# MULTIMODAL VISUALIZATION USING AUGMENTED REALITY AND LARGE DISPLAYS

by

Erik Butler

A thesis submitted to the faculty of
The University of North Carolina at Charlotte
in partial fulfillment of the requirements
for the degree of Master of Science in
Computer Science

Charlotte

2017

Approved by:

______________________________

Dr. Aidong Lu

______________________________

Dr. Samira Shaikh

______________________________

Dr. Weichao Wang

ABSTRACT

ERIK BUTLER. MultiModal Visualization using Augmented Reality and Large Displays. (Under the direction of DR. AIDONG LU)

We are living in an age where technology has allowed data collection to encompass every aspect of our day to day lives. With quintillions of bits of data being created daily, we are under constant pressure to advance our supercomputers and machine learning techniques, as we search for answers to questions that we may not have even known to ask yet.

However, simply connecting the hidden dots within the data is not enough. We need to be able to represent the connections we find in ways that allow us to derive meaning from them. As we push the boundaries of data processing techniques, we should advance our visual analytic techniques as well. Breakthrough advancements in augmented reality (AR) technology have provided us the opportunity to do just that.

In this thesis, I have developed a platform that uses augmented reality to represent data in a way which is more intuitive, and bound by less restrictions. The platform incorporates a multimodal approach, utilizing augmented reality devices (Microsoft's HoloLens in our case) and several large displays. The platform uses AR technology to display interactive data visualizations corresponding to information on the displays, with the perception that these visualizations exist in the real world. With the world as your rendering space, you gain tremendous scale-ability, dimensionality, and the ability to incorporate any real world object into the scope of your visualization.

# ACKNOWLEDGMENTS

I would like to thank my advisor, Dr. Aidong Lu, and my committee members Dr. Samira Shaikh and Dr. Weichao Wang, for giving me such a great opportunity to work on this project. I would also like to thank Zhongyu Li, Dr. Shuiwang Ji, Dr. Shaoting Zhang, and Dr. Jian Huang, for the opportunity to take part in their exciting research on neuron morphology. Lastly, I would like to thank Boyu Fan for developing the D3 applications used for the GSMNP project.

This thesis is based upon work supported by UNC Charlotte and the National Science Foundation under Grant No. 1629913.

# TABLE OF CONTENTS

# LIST OF FIGURES

# CHAPTER 1: INTRODUCTION

For decades now, the scientific community and film industry have imagined a future where virtual images can be displayed in the real world as if they were just another object among us. We have in fact made strides towards reaching this future. We can use light fields to create real world virtual images, a process known as holography. However, creating these images is a very complex process that must take place in a designed environment. The imagined future where we raise our watch to accept a call, and a 3D image of a person pops out of the watch face before our eyes, still feels a long way off. However, technology moves fast, and today's tech has created the opportunity to trick our minds into believing that this future has arrived.

With the advancement of processing power, battery capabilities, sensory technology, display techniques, and of course algorithms; we have seen the rise of augmented and mixed reality devices (AR and MR respectively). These wearable or portable devices either overlay digital images on the real world (in the case of AR), or create the perception that real world and virtual objects coexist and interact in real time (in the case of MR). At least, that is the case which Microsoft is making to differentiate their "MR" technology. We will use the more general term, augmented reality, when referring to mixed reality throughout the thesis.

Initially, developers use this young technology to create applications such as games and interactive educational applications. However, as AR technology has started to mature, there has been greater interest in exploring its use for analytic purposes. This thesis describes a system that aims to enhance data analytic capabilities by introducing a more granular and intuitive relationship with data visualizations, in addition to opening the doors to a multimodal approach, which broadens the scope of our interaction with data.

## 1.1 Technology Background

In 1968, Ivan Sutherland created a head-mounted display named “The Sword of Damocles”. It was the first VR headset, and the first attempt to blend the virtual world with our own. The technology was primitive, doing little more than moving the display closer to the users face, and tracking user movement as a source of input. The technology was not very effective at replicating a 3D environment, which has been the case for decades.

Fast forward to around 2010. Rumors of the mysterious massively funded start-up company, Magic Leap, gave us the first sign that technology was ready to create what Ivan and many others had long dreamed of. Two years later, another start-up announced its plans for a VR headset called the Oculus Rift, and an arms race began. Everyone from tech giants such as Google and Samsung, to smaller start-ups, seemed to be working on some form of virtual, augmented, or mixed reality. Microsoft announced its arrival to the scene in 2015 with the Windows Mixed Reality platform, along with a MR headset named the HoloLens. Apple finally joined the scene in 2017 with their announcement of an augmented reality toolkit to be used for development on their mobile devices. Despite arriving late, Apple leapfrogged many of its competitors in many regards, showing an evident dedication to the new technology. The arrival of the last tech giant, along with the manner in which they showed up, shows us that the industry believes that this technology will play a major role in our future.

## 1.2 Scope and Contribution

This thesis presents early results of multimodal visualization using augmented reality and large displays with the latest HoloLens. We provide results using two important applications, medical imaging and biodiversity. The case studies demonstrate several prototype designs and reveal the advantages over standard desktop visualization.

# CHAPTER 2: RELATED WORK

## 2.1 Multimodal AR Technologies

Two decades ago, in 1997, Ronald T. Azuma released a paper surveying the field of augmented reality [2, 3]. The paper went into an in-depth background discussion of the concepts and technology behind augmented reality, in addition to the many hurdles AR faced during that time. The paper went on to discuss the limited applications of AR during that era, most notably the use of Head-Up Displays (HUDs) by military pilots, which superimposed flight, navigation, and targeting information on the pilots helmet mounted sights. Aside from providing rich background on the field of AR, this paper does an excellent job pointing out that AR has had a number of practical applications for decades now, with many more waiting on technology to untether AR equipment from larger computers.

Another recent study, *A Survey of Augmented Reality*, lead by Mark Billinghurst in 2014 [10], goes into in-depth analysis on the field of AR. After discussing the various technologies related to AR, the study points out the significant progress AR is making today, including the use of AR in various fields of work. One such example is Architecture, where the ability to superimpose virtual information over the real world poses a unique benefit for an industry interested in viewing structures which have yet to be built. Another example is the field of Marketing, which focuses on capturing peoples attention, and motivating them to learn more about a product. The study notes that AR can be used to create more memorable experiences compared to traditional marketing techniques. Lastly, the study discusses the use of AR in education, where in some situations AR has been shown to help students learn more effectively and have increased knowledge retention. One specific case discussed uses AR to overlay interactive 3D digital content on the pages of real books.

To provide a specific example, Fiorentino et al. presented a multimodal method for the augmented visualization on a large screen and a combination of multiple fixed and mobile cameras [7]. An empirical study was presented to evaluate the effectiveness of technical maintenance assisted with interactive augmented reality instructions. The approach combines a large screen and multiple fixed and mobile cameras. A set of maintenance tasks based on manual inspections of a motorbike engine, including tool selection, removal of bolts, and part dis-assembly, are supported by visual labels, 3D virtual models and 3D animations. The statistical analyses of results using paper manuals and augmented instructions proved that augmented instructions reduced significantly participants overall execution time and error rate.

Skipping forward to more current research, a paper published in 2016 by Hanna Schraffenberger [11] suggests that multimodal augmented reality is “the norm rather than the exception.” The paper explores an approach that is centered around the participants perspective in AR environments. In doing so, they find that AR is about much more than simply what the participant can see, which they summarize with three points. First, virtual content can exist in non-visual forms. Second, AR experiences include our experiences of the real world, which is multimodal itself. Third, the illusion of virtual and real world objects being part of the space may be influenced by how they interact with each other. The paper goes on to suggest that realizing such multimodal influences may open the door to a number of possibilities.

## 2.2 Multimodel AR Applications

Multimodel AR approaches have been applied to a number of applications. The following provides several examples to demonstrate that these technologies have been utilized in different fields at different levels.

- An ultrasound guided surgical microscope is developed to take the advantage that surgical microscopy and ultrasound systems are already used in neurosurgery, so it does not add more complexity to the surgical procedure. [8].
- An augmented reality system was developed for knowledge-intensive location-based expert work [12]. The multimodal interaction system combines multiple on-body input and output devices: a speech-based dialogue system, a head-mounted augmented reality display (HMD), and a head-mounted eyetracker.
- Multimodal interaction provides the user with multiple modes of interaction for navigation display [6].
- Multimodal head-mounted display for multimodal alarms in intensive care units [4].

## 2.3 Immersive Analytics

A number of recent studies on immersive analytics have been performed and they provide favorable results for stereoscopic techniques. For example, Alper et al. [1] presented stereoscopic highlighting to help answer accessibility and adjacency queries when interacting with a node-link diagram, and the evaluation results showed that stereoscopic highlighting could significantly enhance graph visualizations for certain use cases. Ware and Mitchell [13] studied the perception of variations of 3D node-link diagrams and showed that stereoscopy reduces errors and response time in a very high resolution stereoscopic display for both skilled and unskilled observers.

The effectiveness of immersive analytics needs to be explored The differences are dependent on the approaches as well as the applications. For example, studies of performances on collaborative immersive visualization using the recent HMDs, such as Oculus Rift and HTC Vive, have shown no difference with expensive equipment such as cave-style environments [5]. Differently, Kwon et al. [9] investigated the effectiveness of graph visualization and the impact of different layout techniques on readability in an HMD, and they concluded that the 3D stereoscopic graph visualization with an Oculus Rift out-performed traditional 2D graph visualizations.

