



Lights, Headset, Tablet, Action: Exploring the Use of Hybrid User Interfaces for Immersive Situated Analytics

XIAOYAN ZHOU*, Colorado State University, USA

BENJAMIN LEE*[†], University of Stuttgart, Germany and JPMorganChase, USA

FRANCISCO R. ORTEGA, Colorado State University, USA

ANIL UFUK BATMAZ, Concordia University, Canada

YALONG YANG, Georgia Institute of Technology, USA

While augmented reality (AR) headsets provide entirely new ways of seeing and interacting with data, traditional computing devices can play a symbiotic role when used in conjunction with AR as a *hybrid user interface*. A promising use case for this setup is *situated analytics*. AR can provide embedded views that are integrated with their physical referents, and a separate device such as a tablet can provide a familiar situated overview of the entire dataset being examined. While prior work has explored similar setups, we sought to understand how people perceive and make use of visualizations presented on both embedded visualizations (in AR) and situated visualizations (on a tablet) to achieve their own goals. To this end, we conducted an exploratory study using a scenario and task familiar to most: adjusting light levels in a smart home based on personal preference and energy usage. In a prototype that simulates AR in virtual reality, embedded visualizations are positioned next to lights distributed across an apartment, and situated visualizations are provided on a handheld tablet. We observed and interviewed 19 participants using the prototype. Participants were easily able to perform the task, though the extent the visualizations were used during the task varied, with some making decisions based on the data and others only on their own preferences. Our findings also suggest the two distinct roles that situated and embedded visualizations can have, and how this clear separation might improve user satisfaction and minimize attention-switching overheads in this hybrid user interface setup. We conclude by discussing the importance of considering the user's needs, goals, and the physical environment for designing and evaluating effective situated analytics applications.

CCS Concepts: • **Human-centered computing** → **Empirical studies in visualization**; *Visualization theory, concepts and paradigms*; **Empirical studies in HCI**; **Virtual reality**; *Mixed / augmented reality*.

Additional Key Words and Phrases: Virtual and augmented reality, Visualization, Situated analytics, Hybrid user interfaces, Immersive analytics

ACM Reference Format:

Xiaoyan Zhou, Benjamin Lee, Francisco R. Ortega, Anil Ufuk Batmaz, and Yalong Yang. 2024. Lights, Headset, Tablet, Action: Exploring the Use of Hybrid User Interfaces for Immersive Situated Analytics. *Proc. ACM Hum.-Comput. Interact.* 8, ISS, Article 547 (December 2024), 23 pages. <https://doi.org/10.1145/3698147>

*Both authors contributed equally to this research.

[†]This research was conducted prior to joining JPMorganChase.

Authors' Contact Information: [Xiaoyan Zhou](mailto:xiaoyan.zhou@colostate.edu), Colorado State University, Fort Collins, USA, xiaoyan.zhou@colostate.edu; [Benjamin Lee](mailto:Benjamin.Lee@jpmchase.com), University of Stuttgart, Stuttgart, Germany and JPMorganChase, New York, USA, Benjamin.Lee@jpmchase.com; [Francisco R. Ortega](mailto:fortega@colostate.edu), Colorado State University, Fort Collins, USA, fortega@colostate.edu; [Anil Ufuk Batmaz](mailto:ufuk.batmaz@concordia.ca), Concordia University, Montreal, Canada, ufuk.batmaz@concordia.ca; [Yalong Yang](mailto:yalong.yang@gatech.edu), Georgia Institute of Technology, Atlanta, USA, yalong.yang@gatech.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2573-0142/2024/12-ART547

<https://doi.org/10.1145/3698147>

1 Introduction

The real world is replete with information that we parse and make sense of in our daily lives. This information can oftentimes be “invisible” however, necessitating the use of external tools to help us understand the objects and phenomena in our environment. For example, a viticulturist (i.e., grape growing expert) may rely on aerial imaging data to understand the current state of their vineyards and determine how best to optimize crop yields for the next season [34]. In abstract terms, the *logical world* in which the imaging data resides directly embodies the *physical world* [75] that, in this case, is the vineyards. Yet, while the logical world may inform the viticulturist and drive them to take action, it is only within the physical world that they can tend to the vineyards directly and impart changes. While these two stages are usually performed at different times and spaces, doing both together *in-situ* can prove beneficial or is sometimes necessary. For instance, consider the same viticulturist arriving on-site shortly after only to notice a pest infestation that the aerial imaging could not capture. New determinations would need to be made using up-to-date knowledge derived from the logical and the physical worlds together at once. This approach of analyzing data within its physical context is known as *situated analytics* (SA) [15].

SA can be facilitated by any technology that provides access to data in close spatial proximity to the objects to which it relates; the latter being referred to as *referents* [75]. Immersive technologies, including virtual reality (VR) and particularly augmented reality (AR), have now been used to integrate visualizations with physical referents [43, 63], which may be *situated* in rough proximity to the referents (i.e., situated visualization or *SitVis*), or *embedded* close to or on top of the referents directly (i.e., embedded visualization or *EmbVis*) [75]. The use of *EmbVis* reduces the spatial distance between data and referent, making it easier to establish mental connections between the two. While there may sometimes only be a single referent to which the data relates (e.g., a street [71]), there can also be many referents that all need to be considered together (e.g., products in a supermarket [12, 16], elements of a museum exhibit [64]). Although a *one-to-one* configuration with each referent having its own *EmbVis* is common [43], Sayara et al. [59] point out that this approach becomes impractical when referents are either too plentiful and/or are spread far apart. Thus, a single situated *dashboard* can aggregate this information into a single overview instead.

Combining a *SitVis* as a situated dashboard with multiple separate *EmbVis* in AR conceptually forms a *hybrid user interface* (HUI) [19]. In HUIs, multiple display modalities are combined together to form a “symbiosis of interfaces” [76], with each device compensating for each other’s downsides. HUIs are now associated with AR and spatial computing, and have been demonstrated in visualization (e.g., [37, 69]) and general contexts (e.g., with the Apple Vision Pro [1]). In the case of SA, the AR headset can provide embedded views of each individual referent’s data, while a separate display that does not mediate the AR experience (e.g., tablet, monitor) can provide a situated dashboard overview of *all* referent’s data.

The combination of SA and HUIs presents a novel situation that imposes new constraints that need to be addressed. First, assuming the distance between referents and their visualizations is minimized, the layout of AR visualizations is based on the physical referents, which may be unknown to the user. This is not the case in typical spatial computing applications wherein the user does have control over this layout (e.g., [46, 48]). Second, consideration and interaction with the physical world and its referents is vital in SA, unlike in conventional HUIs, which use the physical space only as a working canvas for AR views. Finally, the use of the non-AR display is likely to be influenced by the two preceding factors, such as the need to navigate a large physical environment while carrying a physical display and to switch attention between AR and non-AR visualizations. Constraints and requirements such as this would likely affect the design choices made and the strategies that users employ. While prior work by Whitlock et al. [74] has investigated

the combination of mobile and AR devices in-situ, they focus primarily on supporting data collection and analysis in fieldwork, and did not observe how people actually make use of such a setup. We believe that this intersection between SA and HUIs will be a standard approach for integrating AR devices, smartphones and tablets, and Internet of Things (IoT) devices and sensors together to facilitate SA *anywhere and everywhere* [13].

To investigate the challenges and opportunities that lie at this intersection of SA and HUIs, we conducted a VR-based exploratory study on the design and use of a SA application, wherein users have access to a SitVis dashboard displayed on a tablet alongside multiple EmbVis close to their referents. The use of VR allows us to mimic AR and the HUI in a manner more resistant to technical issues such as tracking drift [64]. Our prototype was designed to resemble a simple baseline scenario: optimizing energy usage of lights in a smart home setting [32]. This approach allows us to observe how people interact with the system in the absence of helper functions such as occlusion mitigators [2] and navigational aids [52]. From this, we seek to understand what aspects of the combination of SA and HUIs users like and dislike, and as a result, identify relevant research challenges and opportunities. With this, we hope to build towards a formal set of guidelines for spatially situated visualization and analytics [17].

Our work contributes to both SA and HUIs with the following:

- Identification of the relevant physical environments that hybrid SA can be deployed in (Section 3)
- A prototype (Section 4) and user study (Section 5) using simulated AR in VR wherein 19 participants used situated visualizations on a tablet with embedded visualizations in the environment for a self-directed optimization task
- Analysis of user preferences and behaviors in using such a system (Section 6)
- Takeaways, guidelines, and recommendations for both SA and HUIs (Section 7)

2 Related Work

This work lies at the intersection of three research fields: the use of virtual and augmented reality (VR/AR) for data visualization, also known as immersive analytics (IA) [50], and the aforementioned situated analytics (SA) and hybrid user interfaces (HUI).

2.1 Immersive and Situated Analytics

IA has become increasingly relevant as consumer VR and AR devices like the Apple Vision Pro and Meta Quest 3 become commonplace. Marriott et al. [50] define IA as “the use of engaging, embodied analysis tools to support data understanding and decision making.” A potential benefit of VR and AR is that they enable stereoscopic 3D visualizations such as space-time cubes [20], 3D scatterplots [42, 78], and 3D trajectories [9]. While perceiving depth for 3D visualizations in VR/AR is easier than on desktops [73], caution should still be exercised in their use, especially if it is unjustified [51]. As a compromise, research has advocated using 2D and 3D visualizations together [27], whether they be juxtaposed side-by-side [30] or transformed between 2D and 3D through animations [40]. A third option is to simply forgo the use of 3D visualizations entirely. Through this, the 3D space serves primarily as a workspace for spatially registered 2D visualizations, whose layout may be user controllable [31, 58] or (partially) automated [28, 47], and potentially influenced by the surrounding environment [42, 48].

