Walking with Adaptive Augmented Reality Workspaces: Design and Usage Patterns

Wallace S. Lages
Center for Human-Computer Interaction
Blacksburg, USA
wlages@vt.edu

Doug A. Bowman
Center for Human-Computer Interaction
Blacksburg, USA
dbowman@vt.edu

Figure 1: Adaptive Augmented Reality Workspace in use while walking and approaching walls.

ABSTRACT

Mobile augmented reality may eventually replace our smartphones as the primary way of accessing information on the go. However, current interfaces provide little support to walking and to the variety of actions we perform in the real world. To achieve its full potential, augmented reality interfaces must support the fluid way we move and interact in the physical world. We explored how different adaptation strategies can contribute towards this goal. We evaluated design alternatives through contextual studies and identified the key interaction patterns that interfaces for walking should support. We also identified desirable properties of adaptation-based interface techniques, which can be used to guide the design of the next-generation walking-centered augmented reality workspaces.

CCS CONCEPTS

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

KEYWORDS

augmented reality, adaptive interfaces, wearable, user interface

ACM Reference Format:
Wallace S. Lages and Doug A. Bowman. 2019. Walking with Adaptive Augmented Reality Workspaces: Design and Usage Patterns. In *24th International Conference on Intelligent User Interfaces (IUI '19), March 17–20, 2019, Marina del Rey, CA, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3301275.3302278

1 INTRODUCTION

Wearable augmented reality (AR) devices give us the unprecedented ability to visualize digital information anywhere. However, the way we currently manage information in AR does not allow us to to take full advantage of this fact. Imagine you are working on a project using a virtual AR workspace. You place a diagram on the wall, notes and images on the table, and your e-mail client at your right-hand side. Eventually, you leave your office to get a coffee and someone asks you a question about your project. What do you do? The information you need is on your office wall. You could, of course, go back and pick it up. Or you could open it again from a system menu. However, this would quickly become tedious and inefficient to do every time you walk.

As the previous scenario illustrates, a fixed layout in the world does not allow mobility, and manually moving content is not practical every time you move. Unlike applications in a mobile phone, the physical environment does matter in AR. The new room might not have not a table over which to spread your documents, or it might have a different size and position. Other common solutions, such as attaching the information to your body or head, also have issues. They can, for instance, occlude the real environment and limit your awareness of the surroundings. In addition, it negates the benefits of a closer integration between AR and the world.

If AR is to be truly mobile, interfaces should reflect the way we seamlessly move around in the world. This entails not only being able to access information in different places but also during short or long walks. What if your workspace could follow you around, and when desired, quickly adapt to your new location and task? A good design would also need consider new risks coming from increased multi-tasking, divided attention, and information overload.

The main goal of this research was to investigate what such a system might look like. Instead of designing static interfaces (such as a head-up display or world-fixed information displays), we considered dynamic interfaces that adapt to the user's movement and to the physical environment. These adaptive behaviors could

---

IUI '19, March 17–20, 2019, Marina del Rey, CA, USA

ACM ISBN 978-1-4503-6272-6/19/03…$15.00
https://doi.org/10.1145/3301275.3302278

be triggered by changes in the user's position and orientation, and also by changes in the physical space (Figure 1). By sensing when the user walks, interfaces could move to a more suitable position, change its layout, or switch the primary interaction technique. By sensing the environment, an adaptive interface could maintain its consistency with the world and take advantage of surfaces such as tables and walls.

We designed a variety of behaviors to provide adaptive content management in mobile AR, implemented those behaviors in a modular system that can combine those individual behaviors, and proposed a final minimal set of useful behaviors that can be easily controlled by the user in a variety of mobile and stationary tasks. We also used this system to capture user experience during contextual studies and improve our general understanding of the design requirements for mobile AR workspaces.

2 RELATED WORK

Here we review some of the prior work on AR information displays, 3D workspaces, and walking performance under dual-tasking.

2.1 AR Information displays

Feiner et al. [12] describe an X11 window system for AR that supports three configurations for windows: 1) fixed relative to the HMD; 2) fixed to locations and objects; or 3) fixed to a sphere that surrounds the user. Our studies included behaviors similar to those. Unlike the authors, though, we evaluated translation and rotation components independently. Another early desktop manager was ARWin. Although still based on the X11 manager, ARWin could render the content of each window into a polygon mesh, creating fully 3D windows. The position of each mesh was determined in the world coordinates using markers [11]. Like ARWin, our system makes use of 3D windows. However, our windows can move anywhere in space freely.

No matter which configuration is used, one can potentially have a large amount of information displayed in a limited field of view. Bell et al. [3] introduced algorithms for view management, which had the goal of automating the layout of annotations in the viewing plane as the user moves. This allows the system to optimize the layout for given set of constrains, such as maintaining the visibility of real objects. All the behaviors we explore in this work consider the initial position of the elements and avoid self-overlap. Our goals, though, were different. Our intent was to respect the user preferences for the layout as much as possible. The initial arrangement is 3D and the elements may move inside or outside the view according to the context.

Based on these and prior work, Müller & Dauenhauer [22] proposed a taxonomy for information annotation in AR. However, for the coordinate dimension, they adopt only two coordinates: a world coordinate system (WCS) and a spectator coordinate system (SCS). The taxonomy of Müller & Dauenhauer is similar to ours, in the sense of considering the separate effect of translations and rotations. However, since they use a single reference, they cannot differentiate between the same pose in different coordinate frames.

Finally, Wither et al. [30] describe how AR annotations are related to the environment with the concepts of "location complexity" and "location movement". Considering movements within a coordinate frame can increase the descriptive power of layout annotations. However, only measuring the freedom distance in 3D space does not describe what the movement does. Our adaptation taxonomy tries to improve on this idea by expressing how an element moves (e.g., employing user movement, or using world surfaces and normals as references).

2.2 Walking, Perception, and Attention

Walking is a complex activity that requires one to be cognizant of the destination, the surrounding environment, and to be able to coordinate the limbs to successfully reach the destination. For this reason, cognitive tasks such as talking, doing arithmetic, counting, etc., impact overall walking performance [31]. In a meta study, Al-Yahya et al. [1] concluded that these dual-tasks cause changes in walking speed, cadence and stride characteristics. AR devices can also make walking harder if the content is displayed in way that reduces the visibility of obstacles or other features required for successful navigation.