# CHAPTER 3: MULTIMODEL VISUALIZATION USING UNITY

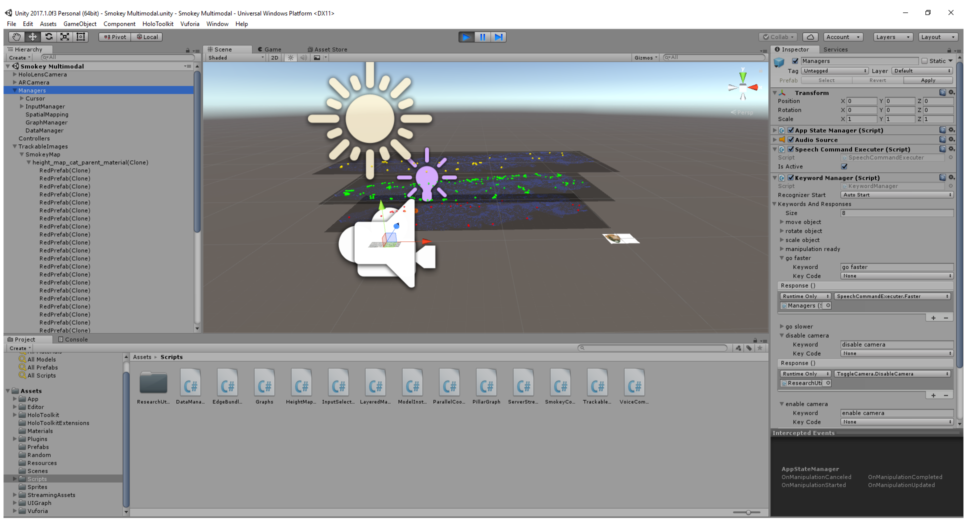

Figure 1: Augmented reality application development using Unity.

With processing power of a Surface Pro 2 tablet packed into an un-tethered headset, Microsoft's HoloLens is one of the first augmented reality devices to hit the market with truly marketable capabilities. The HoloLens uses an array of sensors and cameras to perform spatial mapping and gaze tracking, building a virtual environment which corresponds to the real world around the user. Stereo imaging, gesture recognition, and spatial sound are then used to trick the users' senses into believing that virtual objects coexist with the real world. This, in a general sense, is how Microsoft's HoloLens accomplishes the task of augmenting the users reality in such convincing form.

The cross platform game engine *Unity* has become extremely popular for the development of 2D and 3D applications across an array of devices. Knowing that they needed powerful software development tools to help facilitate the development of interactive 3D holographic applications, Microsoft provided deep integration between the HoloLens and the Unity game engine. Thanks to this integration, developers can build AR applications centered around a series of coordinate systems which lie at the heart of the HoloLens's success.

The default coordinate system, or the *stationary frame of reference*, is typically established when the application starts, and is sustained through-out the course of the application. Once created, the coordinate system attempts to establish a permanent frame of reference in the real world. The idea is to center virtual objects around this coordinate system instead of the user so that they remain stationary as the device moves around. Additional coordinate systems can be established based on an applications needs, however they may suffer from drift and inaccuracies.

To address the drift issues, HoloLens utilizes another set of coordinates called *spatial anchors*. Spatial anchors are data structures which contain a world coordinate and information about the surrounding environment. This allows the HoloLens to accurately locate the coordinates corresponding location in the real world. Spatial anchors can also be serialized and shared between devices, allowing multiple users to share visualizations and collaborate on the interactions.

For our two applications, we utilize the stationary frame of reference in slightly different ways, but overall, our approach is quite similar.

The majority of our visualizations are based on data points provided by our collaborators. These data points may have individual significance, such as providing a reference to a location on a map, or they may be one of many data points mapping out a larger object.

In either case, we begin mapping these data points using our stationary frame of reference as a point of reference. These data points are represented with unity objects known as Transforms, which can be represented as virtual images or simply a reference to a location on a coordinate system.

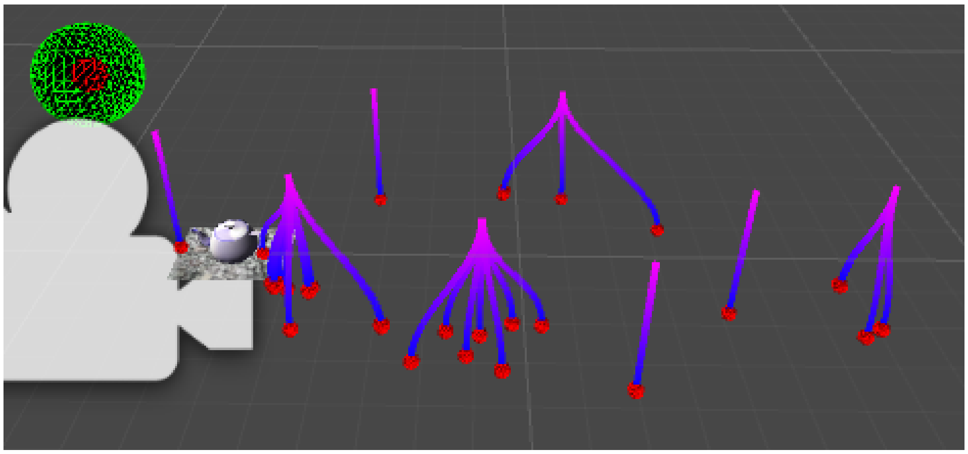

Figure 2: Testing the implementation of our graphing functionality. The altered UI Graph package is being utilized to connect spherical transforms to "empty" transforms, creating bundles.

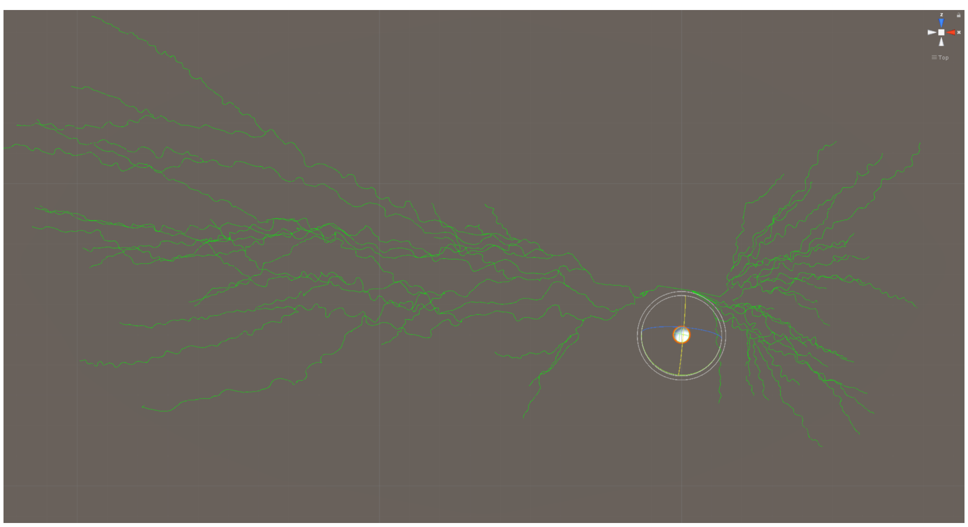

Figure 3: Example using the UI Graph functionality to connect over 11 thousand data points (Transforms), mapping out the structure of a complex neuron.

With our data points mapped in the virtual environment, we often find the need to connect these data points with lines. This may be to represent some correlation between individual data points, or to map the larger objects as mentioned previously. To accomplish this, we utilize Unity's built in line render class. This process is made easier by a package developed by *Broken Machinery* called *UI Graph*. The package uses a parametric curve methodology, called *Bzier curves*, to create curved lines connecting nodes on a graph. With a little re-writing, we were able to use this package to connect the Transforms representing our data points with adjustable curved lines. This proved to not only be an important feature for developing interconnected graphs and mapping out larger objects, but for doing so in a manner that is cleaner and thus easier to visualize. In Figure 2 you can see an example of the altered UI Graph package being utilized to connect spherical transforms (data points) to "empty" transforms (bundle nodes), as part of the construction of an edge bundle graph. Figure 3 shows an example of using the same functionality to map out a larger object, in this case a neuron.

With our objects initialized in the virtual environment, Unity uses the HoloLens's spacial mapping and gaze control information to create a virtual representation of HoloLens's camera, along with its field of view. As virtual objects move into the field of view of the virtual camera, the HoloLens uses stereo imaging to render images of the objects with the perception that they exist in their corresponding locations in the real world. Some of these objects are cable of interaction. Gesture control, gaze tracking assisted by a cursor, and voice commands can be used to interact with our objects. Some of the functionality incorporated includes rendering information, removing it, scaling objects, rotating objects, and of course moving them. The previously mentioned spatial anchors play a big role in many of these functionalities, particularly when it comes to moving objects about. Not only do they assist with simulating accurate movement of virtual objects on a 3D plain, but Unity's collision detection paired with spatial mapping and anchors allow a virtual object's movement to be restricted when real world objects block their path. Alternatively, there is also the option to place virtual objects on real world surfaces thanks to Unity's integration of these various functionalities.