In contrast, SA is characterized by the physical environment. Rather than being just a backdrop or a means of organizing visualizations, the environment shares relevance and meaning with the data [72]. As a consequence, how SitVis are designed and where they are placed becomes both guided by and constrained by the involved referents [43]. This is demonstrated in a recent survey by Shin et al. [63], where the majority of surveyed AR systems for SA situate views in

a world-absolute manner. In other words, visualizations are commonly *embedded* [75] onto the referent directly such that they visually overlap, such as on the front of grocery products [16] or on the façades of buildings [21]. While scenarios involving few referents are trivial, Lee et al. [43] note the difficulties in managing referents that are in large quantities and/or are spread far apart. For instance, comparing the EmbVis of buildings at different locations would necessitate walking between them. It has therefore been recommended that in such cases, only *situated* visualizations be used instead [43, 75]. While this no longer establishes direct spatial connections between the data and referents, it allows for a single representation that provides an overview of all referents. We note, however, that the two are not mutually exclusive: it is possible to provide both SitVis and EmbVis simultaneously—the former serving as an overview and the latter as the detail [79]. Our work shares similarities with three others along this vein. First is by Jahn et al. [32] who proposed a system for showing energy usage data from smart home sensors on both a stationary display and a handheld mobile phone AR display. Pointing the mobile phone at the appropriate object (e.g., a desk lamp) overlays a label with the name and energy consumption of the object in AR. They did not conduct a user study, however, nor did they use an AR headset. Second is by Doerr et al. [12] who use a similar configuration with a (virtual) tablet and referents, but focus primarily on visual highlighting in low-level brushing and linking tasks that do not adequately capture the analytics part of SA. Third is by Whitlock et al. [74] who conducted an exploratory investigation into how SitVis (from mobile devices) can be used in conjunction with EmbVis (from AR headsets) to support data collection and analysis in field work. While they elicited feedback from 10 field analysts in a design probe and guided walkthrough, they did not observe how the analysts would use the system by themselves and the challenges they might have faced in doing so.

To the best of our knowledge, no work has adequately explored how users naturally behave when using HUIs for SA. We, therefore, explore the benefits, challenges, and potential strategies in using a single SitVis dashboard [59] with multiple EmbVis together in a SA context. In particular, we investigate the needs of users in SA when referents and their visualizations are spread far apart in different fields of regard, in a layout that is pre-defined by the environment.

2.2 Hybrid User Interfaces with Mobile Displays and AR

2D visualizations are oftentimes still employed in AR, with the go-to approach being floating virtual panels [31, 47, 58]. While virtual panels provide a highly flexible spatial workspace, using physical displays can still prove beneficial. In particular, mobile devices have become a cornerstone of visualization research [39] due to their ubiquity [14, 56], mobility, and portability [41]. Research in IA has subsequently integrated these mobile devices with AR to serve various purposes in HUIs. A common approach is to use the mobile device as an input method to control AR visualizations, leveraging their touch and/or spatial tracking capabilities (e.g., [10, 28, 62]). Mobile devices can also provide 2D views of data, with AR providing spatially registered 3D views. For example, MARVIS by Langner et al. [37] demonstrates how AR can augment data visualizations on tablets in a tabletop setting. STREAM by Hubenschmid et al. [29] instead focuses primarily on 3D visualizations facilitated by AR, with a spatially-aware handheld tablet serving as both an input device and a display for alternate 2D visualizations. In this work, we utilize a handheld tablet as a familiar and portable interface in which 2D visualizations can be accessed and interacted with by users.

Whenever a user looks at digital content that is distributed across multiple displays and locations, their attention needs to switch between them. This switch has been observed to incur performance overheads, wherein the user needs to reorient to new visual objects that are at different distances, sizes [66], or that even have different visual representations [53]. Rashid et al. [54] describe several factors that influence display switching cost, including the extent to which the user needs to turn to see all relevant displays (*angular coverage*), and whether the content across displays is semantically

connected (*content coordination*). Grubert et al. [24] later provide a summary of further factors identified in the literature, including the ability for users to spread their attention across multiple devices (*divided or split attention* [8]). Systems with HUIs can, therefore, be designed to minimize switching cost. For example, Wang and Lindeman [68] found that synchronizing 3D interaction modes between a tablet view and an immersive VR view reduces this cost in a virtual world building task. In a visualization context, this is similar to a coordinated and multiple view problem [55], wherein data in different views need to be looked at and considered all at once. In a SA context however, as it is grounded in physical reality, the available parameters to accommodate context switching are fewer. For instance, the position of EmbVis may not be easily modified as they then lose their spatial context. Before we seek to minimize this switching cost however, we first seek to understand the degree to which it is a problem in SA.

To the best of our knowledge, no work has focused on situations where 3D views mediated by AR in a HUI are spatially distributed and anchored to physical referents, and studied if and how users are able to handle this. That is, what are the trade-offs between optimizing visualization placement for situatedness with referents, versus for viewability to support comparisons between multiple views.

3 Scenarios for Hybrid Situated Analytics

SA is already a regular activity in the present day as people make sense of and understand data in physical contexts to help make decisions. This data comes both from the logical world [75], usually mediated by mobile devices, and the physical world, based on the user's perception of it. AR headsets are well poised to further bridge the gap between data displayed in the logical world and phenomena present in the physical world. In doing so, AR serves as an *extension* to how we presently view and interact with data. We believe this approach has advantages that make it more feasible for it to become a reality in the near future. First, it leverages users' prior intuition with mobile devices that can ease adoption compared to fully immersive setups [69] such as when performing inputs (e.g., [28, 29]), and therefore builds upon the decades of knowledge on mobile computing rather than seeking to replace it outright. And second, it retains access to the mature mobile ecosystem to allow, for example, existing applications to simply be extended using AR instead of re-building them from the ground up for immersive devices, or for users to easily context switch to other tasks like doing a web search or responding to instant messages.

3.1 Scope

The fundamental research question of this work is the following: *What are the opportunities and challenges of using AR headsets to support situated analytics with mobile devices?* Shin et al. [63] note, however, that there exist at least four archetypal designs of SA that aid the user in different ways and for different purposes depending on the context. In other words, the design of SA and the visualizations within it are dependent on the goals of the user and the configuration of the physical environment [43]. To better define the scenarios to which our work applies, we first describe three main characteristics of the physical environments we are interested in. We chose these to be representative of what we believe to be commonplace in the future.

- **Many physical referents.** There are oftentimes many referents of interest, such as books in a library [3, 67]. As Sousa Calepso et al. [64] point out, a very small number of physical referents are typically examined in SA research. To further explore this gap, we investigate a scenario in which many referents are considered together at once.

- **Spatially distributed referents.** Referents can also be spread far apart in sub-optimal viewing arrangements, such as in two separate rooms, thus making tasks such as pairwise comparisons [79] challenging. The layout of these referents might not even be adjustable by the user, particularly if they are stationary in nature.
- **Interaction with the environment.** Another key aspect of SA, as Sousa Calepso et al. [64] again argue, is the combination of both physical and analytical tasks. That is, the environment should be considered and/or interacted with for SA to have proper value, otherwise it is simply IA.

These three environmental characteristics inform the baseline setup that a SA application can be designed around. These effectively form the basis of our exploratory prototype in Section 4.

- **Situated visualizations on the mobile device.** *Situated* visualizations (SitVis) are data representations viewed in close spatial proximity to the physical referents that they relate to, and are typically displayed on a screen [75]. As a baseline, we assume that SitVis are accessible on a mobile device (e.g., tablet, smartphone). The design and layout of these visualizations may either be bespoke (e.g., [74]), similar to a dashboard [59], or simply an existing mobile application.
- **Embedded visualizations in AR.** *Embedded* visualizations (EmbVis) are data representations that are physically aligned with the physical referents that they relate to. This is typically accomplished using projectors, see-through video, AR, etc. in a manner so that the data appears “on top” of their referents [75]. EmbVis tightly integrate data and referent together such that the two appear one and the same—the data representation becomes a “part” of the referent. Through the use of AR, these visualizations can be viewed in a hands-free manner, thus allowing the referent to be physically interacted with and manipulated.

Under the classification of Shin et al. [63], this form of SA application would be considered an *Assistant*. Each referent has some associated data that supports the user’s decision-making process. The *situating trigger* is, therefore, the referents themselves as they come into view, the *view situatedness* a combination of device-relative for the mobile device and world absolute for AR, and the *visual encoding* and *data abstraction* dependent on the type of data that is being visualized.

3.2 Working Scenario

We now describe a possible scenario in the not-so-distant future that illustrates how the two components of the HUI—mobile device and AR—might be used in a SA application in an environment with the aforementioned characteristics. We use the same scenario for our following prototype in Section 4 and exploratory study in Section 5.

- **The environment.** Alice has recently moved into a new five-room apartment at the beginning of the month. Each room has several light bulbs that are all connected to a pre-installed smart lighting system, which supports manual adjustment of brightness levels. The system provides the estimated energy usage per month (in kWh) of each light bulb to inform the user of their energy consumption. The previous tenants have also consented to provide Alice the historical energy consumption of each light bulb in the previous six months.
- **The task.** Alice, being environmentally conscious, optimizes the brightness levels of each light bulb based on their current and historical energy usage within each room and across the entire apartment. She also takes into account other real-world factors not typically captured as data, such as the presence of natural lighting, the actual perceptible brightness of each different light bulb, and her own subjective preferences of what she considers pleasant to live in.
- **The visualizations.** Alice installs the smart home application on her mobile device, which summarizes the energy usage of all the lights in each room and the overall apartment in a convenient dashboard. After an initial glance at the dashboard, Alice decides to put on her



Fig. 1. Images of our prototype for situated analytics using hybrid user interfaces, simulating AR in VR. Left: In the foreground, the user holds a tablet at chest height with a dashboard of situated visualizations that show energy usage data of lights in the apartment. In the background, embedded visualizations of energy usage are next to their respective lights. Right: The apartment’s living room and kitchen with more embedded visualizations. Leader lines connect each embedded visualization to one or more lights that it relates to.

AR headset to get a more direct sense of how it relates to the lights around her. She now sees several small EmbVis overlaid on the light bulbs throughout her apartment, which indicates the energy usage of each light. With this, Alice can make fine-grained adjustments to each light bulb individually, and see both how bright the light shines and how much energy it is now estimated to use in real time. Alice also walks from room to room to check on and adjust each set of lights. Referencing to the mobile application allows Alice to confirm that the overall energy usage is satisfactorily low enough, or if there are other light bulbs she would like to adjust.

This scenario follows the same smart home setting presented by Jahn et al. [32] as described earlier. The task was also inspired by Schröder et al. [61], who also employed heterogeneous devices (i.e., desktop, tablet AR, headset VR) used in an asynchronous manner to solve a spatial optimization task with lamps in a virtual park. Unique to our scenario however is the *synchronous* use of multiple devices, and our focus on single-user contexts and the absence of quantified “score” that the user optimizes for.

