Spatial Context Switches During Knowledge Tasks in Extended Reality

Wolfgang Büschel
VISUS
University of Stuttgart
Stuttgart, Germany
wolfgang.bueschel@visus.uni-stuttgart.de

Verena Biener
VISUS
University of Stuttgart
Stuttgart, Germany
verena.biener@visus.uni-stuttgart.de

Dieter Schmalstieg
VISUS
University of Stuttgart
Stuttgart, Germany
dieter.schmalstieg@visus.uni-stuttgart.de

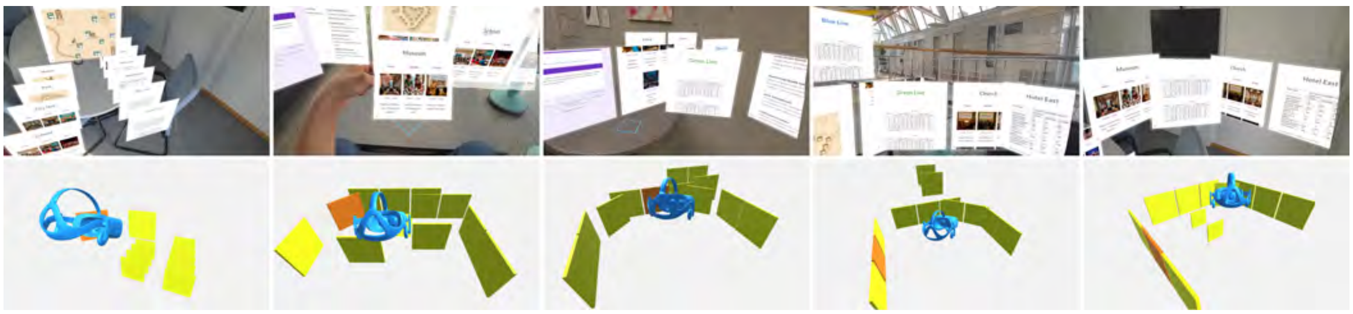

Figure 1: In our study, participants solved a knowledge task in XR while switching between different locations (from left to right: meeting room, seating area, standing table, seating area, meeting room). A selection of layout strategies that we observed is depicted here. The first row shows frames of the recorded video, the second row shows corresponding layout data reconstructed from the logs. Times are not synchronized.

Abstract

Current Extended Reality (XR) devices are increasingly being used as productivity tools. Compared to conventional setups, they allow for more flexibility in dynamic work locations beyond the desktop while providing a large virtual workspace. Recent research has explored how users organize digital documents and how virtual interfaces could be adapted to different locations and scenarios. However, there has been limited research on how location changes affect productivity tasks in XR environments and how users manually adapt virtual content layouts after such task interruptions. To address this, we conducted an exploratory user study (N=17) in which participants worked on a document-centered organization and planning task while changing locations every five minutes. We examined how these spatial transitions interfered with the task and identified layout strategies and patterns. From our observations and participant responses, we derived a set of design guidelines to inform the development of future XR knowledge work systems in mobile contexts.

CCS Concepts

• **Human-centered computing → Mixed / augmented reality**; **Empirical studies in HCI**.

ACM ISBN 979-8-4007-2278-3/2026/04
https://doi.org/10.1145/3772318.3791416

Keywords

extended reality, spatial layout, knowledge work, mobility

ACM Reference Format:
Wolfgang Büschel, Verena Biener, and Dieter Schmalstieg. 2026. Spatial Context Switches During Knowledge Tasks in Extended Reality. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3772318.3791416

1 Introduction

The newest head-mounted display (HMD) generation for Extended Reality (XR, including technologies from VR to AR) promises *spatial computing* [52] as a new way to perform knowledge tasks. Using high-resolution displays and precise head and body tracking supports the creation of virtual workspaces with renewed flexibility of when, how, and where to work. Rather than being confined to a fixed office space, HMD users can work anywhere, even on the go [21]. This option does not make traditional offices irrelevant. Modern work forms such as home offices, co-working spaces, and open floor plans afford more personalization and flexibility [24].

Assuming that XR will be part of this transformation, we must understand how people want to work in dynamic and transitional spaces. One of the emerging challenges concerns working with document collections and layouts, a common foundation of knowledge work tasks [27, 45, 46]. Previous research has shown that using the physical environment to lay out digital content may help organize information, build a mental model of digital structures, and remember where elements are located [32, 33, 53]. However, settling on a particular spatial layout is not trivially compatible with dynamic spaces and mobile use. Previous work has explored how spatial layouts could be automatically adapted to changes in physical location [29, 39]. However, the effects of frequent changes while working on a realistic knowledge task in XR have not yet been well studied.

In this paper, we explore how document-centered tasks are affected by changes in spatial context. We focus on how users arrange or rearrange documents, what types of layouts they prefer, and how this behavior aligns with location changes. To investigate this question, we conducted an exploratory user study with 17 participants. We asked them to work on a complex task with a set of documents in a video-see-through XR (VST-XR) setting. Participants could freely organize documents in layouts, combining world-referenced, body-referenced (i.e., stationary to the user), and surface-referenced layout types. Every five minutes, we asked them to change location, artificially interrupting their tasks and forcing them to adapt to a new physical environment (see Figure 1). We identified distinct layout patterns, several strategies for organizing documents over time, and requirements for document-centric knowledge work in VST-XR. With these insights, our work complements prior research on layout strategies in dynamic work spaces and informs the design of future extended reality knowledge work systems. Furthermore, we also contribute our experimental task design and the associated set of documents, available in the supplemental material, which support researchers conducting future studies in this area.

2 Related Work

Our work is related to existing research investigating how XR can be used for productivity tasks in general and how people lay out digital documents in XR, set up workspaces, and use the environmental context.

2.1 XR for Productivity

In recent years, XR has gained increasing popularity as a tool for productivity. Recent commercial devices have been advertised explicitly for productivity, such as the Apple Vision Pro[1], Lenovo Think Reality glasses[2], or Sightful's Spacetop[3]. In addition, research has demonstrated advantages of using XR as a productivity tool, such as reducing visual distraction by placing virtual boundaries [30, 48], or increasing privacy when, for example, entering passwords [49]. XR has also been shown to directly support productivity tasks such as working with spreadsheets [20] and presentation slides [2], note-taking [47] and sense-making tasks [33], or collaborative tasks such as code editing [12] and data analysis [6]. Specifically, the increased display space has been studied as a major advantage for productivity [33, 46]. More recent work has also shown the feasibility of using XR as a productivity tool in ecologically valid use cases [34, 37], even on a regular basis [3, 8, 22]. However, in these prior studies, participants were generally exposed to a static experience. Although the surroundings changed slightly, such as in public settings [1, 45], the participants remained where they initially set up their virtual content. Therefore, we believe there is a research gap regarding work sessions interrupted by location changes, which are likely to occur frequently in new forms of work.