Figure 4: Testing a holographic augmented reality application on the HoloLens utilizing various discussed functionalities provided by Unity, Vuforia, and UI Graph.

We also used the augmented reality platform Vuforia to help with various functionalities. Vuforia provides a powerful tookit to be used by unity in the development of AR applications. The most notable tool we utilized is a target recognition package which creates unique coordinate systems at the location of predefined images, once they are recognized in the real world by the AR device's camera. These images can be used to establish the users position in scenarios where the surrounding environment of the user is already mapped and stored in the device, possibly an entire building, or to simply provide a reference for where visualizations will be rendered corresponding to the real world. For our purposes, we used the latter of these two scenarios. Figure 4 shows a test application which applies the various functionalities discussed in this section.

# CHAPTER 4: MULTIMODEL VISUALIZATION OF NEURON NETWORKS

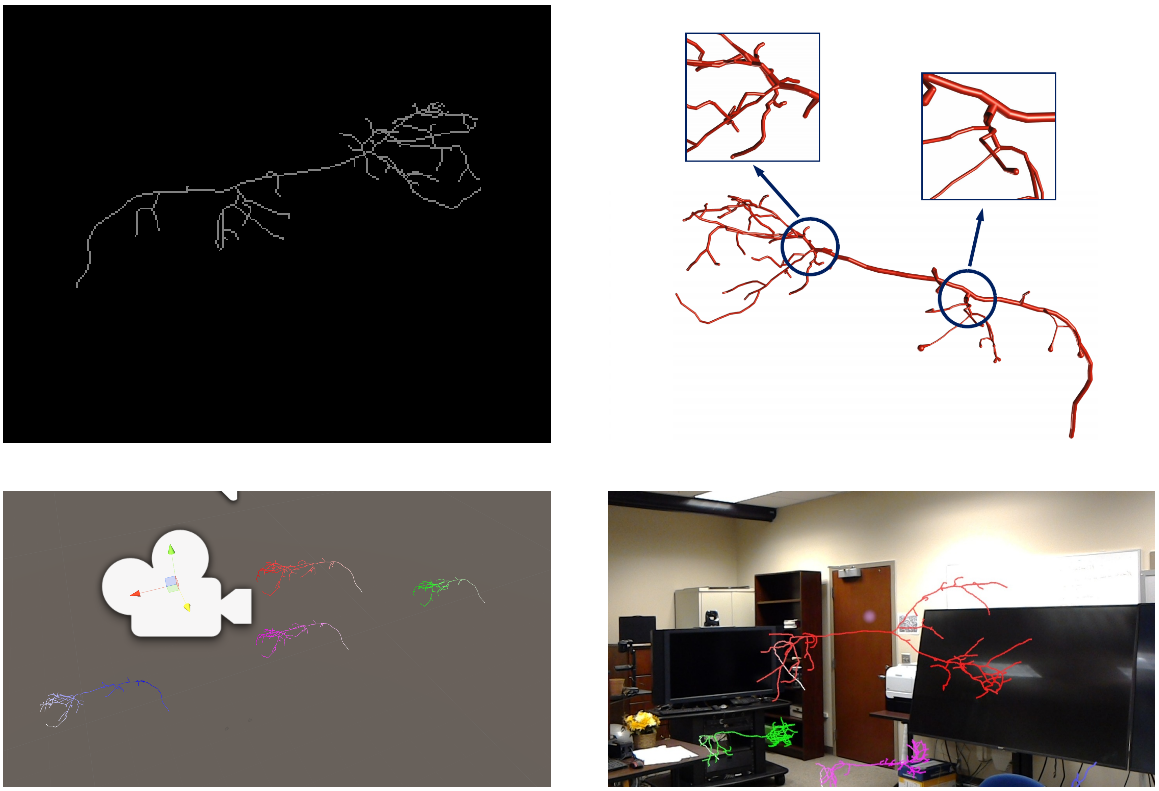

Figure 5: Transition from mapped neuron to augmented reality visualization.

## 4.1 Exploration of Neuronal Morphologies

In the field of Neuroscience, one of the most important topics of interest has been neuron morphology. With strong relevance to neuron properties, such as brain regions, cell types, developmental stages, etc., it is easy to understand why researchers would strive to better understand the morphology of neurons. In recent years, advancements in microscopy and tracing techniques have greatly improved the study of neuron morphology. We are now capable of reconstructing neurons in three dimensions with high precision. With three dimensional neuron data, analyzing the tree-like structures of the neurons makes it possible to describe them using some pre-defined quantitative measurements.

These three dimensional mapping techniques continue to play an important role in neuroscience going forward. With the role neuron morphology plays in determining a neurons connectivity and functional properties, the information that can be gained from new tracing and reconstruction techniques is invaluable to neuroscientists. However, due to the advancements of 3D neuron mapping techniques, there is an ever increasing number of reconstructed neurons being added to public repositories. The large scale data science and complexity of neuron morphology has posed significant challenges. Current techniques has prevented the realization of the full potential of the data. For example, effectively identifying neuron types, similarities in morphology, and uncovering correlations between neuron morphologies and their properties, have been stifled by the complexity of the data. With this data, deep and exhausted retrieval of neuron morphologies is needed to gain any new insight into the understanding of neuron morphology, and its overall roll in neuron function.

## 4.2 Visualizations of Neuronal Morphologies

We have developed an interaction-based adaptive visualization method to assist the interactive exploration process. The goal is to provide an effective visualization approach to study the 3D structures of neuron networks.

We developed our program to visualize the neurons using the Vuforia AR platform, which utilizes the Unity game engine. The combination of the two creates a powerful tool to be used in application development for Microsoft's Windows Holographic Platform.

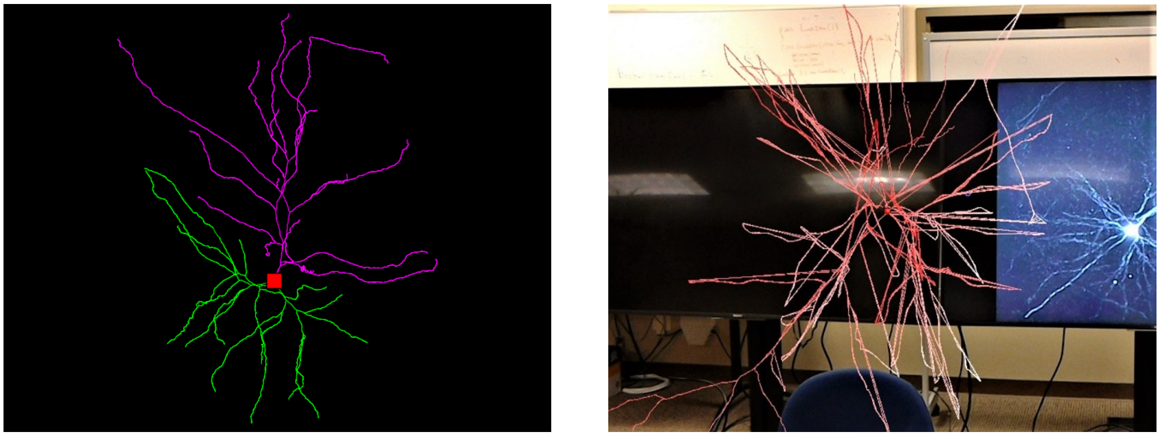

Figure 6: 2D representation of mapped neuron (left) and 3D interactive AR holographic representations of same mapped neuron (right).

When visualizing the neurons, we start by mapping the data points for each neuron to the real world. We do this by creating anchor points in the virtual environment that the HoloLens has created in correspondence with the real world. These anchor points each have their own coordinate systems in which the data points are then rendered. To complete the visualization, we then use Unity's built in line renderer to create lines connecting each data point to their parent data point. The package used to connect these lines uses a Bezier curve methodology, which is a type of parametric curve which can be scaled indefinitely. Using this methodology to direct the curvature of lines in 3D graphs has been extremely valuable in creating a more intuitive visualization experience. However, due to the volume and proximity of the data points for these neurons, no curvature was used on the smaller neurons. Once the neurons are mapped in the virtual environment, the HoloLens uses stereo images to visualize the neurons on the transparent screens of the headset, so that they appear as if they existed in real world in front of the user. As the anchor points correspond to the real world, the visualizations will appear to be fixed to a given location before the user unless moved. This effect persists, even if the user were to leave the room and then return. This is possible thanks to the HoloLens's ability to map the surrounding space, and keep track of the users location and gaze. The neuron visualizations are rendered with depth-based coloring, which adjusts based on the users proximity to the neuron.

We also made the interactible. We placed a transparent transform at the center of each neuron, which also acts as the parent for all the data points of an individual neuron. Upon placing the visually directed cursor of the HoloLens on the center of the neuron, the cursor gives a visual cue, which indicates that an object is capable of interaction. As the transform is transparent, this gives the impression that the neuron itself is the object your are interacting with. The user can then use the HoloLens's the built in "finger tap" gesture select a neuron. Once selected, various voice commands can be used to identify the type of interaction to be performed on the neurons. We included functionality allowing the user to move, scale, and rotate the neurons. With a neuron selected, and interaction type identified, the user would use the built in "pinch" gesture along with hand movement to manipulate the neurons.