4 A Prototype for Situated Analytics with Mobile Device and AR

The purpose of this work is to build a preliminary understanding of how people use a SA application in the form of an HUI. To do this, we develop a SA prototype as shown in Figure 1 which we later use in an exploratory study in Section 5, employing the scenario described in Section 3.2. We intentionally keep the system’s functions simple for it to serve as a baseline. This is to understand the extent that known factors such as occlusion and the need to attention switch (Section 2) negatively affect people’s ability to use the baseline system.

4.1 Technical Setup and Environment

The prototype is developed using Unity 2020.3.4f1 and tested and deployed for use with the Meta Quest Pro in VR with two tracked handheld controllers. We made the decision to simulate AR using VR for several key reasons. VR offers a controlled environment that is unconstrained by technological limitations such as poor tracking, calibration, network connectivity between devices (e.g., IoT sensors), ergonomics, and field-of-view (e.g., with the HoloLens 2). Previous studies have cautioned against the use of AR as it may mask the potential benefits of SA due to these limitations [64, 74], whereby the benefits (and drawbacks) of SA are precisely what we hope to understand from our later exploratory study. Moreover, research has suggested that insights from AR studies can be

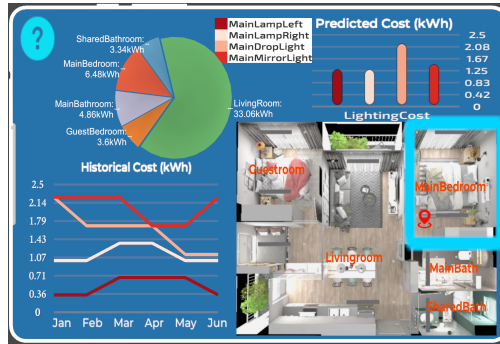


Fig. 2. The dashboard of situated visualizations (referred to as SitVis) that appears on a virtual tablet attached to the user's left hand with four linked views: a pie chart, bar chart, time series, and map.

replicated when simulating AR in VR [22, 44, 45]. Recent AR headsets like the Apple Vision Pro may have overcome some of these technical challenges, though the headset was not available when this prototype was developed and the study conducted. The use of VR also avoids the logistical issues of conducting a study in a real apartment, be it owned by the authors or temporarily rented, which would be personally invasive or expensive. While a less personal scenario would be an alternative (e.g., lights in an office), we wanted a scenario that would be relatable to almost any participant.

The virtual environment is an apartment, as seen in Figure 1 right, that is approximately 10 m by 12 m and consists of five rooms: a living room, two bedrooms, and two bathrooms. As per our scenario, the referents in question are the light bulbs that are spread throughout the apartment. There are 9 ceiling lights, 5 drop lights, and 9 counter lights and lamps. Note that we did not artificially distribute the lights to keep their locations as organic as possible (e.g., no drop lights in the middle of a bathroom). To adjust the brightness of each light, the user points at the desired light with their right-hand controller and presses the trigger button, cycling between 4 brightness levels. The lowest brightness level turns the light off. A raycast beam indicates what the controller is pointed at, and changes from red to green when the pointed object is a valid target. This can be seen in Figure 3 right. Some lights are treated together as a group where logical, such as the adjacent ceiling lights in a room, and thus cycle their brightness levels together. All light sources update in real time, illuminating their surrounding areas accordingly.

Using VR for large virtual environments necessitates using locomotion techniques, such as teleportation, which can cause cybersickness [38] and distract from the true aims of our study. We therefore ensured our prototype worked in a large real-world space, allowing users to physically and naturally walk from one end of the apartment to the other, thus more accurately simulating AR.

4.2 Situated Visualization on Tablet

On the left-hand controller of the user is attached a virtual tablet, as seen in Figure 2, which represents the SitVis used for the task. This mimics a person carrying a tablet with them with one hand. The tablet is 24 cm by 16 cm—large enough for text to be easier to read, but small enough to be similar in size to a real tablet. It can be held in any manner the user wishes, such as in Figure 1 left at chest height. The tablet's interface takes on a dashboard-style layout and consists of four separate views in the four separate corners.

The **pie chart** on the top left provides an overview of the energy consumption of lights across the five rooms of the apartment. It allows the user to understand the part-to-whole relationship



Fig. 3. Left: An example of an embedded visualization (referred to as EmbVis) that is attached to a bedside lamp. Right: The user points at a light with their right controller and presses the trigger to adjust the brightness level at incremental stages.

between each individual room and the rest of the apartment. Using the raycast controller, the user can also select a slice (i.e., room) to filter the other three linked views accordingly. The **bar chart** on the top right therefore shows the breakdown of each set of lights and their individual predicted energy usage per month. This is color-coded with the **time series** on the bottom left, showing the historical energy usage of each set of lights over the previous six months (as per the scenario in Section 3.2). Lastly, the **map** on the bottom right provides a spatial reference for which room the user had selected on the pie chart as indicated by the blue outline, and where they are currently located in the apartment as indicated by the red “location” icon which moves in real-time.

4.3 Embedded Visualizations in AR

Next to each light is a simple **donut chart**, an example seen in Figure 3 left, which are the EmbVis that would be seen in AR. It shows the predicted energy usage per month of the light. The donut, in particular, encodes the part-to-whole relationship of its current usage at the active brightness level versus its maximum possible usage. In cases where lights are logically grouped together (e.g., all illuminating the same area), the donut chart shows the sum of all lights, with leader lines indicating the respective lights that it is linked to (Figure 1 left). As this work is preliminary in nature, we intentionally chose not to include more complex visualizations, particularly those that are multi-dimensional, to focus on the usability of the overall SA setup and not on the understandability of the chosen visualizations.

5 Exploratory User Study

We now describe the exploratory study that we conducted using the SA prototype described in Section 4. The goal was to observe how people use the prototype, elicit the perceived benefits that such a SA application can have, and understand the drawbacks and challenges that need to be resolved to support future SA systems.

5.1 Study Design and Data Collected

Being exploratory in nature, we gave participants the same scenario and task in Section 3.2 that we believe to be common in the future. Unlike prior work (e.g., [61, 64]), we intentionally kept the task open-ended to mimic a realistic scenario in which the user themselves determines what they want to achieve with the tool, rather than force a predefined goal on them that they may feel

to be arbitrary. We discuss other possible task types later in Section 7 that would further explore the breadth of SA. The data provided to participants during the study was artificial, but chosen within ranges that would be plausible. This was particularly relevant for the time series data that needed to be created per light. Else, the energy usage of each light was dependent on the brightness level which the participant was tasked with adjusting. These levels were set to their maximums by default.

We collected subjective quantitative data in the form of three questionnaires. The Simulator Sickness Questionnaire [33] (SSQ) was used to determine if our use of VR had inadvertently induced cybersickness in our participants, which may bias their feedback. The NASA Task Load Index [25] (NASA-TLX) and System Usability Scale [7] (SUS) was used to compare our prototype with known benchmarks [4, 23] and help validate our prototype as a usable SA system. The NASA-TLX questionnaire included the pairwise weighting procedure. We also conducted semi-structured interviews after the task to obtain subjective qualitative data about their general thoughts on the prototype, suggestions for improvement, and whether they could see themselves using a similar system in AR in the future. Several questions were also guided by the considerations raised in Sections 1 and 2. These questions relate to the strategies participants used to: solve the task as a whole; to look at and use both SitVis and EmbVis together; and to navigate the apartment. Lastly, we recorded the first-person perspective of participants' movements and interactions with the system. The study was approved by our university's ethics review board.

5.2 Apparatus

The same setup described in Section 4.1 was used to conduct the study. The study was conducted in a university classroom that had a usable space of approximately 12 m by 14 m—larger than the virtual apartment. We used the Oculus Air Link feature to stream from the experimenter's PC running Unity to the Meta Quest Pro, allowing participants to freely walk around the virtual apartment without the risk of tripping over a tethered cable. The PC was equipped with an Intel Core i9-13980HX 2.20 GHz with 32 GB of RAM, with an NVIDIA RTX 4070 Laptop GPU with 8 GB of VRAM. When answering the questionnaires, participants entered their responses on the experimenter's tablet. We recorded the semi-structured interviews with a GoPro Hero7 Black, set up in such a manner to ensure a high-quality audio recording to be transcribed afterward.

5.3 Study Procedure

The study comprised five sections, totaling approximately 50 minutes.

Pre-study (5 minutes). Participants first signed a consent form and filled in a demographics questionnaire. The experimenter then helped the participant put on and adjust to the Meta Quest Pro.

Training (10 minutes). The experimenter then introduced the scenario of the study as per Section 3.2. They then guided the participant through the virtual environment and its visualizations and features. Participants were then given five minutes to freely explore, use the system, and ask questions. While the same apartment and lights were used during training, the energy usage data shown to participants was heavily simplified compared to the actual task.

Short break (5 minutes). Participants were then allowed to take a short break before the study began if requested.

Study (15 minutes). Participants were then instructed to complete the task as per Section 3.2. They were instructed to think aloud during the experiment and were also told that there was no correct or incorrect answer to the task. A text version of the task could be accessed at any time by selecting the “?” button on the top left of the tablet (Figure 2). The participants had to indicate to

the experimenter when they felt like they had adequately completed the task, which ended the study.

Post-study (15 minutes). Participants answered the SSQ, NASA-TLX, and SUS questionnaires on a physical tablet. They were then interviewed by the experimenter. At the end, participants received a US\$20 gift card as compensation for their time.

5.4 Participants

We recruited 19 participants (six female, 13 male) in total. To mitigate against the effects of novelty bias, especially as the majority of our data collection is subjective in nature, we sought to recruit participants who already had experience with VR and/or AR. 12 experienced participants (four female) were thus recruited from both a VR/AR course taught at the local university and from our VR/AR research group via internal mailing lists and word of mouth. These participants are labelled P1 to P12. To still capture a broad range of opinions however, we also recruited seven novice participants (two female) from the same university who had next to no prior experience with VR or AR via the same advertising methods. These participants are labelled P13 to P19. All participants had normal or corrected-to-normal vision. One participant was red-green color blind and had difficulties seeing the red and green raycast from the right-hand controller, though they did not report any major difficulties in reading the visualizations. One other participant was left-handed, though they reported no difficulties in using the prototype.

