A Usability Study of FaciemAR: Assessing a Facial Recognition App in the Mixed Reality of Microsoft's HoloLens

by Elizabeth O. Hammond,

November 29, 2017

University of North Carolina, Chapel Hill

Steven King, Thesis Chair

Laura Ruel, Committee Member

Jason Cooper, Committee Member

Table of Contents

Introduction

Imagine standing in a room, gesturing toward a wall and seeing a little round window open up to reveal a lovely blue sky. You walk over, look through the window and see a majestic mountain range below the blue sky, much like the view from an airplane. You gesture toward the ceiling to open another window through which you see more blue sky, then you open a window in the floor. Walking over to inspect the hole in the floor you lean over it only to see the same vast mountain range, but you approach it carefully because you're slightly concerned that you could fall through. This is the magic of mixed reality (MR) through a head worn display (HWD) such as Microsoft's HoloLens. You know the holes and the mountains aren't really there, but they seem as real as any other part of your room.

Mixed Reality can be defined as an overlay of digital information on the real world. The user processes information from and performs actions in both the real environment and the virtual environment (Wagner et al., 2009, p. 250). In 1994 Milgram and Kishino described MR as a blending of real and virtual worlds along the virtuality continuum that separates the real from the virtual environments (see figure 1).

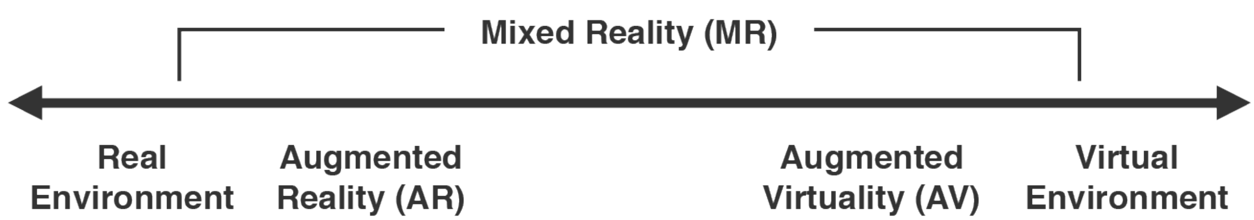

Figure 1. Milgram's virtuality continuum (Milgram, P., & Kishino, F., 1994, p.77)

For instance, using a vehicle's rear parking assist to exit from a parking space is closer to the real environment end of the spectrum while playing RoboRaid (see figure 2), in which the user makes gestures to shoot virtual robots that are emerging from the real walls, is closer to the virtual environment end of Milgram's virtuality continuum: the user still sees the real walls around them, but it looks and feels like the real walls are cracking open with robots.

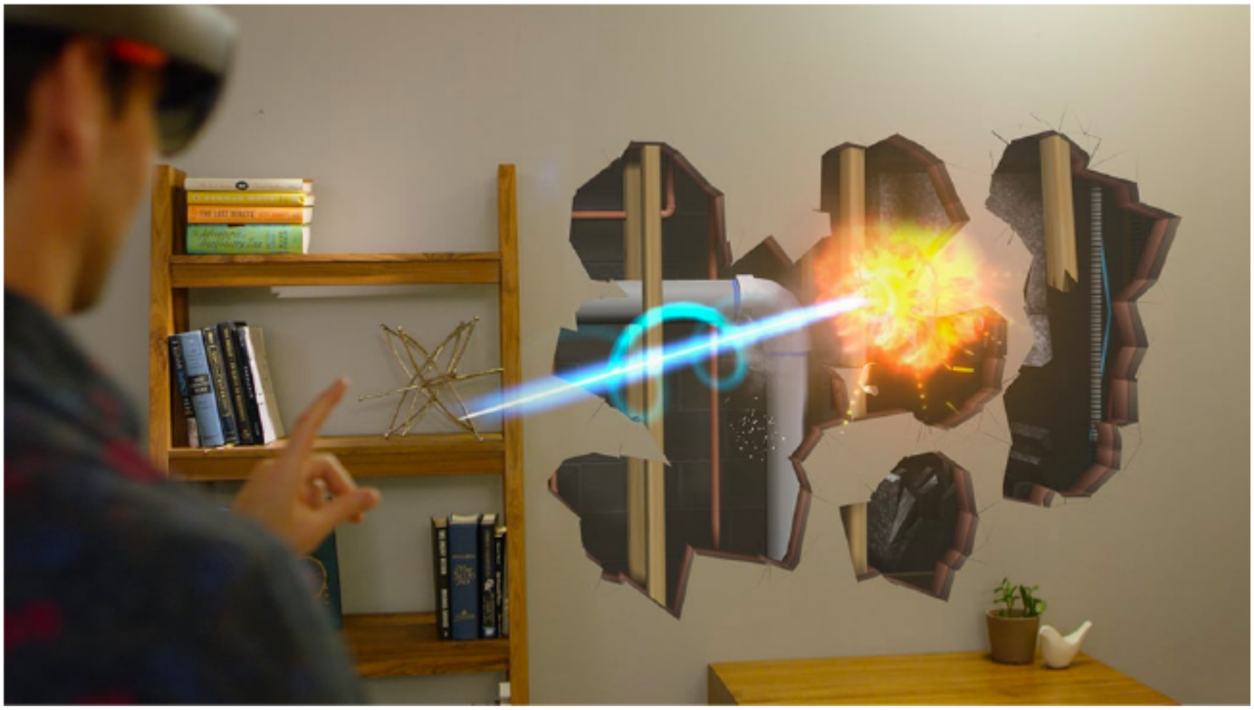

Figure 2. The game RoboRaid ("RoboRaid", n.d.)

The HoloLens, Microsoft's mixed reality device (see figure 3), is the device that will be used in this study. The HoloLens is a cordless HWD that shows the user digital information in the form of 2D objects, and 3D objects that Microsoft calls holograms. The HoloLens looks like a bulky pair of glasses with normal transparent lenses, but it's actually a very powerful wearable computer that allows users to see their real environment as if looking through glasses, while seeing holograms layered on top of their real world. Microsoft derived the term "Mixed Reality" from Milgram's study and uses it to describe the HoloLens as a device that sits between AR on the far left of the virtuality continuum, and VR on the far right. The HoloLens has the capacity to map a user's physical surroundings which make holograms and digital representations of other people, or avatars, appear as though they're really in the environment ("Mixed Reality" n.d.).

Figure 3. Microsoft HoloLens ("Microsoft Hololens", n.d.)

At the time of this writing the HoloLens is only available to developers and commercial organizations. Few applications have been developed for the HoloLens and it has not been adopted by the mainstream. With a limited sphere of use there is scarce knowledge about designing applications for the HoloLens in a way that offers the best user experience. The following literature review will explore the emerging guidelines for user interface design of MR experiences. Relying on the early experiences of designers and Microsoft's own research, it will attempt to cover what we know about the burgeoning field of interface design for MR. In addition, well-established and broadly applicable 2D user interface design guidelines will be woven in since they remain applicable to spatial experiences.

Literature Review

User Interface Design for Mixed Reality in Relation to Established Best Practices

Until recently we've primarily interacted with our computers through windows, icons, menus and pointers, or WIMP-based interfaces. Building upon the innate knowledge that many of us have of real world desks, files and folders, WIMP interfaces employ the metaphor of a desktop to make computers much easier to use than their predecessor, the command line (Kimani, 2009). Ever since the Macintosh popularized the WIMP interface style in 1984 referencing real world experiences has helped interface designs to bridge the gap between analog and digital worlds (Dix, 2009). These metaphors are part of a concept called affordance, a term coined by psychologist J.J. Gibson in the 1970's and applied to the field of design by Don Norman in the 1980s. An affordance is an indication of how something can be used. For instance, the design of draggable volume sliders indicate that the sliders can be moved, and the design of raised buttons and tabs indicate that they can be clicked. An affordance can be indicated by any number of characteristics like language, physical appearance, or patterns that users have grown accustomed to. For instance, Internet users expect that logos that are positioned in the top left of websites will be clickable (Kaptelinin, V., 2013).

It's broadly thought that MR interface design must also be informed by our real world experiences in order to succeed (Gribetz, 2016; Klemmer, Hartmann & Takayama, 2006; Nielsen, 1995; Sepp, 2017). With the advent of spatial realities Gribetz says, "The future of computers isn't locked inside a screen...it's locked inside of us" (2017), and Sepp explains, "While before, we have designed for screens that exist in the real world, now we are tasked with designing the world itself" (2017). Cordless HWDs like the Hololens give users the freedom to move around in the world in a way that flat screens can't. The overlay of digital information on the real world is a very different experience than what most of us have known before and will require new types of designs. User Experience designers who have been working solely in 2D environments will be challenged with designing for spatial experiences, which includes the additional responsibility of physical aspects like user fatigue and safety. To meet this challenge designers will need to experiment, create many prototypes and have an open mind. We're just beginning to write the rules for these spatial experience and don't yet have the research to inform our practices so this is a period of experimentation and development. (Hollomon, Kratchounova, Newton, Gildea, & Knecht, 2017; Howard, 2016; Rizzotto, 2017). As new user interfaces emerge those who perform user research will contribute work to the broad understanding of this new user experience (LaViola, Kruijff, McMahan, Bowman & Poupyrev, 2001, p. 95).

In this nascent stage of user interface design for spatial interfaces, designers are already taking cues from real world experiences. For instance the Meta 2, an AR headset by Meta, uses natural gestures to manipulate objects. To move an object the user simply grabs it and moves it. Pulling an object apart to inspect it requires a pulling apart motion (Gribetz, 2016). Sepp suggests pushing a screen to move it farther away, or pointing to a section of a virtual motorcycle that you want to paint and then using your hand like a brush (2017).

In addition to referencing the real world to design our interfaces there are a number of well-established principles for 2D designs that are also relevant to spatial experiences (Sepp, 2017). These best practices include subjects such as consistency of actions, terminology and visual design, reduction of cognitive load, error prevention and recovery, user control, and communicating system status (Nielsen, 1995; Shneiderman, 2005) (see Appendix A for Nielsen and Shneiderman's full guidelines). In lieu of empirical data based on user research of HoloLens applications, these seminal guidelines can help provide structure for our experimentation, prototyping and research.

Multimodal inputs. The HoloLens is primarily built to use gaze, gesture and voice but it's also capable of receiving input from a keyboard and mouse. Devices that allow two or more combined modes of inputs, or multimodal input, can be effective at giving the user more control and flexibility of use, lowering their cognitive load and reducing errors (LaViola et al., 2017, p. 409-411). Multimodal inputs also support some of the core concepts of user interface design guidelines, like user control and freedom, and flexibility and efficiency of use (Nielsen, 1995). For instance, choosing to use gestural input rather than voice commands in a noisy environment could reduce errors. Using the familiar keyboard instead of voice commands for a productivity application may result in more precise work (Howard, 2017). Different spatial experiences will benefit from different types of inputs. Instead of assuming that applications will use mice and keyboards, designers will need to consider the strengths and weaknesses of gaze, gestures and voice commands for their specific application and user base, and follow up with testing to validate.

Gaze. Gaze is the primary form of targeting in mixed reality and it matches the real world experience of looking at an object before interacting with it. Gaze is like a laser pointer emanating from directly between the user's eyes, and within the HoloLens MR environment it appears as small white dot, or as a donut shape when over an object that can be targeted, though its appearance can be customized. Like the movement of a mouse in a 2D system, a user's gaze is typically the first step to an interaction, and is often followed by a gesture that causes something to happen. Gaze is a powerful tool for creating dynamic experiences because it captures where the user is looking, or not looking, which gives interface designers the opportunity to cue the user to look in a certain direction, or present particular tools based on the targeted object ("Gaze" n.d.).

Using gaze to target an object is often more successful when working with larger targets ("Gaze Targeting", n.d.; LaViola et al., 2001, p. 401; Wilcox, Evans, Pearce, Pollard & Sundstedt, 2008). Microsoft writes that a target size of 1° is large enough for gaze targeting to be successful, but targets at 3° can be selected faster (see figure 4).

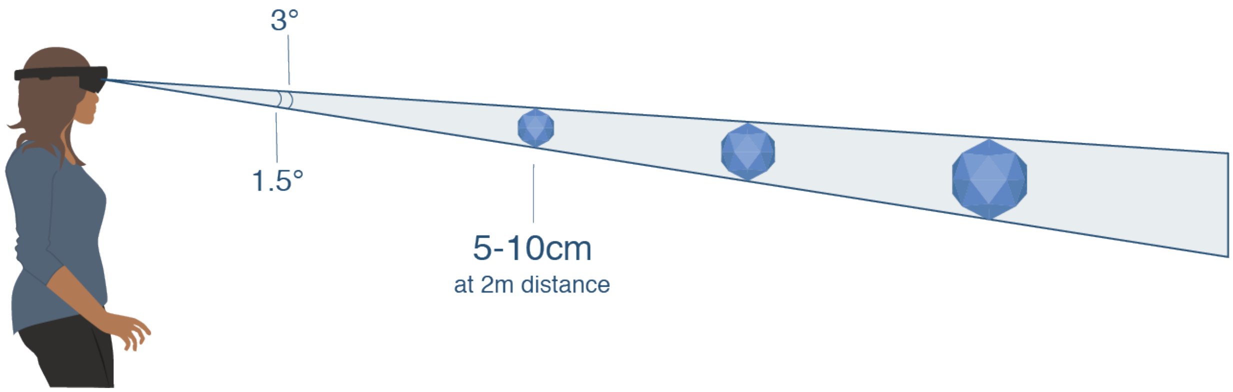

Figure 4. Suggested target sizes for successful gaze targeting in mixed reality (“Gaze Targeting, n.d.).