Sedighi et al. [28] compared the effect of using smart glasses, smartphone, and a paper notebook on gait variability. Participants were asked to perform three different cognitive tasks while sitting in a chair or walking in a treadmill. The authors found that the risk of fall was higher during the dual-task, but participants used more adaptable gait strategies with the head-up display, which might help decrease the risk of falling. Our study prioritized ecological validity, so that we could observe strategies adopted in more realistic settings. Participants walked inside indoors office space, and were free to adjust speed or stop when required. Our system also allowed participants to try not only a classic (display-fixed) head-up display configuration, but also other dynamic layouts. In addition, they could adjust the content layout to match personal preferences and reduce occlusion.

2.3 Adaptation in Desktop Interfaces

Adaptation is a strategy to improve user interaction by optimizing the interface to the way it is being used. Park & Han [23] describe four categories of adaptation: Adaptable interfaces allow the user to manually customize the interface so that items of interest are more visible and accessible. In adaptive interfaces, however, this role is delegated to the system, which automatically modifies the interface according to a policy (e.g., prioritize last used item or most frequently used one). Both options have advantages and disadvantages. An adaptable interface, for example, is less confusing and easier to remember, but creates an additional effort to manually modify the interface [13, 18, 21, 24]. The authors also describe mixed approaches, in which the responsibility for determining and carrying out the adaptation is shared between the user and the system. In adaptable interfaces with user support, the system recognizes when adaptation is required, but the user is responsible for determining and carrying out the adaptation. Alternatively, in adaptive interfaces with user control, the system recognizes and carry out the adaptation, but the role of determining the adaptation is shared by both the user and the system. Our approach constitutes a third mixed approach, in which the adaptation is recognized and

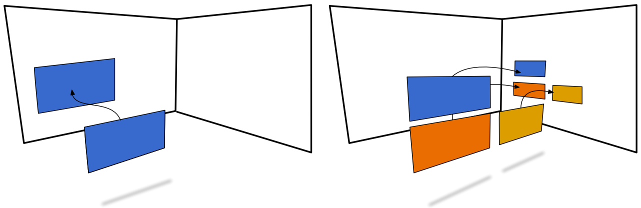

Figure 2: Windows adapting to nearby walls. Left - Window attaching to the nearest wall. Right - Adapting while respecting the workspace layout.

determined by the user, however it is carried out by the system. In this work we refer to it generally as an adaptive interface.

Bouzit et al. [5] propose a design space for adaptive menus, the most studied area in interface adaptation. By looking into different visual aspects of user interfaces, they discuss how each visual attribute can be used to construct adaptive 2D interfaces. Position-changing menus, for instance, refer to adaptations manifested by changes in position. Although not discussed in detail, the authors imply that there is a stability/adaptation tradeoff: more stability allows the system to be more predictable for the user, at the cost of the ability to adapt. In our work, we try to preserve relative spatial stability: the interface elements change position, but keep their relative position as constant as possible. All other properties are preserved: there are no time discontinuities and no changes in size or format.

3 DESIGNING ADAPTIVE AR WORKSPACES

Our initial vision was to use adaptation as a way to: 1) support the use of AR interfaces while walking; and 2) minimize the effort required to adjust the layout to a new environment. For example, a user watching a game in the living room may decide to have some food in the kitchen. Instead of staying in the living room, the window displaying the game could follow the user around the house and as he stops in the kitchen, enlarge and attach to the nearest wall (Figure 2 left). A layout with several windows could work in a similar way, with the additional constraint of considering both the initial layout and the new surrounding physical space to find a good match (Figure 2 right). In this concept, the entire adaptation could be defined by a set of behaviors, such as follow, or attract, that when combined would produce the desired outcome.

3.1 The Design Space

To address the design problem in a more systematic way, we proceeded to expand and detail the elements of our initial vision. Following the work of Zimmermann et al. [32] on contextual adaptation, we built upon the framework of Brusilovsky et al. [9] of adaptive hypermedia, which classifies adaptive systems on four dimensions: adaptation goals, features to adapt, features used for adaptation, and adaptation method. We extend it to include information about the context (in our case, the physical space around the user) and the effect of the adaptation on the visual aspects (Figure 3).

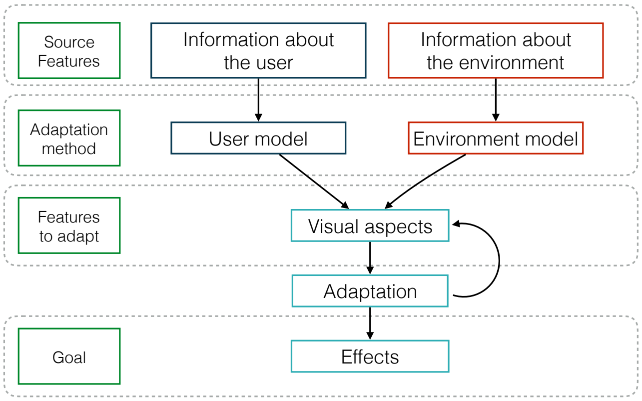

Figure 3: Adaptive System Model. It was extended from Brusilovsky et al. to include the environment and the reflexive term from the system. Green boxes indicate the corresponding dimensions of adaptation.

Although the goal of most adaptive systems is to improve visual-motor performance on the interface component level (e.g., menus), our adaptation goal was to improve the experience within the context changes caused by walking. Our first choice was to decide which features to adapt. Although several visual aspects could be used, we focused on position-adaptation [5]. Position has the highest generality across applications, since it is a feature common to all spatial content.

Within all the possible sources of position adaptation, we focused on those that could be easily obtained from sensors available in AR headsets, in particular:

(1) information about the user position in space (e.g., body position, orientation, or speed)
(2) information about the physical surroundings (e.g., position of walls and other surfaces)
(3) information about the current layout (relative position of the windows to each other and in the field of view)

This taxonomy can describe many layout categories discussed in the literature while offering a practical way to include interactions with the environment as well. For example, adapting to compensate for both head rotation and head position keeps the windows fixed relative to the head. This is equivalent to a display-fixed interface (or SCS positioned & SCS oriented in Müller & Dauenhauer’s [22] taxonomy). Alternatively, adapting to users by matching their changes in position but ignoring rotation is equivalent to Feiner’s [12] spherical surround configuration (also known as User-fixed or SCS positioned & WCS oriented). It is also possible to generate behaviors that are not easily described by a single reference frame. For example, interfaces that match simultaneously both the position of the user and of vertical physical surfaces will follow the user around while staying on the walls. This behavior cannot be described by a fixed position in the user coordinate frame, nor to a fixed position in the world coordinate frame.