5.5 Data Analysis

For the semi-structured interviews, the audio recordings were first transcribed using Otter.ai and then cleaned by the author who conducted the interviews (the experimenter). The same author and a separate author who was not present for the interviews (the non-experimenter) then used thematic analysis [6] on the transcripts together using an inductive approach. The transcripts were read and discussed by the two authors together and, whenever a statement was deemed relevant, a code was assigned that the two authors agreed on. Codes were progressively merged as similarities were found throughout the analysis and were finally assigned to overall themes. The non-experimenter then reviewed, analyzed, and structured a narrative report of the themes, which was later reviewed and validated by the experimenter. For the video coding, a technical issue caused the video footage of P14 to not be saved. Regardless, the same experimenter first compiled an initial set of codes based on handwritten notes and a review of six of the 18 videos. The codes were then refined after discussion with the non-experimenter. The experimenter then applied the refined codebook to all 18 videos.

6 Results

We present our results from participant demographics (Section 6.1), questionnaire responses (Section 6.2), interview feedback (Section 6.3), and observations from video coding (Section 6.4).

6.1 Demographics

The mean age of our participants was 25.05 years old ($sd = 4.21$). Regarding their experience with AR and VR, six reported using either daily, five weekly, and two monthly, and the remaining six reported N/A. For using 2D data visualizations, four participants reported using them daily, seven weekly, four monthly, and the remaining four reported N/A. For using 3D data visualizations, only seven participants reported any prior experience.

6.2 Questionnaire Responses

The results from the three administered questionnaires are as follows.




Fig. 4. The distributions of the overall weighted scores and the six individual unweighted subscales from the NASA-TLX questionnaire [25].

SSQ. The total simulator sickness score [33] has a mean value of 20.47 ($sd = 17.11$). According to Stanney et al. [65], this puts the prototype slightly above the threshold of 20 for “a bad simulator”.

NASA-TLX. Figure 4 shows the distribution of scores for each of the six unweighted subscales and the overall weighted workload score. The weighted workload had a mean score of 31.65 ($sd = 18.07$), which according to Grier [23] places the prototype at slightly below the 20th percentile of global NASA-TLX scores which is 33.00.

SUS. The mean overall score for the SUS was 80.66 ($sd = 10.27$). When breaking this down by participants’ experience level to account for novelty bias, experienced participants gave a lower average score of 77.50 ($sd = 10.50$) than novice participants with 86.07 ($sd = 7.75$). According to Bangor et al. [4], the overall and experienced scores correspond to “Good” and the novice scores correspond to “Excellent”.

6.3 Interview Feedback

We now report on the relevant topics raised by our participants during the interview, categorized by the coded themes from our analysis. Note that in this section, quotes have been modified to remove filler words for readability, and that “tablet” is used synonymously for SitVis, and likewise for “donut chart” and EmbVis.

General impressions of the prototype. Participants provided positive feedback on the system. They thought it was “straightforward and easy to use” [P15], “seems like something that would be useful” [P3] and that “it was very neat, about the tablet [and] light kind of combination” [P5]. They also reported feeling immersed in VR, admitting that “I know I could walk through the bed or walk through the table, but I just kept walking around” [P1] and that they were “definitely immersed” [P17]. However, the use of VR did make it “hard to read [the tablet]” [P14], with at least eight participants making similar statements. There were also some technical issues as when “you turn too fast or walk too fast it kind of lags” [P16]. From a task perspective, several participants reported that the SA prototype was well suited as “the system allowed me to have the information I need when I needed it” [P7], with one participant acknowledging that “we have a tablet that can also display data, it makes sense to potentially capitalize on that” [P3]. Some comments also signified the benefits of situatedness, such as one making the observation that “some rooms will need more energy just because there’s more lights in them and they’re bigger” [P12] and another saying that “it’s good to see that the bedside lamp took a lot of energy, but it wasn’t that bright” [P17].

Impressions of the visualizations. Regarding the SitVis on the tablet, feedback was mixed and sometimes contradictory. For instance, P3 said that the pie chart was the most useful, while P12 did not see the purpose of it. Likewise, four participants found the bar chart helpful, while P15 “never even looked at [it]”. Only the line chart received clearly negative feedback, summed up by: “I was more focused on how I wanted to set up my space, not what someone else did before me” [P3]. In contrast, eight participants gave clearly positive remarks about the EmbVis, “I really liked

that although I had the tablet, there were [donut] charts all by the lights as well” [P15]. It is possible however that they interpreted the EmbVis not as a data visualization, but as a “gauge” that signifies the current lighting level of each light. Four participants made statements that suggest this, such as “it was nice to remember that I liked the lighting at a quarter, and I could go back to it easily” [P10]. This is reinforced by a fifth participant who thought the donut charts “were useful to know what lights I could control, but other than that, I didn’t really use the power numbers too much” [P5].

Navigating the space and its referents. This theme focuses on how participants decided **where** in the apartment they should go to first. While not all participants gave clear responses, we coded seven participants who did zero prioritization and would “just [walk] around and into the rooms exploring” [P17], which was the most common approach. Two participants said the opposite, “I would look at [the map] and say, okay, I’m gonna go there next” [P19]. Many participants also described how the map helped with the task. From seven participants, this was to “identify which room I’m in, either the main bedroom or the guest bedroom” [P5] as the two looked alike, which then let them “figure out which [room] to click on the pie chart” [P1]. Five participants used the map as a reminder tool, “during the training, I almost missed the bathroom next to the door, but then the map helped me realize there’s a bathroom right there” [P4]. On the other hand, four participants reported not using the map at all. Several comments hint as to why: P15 applied her intuition of how apartments are typically laid out, and P4 and P16 became familiarized with the apartment’s layout after the training phase.

Workflows for using visualizations. This theme focuses on **what** visualizations participants used and **how** they used them for the task, after having entered a room with lights in it. We coded four distinct workflows for how the SitVis and EmbVis were used, though we could not confidently assign one to each participant. The most common workflow from six participants was to first focus on the EmbVis, adjusting the lights as necessary, then validating their changes on the SitVis on the tablet. For example, “when I was when I thought I was done with the room, then I would check the chart on the tablet and see if I felt like that was good enough” [P12] and “I used the donut graph [...] and then I used the [tablet] a lot to see the actual statistical differences of changing lighting levels in the room [...] this was mostly just seeing if I had half lighting that I could only get maybe a quarter more lighting down” [P17]. Two participants reported the opposite workflow, with one saying “as I walked into a room, I would click on the pie chart, and I would look at what areas were taking more power and especially at anything that was at the top of those charts, and then I would look around the room and find the light that was presenting the highest amount of energy usage” [P7]. In contrast to these two workflows, several participants claimed not to have even bothered with one set of visualizations or the other. Three participants suggested they only used the tablet, as they “haven’t concentrated much on the graph close to the light” [P6]. Two others instead felt that they only needed the EmbVis, as “for me to accomplish what you told me to do, I don’t think the tablet is needed” [P18].

Priorities in decision-making. This theme now focuses on **why** participants adjusted the lights the way that they did. As with the preceding two themes, we identified three distinct mindsets based on participant responses. Five participants would first adjust the lights based on their own personal preferences, and then make further adjustments and refinements by referring to the energy data. For instance, “I was definitely prioritizing my preference over the cost. If I prefer [the] room to look like this, I’d make the cost work” [P1]. Four participants would instead find a balance between their personal preferences and energy usage, “I just tried to make it comfortable, but still use as little lighting as I could” [P17]. Perhaps the most extreme option, four participants reported having ignored the data entirely, only following their personal preferences. For instance, “I just went with my personal preference, whether I want the lights or not. And mostly, I keep my lights off” [P13].

Further improvements to the prototype. The most common suggestions relate to the design of the SitVis on the tablet and the interactions that it can support. This includes the ability to




Fig. 5. Left: P9 standing close to the light and EmbVis and needing to hold the SitVis high. Right: P6 standing far away from the light and EmbVis while holding the SitVis at a comfortable height.

click on the room in the map to link views together [P2, P3, P7, P10], showing the total energy consumption of the room on the bar chart [P8], showing dollar prices instead of kWh [P17], the ability to make pairwise comparisons of rooms [P5], or even using 3D visualizations [P6, P8]. The only suggestion specific to the EmbVis was to have them always rotate to face the user [P8]. One participant suggested using the position of the user as a means to control which information is shown, “so when I walk into the living room, then I see the pop ups, because there’s not really a reason for me to see what light in the other room is doing, it’s more about where I currently am” [P3]. Lastly, one participant suggested showing the location of the lights on the map itself, “if it was a bigger map, in each room you can see on the tablet, and it shows me where the light is on the map, that would be useful” [P8].

Future use of SA. While participants hinted at being willing to use an AR version of the prototype in the future, this came with caveats. An obvious caveat was due to the high price [P17] and relatively poor ergonomics of current AR headsets [P8]. This also depended on the actual situation, such as “if I had a house though, I would be willing to spend a little bit of money on an app that allowed me to visualize my cost like that” [P7]. Two participants did state that they would likely only use a similar prototype in more complex scenarios. For example, “if it was like an industry where there are multiple machines and there can be emergency hazards and stuff, where you have to actively go and interact with the equipment, having that information [...] can be really helpful” [P11]. This may also be the case in scenarios where “there was a temporal demand where I’ll have to look around or look for that information [...] I would like to have information right next to the object itself, but if it’s not, then I can just refer to it on my tablet or my phone and not have some stuff floating” [P11].

6.4 Observations and Video Coding

To deepen our analysis, we now report on observations and results from our video coding of participant behaviors and interactions with the system. We only highlight the key results from our analysis. Please see the supplemental material for information regarding our codes.