2.2 Spatial Layouts

Prior work has examined various strategies for placing content in a 3D space. Ens et al. [14] describe different dimensions for designing 2D workspaces in 3D. Among others, the dimensions include proximity, input mode, discretization, and perspective. The perspective can be user-referenced (egocentric), meaning that content moves along with the user, or world-referenced (exocentric), meaning that content is anchored in a world-based reference frame.

Luo et al. [42] identified three spatial layout patterns (panoramic strip, semi-cylindrical, furniture-based) in a study in which participants should collaboratively categorize images. Furthermore, Luo et al. [41] found that participants preferred to anchor documents at specific locations, including physical objects and the participants' bodies. Cho et al. [9] conducted an elicitation study in which participants created layouts in four different environments. However, participants did not actively work on a realistic task and thus just tried to imagine a good layout. Lucero and Vetek [40] conducted a study in which participants walked on campus and a busy street wearing AR glasses that provided them with real notifications. Davidson et al. [11] investigated how knowledge work in immersive VR spaces evolves over several sessions. None of these works considers the change of location. In contrast, our work focuses on frequent location changes.

2.3 Adaptive Layouts

Automatic interface adaptations may make relocations easier for XR users. Previous work has looked at how pre-defined states can be used to place virtual content accordingly [19]. It is also a common approach to consider the user's movement and the affordances of the physical environment when optimizing interface placement [28, 29]. Cheng et al. [7] propose a system that considers the semantics between virtual content and possible placement locations.

Evangelista et al. [17] presented a toolkit for combining different adaptation objectives. Based on this toolkit, and by employing a vision-and-language model, Li et al. [31] present a system that considers environmental and social factors when placing virtual interfaces in real-world surroundings.

Other researchers are taking user behavior over a certain time frame into account to optimize the positioning of virtual displays [18, 36] or use reinforcement learning to automatically adapt content placement [35].

Previous work has also investigated different adaptation strategies. For example, Lu and Xu [39] found that a semi-automated approach to transitioning the UI to a different location was preferred over a manual and fully automated approach. Similarly, Lu et al. [38] compared systems that open all, some, or none of the previous apps after switching locations, and found no clear choice that works for all scenarios. Jannat et al. [26] explored whether a manual, semi-automatic, or automatic approach was preferred to transition interfaces between mid-air, arm, and table, and found that semi-automatic was preferred. Ens et al. [15] propose a transition from a body-centric layout to a world-referenced layout after

[1]https://www.apple.com/de/apple-vision-pro/, last accessed August 27th, 2025
[2]https://smartsupport.lenovo.com/us/en/products/smart/arvr/thinkreality-a3/downloads/ds549422, last accessed August 27th, 2025
[3]https://www.sightful.com/product, last accessed August 27th, 2025

transitioning to a new location. Niyazov et al. [43] used natural hand gestures to allow users to guide layout optimization, such as defining attractive edges or exclusion surfaces.

However, this work does not focus on making a better adaptive interface. Instead, we use a prototype that provides simple adaptation to new locations (body-centric moves with the user, world- and surface-centric reappears when the surface is selected). We primarily examine how participants react to location changes in a real-world task. Specifically, we are looking at whether they rearrange content, if their arrangements depend on the physical surroundings, and how they imagine the system should adapt.

3 User Study

Our study investigates how people interact and work with documents in VST-XR when they are forced to change their location frequently. Specifically, we are interested in which types of layouts participants prefer and how they adapt to the location changes. To investigate this in detail, we conducted an exploratory user study. The study was approved by the Committee for Responsibility in Research of the University of Stuttgart and took place in several locations within our institute building.

3.1 Goals and Overall Study Design

Our study design was informed by a set of guiding research questions:

RQ1: What layout types are used to arrange documents? Prior works have shown that users may prefer document arrangement in circular or spherical layouts around them [10], but also that surfaces and furniture are used to group documents by content [42]. Consequently, we are interested in the layout strategies for our setup. We support *world-referenced*, *body-referenced*, and *surface-referenced* layouts of documents in our study.

RQ2: How do location changes and interruptions influence workflows? Dynamic, mobile work may lead to an increased frequency of location changes. One of the main research questions guiding our study is how these transitions affect the work. Thus, the procedure includes frequent location changes every five minutes, forcing participants to adapt to various new environments.

RQ3: How do layout strategies change over time? Unlike most previous work, we are specifically interested in how layout strategies develop over time, particularly when users have changed locations and might realize that their initial strategy is not optimal. We designed our study so that locations would be revisited, allowing participants to rethink their strategies continuously over time and in regard to specific locations.

To answer these questions, our study design has individual users working with sets of digital documents in VST-XR. Their continuous work session is interrupted by forced location changes that do not or only incidentally coincide with finished sub-tasks. The participants worked in three different locations: a meeting room, a seating area with a coffee table, and a standing table (see Figure 2). These locations were chosen as realistic representations of the workplaces people might use in future XR productivity tasks. Through observations, logging of interactions, and post-study interviews, we gain insights that further our understanding of knowledge work in novel XR workspaces.

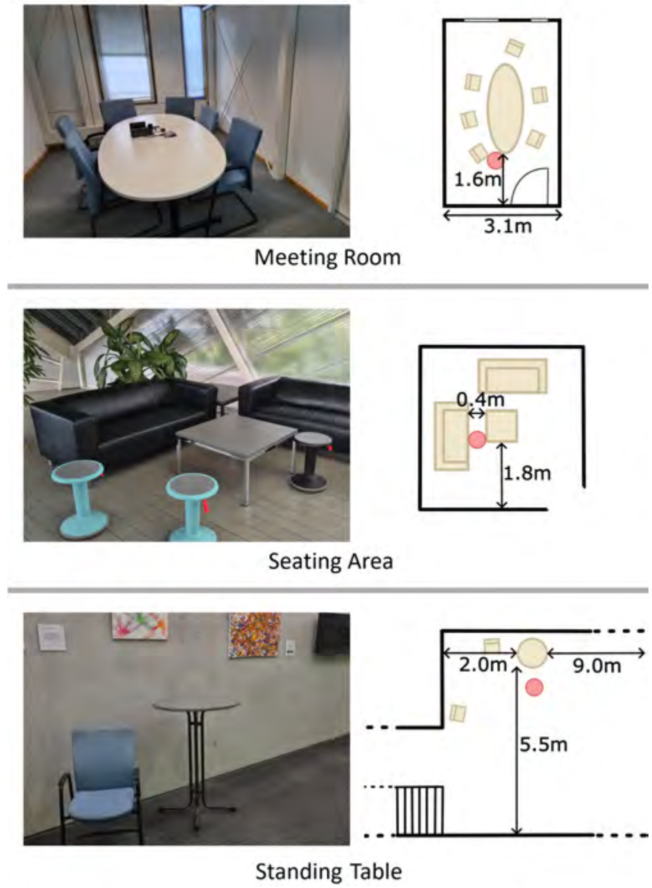

Figure 2: The three locations that we used in our study, including sketches of the floor plans. The red dots indicate typical positions of the participants during the study.

