations and place charts in the physical space around the speakers as desired.

**Q 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 ( $3V$ ) (repeated three times). Here the participants were free to change and make new organizational layouts.

## 5.7 Procedure

Upon arrival, participants were first screened on their expertise with data visualization. All participants reported having at least “good” expertise on the scale of no experience, passing knowledge, good, and expert. 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



**Figure 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.

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 familiar with the system.

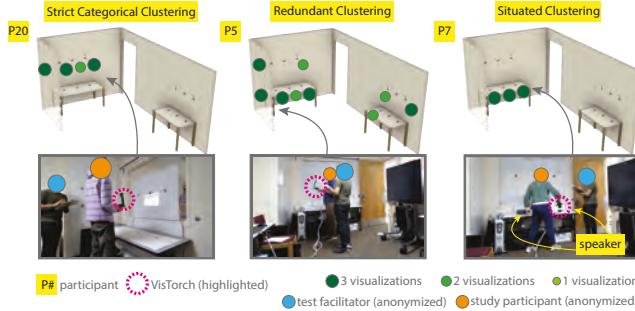
The study began with the non-situated data condition where the participant first explored the list of visualizations and created a layout by arranging them in physical space. However, at this point the type of questions they would be answering was not yet explained. The next phase consisted of three 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: involving one view, two views, and three views. Here the participant was free to make changes to the existing layout or place new charts.

The average time for completion of each of these study conditions was recorded. A typical session lasted between 50–60 minutes. At the end of both the non-situated and situated phase, participants were asked to fill a NASA TLX (task load index) assessment. Participants were compensated with a \$10 gift card.

## 5.8 Results: Overview

Here we give an overview of the results and then dive into the details by data type. Overall, the average completion time for tasks across all conditions were 23.81 minutes (S.D. = 4.19). For both conditions, we collected open-ended qualitative feedback on user



**Figure 5: SPATIAL ORGANIZATIONAL STRATEGIES.** The strategies observed are shown (top) coupled with a snapshot of the corresponding study sessions (bottom). We observed three distinct strategies: *strict categorical clustering* (left), *redundant clustering* (center), and *situated clustering* (right). The green dots show the placement of 1/2/3 visualizations.

experience. We noted the 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.

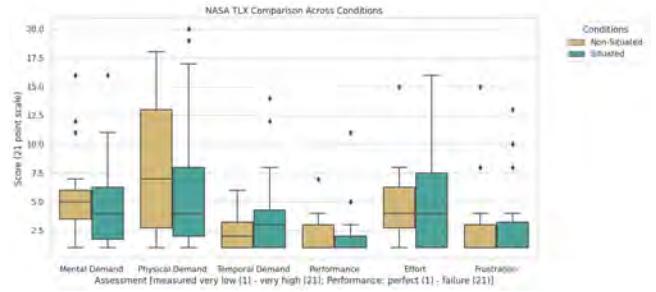
**5.8.1 Non-situated Data.** For the non-situated data condition, we perform a detailed analysis of how the visualizations were organized in space. Then we report on the time taken to complete the task and show the results of the NASA-TLX assessment.

**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 proportion 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.

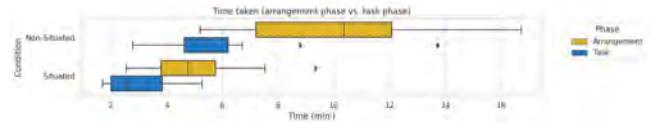
All participants P1-20 generated clusters of visualization based on similarities, attributes, or features that they felt to be of importance. For instance, P1 stated “*I arranged the visualizations by similarity to one another.*” Said P2, “*I made the themes and I remember where the themes were,*” “*This [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 stated that they did not follow any specific strategy to arrange all the visualizations and went through the markers one by one to answer questions.

**1 Completion Time:** The time taken to complete the non-situated data condition across the layout generation (arrangement phase) and identifying data items phase (task phase) (average completion time = 15.88 minutes, S.D. = 3.61) is shown in Fig. 7. P12 took the longest (23.39 minutes), as they had no specific strategy for the layout and went through multiple visualizations to find answers.

**2 NASA-TLX Assessment:** Assessment results for the participants are shown in Fig. 6. We observe that the *mental demand* and *effort* are rated high. From the feedback, we inferred that 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 this issue below.