When and how to use visualizations. Aligned with the interview responses, participants showed great variation in the extent they used the visualizations. For the use of the SitVis, on one extreme, two participants [P13, P18] used the SitVis for less than 6% of the experiment. On the other extreme, seven participants [P1, P5, P6, P8, P10, P12, P17] used the SitVis for at least 40% of the experiment. For the participants who used the SitVis, when and how they did so varied. For

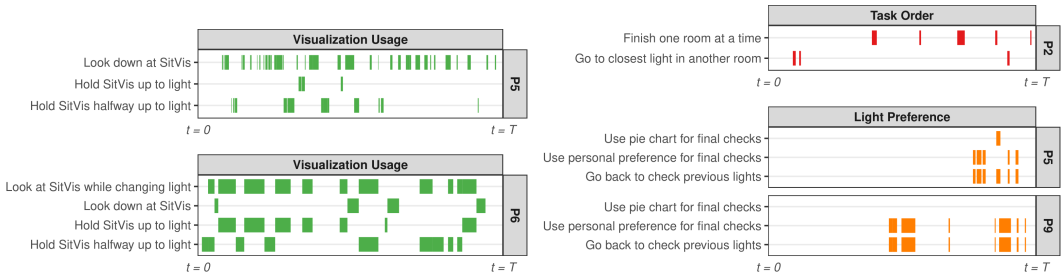


Fig. 6. Timelines of coded behaviors. See supplemental material for complete timelines of all participants.

instance, P3, P4, and P7 only looked at the SitVis after having adjusted a light. P6 in particular instead mainly used the SitVis while in the middle of adjusting a light. Others, such as P1, P5, and P10, would constantly refer back to the SitVis throughout the entire experiment, regardless of where or what they were doing (e.g., in Figure 6 top left).

Relative positioning of the visualizations. Participants also varied in how the SitVis was held. In general, when not interacting with a light, the SitVis was held at a relaxed hip height. When they were interacting with a light, participants held the SitVis at different heights, influencing the relative angular distance between the SitVis and EmbVis. To summarize, some participants, including P1, P5, and P17, would hold the SitVis relatively close to but not visibly aligned with the EmbVis. For instance, the SitVis would be held at chest height while looking at an EmbVis that is located close to the ceiling. In contrast, some participants, including P6, P8, and P10, would hold the SitVis as visibly close to the EmbVis as possible, even if the latter was close to the ceiling (e.g., in Figure 5 left), for at least 13% of the experiment. However, how they held the SitVis changed throughout the experiment (e.g., in Figure 6 bottom left). One developed strategy to minimize physical effort was to simply stand further away from the EmbVis, thus also minimizing the angular coverage needed while the SitVis was held at a more comfortable lower height (e.g., in Figure 5 right).

Navigation strategy. Aligned with our interview results, we observed two general approaches for prioritizing which lights to navigate to. Five participants [P2, P12, P13, P15, P19] appeared to be opportunistic, going from one light to the next that was visible in their peripheral vision, regardless of whether it was in the same room or not (e.g., in Figure 6 top right). This was especially so in instances where the lights were spread apart in a room, causing participants to potentially be unaware of them. On the other hand, four participants [P2, P3, P4, P19] had clearly first adjusted the SitVis to show the data of a specific room, before then navigating to it. However, this did not happen for every room, nor was there consistency between participants on when during the task this behavior was exhibited. For instance, P4 did it twice during the first half of the experiment, whereas P3 also did it twice but in the second half.

Refining light settings. We observed that at later stages of the experiment, many participants would return to lights that they had previously adjusted to make further refinements. Whether or not they had consciously referred to the data while doing so appeared to vary among them. Four participants [P5, P6, P17, P19] had deliberately checked the SitVis and its data while making these refinements. In contrast, nine participants [P2, P3, P4, P5, P9, P10, P11, P13, P18] did not directly look at the SitVis, but only at the light itself and its associated EmbVis. P5 was the only participant who did both, as shown in Figure 6 bottom right comparing them and the just mentioned P9.

Beyond the lights and visualizations. Throughout our coding, we also took note of several instances where participants interacted with the prototype outside of the lights and visualizations.

For example, five participants [P10, P13, P17, P18, P19] had looked at the furniture and décor that were in the background during the experiment. Some others also attempted to interact with the environment itself, such as trying to open drawers and closets. More importantly, 17 participants had, at some point, clearly panned their heads around after adjusting a light, presumably to observe the appearance of the room with the new lighting level.

7 Main Takeaways, Discussion, Guidelines, & Future Work

Based on our study, we now discuss key takeaways that serve as guidelines for HUIs and SA, and highlight avenues for future work.

The prototype was usable and valid. As a whole, the prototype that we had developed to embody the baseline of using HUIs for SA was functional and fit for purpose. While the SSQ results are poor by some measures [65], we note that VR studies frequently report average scores above 20 [5], thus we can reasonably say that the use of VR did not impair participants significantly. In contrast, the NASA-TLX and SUS scores were favorable, indicating that the prototype was functional and fit for purpose. The low task load however was likely influenced by the task aims being self-determined, with some participants more actively involved in the task than others. The high SUS scores may also be the result of the prototype’s novelty, though the average score of 77.50 from only the experienced participants is still promising. As a result, we believe that the combination of HUIs and SA has potential and warrants further investigation, though we acknowledge the potential differences when deploying in a physical mobile device and AR. In particular, future work should likely seek to understand the usability of needing to carry a physical device while interacting with physical objects—all while analyzing and making sense of data.

The needs and goals of users in SA are varied, and so too should their design and evaluation. We observed participants to have greatly contrasting interpretations of the presented task and scenario, which affected what aspects of the prototype they used and how they used them. For instance, some participants relied on the visualizations to help make decisions while others simply ignored them. This is because we chose not to provide strict instruction on what was deemed correct, such as through gamification [61, 64], which allowed our participants to make their own choices much like in the real world, thus bolstering our study’s ecological validity. Yet, no participant reported feeling unable to complete the task to a degree they were satisfied with. While perhaps unsurprising in hindsight, it does raise questions about the needs of users and how we test and validate SA applications. For one, it is clear that not everyone has an innate “analytical need” that a sophisticated SA application might accommodate. One person might only care to look at the summarized health rating of a grocery product, while another might look at and consider the breakdown of nutritional values. As our participants spanned this entire range, it is to be expected that the same be so in reality. Therefore, SA applications should be able to show as much or as little information as necessary for users to still meet their goals. By extension, when evaluating SA applications in a usability setting, the typical performance measures in visualization such as time, accuracy, or even “insights gained” are arguably less important, especially in scenarios where users set and define their own goals based on their own expertise and preferences. While we intentionally do not claim to provide a better measure for this—see a related article by Wang et al. [70] for further discussion—we do suggest that future work takes into account the varied goals of users when both designing and evaluating SA applications. We note, however, that our work takes the perspective of SA as an *Assistant* [63], and this may therefore not apply to other archetypes with more clearly defined use cases and goals such as *Simulators* or *Planners*.

SitVis for cognitive tasks, EmbVis for physical tasks. The EmbVis used in our study were intentionally simple and one-dimensional. Surprisingly, no participants suggested increasing the amount of information shown on the EmbVis, and instead gave positive remarks to its current design.

This indicates that they appreciated the simplicity of the EmbVis, likely as it directly supported the task that they were doing which was adjusting light levels. Any feature suggestions to improve the visualizations were only about adding further functionality to the SitVis, such as allowing pairwise comparisons of rooms [P5]. Given the hybrid setup of our prototype, we hypothesize that the SitVis embodied the role of the “data analysis” tool that supports the cognitive task, and the EmbVis the role of the “helper” tools that support the physical task. In this sense, a SA designer may consider offloading more complex data and visualizations onto the SitVis, while keeping the EmbVis straightforward and simple. The EmbVis may also take on a visual guidance role instead, instructing users on where and how to interact with referents through the use of glyphs and trajectories for instance [43]. Such an asymmetric distribution of visual elements and information has also been recommended by Rashid et al. [53] as a way to minimize attention-switching cost in HUIs. Of course, further research is needed to confirm this hypothesis for SA, particularly by varying the level of information shown on the SitVis versus EmbVis.

Attention-switching between SitVis and EmbVis was not a problem, though fatigue should still be considered. Results from our video coding showed that, with some exceptions, most participants used both SitVis and EmbVis together for the task. To our surprise, despite it being one of the underlying research interests of this work, no participants reported any difficulties in attention-switching between the two visualizations. Note that we did not directly ask them this to avoid any leading questions in our interviews. This is likely related to the preceding topic in that the EmbVis was simple and different from the SitVis. Regardless, we did observe participants holding the SitVis up to varying heights, whether it be halfway up to the light at a comfortable hip or chest height, or all the way up to head height when the EmbVis was close to the ceiling. Through all of this, no participants reported any fatigue during the interviews, though we note several outliers in the physical demand score in the NASA-TLX. It is also important to note that the Meta Quest Pro controller is comparatively lighter than, for instance, an Apple iPad Pro 11-inch (164 g versus 466 g) while having a more comfortable center of gravity. Nevertheless, it is likely necessary to further investigate possible solutions to mitigate fatigue for similar HUIs. When visual comparisons between SitVis and EmbVis are necessary, an obvious route is to minimize the distance between the two, but this is likely dependent on the task. Moving the EmbVis down next to the SitVis is one option, though this may be pointless if the user needed to, for example, see lighting changes in the environment, and thus keep their attention around the light anyway. Moving the SitVis up to the EmbVis is the other option, likely by creating a mirror image of the SitVis to still allow for touch input on the mobile device. Future research may investigate suitable techniques for combining HUIs and SA, especially when there is a need to support classical coordinated multiple view techniques [55] like brushing and linking [12]. As an aside, in Figure 1 right, the donut chart portion of the EmbVis is visibly distant from the four lights. This may raise the question of whether attention-switching was needed between the EmbVis and referent. This concern was not raised in our interviews, though we could not accurately confirm this from video coding due to lack of eye gaze data. Naturally, if the donut chart were further away (e.g., as far as *viewer space* [21]) it would no longer be considered embedded. Certain visualizations such as glyphs and decals are reliant on sharing the same spacial coordinate system as their referent [43], and thus attention should also be innately shared in these cases.

Support revisiting prior referents through bookmarks or annotations. Video coding confirmed that many participants, in the later parts of the task, would go back to and make final adjustments to the lights. While they did differ in whether these refinements were based on the data or purely on personal preference, the fact that this was a regular occurrence does signify its importance. In our scenario, we assume that participants could easily do this due to the pre-attentive nature of the task. That is, they could easily go from room to room and immediately

identify whether it is slightly too bright or too dark for their liking. This is likely not the case for a more complex task however that requires more than a cursory glance of past referents, such as re-inspecting individual yet similar-looking components of a machine. In these situations, a means to bookmark or annotate referents of interest may prove beneficial, thus supporting the *process and provenance* of SA [26]. Future work might investigate how best to keep track of provenance, whether it be the use of simple 3D drawings in AR, classic 2D annotations on the mobile device, or something completely new, such as the use of world-in-miniature-like representations of the environment for embodied provenance [77].