3.2 Study Task

To get a meaningful idea of how participants use the systems, we decided to employ a task that provides a high degree of ecological validity. Similar studies regarding the public use of XR devices [27, 45] had participants organize conference and lecture schedules, considering floor plans and schedules. We followed this style of task, but increased the number of documents to 11 (compared to six documents in the work of Pavanatto et al. [45]). In our task, participants play the role of event organizers and receive requests from people, such as visitors, that they need to answer by consulting at least two of the documents. These documents include a map of the city that hosts the event, the bus schedules for two bus lines, information on two different hotels, event schedules for four locations (museum, school, city hall, church), and detailed information for two places that host markets (see Figure 3). Both the eleven documents and the requests were provided as websites that could be displayed in our application. Similar to Kaeder et al. [27], we provided participants with only two responses to choose from (mostly yes/no) to avoid the use of a keyboard and extensive text entry.

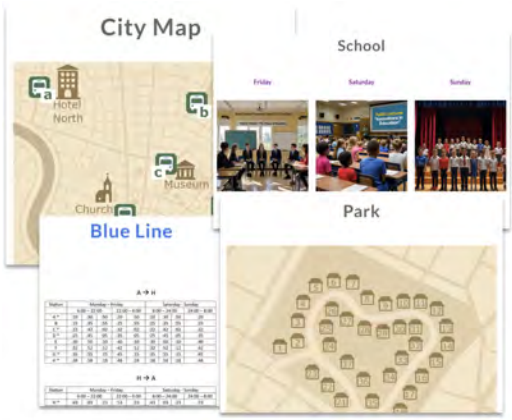

Figure 3: A selection of the documents presented to the participants during the study, such as a city map, the schedule of a bus line, a schedule for events happening in the school, and information about a market taking place in the park.

3.3 Apparatus

For our study, we used the Meta Quest 3. It offers a video-see-through mode with 18 PPD (pixels per degree) and a display resolution of 2064×2208 per eye. We implemented our study prototype in Unity 6[4], using the Meta XR All-in-One SDK[5]. All interaction in the prototype is done through direct manipulation (poke, grab) using the built-in hand tracking. We implemented interactive situated content containers based on the TLabWebView[6] Unity plugin, giving us the flexibility to show arbitrary web content. Users can use touch interaction to scroll through the documents or tap buttons or links. The documents have a size of 18.5 cm × 18.5 cm and are initially placed one behind the other, limiting their visibility (see Figure 1, left). They can be freely moved around by grabbing and placing them in space. Any document placed like this will remain in this position, allowing for *world-referenced layouts*. In addition, we support *surface-referenced layouts*: Similarly to *TriPad* [13], an anchor can be placed by using a three-finger gesture on a surface, which defines a plane on which two layouts can be filled with documents. Finally, to support *body-referenced layouts*, we placed eight layouts in a ring around the user. They are attached to the camera rig of the scene and thus follow the user. However, their rotation is world-locked, allowing the user to turn around to access content behind them.

For each type of layout, additional documents can be added by dragging them onto the existing layout. Drag-and-drop also allows for reordering documents in a layout. To support transitions between locations, the surface anchors are used: When a user places a new anchor, all world-referenced documents are moved from the old anchor to the same position relative to the new one; all surface-referenced documents are transferred to the new surface layouts, and all body-referenced layouts are rotated towards the new anchor. This behavior allows users to continue their work with minimal interruption. We believe that this represents a realistic trade-off between fully automatic layout adaptation and manual content placement for changing spatial contexts.

[4]Unity 3D Engine: https://unity.com
[5]Meta XR SDK: https://developers.meta.com/horizon/documentation/unity/unity-sdks-overview
[6]TLabWebView: https://github.com/TLabAltoh/TLabWebView

Table 1: Order of the three different locations for each participant. ST: Standing Table, SA: Seating Area, MR: Meeting Room

	1st location	2nd location	3rd location	4th location	5th location
P6, P11, P17	ST	SA	MR	SA	ST
P1, P7, P12	ST	MR	SA	MR	ST
P2, P8, P13	MR	ST	SA	ST	MR
P3, P14	MR	SA	ST	SA	MR
P4, P9, P15	SA	MR	ST	MR	SA
P5, P10, P16	SA	ST	MR	ST	SA

3.4 Procedure

At the beginning of each study session, participants completed a consent form and a demographic questionnaire. Then, the system and the task were explained to them, supported by video snippets that showed the interaction possibilities. Afterwards, participants could try out the system in a training task using different documents. Then the actual task started at the first location. Participants had five minutes to work on the task before being asked to move to the next location, where they could continue where they left off. Overall, participants completed five five-minute sessions in three different locations (meeting room, seating area, standing table, as specified in section 3.1). Participants moved from the first to the second to the third location, back to the second, and then to the first (see Figure 1). This order means that participants experienced new locations, old locations they had visited recently, and old locations they had visited a little longer ago. The order of the locations was counterbalanced between the participants to alleviate any effects the different locations might have. These orders are displayed in Table 1. During transitions, each a short 1-2 minute walk, less-experienced participants removed the HMD for safety reasons. After arriving in the new location, participants placed the spatial anchor, resetting the document locations as described above. After all five sessions were completed, participants answered questionnaires, and we conducted a short interview to get more detailed insight into the participants' experiences. The study lasted about 60 minutes, and participants were compensated with 15 Euros.

3.5 Measures

We recorded videos of the study with the Quest's screen-capturing function and tracking data for the head, hands, and all documents, as well as interaction events. Furthermore, we collected subjective data through several questionnaires: We used the System Usability Scale [4] to verify the general usability of our prototype. We also used the NASA TLX questionnaire [23] to gauge the physical and mental demands of the task. In addition, we asked participants a range of questions about their experiences during the study, such as their content placement strategies and their impression of the system. Finally, we conducted a short interview asking for insights about the location changes, the task, and content placement. All additional questions and interview questions can be found in the supplementary material.