They also suggest providing feedback when an element is active (“Gaze Targeting”, n.d.) to make sure the system status is visible, a best practice that 2D guidelines support (Nielsen, 1995; Shneiderman, 2005). For instance, Microsoft recommends using a donut shaped cursor to indicate where the user is currently looking and an effect like a “hover” state, audio feedback, or a clear alignment of the cursor and element for objects that can be selected (“Gaze”, n.d.). There are also a number of ways to accept “near misses” as though an object was correctly targeted, like gaze stabilization, closest link algorithms and Focus Stickiness (“Gaze Targeting”, n.d.). Different methods work best for different scenarios.

Gestural commands. Gestures, motions and postures made with the hands and tracked by the HoloLens' gesture-sensing cameras will likely seem the most novel to new users of HWDs. Making gestures to control a computer may bring to mind the famous pop culture reference in the scene from the movie *Mission Impossible* in which Tom Cruise rapidly commands many screens by using his hands. But for most of us, relying on constant, repeated gestural input can result in poor user experience due to user fatigue (“Gesture Design” n.d.; Howard, 2017). Fatigue leads to relaxed gestures which can be unrecognizable to the gesture-sensing cameras, so designers are encouraged to offer a mix of gesture and speech input, incorporate short breaks, and pay attention to the duration of tasks (“Gesture Design” n.d.; LaViola et al., 2017, p.402).

The HoloLens currently recognizes four gestures: bloom, ready, air tap and hold (see Figure 5). Bloom shows the user the HoloLens home, or the menu that allows access to all

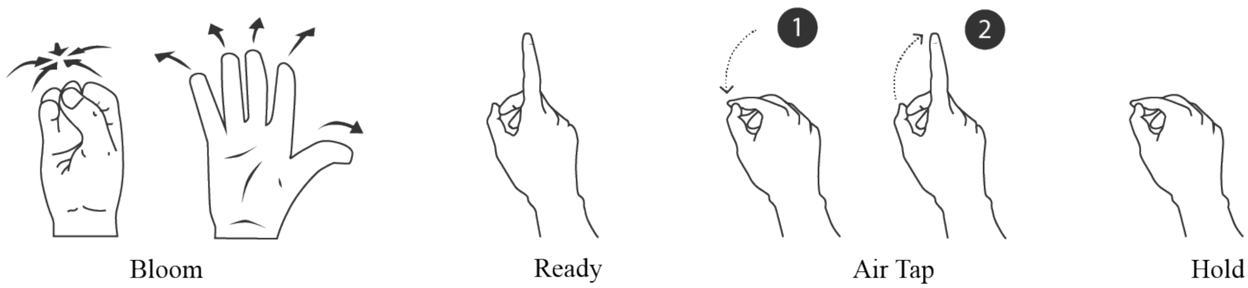

Figure 5. Gestures used to control the Microsoft HoloLens

applications and settings. The Ready gesture lets the HoloLens know that the user is ready to select an object. Air tap is used to select an object, much like clicking a mouse or tapping a touchscreen. Hold, which is performed after tapping an object to select it, allows continuous gestures like dragging, rotating and zooming. User interface designers should be aware of the range that the gesture-sensing cameras track, or the gesture frame, which is roughly from nose to waist and between the shoulders. Both designers and users will benefit from understanding the gesture frame: users will be able to comfortably position their hands in an area that the HoloLens can track, and designers will have the opportunity to notify the user when their hand is almost out of the frame, and make sure there are minimal consequences. For instance, when a user moves out of the gesture frame as they drag a hologram, the hologram should stop at the boundary of the gesture frame instead of returning to its starting position (“Gesture Design” n.d.;).

A well-designed interface will provide feedback to the user so they understand the system's status, a hefty component of 2D design (Nielsen, 1995), and also a subject of 3D design (Rizzotto, 2017). Audio can be a rich source of feedback when a user performs a gesture. John Howard, Co-founder & Creative Director at Loook, Inc., calls this haptic audio. Haptic means touch but Howard uses the term intentionally because in MR and VR an object's audio can offer affordances that affect how a user understands and interacts with it. For instance, when a user performs a continuous gesture like scrolling, she may hear a clicking sound to let her know the scrolling works like a ratchet rather than scrolling smoothly (2017).

Keyboard, mouse and touchscreen. Spatial experiences such as MR in a head-worn display environment can be controlled by the familiar keyboard, mouse and touchscreen. Users can benefit from having access to these input modes in a number of ways. First, mice and keyboards are precise, which is especially helpful when using productivity applications. Second, some users may sooner adopt HWDs if they have access to the tools that currently mesh with their workflow. Third, a user can choose to input data using a keyboard and mouse and then use the HoloLens or other HWD to see the result as a hologram, such as in a product design scenario (Howard, 2017).

Voice commands. Microsoft suggests a see it, say it mentality in which labels on buttons are identical to their associated voice commands, providing an explicit affordance for voice commands. This is in alignment with Nielsen's guideline to rely on recognition rather than recall (1995). If the user sees what they should say, they'll say it without struggling to recall the command. On the other hand, Rizzotto believes that users would be better served by making voice commands more of a conversation. For example, the app should accept both "enlarge the picture" and "make the picture bigger". Microsoft's solution is, as Rizzotto says, like shouting at the “emptiness of space” and he instead suggests a visual point of reference to talk to, like a character or animation that triggers upon sensing speech (Rizzotto, 2017). Different methods of voice command may be appropriate for different applications and users.

Using voice commands can benefit users in a number of ways. They can be more efficient than gestures for some applications (“Voice design”, n.d.). A study by Majaranta shows that gaze control enables the selection of 10-15 words per minute while speech enables 150-250 words per minute (as cited in Hollomon, Kratchounova, Newton, Gildea & Knecht, Appendix B, 2017). Using voice commands can reduce cognitive load because they’re intuitive and easy to learn and remember (“Voice Design”, n.d.), and it can be ideal for those who have physical disabilities or who are using their hands for other tasks. For example, a surgeon in a medical operating room can continue using her hands to operate while issuing voice commands and maintaining a sterile environment (LaViola et al., 2017, p. 398). Like other successful user interfaces, voice command is based on learned behavior from the real world experience of talking to each other, and more recently talking to a wide array of devices that accept voice commands.

When considering input(s) to use for a given application, designers should also be aware of the weaknesses of voice command. It’s not ideal for continuous input such as moving or scaling objects, adjusting volume or any functions that require fine-grained control. Systems can fail to detect voice commands, or interpret them incorrectly, and it may not be acceptable in some public places. When working with the HoloLens designers should take advantage of its ability to display a visual cue when voice input is recognized, which gives the user confidence in using voice as a primary input (“Voice Design”, n.d.). Giving feedback when the system acknowledges a command, in this case a voice command, is a practice that enjoys wide agreement (Rizzotto, 2017; “Voice Design”, n.d.; Nielsen, 1995).

Spatial Mapping: the magic of holograms. One of the most unique features of the HoloLens is its ability to map the user’s environment. This is what makes it possible to place holograms on real surfaces and to occlude them (show them in front of objects located in the background, and hide them behind objects in front of them). In short, this is part of the magic that convinces us that holograms exist in our real world spaces. As a user looks around their environment, the HoloLens scans the surfaces and builds a map of the area. That map persists between application launches and it’s updated when surfaces in the environment change. Spatial mapping can be used to scan parts of rooms, entire rooms, to scan automatically or manually, or not to scan at all depending upon what the application requires. If an application is dependent upon spatial mapping but the environment has not been mapped, designers are encouraged to notify the user to take the appropriate action to map their environment to minimize potential errors (“Spatial Mapping Design”, n.d.).

Spatial Sound. Vision and sound can act in parallel to help us understand our environments (Howard, 2017; Maciel, Carvalho & Melo, p. 518). Spatial sound is technique that can be used to help realistically place an object in the MR environment. The HoloLens’ spatial audio engine attenuates the sound in the MR environment so that a user feels like the sound is coming from a specific place. Designers who typically work in the 2D space may not rely upon sound, but spatial sound in virtual worlds contribute a vast amount of information in a way that the user intuitively understands. First, it audibly grounds objects to make them seem like they really exist, so it supports the feeling of presence of MR, or the feeling of being there. Second, it can act as an audible affordance to help the user find objects outside of their field of view

(Howard, 2017; “Spatial Sound Design”, n.d.). For instance, in the MR game RoboRaid the user often hears the sound of a wall cracking before they see it. The sound alerts them to the location of the crack so the user turn towards it to see it enlarge and turn into a gaping hole from which robots emerge.

Once again, designers will create more successful experiences if they build upon real world knowledge when using spatial sounds. This is in line with Nielsen’s guideline to match between the system and the real world, or to follow real-world conventions in order to make information appear in a logical way (Nielsen, 1995). If a user hears a sound that doesn’t match their experience in the real world it can cause confusion, and generally sound unrealistic to the user (Maciel, et al. p. 518). Using real rather than synthesized sounds and realistically attaching a sound to an object will improve the discoverability of the object in the spatial environment and our instinctual expectation of where a sound should come from will draw our attention more than any spatial cues. For example, the sound of birds chirping will pull a user’s attention upward, even if they hear the sound coming from below them. Making sound emanate from the center of small objects can be satisfactory, but for large objects the sound should come from the part of the object making the noise in order to increase accuracy of the location of the sound (“Spatial Sound Design”, n.d.).

Information Spaces and user interface controls. The HoloLens’ current small field of view (FOV) is a limitation that designers should factor into their interfaces. UC Davis researcher Oliver Kreylos estimates the FOV to be 30° x 17.5° (see figure 6). As a user surveys their surroundings they’ll see the full real world, but the digital overlay of MR is small enough so that the user could easily overlook tools or other objects in an application. Unlike flat screens, there isn’t a static frame in MR and users of HWDs will naturally explore and look in different directions. It can be difficult to know where to anchor information important to the user interface like menus, and informational displays (Howard, 2017).

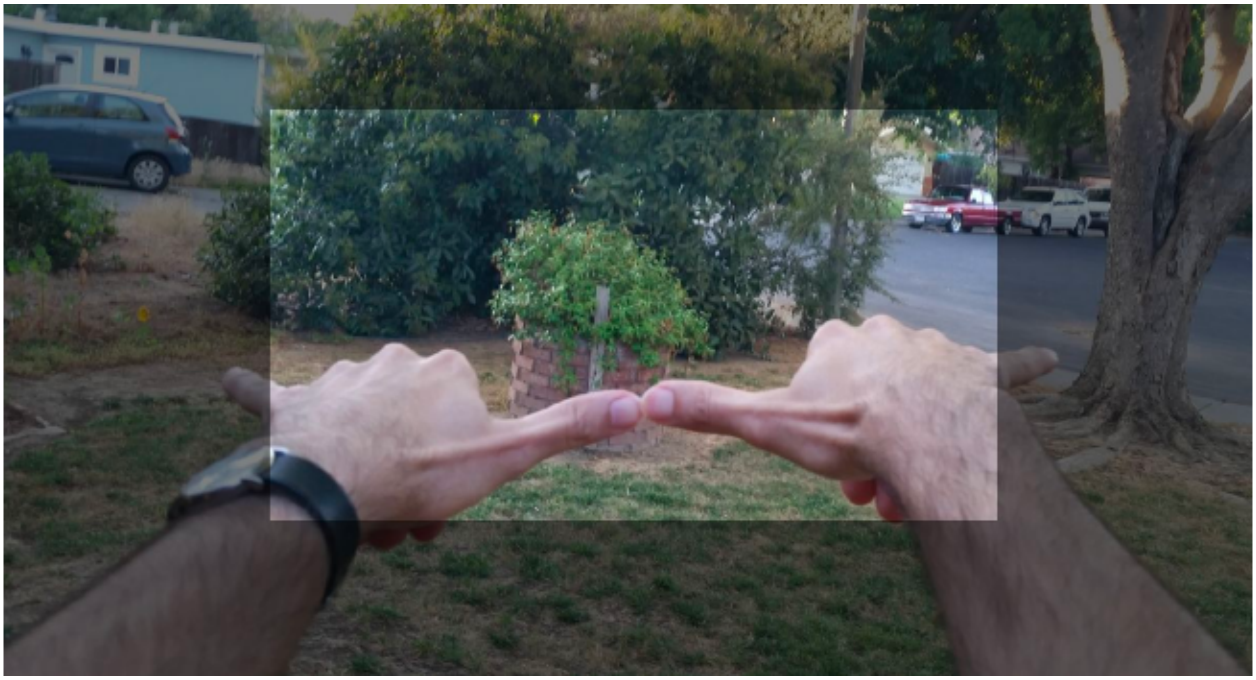

Figure 6. A visual representation of the size of the HoloLens' field of view (Kreylos, 2015).