3.2 Behavior Selection

We began our exploration of the design space by first implementing several alternative behaviors for three major source features: we implemented eight behaviors based on user's spatial information, four based on display layout and field of view, and six involving adaptation to the physical environment (Table 1). We also investigated more complex conditional behaviors that could not be easily described by combinations, such as:

- Adapting to head orientation only if walking speed was greater than a specified value (only adapts when the user is going somewhere)
- Adapting to the body position only if the distance between the user and the element exceeded a threshold (allows the user to walk some distance without disturbing the windows)
- Sticking to the walls only if the distance from the element to a wall was less than a threshold (prevents distant walls from attracting the interface)
- Moving elements to the front of the user when a button was pressed (gives the users the chance to prioritize legibility)

Evaluating all the combinations of these initial behaviors would be unfeasible. To reduce our initial design space to the most promising sets, we evaluated them by conducting an informal formative evaluation [27]. Four virtual environments experts evaluated the usability and potential utility of each one in walking scenarios. This step led us to combine parameters that did not work well independently and to fix default values for a few behaviors. We decided to:

- Combine the yaw and pitch in a single rotation parameter and ignore changes in head roll. After experimentation, we saw that users tend to understand yaw and pitch as part of the same integral space. On the other hand, we expected the up axis of elements to always point up so it was fixed to align with the up vector of the world.
- Always detect collision with physical surfaces. We explored the options of having virtual elements ignore physical walls or to actually collide with them. We decided for the latter, since it would enable users to take advantage of physical surfaces in space and also prevent depth mismatch (seeing one thing inside another).
- Try to always respect the relative layout of elements. One of the options discussed was regarding how much freedom to give to elements as they adapt to changes (in particular changes in the local physical space). If elements have complete freedom, they can always rearrange themselves optimally in relation to the available space. For example, during adaptation a larger element could go to a large wall and a smaller element to a small one. However, arbitrary rearrangement could generate confusion, since the user could not use spatial memory to quickly locate elements. A compromise was to allow adaptation but preserve the relative position of the elements with respect to each other.
- Combine some of the behaviors into two new "smart" behaviors. A few ideas could not be implemented by activating basic options together. Auto-centering was a behavior designed to automatically move and align the windows in front of the user during walking, following a body-centric reference frame. The Attraction behavior moved and placed the elements along the wall closest to the user, if the user was reasonably close. If the user moved away, the elements would detach from the wall.

By the end of the process, we reduced the initial set of options to four basic behaviors that demonstrated high potential to support different goals and interaction styles:

> **[Follow]** Adaptation to the user position. This behavior maintains the relative position between the element and the user but ignores body and head rotation. Elements will follow the user while keeping their initial world orientation.
> **[Rotation]** Adaptation to the user orientation. This behavior maintains the relative position and orientation of the element regarding the user's head. As the user rotates, the element will rotate around the user to keep the same position in the field of view.
> **[Attraction]** Adaptation to surfaces. This behavior moves interface elements to the surface closest to the user while interpolating between the initial orientation and the wall normal. As it gets closer to a wall, the element gradually aligns itself and attach to the closest wall surface.
> **[Auto-Centering]** Adaptation to user movement while walking. Elements follow the user as in Follow, but rotate to align the forward direction to the user's walking trajectory. This allows the user to look around to see different interface elements while walking. The elements behave as if attached to the body when walking and world-fixed when the user is stationary.

In addition, each behavior could also be temporarily activated by using a button in the controller (next section). Figure 4 shows how adaptation happens with each behavior.

4 AR ADAPTIVE WORKSPACE IMPLEMENTATION

Once the final behaviors were selected, we developed a prototype using Unity 2017 and two Microsoft HoloLens optical see-through HMDs. One HoloLens was worn by the user and one other was used by the researcher to monitor the operation of the system. The device was chosen for being untethered and able to reconstruct nearby geometry. In addition, it supports hand gestures for interaction, making it a good platform for our application. We also used an xBox 360 Bluetooth controller for user input.

4.1 Adaptation Approach

Our implementation approach was inspired by behavioral robotics [7, 8], in which agents adapt to external conditions using only local information about the world. This is a useful feature for real-time implementation in a real device, as it is fast and works even with incomplete knowledge of the environment. In this way, interface elements follow a gradient descent over a potential field created by surfaces and other physical structures in space [2, 15]. This approach allowed us to quickly prototype and test several behaviors. For example, weights can be set so that elements can move towards

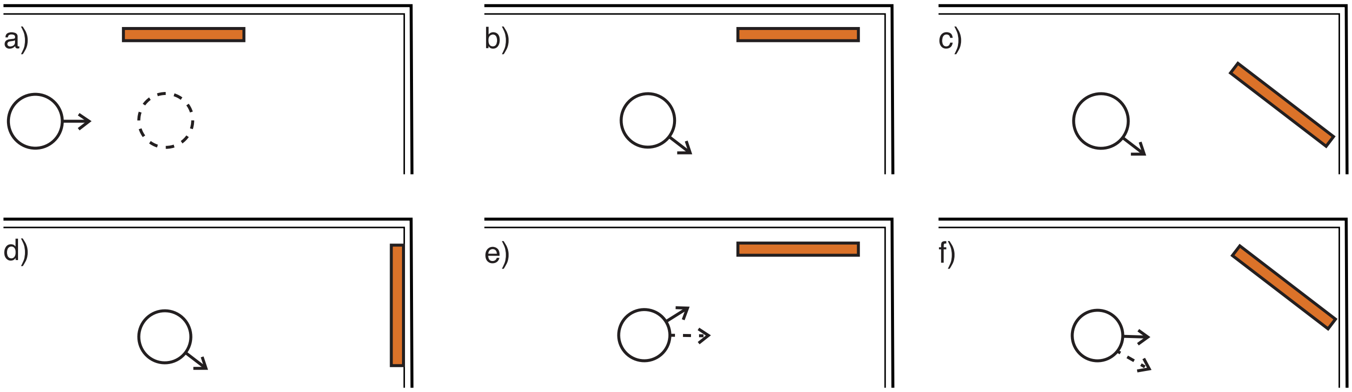

Figure 4: Schematic illustration of the selected behaviors. The circle indicates the position of the user relative to a single window (orange). The solid arrow indicates the view direction and the dashed arrow indicates the walk direction. The user walks from the initial position towards the corner of the room (a). With Follow, the window moves along but ignores any rotation (b). With Rotation and Follow, the windows maintains its relative position and orientation regarding the view direction (c). Attraction moves and aligns the window with the front-facing wall (d). With Auto-Centering, the window adapts to the body orientation and position, giving the user freedom to look around while walking (e and f)

Table 1: Behaviors implemented based on a single source feature

User Spatial Information	Display Parameters	Physical Environment Information
Ignore user position	Maintain the relative arrangement	Ignore physical obstacles
Follow user position	Spread around physical space	Collide with physical surfaces
Maintain original world rotation	Ignore changes in angular size	Stick on walls upon collision
Align with walking direction	Maintain the same angular size with distance	Ignore surface orientation
Ignore head pitch		Align proportionally to distance
Follow user look direction		Align only upon contact
Ignore head roll		
Rotate with head roll		

the user, slide along walls, or to transition between them. A drawback of this approach is that it is susceptible to local minima and can create problems in narrow passages [16]. An interface element can get stuck in an equipotential field created between two walls, or in front of an obstacle.