3.6 Participants

Participants were recruited through university mailing lists. 17 participants (4 female, 13 male) completed the study successfully. The participants' average age was 25.4 years ($sd = 2.0$). All of them were students or employees of the University of Stuttgart with a technical background. On a scale from 1 (None) to 5 (Extensive), participants' average experience with AR was 2.7 ($sd = 0.8$). Three participants reported using XR for productivity tasks before, but only in the context of prior studies they participated in. All participants had normal or corrected-to-normal vision, and five wore glasses during the study.

4 Results

After viewing several video recordings of the study, two of the authors decided on a list of characteristics that can be observed and coded in the videos. These include observed types of layouts, the amount of movement by the participants, their posture, the number of containers in layouts, etc.[7] They then compared results for one video to ensure good inter-coder reliability, before each coder coded half of the remaining videos. Event logs were analyzed using spreadsheet software. We used the MIRIA toolkit [5] to analyze the tracking logs and visualize how the layouts evolved.

In the following paragraphs, we present the combined observations from the questionnaires, video recordings, log data, and participant responses during interviews.

Usability. On average, the participants rated the usability of the system (SUS) [4] with 66.8 ($sd = 16.7$), which can be considered average. Post-study interviews confirmed that touch interaction within the documents was cumbersome, and several participants explicitly mentioned that they had problems scrolling and selecting the answers (P9, P19, P13, P17).

Task Load. The answers to the NASA TLX questions [23] indicated a relatively high perceived task load. Participants rated the mental demand as 79.7 of 100 ($sd = 15.9$), the physical demand 51.8 ($sd = 24.5$), the temporal demand 47.4 ($sd = 26.4$), effort 73.2 ($sd = 16.3$), and frustration 45.3 ($sd = 33.7$). The average performance rating was 47.9 ($sd = 28.5$) (100 being the best). These values indicate that our task was sufficiently demanding to force participants to engage with the document content, and that the location changes likely induced some time pressure.

Effects of Spatial Transitions. In the questionnaire, participants indicated whether they focused more on neatly positioning the windows (1) or more on answering the questions quickly (5). The results are depicted in Figure 4 along with participants' agreement to a range of statements about their experience with the system (1: totally disagree, 5: totally agree). The answers indicate that most participants found it easy to arrange the content, and it also helped them answer the questions.

[7] Please also refer to the supplemental material.

Figure 4: Participants' agreement (1-5) to a range of statements regarding their experience during the study.

From our observations in the videos, we conclude that 15 participants only made minor adjustments to the layouts after transitioning to a new position. These were mainly caused by physical constraints, such as the world-referenced documents being inside a physical object or too high or low (because they are positioned relative to the table in the previous location). However, the results were mixed in the questionnaire when asked whether the content placement was as expected after moving to a new location. This observation indicates that "ideal" content placement could be very subjective. Five participants (P1, P4, P8, P10, P17) specified in the interview that only minor adjustments were needed, and three participants (P12 - P14) had to make some more adjustments because the documents were in different places.

Consequently, we also found mixed results when asking participants whether the transition interfered with their task progression. In the interviews, four participants (P1, P8, P11, P16) had a hard time remembering their progress on the current question, while others (P2, P3, P5, P9, P10, P13 - P15, P17) said it did not influence them at all. P4 and P6 had no problem continuing their train of thought, but had to adapt to the new position of the documents.

Suggested Improvements. When asked about their ideal content placement after transitioning, four participants (P1, P4, P5, P9) mentioned that the relative positions should be kept. P5 specifies that this includes users and documents, but the physical room does not matter. P6 mentions that, at the same location, the positioning should be exactly the same. P16 would like to place everything around him, i.e., body referenced. Others (P1, P3, P8, P10)

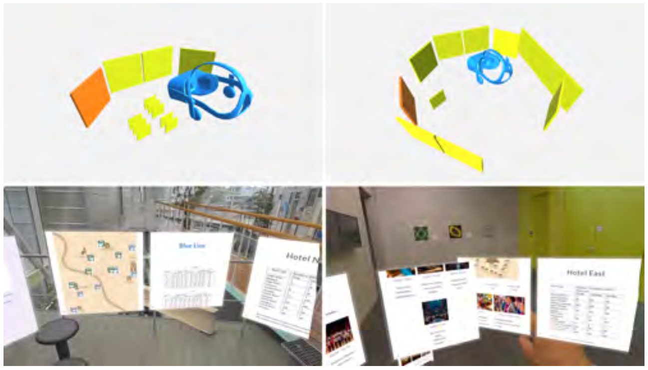

Figure 5: Example of half-circle layout, being extended into a full circle over time (P9). The first row shows data reconstructed from the logs, the second row shows frames of the recorded video. Times are not synchronized.

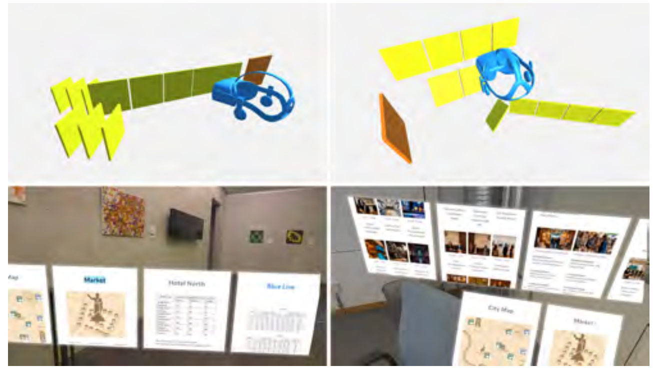

Figure 6: Example of a linear layout, being extended into a multi-row layout over time (P6). The first row shows data reconstructed from the logs, the second row shows frames of the recorded video. Times are not synchronized.

specifically criticized the placement of world-referenced content as sometimes inconsistent.

Features that several participants found missing included 2D layout templates (P1, P3, P15), a resize or zoom functionality (P5, P9, P13 - P15), and moving all objects in a layout at once (P5, P8, P14). Several participants also had problems with mid-air touch interaction (P2, P8, P11), and P8 suggested improving this through audio feedback. Other suggestions included vertical layouts (P5), customizing the distance of body reference layouts (P1), and the possibility to hide everything (P1).

Frames of Reference. According to the log data, body- and world-referenced layouts were the most common, accounting for 48% and 45% respectively, of the average time that documents spent in any layout. Surface-referenced layouts were rarer, with approximately 7%. However, personal preferences are very diverse. For example, six participants did not use surface layouts at all. P5 used almost exclusively (94%) world-referenced layouts, P15 used mostly (around 82%) body layouts, and P17 used the three layout types more balanced (world: 34%, body: 33%, surface: 33%). This finding is in contrast to how participants ranked the layouts in the questionnaire according to their preferences. Eight participants ranked body-referenced as highest; six, surface-referenced, and three, world-referenced.