Organizing information contextually is one strategy when determining where to position user interface elements. Ens, Hincapié-Ramos, & Irani describe this an "information space", an area that maps to a real world space. The information is shown in the context of the surrounding environment (2014). For instance, selecting a character in a gaming application could display an information box positioned above the character's head that contains their health status and game points. Menus can also be tethered to the objects they relate to, but making them large enough so that they extend outside of the FOV will likely make it difficult for users to complete actions, particularly if the menu actions require that the user refers to the original object (Howard, 2017).

The concept of information spaces is an example of the first of three methods used to position content in the HoloLens and other HWDs. Information spaces are considered to be “world-locked”, which means content is anchored to items in the MR environment. This is often the most favored method of positioning content for the HoloLens. The content has no attachment to the user and will be in the same place the next time the application is opened (McDonald, 2016). Another useful way to anchor content is to position an object next to the user, for instance, above their right shoulder. The content will follow them as they move around (McDonald, 2016). The least favored method, and one that many agree should be avoided, is display-locked content, or content that is positionally locked to the display device. McDonald says it results in a feeling of “clingyness” that users want to shake off (2016). LaViola believes it reduces the sense of presence and occludes the environment (2017, p. 388), and Maciel, Carvalho & Melo think it’s distracting, confusing and wastes cognitive resources (2017, p. 514).

Another area of caution when positioning content is user comfort. Viewing stereoscopic 3D content can cause eyestrain, nausea, headaches and motion sickness (“Comfort”, n.d.). One cause of discomfort is vergence-accommodation conflict. Vergence, the movement of our eyes in their sockets to maintain fixation on an object, and accommodation, the change in focal power of the lens within our eyes, are coupled together. The distances that the eyes must converge and accommodate are generally the same, so they work together as a system. Vergence-accommodation conflict can occur when our visual system is unsure what to focus on (“Comfort”, n.d.). Research from Kim, Kane & Banks suggests that the vergence–accommodation system more easily adjusts to conflicts when graphics are introduced slowly so introducing large disparities gradually can help the viewer stay reasonably comfortable (2014, p. 7). Microsoft’s suggested best practice is to place holograms between 1.25 meters and 5 meters from the user. If content must be positioned closer than 1 meter they suggest creating a “depth budget” by estimating how much time a user is expected to view that content and avoid placing them in that situation more than 25% of the time ("Comfort", n.d). Another potential source of discomfort to consider when positioning content is gaze angle. Microsoft's research suggests that designers can mitigate strain from excessive eye and neck movement by positioning content directly in front of our gaze at 0° to 35° degrees below the horizon ("Comfort", n.d.).

Movement. Movement is an integral part of the HWD experience: users are free to move in the physical world, and 3D holograms, information spaces and user interface controls may also move. Virtual reality experiences may call for virtual transportation from one space to another, possibly leading to sickness from a mismatch of visual and vestibular cues of motion, but in mixed reality applications, the user generally moves normally since they can see their real environment. In MR the attention to motion revolves around the digital content overlaying the real world. Microsoft's research gives us a base from which to explore the most effective ways to introduce content. They suggest that designers present content smoothly, avoiding abrupt movements. Users will react to large menus coming towards them so menus should use shorter animations, the animations should be slower, and should be confined to a single plane rather than moving on the Z axis, or through depth ("Interaction Fundamentals", n.d.).

Executive Summary

This usability study was created and conducted in order to assess the usability of FaciemAR, a facial recognition application created for use in the Microsoft HoloLens. By reviewing usability guidelines compiled from years of designing for flat devices, and surveying the burgeoning field of spatial design for mixed reality it provides a foundation from which to assess the application. The primary goal of this study is to find and document as many usability issues as possible so FaciemAR can be modified to give users a better experience. The secondary goal is to contribute to the small but expanding knowledge base of designing for spatial applications.

This study was administered to six undergraduate students in the School of Media and Journalism at UNC Chapel Hill on October 18, 2017. Unlike a usability study conducted on a flat screen in which it can be assumed that participants know how to use the device, no similar assumption could be made for this study. Since HoloLens devices are not available for mass consumption at this time of this writing it's unlikely that any of the participants had used one. To ensure that all participants were able to confidently use the HoloLens the study included a pre-experiment interview to measure the participant's use of technology, including the HoloLens, a 20 minute training session, and post-training questions to measure participant confidence with the HoloLens after training. Next, they were asked to complete four tasks using FaciemAR, followed by four open-ended questions, and a final questionnaire.

The study found that, in addition to a high novelty factor, the participants generally thought FaciemAR would be helpful if they needed to learn more about people in their environment. They were satisfied with the time it took for FaciemAR to recognize a face, and with the information that appeared once the scan was complete. Most felt like the distance at which faces could be recognized would most likely prevent them from using the app: scans from 5 to 10 feet away from the subject were successful, but at 15 feet away from the subject only two of eleven scans were successful. Interestingly, the fact that wearing a HoloLens in public would attract attention was not mentioned as an issue. Most of the interface issues revolved around giving users more instruction and more control of the interface.

Methodology

A usability study was conducted to assess the user experience of FaciemAR, a facial recognition application built for the HoloLens environment (see Appendix C for a representation of FaciemAR's interface). The study was moderated in-person by a single evaluator, and performed on UNC Chapel Hill's campus using a HoloLens provided by the Emerging Technology lab in the School of Media and Journalism. The goal of this study was to capture both qualitative and quantitative data to find and document as many usability issues as possible so FaciemAR can be modified to give users a better experience.

This study employed the formative evaluation method of usability studies in which the evaluator collects qualitative data through observations of users performing pre-determined tasks (LaViola et al., 2017, p. 458). Some quantitative data, like level of HoloLens experience of the user, task timing and errors were also captured by pre-experiment questions and observation. It should be noted that FaciemAR does not use spatial sound so this study did not address that interface component. Observing participants as they attempt a few realistic tasks can uncover the most common issues in an application. This type of study is conducted in order to improve a product based on user feedback and is typically administered multiple times during the development of the product (Sauro, 2015; Krug, 2014, p. 115, Nielsen, 2000).

Opinions vary on the ideal number of participants but since statistical differences aren't relevant for finding interface problems it's broadly agreed that a small group of participants is sufficient to provide actionable insights in the context of a usability study. Norman writes that 5-20 participants should be observed, (2008, p. 108). Krug suggests that three participants is the ideal number for each round of user testing since three people can find more usability issues than can be fixed in a reasonable amount of time (2014, p. 119). Nielsen argues that that 5 participants will uncover 85% of the problems and that iterative testing is more important than finding 100% of the problems during the first test (2000). Since using an application via HWD such as the HoloLens is a different experience than using a flat screen, this usability study requires special consideration. LaViola, Kruijff, McMahan, Bowman & Poupyrev write that when using novel technology results may be highly variable so it may be better to work with a larger group of participants. In addition some participants may develop motion sickness or fatigue during the study which could skew the results (2017, p. 465-466). Given these two potential user issues this study proposed to test seven subjects to increase the likelihood that at least five participants would complete the study.

Krug writes that testing with participants who reflect the appropriate user base is valuable but not necessary. Particularly for a first test, any participants will find usability flaws (2014, p. 120). FaciemAR was conceived as a tool for conference attendees to use in order to discover information such as the names and occupations of other attendees. The demographic consists of adults of all races and genders with enough technological savvy to be interested in the HoloLens, and the wealth to buy it, or the good fortune to have access to one. Given that this is the first usability study of FaciemAR, the demographic for the app is broad, and the study was conducted in the School of Media and Journalism at UNC Chapel Hill, seven students aged 18 and older were recruited from the school's research participation pool. Due to scheduling issues, one subject was not able to participate so six subjects completed the study.

In usability studies that test software on flat-screen devices an assumption is made that the user has some experience using the device. No similar assumption could be made for this study since HoloLens devices have not yet moved into the consumer market. Becoming comfortable wearing the HoloLens headset, looking through the limited field of view and using gestures and gaze to navigate must be learned. Since some students in the School of Media and Journalism have access to the Emerging Technology Lab that houses the HoloLens, this study was designed to accommodate participants with and without experience with the device. In order to test FaciemAR rather than the HoloLens itself all participants would need to achieve a certain level of comfort with the device before performing tasks in the application. As part of this user study participants were guided through a 20 minute training routine to become comfortable with the device.

Sepp writes that human learning stems from our brain's Mirror Neuron System (MNS) which helps us to learn gestures and actions by observing them in other people and suggests that people need demonstrations and training before they're thrown into an unfamiliar reality such as mixed reality. In the Maciel, et al. study of a VR application users told them they'd like a tutorial that explains how to move around and interact with items in the app (2017). This study proposed to lead participants through the following training exercises. First, because we learn by demonstration the evaluator showed each participant a short video of a the HoloLens menu, and how to use gaze and gestures to navigate since these are skills necessary to move to the second step in training. A video was chosen instead of a live demonstration to ensure that each participant saw an identical presentation (see the video at https://vimeo.com/239680286). After the video the evaluator and participant practiced the gestures together, which gave participants another opportunity to learn the gestures by observe someone else. Next, each participant put on the HoloLens and was instructed to complete the *Learn Gestures* app, the Microsoft HoloLens tutorial which gives users practice with HoloLens navigation. Finally, the participants were directed to spend 5 minutes playing the game "Crossy Road". The game's simple premise is to make a chicken character cross a number of busy roads without being hit by oncoming traffic.

This game relies heavily on the air tap gesture, as does FaciemAR, which gave participants additional experience using the device before starting the test. The participants were asked to assess their skill level with the HoloLens, and overall confidence and physical comfort using the device before and after the training. The study found that no students had used the HoloLens so their skill level was only assessed after the training.

FaciemAR has several requirements that affected the location of the study. First, since the application's primary function is to scan the face of another person there had to be a face to scan. In this case, the evaluator acted as the subject to be scanned. Second, the application is meant to scan faces of conference attendees in a variety of spaces so the participants were tasked with scanning the evaluator's face from a number of distances. This means the room had to be at least 15 feet long. Third, the HoloLens must be connected to a wireless Internet connection in order for the app to query its database so the testing location had to have reliable wifi. The test was performed in the Observation Lab in the School of Media and Journalism at UNC Chapel Hill.

From experience testing an AR application, Hunsucker advises that in AR it's not easy to see what the user is seeing so evaluators should plan ahead to record the user view (2017). In addition, it would be difficult for one evaluator to observe and document gestures, movements and voice when performing a usability test in a non-traditional UI (LaViola et al., 2017, p. 465). For this study the HoloLens was configured to use the Windows Device Portal, a companion application that facilitated live streaming of the user's view to a computer that captured the audio and video for review at a later time. A separate audio recording was also made of each session to ensure that every part of the study was captured since the HoloLens wasn't on and streaming during the entire session.

Each participant was guided through a number of activities during the test:

- First, the evaluator greeted the student, explained the testing procedure, and asked the participant to sign a consent form (see Appendix B for all scripts and forms).
- Next, each participant was asked pre-experiment questions to assess their skill level and comfort with technology, and the HoloLens in particular if they'd used it before.
- The participants were led through a training session to make sure they had a minimum level of skill with the HoloLens. First, they watched a short demonstration video of the evaluator using the HoloLens, then they were asked to practice gestures with the evaluator. Next they spent about 5 minutes using the Microsoft *Learn Gestures* tutorial, and after that they practiced what they'd learned by playing a game inside the HoloLens for 5 minutes.
- After training they were asked a few post-training questions to establish their skill level with the HoloLens.
- The evaluator observed the participants as they completed four assigned tasks in the HoloLens environment.
- Next, the participants were asked to answer four open-ended post-experiment questions meant to elicit their final thoughts about FaciemAR and the HoloLens.
- Last, quantitative data was captured by a post-experiment questionnaire covering more specific aspects of the app such as navigation, aesthetic design, and satisfaction and frustration levels regarding different aspects of the app.

Participant Profiles

Six undergraduates from the School of Media and Journalism at UNC Chapel Hill participated in this study. Five were female and one was male and their ages ranged from 19 to 22. None had previously used the HoloLens and the only participant who reported playing video games said they played several times per week. All participants were comfortable with technology, using handheld devices (phones, tablets) an average of 2-5 hours/day and desktop computers (including laptops) an average of 3.5-6.5 hours per day (see figure 7).