**Figure 6: NASA TLX.** NASA TLX assessment for non-situated and situated conditions for all 20 participants is shown.



**Figure 7: TIME.** The time (min.) taken to complete the arrangement phase (layout generation) and task phase (identifying data items) across both the non-situated and situated conditions is shown for all 20 participants.

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

**Qualitative Feedback:** We asked the participants about their experience in enacting a sales person and answering questions about the products. P1 said that “[it] helps build mental connection with the product.” P20 said that “it felt very natural to want to attach the data to the physical artifacts.” P2 found the scenario interesting and noted that “you can

*see the product and the images[visualizations] together.*" P4 and P7 both noted that the device is helpful in shopping experiences when the customer shares the same view as the sales person as opposed to looking at a separate screen.

⌚ **Completion Time:** The time taken to complete the situated data condition across the layout generation (arrangement phase) and data identification (task phase) (average completion time = 7.92 minutes, S.D. = 1.83) is shown in Fig. 7.

☰ **NASA TLX Assessment:** These results are shown in Fig. 6. Participants mentioned that physical demand was rated higher because of the weight of the system, which over time got a little tiring.

**5.8.3 Qualitative Experiences.** Here we summarize all the experiences and open-ended feedback of all the participants based on the features of the VisTorch device. We also present the participants' ideas on how they could use the device.

👉 **Deixis in Guided Presentation:** We observed that the participants often used deictic metaphors when presenting an overview of their layouts to the experiment administrator. They not only used point and reveal gestures with the VisTorch to direct the viewer's attention, but also used deictic words such as *there, here, these, and those* to refer to charts. Understanding such deictic metaphors was easy for anyone co-located in the space, thus making it easy to narrate an overview of the data. P3 noted how the deictic affordance of the device would help them create guided presentations, saying "*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, [the 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]*." P13 echoed the idea saying physical movement and spatially distributed information would help create engaging interaction with an audience. P5 enjoyed pointing to reveal data: "*It was easy to point in a direction and see a visualization*."

👉 **Spatial Arrangement and Mental Models:** All participants mentioned the importance of self-authored layouts of visualizations, which helped memorization. Although the layouts generated by the participants differed from one another, it was interesting to see participants being accurate in finding charts while completing tasks. In general, participants preferred to organize layouts on vertical surfaces as opposed to horizontal ones. This is probably because it is relatively inconvenient to aim VisTorch at an angle onto horizontal surfaces compared to vertical ones. Participants also preferred to use physical corners, as we see in Fig. 4, where *Wall A* and *Wall B* were heavily used. The rationale may be that corners provide two surfaces to interact with. We also observed that the participants preferred to use the surfaces that were away from bright light sources. The low brightness of the projector is clearly a limiting factor. P1 said "*I like that you can place [visualizations] in your surroundings,*" and added, "*[situated visualizations] help build mental connection with the product.*" Said P2: "*I made the themes [of*

*organizing visualizations in the layout] and I also remember where the themes were.*" They also commented on how the experience of overlaying digital information on physical objects helps enhance the experience of making decisions about the physical products. P7 stated, "*As long as I spend time to organize [visualizations], it's easy for me to quickly grab the information.*" P14 felt that using space to embody data helped with recall.

📌 **Authoring Ad-hoc Dashboards with Direct Manipulation:** Participants stated that self-authoring visualization dashboards gave them freedom in arranging in space, sometimes even on the fly. Paired with the metaphor of manipulating visualizations by pointing made the user experience seamless, as if the charts were physical artifacts. We speculate that VisTorch facilitates this experience by being portable and handheld, and by affording overlaying digital information on physical space as well as direct manipulation of that information. This helps create a virtually infinite dashboard that can be overlaid on physical space where objects of everyday life act as anchors. P2 said, "*I did like to make my own dashboard on the go.*" Similarly, P3 said, "*I like that I can place the same visualization multiple times*" while talking about making dashboards to do comparison tasks. Both P9 and P11 stated 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. Said P3: "*I feel like I am holding [a 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 felt that the interaction was seamless.

## 5.9 Results: Spatial Organization Strategies

We observed three broad organizational strategies employed by the participants in the study.

**5.9.1 Strict Categorical Clustering:** We observed 70% of the participants strictly following a categorical clustering technique, where they placed one copy of all the visualizations in specific self-authored clusters. Here the participants made clusters of distinct categories and preferred to move between these clusters while performing tasks, such as referring to a visualization or comparing between two visualizations from different clusters. An example of such a participant authored layout is shown in Fig. 5. We observed that participants using this strategy were able to distinctly remember the position of each cluster and physically moved between clusters to perform any comparison tasks.