Posture and Movement. In total, participants completed 85 five-minute sessions. In 67 sessions (79%), participants were only standing, even though seating was available in both the meeting room and the seating area. In 46 sessions (54%), participants only performed small movements (no obvious steps or unnatural rotations). In 31 sessions (36%), they performed rather large rotations, which most likely involved at least the upper body, and, in eight sessions (9%), participants walked back and forth to inspect documents. Interestingly, several participants (P3, P5) reported in the interview that they moved their body rather than repositioning the content.

Layout Arrangement. Generally, we observed five main layout strategies: a half-circle (Figure 5, left), a full-circle (Figure 5, right), linear layouts (Figure 6, left), multi-row linear layouts (Figure 6, right), which are sometimes combined with half-circle layouts, and loose or irregular layouts (Figure 7). Circular layouts roughly correspond to cylindrical layouts [10] or panoramic strip patterns [42] observed in prior works. Most of the participants (12) started with a half-circle layout, two started with a linear layout, two with a loose layout, and one participant with a full-circle layout. However, in the last session, only four participants still used the half-circle layout. Of the ones who started with a half-circle, five transitioned their layout into a full-circle, two shifted to a loose layout, and one to a linear one. The two participants starting in the linear layout both transitioned to a multi-row or compound layout. One of the two who started with a loose layout kept it; the other ended up with a linear layout. The person with the full-circle layout stuck with it until the end.

Figure 7: Example of a loose or irregular layout over time (P8). The first row shows data reconstructed from the logs, the second row shows frames of the recorded video. Times are not synchronized.

The number of documents placed in a single layout varied considerably between participants, but also between sessions, with a minimum number of 1, a maximum of 8, and an average of 3.3 ($sd = 1.3$).

Grouping. Participants reported different strategies for content placement. Some participants reported that they arranged content similarly in all locations, while others did not (see Figure 4). During interviews, participants generally described two main strategies: sorting by *content* type (P1, P3 - P5, P7, P9, P13, P14, P16) and sorting by currently *needed* documents (P1 - P3, P6, P9, P12, P16, P17). As indicated in the questionnaire, these strategies could change over time. For example, P1 changed from *content* to *needed* strategy, while P9 and P16 changed from *needed* to the *content* strategy, and P3 employed a mixture of both.

In the videos, we observed that two participants clearly grouped by *content*, 10 participants seemed to group content occasionally, while five participants apparently did not group the content in any way. However, there is significant uncertainty in these observations, as we cannot infer from the videos whether the participants had an underlying plan, specifically for the *needed* strategy.

It was apparent that 10 participants kept the questionnaire (orange document in Figures 5, 6, 7, 8, and 9) in a separate layout throughout the study, probably to keep it always in sight. Two kept it separate for a while, and five did not handle it differently from the other documents.

Physical Environment. From the questionnaire, we also see mixed results on whether participants considered the physical environment when placing content (see Figure 4). However, during the interviews, most of the participants stated that they did not consider the properties of the physical environment in any way (P4 - P10, P13 - P15, P17), which may be partially explained by the lack of semantic coupling between the documents and the environments. Also, the observations mentioned above indicate that layout strategies were more task-driven than environment-driven. P1, P11, and P12 only mentioned considering their own ergonomics, and P1, P2, and P16 liked the table. Over time, P16 forgot about the table and placed the documents everywhere. P3 first sat down in the seating area but later stood up again, when the space between the sofa and the table became too restrictive.

5 Discussion

In this section, we summarize the results of our study in relation to our guiding research questions, derive an initial set of design guidelines, and discuss the limitations of our work.

5.1 Research Questions

RQ1: Overall, we saw a large variety of distinct layouts. The layouts used by our participants are mostly consistent with what has been reported in prior work [10, 42], e.g., regarding circular and half-circular layouts (Figure 5), which have been investigated, among others, by Ens et al. [16]. However, we also saw linear layouts (Figure 6), which may have been supported by our automatic snapping of documents. In contrast, we observed fewer surface-referenced or environmental layouts than expected, although some loose or unstructured layouts could be interpreted as this (Figure 7). While it is clear that the chosen layouts were influenced by the prototype design (e.g., snapping of documents to arbitrary surfaces [44] might motivate people to use the environment more), this observation could also indicate that for mentally demanding tasks under time pressure, users care little about surfaces as document anchors, even in augmented reality. This behavior is corroborated by questionnaires and interviews, which show that many participants ignored the environment, and contrasts with some earlier findings, such as those by Luo et al. [42] who found that participants used physical surfaces to organize content.

RQ2: As expected, we observed some individual differences in how location changes interfered with task completion. Overall, a majority of participants neither reported that they had to start over with the current question nor that location changes had strong interference with the task in general. However, some mentioned the need to adapt the layout to the new environment, such as when table heights did not fit their preferred position or if documents were positioned too close to physical objects. These findings indicate that, for the short interruptions in our study, no "mental reset" was necessary. Still, it underlines the importance of rapid setups or adaptive layouts [31, 35, 39] to continue tasks promptly after location changes.

RQ3: Over time, many participants "completed" layouts that they started earlier. For example, Figure 8 shows how P2 started with a half-circle layout and some smaller linear layouts, but gradually changed to a full-circle layout and later partially broke this layout again. In contrast, Figure 9 shows P17 building and refining a multi-row linear layout over time. This observation suggests that people iterate or optimize their chosen layouts over time. This is in line with findings for longer, multiple-session work in VR, investigated by Davidson et al. [11]. Interestingly, even in our VST-XR setting, such changes were less motivated by the environment, but more by the relevance of a document at a given time or the general level of clutter.

5.2 Recommendations

Support both world- and body-referenced arrangements. As body- and world-referenced layouts are used the most, we recommend explicitly supporting these layout container styles. While this may seem obvious, many prior works investigating layout strategies have not considered body layouts [9, 10, 42]. These layouts follow the user, providing convenient, location-independent access. Still, they do not scale as well with many items and may leave documents at uncomfortable distances to the user for interaction, or lead to occlusions. We saw the use of body layouts even when moving between clusters of world-referenced documents at a single location, suggesting a need for them even if world-referenced layouts are precisely and reliably tracked. This use of body layouts may suggest quick-access usage, similar to virtual toolbelts [51], and future work in multi-user settings should investigate the interplay with the notion of personal territories, a concept also known from other technical setups, including tabletops [50] and AR-extended display walls [25].