4.2 Implementation Details

Our adaptation strategy considered three information sources: user movement, environment surfaces, and the relationship between the windows in a given layout. Our general approach involved:

1. Assigning an attraction field to each wall and to the window initial position
2. Restricting the degrees of freedom of the window according to the behavior configuration
3. Assigning weights to the attraction fields according to the behavior configuration
4. Computing the potential field gradient at the window location
5. Setting the speed vector of the window to the negative of the gradient

The user pose was obtained directly from the HoloLens in real time. Since the HoloLens is head-worn, body orientation was estimated using the velocity vector. In addition, the position was filtered by a running average to remove head bouncing.

The environment information could also be obtained from HoloLens, by accessing the spatial mesh data created by the self-localization system. However, the scanned mesh can include people walking nearby, creating collision "ghosts". The mapping is also slower than average walking speeds and has a limited range. Our solution was to manually indicate the location and extent of the walls and other surfaces of interest by drawing them in an editing mode. This allowed us to prototype and evaluate adaptive behaviors that involved surfaces despite the limitations of current mapping technology. An ideal future implementation of our approach would not require manual specification of surfaces. These surfaces are then saved as a HoloLens spatial anchor, and can be reused again when the room is recognized by the device. Each rectangular shape is associated with a finite equipotential surface during runtime (Figure 5).

The two HoloLenses were networked using a text-based UDP/IP protocol. The reference frames were synchronized by setting spatial anchors at specific positions on the floor. We used several anchors distributed around, so that the system was not forced to use an anchor which was potentially too far to be reliably detected and

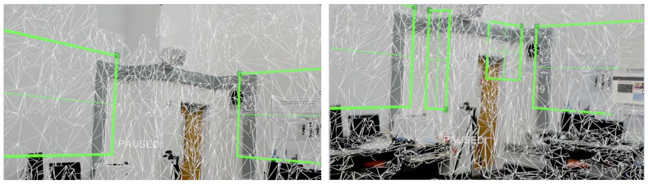

Figure 5: Wall Editor. Green rectangles define the position and extent of attractive surfaces.

tracked. We also added support for adjusting the position of the windows using the hands. The user can do the tap-and-hold gesture to grab a window and move it around. When the user opens the hand, the window is released at its current position, reoriented to face the user, and its original position is reset to the new location.

4.3 Behavior Configuration

To explore user preferences, it was necessary to individually disable or enable the four behaviors implemented (follow, rotation, attraction, and auto-centering). We also felt that the ability to temporarily enable a behavior could lead to interesting usage strategies. To achieve both goals, we implemented a configuration menu. The menu could be opened or dismissed by pressing the X button on the xBox controller. Once opened, users could air-tap the corresponding option to activate, deactivate, or assign a behavior to the controller's B button. If a behavior was assigned to this button, it would be active as long as the B button was pressed. Multiple behaviors could be assigned to the button simultaneously. The configuration parameters were used for all the windows.

5 EVALUATION OF THE AR WORKSPACE

We conducted an exploratory research to investigate the use of adaptive AR workspaces for walking. Our goal was to learn which combinations were preferred and why. We also wanted to find out how participants would use the available behaviors, as they can be used in different ways. To do so, we asked participants to try various combinations of adaptive behaviors and use them to perform realistic everyday tasks while walking.

The qualitative studies were semi-structured [29], and we applied thematic analysis to identify, analyze, and report the patterns from the data [6]. The data consisted of system logs, interview questions, and observations. The interview questions were adapted or modified when necessary to keep a productive dialogue with the researcher. During the analysis, we adopted an inductive approach in which themes were identified from the data without the use of a-priori categories. After repeated patterns were aggregated into categories, those were interpreted in the light of known principles and theories. Some codes, though, reflected the specific behaviors implemented in our system. So, our approach was both inductive and deductive, starting from the phenomenon, capturing the data and proceeding to the explanation [10, 25].

During the contextual studies, our system was used as a technology probe [14] that allowed us to: 1) identify the strengths and limitations of our adaptation approach to walking conditions, 2) investigate how participants would appropriate AR to their needs and interaction styles in real settings, and 3) engage participants in co-creative design to generate new ideas and directions for walking interfaces.

5.1 Scenario Design

In order to evaluate the adaptive workspace system, we designed two scenarios. Our goal was to create a setting that was realistic, and involve participants in tasks that were reasonably complex and believable. In the first scenario, participants were asked to assume the role of an interior designer. They would then meet the owner of an apartment (experimenter) and review some of what was discussed in a previous meeting. In this scenario, walking was introduced as an inherent part of the activity. In the second scenario, the participants would pretend to be talking to a friend (experimenter) about news and trying to schedule a day to meet at the beach. This scenario captures concurrent activities that we perform daily (e.g., taking a call while walking to the garage or checking the calendar while leaving for lunch).

The scenarios were designed so that a variety of content could be used. The information for both scenarios was distributed in the following applications: web page, list of prices, floor plan, notes, calendar, and weather forecast. Each one consisted of a single static window pane (an image). Each scenario required the use of three windows, two in conjunction and one isolated. The interior design scenario required the notes, floorplan and price list (Figure 6, first three windows). The experimenter would ask questions about changes in the apartment, prices of furniture, and their location in the apartment. The second scenario required the weather forecast, calendar and webpage (Figure 6, last three windows). The experimenter would ask questions about news, and then try to schedule a day to meet at the beach (a day with good weather and when the participant was also free).

The tasks, then, required the participant to walk and: 1) listen to and interpret the experimenter's question, 2) find the window(s) with the required information in the layout, 3) retrieve the information, and 4) verbally express the answer. In most cases what followed was a natural conversation between the experimenter and the participant.

5.2 Physical Environments

The study was conducted in two spaces in which the participants could walk alongside the experimenter. Each scenario was run separately. In the first scenario, we asked participants to walk around tables inside a large lab. The path was approximately 40m in length . In the second scenario, participants walked around 60m along a corridor.

