STUDYING THE PERCEIVED INTELLIGENCE BETWEEN AGES USING VIRTUAL SPACES

Joshua Wagner

A Capstone Project Submitted to the
University of North Carolina Wilmington in Partial Fulfillment
of the Requirements for the Degree of
Master of Science

Department of Computer Science
Department of Information Systems and Operations Management

University of North Carolina Wilmington

2023

Approved by

Advisory Committee

Geoff Stoker

Brittany Morago

Elham Ebrahimi, Chair

Accepted By

Dean, Graduate School

TABLE OF CONTENTS

ABSTRACT

Studying the Perceived Intelligence Between Ages Using Virtual Spaces
Wagner, Joshua, 2023. Capstone Paper, University of North Carolina Wilmington.

Virtual environments have become an increasingly used form of communication in the modern world. Whether they are text-based, illustrated, or fully immersive, human interaction is now regularly done using a computer. People today will interact with real or virtual people in the same manner through a screen as they would if they were physically in the same room. This provides a strong opportunity for psychologists to study human interactions in a perfectly controlled environment. A simulation was developed in the Unity gaming engine and deployed into human trials with the purpose of studying how subjects will perceive and interact with virtual humans, avatars, of varying ages. Specifically, if they perceive any age group as appearing as more or less intelligent when compared to their peers. The simulation has participants in a boardroom-like setting where they must work with the avatars to solve simple mathematical problems. Participants are asked to judge the intelligence or abilities of the avatars twice, once early and once later in the simulation. Testing with 17 participants, it was found that people overwhelmingly believed that the older avatar was smarter when asked early in the simulation. However, more people believed that the younger man would be more helpful in solving the final math problem near the end of the simulation. The trends show that this simulation can collect valuable data for researchers, and reinforces that virtual environments are a key tool in psychological trials.

LIST OF TABLES

LIST OF FIGURES

Figure Page

CHAPTER 1: INTRODUCTION

The advancement and widespread use of computing technology has led to virtual environments becoming a regular tool for people to interface with. Whether the form is text-based, graphical, or fully immersive (utilizing virtual reality devices), much of human interaction in the modern age takes place utilizing a computer.

”Virtual environment” is a broad term which, in the context of this study, will refer to some abstract space accessed through the use of a computing device. While these virtual environments may be used to communicate with another real user, they may also be used to facilitate interactions between a real user and fictional, autonomous characters. In reference of the latter, if given in the right context, a person may be willing to suspend their disbelief enough to treat these fictional avatars in ways such as if they were a real human being. This opens up a major opportunity in the fields of psychological research, as these avatars and virtual environments can be used to conduct studies in perfectly controlled conditions.

This research project looks into one scenario where a virtual environment can be used to study and potentially alter human psychology. This is done through the context of a modern social phenomenon: the idea that someone is less capable due entirely to their age, or ageism. Age-based prejudice as a concept to look at in-depth is relatively new when compared to other forms of bias such as racism and sexism. Along with other possible factors, it has been studied considerably less than other forms, despite how pervasive it is.

The project aims to answer two questions:

1. Does age play a role the perceived intelligence of a person?
2. Can a virtual environment be used to challenge that?

This is done through the development of a simulation where volunteers converse with virtual people of varying ages, being asked multiple times throughout which of the avatars they believe to be most intelligent or capable of solving some problem. A trend of discrepancies between users’ decisions may reveal that participating in the simulation has some affect on participants way of thinking, and could reveal or challenge inherit biases.

CHAPTER 2: REVIEW OF LITERATURE REVIEW AND ANALYSIS

The following section shares and explains the studies and papers that functioned as preliminary research for this project. To justify this work, research was collected aiming to prove three factors.

1. Ageism is understudied for how uniquely pervasive it is.
2. Virtual environments are not only practical for behavioral studies, but are ideal environments that eliminates the most common issues found in previous experiments.
3. Virtual environments has been used previously to study and influence human behavior. Especially in cases of environmental, social, and philanthropic causes.

Table 1.

Types of prejudice and the number of results found by Google Scholar when searched for.

Search Term	Number of Results
Racism	2,320,000
Sexism	542,000
Xenophobia	276,000
Homophobia	238,000
Ageism	137,000
Classism	85,300
Ableism	48,700

2.1 Ageism is Prevalent Yet Understudied

Prejudice, discrimination, and biases based on social characteristics are common focuses for psychological research. What is unique about age compared to other social groups based on sex, race, or ethnicity is that it is the only characteristic where the group people belong to changes. Throughout a person's life, they may witness or experience age-based prejudice as an adolescent, young adult, middle-aged adult, or elder.

Compared to other forms of discrimination, ageism has been the focus of considerably less studies and scholarly papers. Table 1 compares the number of results found by Google Scholar when each of the discriminatory terms are searched for. This could be due to a number of possible factors, including:

- ”Ageism” is a relatively new term, coined by Robert N. Butler in 1969. It should be noted that searching for the exact phrase ”discrimination against the elderly” still produced considerably fewer results than similar phrases (e.x. ”discrimination against women”).
- Age-based prejudice is less commonly noted due to some societal reason. It may be less common to begin with, harder to detect, more likely accepted as inoffensive speech/behavior, or some combination of these factors.
- Ageism is seen as less harmful than other forms of discrimination, and therefore in less need of critical examination. This could be due to historical and cultural factors.

Erdman Palmore's article ”The Ageism Survey: First Findings” [3] sheds light on how widespread ageism is as well as which populations are most affected by it and in what forms it tends to be presented. 84 seniors (aged 60 and up) were surveyed on how many times they had experienced certain forms of ageism (either zero, one, or two or more times). Over 77% of the respondents identified at least one form of ageism they had encountered, with over half of these incidents having occurred multiple times. The most commonly reported forms identified included ”jokes that poke fun at old people”, being ignored, being patronized, and being shown less dignity or respect. Assumptions about ability due to age were also common, such as the assumption that they could not hear well or that they had a harder time understanding something due to their age. Cases of severe discrimination, such as having property vandalized or being victimized by a criminal, were the least commonly reported. Interestingly, there was little discrepancy in ageism experienced by respondents above the age of 75 compared to below, nor was there a large difference in reported ageism from men and women. Rather, there seemed to be a negative correlation between the level of education and number of reported experiences. Those with education listed as high school or less experiencing an average of five different incidents, while respondents with college or postgraduate education reported an average of four and three and a half incidents respectively.