**5.9.2 Redundant Clustering:** We observed a pattern where 30% of participants made not only unique categorical clusters, but also used multiple copies of the same visualization that they believed were relevant in context. For instance, Fig. 5 shows multiple copies of the same visualization created by P5 because they believed that each cluster was unique. We also observed this while analyzing the open ended feedback, where the participant (P3) said, "*I had placed things by groups and knew exactly where to find the answer*" about creating clusters with multiple views of the same visualization may help answer question faster.

**5.9.3 Situated Clustering:** We noted that all participants prioritized the immediate vicinity of the artifacts in placing visualizations that described the artifacts, as shown in Fig. 5. This was observed while working with the situated condition where the visualizations described the speakers. The visualizations about the speakers were aligned with their physical placement.

## 5.10 Results: Temporal Organization Strategies

Upon analyzing the data we observed that during free exploration and authoring of layout clusters, participants spent more time compared to when performing tasks. In Fig. 7, we show a comparison of time spent by the participants in authoring layouts (arrangement phase) vs. answering questions (task phase) during the task. We observe that once the participants completed the exploration phase, the time spent in answering questions was significantly less.

A minor temporal pattern we noted was that, unsurprisingly, participant recall improved over time as they learned the location of charts with use. However, this is not an organizational strategy, so we do not offer it as a finding.

## 6 EXPERT REVIEW

To complement the user study, which was conducted with a convenience population of graduate students, we also performed an expert review with four professionals drawn from varying fields with situated data. This study was intended to both yield empirical data for VisTorch as well as derive realistic and practical usage scenarios for the device.

### 6.1 Participants

We conducted a total of 4 reviews with experts from different disciplines (E1-E4):

- **Expert 1:** a creative technologist and educator at the Smithsonian Museums who designs AR museum guides.
- **Expert 2:** a projections-media designer who specializes in creative system integration, cinematography, analog-digital hybrid puppetry and immersive performances at a performing arts center.
- **Expert 3:** an environmental education coordinator who works with a school of conservation that fosters environmental knowledge through education programs delivered in natural settings.
- **Expert 4:** an economist and information science Ph.D. who develops statistical software.

These were all working professionals in data-driven fields with at least 3 years of experience. They were compensated for their time and labor.

### 6.2 Procedure

We first introduced VisTorch and all its features (including translation and scaling) to the experts. We then let them use the device to explore the Nobel laureates dataset and solve a few tasks. These tasks were based on the attributes of the dataset that could be answered by referring to one, two, or three visualizations, respectively. This was done to make sure that they completely understood the

functionality of the device. After this, we asked if they would speculate on potential use cases for the device in their everyday work and to demonstrate how it might be used in such settings.

## 6.3 Results

All experts said that they would use deictic gestures while interacting with data, such as referring to information about an artwork in museums, set props in a theater, brainstorming on large whiteboards, and presenting data about environmental impacts of practices in (metalsmithing) workshops. The affordance of a shared display of VisTorch between the presenter and their respective patrons co-located in space facilitates deictic metaphors about the presented data and any related artifact. The experts said they would use VisTorch to create a spatial arrangement of visualizations in context of their respective artifacts for a better understanding of the group they are presenting to. VisTorch helps establish mental connections between data and physical artifacts by situating the data. The experts also said that the ability to manipulate digital information in physical space and create ad-hoc dashboards would help compare ideas such as comparing between the underlying composition principles between two different artworks in museums, comparing the position of similar theater props on set plots while arranging sets, contrast ideas between multiple approaches in a brainstorming, and contrasting the impacts of metalsmithing practices in a shed.

**6.3.1 Spatial Reasoning.** Several experts remarked on the use of VisTorch to support spatial reasoning in situated data. Said E1 on the topic of seeing shapes and patterns in paintings that may be invisible to the untrained eye: *“With a lot of traditional artwork there are underlying compositional shapes and patterns... For example, triangles will be used for stability... On top of that you have things such as implied lines that will guide the viewer... [...] So the viewer’s eye can be guided and someone that is not experienced looking at art might not recognize this, at least not consciously.”*

Similarly, E2 remarked on spatial arrangements on the theater stage: *“To talk about set walls, I have 16 flats and most of them look the same. I need to be able to quickly assess with the crew of the touring house where this flat goes. We have this application and I scan the back of the flat [set walls], and instead of having entire charts pasted to the back of flats that will surely get damaged... I see the [stage plot] drawing and this particular component highlighted in it.”*

Finally, E4 speculated that the device would enable in-situ brainstorming where visualizations can be placed alongside hand-written ideas/algorithms on the whiteboards/frosted glass.