Allow for arbitrary changes of layouts. The personal workspaces that the participants in our study constructed continued to evolve. In many cases, they changed to entirely different layout structures. Consequently, for work with medium or large amounts of documents, we recommend not forcing users into predefined or inflexible layouts. Furthermore, automatic layout adaptations should consider not only changes to physical environments, but also that users will

Figure 8: Layout progression of a circular layout over the five locations (P2).

Figure 9: Layout progression of a multi-row linear layout over the five locations (P17).

create and modify layouts over time, which prior research often does not take into account (e.g., [29, 31]). Still, as we saw with our linear layouts and snapping, users like assistive layout techniques, and in fact, some asked for more layout options. Balancing between the expressiveness of manual layouts and the efficacy of automation remains a worthwhile research direction.

Ensure user control in adaptive layouts. Unless virtual content is strongly semantically coupled to physical objects, typical workspaces of knowledge workers need to "follow" users to new locations. That means they have to adapt to the available physical restrictions (for example, ensuring that the user can see and reach all layouts), while ensuring the continuity of the previously created layout. This highlights the need for powerful adaptive layout techniques. However, in our study, we saw that users want to use their workspaces from whatever position they find suitable, not limiting themselves to a predetermined main viewing direction. Therefore, we emphasize that adaptive layout techniques must provide easy ways for the user to change the orientation of documents and arrangements. Even though future automatic content placement strategies may become better at inferring a user's intentions, the fine nuances of highly task and context-dependent user preferences will remain elusive. For example, strategies such as transitioning from body-referenced to world- or surface-referenced automatic layouts, as proposed by Ens et al. [16], do not align with how our participants actively worked with multiple documents.

5.3 Limitations and Future Work

Our study has some limitations that can be addressed in future work. First, the number of participants was rather low, and, unfortunately, our sample was not fully gender balanced. In addition, all participants were of similar age. Also, as suggested by the participants (see Section 4), in future work, we will further improve our prototype. For practical reasons, all study locations were within our institute building to minimize travel time, and task sessions were limited to five minutes to keep the study duration at a reasonable length. In future work, we will explore such a system in more diverse locations and extend the study to longer sessions, possibly expanded over several days. Additionally, our findings are potentially influenced by the task characteristics, so in future prolonged studies, we will test various task types. Finally, the data from our log files are only approximations, as, due to technical restrictions, transitions between locations are not fully removed. Furthermore, there might be a slight bias, because documents always start in a default location in world-referenced layouts.

6 Conclusion

In this paper, we present insights from an exploratory study in which participants worked on a document-centered organization and planning task in VST-XR while repeatedly changing their locations. We observed that, similar to prior work in static scenarios, participants used a variety of layout strategies for placing the documents, such as circular or linear layouts in a world- and body-referenced manner. Yet, they rarely placed documents in surface-based layouts. In contrast to prior studies the questionnaires and interviews suggest that the physical environment played no major role in their placement strategies and that the layout modifications were more likely caused by the task requirements than the physical environment. Although the task was rated somewhat demanding, most participants did not find the transition very disruptive to the task and only required minor adjustments to the layout. Overall, layout strategies and preferences seem highly subjective, and systems should therefore provide different options for placing content, allow further changes to adapt to evolving task requirements, and provide the features to restore a previous layout quickly.

Acknowledgments

This work was supported by the Alexander von Humboldt Foundation, funded by the German Federal Ministry of Research, Technology and Space, the German Research Foundation DFG (495135767), and the Austrian Research Funds FWF (I5912).

References

[1] Laura Bajorunaite, Julie R Williamson, and Stephen Brewster. 2024. VR Headsets In-The-Wild: Qualitative Insights on Safety, Awareness, and Social Challenges from Real Train Journeys. In *Proceedings of the ACM MUM*. Association for Computing Machinery, New York, NY, USA, 82–94. doi:10.1145/3701571.3701576

[2] Verena Biener, Travis Gesslein, Daniel Schneider, Felix Kawala, Alexander Otte, Per Ola Kristensson, Michel Pahud, Eyal Ofek, Cuauhtli Campos, Matjaž Kljun, Klen Čopič Pucihar, and Jens Grubert. 2022. PoVRPoint: Authoring Presentations in Mobile Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 28, 5 (2022), 2069–2079. doi:10.1109/TVCG.2022.3150474