2.2 The Practicality of virtual environments in Psychological Research

The utilization of virtual environments in psychological studies is a practice that has become more common in use. Virtual environments is not only viable as a space to research human behavior, but remedies many challenges presented by the traditional method of conducting these studies: carefully selected actors in meticulously crafted environments. Xueni Pan and Antonia F. de C. Hamilton wrote their article ”Why and how to use virtual reality to study human social interaction: The challenges of exploring a new research landscape” [4] as a pitch to psychologists to consider adoption of virtual reality in their human trials. This is due to their key desire: theoretically perfect experiment control allowing for consistent research methods. In a virtual environment, researchers are able to change one variable to their exact need without certain complications. They give the example of wanting to examine how race and gender influence a subject’s perspective. ”[Physical trials] would require four different actors of different races/genders – it is hard to assemble such a team, and even harder to match them for facial attractiveness, height, or other social features. With virtual characters, it is possible to create infinitely many combinations of social variables and test them against each other.” The largest challenge faced, one the authors liken to scaling the Olympus Mons, is the creation of a fully-intractable human capable of passing the VR Turing Test. Other challenges faced by using virtual reality include the uncanny valley effect, embodiment, and presence within a virtual space.

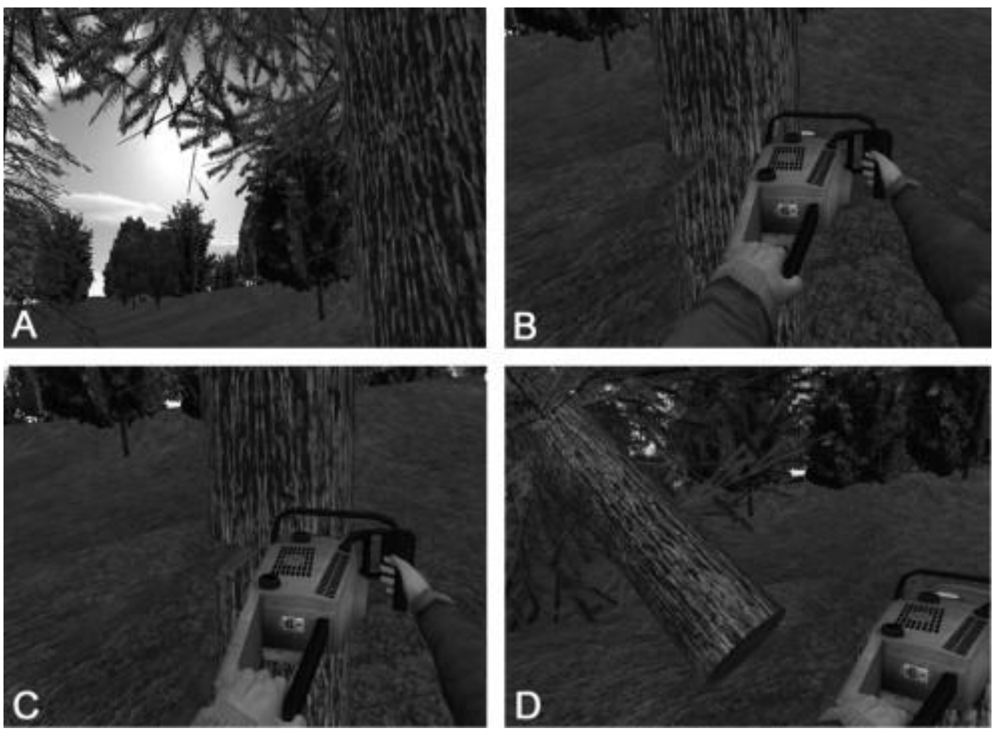

Figure 1: Screenshots of the tree-cutting simulation by Ahn, Bailenson, and Park.

The 2008 Media Psychology article ”What Can Virtual Reality Teach Us About Prosocial Tendencies in Real and Virtual Environments?” [2] suggests that real people are able to connect with and recognize virtual people in manners consistent with real life situations. The study had participants experience a simulation where they would witness a virtual person in need; such as a blind man who had lost his cane. Responses to these situations matched with previous findings about the proportion of those who help people in need in the physical world. In a second study testing the same hypothesis, participants again reacted in ways consistent with their reported level of empathy. Within a virtual environment, people are willing to suspend disbelief to the point where they acknowledge virtual people the same they would real people, suggesting the viability of virtual reality's use in future behavioral research.

2.3 Similar Studies

The basis of this project: using a virtual environment to study and influence human behavior, has been done before and these past studies served as heavy inspiration for this work.

A 2014 study titled: ”Short- and long-term effects of embodied experiences in immersive virtual environments on environmental locus of control and behavior” [1] examined how virtual reality effectively could be used to encourage environmentalism compared to traditional print media. Participants either experienced a virtual environment simulation of cutting a tree down or were asked to describe the cutting of a tree after reading about it in print. Later on, a researcher would (unknowingly to the participant) intentionally spill water, and ask them to help clean it up using napkins on a nearby table. Participants who experienced the virtual environment used around 20% less paper in cleaning up the spill. The conclusion reached was that a perceived personal impact on the environment – a locus of control – had a great effect on environmentalist behaviors.

In 2012, a group in Stanford conducted the study ”Virtual Superheroes: Using Superpowers in Virtual Reality to Encourage Prosocial Behavior” [5] examining if virtual reality can be used encourage prosocial or altruistic behavior depending on the situation presented or actions taken in the virtual environment. The simulation involved participants flying around a virtual city and completing some task. They were divided into groups two-by-two: they either flew around a city in a helicopter or by sticking their arms out akin to Superman; and they would either tour the city or assist a lost diabetic child who was in need of insulin. Similarly to the previous study, researchers would then knock over a cup, this time full of pens, and would ask the participant to help pick the pens up. The groups who helped the child and the ones who flew around as Superman picked up considerably more pens than those in the other groups and those who experienced super flight were also quicker to help. Six participants total did not help pick up the pens, all of which were in one of the helicopter groups.