**6.3.2 Engagement and Efficiency.** Another common theme was the use of spatial and situated visualization for increasing the engagement of participants and audience members. Expert E3 noted that *“[VisTorch] shows a little bit more of a dynamic way of showing data... they (students) will be looking at different parts of the room that will keep [students] more engaged.”* E4 felt that *“With printed stacks of paper notes, you are not looking at a shared artifact right in front of all of you at the same time, some people have to flip through their stuff to catch up. Being able to remove something and replace it something else would speed things along.”*

**6.3.3 Collaboration.** Finally, all but one participant explicitly noted the utility of VisTorch for collaboration. E1 said “*if a tour guide was taking a group on tour... You know everyone crowds around the artwork and so its already being observed by a focus group. The tour guide could bring up more information projected on to the art, point out interesting things using (projected) overlays on the art.*” E2 had similar sentiments: “*It's not just me who is seeing [the projected set plot on the set components], it but also the crew... Because it's a shared activity. It promotes a thorough understanding of the design.*” And E4: “*Collaborative brainstorming would be how I would use it. I can see myself doing that...*”

## 7 DISCUSSION

Here we provide our interpretation of the results and discuss how VisTorch can be generalized across various scenarios.

### 7.1 Explaining the Results

Our results indicate that VisTorch helps individuals across disciplines make efficient and intelligent use of physical space—as previously proposed by Kirsh [30]—to simplify choice (by allowing users to organize charts in a sequential order for a specific task), perception (by placing prominent charts at a relevant landmark, like a speaker), and computation (by duplicating charts as needed). Our findings are also analogous to organization strategies observed in large display environments for sensemaking, such as using physical space to form external memory [1, 26], establishing spatial semantics to describe relationships in large display walls and AR [34], and using physical navigation to enhance performance [3, 27]. They also support Calepso et al.’s findings on people’s preferences for physical referents [8].

It is worth noting that for the non-situated data in our user study, the physical landmarks in our test environment are convenient “proxies” used merely as placeholders for visualizations; this means that the visualizations are what Satriadi et al. [45] call “proxsituated.” For our situated data condition, the visualizations were mostly (but not always) attached to the physical referents—the Bluetooth speakers—themselves.

Finally, an interesting observation is that even though VisTorch employs peephole interaction, our results are comparable to situations involving large displays and AR where the contents are always visible. This may be attributed to our participants’ ability to use physical landmarks in lieu of actual contents to aid recall. While the dataset used in our study was smaller compared to the related work, it remains interesting to see that spatial organizational strategies are consistent across platforms.

### 7.2 Generalizing the Results

We feel that one of the outcomes of this work is that situated data is central to many applications. While our 20-participant user study provided *deep* insight into spatial organization for data analysis, our expert review showed the *breadth* and *variety* of such tasks. The review engaged four different professionals from widely disparate fields—education, statistics, cinematography, stage production, and museums—who were all readily able to discuss how a device such as the VisTorch could aid them in their everyday tasks. Just like time is a part of all data, even if it may not always be explicitly captured,

so is space: not just *when* but also *where* the data was collected, how it spatially *relates* to other data, and how its constituent parts are *connected*. Space is a canvas in more ways than one.

### 7.3 Limitations and Future Work

While VisTorch has all the affordances needed to author situated visualization dashboards in the real world, our implementation has room for improvement. For one thing, participants from our study felt that the device is heavy. Some participants even put the device down for a few seconds in between tasks for a break. 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 fairly heavy. The fact that VisTorch must be held at all times not only means that weight is a factor, but also that one hand is captive and thus not available for bimanual interaction. Furthermore, while onboard computation can be integrated into the device in the future, the current implementation is tethered.

An alternative to the handheld projection enabled by VisTorch would be the use of fixed projectors and cameras to turn an entire room into an intelligent display environment, as in the RoomAlive [28] project. Such environments would enable mixed reality experiences without tying up one of the user’s hands, for example. However, the VisTorch has the benefit of being fully mobile—at least in its future untethered state—and usable anytime and anywhere.