5.3 Data collection

We collected different types of information during the study, and most of it was converted to text format (e.g., questionnaires and observations) so that it could be included in the coding process. The experimenter used a second HoloLens to monitor what the participant was doing during the study. Besides the position of each window, the experimenter was also able to see the current behavior configuration and button presses from the participant. After each trial, we interviewed the participants regarding their choice of behaviors, layout of the windows, how much their expectations were

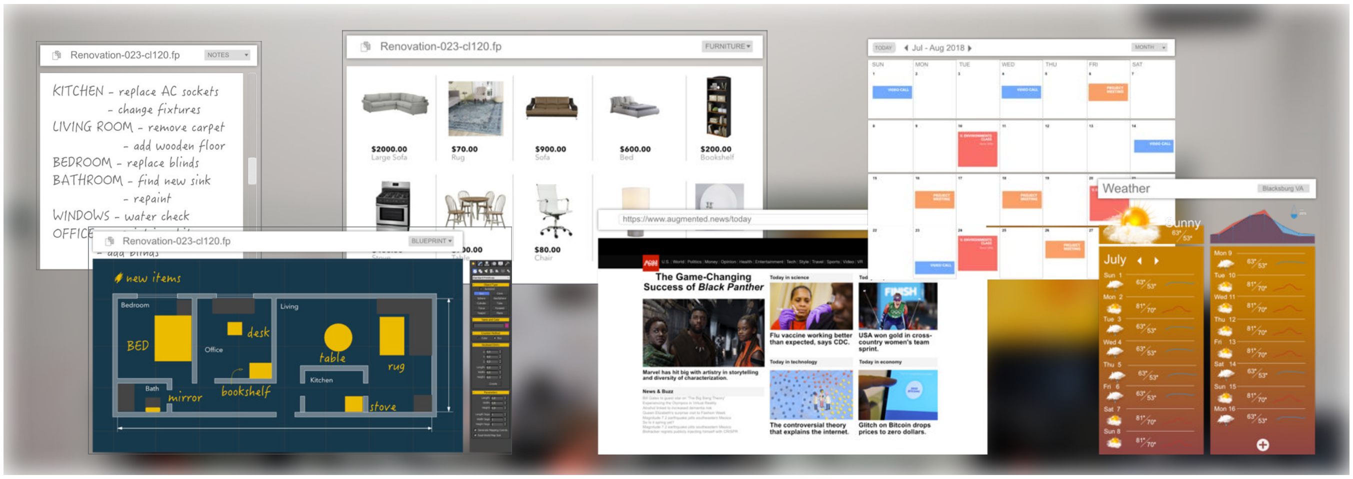
KITCHEN - replace AC sockets
- change fixtures
LIVING ROOM - remove carpet
- add wooden floor
BEDROOM - replace blinds
BATHROOM - find new sink
- repaint
WINDOWS - water check
OFFICE

The Game-Changing Success of *Black Panther*

Marvel has hit big with artistry in storytelling and diversity of characterization.

Figure 6: Windows used in the study, right to left: notes, floor plan, price list, web page, calendar, and weather forecast

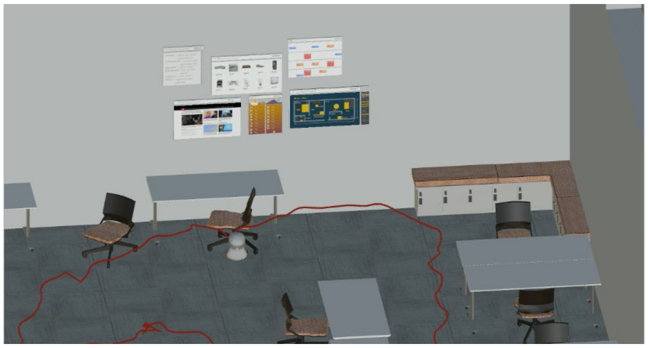

Figure 7: Screenshot from the Playback system. Red lines indicate the user path in the space.

met during the trial. We also asked them to elaborate on differences between trials, what they liked, what they did not like and what they would like to see changed or added. The interviews were audio recorded and transcribed. We also recorded the trajectory and head orientation of the participants in space, the configuration selected, the activation of behaviors using the joystick, and the position of all the windows throughout each trial. We used this log during analysis to replay each trial in 3D as a supplementary source to understand and confirm events (Figure 7).

5.4 Participants and Procedure

We recruited 14 university students with corrected or normal vision. Ten participants were males and four females. Eight of them had used AR just once or twice. Upon arrival, participants were asked to read and sign the consent form. Next, they were asked to fill out a background questionnaire containing demographic questions, a question about experience with VR, and another about experience with computer games. Next, the experimenter explained the goals of the study, described the scenario, and explained the interaction method. Next, they were introduced to the HoloLens, guided on how to adjust it, and on how to perform the HoloLens calibration procedure. The experimenter then walked the participants through each behavior, activating one at a time and explaining how each one worked. The experimenter also trained the participants to grab and move the windows around to configure the layout. Participants then performed at least two trials with each scenario. Each trial consisted of thinking and deciding which combination of behaviors to try, adjusting the windows, walking and performing the tasks, and joining in a semi-structured interview. The participants were instructed to follow the experimenter as he walked along the path at a comfortable pace. If the participant stopped to answer, the experimenter asked the participant to start walking again (if they did not do so after answering). During the final interview, we asked participants to order each behavior in terms of decreasing utility.

5.5 Analysis

We used the qualitative analysis software AQUAD 7 to code approximately 400 lines of text. Coding of the data was performed in two cycles. In the first cycle, we mostly employed Open Coding, breaking down the data into parts and looking for commonalities and differences [26]. The codes were developed following the principle of "constant comparisons" to ensure consistency [10]. Some of them were divided into a positive and a negative aspect to facilitate analysis.

We also applied Holistic Coding to identify passages related to our probe design (e.g., behaviors, configurations, etc.) and Structural Coding to identify answers from recurrent interview questions. We applied descriptive codes for other passages (usability of the behaviors, preferences, strategies adopted by the participants, reading, content, safety concerns, etc.). The codes for demographic information were singular and were used to characterize each participant file (Attribute Coding). Some super-ordinate codes were also created to identify passages as belonging to different trials and locations. These codes were used later to facilitate sub-coding.