Both studies utilize virtual reality to encourage positive behavior. Virtual reality offers a unique advantage over other mediums: a sense of presence. Participants can be put in specialized or unrealistic situations with minimal complications outside of the development of the technology. By being the one who cuts a tree down or by becoming a super hero that helps those in need (or even just admires the city) users were able to understand the weight of their actions and believe they were a part of the virtual world. While virtual reality is advantageous for its ability to fully immerse a user, it is not necessary for an environment to be presented as believable for the user.

CHAPTER 3: METHODOLOGY

3.0 The Simulation

The simulation, developed in Unity, serves as a virtual environment where human volunteers interact with virtual humans (avatars). They decide which they perceive to be most and least intelligent. Afterwards they will interact with and watch these avatars before finally being prompted to choose one to help with a math problem. The variables most closely monitored are as follows:

- The choices the user makes in regards to selecting the most capable of the avatars
- How long it takes the user to arrive at the conclusions
- The discrepancy, if any, of their choices and how long it takes them to make them.

If the user chooses a different avatar to help them with the final problem, or if there is a significant difference in the amount of time taken to arrive at that conclusion, then it can be reasonably assumedthat the simulation had some effect on their thinking.

3.1 Human Trials in Three Stages

Trials are done with volunteers using a provided machine and engaging with the simulation in three stages, followed by a questionnaire.

1. The user holds simple one-on-one conversations with three virtual people: a younger man (male one), a middle-aged man (male two), and an older man (male three). The conversations are quick and simple question-and-answer style exchanges where the avatars introduce themselves, comment on something, and then ask the user for their opinion. Male one and three ask the user for their opinion on hot vs cold weather and cats vs dogs. The answers to

these questions are not important, but serve as a way for the user to interact with the avatars while also getting used to how selecting an option will go.

Male two is introduced last, and he will ask the user if they believe the younger or older man is more intelligent. The user's choice and how long it takes for an answer to be selected on prompted are recorded. None of the user's answers in this section affect the simulation in any way, as the avatars will all respond identically no matter what option the user picks for any of the three questions.

2. The user participates a round-table discussion between the three avatars where they discuss simple math problems.

   There are three problems in total, with the discussion of the first two following a pattern: one where male one is correct and male three is wrong, and another where the roles are reversed. Male two will passively partake in the discussions, but the bulk of the conversation is carried by male one and three. The user is prompted to select between the answer proposed by male one and three.

3. Before the third and final question is read, male two receives a phone call and must leave the room. The user and the avatars are instructed that the final problem must be solved in pairs. The user must decide who they would rather work with for the final problem between male one and two. As with the question asking who the user perceives as smarter, the answer chosen and time taken to select it is recorded.

Depending on the user’s choice, either male one or three will leave the room, and the remaining avatar will work with the user to solve the third and final problem. Upon completing this problem, the simulation ends. Data collected, the answers chosen by the user and time taken for each, is exported to a private spreadsheet upon the completion of the simulation.

3.2 The Avatars

The models for all three avatars are provided by the Microsoft Rocketbox Avatar library and were chosen to share as many physical features as possible beyond the difference in age. They remain unaltered with the exception of male one, whose clothes have been changed to a more formal outfit to match that of the other avatars. Figure 2 shows the three models used, with male one, two, and three standing from left to right.

Figure 2. Models for the avatars as taken from Microsoft Rocketbox.

Figure 3 shows male one after receiving an updated outfit.

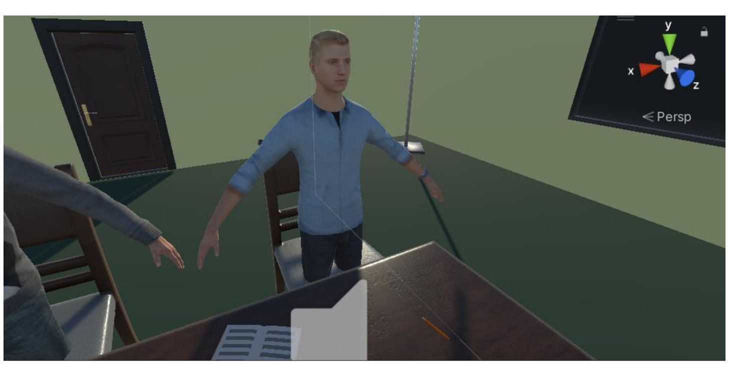

Figure 3: Male one, with updated clothes to more closely match the other avatars in formality.

All three models are properly rigged and animated. The animations necessary for the project include: sitting down, simple hand gestures, and lip-syncing. Sample animations for sitting and gestures were included alongside the models in Microsoft Rocketbox, while the lip-syncing was accomplished using Oculus Lipsync for Unity Development: a Unity package that automatically animates mouth movement in accordance to some provided audio. Overall, Oculus Lipsync proved to be an effective free solution for lipsyncing, and allowed the avatars to appear considerably more lifelike than they would have otherwise.

With a heavy focus on dialogue, there was a heavy need for the avatars to be voiced. The voices were provided by three actors and are internally referred to as ”Voice B”, ”Voice T”, and ”Voice L”. Voice B is deeper and somewhat monotone while Voice L is more soft-spoken. To account for the possible effect this would have on a user’s perception of the characters, the younger and older man can speak the lines provided by either Voice B or Voice L. This is decided randomly before each simulation instance loads. The middle-aged man always uses Voice T, regardless of which voices the other avatars are assigned. All results are collected noting which voice is assigned to which avatar for that specific simulation instance.

With a heavy focus on dialogue, there was a heavy need for the avatars to be voiced. The voices were provided by three actors and are internally referred to as ”Voice B”, ”Voice T”, and ”Voice L”. Voice B is deeper and somewhat monotone while Voice L is more soft-spoken. To account for the possible effect this would have on a user’s perception of the characters, the younger and older man can speak the lines provided by either Voice B or Voice L. This is decided randomly before each simulation instance loads. The middle-aged man always uses Voice T, regardless of which voices the other

avatars are assigned. All results are collected noting which voice is assigned to which avatar for that specific simulation instance.