A few participants (P3, P7, P8, and P18) suggested improvements to VisTorch, such as increasing the resolution of the device and addressing the occasional cropping of the display when the projector is held at oblique angles to a surface. This problem can be attributed to the projector’s low resolution and narrow throw angle. On the other hand, the current projector does offer other advantages such as being focus free, low power, lightweight, and extremely portable. At the time of design, this was the best option available.

We stabilize the image of the display and smooth out any small movements by employing a stabilization technique based on moving averages. This induces a small latency of a few seconds to stabilize the image after pointing the device. This was pointed out by P3, P6, P8, P12, and P15. In future work, we will reduce this latency by parallelizing the stabilization as well as dynamically changing the window of moving averages based on user actions.

Finally, we envision incorporating an IR camera where the fiducial markers in the environment are painted with IR ink, making them invisible to the human eye. Alternatively, instead of markers, we may conceivably use camera-based tracking of real-world landmarks to virtually place objects in the physical environment.

## 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 and expert reviewing involving four professionals in data-driven professions. A generalization of our results shows the utility of the

device as a projected data display for arranging and analyzing data representations in physical space.

## ACKNOWLEDGMENTS

This work was partly supported by grant IIS-1908605 from the U.S. National Science Foundation and Villum Investigator grant VL-54492 by Villum Fonden. 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