[3] Verena Biener, Florian Jack Winston, Dieter Schmalstieg, and Alexander Plopski. 2025. Long-Term Experiences from Working with Extended Reality in the Wild. In *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 667–676. doi:10.1109/ISMAR67309.2025.00076
[4] John Brooke. 1996. SUS—A quick and dirty usability scale. *Usability Evaluation in Industry* 189, 194 (1996), 4–7. doi:10.1201/9781498710411-35
[5] 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 *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 470, 15 pages. doi:10.1145/3411764.3445651
[6] Simon Butscher, Sebastian Hubenschmid, Jens Müller, Johannes Fuchs, and Harald Reiterer. 2018. Clusters, Trends, and Outliers: How Immersive Technologies Can Facilitate the Collaborative Analysis of Multidimensional Data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3173664
[7] 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 *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '21)*. Association for Computing Machinery, New York, NY, USA, 282–297. doi:10.1145/3472749.3474750
[8] Yi Fei Cheng, Ari Carden, Hyunsung Cho, Catarina G. Fidalgo, Jonathan Wieland, and David Lindlbauer. 2025. Augmented Reality Productivity In-the-Wild: A Diary Study of Usage Patterns and Experiences of Working With AR Laptops in Real-World Settings. *IEEE Transactions on Visualization and Computer Graphics* 31, 10 (2025), 9195–9212. doi:10.1109/TVCG.2025.3592962
[9] Hyunsung Cho, Yukang Yan, Kashyap Todi, Mark Parent, Missie Smith, Tanya R. Jonker, Hrvoje Benko, and David Lindlbauer. 2024. MineXR: Mining Personalized Extended Reality Interfaces. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 609, 17 pages. doi:10.1145/3613904.3642394
[10] Kylie Davidson, Lee Lisle, Ibrahim A. Tahmid, Kirsten Whitley, Chris North, and Doug A. Bowman. 2023. Uncovering Best Practices in Immersive Space to Think. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 1094–1103. doi:10.1109/ISMAR59233.2023.00126
[11] Kylie Davidson, Lee Lisle, Ibrahim A. Tahmid, Kirsten Whitley, Chris North, and Doug A. Bowman. 2025. Investigating Professional Analyst Strategies in Immersive Space to Think. *IEEE Transactions on Visualization and Computer Graphics* 31, 9 (2025), 5364–5378. doi:10.1109/TVCG.2024.3444594
[12] James Dominic, Brock Tubre, Charles Ritter, Jada Houser, Colton Smith, and Paige Rodeghero. 2020. Remote pair programming in virtual reality. In *IEEE ICSME*. IEEE, 406–417. doi:10.1145/3387904.3389287
[13] Camille Dupré, Caroline Appert, Stéphanie Rey, Houssem Saidi, and Emmanuel Pietriga. 2024. TriPad: Touch Input in AR on Ordinary Surfaces with Hand Tracking Only. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 754, 18 pages. doi:10.1145/3613904.3642323
[14] Barrett Ens, Juan David Hincapié-Ramos, and Pourang Irani. 2014. Ethereal planes: a design framework for 2D information space in 3D mixed reality environments. In *Proceedings of the 2nd ACM Symposium on Spatial User Interaction* (Honolulu, Hawaii, USA) *(SUI '14)*. Association for Computing Machinery, New York, NY, USA, 2–12. doi:10.1145/2659766.2659769
[15] Barrett Ens, Eyal Ofek, Neil Bruce, and Pourang Irani. 2015. Spatial Constancy of Surface-Embedded Layouts across Multiple Environments. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction* (Los Angeles, California, USA) *(SUI '15)*. Association for Computing Machinery, New York, NY, USA, 65–68. doi:10.1145/2788940.2788954
[16] Barrett M. Ens, Rory Finnegan, and Pourang P. Irani. 2014. The personal cockpit: a spatial interface for effective task switching on head-worn displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 3171–3180. doi:10.1145/2556288.2557058
[17] João Marcelo Evangelista Belo, Mathias N. Lystbæk, Anna Maria Feit, Ken Pfeuffer, Peter Kán, Antti Oulasvirta, and Kaj Grønbæk. 2022. AUIT – the Adaptive User Interfaces Toolkit for Designing XR Applications. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) *(UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 48, 16 pages. doi:10.1145/3526113.3545651
[18] Andreas Fender, David Lindlbauer, Philipp Herholz, Marc Alexa, and Jörg Müller. 2017. HeatSpace: Automatic Placement of Displays by Empirical Analysis of User Behavior. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (Québec City, QC, Canada) *(UIST '17)*. Association for Computing Machinery, New York, NY, USA, 611–621. doi:10.1145/3126594.3126621
[19] Andreas Fender and Jörg Müller. 2019. SpaceState: Ad-Hoc Definition and Recognition of Hierarchical Room States for Smart Environments. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces* (Daejeon, Republic of Korea) *(ISS '19)*. Association for Computing Machinery, New York, NY, USA, 303–314. doi:10.1145/3343055.3359715
[20] Travis Gesslein, Verena Biener, Philipp Gagel, Daniel Schneider, Per Ola Kristensson, Eyal Ofek, Michel Pahud, and Jens Grubert. 2020. Pen-based Interaction with Spreadsheets in Mobile Virtual Reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 361–373. doi:10.1109/ISMAR50242.2020.00063
[21] Jens Grubert, Eyal Ofek, Michel Pahud, and Per Ola Kristensson. 2018. The Office of the Future: Virtual, Portable, and Global. *IEEE Computer Graphics and Applications* 38, 6 (2018), 125–133. doi:10.1109/MCG.2018.2875609
[22] Eugy Han, Mark R Miller, Cyan DeVeaux, Hanseul Jun, Kristine L Nowak, Jeffrey T Hancock, Nilam Ram, and Jeremy N Bailenson. 2023. People, places, and time: a large-scale, longitudinal study of transformed avatars and environmental context in group interaction in the metaverse. *Journal of Computer-Mediated Communication* 28, 2 (2023), zmac031. doi:10.1093/jcmc/zmac031
[23] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183. doi:10.1016/S0166-4115(08)62386-9
[24] Julia Hobsbawm. 2022. *The Nowhere Office: Reinventing Work and the Workplace of the Future*. Basic Books.
[25] Raphaël James, Anastasia Bezerianos, and Olivier Chapuis. 2023. Evaluating the Extension of Wall Displays with AR for Collaborative Work. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 99, 17 pages. doi:10.1145/3544548.3580752
[26] Marium-E Jannat, Prateek Dhaka, Keiko Katsuragawa, and Khalad Hasan. 2024. Exploring Augmented Reality User Interface Transitions Across Mid-Air, On-Body and Physical Surfaces. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 1177–1186. doi:10.1109/ISMAR62088.2024.00134
[27] Janne Kaeder, Maurizio Vergari, Verena Biener, Tanja Kojić, Jens Grubert, Sebastian Möller, and Jan-Niklas Voigt Antons. 2024. Working with Mixed Reality in Public: Effects of Virtual Display Layouts on Productivity, Feeling of Safety, and Social Acceptability. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 740–748. doi:10.1109/ISMAR62088.2024.00089
[28] Wallace Lages and Doug Bowman. 2019. Adjustable Adaptation for Spatial Augmented Reality Workspaces. In *Symposium on Spatial User Interaction* (New Orleans, LA, USA) *(SUI '19)*. Association for Computing Machinery, New York, NY, USA, Article 20, 2 pages. doi:10.1145/3357251.3358755
[29] Wallace S. Lages and Doug A. Bowman. 2019. Walking with adaptive augmented reality workspaces: design and usage patterns. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (Marina del Ray, California) *(IUI '19)*. Association for Computing Machinery, New York, NY, USA, 356–366. doi:10.1145/3301275.3302278
[30] Hyelip Lee, Seungwoo Je, Rachel Kim, Himanshu Verma, Hamed Alavi, and Andrea Bianchi. 2022. Partitioning open-plan workspaces via augmented reality. *Personal and Ubiquitous Computing* 26 (2022), 609–624. doi:10.1007/s00779-019-01306-0
[31] Zhipeng Li, Christoph Gebhardt, Yves Inglin, Nicolas Steck, Paul Streli, and Christian Holz. 2024. SituationAdapt: Contextual UI Optimization in Mixed Reality with Situation Awareness via LLM Reasoning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) *(UIST '24)*. Association for Computing Machinery, New York, NY, USA, Article 43, 13 pages. doi:10.1145/3654777.3676470
[32] Lee Lisle, Xiaoyu Chen, J.K. Edward Gitre, Chris North, and Doug A. Bowman. 2020. Evaluating the Benefits of the Immersive Space to Think. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. 331–337. doi:10.1109/VRW50115.2020.00073
[33] Lee Lisle, Kylie Davidson, Leonardo Pavanatto, Ibrahim A. Tahmid, Chris North, and Doug A. Bowman. 2023. Spaces to Think: A Comparison of Small, Large, and Immersive Displays for the Sensemaking Process. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 1084–1093. doi:10.1109/ISMAR59233.2023.00125
[34] Feiyu Lu and Doug A. Bowman. 2021. Evaluating the Potential of Glanceable AR Interfaces for Authentic Everyday Uses. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. 768–777. doi:10.1109/VR50410.2021.00104
[35] Feiyu Lu, Mengyu Chen, Hsiang Hsu, Pranav Deshpande, Cheng Yao Wang, and Blair MacIntyre. 2024. Adaptive 3D UI Placement in Mixed Reality Using Deep Reinforcement Learning. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 32, 7 pages. doi:10.1145/3613905.3651059
[36] Feiyu Lu, Mengyu Chen, Hsiang Hsu, Pranav Deshpande, Cheng Yao Wang, and Blair MacIntyre. 2024. Adaptive Content Placement in Mixed Reality Through