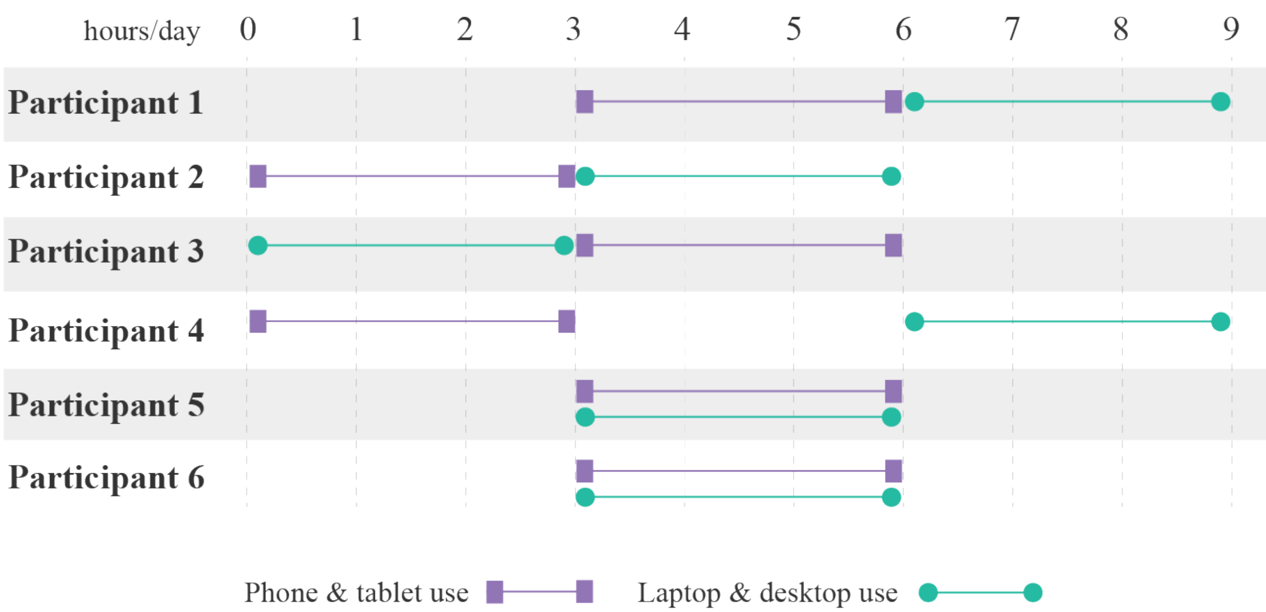

Figure 7. Average self-reported daily use of technology.

Two participants reported attending one conference within their field of study. One was not a native English speaker and one was hoarse from a cold, of note because the *Learn Gestures* application that each participant used during training required them to use speech commands.

Since none of the participants had used the HoloLens before the study, they were all considered to have had no experience before training. After training, students self-reported that their HoloLens skills were approaching expert level, indicating at least 3.5 points on a 5-point scale for each skill that was measured (see figure 8). In addition, on average their reported

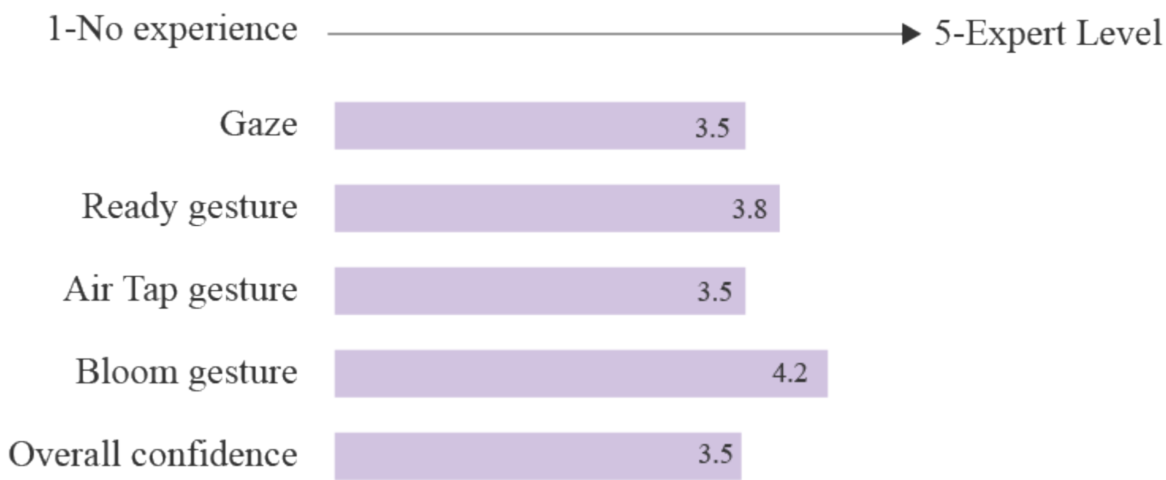

Figure 8. Average self-reported daily use of technology.

physical comfort using the HoloLens was 3.8 on a 5-point scale from uncomfortable to very comfortable. This confidence and comfort using the HoloLens made it possible to rely on their feedback while using FaciemAR. In essence, they were qualified to provide feedback on the application rather than the device.

Limitations

Using the Window's Device Portal, the evaluator was able to stream and record live video from the HoloLens to a laptop. While working through the training portion of the experiments it was quite helpful to be able to use a laptop to see what the participants were seeing through the HoloLens. Even though there is a delay in streaming HoloLens video to a computer, this setup allowed the evaluator to answer questions or give instruction related to the use of the HoloLens in a relatively timely manner during training. Unfortunately FaciemAR does not function when the Window's Device Portal is used to stream live video from the HoloLens so the evaluator could not see what the participants were seeing. It would have been ideal to have had eyes on their progress. Instead, audio recordings and clarifying questions were used to understand each subject's progress. Since FaciemAR is a single-purpose-app it was usually possible to understand a participant's actions based on their feedback. Still, there were several instances in which the evaluator could not extrapolate the participant's actions. Those instances were noted in the data and were not included in the final results.

At the time of this writing there are no known examples of training methods that are empirically proven to prepare participants to test an application in the HoloLens. While five of the participants were up to the challenge of using the HoloLens, one in particular had problems with the air tap gesture. It took her much longer to work through the tasks than the others. It's possible that using a proven training method would have reduced the time that participants took to complete various tasks.

This study's tasks directed participants to scan the same person from 5 feet, 10 feet and 15 feet but practically speaking, there are few reasons that a user would scan the same person multiple times from increasing distances. It would be unnecessary to immediately re-scan a face that had just been successfully identified. In a real life scenario it's more likely that the user would scan the same person from decreasing distances (in the case that the scan failed from far away), or scan different people in succession. The results of this study are tangentially related to scanning different people in succession, but there is opportunity for additional research to test that use case. Similarly, there is opportunity to test the functionality of the app when attempting to scan faces that are looking not directly at the HoloLens.

Participants were asked about their physical comfort as part of the Post Training Interview and Post Experiment Questionnaire. One participant said she felt fine during the interview, "It was like a sensation around my forehead. But I don't think it was super uncomfortable. It was just something I noticed." On the questionnaire the same participant answered 1 on a 5-point scale from 1-feeling physically ill to 5-not feeling physically ill. Since she previously stated that she felt ok it's possible that she answered 1, physically ill, by mistake. If that was the case, the average physical comfort level of 3.8 as reported by the participants should be higher.

Findings

Positive Findings

1. It's helpful. When asked what they liked the best about FaciemAR, three of the participants said they liked being able to see information about other people. One said, "That was pretty cool...I can just click on you and then it can give me information." Two of them further explained that they liked that you could see limited information in the primary information box that appears when you scan someone, and then select that box to see the secondary information box which offers more about them.

2. It's easy. The other three participants focused on how easy it was to use. One said, "It was really easy for me to figure out...I just click on you and then it can give me information, which is neat." Another said, "If I was in the situation where I needed to know information about people it would be super simple."

3. Discovering the secondary information box is simple. Participants found task 2, in which they were instructed to find the subject's employer, the simplest to complete. Participants had to air tap the primary information box which would then present the secondary information box. On average the task took 10 seconds, but four of the six participants were able to complete the task in 10 seconds or under. It was likely so successful because the primary information box turns orange when a user targets it with their gaze. One participant said, "When you enter this area it becomes orange so you know that you can click." The affordance of an item changing appearance when it can be air tapped is very much like the "hover" state that participants know from using 2D devices, so they intuitively knew what to do once the primary information box turned orange even though there were no explicit instructions.

4. Aesthetic design, facial recognition speed, and overall impression received the highest marks. At the end of each session the participant was asked to fill out a questionnaire used to gather quantitative data on a variety of aspects of using FaciemAR, from the strength of the navigational controls to physical illness encountered while using it. The highest average ratings were given to the strength of the aesthetic design, the speed at which the app recognized a face, and the positivity of their overall impression of the app (see figure 9). Half of the participants said they'd recommend the app to a friend.

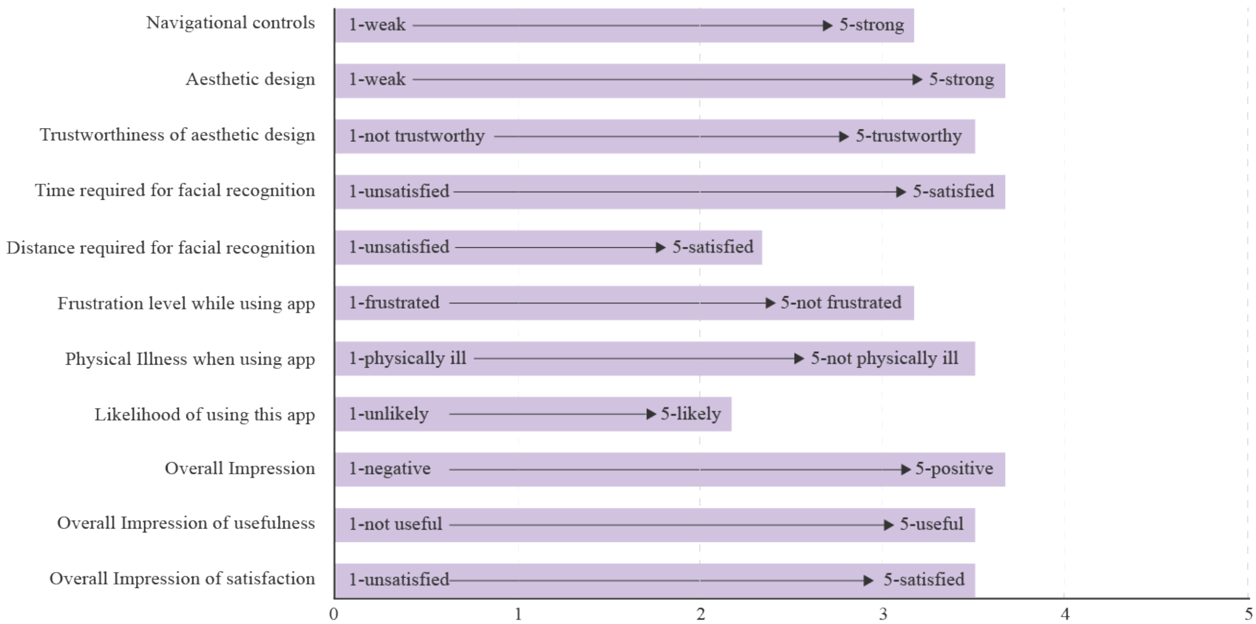

Figure 9. FaciemAR final impressions questionnaire findings

5. Notification while facial recognition is working is a success. All participants understood the message that is shown while the facial recognition is working.

Minor Problems and Recommendations

A minor problem is defined as any negative finding that was observed by less than half of the participants tested. Minor problems do not prevent the participants from completing their tasks.

1. Use spatial sound to guide users. Several participants failed to notice the FaciemAR logo when it appeared, and slowly spun around in place to locate it. Taking advantage of spatial sound by adding an audio cue once the app opens would help users to locate it in their environment.

2. Give users some indication that selecting the primary information box will result in additional information. Task 2, which asked participants to find more information about the subject they'd just scanned, was the easiest task for most. All were successful within a short period of time. However, a few participants asked whether they should air tap the primary information box before selecting it. There was no indicator, other than the change in color upon targeting with gaze, that it could be air tapped. It's recommended that the primary information box include text or an icon that indicates that selecting it will result in more information (see figure 10).

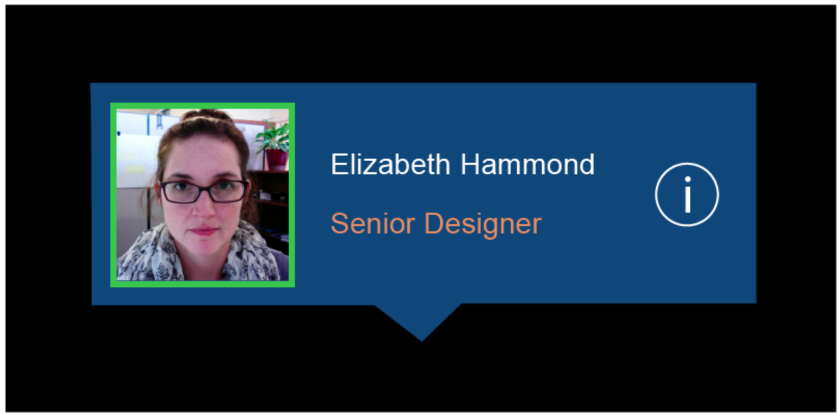

Figure 10. A representation of FaciemAR's primary information box showing an indicator that air tapping will display more details.