3.3 The Virtual Environment

The environment that the trials take place within is a small room with furniture resembling that of an office boardroom. In the middle, there is a large table where four chairs are seated. One chair is placed under the user's point-of-view, while the other three are positioned so that the avatars will sit in them. Figure 4 shows this space. On the wall to the right of the user, there is a television. This is used to display graphics relating to the current math problem being discussed. In front of the user in a corner there is a lamp acting as a light source, and behind them there is a fan that creates a soft noise. The furniture was selected to make sense in context and to not be distracting nor overwhelming for the user. The focus must remain on the avatars, rather than the environment.

3.4 The Math Problems

The user must work with the avatars to solve three math problems. The questions were written to be around a seventh-grade level of difficulty, in hopes that would be a sweet spot where they are simple enough to solve mentally while still being understandable to make a mistake. The problems are also intended to be the same level of difficulty. The math problems as they are presented are as follows:

1. Jill is going to a picnic with her friends this weekend, and she has decided to bring a pie. She visits her local bakery and notices that they only have two left. The apple pie has a radius of 5 inches, while the blueberry pie has a diameter of 9 inches. If both pies cost the same amount, which pie is the better value?

2. Jill has decided to buy the apple pie. The pie costs $11, plus an additional 5% tax. Jill decides to pay in cash, using a $20 bill. How much money will Jill receive in change?
3. Jill brings the apple pie to the picnic to share with her three friends: Tom, Nancy, and Robert. Jill cuts the pie into four equal slices, one for each person. However, Robert mentions he is on a diet, so he gives half of his slice to Tom. What fraction of the pie was eaten by Tom?

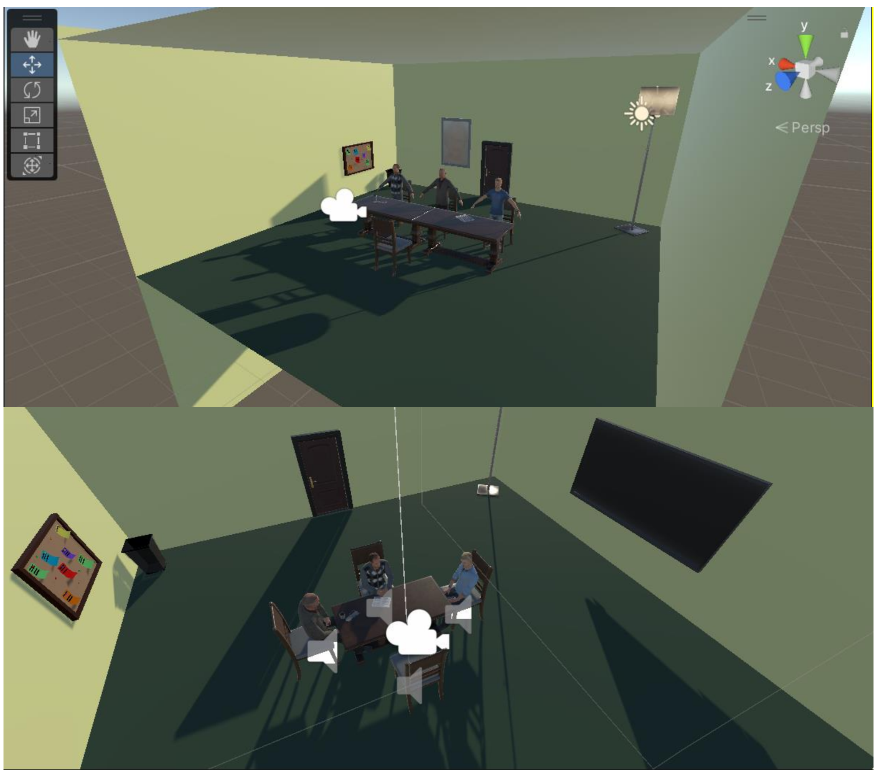

Figure 4: An early prototype for the virtual environment (top) vs the final rendition (bottom), rendered within Unity.

3.5 The Post-Simulation Questionnaire

After the simulation, the user is asked to fill out a questionnaire. In this, the user is asked to choose words to describe male one and three, if their perception of the avatars changed over the duration of the simulation, and to rate the difficulty of each math problem. These answers along with nonidentifiable demographic information (Age, gender identity, and racial identity) are matched with the corresponding row in the spreadsheet to complete the data collection for each user.

Users were asked to choose any and all of the following words they would use to describe the younger man, and again for the older man.

- Confident
- Intelligent
- Confused
- Intimidating
- Timid
- Weak
- A Leader
- Unintelligent
- Calming
- Friendly
- Rude
- Condescending

Participants were asked to rate the difficulty of each math problem on a scale from 1 (very easy) to 5 (very difficult). Because Voice B and Voice L get the first and second questions incorrect, examining if participants felt one question is considerably more or less difficult than the other could reveal an incongruity with how the avatars are presented.

Participants are also asked to report how confident they were in selecting which avatar appeared more intelligent and in selecting which avatar would stay for the final question. In addition, they select if their impression on the younger and older avatar became more or less positive (if any change at all) over the course of the simulation.

CHAPTER 4: OUTLINE OF DEVELOPMENT PROCESS

4.1 Initial Proposed Timeline

Development of the simulation began in spring of 2023. A simple environment was been created and much of the work involving the character models had been completed on an expedited timeframe.

The bulk of development took place over the summer of that year. Priority additions were user interaction from the in-built UI tools, convincing animations that avoided the uncanny valley, and VR support.

A Meta Quest 2 headset was loaned from the university with intention of the simulation running through the virtual reality device. Over time, this would be dropped as the scope of the project needed to be reduced and difficulties were faced in producing a working product capable of VR by the deadline.

Desired data and collection methods were decided upon early and remained consistent throughout the project. The goal was to have at minimum twenty participants of varying ages, and to analyze the data to find the following conclusions:

- How often each avatar was ranked as most or least intelligent.
- How often the user changed their mind from the start to the end of the trial.
- How long it took for the users to come to these conclusions.
- If there are any trends with specific demographics choosing specific answers. See Figure 5 for a listing of demographics of participants by age.