Empirical User Behavioral Patterns. In *2024 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 199–204. doi:10.1109/ISMAR-Adjunct64951.2024.00049

[37] Feiyu Lu, Leonardo Pavanatto, and Doug A. Bowman. 2023. In-the-Wild Experiences with an Interactive Glanceable AR System for Everyday Use. In *Proceedings of the 2023 ACM Symposium on Spatial User Interaction* (Sydney, NSW, Australia) *(SUI '23)*. Association for Computing Machinery, New York, NY, USA, Article 11, 9 pages. doi:10.1145/3607822.3614515

[38] Feiyu Lu, Leonardo Pavanatto, Shakiba Davari, Lei Zhang, Lee Lisle, and Doug A. Bowman. 2025. "Where Did My Apps Go?" Supporting Scalable and Transition-Aware Access to Everyday Applications in Head-Worn Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 31, 9 (2025), 6112–6129. doi:10.1109/TVCG.2024.3493115

[39] Feiyu Lu and Yan Xu. 2022. Exploring Spatial UI Transition Mechanisms with Head-Worn Augmented Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 550, 16 pages. doi:10.1145/3491102.3517723

[40] Andrés Lucero and Akos Vetek. 2014. NotifEye: using interactive glasses to deal with notifications while walking in public. In *Proceedings of the 11th Conference on Advances in Computer Entertainment Technology* (Funchal, Portugal) *(ACE '14)*. Association for Computing Machinery, New York, NY, USA, Article 17, 10 pages. doi:10.1145/2663806.2663824

[41] Weizhou Luo, Mats Ole Ellenberg, Marc Satkowski, and Raimund Dachselt. 2025. Documents in Your Hands: Exploring Interaction Techniques for Spatial Arrangement of Augmented Reality Documents. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1218, 22 pages. https://doi.org/10.1145/3706598.3713518

[42] Weizhou Luo, Anke Lehmann, Yushan Yang, and Raimund Dachselt. 2021. Investigating Document Layout and Placement Strategies for Collaborative Sensemaking in Augmented Reality. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 456, 7 pages. doi:10.1145/3411763.3451588

[43] Aziz Niyazov, Barrett Ens, Kadek Ananta Satriadi, Nicolas Mellado, Loic Barthe, Tim Dwyer, and Marcos Serrano. 2023. User-Driven Constraints for Layout Optimisation in Augmented Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 35, 16 pages. doi:10.1145/3544548.3580873

[44] Benjamin Nuernberger, Eyal Ofek, Hrvoje Benko, and Andrew D. Wilson. 2016. SnapToReality: Aligning Augmented Reality to the Real World. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 1233–1244. doi:10.1145/2858036.2858250

[45] Leonardo Pavanatto, Verena Biener, Jennifer Chandran, Snehanjali Kalamkar, Feiyu Lu, John J. Dudley, Jinghui Hu, G. Nikki Ramirez, Per Ola Kristensson, Alexander Giovannelli, Luke Schlueter, Jörg Müller, Jens Grubert, and Doug A. Bowman. 2025. Working in Extended Reality in the Wild: Worker and Bystander Experiences of XR Virtual Displays in Public Real-World Settings. *IEEE Transactions on Visualization and Computer Graphics* 31, 10 (2025), 9104–9123. doi:10.1109/TVCG.2025.3589283

[46] Leonardo Pavanatto, Feiyu Lu, Chris North, and Doug A. Bowman. 2025. Multiple Monitors or Single Canvas? Evaluating Window Management and Layout Strategies on Virtual Displays. *IEEE Transactions on Visualization and Computer Graphics* 31, 3 (2025), 1713–1730. doi:10.1109/TVCG.2024.3368930

[47] Leping Qiu, Erin Seongyoon Kim, Sangho Suh, Ludwig Sidenmark, and Tovi Grossman. 2025. MaRginalia: Enabling In-person Lecture Capturing and Note-taking Through Mixed Reality. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 141, 15 pages. doi:10.1145/3706598.3714065

[48] Anastasia Ruvimova, Junhyeok Kim, Thomas Fritz, Mark Hancock, and David C. Shepherd. 2020. "Transport Me Away": Fostering Flow in Open Offices through Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3313831.3376724

[49] Daniel Schneider, Alexander Otte, Travis Gesslein, Philipp Gagel, Bastian Kuth, Mohamad Shahm Damlakhi, Oliver Dietz, Eyal Ofek, Michel Pahud, Per Ola Kristensson, Jörg Müller, and Jens Grubert. 2019. ReconViguRation: Reconfiguring Physical Keyboards in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 25, 11 (2019), 3190–3201. doi:10.1109/TVCG.2019.2932239

[50] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori Inkpen. 2004. Territoriality in collaborative tabletop workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA) *(CSCW '04)*. Association for Computing Machinery, New York, NY, USA, 294–303. doi:10.1145/1031607.1031655

[51] Henry A. Sowizral. 1994. Tools and metatools for interacting with a virtual environment. In *Stereoscopic Displays and Virtual Reality Systems*, Scott S. Fisher, John O. Merritt, and Mark T. Bolas (Eds.), Vol. 2177. International Society for Optics and Photonics, SPIE, 409 – 416. doi:10.1117/12.173904

[52] Jiangnan Xu, Konstantinos Papangelis, Garreth W. Tigwell, Nicolas Lalone, Pengyuan Zhou, Michael Saker, Alan Chamberlain, John Dunham, Sanzida Mojib Luna, and David Schwartz. 2024. Spatial Computing: Defining the Vision for the Future. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 583, 4 pages. doi:10.1145/3613905.3643978

[53] Fumeng Yang, Jing Qian, Johannes Novotny, David Badre, Cullen D. Jackson, and David H. Laidlaw. 2021. A Virtual Reality Memory Palace Variant Aids Knowledge Retrieval from Scholarly Articles. *IEEE Transactions on Visualization and Computer Graphics* 27, 12 (2021), 4359–4373. doi:10.1109/TVCG.2020.3009003