

Holographic Rovers: Augmented Reality and the Microsoft

HoloLens

Laura Toler

Kennedy Space Center

B.S. Visualization

AVR Intern: 2017 Spring Session

04/14/2017

Holographic Rovers: Augmented Reality and the Microsoft

HoloLens

Laura E. Toler

Texas A&M, College Station, TX, 77840

Acronyms

API - Application Programming Interface

AR - Augmented Reality

FOV - Field of View

HMD - Head Mounted Display

HUD - Heads-Up Display

VR - Virtual Reality

I. Introduction

One frequent trope in science fiction is that one day the digital world will be seamlessly integrated with the physical, or ‘real’ world. Today everyone has a small computer in their pocket and Virtual Reality and Augmented Reality are no longer mere pipe dreams. The universal adoption and implementation of VR and AR, however, is a work in progress.

Mixed Reality, as defined by Dr. Paul Milgram in 1994, refers to any location on the ‘virtuality continuum’ where the user perceives both physical reality and digitally simulated reality[6]. Virtual Reality refers to when the user only perceives a simulated world, and exists at the extreme of the ‘virtuality continuum’. This is usually achieved using a Head Mounted Display and earphones, although different research centers are seeking ways to manipulate other senses. Augmented Reality refers to the user perceiving all of the physical world, with some digital additions. These additions can range from flat text in a Heads-Up Display to holograms that seem to actively exist in the environment.

The current AR and VR technology is quite young, so it doesn’t always work as well as they should and many people simply do not know how AR and VR can potentially be used outside of video games. Augmented Reality in particular has so many potential use cases. It’s simply a matter of evaluating the technology available and the needs of the company.

II. Types of AR

A. Head Mounted Displays

Most research and development in Augmented Reality is focused on the use of Head Mounted Displays, or HMDs. There are four major HMDs that have been announced, but only two of them are currently available.

The Microsoft HoloLens is the HMD currently in use at Kennedy Space Center. Using several cameras on the front of the headset, it maps the world around the user and then projects digital images onto a transparent visor so that the images seem to exist in the physical world. The user can walk around and view the holograms from all sides, and can interact with them using two distinct gestures that are recognized by the HoloLens software. The HoloLens is approximately as powerful as a smartphone, but it uses a proprietary chip called the HPU (Holographic Processing Unit) to handle real-time spatial mapping and processing. The largest drawback to the HoloLens available right now is the very limited Field of View, about 30° by 17.5°[5]. This works out to be about the size of a smartphone held 6-8 inches away from the user’s face. Appearance-wise, the HoloLens “looks like a halo and a pair of sunglasses birthed an overweight child”[4]. The current version is a developer version, and the consumer version is rumored to be released in 2019.

The Meta 2 also works by projecting digital images onto a transparent visor. Unlike the Microsoft HoloLens, it is tethered by a 9 foot cord to a computer. This allows for the use of more powerful processors, but it limits the mobility of the user. The Field of View is larger, although exact numbers are hard to come by. The Meta 2 is designed to work with the user’s natural movements, rather than requiring the user to interact with programs using pre-determined, often

non-intuitive gestures. Like the HoloLens, the Meta 2 is also somewhat bulky and awkward-looking.

The Magic Leap is not currently available, so most information about it is from advertising and press releases rather than from actual users. It will purportedly work using ‘light field technology’ to simulate depth. This allows for the user’s eyes to naturally focus on a digital object where it appears to be, rather than focusing on the visor where the images is being displayed. The Magic Leap will not require a physical connection to a computer when the consumer version is released, but current prototypes still seem to require a connection. Like the Meta 2, the Magic Leap will be able to recognize and interpret natural gestures and use them to interact with the programs.

Avegant Light Field was only recently announced, so information is still very sparse. Avegant has already produced Glyph, a HMD for watching movies and TV. Light Field appears to be an AR HMD, similar to the Magic Leap. As the name suggests, it will also use light field technology in its display.

Virtual Reality HMDs, such as the Oculus Rift and the HTC Vive, can be modified to work as AR HMDs by adding a camera component and setting up the program to process and display that input as the world. This is not considered ideal, since that is a lot of processing power that would be allocated just to displaying the room as it is, with fewer resources left to add any digital images to the room.

B. Smartphones

Any smartphone with a camera can be used as an Augmented Reality display device. Pokémon GO is one very famous instance, where players used their phone to track and ‘catch’ fictional creatures from the Pokémon universe. Yelp has also produced an AR app, Monocle, that allows Yelp users to see markers for local businesses through the camera. In both cases, the app uses GPS to place the phone in relation to an internal map of existing data and then uses the camera to create a display where the digital data is overlaid onto the real world.

Most AR programs designed to be used with a smartphone will be apps, packaged applications optimized for a particular phone and operating system. However, Jerome Etienne is working on an Application Programming Interface, or API, that will allow phones to run AR applications from their web browsers. The Javascript API is open source and functional, with updates being pushed to the GitHub page frequently.