4.2 Using Unity

The simulation was developed using the engine Unity. Unity was chosen for the ease of use, lack of paywalls, and wide variety of hardware compatibility in the case that the simulation would be run on different machines. One particularly helpful aspect of

Unity was the asset store, which provided free models used for the furniture. The ease of access allowed several high quality models to be inserted directly into the scene with little to no overhead.

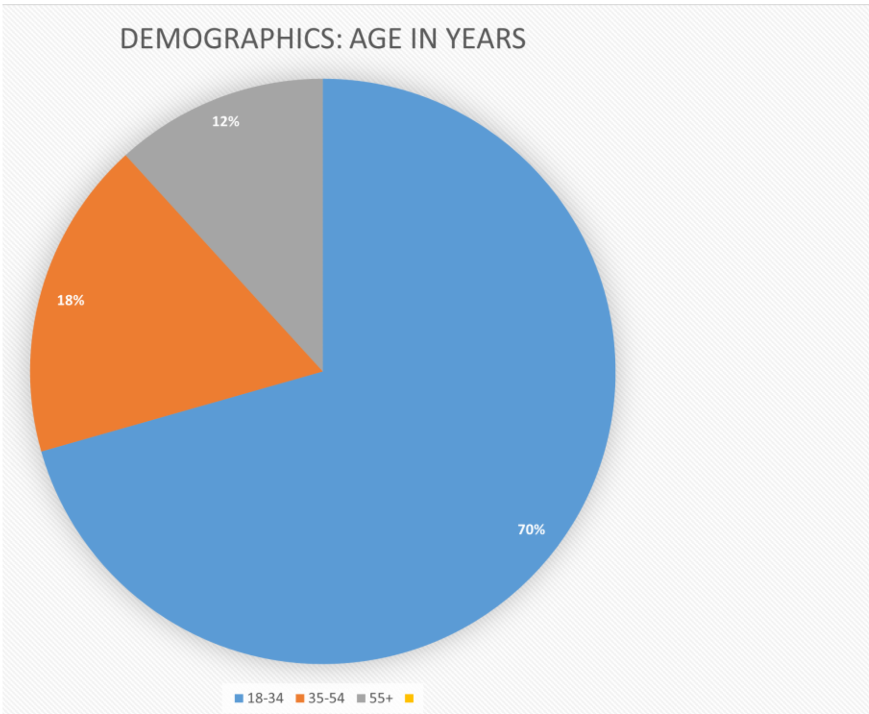

Figure 5: Chart of the demographics of participants by age.

Unity scripts are written in C#, and are used in conjunction with the editor to apply logic to objects and scenes. Because the simulation takes place in a single scene with largely stationary props, the coding necessary was relatively light. Most of what was needed are basic aspects of the engine, such as user interface, audio sources, and the loading/unloading of objects within a scene.

A considerable amount of work went into the dialogue systems. When the simulation is initiated, a random variable assigns voice B and voice L to the audio sources on the younger and older male. To make a conversation between the avatars flow naturally, the simulation needs to keep track of what the previous, current, and next line for each avatar is at any given moment, in addition to a variable tracking the overall progress throughout the full conversation. The audio files for each of the three voices are saved in arrays and each array has an associated integer that increments whenever that line is spoken. Whenever a line is spoken, the corresponding integer for that array (along with a single, global counter) increments by one. The simulation unloads that audio clip and loads in the next one for that set of voice lines. Using the global counter, it references another array that contains the order each voice line is played to decide which line will be spoken next. At any given point, every audio source in the scene has the next audio clip to play loaded, and will wait for the global counter to match an index in the speaking orders array in order to play. The result is that the lines are spoken in order without awkward delays and without any abrupt cut offs. This process for playing the dialogue runs asynchronously as a coroutine, and is also used to display subtitles at the bottom of the screen.

Hard coded into the dialogue coroutine are exceptions when the global counter reaches certain numbers. When this occurs, special instructions are read in addition to playing the next dialogue. These special instructions can include loading/unloading objects (such as the avatars when they enter or leave the room), displaying choices for the user while waiting for input, and for displaying math problems on the television.

Another involved script was the one used to record and export user input. Whenever the user is prompted to make a choice, a canvas containing two buttons is displayed. All coroutines are paused until the user selects an option. The buttons each correspond to one of two options relating to the question that was asked, and each contain text that changes between questions. When the canvas is displayed, an internal timer starts. When the user selects an option, the timer is stopped, and the time taken along with the option they chose are saved to a GameObject in the scene. When the final question is answered, the simulation exports all data saved to this GameObject to a spreadsheet in a .csv format. This was done using the built-in TextWriter and StreamWriter classes, allowing variables to be appended to an existing spreadsheet or create a new one if no such file exists.

4.3 Troubles in VR Development

Originally, the simulation was built utilizing the OpenXR framework to allow compatibility for VR devices. Namely, it allowed motion tracking for the hardware and enabled input using the controllers. While OpenXR greatly streamlines the process, development for VR still proved more challenging than for a traditional display. One large factor was physical space. To accurately test the simulation in VR, a large amount of open space was necessary to allow for full, unrestricted movement. Despite the final simulation requiring very little movement from the user, and asks that they are seated, the need for an appropriate environment remained.

The headset used in development and the target platform was the Meta Quest 2. To begin development for the platform and to allow loading of custom software, the headset must be signed into with a Meta developer account. The process for creating a developer account involves the creation and assignment of some business that applications would be developed under. The device must also have developer mode enabled, which requires the use of the Meta Quest app on a mobile device and for the headset to be paired to the device. Putting the headset into developer mode and selecting an .apk from the file explorer does not work, as this only allows the user to view the properties of the .apk rather than load it despite the headset running on Android firmware and being compatible with them. It should work, but it did not. Eventually, a solution was found, as use of a third-party PC application allows the .apk to be injected into the headset, allowing the simulation to run on the headset without need to be connected to a PC. The documentation for this feature as it was during development was largely either outdated or inaccurate, as it appeared to have been written while the technology was still in some iterative phase.