Our next step in increasing the intractability of the neurons would be to incorporate the ability to select individual branches of the neurons, visually accenting them, then rendering a larger copy of the individual branch for a more granular view. We are also in the process of creating multi-user support for our visualization programs on the HoloLens. This would allow multiple users wearing HoloLenses to collaborate on information in real time. We are including private workspaces for each individual, and a shared workspace, where selected visualizations can be viewed by all members involved in the collaboration. This allows for a shared perspective where users can analyze information in the augmented reality environment as if they were standing around an actual object in the real world.

## 4.3 Interactive Walk-Through

- A neuron network is placed in a large empty space and a user can walk around the data for observation. The colors of the network are adjusted according the distances to the user. See Figure 7.

- The Neurons can be interacted with in various ways, including: airtap gesture, pinch gesture, voice, gaze. This includes scaling their size, rotation, and moving them. See Figure 8, Figure 9, and Figure 10.

- Multiple neurons can be rendered for easy visual comparison. See Figure 11.

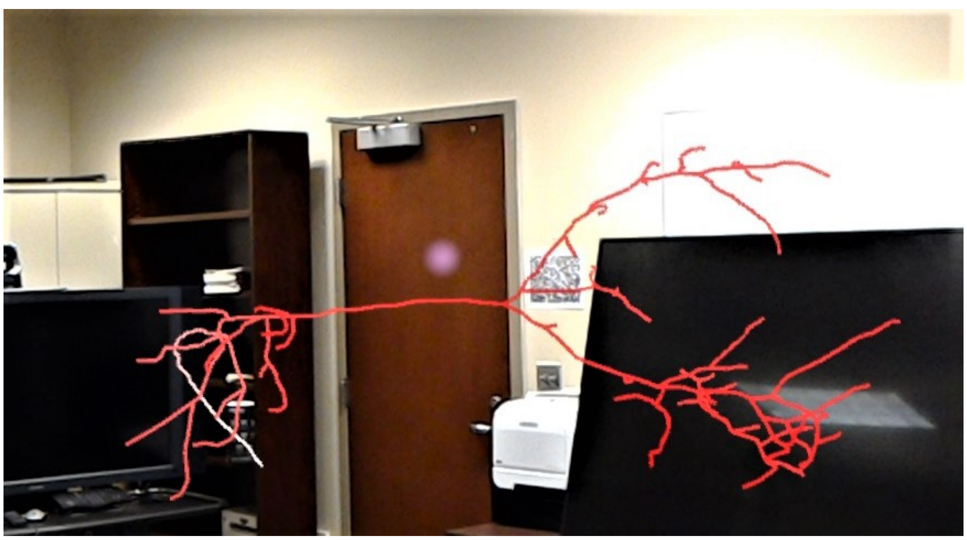

Figure 7: Rendering of single neuron.

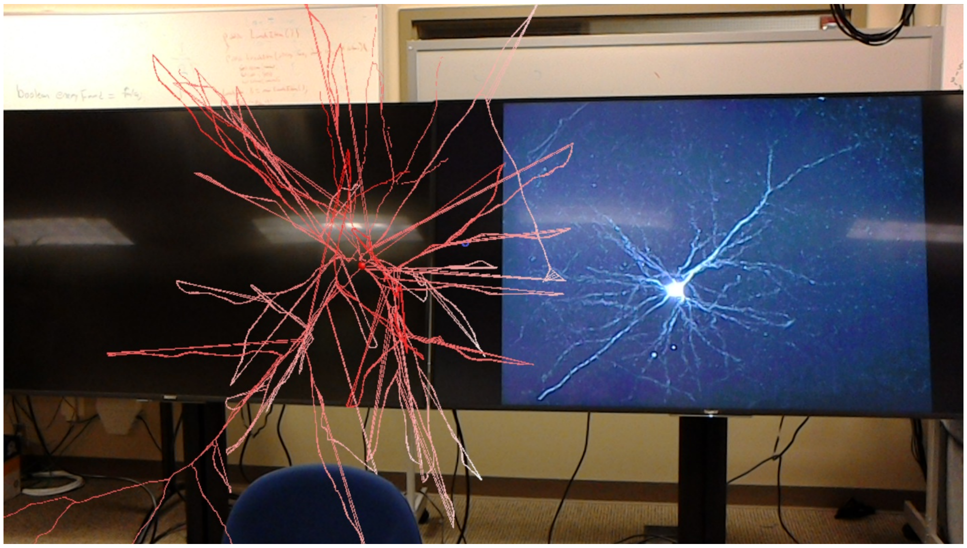

Figure 8: Original Rendering of neuron.

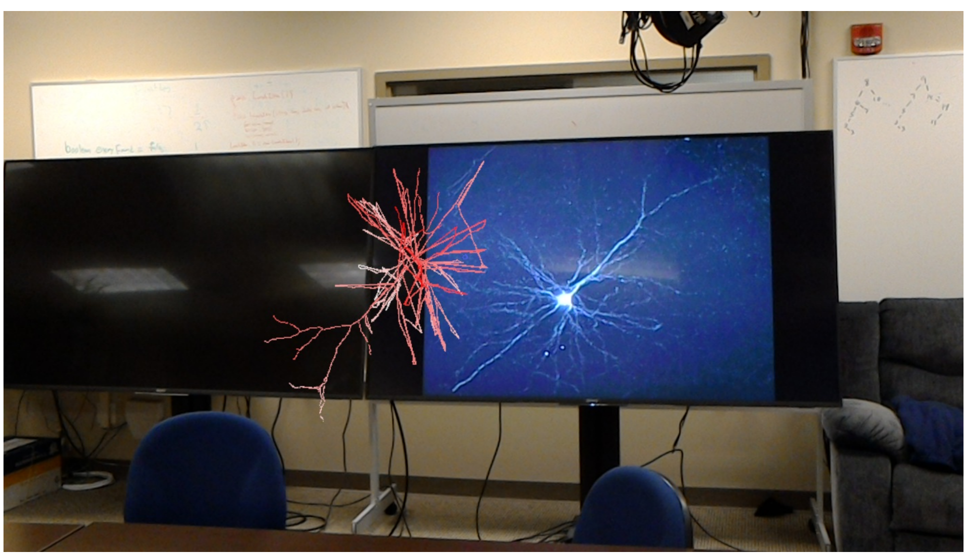

Figure 9: Previous neuron scaled down in size.

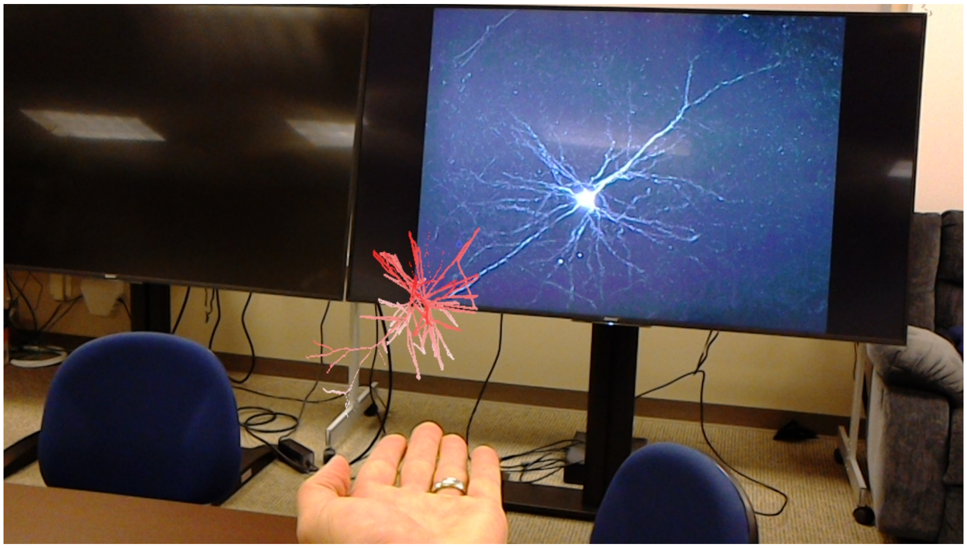

Figure 10: Previous neuron scaled down even smaller in size.