Personalizing SA based on the user’s familiarity with the environment. Because participants needed to physically navigate the apartment, we naturally decided to include a map to help them do so. As our interviews showed, many participants ended up not using it to navigate. While the apartment was initially new to them, they quickly grew familiar from both the training phase and their own understanding of modern apartment design. This allowed them to complete the task without much navigational or attentional assistance. This familiarity factor we now argue should be considered in SA as a means of personalization. While the homeowner themselves would likely not need any navigational or attentional guidance, a tradesperson who has never seen the place before likely would. We note that just providing this information “just in case” may ultimately cause visual and information overload, especially when using AR cues [12]. Future work should therefore seek to investigate how SA applications can adapt to the changing needs of different individuals, whether it be based on their familiarity or something else.

There is a need to understand the relationship between physical environment and SA. It goes without saying that our choice of scenario and environment directly influenced how our participants behaved, which in turn influenced our observations and takeaways. We had, in Section 3.1, characterized our scope and setting to illustrate the types of real-world contexts to which we think our work applies. Despite this, it is very likely that just by changing one or two of these parameters would user behavior also change. For example, not needing to interact with the environment may have participants complain about needing to move, thus making a *ProxSituating* approach more appealing [57]. Even changing the environment itself may alter user behavior and preference, like how our participants relied on the map to differentiate between similar-looking rooms; rooms with visually distinct appearances would not face this issue. From a research perspective, this suggests the need to empirically evaluate how the differing characteristics of physical environments influence the usability of SA, akin to the evaluation of highlighting techniques based on referent layouts by Doerr et al. [12]. From a design perspective, this suggests the need for SA applications to intelligently adapt to changing environments and goals, as no two settings will be the exact same. This could be through integration with artificial intelligence (as suggested by Shin et al. [63]), through visualizations whose layout [11, 18] or even appearance [36] adapts to the real world, or something else entirely.

Limitations regarding the use of VR. While we do believe our results to be transferrable to AR as suggested in other research [22, 44, 45], it is undeniable that VR has had at least some influence on our results. For example, several participants were distracted by trying to interact with virtual furniture, though this could also be seen as an indication that they were truly immersed. In such cases however, this did not significantly impact our ability to elicit relevant feedback. Using VR had also meant the omission of several factors that would have needed to be considered when using AR. Chief among them is the need to physically hold and touch a tablet to begin with, especially considering the aforementioned weight and balance differences between it and a VR controller. It is likely that using a physical tablet with AR would have users behaving in ways not observed in our study, such as them setting the tablet down on a table, or them resting on a sofa to look at the SitVis before deciding to get up and walk around to look at the EmbVis, or even them opening

a web browser to look up real-world energy prices. Using VR also resulted in us being unable to include a task that required the physical manipulation of referents, such as the scenario used by Sousa Calepso et al. [64]. It may be that a task that is more *pragmatic* [35] (i.e., working towards a physical goal) in nature would influence the usability of SitVis and EmbVis in this setup, such as if the user needed to manually install light bulbs of different color temperatures. Likewise, a task that is more *epistemic* [35] (i.e., discovering information) in nature may also see different feedback, such as if the user needed to diagnose an underlying problem. Despite these limitations, we still believe that our study provides insights into the roles that both SitVis and EmbVis have between each other and how users might perceive their utility for their own goals, even if users' exact behaviors and movement patterns differ in reality. We also reiterate the technical challenges that Sousa Calepso et al. [64] described in their study, as these mask any potential benefit that SA might have.

Limitations regarding the study design. As described in Sections 5.1 and 7, the task was deliberately open-ended to facilitate a realistic scenario where participants decided for themselves what they considered to be “correct”. Consider a smart home or health monitoring app currently on the market, for example—the user is free to make what they will of the presented information and is not forced into a single expected outcome. Regardless, we acknowledge that our study only captures one particular variety of SA application, as ours falls under the category of an *Assistant* [63] that simply provides information about physical referents at their respective locations. We initially considered but ultimately opted not to include other tasks that would be characterized more as *Simulators* or *Planners*, which would also include more complex visualization types. One possible task to optimize the placement of network routers to maximize for Wi-Fi network coverage (as a simplified version of the task used by Schröder et al. [61]), which would have used a birds-eye heatmap of signal strength as the SitVis and an equivalent heatmap on the floor as the EmbVis (akin to that by Luo et al. [49]). Another example was a task to diagnose and resolve an electrical wiring fault in the house using an X-ray style EmbVis of cables (akin to that by Schall et al. [60]). Their omission was in part due to the lack of a taxonomy for SA that would allow us to confidently and systematically investigate a variety of task types, especially as user behavior and preferences would likely vary drastically between task and visualization designs. Future work should first seek to establish this task taxonomy for SA to facilitate the design and comparison of different SA applications and scenarios that goes beyond that of Shin et al. [63], such as whether the task is pragmatic or epistemic [35, 64].

8 Conclusion

Motivated by the potential of HUIs to support SA scenarios, we investigated how people may organically use both SitVis and EmbVis together in an optimization task, with an aim to understand the root challenges and opportunities. We developed a prototype that used VR to simulate AR in an apartment setting, which contained smart lights scattered across multiple rooms. The task was to optimize the energy usage of lights by adjusting their lighting levels. The SitVis comprised of a dashboard of multiple views that appeared on a handheld virtual tablet, and the EmbVis were donut charts attached to their respective lights. We used this prototype in an exploratory study of 19 participants, who could walk around freely in the apartment and adjust the lights in any way they wished based on the data that was available. Our findings suggest that when users could decide what they considered to be the “correct” solution to the task, how they approached it varied greatly, with some taking the data into close consideration with others simply ignoring it. Yet, they were still satisfied with their outcomes and thus both approaches we see as equally valid. We also saw evidence that the hybrid user interface setup was indeed usable by participants, with known issues such as attention-switching cost not being identified as a concern. From this, we believe our work shows the potential of further combing heterogeneous devices for SA in the future. Our results

also demonstrate the need to widen perspectives on designing and evaluating SA applications, taking into account the myriad differences in user aims, goals, and physical environments. Future work should seek to understand how any given SA system might still fare should any one of these variables change, both in terms of human factors and the system's capability to adapt.

Acknowledgments

We acknowledge the National Science Foundation awards 2327569, 2238313, 2223432, 2223459, 2037417, and 1948254. This work was also funded by the German Research Foundation (DFG) project 495135767 and the Austria Science Fund (FWF) project I 5912-N (joint Weave project), and partially supported by Germany's Excellence Strategy – EXC 2120/1 – 390831618 (DFG).

This paper was prepared for informational purposes by the Global Technology Applied Research Center of J.P. Morgan Chase & Co. This paper is not a product of the Research Department of J.P. Morgan Chase & Co. or its affiliates. Neither J.P. Morgan Chase & Co. nor any of its affiliates make any explicit or implied representation or warranty and none of them accept any liability in connection with this paper, including without limitation, the completeness, accuracy, reliability of the information contained herein and the potential legal, compliance, tax or accounting effects thereof. This document is not intended as investment research or investment advice, or as a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction.

References

- [1] Apple. 2024. Use Your Mac with Apple Vision Pro. Accessed 2024-10-01. <https://support.apple.com/en-us/HT213971>
- [2] Ambre Assor, Arnaud Prouzeau, Martin Hachet, and Pierre Dragicevic. 2024. Handling Non-Visible Referents in Situated Visualizations. *IEEE TVCG* 30, 1 (Jan. 2024), 1336–1346. <https://doi.org/10.1109/TVCG.2023.3327361>
- [3] Benjamin Bach, Ronell Sicat, Hanspeter Pfister, and Aaron Quigley. 2017. Drawing into the AR-CANVAS: Designing Embedded Visualizations for Augmented Reality. In *Workshop on Immersive Analytics, IEEE Vis*.
- [4] Aaron Bangor, Philip Kortum, and James Miller. 2009. Determining What Individual SUS Scores Mean: Adding an Adjective Rating Scale. *Journal of Usability Studies* 4, 3 (May 2009), 114–123.
- [5] Pauline Bimberg, Tim Weissker, and Alexander Kulik. 2020. On the Usage of the Simulator Sickness Questionnaire for Virtual Reality Research. In *Proc. VRW*. 464–467. <https://doi.org/10.1109/VRW50115.2020.00098>
- [6] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- [7] John Brooke. 1996. SUS: A 'Quick and Dirty' Usability Scale. In *Usability Evaluation In Industry*. CRC Press.
- [8] Andy Brown, Michael Evans, Caroline Jay, Maxine Glancy, Rhianne Jones, and Simon Harper. 2014. HCI over Multiple Screens. In *Proc. CHI EA*. ACM, 665–674. <https://doi.org/10.1145/2559206.2578869>
- [9] Wolfgang Büschel, Anke Lehmann, and Raimund Dachselt. 2021. MIRIA: A Mixed Reality Toolkit for the In-Situ Visualization and Analysis of Spatio-Temporal Interaction Data. In *Proc. CHI*. ACM, 1–15. <https://doi.org/10.1145/3411764.3445651>
- [10] Wolfgang Büschel, Annett Mitschick, Thomas Meyer, and Raimund Dachselt. 2019. Investigating Smartphone-based Pan and Zoom in 3D Data Spaces in Augmented Reality. In *Proc. MobileHCI*. ACM, 1–13. <https://doi.org/10.1145/3338286.3340113>
- [11] Yifei Cheng, Yukang Yan, Xin Yi, Yuanchun Shi, and David Lindlbauer. 2021. SemanticAdapt: Optimization-based Adaptation of Mixed Reality Layouts Leveraging Virtual-Physical Semantic Connections. In *Proc. UIST*. ACM, 282–297. <https://doi.org/10.1145/3472749.3474750>
- [12] Nina Doerr, Benjamin Lee, Katarina Baricova, Dieter Schmalstieg, and Michael Sedlmair. 2024. Visual Highlighting for Situated Brushing and Linking. *Comput. Graphics Forum* 43, 3 (June 2024). <https://doi.org/10.1111/cgf.15105>
- [13] Niklas Elmqvist. 2023. Data Analytics Anywhere and Everywhere. *Commun. ACM* 66, 12 (Nov. 2023), 52–63. <https://doi.org/10.1145/3584858>
- [14] Niklas Elmqvist and Pourang Irani. 2013. Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime. *Computer* 46, 4 (April 2013), 86–89. <https://doi.org/10.1109/MC.2013.147>
- [15] Neven ElSayed, Bruce Thomas, Kim Marriott, Julia Piantadosi, and Ross Smith. 2015. Situated Analytics. In *Proc. BDVA*. 1–8. <https://doi.org/10.1109/BDVA.2015.7314302>