Alternatively, to get a custom program running on the headset, the use of Quest Link allows the headset to be paired with a PC running the Oculus Desktop App. This can be done using a USB cable or wirelessly through Air Link. Quest Link is a function of the headset that allows a user to view and use their desktop computer within virtual reality. This method could've been used to test VR capability during development, however, Unity using OpenXR natively supports use of VR hardware within the editor without the need of the Oculus app, although it will work with it.

This set up worked until some time late in development when, without warning, the headset was no longer recognized by Unity. After searching for the solution, it appeared that use of Quest Link was suddenly required. After signing into the Oculus app on the PC, Quest Link was unable to connect the headset to the PC. Whether using the cable or Air Link, the devices would search endlessly and never detect one another. After searching online, no suitable solution was found. The backup plan was to build the app directly onto the headset and run it without the use of a PC. However, the build was unable to complete due to unknown errors within the Oculus Lipsync script. This issue, once again, has no documented solution. The following actions were taken in attempt to find a single way to run the simulation on the Meta Quest 2:

- Factory resetting the headset.
- Reinstalling the Oculus Desktop App
- Reconnecting using the Meta Quest Mobile App

- Creation of a new Meta business account
- Reinstallation of the Unity Editor
- Manual combing of the Oculus Lipsync script to find the issue
- Use of the Meta Quest Developer Hub Desktop App, separate from the Oculus Desktop App. The Developer Hub was able to detect the headset, but cannot utilize Quest Link. Even if both the regular and developer hub apps were running and the developer hub was actively recognizing the headset, the regular app never could.

Each of these solutions were deployed several times each, with no progress made despite approximately a week and a half dedicated solely to troubleshooting. With the deadline approaching, the decision was made to drop the VR aspects of the project and transition to a desktop simulation, focusing more on the data collection and analysis.

While the project was initially pitched as one utilizing virtual reality, the technology is not wholly required for the collection of data. Human beings are willing to suspend disbelief when interacting with virtual characters even in non-immersive virtual environments. In a desktop format, the simulation still serves as a tool where participants interact with and must choose between avatars. Any patterns or biases within the data will still be evident regardless of the format of collection. It was decided to continue the project without use of VR for these reasons.

4.4 Initial Concerns

The following is a list of possible concerns for the project that were had going in and some considerations for how they may have been tackled.

- Given that trials will be held in a university space, the pool of volunteers is overwhelmingly comprised of young adults. Ideally, all age groups are represented in equal numbers, as to mitigate a skew where college students will likely favor the younger avatar over the older one.

- It is unsure if participants would still interact with the avatars in a manner consistent to their behaviors in reality after the pivot away from virtual reality.

CHAPTER 5: CONCLUSIONS AND FUTURE WORK

5.1 Trials and Recruitment

Participants were found by simply approaching and asking if they would like to be part of the study. Participants were gathered primarily in Randall Library on UNCW's campus. All trials were completed without incident.

In total, 17 subjects participated. Participants were divided into three different age brackets: younger adults (aged 18 - 34), middle adults (aged 35 - 54), and older adults (aged 55+). Of all participants, 12 were younger, 3 were middle-aged, and 2 were older. These brackets were chosen to relate to the three avatars present within the simulation.

The total number of subjects is lower than initially hoped. It's possible that the research may be continued at another time to collect further data. While the sample size is too small to make conclusions about any group at large, some inferences can be made from the gathered results.

5.2 Inferences From the Data

It was found that participants across all age groups were considerably more likely to initially claim the older man appears more intelligent. However, people were more likely to choose the younger man to assist with the final math problem. There was a noticeable preference for voice B when selecting which of the two avatars was smarter, but no preference between the voices when choosing who to stay. Racial and gender identity of the participants had no discernible impact on their decisions. The sample size of middle and older-aged participants is too small to make any meaningful comparisons, as the subjects were overwhelmingly classified under young adults.

When responding to the questionnaire, participants were more likely to refer to the older man as ”Intimidating”, ”Weak”, and ”A Leader” while the younger man was more commonly referred to as ”Calming”. Both avatars had an equal amount of participants refer to them as ”Intelligent” (used by thirteen participants each) as well as ”Unintelligent” (used by one participant each). When rating the difficulty of the three math questions on a scale from one to five, the final scores averaged out to question one being the easiest (1.82 average rating), followed by question two (2.12 average rating), and question three being the most difficult (2.47 average rating). Out of the 17 participants, two answered question one incorrectly, two answered question two incorrectly, and four answered question 3 incorrectly. Question two being viewed as slightly more difficult on average compared to question one seems to have had little to no effect on the perception of the intelligence of the avatars.

On average, participants were slightly quicker in making choices whenever selecting male one. Nearly across the board, the decision for which avatar stays was made faster than the decision for which appeared smarter. This could suggest increased confidence in answers, as participants were more likely to respond ”Agree” or ”Strongly Agree” when asked if they were confident in choosing which avatar stays when compared to when asked who they believed was more intelligent.

5.3 Addressing Concerns

- Subject pool is lacking in size and diversity, can not represent a larger population as it is currently. The age range of participants is especially unsatisfactory. A longer, more structured testing process would have helped in this regard. Especially if held in multiple locations, including offcampus sites. For a study focused on perceptions of age, having 70% of participants fall under ”young adult” is wholly unsatisfactory.
- While VR is not necessary for a psychological trial, it would have benefited

the study and allowed for more confidence in participants suspending disbelief. A future iteration of the simulation could be made to run on VR to act as a closer match for a traditional psychological study.

- The position of the avatars may have had an averse effect on the subjects. The younger man was always positioned to the right of the user, and the television that displayed graphics relating to the math problems was positioned behind him. These may have impacted user's views of the avatars and could explain how the younger man was preferred when asking for help on the final problem.

5.3 Future Works

Fundamentally, this project serves as a proof of concept including some preliminary research. The software as it currently stands is capable of producing valuable data that can be used to identify age-related biases. Ideally, human trials should continue with a wider range of subjects to increase confidence in the validity of the findings.

Further development of the simulation itself should also be beneficial. Implementation of VR and shuffling of the avatar's positions (in conjunction with moving the TV prop to a more neutral location in the office) addresses the primary concerns with the final study. The simulation as it stands is to be made publicly available, and further research utilizing it is wholly encouraged.