3. Support voice commands or let the user know that they're not supported. One participant tried using voice commands in 3 of the 4 tasks, realized it wasn't working and used gestures to complete the tasks instead. The recommendation is that FaciemAR support voice commands. If that's not possible, or until that happens, it's suggested that the app inform the user that voice commands aren't supported. A simple instruction screen (see figure 11) could be presented after the FaciemAR logo and before the facial recognition circles appear. The screen would alert users to air tap to scan, and that speech recognition is not available, which would prevent the user from attempting speech. This follows both Nielsen and

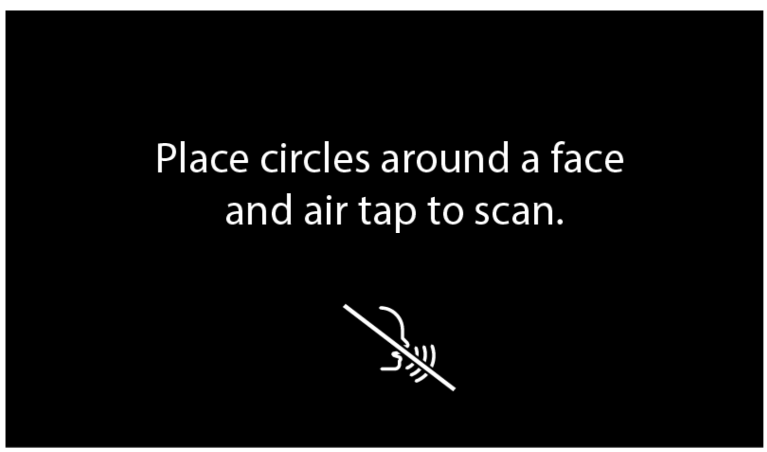

Figure 11. Instruction screen recommended to inform user that there is no speech recognition support.

Shneiderman's guidelines to design interfaces, as much as possible, so that users can't make errors (see appendix A).

4. Make it easier to manipulate the secondary information box from a distance. The primary information box is world-locked, meaning it's anchored to an item within the MR world. In this case it's anchored slightly above the subject who has been recognized. The secondary information box also appears to be world-locked in the foreground of the primary information box, in relation to the user instead of the subject who has been scanned. Being world-locked means that if the user turns their head, they no longer see the information boxes since they're anchored to a particular place in the environment. And if they walk away from the subject who has been scanned, both information boxes will appear smaller, thus more difficult to read, and manipulate using gestures. Two of the participants noted that once they moved away from open information boxes the content in them became more difficult to read. One participant said, "Because it's further away it looks sort of pixelated. The text is a little fuzzy." With some perseverance, both participants were able to target and air tap the remove button on the secondary information box, and air tap again

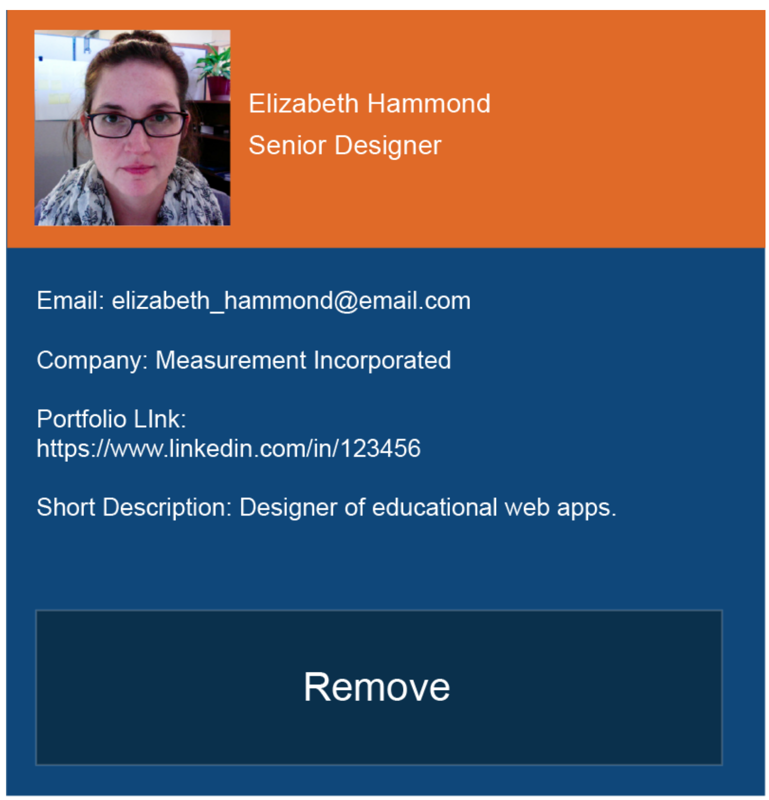

Figure 12. Representation of a larger remove button on the secondary information box

on the primary information box, but it wasn't easy. Both were very pleased that upon re-opening the secondary information box from the farther distance, it was positioned closer to them, which made it easier to read. Speaking of the secondary information box, one said, "when I was at the 5 foot mark I pulled it up and then I walked back to the 10 foot mark with it still up and at that point it was...unreadable. And then I closed it and then reopened it and it appeared closer to me at the 10 foot mark and it was readable."

It's possible that a user would want to remove the secondary information box from a distance farther than the distance at which they had originally scanned the subject. The recommendation is to enlarge the remove button on the secondary information box and enable voice commands so that a user could simply say, "remove" to close the secondary information box (see figure 12). Using the word, "remove" is in line with Microsoft's see it, say it suggestion for voice commands, and also Nielsen's guideline to rely on recognition rather than recall (1995).

Major Problems and Recommendations

A major problem is defined as any negative finding that was observed by more than half of the participants tested. Because of the issue, participants had difficulty completing the task.

1. Instruct the user how to scan a face. Once a face is in the line of sight of the HoloLens vibrating circles appear, but there are no instructions (see Appendix C for a visual representation of FaciemAR's interface). All participants figured out that they were supposed to position the circles over the subject's face and then air tap the face to begin a scan. One did it without questions or hesitation, and a couple only took 10-13 seconds to figure it out. Four of the participants were verbally unsure about what to do. One attempted speech recognition, "Do I say something? Who is this person?" After nothing happened they said, "I guess I'll just tap on you", and successfully scanned the subject. Another asked the evaluator what to do, and once the evaluator asked what she would do if no one was there to help she air tapped the circles. She reported, "Once I knew that I was supposed to tap you, [it wasn't] hard." It's suggested that an instruction screen be added to tell users to position the circles around a face and air tap to scan (see figure 11). Other apps built for the HoloLens do this, RoboRaid in particular gives very brief instructions before the game begins. It should be noted that the circles only appear once a face is in view so that in itself is a strong indication that the app is ready to scan.

2. Give an indication of when circles are properly aligned to a face. Three of the participants reported difficulty putting circles around the subject's face. From 5 feet away participant 1 said, "I'm trying to put it around your face and I'm having a hard time putting it exactly around your face." Participant 6 was not sure when the circles were focused, and when she should air tap them. She said, circles are "Jumping around your face and head. I'm confused [as to] what I should do." At the 10 foot mark Participant 4 said, "Ooh, this is hard...it's so fidgety, so sensitive." A couple of issues were experienced when attempting to target the subject's face. First, Participant 1 noted that the circles were partially in and partially out of frame, so she may have been wearing the HoloLens at an angle that didn't give her access to the full FOV. It can actually be challenging to position the HoloLens in a way that does give the user a full FOV. Designers should consider this to be a common problem and attempt to design an interface to accommodate the issue. The evaluator asked if she wanted to try repositioning the HoloLens. She reported that she could still only see partial circles, but facial recognition did work after that. The second issue is that, the other two participants seemed to be confused about the continuously vibrating circles. They weren't sure when to air tap to start a scan because the circles didn’t give any indication of having successfully targeted a face to scan. Giving users an indication of when the circles have accurately targeted a face would make the user more confident about when to air tap to scan, even if they don’t have access to the full FOV. A significant indication would be for the circles to stop moving when they find a target, but it seems likely that the Face API, which powers the facial recognition process, uses moving circles to target a potentially moving face, indicating that it’s constantly re-targeting. Other indicators like changing the color of the circles, using an audio cue, and changing how the user’s gaze appears could be used.

3. Give the user the ability to start and stop scans. Not knowing how to close the primary information box and/or circles was something that most participants noticed during tasks 2, 3 or 4. Since the study called for them to scan the same person three times in succession, the information box used in one task generally remained open for the second task. Often, once they were ready to scan again they saw either the circles or the primary information box or both. Participants weren’t sure what to do. One said, “How do I get out of the main, like, circle, because you’re still identified.” She closed and reopened Faciem and performed a successful scan. Another participant, as she was trying to begin a scan, said, "There [were] no options, like that were like visible, like oh click here but I just kind of did my hand and clicked on it to cancel it. But now I'm still trying to figure out how to X out of it cause all I see is the option of the circle. I don't know how to get out of that." One participant realized that air tapping anywhere outside of the primary information box removed the box, but other participants thought another scan had occurred, gave up, or closed and re-opened the app. In real world use it’s unlikely that a user would successfully scan someone, but it’s always best to give users control of their application. One of Nielsen's 10 User Interface Design Guidelines is to give users control and freedom (see appendix A), which makes undoing and redoing previous actions possible.

The first recommendation for this issue is to give the user more control over when the circles appear. Currently they seem to appear whenever the app senses a face to scan, but there should be an obvious way to start and stop them. In a real life context it's possible that a user would want to look around for a particular face to scan without seeing the circles whenever they land on a face. When ready they could air tap once to bring up the circles, then air tap again to begin facial recognition. In that case, the instructions that appear after the logo screen might look like figure 13.

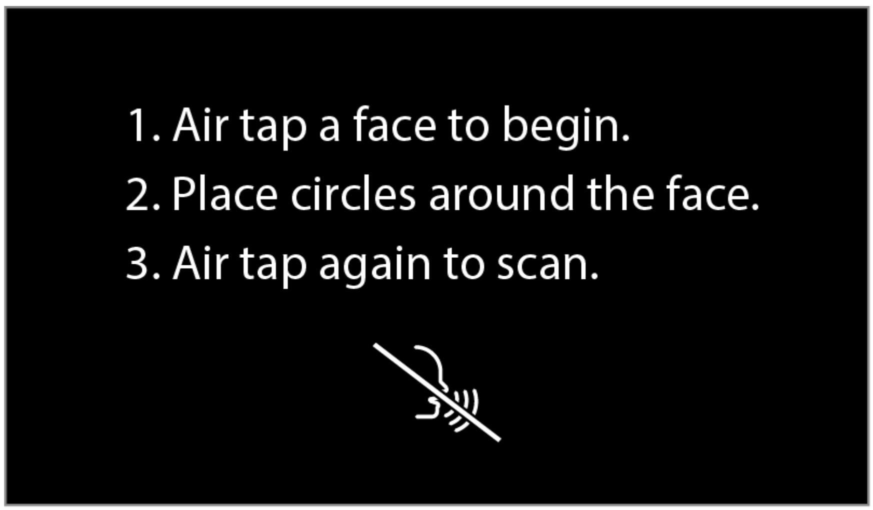

Figure 13. Version two of recommended instruction screen

In case the user changes their mind once the circles are visible, they'll need a way to close out of them. If technically feasible, air tapping outside the circles to close them would provide consistent functionality with the current (though undisclosed) ability to air tap outside the primary information box to remove it. Thus, the second recommendation is to make users aware that they can air tap anywhere outside of the vibrating circles to close them. It would be redundant to alert users to this functionality each time they attempt to scan a face. Shneiderman's second rule is to "Seek universal accessibility" which means offering an experience appropriate for novice users and power users alike. One solution is to display a short message to the user the first time they attempt a scan. It would instruct the user to tap outside the circles to close them. To satisfy novice users and power users it could include two options that could be selected by air tapping or by voice commands. The first option would be something like “Close” and would close the message, which would appear again upon the next scanning attempt. The second could be “Don’t show this message again.”

To give the user control over the display of the primary information box it would be helpful to add a “Remove” tool to the box. A user could air tap it or say “remove” to close it, just like the controls recommended for the secondary information box.

4. Facial recognition is awkward to use in real life from short distances Though this didn’t cause any difficulties when performing tasks, the consensus was that using FaciemAR in the HoloLens to identify people in real life would be awkward if scanning a face was only possible in close proximity. One said, “You had to be a certain distance to do it...so you might as well go up to the person and ask them if you’re going to be that close.” Another remarked after completing the scan from 5 feet away, “It would be really awkward if I was like", and then mimes an air tap.” This opinion was so prevalent that it can be considered a major issue. Participants gave the likeliness of using the app a 2.2, the lowest average score of any other item on the Post Experiment Questionnaire (see figure 8). However, it should be noted that participants had more of an issue with the distance at which the scanning worked than the attention they would attract by wearing the HoloLens in public.

Catastrophic Problems

A catastrophic problem is defined as any negative finding that was observed by more than half of the participants tested. Because of the issue, users could not complete the task.