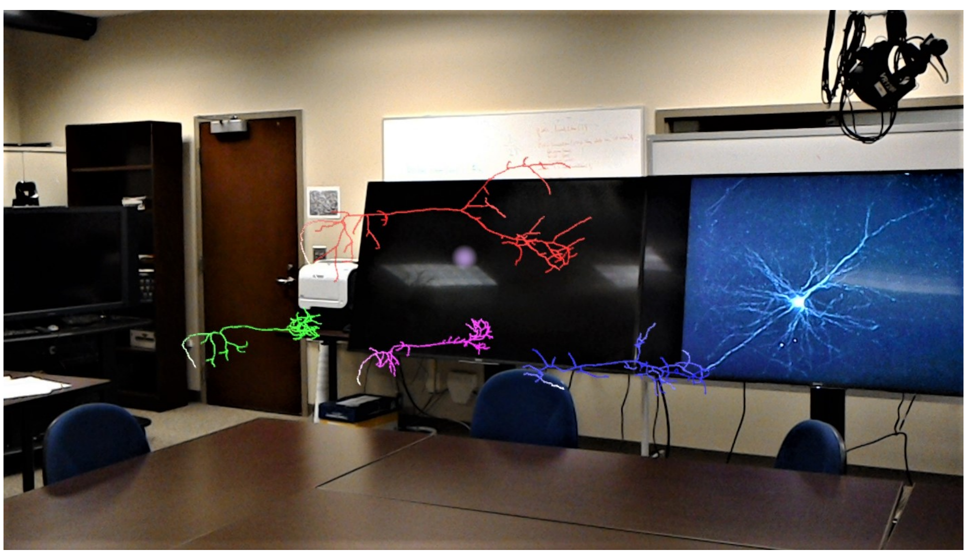

Figure 11: Multiple neurons displayed for visual comparison

# CHAPTER 5: MULTIMODEL VISUALIZATION OF BIODIVERSITY DATA

## 5.1 Biodiversity Hotspot Management

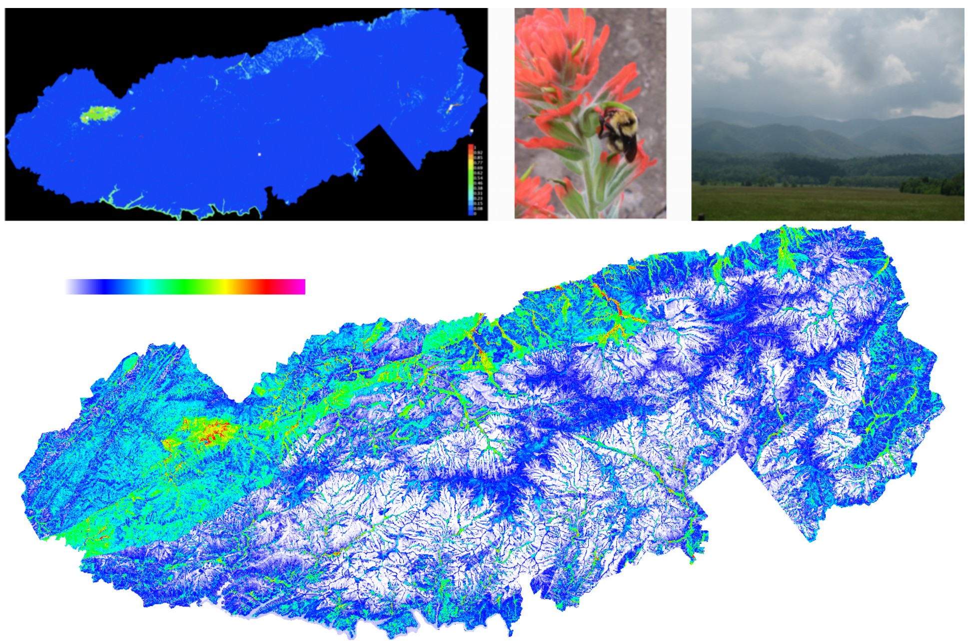

Figure 12: Example habitat range heat map for a species in the Smoky Mountains.

As a World Heritage Site and International Biosphere Reserve, The Great Smokey Mountain National Park (GSMNP) is one of the most important national parks in The United States. The GSMNP contains one of the largest remnant never-cut forests in the eastern US, and it is the most biological diverse park in the United States National Park system. There has been over 19,000 species documented in the GSMNP, and scientists believe there may be an additional 80,000-100,000 species yet to be discovered there.

The GSMNP has teamed up with Discover Life in America (DLIA), to make an extensive catalog of the biodiversity in the park. The DLIA is a non profit group committed to creating and maintaining scientifically sound biodiversity inventory. Together, they created the Taxa Biodiversity Inventory (ATBI). This unique data set holds a distribution of the macroscopic biodiversity in the GSMNP. This would include the geo-referenced mapping of thousands of species occurrences, in addition to more than 40 high quality environmental variables such as geology, soil, terrain, vegitation, and climate. Figure 13 is an example heat map for the habitat range of a single species in the GSMNP.

Currently, due to the wealth and complexity of the data, investigating distribution predictions in the context of the natural habitat is unfeasible. The **goal** of this study is to identify groups of species which have overlapping habitats, and have similar dependencies on their natural environment on an ecosystem scale.

Answers to these questions require an ability to make efficient and comprehensive comparisons of the 1000s of high-priority species, and to do so in a exponentially complex space of how the species depend on the many different environmental variables.

## 5.2 Visualization of Biodiversity Data

Figure 13: The example map for Smoky Mountain.

To accomplish our goals, we have started to develop an application that is centered around the previously mentioned target recognition functionality provided by the Vuforia platform. The target image, or *marker*, is a physical or digital image displayed in the real world, which is recognizable by the users AR device. This marker acts as a point of reference in which digital content can be displayed in relation. A point of reference can be a static location, such as a location in a building which associates the user with a digital structural map, or it could be as simple as an image on a phone, allowing for a more mobile experience. For our purposes, we are using our marker(s) in association with several 75” displays where digital content will be displayed in relation to the information on the screens. Figure 13 shows the map of Smoky Mountains with the Vuforia marker, which the majority of our visualizations will be centered around. We have one display, that for the time being, is dedicated to displaying this image, while a second display is used to provide additional information which is often associated with the images being produced by the HoloLens.

Upon recognition of a marker, the HoloLens will render augmented reality content in relation to the identified marker’s position. The nature of this content will depend on a series of voice commands, and/or the selections made on our second display. For example, the user may pull up a scatter plot on the second display and select a range of data. Species geo-referenced locations or environmental information pertaining to that selection may then be displayed by the HaloLens in relation to images on one of the two screens, or independently. Voice commands can then be used to change how the AR content is being represented, often providing the user with interactive visualizations.

We are currently exploring a number of different ways we can display information which represents various aspects of the GSMNP biodiversity data. The most obvious example would be the geo-referenced species locations, which we display in relation to a map of the park. The primary map we use for this purpose is seen in the previously mentioned Figure 13. Using recorded coordinates for sightings of a given species, the HoloLens displays digital AR images representing species locations based on the Vuforia markers location in relation to the map. Location data for a number of species can be displayed by the HoloLens simultaneously, as the device is capable of rendering thousands of data points.

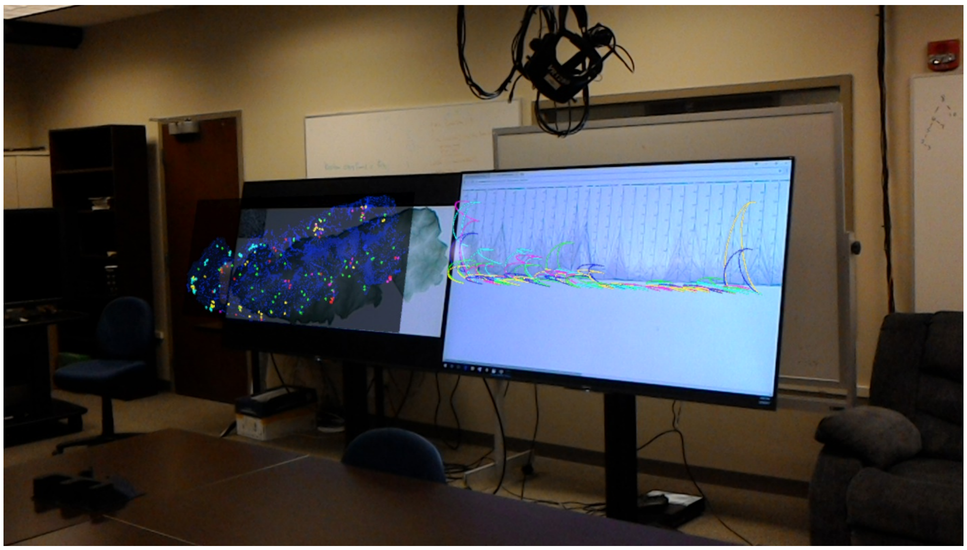

Figure 14: Application of augmented map, geo-referenced data points, and parallel coordinate graph.

To expand on mapping location based information with AR equipment, we also developed the functionality to map this data to digital maps also produced by the HoloLens. By additionally incorporating augmented reality maps, we gain a number of benefits. One such benefit is the capability of using maps that incorporate environmental information, such as geology or rainfall. This allows us to add another layer of information to our visualization. Additionally, we can render any number of maps to take better control of how we would like to make visual comparisons of different aspect of the data. For example, we find it helpful to create layers of semi-transparent maps, each map (or layer) having a different species mapped to it. From here, we can manipulate the locations and orientations of these maps with voice commands and hand gestures. The maps can be aligned with each other, having enough space that the user can see clear distinctions between each map, while still being able to identify each layer due to their individual transparency. This makes comparative analysis and additional manipulation of the visualizations easy and intuitive. The user can easily use hand gestures to separate various layers of the data if the user feels it may give them a cleaner look at a specific layer. As the user moves a central layer, any additional posterior layers will move with the layer the user is manipulating, as they are parented. Additionally, with a quick voice command, all the layers can be collapsed on themselves allowing the data from every layer to be directly compared on a single plane. If the user wishes to compare maps side by side, or one above the other, this is an option as well. This type of functionality greatly increase the scope of information which the user can have access to at one time. Without being restricted to a predefined interface, the user can continuously increase the amount of information simultaneously in front of them as they see fit.