5.4 Conclusion

Ageism is understudied compared to other forms of bias. A simulation to identify these biases was developed and has revealed signs of a skewed perception based on age. Participants assumed the older avatar was more intelligent on a brief impression, but preferred the younger avatar when asked to work with one to solve the math problem.

This pattern is clear even when the voices used and sentences spoken were switched between the two.

The sample size for and diversity of participants is somewhat underwhelming, so further testing would be necessary to determine more accurate results. However, the utility of the simulation is undeniable, and the gathered data does point to some consistency. Ideally, this study acts as a signal boost for research into age-based discrimination, and as a sign that virtual environments are viable methods of conducting psychological studies.

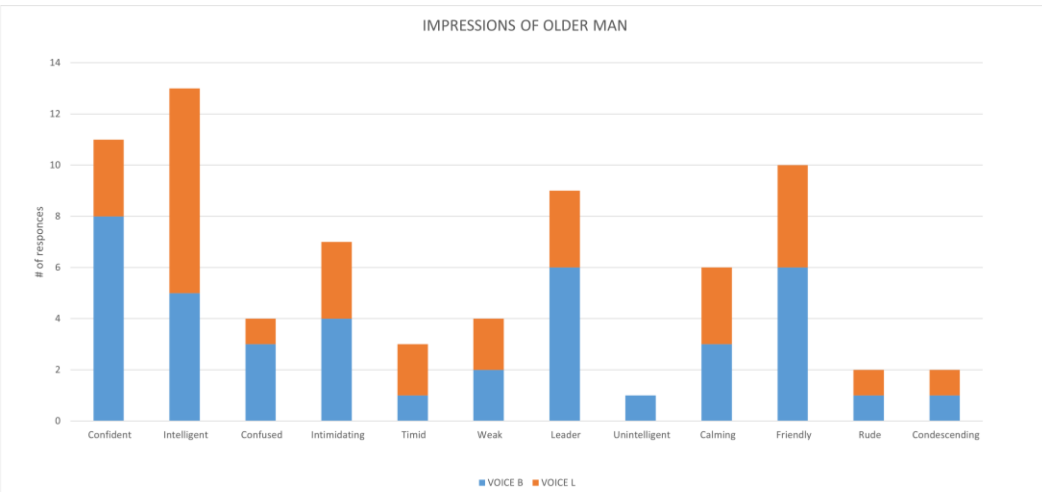

Figure 6: Count of each time participants referred to Male Three as one of the included impressions.

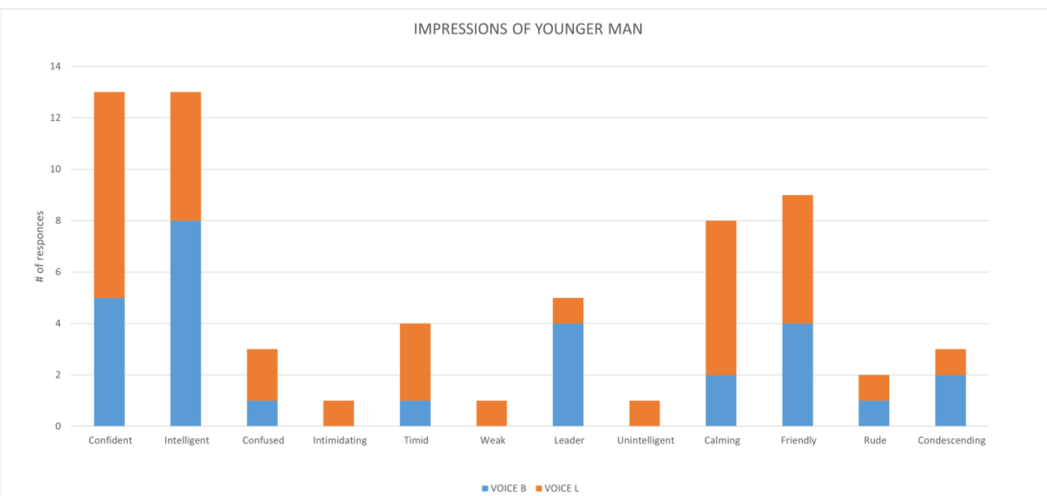

Figure 7: Count of each time participants referred to Male One as one of the included impressions.

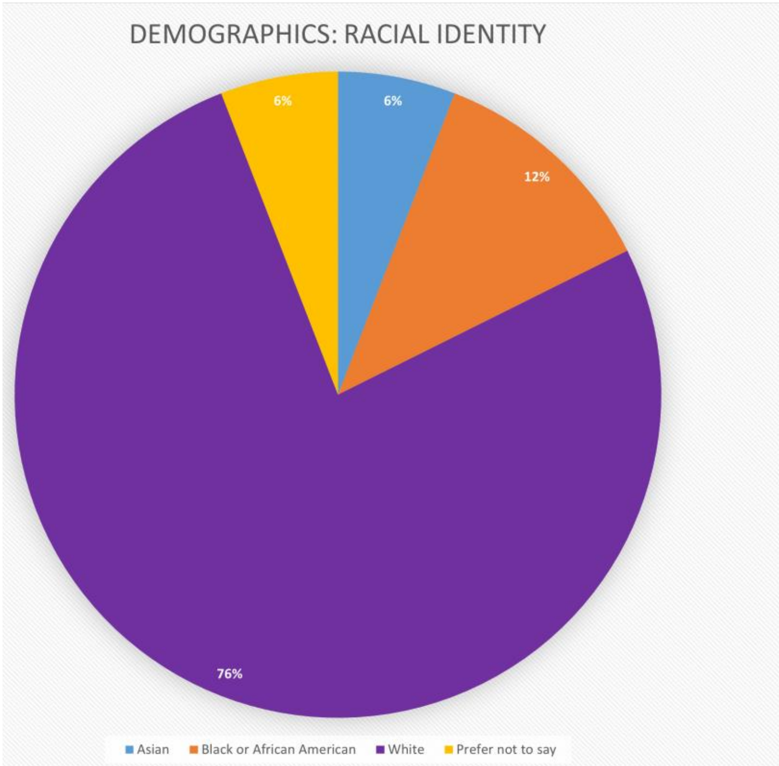

Figure 8: Reported racial identity of participants.