1. Participants noticed inconsistent recognition from 15 feet away. In total, the participants ran eleven facial recognition scans while standing 15 feet away from the subject. The scans were successful in two of the attempts. The first successful scan occurred after participant 1 had closed and re-opened the app. She scanned four times, and on the fourth try she said the evaluator was recognized. When asked about the difference between this and the other tries she said, "I honestly don't know how that happened. It just popped up." Participant 3 was also able to successfully scan the evaluator. He did not close the application before scanning so there is a possibility that he was seeing the successful scan that was left open from task 3. His second try, after closing and re-opening the app, failed. This difficulty explains why participants were generally unsatisfied with the distance at which facial recognition worked (see figure 9).

This is a prime scenario in which it would have been helpful to see video of what the participants were seeing. Seeing what participants were doing may have helped to clarify whether there was a common error that could be remedied with interface changes, or if scanning at certain distances is a limitation of the Face API. If scanning from this distance is a desired feature, it will require further testing.

Discussion

Voice recognition for controlling the HoloLens didn't work well for a non-native English speaker or for a participant who was hoarse from a cold. Regarding the former, Microsoft's documentation states that speech recognition is, at this time, optimized for native speakers of United States English ("Voice Input", n.d.). At one point the *Learn Gestures* HoloLens tutorial only offered a way to proceed by using speech, which was dicey for these two participants. Both had to try over and over and over again to move forward in the tutorial. In the 2D world

Shneiderman's second rule is to seek universal usability, to recognize the needs of diverse users. This is a valuable reminder that inputs in addition to speech recognition must be offered in order to build inclusive software.

Beyond the realm of this study, but ripe for further research, is what the user of FaciemAR would want to do with the secondary information box once it's visible. Would a user scan one person, read their information, close the information and approach that person to meet them or would they want to "save" the secondary information boxes in order to recall information about subjects when needed? It's possible that they would want to scan various people and organize those secondary information boxes in some way to facilitate contact with multiple people. Understanding the real world use cases would be helpful when creating user controls for the secondary information box in particular.

Conclusion

FaciemAR's purpose is to scan a face, query an opt-in database, and present information about the subject if they're in the database. It's envisioned as an application that could be useful at a conference in which the attendees have been asked to enter their professional information in a database. The participants in this study found it easy and fun to use and thought that being able to quickly access professional information about a person would be helpful. They were satisfied with the speed of the facial recognition.

With a few interface changes FaciemAR could offer users the guidance and control they need in order to use the application more confidently. The addition of spatial sound, enabling voice commands, and adding text instructions would be very helpful in preventing user error. Giving the user control over starting and stopping scans, and alerting them when the scanning circles are aligned properly to a face would help the users to feel much more confident while performing scans. Finally, allowing users to close primary and secondary information boxes from afar will help users to manage their digital environment in a way that makes sense to them.

The biggest roadblock this study found for real-world use of FaciemAR is that facial recognition scans cannot be performed as discreetly as the participants desired. The user would have to be within about 10 feet of the subject in order to successfully recognize the subject. Most participants felt that it would be impractical to scan a person from this distance, thinking that it would be less awkward and more polite to just say hello. One participant said, “I guess if I was going to be using it a conference I'd be a little uncomfortable, like staring at them and clicking at them. Um, but yeah I guess if it was from the farthest mark (15 ft) it wouldn't be that bad…but I can't imagine doing it from like 3 feet away.”

As mentioned in the limitations section, this study did not test the success of scans from any position other than directly in front of a subject. Being able to scan a subject while they aren’t looking directly at the user would be another way to make facial recognition more discreet. A participant noted, “If I was going to figure out who they were I wouldn't want them to be like staring directly at me. I think I'd want them to be like a profile or something. Um, like as discreet as possible."

FaciemAR is one example of a the compelling new world of Mixed Reality, which will call on designers to creatively consider a vast new terrain of design challenges that haven’t existed before. Working in MR, in a space in which the user is the camera, the environment is part of the experience, movement is expected, and gaze and gesture are primary forms of input, will result in creative solutions and applications that have yet to be discovered. The most successful, most useful applications will build upon our shared human experience, the broad, well-established guidelines that our predecessors have contributed and plenty of prototyping and user testing.

References

Dix, A. (2009). Human-Computer Interaction, 1327–1331. https://doi.org/10.1007/978-0-387-39940-9_192

Ens, B., Hincapié-Ramos, J. D., & Irani, P. (2014). Ethereal planes: A Design Framework for 2D Information Spaces in 3D Mixed Reality Environments. *Proceedings of the 2nd ACM symposium on Spatial user interaction - SUI 14*. doi:10.1145/2659766.2659769

Comfort. (n.d.). Retrieved July 11, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/comfort

Gaze. (n.d.). Retrieved July 11, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/gaze

Gaze Targeting. (n.d.). Retrieved July 11, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/gaze_targeting

Gesture Design. (n.d.). Retrieved June 25, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/gesture_design

Gonzalez, J.G. (2016, December 02). The UX Workflow for Hololens & Mixed Reality – Hacker Noon. Retrieved July 06, 2017, from https://hackernoon.com/the-ux-workflow-for-hololens-mixed-reality-3bf59192e577

Gribetz, Meron. "A glimpse of the future through an augmented reality headset." *Meron Gribetz: A glimpse of the future through an augmented reality headset | TED Talk | TED.com*. Ted Talks, Feb. 2016. Web. 04 July 2017.

Hollomon, M., Kratchounova, D., Newton, D., Gildea, K., & Knecht, W. (2017, January). (Tech. No. DOT/FAA/AM - 17/ 4). Retrieved June 25, 2017, from https://www.faa.gov/data_research/research/med_humanfacs/oamtechreports/2010s/media/201704.pdf

Hologram. (n.d.). Retrieved June 25, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/hologram

HoloLens Design Patterns. (2016, May 12). Retrieved July 08, 2017, from http://8ninths.com/hololens-design-patterns/

Howard, J. (2017, July 4). *Spatial UX and UI*. Speech presented at Augmented World Expo in Santa Clara Convention Center, Santa Clara, CA.

Hunsucker, A. J. (2017, February 20). Usability Testing In Augmented Reality – Andrew J Hunsucker – Medium. Retrieved August 12, 2017, from https://medium.com/@AndrewJHCI/usability-testing-in-augmented-reality-df8f6c8d0d71

Interaction Fundamentals. (n.d.). Retrieved June 25, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/interaction_fundamentals

Jacob, R. J. K., Girouard, A., Hirshfield, L. M., Horn, M. S., Shaer, O., Solovey, E. T., & Zigelbaum, J. (2008). Reality-based Interaction: A Framework for post-WIMP Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 201–210). New York, NY, USA: ACM. https://doi.org/10.1145/1357054.1357089

Kaptelinin, V. (2013). Affordances. Retrieved September 21, 2017, from https://www.interaction-design.org/literature/book/the-encyclopedia-of-human-computer-interaction-2nd-ed/affordances#chapter_start

Kim, J., Kane, D., and Banks, M. S. (2014). The rate of change of vergence-accommodation conflict affects visual discomfort, *Vision Research*, 105, 159-165.

Kimani, S. (2009). WIMP Interfaces. In L. LIU & M. T. ÖZSU (Eds.), *Encyclopedia of Database Systems* (pp. 3529–3533). Springer US. https://doi.org/10.1007/978-0-387-39940-9_467

Kreylos, O. (2015, June 18). On the road for VR: Microsoft HoloLens at Build 2015, San Francisco. Retrieved July 13, 2017, from http://doc-ok.org/?p=1223

Krug, S. (2014). *Dont make me think, revisited: a common sense approach to Web usability*. Berkeley, CA: New Riders Pub.

LaViola, J. J., Kruijff, E., McMahan, R. P., Bowman, D. A., & Poupyrev, I. (2017). *3D user interfaces theory and practice*. Boston: Addison-Wesley.

Loup-Escande, E., Frenoy, R., Poplimont, G., Thouvenin, I., Gapenne, O., & Megalakaki, O. (2017). Contributions of mixed reality in a calligraphy learning task: Effects of supplementary visual feedback and expertise on cognitive load, user experience and gestural performance. *Computers in Human Behavior*, *75*, 42–49. https://doi.org/10.1016/j.chb.2017.05.006

Maciel, F., Lourenço, A., Carvalho, P., & Melo, P. (2017). Visual and Interactive Concerns for VR Applications: A Case Study. *Design, User Experience, and Usability: Designing Pleasurable Experiences,DUXU 2017*(Part II), lncs 10289, 510-523. doi:10.1007/978-3-319-58637-3_40

McDonald, C. (2016, November 22). Exploring the HoloLens' Content Lock Options – Mission Data Journal. Retrieved June 25, 2017, from https://journal.missiondata.com/exploring-the-hololens-content-lock-options-6292531cc8d6

Microsoft HoloLens. (n.d.). Retrieved September 21, 2017, from https://www.microsoft.com/en-us/hololens

Milgram, P., & Kishino, F. (1994). A taxonomy of mixed reality visual displays. IEICE Transactions on Information Systems, 77, 1321–1329.

RoboRaid. (n.d.). Retrieved November 11, 2017, from https://www.microsoft.com/en-us/hololens/apps/roboraid

Mixed Reality. (n.d.). Retrieved September 21, 2017, from https://developer.microsoft.com/en-us/windows/mixed-reality/mixed_reality

Nielsen, Jakob (1995, January 1). 10 Usability Heuristics for User Interface Design. Retrieved June 30, 2017, from https://www.nngroup.com/articles/ten-usability-heuristics/

Nielsen, Jakob (2000, March 19). Nielsen Norman Group. Retrieved August 12, 2017, from https://www.nngroup.com/articles/why-you-only-need-to-test-with-5-users/

Norman, Kent L. *Cyberpsychology: an introduction to human-computer interaction*. Cambridge: Cambridge U Press, 2008. Print.

Poushneh, A., & Vasquez-Parraga, A. Z. (2017). Discernible impact of augmented reality on retail customer's experience, satisfaction and willingness to buy. *Journal of Retailing and Consumer Services*, *34*, 229–234. https://doi.org/10.1016/j.jretconser.2016.10.005

Rizzotto, L. (2017, January 16). Best Practices For Mixed Reality Design in 2017 – FuturePi – Medium. Retrieved July 04, 2017, from https://medium.com/futurepi/best-practices-for-mixed-reality-design-in-2017-7dab602574eb

Sauro, Jeff. "5 Types of Usability Tests." *MeasuringU*. N.p., 11 Aug. 2015. Web. 12 Aug. 2017.

Scott R. Klemmer , Björn Hartmann , Leila Takayama, How bodies matter: five themes for interaction design, Proceedings of the 6th conference on Designing Interactive systems, June 26-28, 2006, University Park, PA, USA [doi>10.1145/1142405.1142429]

Sepp, S. (2017, March 27). Toward User Interface and Experience Guidelines for Virtual, Augmented and Mixed Realities. Retrieved July 05, 2017, from https://medium.com/usable-or-not/toward-user-interface-and-experience-guidelines-for-virtual-augmented-and-mixed-realities-939eec8394e4

Shneiderman, B. (2005). Shneiderman's eight golden rules of interface design. *Retrieved July*, *25*, 2009.

Spatial Mapping. (n.d.). Retrieved July 12, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/spatial_mapping_design

Spatial Sound Design. (n.d.). Retrieved July 13, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/spatial_sound_design

Voice Design. (n.d.). Retrieved June 25, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/voice_design

Voice Input. (n.d.). Retrieved October 23, 2017, from Microsoft: https://developer.microsoft.com/en-us/windows/mixed-reality/voice_input

Wagner, I., Broll, W., Jacucci, G., Kuutii, K., McCall, R., Morrison, A., … Terrin, J.-J. (2009). On the Role of Presence in Mixed Reality. *Presence: Teleoperators and Virtual Environments*, *18*(4), 249–276. https://doi.org/10.1162/pres.18.4.249

Wilcox, T., Evans, M., Pearce, C., Pollard, N., & Sundstedt, V. (2008). Gaze and voice based game interaction: the revenge of the killer penguins. Poster presented at ACM SIGGRAPH 2008 (p. 1). New York, NY: ACM Press.

Wong, E. (2017, July 7). Shneiderman's Eight Golden Rules Will Help You Design Better Interfaces. Retrieved July 14, 2017, from https://www.interaction-design.org/literature/article/shneiderman-s-eight-golden-rules-will-help-you-design-better-interfaces

Appendix A

Ben Shneiderman's Eight Golden Rules of Interface Design

1. **Strive for consistency.** Consistent sequences of actions should be required in similar situations; identical terminology should be used in prompts, menus, and help screens; and consistent color, layout, capitalization, fonts, and so on, should be employed throughout. Exceptions, such as required confirmation of the delete command or no echoing of passwords, should be comprehensible and limited in number
2. **Seek universal usability.** Recognize the needs of diverse users and design for plasticity, facilitating transformation of content. Novice to expert differences, age ranges, disabilities, international variations, and technological diversity each enrich the spectrum of requirements that guides design. Adding features for novices, such as explanations, and features for experts, such as shortcuts and faster pacing, enriches the interface design and improves perceived quality.
3. **Offer informative feedback.** For every user action, there should be an interface feedback. For frequent and minor actions, the response can be modest, whereas for infrequent and major actions, the response should be more substantial. Visual presentation of the objects of interest provides a convenient environment for showing changes explicitly (see the discussion of direct manipulation in Chapter 7).
4. **Design dialogs to yield closure.** Sequences of actions should be organized into groups with a beginning, middle, and end. Informative feedback at the completion of a group of actions gives users the satisfaction of accomplishment, a sense of relief, a signal to drop contingency plans from their minds, and an indicator to prepare for the next group of

actions. For example, e-commerce websites move users from selecting products to the checkout, ending with a clear confirmation page that completes the transaction.