We are also working on various types of augmented reality graphs to help represent the biodiversity data available to us. Some of these graphs exist to help associate information between various sources of visualizations. For example, we may use what we refer to as a *pillar graph*, which uses lines to associate data points on augmented reality maps to their associated locations on maps on our display, or other AR maps. This allows us to clearly associate similar data to multiple maps where each map may visualize some additional, but unique, information. Another graph we use to show similarities of data on our maps is an edge bundle graph. This type of graph bundles data that is close in proximity and

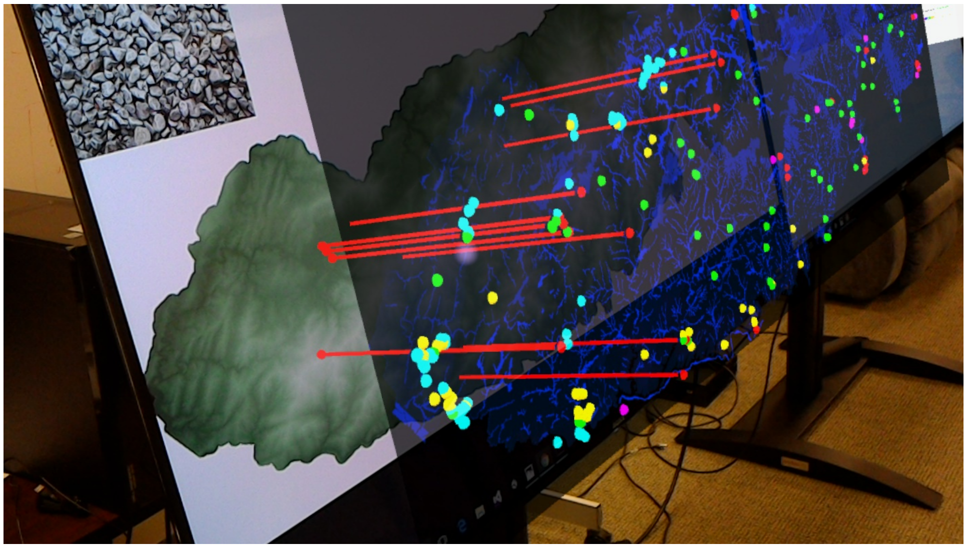

Figure 15: Augmented reality pillar graph example.

shares some type of similarity, then networks those bundles together. This visualizes an association among a subset of data, and does so in a way that creates a cleaner and easier to fallow visualization. Associations among subsets of data can be anything from simply identifying a single species, to bundling species which share an array of environmental variable similarities. Additionally, we use similar graphing techniques to enhance graphs that may be shown on our second display. For example, our parallel coordinate graph compares the 40 plus environmental variables for a large number of species at one time. Each axis of the graph represents one environmental variable, with the height of the axis representing the range of values for each variable. Each species is represented with its own line that intersects each axis at a height based on that species recorded data for that environmental variable(axis). The user can select ranges of data for individual axes, which will fade the lines of species which do not fall within the selected range for the specific axis. This allows the user to isolate species which have similarities among their environmental variables. We then use the HoloLens to create arching lines out of the screen, matching the active species from the graph. This makes visualizing this information much easier as the user moves about the room. Additionally, the selections made on this graph often dictate what species data is rendered by the HoloLens. For example, if the user narrows down three species that are closely correlated on the parallel coordinate graph, the user can have locational data for those species rendered by the HoloLens on the neighboring map. The lines on the parallel coordinate graph and the geo-referenced data points are color coordinated for each species.

### 5.3 Interactive Walk-Through

I would like to start the walk through with a quick disclaimer. There are several challenges we currently face with taking pictures using the HoloLens. One, the HoloLens seems to struggle when using resources on taking pictures. The images often degrade in quality and can be hard to see. Their spatial anchors also seem to suffer a bit with performance as the visualizations tend to be out of place in the photos. Also, the background for our sprites are visible in the photos we take, and unfortunately tend to obscure other content behind them. You will notice this with images of the digital maps produced by the HoloLens. The backgrounds are transparent when actually using the HoloLens.

- Upon making feature selections on the parallel coordinate graph, the arcing lines on the graph, in addition to the recorded locations for the species that fit the feature criteria are rendered by the HoloLens. See Figure 16.
- The user may then chose display the data for each species on individual maps, or layers. This can be done with a simple voice command. See Figure 17.

- From here, there is a number of options available to the user. One would be the option to use hand gestures to adjust the positioning various layers. See Figure 18 and Figure 19.

- Alternatively or afterwards, the user may choose to use the flatten voice command to adjust the layers so they all exist on the same plane, making data points for all relevant species easy to visualize. See Figure 20.

- Another option would be to displace some of the layers to another location, and even flatten them from there. See Figure 21 and Figure 22.

- Finally, the user also has the option to activate various graphs to help view relationships among the data. See Figure 23, Figure 24, and Figure 25.

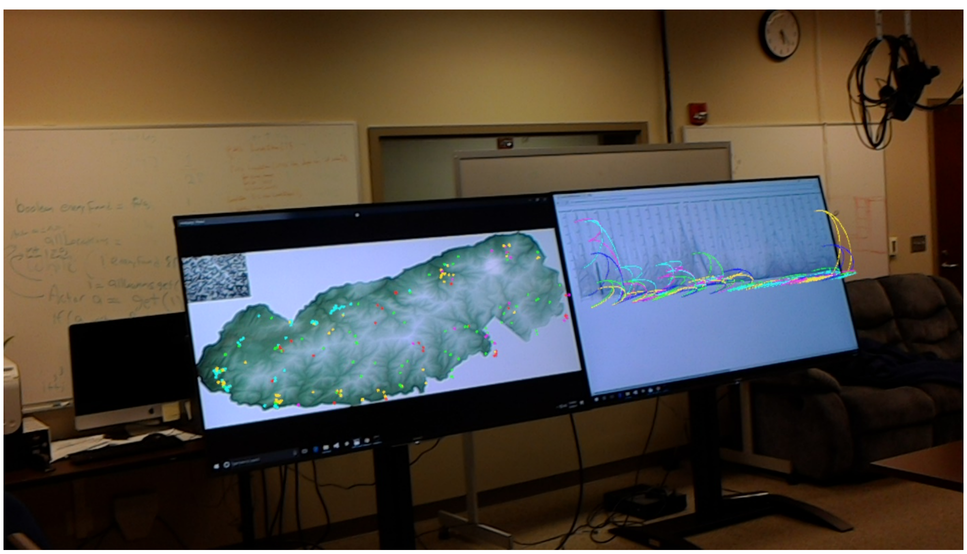

Figure 16: Pillar graph selection and default rendering of associated data points. Note the drift occurring while taking pictures on the HoloLens.

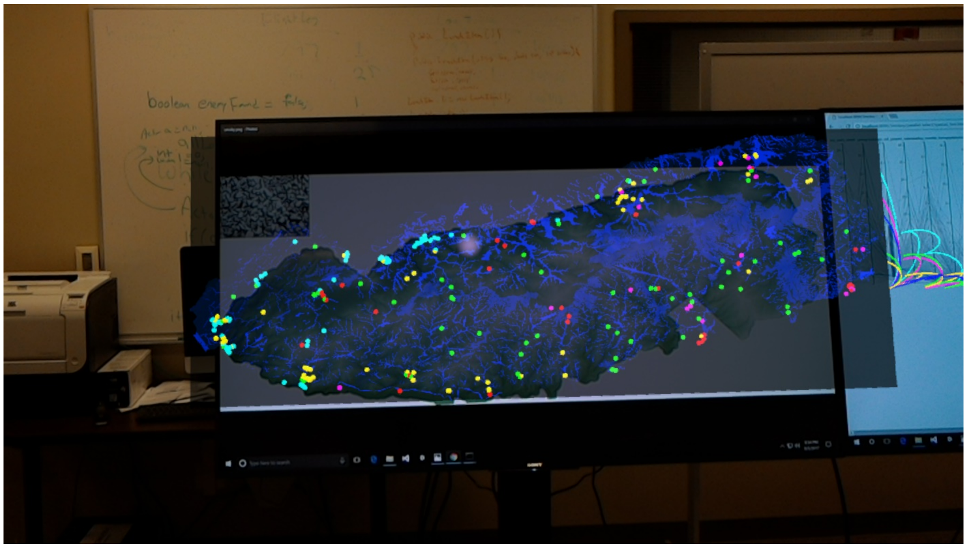

Figure 17: Front view of layered mapping layout.

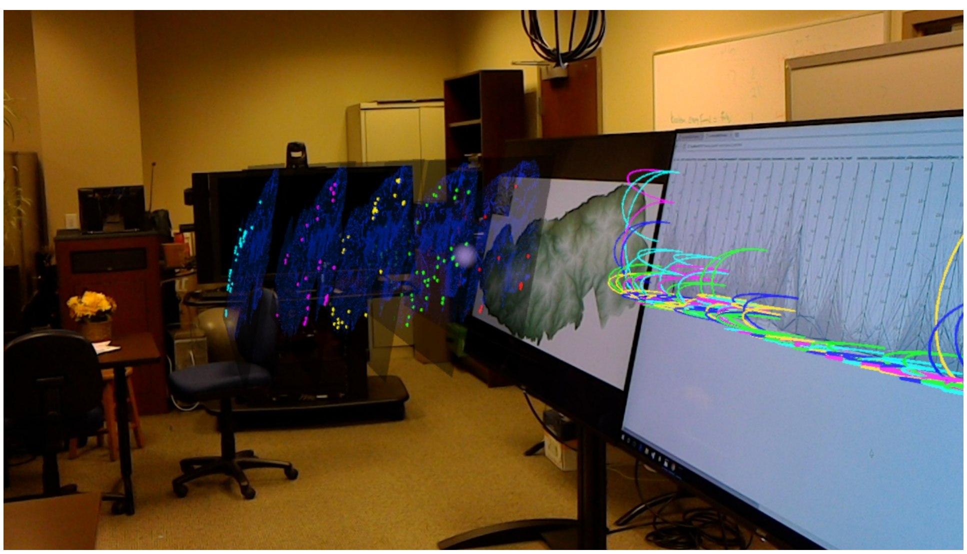

Figure 18: Layered mapping prior to gesture manipulation.