In a second cycle, the L1 codes were grouped into categories. To investigate characteristics of the behaviors, we applied Axial Coding . In this method, the fractured data from the first cycle is recombined to retrieve larger categories [26]. We grouped codes along attribute dimensions and applied the C-family coding [4] (causes, contexts, contingencies, consequences, and conditions) to identify possible causal conditions, contexts and consequences of phenomena. We also looked for patterns expressed by expressions such as "if/else", "because" that could help to understand participant's recurrent patterns (Pattern Coding).

6 FINDINGS

To focus the analysis, we first compared the combination codes in importance. We looked at how frequently each combination code appeared in data from different participants in the first and in the second scenarios. The most used combinations were [Follow-on, Auto-Centering-On] (twelve times), [Follow-on, Rotation-Button] (twelve times), [Rotation-Button, Auto-Centering-On] (four times), and [Follow-On, Attraction] (four times).

We also looked into the individual ranks given by the participants and computed the average rank, assigning weights 3, 2, and 1 to the first, second, and third behaviors (follow was not considered since it was necessary to complete the task). Rotation was considered the most useful, with rank 26, followed by Auto-Centering and Attraction, with rank 20.

6.1 Behavior Perceptions - Which behavior characteristics participants deemed useful?

To investigate the reasons for the preferences, we sub-coded all passages describing those behaviors. This allowed us to capture specific meanings, to keep track of the bigger picture, and to establish dimensions for those codes. Table 2 list the codes which appeared at least twice during the analysis. The Auto-Centering behavior was deemed convenient and easy, but had the perceived disadvantage of taking some time to respond and not allowing precise control of the windows. On the other hand, the combination [Follow-On, Rotation-Button] was judged as fast and precise, at the expense of requiring manual control. The attraction behavior was considered safe and appropriate for stationary conditions. However, it was hard to predict where the windows would attach and it was not useful on lateral walls.

In the second cycle, we engaged in axial coding, searching for relationships between the categories and identifying thematic axes. We found three major dimensions which characterize the appropriateness of the behaviors for the scenarios we studied: 1) compatibility with other real-world activities, 2) compatibility with vision-body coordination involved in walking, and 3) the tradeoff between usability and expressiveness of the interface.

1. **Compatibility with other real-world activities** - Participants noted, during the study, that having many windows too close blocked the view of the surroundings. It also made interaction with the world harder. For example, when talking to someone or manipulating object on a desk: "if I need to be doing some delicate operation it would be more user-friendly keeping them on the closest wall and when I require I could just look at that." Participants also mentioned during the last phase of Scenario 1 that it was hard to get the windows in a good place when using auto-centering: "auto-centering is not very useful mostly because it can get between you and other people very easily and prevent you from seeing what is happening at close quarters".
2. **Compatibility with vision-body coordination involved in walking** - One of the challenges of the task was to read the content while navigating around desks and other obstacles in the room. Even for those with some familiarity with the space, walking was hard if they could not look where they were walking. "if the windows takes [sic] the whole view of the corridor, you would not be able to see if someone was coming". Even though most participants claimed they could see clearly through the windows, a few participants mentioned that they would like the windows to "occlude less of the actual visual space." "The good side [of attraction] is that it was not always popping in front of me. I could see the space more clear in front of me ... where I am going." Participants also mentioned that it was hard to read a window on the side (which was positioned there, in the process of moving to another location, or attracted to a lateral wall). "Attraction is not very useful if you are not moving towards the wall, because you need to keep your head at 90 degrees' angle, and you cannot see what is front of you. You are more likely to bump into objects."
3. **Good tradeoff between usability and expressiveness** - The main advantage recognized by the users for the auto-centering and attraction modes was the fact that they did not require any input from the user. Some participants felt that using the button was an additional overhead: "I liked auto-centering better, because with the rotation I had to keep pressing the b-button." Occasionally, they forgot to press the button and one expressed the concern that it could become tiring over time. The ability to perform multiple functions was recognized for the [Follow-on, Rotation-button] configuration as a benefit.

6.2 Interaction Patterns - How and why participants used different behaviors?

To find out how participants were using the behaviors, we looked into the L1 codes for those related with participants' actions. We applied the C-Family coding (context, causes, contingencies, consequences and conditions) to reconstruct the course of actions used to achieve a specific purpose (Glaser 1978; Böhm 2004; Strauss & Corbin 1990).

We considered "actions" the participant's response to a phenomenon. Following the recommendation of Strauss and Corbin we did a functional analysis and considered both conscious and unconscious actions [4]. The events or conditions that led to the phenomenon and participant action were considered "causes". If a condition acted as a modifier to the phenomenon, or delimited the set of possible user action, they were classified as a part of the "context" [10]. For example, a commentaries such as: "if you put them on the walls, you are not going to walk through a wall anyway, so it is safer" or "If I need to be doing some delicate operation it would be more user-friendly keeping them on the closest wall" indicate that given a stationary or moving context, if the participant could not see the world an action would be "enable attraction" which resulted in the clear surroundings.

Table 2: Most Frequent Codes Related to Each Behavior

Auto-Centering	Rotation on Button	Attraction
Takes too long to align	Allows precise control (with button)	Is hard do predict / control
Does not allow control of position	Restrictive if always on	Is good on frontal but not lateral walls
Is convenient	Fast	Is appropriate for stationary conditions
Works as long there is no interaction	Requires users to automate correct button pressing	Makes walking safer
Allow freedom to look around		Requires Stopping

Grouping the codes, we found four core categories related to the interaction patterns participants adopted during the study.