- Christopher Andrews, Alex Endert, and Chris North. 2010. Space to Think: Large, High-Resolution Displays for Sensemaking. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 55–64. <https://doi.org/10.1145/1753326.1753336>
- Benjamin Bach, Ronell Sicat, Johanna Beyer, Maxime Cordeil, and Hanspeter Pfister. 2018. The Hologram in My Hand: How Effective is Interactive Exploration of 3D Visualizations in Immersive Tangible Augmented Reality? *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 457–467. <https://doi.org/10.1109/TVCG.2017.2745941>
- Robert Ball, Chris North, and Doug A. Bowman. 2007. Move to improve: promoting physical navigation to increase user performance with large displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 191–200. <https://doi.org/10.1145/1240624.1240656>
- Andrea Batch, Sungbok Shin, Julia Liu, Peter W. S. Butcher, Panagiotis D. Ritsos, and Niklas Elmqvist. 2023. Evaluating View Management for Situated Visualization in Web-based Handheld AR. *Computer Graphics Forum* 42, 3 (2023), 349–360. <https://doi.org/10.1111/cgf.14835>
- Gábor Blaskó, Steven Feiner, and Franz Coriand. 2005. Exploring Interaction with a Simulated Wrist-Worn Projection Display. In *Proceedings of the IEEE International Symposium on Wearable Computers*. IEEE Computer Society, Los Alamitos, CA, USA, 2–9. <https://doi.org/10.1109/ISWC.2005.21>
- Nathalie Bressa, Henrik Korsgaard, Aurélien Tabard, Steven Houben, and Jo Vermeulen. 2022. What's the Situation with Situated Visualization? A Survey and Perspectives on Situatedness. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 107–117. <https://doi.org/10.1109/TVCG.2021.3114835>
- Peter William Scott Butcher, Nigel W. John, and Panagiotis D. Ritsos. 2021. VRIA: A Web-based Framework for Creating Immersive Analytics Experiences. *IEEE Transactions on Visualization and Computer Graphics* 27, 7 (2021), 3213–3225. <https://doi.org/10.1109/TVCG.2020.2965109>
- Aimée Sousa Calepsø, Philipp Fleck, Dieter Schmalstieg, and Michael Sedlmair. 2023. Exploring Augmented Reality for Situated Analytics with Many Movable Physical Referents. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*. ACM, New York, NY, USA, 6:1–6:12. <https://doi.org/10.1145/3611659.3615700>
- Tom Chandler, Maxime Cordeil, Tobias Czauderna, Tim Dwyer, Jaroslaw Glowacki, Cagatay Goncu, Matthias Klapperstueck, Karsten Klein, Falk Schreiber, and Elliot Wilson. 2015. Immersive Analytics. In *Proceedings of the International Symposium on Big Data Visual Analytics*. IEEE Computer Society, Los Alamitos, CA, USA, 1–8. <https://doi.org/10.1109/BDVA.2015.7314296>
- Zhutian Chen, Yijia Su, Yifang Wang, Qianwen Wang, Huamin Qu, and Yingcai Wu. 2019. MARViST: Authoring Glyph-based Visualization in Mobile Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 26, 8 (2019), 2645–2658. <https://doi.org/10.1109/TVCG.2019.2892415>
- Maxime Cordeil, Andrew Cunningham, Benjamin Bach, Christophe Hurter, Bruce H. Thomas, Kim Marriott, and Tim Dwyer. 2019. IATK: An Immersive Analytics Toolkit. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*. IEEE Computer Society, Los Alamitos, CA, USA, 200–209. <https://doi.org/10.1109/VR.2019.8797978>
- Maxime Cordeil, Andrew Cunningham, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2017. ImAxes: Immersive Axes as Embodied Affordances for Interactive Multivariate Data Visualisation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 71–83. <https://doi.org/10.1145/3126594.3126613>
- Raimund Dachselt, Jonna Häkkilä, Matt Jones, Markus Löchtefeld, Michael Rohs, and Enrico Rukzio. 2012. Pico projectors: firefly or bright future? *Interactions* 19, 2 (2012), 24–29. <https://doi.org/10.1145/2090150.2090158>
- Mustafa Doga Dogan, Ahmad Taka, Michael Lu, Yunyi Zhu, Akshat Kumar, Aakar Gupta, and Stefanie Müller. 2022. 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*. ACM, New York, NY, USA, 269:1–269:12. <https://doi.org/10.1145/3491102.3501951>
- Niklas Elmqvist. 2011. Embodied Human-Data Interaction. In *Proceedings of the ACM CHI Workshop on Embodied Interaction: Theory and Practice in HCI*. ACM, New York, NY, USA, 104–107.
- Niklas Elmqvist. 2023. Data Analytics Anywhere and Everywhere. *Commun. ACM* 66, 12 (2023), 52–63. <https://doi.org/10.1145/3584858>
- Niklas Elmqvist and Pourang Irani. 2013. Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime. *IEEE Computer* 46, 4 (2013), 86–89. <https://doi.org/10.1109/mc.2013.147>
- Neven A. M. Elsayed, Bruce H. Thomas, Kim Marriott, Julia Piantadosi, and Ross T. Smith. 2016. Situated Analytics: Demonstrating immersive analytical tools with Augmented Reality. *Journal of Visual Languages & Computing* 36 (2016), 13–23. <https://doi.org/10.1016/j.jvlc.2016.07.006>
- Steven Feiner, Blair MacIntyre, Tobias Höllerer, and Anthony Webster. 1997. A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment. *Personal Technologies* 1, 4 (1997), 208–217.
- Fernando Garcia-Sanjuan, Javier Jaen, and Vicente Nacher. 2016. Toward a General Conceptualization of Multi-Display Environments. *Frontiers in ICT* 3 (2016), 20. <https://doi.org/10.3389/fict.2016.00020>
- Sergio Garrido-Jurado, Rafael Muñoz-Salinas, Francisco José Madrid-Cuevas, and Manuel Jesús Marín-Jiménez. 2014. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition* 47, 6 (2014), 2280–2292. <https://doi.org/10.1016/j.patcog.2014.01.005>
- Steven Henderson and Steven Feiner. 2011. Exploring the Benefits of Augmented Reality Documentation for Maintenance and Repair. *IEEE Transactions on Visualization and Computer Graphics* 17, 10 (2011), 1355–1368. <https://doi.org/10.1109/TVCG.2010.245>
- Paul Hine, Loza Tadesse Mamo, and Narcís Parés. 2022. AR Magic Lantern: Group-based Co-Located Augmentation Based on the World-as-Support AR Paradigm. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 207:1–207:5. <https://doi.org/10.1145/3491101.3519918>
- Edwin Hutchins. 1995. *Cognition in the Wild*. MIT Press, Cambridge, MA, USA.
- Petra Isenberg, Niklas Elmqvist, Jean Scholtz, Daniel Cernea, Kwan-Liu Ma, and Hans Hagen. 2011. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization* 10, 4 (2011), 310–326. <https://doi.org/10.1177/1473816111412817>
- Petra Isenberg, Danyel Fisher, Sharoda A. Paul, Meredith Ringel Morris, Kori Inkpen, and Mary Czerwinski. 2012. Co-Located Collaborative Visual Analytics around a Tabletop Display. *IEEE Transactions on Visualization and Computer Graphics* 18, 5 (2012), 689–702. <https://doi.org/10.1109/TVCG.2011.287>
- Yvonne Jansen, Jonas Schjerlund, and Kasper Hornbæk. 2019. Effects of Locomotion and Visual Overview on Spatial Memory when Interacting with Wall Displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 291:1–291:12. <https://doi.org/10.1145/3290605.3300521>
- Brett J. Jones, Rajinder Sodhi, Michael Murdock, Ravish Mehra, Hrvoje Benko, Andrew D. Wilson, Eyal Ofek, Blair MacIntyre, Nikunj Raghuvanshi, and Lior Shapira. 2014. RoomAlive: magical experiences enabled by scalable, adaptive projector-camera units. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 637–644. <https://doi.org/10.1145/2642918.2647383>
- KyungTae Kim, Waqas Javed, Cary Williams, Niklas Elmqvist, and Pourang Irani. 2010. Hugin: A Framework for Awareness and Coordination in Mixed-Presence Collaborative Information Visualization. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. ACM, New York, NY, USA, 231–240. <https://doi.org/10.1145/1936652.1936694>
- David Kirsh. 1995. The Intelligent Use of Space. *Artificial Intelligence* 73, 1–2 (1995), 31–68. [https://doi.org/10.1016/0004-3702\(94\)00017-U](https://doi.org/10.1016/0004-3702(94)00017-U)
- Bongshin Lee, Raimund Dachselt, Petra Isenberg, and Eun Kyoung Choe. 2022. *Mobile Data Visualization*. Chapman and Hall/CRC Press, Boca Raton, FL, USA.
- Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. 2022. Effects of Display Layout on Spatial Memory for Immersive Environments. *Proceedings of the ACM on Human-Computer Interaction* 6, ISS (2022), 468–488. <https://doi.org/10.1145/3567729>
- Qiang Liu and Erik Jorgensen. 2011. Muscle memory. *Journal of Physiology* 89, Pt 4 (2011), 775–776. <https://doi.org/10.1113/jphysiol.2011.205088>
- Weizhou Luo, Anke Lehmann, Hjalmar Widengren, and Raimund Dachselt. 2022. Where Should We Put It? Layout and Placement Strategies of Documents in Augmented Reality for Collaborative Sensemaking. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 627:1–627:16. <https://doi.org/10.1145/3491102.3501946>
- Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H. Thomas (Eds.). 2018. *Immersive Analytics*. Lecture Notes in Computer Science, Vol. 11190. Springer, New York, NY, USA. <https://doi.org/10.1007/978-3-030-01388-2>
- Miguel A. Nacenta, Regan L. Mandryk, and Carl Gutwin. 2008. Targeting across displayless space. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 777–786. <https://doi.org/10.1145/1357054.1357178>