C. Projections

Augmented Reality does not automatically require interrupting the user’s vision with some piece of hardware. Lightform is working on Projection Mapping technology, which is projecting images onto the surfaces of objects to create illusions or to change the appearance of the environment. Like an AR HMD, Projection Mapping first scans a room with at least two cameras and then produces an internal map of the space. Based on this internal map, the system then projects images onto the objects in the room to create the desired effects. Brett Jones,

founder of Lightform, described Projection Mapping as “Augmented Reality without the headset” [2]. His stated goal for Lightform, and for projection mapping in general, is to make AR technology available for everyone, rather than AR being restricted to everyone who has a HMD that plays well with the program in question.

III. HoloLens Evaluation

A. Existing Functionality

The HoloLens is designed to be a computer, not to be a gaming console or a television. There are games and movie applications available for the HoloLens, but Microsoft is developing the hardware and software specifically to make the HoloLens a personal computer. Microsoft Outlook Calendar and Outlook Mail are already available within the HoloLens, and preview versions of Word, Excel, Powerpoint, and OneNote are available. Skype for HoloLens is also already available, although it cannot be used for face-to-face telecommunications the way it normally is since there is not a camera pointed at the user. Physical keyboards can be paired with the HoloLens, since it is designed to function as a Windows 10 PC, or the user can simply speak to enter text. Cortana is already integrated into the operating system, and she is fully functional.

HoloLens applications created with Unity3D can integrate a software called Vuforia. Vuforia allows the application to recognize and respond to 2D images such as QR codes and signs. Microsoft has not released a proprietary image recognition software for the HoloLens yet, nor have they endorsed any 3D recognition software. The current spatial mapping software embedded in the HoloLens is too imprecise to be used to recognize specific physical objects. However, the general attitude in the industry is that AR HMDs should be able to recognize objects in the environment and react to them so this may change as the consumer release of the HoloLens approaches.

B. HoloLens Academy by Microsoft

Microsoft has only released a developer version of the HoloLens, and they are encouraging developers to create software and applications for the HoloLens before the consumer release. As part of that goal, Microsoft is offering support for those developing in the Unity Game Engine for the Microsoft HoloLens. This support include two series of text tutorials, a series of video tutorials synced up with the text tutorials, documentation, a forum for HoloLens developers, and Microsoft staff who go into the forums to help answer questions.

The HoloLens Academy tutorials come in two series, the 101 and the 200s. The two 101 tutorials touch on 7 topics, one teaching how to develop with the use of a HoloLens, and the other focusing on the use of a HoloLens Emulator. The 200 series has 6 separate tutorials, each focusing on a separate aspect of Augmented Reality development, including gaze, gesture, and spatial mapping.

While these tutorials are an excellent place to start learning about the concepts behind AR applications, the tutorials overwhelmingly use pre-made game objects, known as Prefabs, that

the user downloads and uses. While this helps “to minimize changes to the course instructions and keep the tutorial in-sync with the video recordings”[1], it can hinder the user from understanding how to construct applications from scratch. If the user is already experienced in developing games or applications with Unity3D, then the HoloLens Academy is an excellent starting place.

C. Use Cases

The primary use case for the HoloLens as designed by Microsoft is simply to replace a standard PC with a HoloLens. It allows for more screens and a different dynamic with the office. HoloLenses can be synced up to share holograms or to share a single hologram while each individual is seeing different relevant data on their screen. Personal spaces such as offices can be customized, with screens and holograms placed wherever they work best for that particular user.

Another use case for the HoloLens, or for other AR HMDs, is to augment skilled workers and technicians in the field. Instructions and data can be overlaid on the screen, and if the system has good object recognition software then the headset can recognize what needs to be done and assist the user. This is especially helpful when the user is working in an area with too many systems for them to effectively learn, such as with astronauts. Rather than memorizing a series of pictures and instructions, the headset can show how things need to be manipulated to achieve the desired goal.

D. Health Risks

Given how new Augmented Reality technology is, there are no known long-term health risks specifically associated with it. However, some risks are immediately apparent when using the HoloLens. There is a risk of eyestrain due to the eyes having to frequently focus on a plane within an inch of the eyes for extended periods of time. While this is also a risk with any kind of digital display, the location of the HoloLens display does raise some concerns. Lightfield technology, such as that being developed by Magic Leap and Avegant, should theoretically alleviate these issues. Another risk with the HoloLens is neck strain. The headset is amazingly lightweight given how powerful it is, but if the user doesn't get the head straps adjusted perfectly after putting it on, the weight can cause neck strain after extended use.

IV. Application - Curiosity Rover

A. Design

One of the primary goals of the project was to assess the usability of the Microsoft HoloLens by producing a functioning application within the time allotted by the internship. It was decided that, while it was unlikely that the app would be used, an outreach-style application

designed to educate the user about the Curiosity Rover would work very well for testing the HoloLens.

The application displays a hologram of the Curiosity Rover, scaled down so that it fits comfortably in an office. The user is able to select various parts of the rover to trigger different functions.