1. **Obtaining Information** This category groups the strategies used to obtain the information necessary to answer the questions asked during the study. Participants stated and were observed using the combination [Follow-On, Rotation-Button], [Follow-Button, Rotation-Button], and Auto-centering to achieve this goal. All these combinations brought the windows to the locations the participant previously established. The tradeoff was the occlusion of the surrounding by the windows. In addition, when they were close enough to read, it was not possible to fit all the windows in the field of view, which required the participants to turn their heads.
2. **Managing Information Layout** This category groups the actions taken to organize the windows so that they could be clearly visible and help participants find where to obtain the information they needed. If the target window was close, but in an undesired position, some participants would drag to rearrange them again. This often happened because the initial configuration set the windows was too far away or on one of the sides. Hand adjustment of the layout also happened when windows would accidentally overlap each other. Another strategy used to organize information was to open the Menu and enable Attraction so that the windows would move to the nearest wall. Since switching between [Follow-Button, Rotation-Button] or Auto-centering and Attraction was not possible without opening the Menu, this was not a common solution. However, participants did express the desire to switch between those (see suggestions analysis in this section).
3. **Compensating for Reduced Awareness** This category groups the actions taken to compensate for the lack of visibility of the physical world or the toll of attention caused when focusing on the windows to obtain information. When faced with this issue, participants either slowed down the pace or stopped. The lack of sensory data about the world happened for two reasons: 1) when users tightly packed the windows so that they could not see very well through it 2) when users were not looking in the direction they were walking (because the information was on the side or on a lateral wall). In addition, participants also indicated they slowed down or stopped because of the attention spent on reading / thinking about the content (even if it was right in front). In all cases, the act of slowing down and eventually stopping seem to have happened both consciously and unconsciously, as some participants reported to be unaware of doing so. Enabling attraction was also used as a way to cope with the lack of awareness of the environment. By moving the windows to the wall, a participant reported he was more comfortable walking.
4. **Harmonizing Augmented Reality with the World** This category groups the strategies used with the goal of making other interactions with the world more sensible or efficient. For example, when writing on the physical board one participant wanted to align the virtual floorplan with the physical floorplan to compare. Participants also expressed the desire to move windows out of the way to talk to someone or otherwise see the environment to perform other tasks. Two strategies were used or mentioned: 1) use the combination [Follow-On, Rotation-Button] to manipulate the windows and adjust their position in space, and 2) move the windows to the walls so that they naturally integrate with the physical surfaces.

7 DISCUSSION

Overall, our research showed that position adaptation is a promising way to tackle the challenges of walking. Participants had no issues learning our final set of behaviors or using them to accomplish the tasks. However, we learned that some properties are important in the performance of these behaviors, such as stability, predictability, speed, and control. Our implementation could still be optimized in those aspects. Due to the use of potential fields, our windows moved smoothly and with low acceleration. Thus, they took some time to rest at final positions. The auto-centering behavior also had a small latency due to the way we decoupled the body movement from the head movement (simple average). We believe that the smaller accelerations helped users to track the windows before they left the field of view. However, quickly moving the windows (instead of using a potential field) should allow techniques to perform better.

The primary impact of the dual-task was a performance penalty on walking. This result is consistent with previous reports of reduced walking speed in dual-task conditions, which was observed for diverse cognitive tasks and in different populations [1]. We observed that the task required high attentional demand, which likely caused a lack of environmental awareness. This observation is consistent with the participants' notes about "lack of attention" which eventually led to a few trips and bumps. It has been shown that users talking on the phone or texting have reduced awareness [19, 20]. However, it is also possible that they may have been caused by the direct impact in the gait parameters by the dual-task [31].

Some participants reported difficulty seeing the environment through the windows due to occlusion. In some cases, participants compensated the lack of awareness by moving the windows away. In others, they slowed down. Although locomotion does not require continuous input from vision [17], the study design placed high emphasis on the cognitive tasks. Even a good layout required users to switch the gaze between the windows and the environment. For the participants who did try to observe the environment, but did not engage in any active strategy to manage the interface, the layout of the windows might have helped. As users were free to adjust the windows, they could opt to leave more of the visual field empty.

7.1 Designing For Walking

Our study revealed three major features that techniques for walking-centered AR workspaces should have and four fundamental interaction patterns associated with them. These can be combined into the following design guidelines:

Incorporate support to compensate for reduced awareness and matching with the real world by:

- Allowing content positioning - Users may need to move AR content to a specific position in space to reflect new priorities, increase the comfort of the current task, or increase the semantic association with the world.
- Allowing physical interactions - Unlike VR, AR users are still engaged in the world. They may want to interact with objects (e.g., write something), engage in a conversation, etc. Make sure that AR interface allows unobstructed view when necessary.
- Not occluding the walking path - As part of world activities, users will walk around. To give users the confidence to walk, the interface should not block the user path. For example, allow users to look around the interface if possible.
- Supporting multiple uses - Users will be performing multiple tasks, so the cognitive and motor load of the interface should be kept at minimum. One possibility is to use automatic behaviors or provide a single operation with multiple uses.

7.2 Limitations and Future Work

Even though interface interaction was minimal during the tasks (at most pressing a button), some cognitive effort was needed to understand and predict the behavior of the windows in the environment. In future studies, an assessment of the user's perceived workload could be used to separate the impact of managing walking-centric AR UIs from the load caused by the task itself.

We also noticed that several participants intentionally slowed down or stopped to answer questions. Some could be afraid to trip and some mentioned that visual artifacts made reading uncomfortable when they were moving too fast (e.g., color separation). Additional studies are needed to separate these different causes and also separate more clearly intentional and unconscious walking performance degradation.

Regarding the behaviors, a few participants mentioned that the attraction to the nearest wall was not intuitive. Many also expressed the desire to switch between different combinations while walking, which our system did not support. Even though the experimenter walked through every behavior to make sure participants understood each one, we did not enforce them to try specific combinations. As a result, some participants might have given preference to behaviors that were easier to understand. However, most participants were eager to try different combinations on the subsequent trials, which allowed most useful combinations to be evaluated.

To simplify the configuration of the behaviors and to highlight the difference between them, all windows in the study followed the same behavior. In a fully developed AR workspace, though, each application could adopt the behavior more appropriate to its function. For example, interface elements that reference physical objects could orbit around a specific world position. Other elements could prioritize following the user around or occupying specific spots in each room.

The current study was also focused on 2D content, such as text and images in a traditional window layout. While this represents a likely transition scenario to future AR workspaces, adaptation involving 3D elements may require additional information about the environment (e.g., horizontal surfaces and empty volumes).

Finally, our research was conducted with young university students. It would be interesting to investigate whether other groups (e.g., different ages, background, or cognitive skills) display different preferences or reveal additional aspects that could be considered during adaptation.

8 CONCLUSIONS

In summary, participants appreciated the ability to have information with them while walking, in particular to quickly look up information. The majority also opted for automatic following, instead of moving the windows on demand. Auto-centering was considered convenient and straightforward to use, and the combination [Follow-On, Rotation-Button] was deemed more precise.

We also identified four fundamental interaction patterns in AR walking scenarios. Those comprehend the process of 1) managing the information layout, 2) obtaining the desired information, 3) compensating for the lack of awareness of the environment, and 4) adjusting virtual information so that it is presented harmonically with the world.

In addition, we were able to isolate the ideal set of properties interfaces should have to be effective. We found that interfaces should 1) be compatible with the vision-body coordination involved in walking, 2) be compatible with secondary real-world activities, and 3) display a good tradeoff between usability and expressiveness. Interfaces designed for walking should try to incorporate support to compensate for reduced awareness and matching with the real world.