- [16] Neven 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 (Oct. 2016), 13–23. <https://doi.org/10.1016/j.jvlc.2016.07.006>
- [17] Barrett Ens, Benjamin Bach, Maxime Cordeil, Ulrich Engelke, Marcos Serrano, Wesley Willett, Arnaud Prouzeau, Christoph Anthes, Wolfgang Büschel, Cody Dunne, Tim Dwyer, Jens Grubert, Jason H. Haga, Nurit Kirshenbaum, Dylan Kobayashi, Tica Lin, Monsurat Olaosebikan, Fabian Pointecker, David Saffo, Nazmus Saquib, Dieter Schmalstieg, Danielle Albers Szafir, Matt Whitlock, and Yalong Yang. 2021. Grand Challenges in Immersive Analytics. In *Proc. CHI*. ACM, 1–17. <https://doi.org/10.1145/3411764.3446866>
- [18] João Marcelo Evangelista Belo, Mathias N. Lystbæk, Anna Maria Feit, Ken Pfeuffer, Peter Kán, Antti Oulasvirta, and Kaj Grønbaek. 2022. AUIT – the Adaptive User Interfaces Toolkit for Designing XR Applications. In *Proc. UIST*. ACM, 1–16. <https://doi.org/10.1145/3526113.3545651>
- [19] Steven Feiner and Ari Shamash. 1991. Hybrid User Interfaces: Breeding Virtually Bigger Interfaces for Physically Smaller Computers. In *Proc. UIST*. ACM Press, 9–17. <https://doi.org/10.1145/120782.120783>
- [20] Jorge A. Wagner Filho, Wolfgang Stuerzlinger, and Luciana Nedel. 2020. Evaluating an Immersive Space-Time Cube Geovisualization for Intuitive Trajectory Data Exploration. *IEEE TVCG* 26, 01 (Jan. 2020), 514–524. <https://doi.org/10.1109/TVCG.2019.2934415>
- [21] Zeinab Ghaemi, Kadek Ananta Satriadi, Ulrich Engelke, Barrett Ens, and Bernhard Jenny. 2023. Visualization Placement for Outdoor Augmented Data Tours. In *Proc. SUJ*. ACM, 1–14. <https://doi.org/10.1145/3607822.3614518>
- [22] Jerônimo G Grandi, Zekun Cao, Mark Ogren, and Regis Kopper. 2021. Design and Simulation of Next-Generation Augmented Reality User Interfaces in Virtual Reality. In *Proc. VRW*. 23–29. <https://doi.org/10.1109/VRW52623.2021.00011>
- [23] Rebecca A. Grier. 2015. How High Is High? A Meta-Analysis of NASA-TLX Global Workload Scores. *Proc. Hum. Factors Ergon. Soc. Annu. Meet.* 59, 1 (Sept. 2015), 1727–1731. <https://doi.org/10.1177/1541931215591373>
- [24] Jens Grubert, Matthias Kranz, and Aaron Quigley. 2016. Challenges in Mobile Multi-Device Ecosystems. *mUX: The Journal of Mobile User Experience* 5, 1 (Aug. 2016), 5. <https://doi.org/10.1186/s13678-016-0007-y>
- [25] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*. North-Holland, 139–183. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9)
- [26] Jeffrey Heer and Ben Shneiderman. 2012. Interactive Dynamics for Visual Analysis. *Commun. ACM* 55, 4 (April 2012), 45–54. <https://doi.org/10.1145/2133806.2133821>
- [27] Jiayi Hong, Rostyslav Hnatyshyn, Ebrar A. D. Santos, Ross Maciejewski, and Tobias Isenberg. 2024. A Survey of Designs for Combined 2D+3D Visual Representations. *IEEE TVCG* 30, 6 (June 2024), 2888–2902. <https://doi.org/10.1109/tvcg.2024.3388516>
- [28] Jinbin Huang, Shuang Liang, Qi Xiong, Yu Gao, Chao Mei, Yi Xu, and Chris Bryan. 2022. SPARVIS: Combining Smartphone and Augmented Reality for Visual Data Analytics. In *Proc. ISMAR-Adjunct*. 111–117. <https://doi.org/10.1109/ISMAR-Adjunct57072.2022.00030>
- [29] Sebastian Hubenschmid, Johannes Zagermann, Simon Butscher, and Harald Reiterer. 2021. STREAM: Exploring the Combination of Spatially-Aware Tablets with Augmented Reality Head-Mounted Displays for Immersive Analytics. In *Proc. CHI*. ACM, 1–14. <https://doi.org/10.1145/3411764.3445298>
- [30] Christophe Hurter, Nathalie Henry Riche, Steven M. Drucker, Maxime Cordeil, Richard Alligier, and Romain Vuillemot. 2019. FiberClay: Sculpting Three Dimensional Trajectories to Reveal Structural Insights. *IEEE TVCG* 25, 1 (Jan. 2019), 704–714. <https://doi.org/10.1109/TVCG.2018.2865191>
- [31] Sungwon In, Tica Lin, Chris North, Hanspeter Pfister, and Yalong Yang. 2023. This Is the Table I Want! Interactive Data Transformation on Desktop and in Virtual Reality. *IEEE TVCG* (2023), 1–16. <https://doi.org/10.1109/TVCG.2023.3299602>
- [32] Marco Jahn, Marc Jentsch, Christian R. Prause, Ferry Pramudianto, Amro Al-Akkad, and Rene Reiners. 2010. The Energy Aware Smart Home. In *Proc. FutureTech*. 1–8. <https://doi.org/10.1109/FUTURETECH.2010.5482712>
- [33] Robert S. Kennedy, Norman E. Lane, Kevin S. Berbaum, and Michael G. Lilienthal. 1993. Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness. *The International Journal of Aviation Psychology* 3, 3 (July 1993), 203–220. https://doi.org/10.1207/s15327108ijap0303_3
- [34] G.R. King, W. Piekarski, and B.H. Thomas. 2005. ARVino - Outdoor Augmented Reality Visualisation of Viticulture GIS Data. In *Proc. ISMAR*. 52–55. <https://doi.org/10.1109/ISMAR.2005.14>
- [35] David Kirsh and Paul Maglio. 1994. On Distinguishing Epistemic from Pragmatic Action. *Cognitive Science* 18, 4 (Oct. 1994), 513–549. [https://doi.org/10.1016/0364-0213\(94\)90007-8](https://doi.org/10.1016/0364-0213(94)90007-8)
- [36] Ari Kouts, Lonni Besançon, Michael Sedlmair, and Benjamin Lee. 2023. LSDvis: Hallucinatory Data Visualisations in Real World Environments. In *Alt.VIS 2023, an IEEE VIS Workshop*. arXiv. <https://doi.org/10.48550/arXiv.2312.11144> arXiv:2312.11144