- Main Chassis: Selecting the main chassis allows the user to move the rover around in the room. Since the application uses spatial mapping, the user can set the rover on any flat surface in the room.
- ChemCam: Selecting on the ChemCam triggers an animation of a laser being shot at a rock. A large text panel describing the function of ChemCam also appears next to the rover, and will rotate to face the user at all times so that it is easier to read.
- MAHLI: Selecting on the 'hand' of the Curiosity Rover triggers an animation of the arm extending the 'hand' out. 5 images taken by MAHLI pop up behind the rover, and a text panel like the one for ChemCam appears next to the rover.
- Wheel: Selecting the forward right wheel of the Curiosity Rover triggers an animation of the wheels rolling and moving over an imaginary hill. A text panel like the one for ChemCam appears next to the rover.

B. Resources and Assets

Given that the primary focus of the project was on the application itself, it made more sense to use some publically available assets as well as creating some assets, rather than creating everything completely from scratch.

- Models: The initial models for the Curiosity Rover were sourced from NASA 3D Resources, a publicly available collection of 3D models of various types hosted by Ames Research Center. The models were subsequently modified and optimized for use within a digital environment, as the original models were designed to be 3D printed. The models were modified using Blender, an open-source 3D modelling and animation software.
- UV Textures: Each UV texture was created from scratch using reference images of the relevant part of the rover. GIMP (GNU Image Manipulating Program) was used to create the textures based on the UV maps generated by Blender.
- Scripts: Some of the scripts are modified versions of the tutorial scripts from the HoloLens Academy. These are the scripts that handle the functions of the HoloLens rather than the unique functions of the application. Most of the scripts, however, were created specifically for this application.
- Animations: Unity3D has a built-in animation system designed for creating things like cutscenes and animation cycles for video games. All of the animations were created for this project within Unity3D.

C. Issues

Some of the issues encountered over the course of the project were entirely unrelated to the HoloLens. Unity3D's Animator does not work as well as it could, and the animations did not always play well with the scripts. This could be due to the scripts themselves or due to the particular version of Unity3D that was used. Other issues were specific to creating an application for the HoloLens, such as having to deploy the code twice to send it to the HoloLens, but this will likely be fixed as new patches are pushed for the Unity HoloLens Developer SDK. The primary issues found with the HoloLens were text displays and frame of reference.

The creation of holographic text that is integrated into the environment requires new techniques to display text, since free-floating text is often quite jarring. The text for this project was attached to a cuboid which functioned as a display board for the text. The text and boards also had an attached script that constantly rotated the text to face the user. This enabled the user to be able to read the text from any angle.

An early version of the project had to be scrapped because the system had switched to using an attached frame of reference instead of a static frame of reference. This manifested as the hologram constantly following the user's gaze and location rather than staying in one place, and persisted after the implementation of spatial mapping. Microsoft only has one piece of documentation on the topic, and nothing about how to switch from one frame of reference to the other. The attached frame of reference could theoretically be used for HUDs, but only if it can be deliberately controlled. Right now, there does not appear to be a way to control which frame of reference is used, and the system seems to typically default to a static frame of reference.

V. Conclusions

Augmented Reality, especially AR for industry, is worth investing in and continuing to research. From aiding in construction and maintenance to encouraging interest from younger generations, AR can greatly benefit NASA. Non-enterprise applications for AR don't necessarily need to be designed for Head Mounted Displays. The average American adult has a smartphone that is more than capable of running a simple AR game or digital tour. HMDS, like the Microsoft HoloLens, should be reserved for cases such as astronauts and field technicians who need instructions as they work or for internal proposals and assessments such as a digital demo of a proposed rocket design.

Of the AR HMDs currently available, the Microsoft HoloLens is the most functional. However, it is still a developer version, with the consumer version not likely to be released until 2019. The current version functions very well when used for 3D demos, and the consumer version will work as a personal computer. There could be some issues regarding the use in a cubicle environment, given the need for space for gestures and the need for voice controls. If the object recognition and spatial mapping software are improved, then it will be an incredible asset in any skilled laborer or technician's hands. There would need to be less training on specific types of hardware and there would be fewer mistakes made out in the field.

¹Ahilier, Forum Response, Windows Holographic Developer Forum,

URL:<https://forums.hololens.com/discussion/1912/spatial-mapping-from-scratch>

²Bird, G., “Avegant Introduces Light Field Technology for Mixed Reality”, *PR Newswire*,

URL:http://www.prnewswire.com/news-releases/avegant-introduces-light-field-technology-for-mixed-reality-300423855.html?tc=eml_cleartime

³Curiosity Rover, 3D Model, Ver 1.0, NASA Jet Propulsion Laboratory,

URL:<https://nasa3d.arc.nasa.gov/>

⁴Dachis, A., “What's the Difference Between HoloLens, Meta & Magic Leap?”, *mixed.reality*

URL: <https://mixed.reality.news/news/whats-difference-between-hololens-meta-magic-leap-0171361/>

⁵Kreylos, O., “HoloLens and Field of View in Augmented Reality”, *Doc-Ok*, URL:<http://doc-ok.org/?p=1274>

⁶Milgram, P. and Kishino, F., “A TAXONOMY OF MIXED REALITY VISUAL DISPLAYS”,
IEICE Transactions on Information Systems, Vol E77-D, No.12, December 1994,
URL:http://etclab.mie.utoronto.ca/people/paul_dir/IEICE94/ieice.html