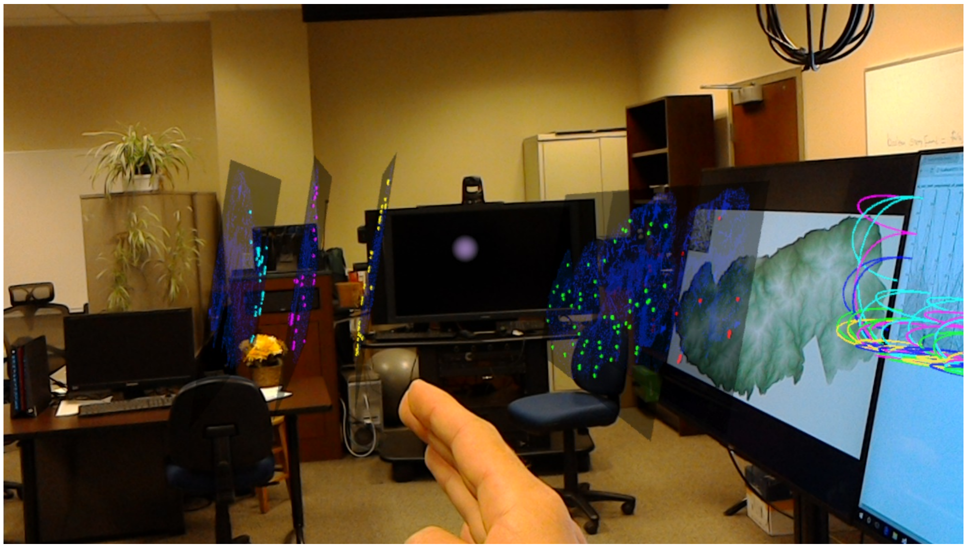

Figure 19: Layered mapping manipulation by gesture.

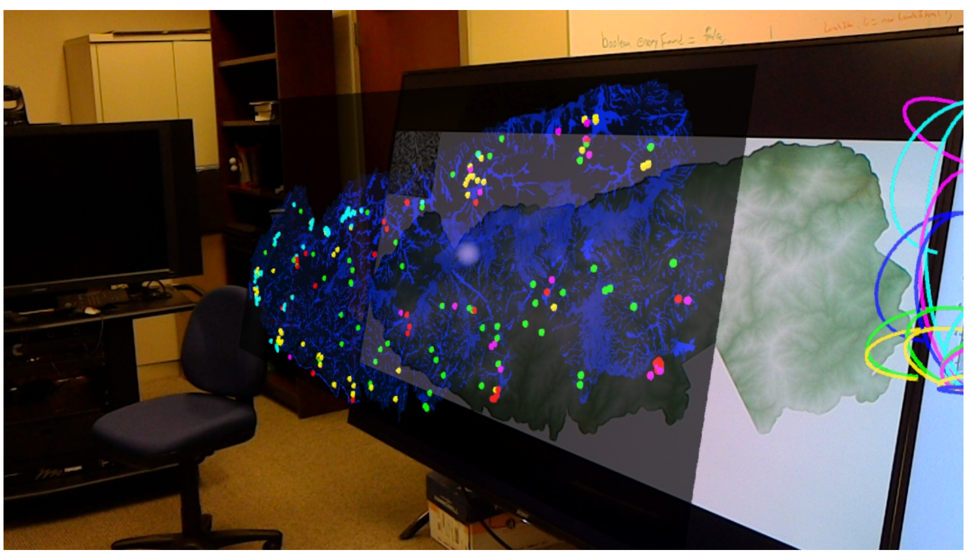

Figure 20: Layered mapping flattened.

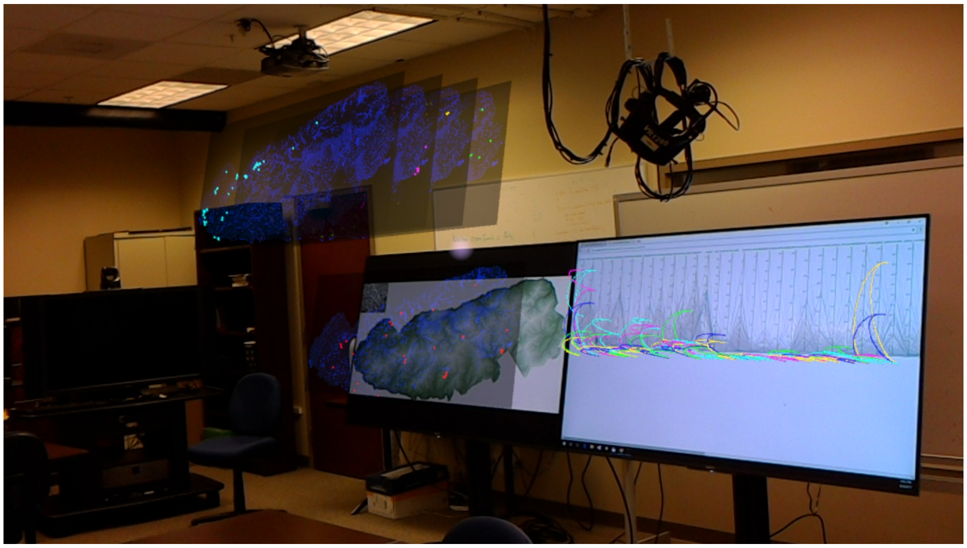

Figure 21: Layered mapping displaced upwards.

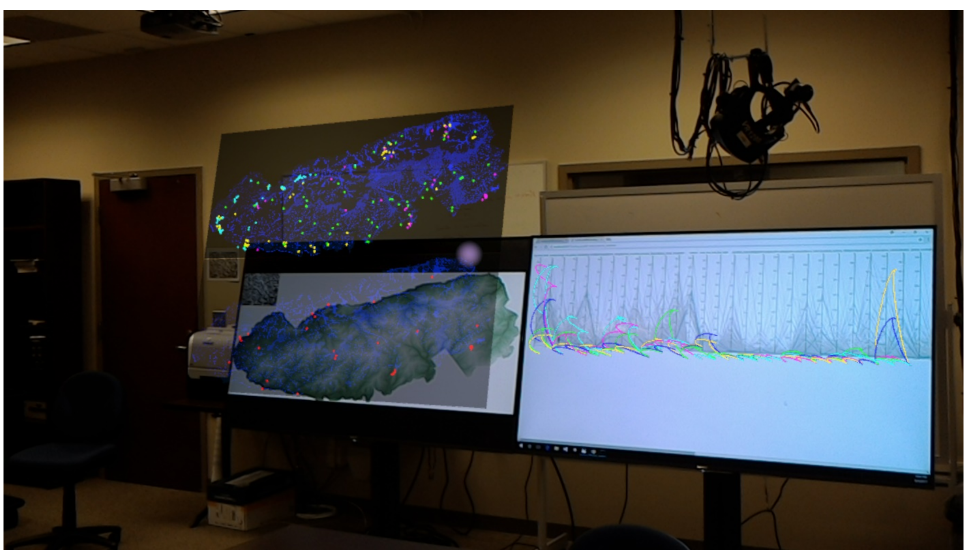

Figure 22: Layered mapping displaced upwards and flattened.

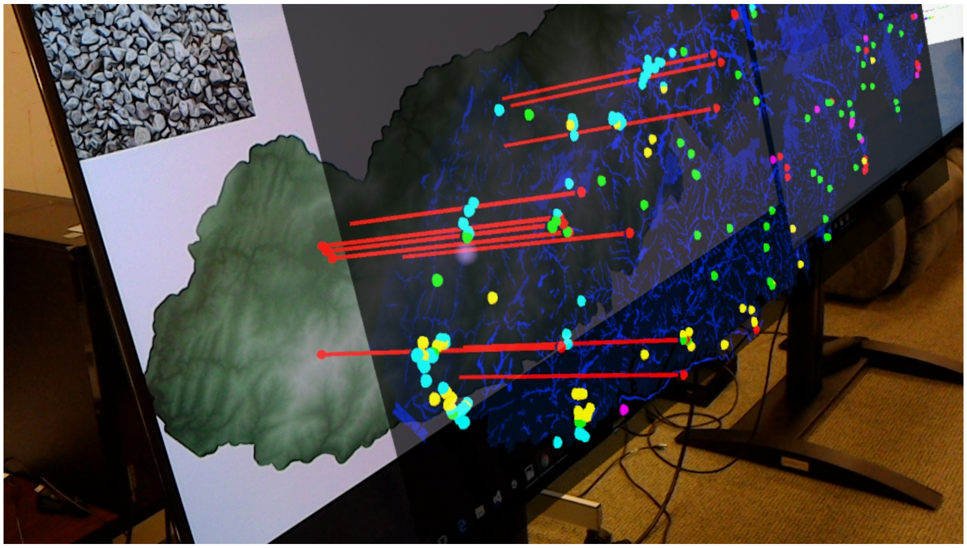

Figure 23: Pillar Graph.