[37] Miguel A. Nacenta, Satoshi Sakurai, Tokuo Yamaguchi, Yohei Miki, Yuichi Itoh, Yoshifumi Kitamura, Sriram Subramanian, and Carl Gutwin. 2007. E-conic: a perspective-aware interface for multi-display environments. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 279–288. <https://doi.org/10.1145/1294211.1294260>

[38] Miguel A. Nacenta, Samer Sallam, Bernard Champoux, Sriram Subramanian, and Carl Gutwin. 2006. Perspective cursor: perspective-based interaction for multi-display environments. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 289–298. <https://doi.org/10.1145/1124772.1124817>

[39] Claudio S. Pinhanez. 2001. The Everywhere Displays Projector: A Device to Create Ubiquitous Graphical Interfaces. In *Proceedings of the International Conference on Ubiquitous Computing (Lecture Notes in Computer Science, Vol. 2201)*. Springer, New York, NY, USA, 315–331. [https://doi.org/10.1007/3-540-45427-6\\_27](https://doi.org/10.1007/3-540-45427-6_27)

[40] Ramesh Raskar, Greg Welch, Matt Cutts, Adam T. Lake, Lev Stesin, and Henry Fuchs. 1998. 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*. ACM, New York, NY, USA, 179–188. <https://doi.org/10.1145/280814.280861>

[41] J. Roberts, P. Ritsos, S. K. Badam, D. Brodbeck, J. Kennedy, and N. Elmquist. 2014. Visualization beyond the Desktop—the Next Big Thing. *IEEE Computer Graphics and Applications* 34, 6 (Nov 2014), 26–34. <https://doi.org/10.1109/MCG.2014.82>

[42] George Robertson, Mary Czerwinski, Kevin Larson, Daniel C. Robbins, David Thiel, and Maarten van Dantzig. 1998. Data Mountain: Using Spatial Memory for Document Management. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 153–162. <https://doi.org/10.1145/288392.288596>