ACKNOWLEDGMENTS

The authors gratefully acknowledge funding support from the Immersive Sciences program in the Office of Naval Research and from the Brazilian National Council for Scientific and Technological Development. The authors also are thankful for the suggestions from the anonymous reviewers.

REFERENCES

[1] Emad Al-Yahya, Helen Dawes, Lesley Smith, Andrea Dennis, Ken Howells, and Janet Cockburn. 2011. Cognitive motor interference while walking: a systematic review and meta-analysis. *Neuroscience and biobehavioral reviews* 35, 3 (Jan 2011), 715–28.

[2] J Randolph Andrews and Neville Hogan. 1983. *Impedance control as a framework for implementing obstacle avoidance in a manipulator.* Master's thesis. M. I. T., Dept. of Mechanical Engineering.

[3] Blaine Bell, Steven Feiner, and Tobias Höllerer. 2001. View management for virtual and augmented reality. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*. ACM, 101–110.

[4] A Böhm. 2004. Theoretical coding: text analysis in grounded theory. Dans U. Flick, E. von Kardorff, & I. Steinke (Éds), A companion to qualitative research (pp. 270-275).

[5] Sara Bouzit, Gaëlle Calvary, Denis Chêne, and Jean Vanderdonckt. 2016. A design space for engineering graphical adaptive menus. In *Proceedings of the 8th ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. ACM, 239–244.

[6] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.

[7] Rodney A Brooks. 1991. Intelligence without representation. *Artificial intelligence* 47, 1-3 (1991), 139–159.

[8] Rodney A Brooks. 1997. From earwigs to humans. *Robotics and autonomous systems* 20, 2-4 (1997), 291–304.

[9] Peter Brusilovsky, Alfred Kobsa, and Julita Vassileva. 1998. *Adaptive hypertext and hypermedia*. Springer.

[10] Juliet Corbin and Anselm Strauss. 1998. Basics of qualitative research: Techniques and procedures for developing grounded theory. *SAGE Publications, Thousand Oaks, USA* (1998).

[11] Stephen Di Verdi, Daniel Nurmi, and Tobias Hollerer. 2003. ARWin-a desktop augmented reality window manager. In *Mixed and Augmented Reality, 2003. Proceedings. The Second IEEE and ACM International Symposium on*. IEEE, 298–299.

[12] Steven Feiner, Blair MacIntyre, Marcus Haupt, and Eliot Solomon. 1993. Windows on the world: 2D windows for 3D augmented reality. *Proceedings of the 6th annual ACM symposium on User interface software and technology*, 145–155.

[13] Leah Findlater and Joanna McGrenere. 2004. A comparison of static, adaptive, and adaptable menus. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 89–96.

[14] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, and BjÃűrn Eiderbäck. 2003. Technology probes: inspiring design for and with families. *Conference on Human Factors in Computing Systems (CHI)*, 17–24.

[15] Oussama Khatib. 1986. Real-time obstacle avoidance for manipulators and mobile robots. In *Autonomous robot vehicles*. Springer, 396–404.

[16] Y. Koren and J. Borenstein. 1991. Potential field methods and their inherent limitations for mobile robot navigation. In *Robotics and Automation, 1991. Proceedings., 1991 IEEE International Conference on*. 1398–1404.

[17] M Laurent and J A Thomson. 1988. The role of visual information in control of a constrained locomotor task. *Journal of motor behavior* 20, 1 (Mar 1988), 17–37.

[18] Talia Lavie and Joachim Meyer. 2010. Benefits and costs of adaptive user interfaces. *International Journal of Human-Computer Studies* 68, 8 (2010), 508–524.

[19] Jongil Lim, Avelino Amado, Leo Sheehan, and Richard EA Van Emmerik. 2015. Dual task interference during walking: The effects of texting on situational awareness and gait stability. *Gait & posture* 42, 4 (2015), 466–471.

[20] Ming-I Brandon Lin and Yu-Ping Huang. 2017. The impact of walking while using a smartphone on pedestrians’ awareness of roadside events. *Accident Analysis & Prevention* 101 (2017), 87–96.

[21] Jeffrey Mitchell and Ben Shneiderman. 1989. Dynamic versus static menus: an exploratory comparison. *ACM SigCHI Bulletin* 20, 4 (1989), 33–37.

[22] Tobias Müller and Ralf Dauenhauer. 2016. A Taxonomy for Information Linking in Augmented Reality. In *Augmented Reality, Virtual Reality, and Computer Graphics*, Lucio Tommaso De Paolis and Antonio Mongelli (Eds.). Springer International Publishing, Cham, 368–387.

[23] Jungchul Park and Sung H Han. 2011. Complementary menus: Combining adaptable and adaptive approaches for menu interface. *International Journal of Industrial Ergonomics* 41, 3 (2011), 305–316.

[24] Jungchul Park, Sung H Han, Yong S Park, and Youngseok Cho. 2007. Adaptable versus adaptive menus on the desktop: Performance and user satisfaction. *International journal of industrial ergonomics* 37, 8 (2007), 675–684.

[25] Jo Reichertz. 2014. Abduction, deduction and induction in qualitative research. In *The SAGE Handbook of Qualitative Data Analysis*, Uwe Flick (Ed.). SAGE Publications Ltd, London, 123–135.

[26] Johnny Saldaña. 2015. *The coding manual for qualitative researchers*. Sage.

[27] Michael S Scriven. 1967. The methodology of evaluation (perspectives of curriculum evaluation, and aera monograph series on curriculum evaluation, no. 1). *Chicago: Rand NcNally* (1967).

[28] Alireza Sedighi, Sophia M Ulman, and Maury A Nussbaum. 2018. Information presentation through a head-worn display ("smart glasses") has a smaller influence on the temporal structure of gait variability during dual-task gait compared to handheld displays (paper-based system and smartphone). *PLoS one* 13, 4 (2018).

[29] Mads Soegaard and Rikke Friis Dam. 2012. The encyclopedia of human-computer interaction. *The Encyclopedia of Human-Computer Interaction* (2012).

[30] Jason Wither, Stephen DiVerdi, and Tobias Höllerer. 2009. Annotation in outdoor augmented reality. *Computers & Graphics* 33, 6 (2009), 679–689.

[31] Galit Yogev-Seligmann, Jeffrey M Hausdorff, and Nir Giladi. 2008. The role of executive function and attention in gait. *Movement disorders* 23, 3 (2008), 329–342.

[32] Andreas Zimmermann, Marcus Specht, and Andreas Lorenz. 2005. Personalization and context management. *User Modeling and User-Adapted Interaction* 15, 3-4 (2005), 275–302.