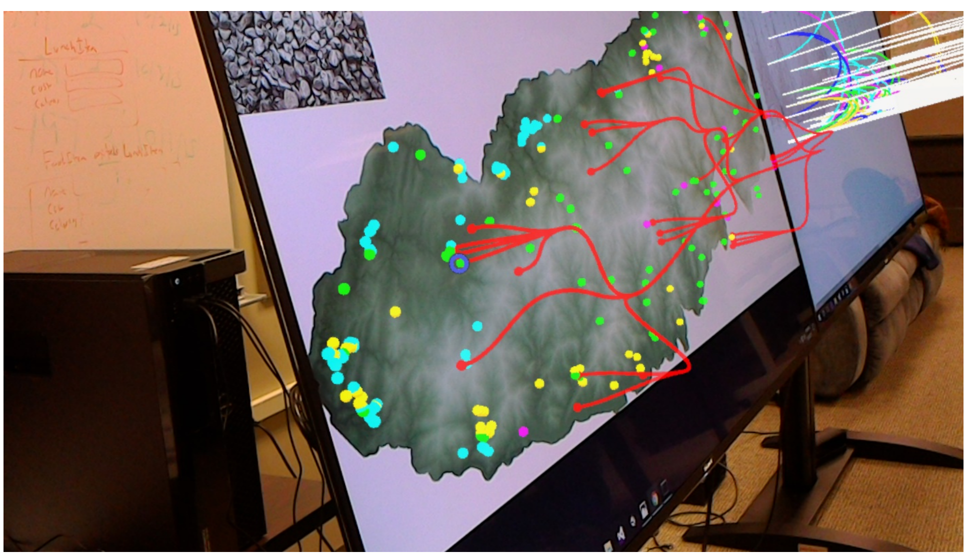

Figure 24: Edge bundle graph example 1.

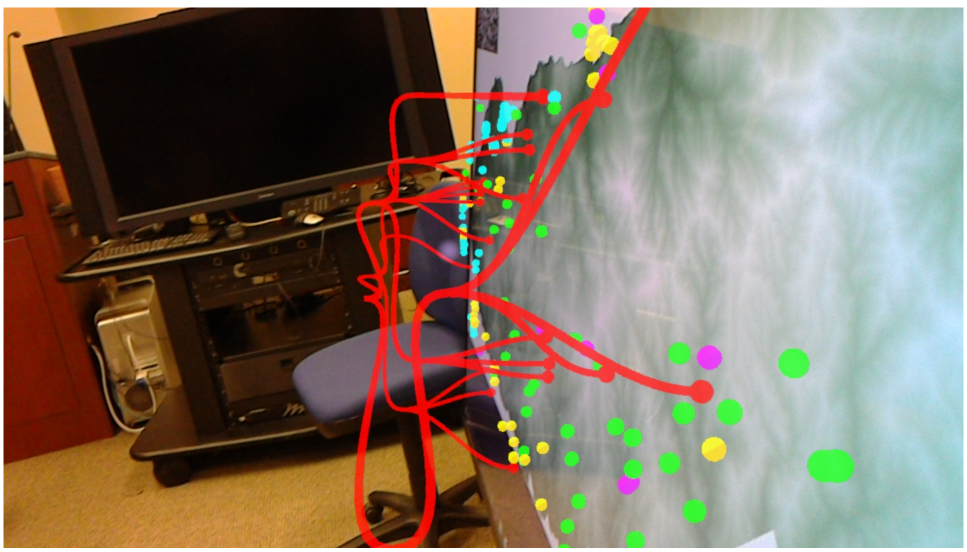

Figure 25: Edge bundle graph example 2.

# CHAPTER 6: CONCLUSIONS AND FUTURE WORK

In our approach to multimodal augmented reality, we used our AR device along side a some what traditional visualization setup. The two 75” displays may be a little out of scale, but overall the dual screen computer set up is quite typical. The idea was to broaden the scope of our visualizations, in addition to creating an intimacy with them, which we hoped would create a more intuitive and impactful experience. We used the HoloLens to build upon traditional visualization tools, and in doing so, greatly increased the scope of information available to the user at one time. More than that, we have started to develop an experience which I feel resonates more with the user.

Sadly, the benefits of AR are largely lost when images are translated back to traditional screens, and explaining interaction is not the same as experiencing it. This makes conveying the impact of the sense of presence brought by AR rather difficult. However, I do feel confident saying, even in the early stages of a much bigger project, that interacting with digital content in this manner feels much more meaningful than traditional methods. What is more, we are just scratching the surface of what AR can be. Using hand gestures to interact with visualizations was probably the biggest step towards making them feel as if they were a part of the world around me. Yet, there are only a couple of hand gestures which the HoloLens is currently capable of supporting. The limited ways we can interact with visualizations is something that will quickly change though. As it does, I believe AR will be propelled forward even faster.

The projects for this thesis utilized a multimodal approach which was built around a rather stationary setting, as you would typically find in a typical work environment involving visualization. I feel it is important to note, just as it was mentioned in the related work, neither multimodality nor augmented reality is restricted by such confines. Real world objects can encompass a multimodal approach, and you can bring today's AR devices to them. This is one of the biggest benefits of AR, and we have yet to fully realize it. While today we are using an AR device to represent geo-referenced information in a lab, tomorrows scientists may very well have AR devices in the field, displaying environmental information based on their current location.

I think we have made some tremendous strides in augmented reality, and our work has demonstrated how a multimodal approach to AR visualization can potentially improve productivity and the users overall experience. However, AR is still in it's infancy as a mainstream product, and still requires much more research if it is to reach, and hopefully exceed, the potential which people see in it. Within the scope of our project, there is still much more we would like to accomplish. Much of the future work planned includes improving the ways in which we interact with our visualizations. Now more than ever, I feel that this may be the most fundamentally important aspect of AR in a mixed-reality environment. We will also look for ways to expand on our visualizations, especially when it comes to representing complex relationships between species and their environmental variables for the GSMNP project. Lastly, we are close to implementing features that will allow multiple HoloLens users to collaborate on projects. Each user with an AR device will have a shared and private work space. Users can work within either work space, in addition to moving content into the shared works space for other users to see and interact with simultaneously.

The idea of two people on different sides of the world being able to interact with the same object, as if they were in the same room together, could have a huge impact and how we collaborate on projects in the future.

# REFERENCES


[1] B. Alper, T. Hollerer, J. Kuchera-Morin, and A. Forbes. Stereoscopic highlighting: 2d graph visualization on stereo displays. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2325–2333, 2011.

[2] R. T. Azuma. A survey of augmented reality. *Presence: Teleoper. Virtual Environ.*, 6(4):355–385, Aug. 1997.

[3] M. Billinghurst, A. Clark, and G. Lee. A survey of augmented reality. *Found. Trends Hum.-Comput. Interact.*, 8(2-3):73–272, Mar. 2015.

[4] V. Cobus, W. Heuten, and S. Boll. Multimodal head-mounted display for multimodal alarms in intensive care units. In *Proceedings of the 6th ACM International Symposium on Pervasive Displays*, PerDis '17, pages 26:1–26:2, 2017.

[5] M. Cordeil, T. Dwyer, K. Klein, B. Laha, K. Marriot, and B. H. Thomas. Immersive collaborative analysis of network connectivity: Cave-style or head-mounted display? *IEEE Transactions on Visualization and Computer Graphics*, PP(99):1–1, 2016.

[6] M. Dostl. Multimodal navigation display. In *2015 IEEE/AIAA 34th Digital Avionics Systems Conference (DASC)*, pages 1–11, Sept 2015.

[7] M. Fiorentino, A. E. Uva, M. Gattullo, S. Debernardis, and G. Monno. Augmented reality on large screen for interactive maintenance instructions. *Computers in Industry*, 65(2):270 – 278, 2014.

[8] J. G. Giraldez, H. Talib, M. Caversaccio, and M. A. G. Ballester. Multimodal augmented reality system for surgical microscopy. In *Proceedings of Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display*, 2006.

[9] O. H. Kwon, C. Muelder, K. Lee, and K. L. Ma. A study of layout, rendering, and interaction methods for immersive graph visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(7):1802–1815, 2016.

[10] B. M. and L. M. *Multimodal Interfaces for Augmented Reality, in Dill J., Earnshaw R., Kasik D., Vince J., Wong P. (eds) Expanding the Frontiers of Visual Analytics and Visualization*. Springer, London, 2012.

[11] H. Schraffenberger and E. van der Heide. Multimodal augmented reality: The norm rather than the exception. In *Proceedings of the 2016 Workshop on Multimodal Virtual and Augmented Reality*, MVAR '16, pages 1:1–1:6, 2016.

[12] D. Sonntag and T. Toyama. On-body ie: A head-mounted multimodal augmented reality system for learning and recalling faces. In *2013 9th International Conference on Intelligent Environments*, pages 151–156, 2013.

[13] C. Ware and P. Mitchell. Visualizing graphs in three dimensions. *ACM Trans. Appl. Percept.*, 5(1):2:1–2:15, Jan. 2008.