[43] Enrico Rukzio, Paul Holleis, and Hans Gellersen. 2012. Personal Projectors for Pervasive Computing. *IEEE Pervasive Computing* 11, 2 (2012), 30–37. <https://doi.org/10.1109/MPRV.2011.17>

[44] Alper Sarikaya, Michael Correll, Lyn Bartram, Melanie Tory, and Danyel Fisher. 2019. What Do We Talk About When We Talk About Dashboards? *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 682–692. <https://doi.org/10.1109/TVCG.2018.2864903>

[45] Kadek Ananta Satriadi, Andrew Cunningham, Ross T. Smith, Tim Dwyer, Adam Drogemuller, and Bruce H. Thomas. 2023. ProxSituated Visualization: An Extended Model of Situated Visualization using Proxies for Physical Referents. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 382:1–382:20. <https://doi.org/10.1145/3544548.3580952>

[46] Joey Scarr, Andy Cockburn, and Carl Gutwin. 2013. Supporting and Exploiting Spatial Memory in User Interfaces. *Foundations and Trends in Human-Computer Interaction* 6, 1 (2013), 1–84. <https://doi.org/10.1561/1100000046>

[47] Johannes Schöning, Markus Löchtefeld, Michael Rohs, and Antonio Krüger. 2010. Projector Phones: A New Class of Interfaces for Augmented Reality. *International Journal of Mobile Human-Computer Interaction* 2, 3 (2010), 1–14. <https://doi.org/10.4018/jmhci.2010070101>

[48] Lawrence A. Shapiro. 2011. *Embodied Cognition*. Routledge, New York, NY, USA.

[49] Sungbok Shin, Andrea Batch, Peter W. S. Butcher, Panagiotis D. Ritsos, and Niklas Elmquist. 2024. The Reality of the Situation: A Survey of Situated Analytics. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 19 pages. <https://doi.org/10.1109/TVCG.2023.3285546> to appear.

[50] Yael Shrager, Peter J. Bayley, Bruno Bontempi, Ramona O. Hopkins, , and Larry R. Squire. 2007. Spatial memory and the human hippocampus. *Proceedings of the National Academy of Sciences* 104, 8 (2007), 2961–2966. <https://doi.org/10.1073/pnas.0611233104>

[51] Ronell Sicat, Jiabao Li, Junyoung Choi, Maxime Cordeil, Won-Ki Jeong, Benjamin Bach, and Hanspeter Pfister. 2019. DXR: A Toolkit for Building Immersive Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 715–725. <https://doi.org/10.1109/TVCG.2018.2865152>

[52] Lucy A. Suchman. 1987. *Plans and Situated Actions — The Problem of Human-Machine Communication*. Cambridge University Press, Cambridge, UK.

[53] Masanori Sugimoto, Kosuke Miyahara, Hiroshi Inoue, and Yuji Tsunesada. 2005. Hotaru: Intuitive Manipulation Techniques for Projected Displays of Mobile Devices. In *Proceedings of the IFIP TC13 Conference on Human-Computer Interaction (Lecture Notes in Computer Science, Vol. 3585)*. Springer, New York, NY, USA, 57–68. [https://doi.org/10.1007/11555261\\_8](https://doi.org/10.1007/11555261_8)

[54] Bruce H. Thomas, Gregory F. Welch, Pierre Dragicevic, Niklas Elmquist, Pourang Irani, Yvonne Jansen, Dieter Schmalstieg, Aurélien Tabard, Neven A. M. ElSayed, Ross T. Smith, and Wesley Willett. 2018. Situated Analytics. In *Immersive Analytics (Lecture Notes in Computer Science, Vol. 11190)*. Springer, New York, NY, USA, 185–220. [https://doi.org/10.1007/978-3-030-01388-2\\_7](https://doi.org/10.1007/978-3-030-01388-2_7)

[55] Jo Vermeulen, Christopher Collins, Raimund Dachsel, Pourang Irani, and Alark Joshi. 2021. Reflections on Ubiquitous Visualization. In *Mobile Data Visualization*. Chapman and Hall/CRC Press, Boca Raton, FL, USA, 263–316.

[56] Mark Weiser. 1991. The computer for the 21st Century. *Scientific American* 265, 3 (1991), 94–104. <https://doi.org/10.1145/329124.329126>

[57] Sean White and Steven Feiner. 2009. SiteLens: Situated Visualization Techniques for Urban Site Visits. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–8. <https://doi.org/10.1145/642611.642613>