- [37] Ricardo Langner, Marc Satkowski, Wolfgang Büschel, and Raimund Dachsel. 2021. MARVIS: Combining Mobile Devices and Augmented Reality for Visual Data Analysis. In *Proc. CHI*. ACM, 1–17. <https://doi.org/10.1145/3411764.3445593>
- [38] Joseph J. LaViola, Ernst Kruijff, Ryan P. McMahan, Doug A. Bowman, and Ivan Poupyrev. 2017. *3D User Interfaces: Theory and Practice* (second edition ed.). Addison-Wesley.
- [39] Bongshin Lee, Matthew Brehmer, Petra Isenberg, Eun Kyoung Choe, Ricardo Langner, and Raimund Dachsel. 2018. Data Visualization on Mobile Devices. In *Proc. CHI EA*. ACM, 1–8. <https://doi.org/10.1145/3170427.3170631>
- [40] Benjamin Lee, Maxime Cordeil, Arnaud Prouzeau, Bernhard Jenny, and Tim Dwyer. 2022. A Design Space For Data Visualisation Transformations Between 2D And 3D In Mixed-Reality Environments. In *Proc. CHI*. ACM, 1–14. <https://doi.org/10.1145/3491102.3501859>
- [41] Bongshin Lee, Raimund Dachsel, Petra Isenberg, and Eun Kyoung Choe (Eds.). 2022. *Mobile Data Visualization* (first edition ed.). CRC Press.
- [42] Benjamin Lee, Xiaoyun Hu, Maxime Cordeil, Arnaud Prouzeau, Bernhard Jenny, and Tim Dwyer. 2021. Shared Surfaces and Spaces: Collaborative Data Visualisation in a Co-located Immersive Environment. *IEEE TVCG* 27, 2 (Feb. 2021), 1171–1181. <https://doi.org/10.1109/TVCG.2020.3030450>
- [43] Benjamin Lee, Michael Sedlmair, and Dieter Schmalstieg. 2024. Design Patterns for Situated Visualization in Augmented Reality. *IEEE TVCG* 30, 1 (Jan. 2024), 1324–1335. <https://doi.org/10.1109/TVCG.2023.3327398>
- [44] Cha Lee, Scott Bonebrake, Tobias Hollerer, and Doug A. Bowman. 2009. A Replication Study Testing the Validity of AR Simulation in VR for Controlled Experiments. In *Proc. ISMAR*. IEEE, 203–204. <https://doi.org/10.1109/ISMAR.2009.5336464>
- [45] Cha Lee, Gustavo A. Rincon, Greg Meyer, Tobias Hollerer, and Doug A. Bowman. 2013. The Effects of Visual Realism on Search Tasks in Mixed Reality Simulation. *IEEE TVCG* 19, 4 (April 2013), 547–556. <https://doi.org/10.1109/TVCG.2013.41>
- [46] Lee Lisle, Kylie Davidson, Edward J.K. Gitre, Chris North, and Doug A. Bowman. 2021. Sensemaking Strategies with Immersive Space to Think. In *Proc. VR*. 529–537. <https://doi.org/10.1109/VR50410.2021.00077>
- [47] Jiazhou Liu, Barrett Ens, Arnaud Prouzeau, Jim Smiley, Isobel Kara Nixon, Sarah Goodwin, and Tim Dwyer. 2023. DataDancing: An Exploration of the Design Space For Visualisation View Management for 3D Surfaces and Spaces. In *Proc. CHI*. ACM, 1–17. <https://doi.org/10.1145/3544548.3580827>
- [48] Weizhou Luo, Anke Lehmann, Hjalmar Widengren, and Raimund Dachsel. 2022. Where Should We Put It? Layout and Placement Strategies of Documents in Augmented Reality for Collaborative Sensemaking. In *Proc. CHI*. ACM, 1–16. <https://doi.org/10.1145/3491102.3501946>
- [49] Weizhou Luo, Zhongyuan Yu, Rufat Rzayev, Marc Satkowski, Stefan Gumhold, Matthew McGinity, and Raimund Dachsel. 2023. Pearl: Physical Environment Based Augmented Reality Lenses for In-Situ Human Movement Analysis. In *Proc. CHI*. ACM, 1–15. <https://doi.org/10.1145/3544548.3580715>
- [50] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. Henry Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas. 2018. *Immersive Analytics*. Springer International Publishing. <https://books.google.com.au/books?id=vaVyDwAAQBAJ>
- [51] Tamara Munzner. 2014. *Visualization Analysis and Design*. CRC Press, Taylor & Francis Group, CRC Press is an imprint of the Taylor & Francis Group, an informa business.
- [52] Wolfgang Narzt, Gustav Pomberger, Alois Ferscha, Dieter Kolb, Reiner Müller, Jan Wiegardt, Horst Hörtner, and Christopher Lindinger. 2006. Augmented Reality Navigation Systems. *Univers. Access Inf. Soc.* 4, 3 (March 2006), 177–187. <https://doi.org/10.1007/s10209-005-0017-5>
- [53] Umar Rashid, Miguel A. Nacenta, and Aaron Quigley. 2012. The Cost of Display Switching: A Comparison of Mobile, Large Display and Hybrid UI Configurations. In *Proc. AVI*. ACM, 99–106. <https://doi.org/10.1145/2254556.2254577>
- [54] Umar Rashid, Miguel A. Nacenta, and Aaron Quigley. 2012. Factors Influencing Visual Attention Switch in Multi-Display User Interfaces: A Survey. In *Proc. PERDIS*. ACM, 1–6. <https://doi.org/10.1145/2307798.2307799>
- [55] Jonathan C. Roberts. 2007. State of the Art: Coordinated & Multiple Views in Exploratory Visualization. In *Proc. CMV*. 61–71. <https://doi.org/10.1109/CMV.2007.20>
- [56] Irene Ros. [n. d.]. MobileVis. Accessed 2024-10-01. <https://mobilevis.is/>
- [57] Kadek Ananta Satriadi, Andrew Cunningham, Ross T. Smith, Tim Dwyer, Adam Drogemuller, and Bruce H. Thomas. 2023. ProxSituational Visualization: An Extended Model of Situated Visualization Using Proxies for Physical Referents. In *Proc. CHI*. ACM, 1–20. <https://doi.org/10.1145/3544548.3580952>
- [58] Kadek Ananta Satriadi, Barrett Ens, Maxime Cordeil, Tobias Czuderna, and Bernhard Jenny. 2020. Maps Around Me: 3D Multiview Layouts in Immersive Spaces. *Proc. ACMHCI* 4, ISS (Nov. 2020), 201:1–201:20. <https://doi.org/10.1145/3427329>
- [59] Anika Sayara, Benjamin Lee, Carlos Quijano-Chavez, and Michael Sedlmair. 2023. Designing Situated Dashboards: Challenges and Opportunities. In *Proc. ISMAR-Adjunct*. 97–102. <https://doi.org/10.1109/ISMAR-Adjunct60411.2023.00028>
- [60] Gerhard Schall, Erick Mendez, and Dieter Schmalstieg. 2008. Virtual Redlining for Civil Engineering in Real Environments. In *Proc. ISMAR*. 95–98. <https://doi.org/10.1109/ISMAR.2008.4637332>

- [61] Jan-Henrik Schröder, Daniel Schacht, Niklas Peper, Anita Marie Hamurculu, and Hans-Christian Jetter. 2023. Collaborating Across Realities: Analytical Lenses for Understanding Dyadic Collaboration in Transitional Interfaces. In *Proc. CHI*. ACM, 1–16. <https://doi.org/10.1145/3544548.3580879>
- [62] Mickael Sereno, Stéphane Gosset, Lonni Besançon, and Tobias Isenberg. 2022. Hybrid Touch/Tangible Spatial Selection in Augmented Reality. *Comput. Graphics Forum* 41, 3 (2022), 403–415. <https://doi.org/10.1111/cgf.14550>
- [63] Sungbok Shin, Andrea Batch, Peter W. S. Butcher, Panagiotis D. Ritsos, and Niklas Elmquist. 2023. The Reality of the Situation: A Survey of Situated Analytics. *IEEE TVCG* (2023), 1–19. <https://doi.org/10.1109/TVCG.2023.3285546>
- [64] Aimee Sousa Calepso, Philipp Fleck, Dieter Schmalstieg, and Michael Sedlmair. 2023. Exploring Augmented Reality for Situated Analytics with Many Movable Physical Referents. In *Proc. VRST*. ACM, 1–12. <https://doi.org/10.1145/3611659.3615700>
- [65] Kay M. Stanney, Robert S. Kennedy, and Julie M. Drexler. 1997. Cybersickness Is Not Simulator Sickness. *Proc. Hum. Factors Ergon. Soc. Annu. Meet.* 41, 2 (Oct. 1997), 1138–1142. <https://doi.org/10.1177/107118139704100292>
- [66] Desney S. Tan and Mary Czerwinski. 2003. Effects of Visual Separation and Physical Discontinuities When Distributing Information across Multiple Displays. In *Proc. INTERACT*, Matthias Rauterberg, Marino Menozzi, and Janet Wesson (Eds.). IOS Press.
- [67] Markus Tatzgern, Valeria Orso, Denis Kalkofen, Giulio Jacucci, Luciano Gamberini, and Dieter Schmalstieg. 2016. Adaptive Information Density for Augmented Reality Displays. In *Proc. VR*. 83–92. <https://doi.org/10.1109/VR.2016.7504691>
- [68] Jia Wang and Robert Lindeman. 2014. Coordinated 3D Interaction in Tablet- and HMD-based Hybrid Virtual Environments. In *Proc. SUI*. ACM, 70–79. <https://doi.org/10.1145/2659766.2659777>
- [69] Xiyao Wang, Lonni Besançon, David Rousseau, Mickael Sereno, Mehdi Ammi, and Tobias Isenberg. 2020. Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools. In *Proc. CHI*. ACM, 1–13. <https://doi.org/10.1145/3313831.3376657>
- [70] Yun Wang, Adrien Segal, Roberta Klatzky, Daniel F. Keefe, Petra Isenberg, Jorn Hurtienne, Eva Hornecker, Tim Dwyer, and Stephen Barras. 2019. An Emotional Response to the Value of Visualization. *IEEE CG&A* 39, 5 (Sept. 2019), 8–17. <https://doi.org/10.1109/MCG.2019.2923483>
- [71] Sean White and Steven Feiner. 2009. SiteLens: Situated Visualization Techniques for Urban Site Visits. In *Proc. CHI*. ACM, 1117–1120. <https://doi.org/10.1145/1518701.1518871>
- [72] Sean Michael White. 2009. *Interaction and Presentation Techniques for Situated Visualization*. Ph.D. Dissertation. Columbia University, USA.
- [73] Matt Whitlock, Stephen Smart, and Danielle Albers Szafr. 2020. Graphical Perception for Immersive Analytics. In *Proc. VR*. 616–625. <https://doi.org/10.1109/VR46266.2020.00084>
- [74] Matt Whitlock, Keke Wu, and Danielle Albers Szafr. 2020. Designing for Mobile and Immersive Visual Analytics in the Field. *IEEE TVCG* 26, 1 (Jan. 2020), 503–513. <https://doi.org/10.1109/TVCG.2019.2934282>
- [75] Wesley Willett, Yvonne Jansen, and Pierre Dragicevic. 2017. Embedded Data Representations. *IEEE TVCG* 23, 1 (Jan. 2017), 461–470. <https://doi.org/10.1109/TVCG.2016.2598608>
- [76] Johannes Zagermann, Sebastian Hubenschmid, Priscilla Balestrucci, Tiare Feuchtnner, Sven Mayer, Marc O. Ernst, Albrecht Schmidt, and Harald Reiterer. 2022. Complementary Interfaces for Visual Computing. *it - Information Technology* 64, 4-5 (Aug. 2022), 145–154. <https://doi.org/10.1515/itit-2022-0031>
- [77] Yidan Zhang, Barrett Ens, Kadek Ananta Satriadi, Ying Yang, and Sarah Goodwin. 2023. Embodied Provenance for Immersive Sensemaking. *Proc. ACMHCI* 7, ISS (Nov. 2023), 435:198–435:216. <https://doi.org/10.1145/3626471>
- [78] Xiaoyan Zhou, Anil Ufuk Batmaz, Adam Sinclair Williams, Dylan Schreiber, and Francisco Raul Ortega. 2024. I Did Not Notice: A Comparison of Immersive Analytics with Augmented and Virtual Reality. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 193, 7 pages. <https://doi.org/10.1145/3613905.3651085>
- [79] Qian Zhu, Zhuo Wang, Wei Zeng, Wai Tong, Weiyue Lin, and Xiaojuan Ma. 2024. Make Interaction Situated: Designing User Acceptable Interaction for Situated Visualization in Public Environments. In *Proc. CHI*. ACM. <https://doi.org/10.1145/3613904.3642049>

Received 2024-07-01; accepted 2024-09-20