5. **Prevent errors.** As much as possible, design the interface so that users cannot make serious errors; for example, gray out menu items that are not appropriate and do not allow alphabetic characters in numeric entry fields (Section 3.3.5). If users make an error, the interface should offer simple, constructive, and specific instructions for recovery. For example, users should not have to retype an entire name-address form if they enter an invalid zip code but rather should be guided to repair only the faulty part. Erroneous actions should leave the interface state unchanged, or the interface should give instructions about restoring the state.
6. **Permit easy reversal of actions.** As much as possible, actions should be reversible. This feature relieves anxiety, since users know that errors can be undone, and encourages exploration of unfamiliar options. The units of reversibility may be a single action, a data-entry task, or a complete group of actions, such as entry of a name-address block.
7. **Keep users in control.** Experienced users strongly desire the sense that they are in charge of the interface and that the interface responds to their actions. They don't want surprises or changes in familiar behavior, and they are annoyed by tedious data-entry sequences, difficulty in obtaining necessary information, and inability to produce their desired result.
8. **Reduce short-term memory load.** Humans' limited capacity for information processing in short-term memory (the rule of thumb is that people can remember "seven plus or minus two chunks" of information) requires that designers avoid interfaces in which users must remember information from one display and then use that information on another

display. It means that cellphones should not require reentry of phone numbers, website locations should remain visible, and lengthy forms should be compacted to fit a single display.

Jakob Nielsen's 10 User Interface Design Guidelines

1. **Visibility of system status.** Users should always be informed of system operations with easy to understand and highly visible status displayed on the screen within a reasonable amount of time.
2. **Match between system and the real world.** Designers should endeavor to mirror the language and concepts users would find in the real world based on who their target users are. Presenting information in logical order and piggybacking on user's expectations derived from their real-world experiences will reduce cognitive strain and make systems easier to use.
3. **User control and freedom.** Offer users a digital space where backward steps are possible, including undoing and redoing previous actions.
4. **Consistency and standards.** Interface designers should ensure that both the graphic elements and terminology are maintained across similar platforms. For example, an icon that represents one category or concept should not represent a different concept when used on a different screen.
5. **Error prevention.** Whenever possible, design systems so that potential errors are kept to a minimum. Users do not like being called upon to detect and remedy problems, which may on occasion be beyond their level of expertise. Eliminating or flagging actions that may result in errors are two possible means of achieving error prevention.

6. **Recognition rather than recall.** Minimize cognitive load by maintaining task-relevant information within the display while users explore the interface. Human attention is limited and we are only capable of maintaining around five items in our short-term memory at one time. Due to the limitations of short-term memory, designers should ensure users can simply employ recognition instead of recalling information across parts of the dialogue. Recognizing something is always easier than recall because recognition involves perceiving cues that help us reach into our vast memory and allowing relevant information to surface. For example, we often find the format of multiple choice questions easier than short answer questions on a test because it only requires us to recognize the answer rather than recall it from our memory.
7. **Flexibility and efficiency of use.** With increased use comes the demand for less interactions that allow faster navigation. This can be achieved by using abbreviations, function keys, hidden commands and macro facilities. Users should be able to customize or tailor the interface to suit their needs so that frequent actions can be achieved through more convenient means.
8. **Aesthetic and minimalist design.** Keep clutter to a minimum. All unnecessary information competes for the user's limited attentional resources, which could inhibit user’s memory retrieval of relevant information. Therefore, the display must be reduced to only the necessary components for the current tasks, whilst providing clearly visible and unambiguous means of navigating to other content.
9. **Help users recognize, diagnose and recover from errors.** Designers should assume users are unable to understand technical terminology, therefore, error messages should almost always be expressed in plain language to ensure nothing gets lost in translation.

10. **Help and documentation.** Ideally, we want users to navigate the system without having to resort to documentation. However, depending on the type of solution, documentation may be necessary. When users require help, ensure it is easily located, specific to the task at hand and worded in a way that will guide them through the necessary steps towards a solution to the issue they are facing.

Appendix B

Usability Script

Hi, I'm Elizabeth, and I'm going to be walking you through this session today.

Before we begin, I have some information for you, and I'm going to read it to make sure that I cover everything.

You probably already have a good idea of why you're here, but let me go over it again briefly. I'm asking people to try using a facial recognition application so I can see whether it works as intended. This is part of my thesis project as a Masters of Technology & Communication student here at UNC. The session should last around an hour.

The first thing I want to make clear right away is that we're testing the *application*, not you. You can't do anything wrong here. Anything you do will simply help us understand how to improve the app.

As you use the app, I'm going to ask you as much as possible to try to think out loud: to say what you're looking at, what you're trying to do, and what you're thinking. This will be a big help to me.

Also, please don't worry that you're going to hurt anyone's feelings by giving your feedback. We're doing this to improve the app, so we need to hear your honest reactions.

If you have any questions as we go along, just ask them. I may not be able to answer them right away, since we're interested in how people do when they don't have someone standing next to them to help. But if you still have any questions when we're done I'll try to answer them then. And if you need to take a break at any point, just let me know.

With your permission, we're going to record audio of our conversation and a video of what you see through the HoloLens. In short, the HoloLens is a mixed reality device that you wear like glasses. Its outward facing cameras will be recording everything you see. You yourself will not be captured on video—only your hands and maybe your arms. The recording will be used to help us figure out how to improve the app.

If you would, I'm going to ask you to sign an **informed consent form (1)**. It will tell you why we're doing the project and ask for your consent to participate. It will also ask for your permission to record audio of what you're saying, and video of what you're seeing through the HoloLens. Do you have any questions so far?

Here's the consent form. Please read it and if you agree, sign on first line on the 2nd page.

Notes: Evaluator hands the participant a the informed consent form and a pen. Upon signature and return of the form, facilitator places the form in the consent form folder and gives the participant an unsigned consent form to keep.

Here's a copy for you to keep. I'm going to start the audio recorder now.

Notes: Evaluator starts the recording. (Press record button twice. Make sure remaining minutes are counting down.)

OK. Before we look at the app, I'd like to ask you a few questions.

Notes: Evaluator pulls a **Pre-Experiment Interview form (2)** out of the participant's folder, asks the questions and records the responses. Once done, places the questionnaire back into the folder.

OK, thanks. Now I'm going to guide you through a process meant to bring all participants in this experiment up to the same skill level with the HoloLens.

First I'm going to show you a short video demonstrating how to use the HoloLens. It's going to show you the main menu, and how to navigate to other applications.

Notes: Evaluator shows the short video to the participant on a laptop.

Let's practice air tap and bloom together a couple of times. You'll need to make sure your gestures can be seen by the HoloLens—try to make the gestures between your shoulders and below eye level.

Now let's take a look at the HoloLens.

Note: Demonstrate while explaining the following…

Point out brightness buttons (left arm), volume buttons (right arm). It's going to enter standby if left alone for 3 minutes. To take wake it up just tap the power button.

To put it on, tilt the band up and loosen the adjustment wheel. Put it on your head and tighten the adjustment wheel until it feels secure. The headband should be positioned right below your hairline.

Note: evaluator demonstrates putting it on.

Next, I'm going to ask you to put this on and go through Microsoft's tutorial on using the HoloLens. You can get to it just like you saw in the video. Just air tap on "Learn Gestures", then air tap the hologram that looks like a hand and the video will start. Take as long as you'd like. I'll be here if you have any questions.

Notes: Evaluator gives the HoloLens to the participant and they put it on. The evaluator may need to help adjust the HoloLens. Evaluator starts capturing video from the Hololens. The tutorial should take about 5 minutes.

Perfect. You can keep the HoloLens on. Next, and last in the training part of this session, you'll play the game Crossy Road for 5 minutes to become more comfortable with gazing and air tapping. It's a simple premise: you're a chicken trying to cross busy roads. You'll gaze and airtap to cross the roads. Again, this is just part of the training, you're not being tested on this.

When you're ready, gaze on the + button and air tap it, then gaze and air tap the Crossy Road icon to get started. I'll let you know 30 seconds before it's time to quit.

Notes: Evaluator sets a timer for 5 minutes, checks again that audio and video are being captured and pulls out the **Post-Training interview (3)** and **Evaluator Observations form (4)**. Once the participant has played the game for 4:30 minutes, evaluator lets them know that they have 30 seconds left.

You have 30 seconds left.

Notes: Once 5 minutes have elapsed the evaluator asks the participant to exit Crossy Roads.

Ok, time's up. Please close Crossy Roads and go back to the menu. I'd like you to take the HoloLens off while I ask you some questions.

Notes: Evaluator puts the HoloLens back on charger, stops the video recording and asks the participant the questions on the **Post-Training Interview (3)** and records the answers.

Thanks for your responses. Now, are you ready to look at FaciemAR? Again, FaciemAR is a facial recognition app meant to be used in the HoloLens. It's envisioned as something you'd wear at a conference in order to identify other conference-goers (but only those who decide to opt in.)

Now I'd like you to stand here at this mark (5ft mark) and I'll stand 5 feet away from you. For this next part I'll just ask you to do a few tasks. While you're doing these tasks it would help me if you'd talk out loud. As much as possible, try to think out loud as you go along. It will help me to understand how to make the app better.

Notes: Evaluator asks Participant to put the HoloLens on

Once you put on the HoloLens, find FaciemAR on the main menu and open it. Let me know once you see something and then I'll ask you to do the first task.

Notes: Evaluator refers to **Evaluator Observations Form (4)** to walk through the tasks.

Thanks, that's the end of what we'll be doing with the HoloLens, you can remove it. (Evaluator puts HoloLens back on the chargers). Now I'd like to ask a few follow-up questions about your experience with FaciemAR and ask you to fill out a short questionnaire.

Notes: Evaluator pulls out the **Post-Experiment Interview (5)** form, asks the questions and takes notes.

The last thing I'll ask you to do is fill out this **Post-Experiment Questionnaire (6)**.

Notes: Evaluator hands the **Post-Experiment Questionnaire (6)** and a pen to the participant.

Thanks again. That was very helpful. you're done with the test. Do you have any questions for me, now that we're done?

Notes: Note their questions and provide answers if possible. Thank them for their participation and escort them out.

Informed Consent Form

FaciemAR User Study Consent Form

You are being asked to take part in a research study to evaluate the design of FaciemAR, an augmented reality (AR) application that runs in the Microsoft HoloLens device. Please read this form carefully and ask any questions you may have before agreeing to take part in the study.

What the study is about: The purpose of this study is to learn how to make the FaciemAR application better.

What I will ask you to do: If you agree to be in this study, I'll ask you some questions about your technology use and whether you've used the HoloLens. I'll ask you to participate in a series of training exercises to learn to use the HoloLens. I'll ask some post training questions about your skill level with the HoloLens, and will then observe you as you complete 4 tasks using FaciemAR. Last, I'll ask you to complete a post-experiment interview and questionnaire about your overall thoughts about FaciemAR. The study will take about an hour to complete.

Risks and benefits: I do not anticipate any risks to you participating in this study other than those encountered in day-to-day life.

There are no benefits to you. Through this study we hope that we'll be able to learn how to improve FaciemAR, and add to the general knowledge of best practices for designing AR applications.

Your answers will be confidential: The records of this study will be kept private. In any sort of report I make public I will not include any information that will make it possible to identify you. Research records will be kept in a locked file; only the researcher will have access to the records.

Taking part is voluntary: Taking part in this study is completely voluntary. You may skip any questions that you do not want to answer. If you decide not to take part or to skip some of the questions, it will not affect your current or future relationship with UNC Chapel Hill. If you decide to take part, you are free to withdraw at any time.

If you have questions: The researcher conducting this study is Elizabeth Hammond. Please ask any questions you have now. If you have questions later, you may contact Elizabeth Hammond at elizabethohammond@gmail.com or at 919-265-4069. If you have any questions or concerns regarding your rights as a subject in this study, you may access the Institutional Review Board (IRB) website at http://research.unc.edu/human-research-ethics/.

You will be given a copy of this form to keep for your records.

Statement of Consent: I understand and consent to the use and release of the audio and video recordings by Elizabeth Hammond. I understand that the information and recordings are for research purposes only and that neither my name nor my image, other than what the HoloLens outward facing cameras capture, will be used.

I have read the information on this form and have had all of my questions answered.

______________________________________ ____________________

Your Signature Date

______________________________________ ____________________

Research Team Member Obtaining Consent (sign) Date

______________________________________ ____________________

Research Team Member Obtaining Consent (print) Date

Pre-Experiment Interview

1. What's your age ______

2. What's the highest level of education you've completed (circle one):

Some college | Bachelor's Degree | Master's Degree | Doctorate

3. On average, how many hours a day do you spend using your phone or tablet as a computer (i.e., for anything but making phone calls) (circle one)

0-3 hours 3-6 hours 6-9 hours 9-15 hours More than 15 hours

4. On average, how many hours a day do you spend using a laptop or desktop computer?
(circle one) 0-3 hours 3-6 hours 6-9 hours 9-15 hours More than 15 hours

5. Do you play video games? (If no, skip to question 7) Yes ____ No ____

6. How frequently do you play video games on any device, i.e., phone, tablet, desktop, game system. (circle one)

Daily | Several times per week | Weekly | Monthly

7. Have you used the HoloLens? (If no, skip to question 12)

Yes ____ No ____

8. How frequently do you use the HoloLens?

Daily | Several times per week | Weekly | Monthly

9. When was the last time you used a HoloLens? (circle one)

Today or yesterday | In the last week | In the last month | Never

10. On a scale of 1-5 (with 1 being no experience and 5 being expert) **please rate your skill with the following aspects of HoloLens use:**

Successfully targeting objects using gaze	1	2	3	4	5
Successfully using the ready gesture	1	2	3	4	5
Successfully using the air tap gesture	1	2	3	4	5
Successfully using the bloom gesture	1	2	3	4	5
Overall confidence using the HoloLens	1	2	3	4	5

11. On a scale of 1-5 (with 1 being uncomfortable and 5 being very comfortable) **rate your physical comfort level while using the HoloLens:** 1 2 3 4 5

If you felt physical discomfort, describe how you felt.

12. What is your field of study or profession?

13. Have you attended a conference related to your field of study or profession? (check one)
Yes ____ No ____

If you answered yes, how many conferences related to your field of study or profession have you attended?

Post-Training Interview

1. What are your overall impressions of using the HoloLens?

2. On a scale of 1-5 (with 1 being no experience and 5 being expert) **please rate your skill with the following aspects of HoloLens use after training:**

Successfully targeting objects using gaze	1	2	3	4	5
Successfully using the ready gesture	1	2	3	4	5
Successfully using the air tap gesture	1	2	3	4	5
Successfully using the bloom gesture	1	2	3	4	5
Overall comfort using the HoloLens	1	2	3	4	5

Post-Training Interview

1. What are your overall impressions of using the HoloLens?

2. On a scale of 1-5 (with 1 being no experience and 5 being expert) **please rate your skill with the following aspects of HoloLens use after training:**

Successfully targeting objects using gaze	1	2	3	4	5
Successfully using the ready gesture	1	2	3	4	5
Successfully using the air tap gesture	1	2	3	4	5
Successfully using the bloom gesture	1	2	3	4	5
Overall confidence using the HoloLens	1	2	3	4	5

3. On a scale of 1-5 (with 1 being uncomfortable and 5 being very comfortable) **rate your physical comfort level while using the HoloLens:** 1 2 3 4 5

If you felt physical discomfort, describe how you felt.

Evaluator Observations Form

Observe Good Quotes:
(Note any good quotes and when they happened during the recording.)

Task Observations

Task 1
Pretend I'm a conference attendee and you'd like to know who I am. Use the app to find out who I am. (From 5 feet away)

Success rate: was the user able to complete the task? (yes/no)

Error rate: how many critical errors were encountered?

Error rate: how many noncritical errors were encountered?

Time to completion (in seconds): how much time did it take the user to complete the task?

Ease question: Ask the user how easy it was to complete this task on a 5-point scale
(1-very difficult to 5-very easy)

Other observations?

Task 2
Use the app to tell me where I work.

Success rate: was the user able to complete the task? (yes/no)

Error rate: how many critical errors were encountered?

Error rate: how many noncritical errors were encountered?

Time to completion (in seconds): how much time did it take the user to complete the task?

Ease question: Ask the user how easy it was to complete this task on a 5-point scale
(1-very difficult to 5-very easy)

Other observations?

Task 3
Note: Evaluator guides participant to the 10 ft mark.

Pretend I'm a conference attendee and you'd like to know who I am. Use the app to find out who I am.

Success rate: was the user able to complete the task? (yes/no)

Error rate: how many critical errors were encountered?

Error rate: how many noncritical errors were encountered?

Time to completion (in seconds): how much time did it take the user to complete the task?

Ease question: Ask the user how easy it was to complete this task on a 5-point scale
(1-very difficult to 5-very easy)

Other observations?

Task 4
Note: Evaluator guides participant to the 15 ft mark.

Pretend I'm a conference attendee and you'd like to know who I am. Use the app to find out who I am.

Success rate: was the user able to complete the task? (yes/no)

Error rate: how many critical errors were encountered?

Error rate: how many noncritical errors were encountered?

Time to completion (in seconds): how much time did it take the user to complete the task?

Ease question: Ask the user how easy it was to complete this task on a 5-point scale
(1-very difficult to 5-very easy)

Other observations?

Post-Experiment Interview

1. What are your overall thoughts about the application?

2. What did you like best about it?

3. What did you like least about it?

4. Would you be interested in using something like the HoloLens in the future?

Post-Experiment Questionnaire

Please rate your overall impression with the following elements of FaciemAR. When answering consider the app rather than the HoloLens. (Circle one: 1-poor, 3-neutral, 5-excellent)

Navigational Controls

Weak	1	2	3	4	5	Strong

Aesthetic Design (colors, fonts, the way the app looks)

Weak	1	2	3	4	5	Strong
Not Trustworthy	1	2	3	4	5	Trustworthy

How satisfied were you with the time it took for the facial recognition to work?

Unsatisfied	1	2	3	4	5	Satisfied

How satisfied were you with the distance at which the facial recognition worked?

Unsatisfied	1	2	3	4	5	Satisfied

How did you feel when using the app?

Frustrated	1	2	3	4	5	Not Frustrated
Physically Ill	1	2	3	4	5	Not Physically Ill

How likely would you be use this app?

Unlikely	1	2	3	4	5	Likely

Overall Impression of the app

Negative	1	2	3	4	5	Positive
Not Useful	1	2	3	4	5	Useful
Unsatisfied	1	2	3	4	5	Satisfied

Would you recommend this app to a friend?

Yes ____ No ____

Appendix C

A visual representation of FaciemAR's interface

1. Users see this "tap to place" window when they open FaciemAR. HoloLens apps often tell users to air tap to place the screen in their environment.

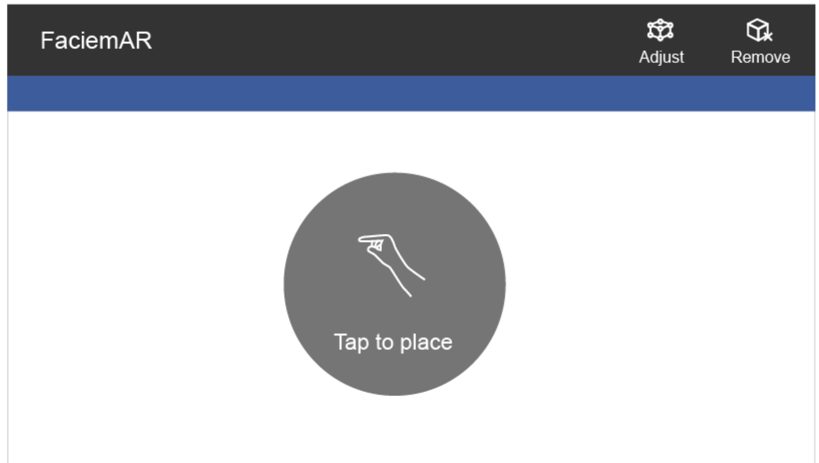

2. The logo appears shortly after the user places the application in their environment.

3. Vibrating white and yellow circles appear when the user points the HoloLens in the direction of a person's face. By placing these around a face, and air tapping them, the facial recognition scan begins.

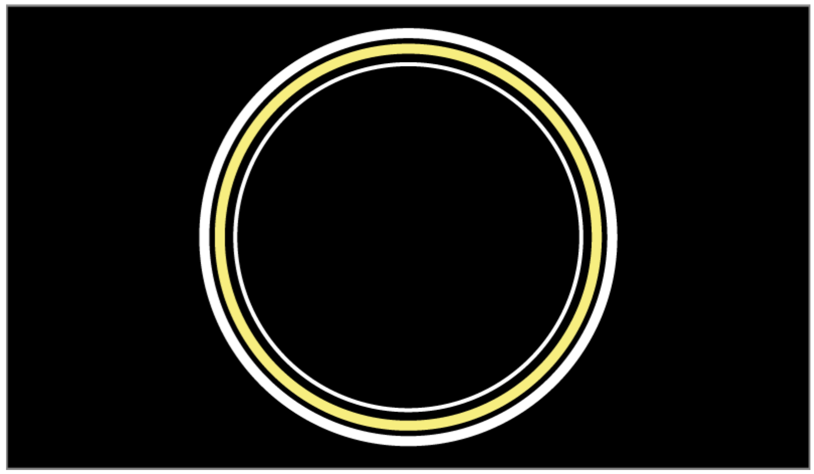

4. This message appears while the facial recognition is working.

5. If the app cannot identify the subject it takes a photo of the subject and presents the message, Not Recognized.

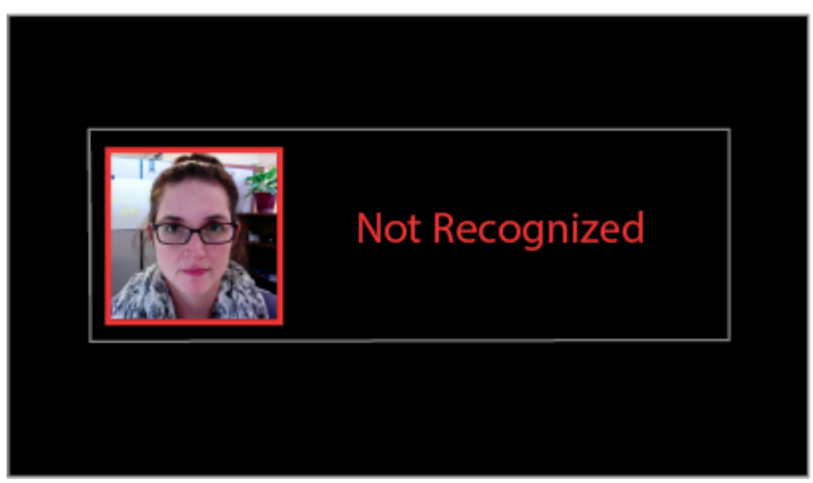

6. If the scan is successful, this primary information box will appear slightly above the subject who was identified.

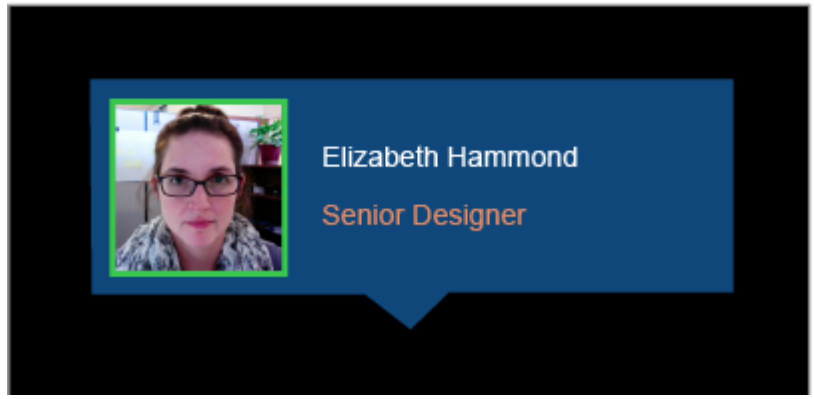

7. When the user gazes on the primary information box it turns orange and the gaze becomes a yellow donut, both indicators that it can be selected. Air tapping anywhere outside of this box will close it.

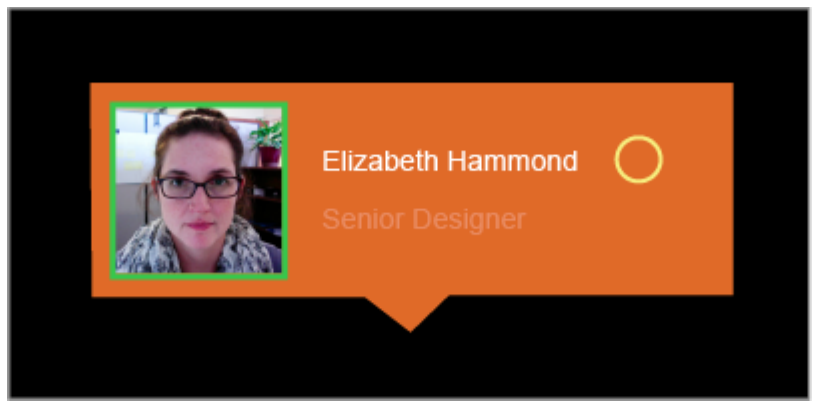

8. When the user air taps the primary information box this secondary information box appears in the foreground, closer to the user. If the user selects the Remove button the primary information box remains visible.