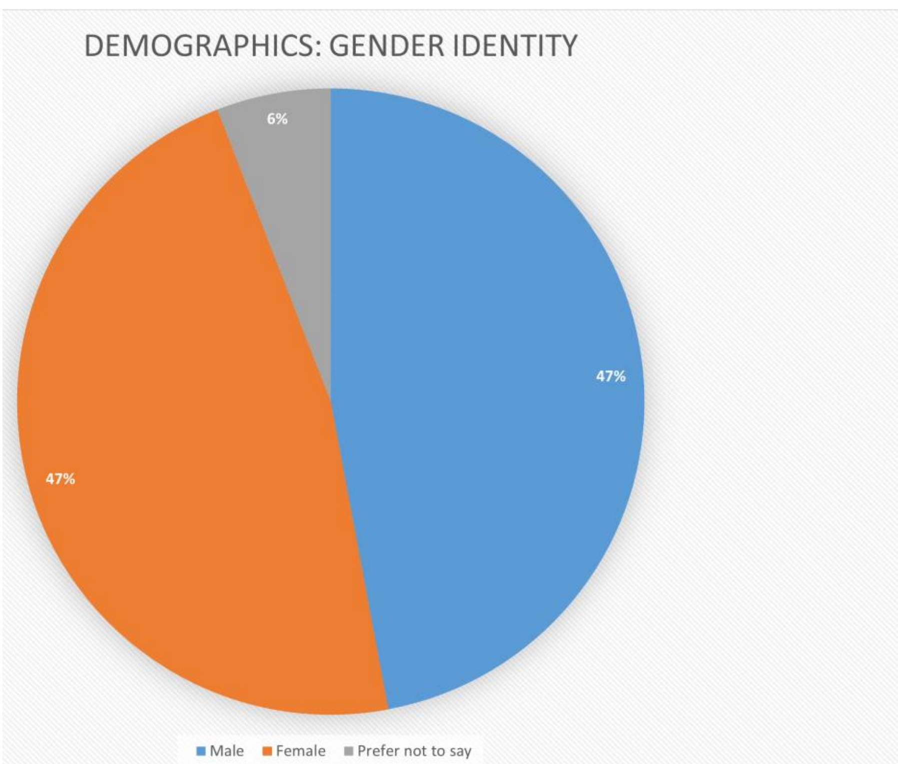

Figure 9: Reported gender identity of participants.

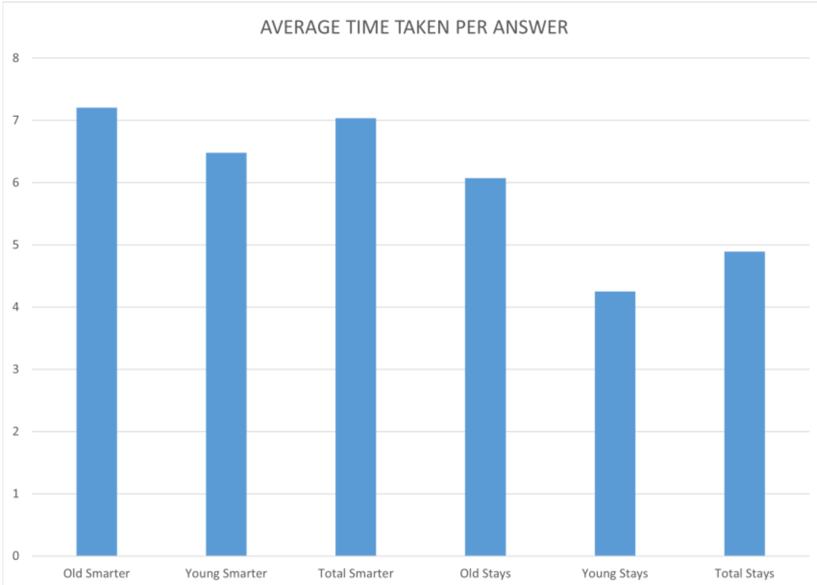

Figure 10: Average time taken to select which avatar appeared smarter and to select which avatar stays (in seconds).

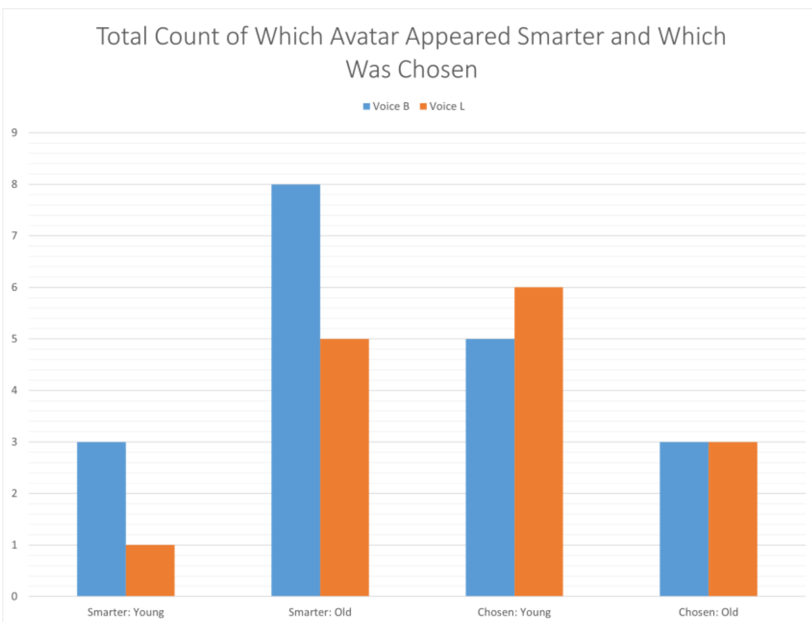

Figure 11: Count of total times each avatar was chosen for the two central questions.

REFERENCES

[1] Sun Joo Grace Ahn, Jeremy N Bailenson, and Dooyeon Park. Short-and long-term effects of embodied experiences in immersive virtual environments on environmental locus of control and behavior. Computers in Human Behavior, 39:235–245, 2014.

[2] Omri Gillath, Cade McCall, Phillip R Shaver, and Jim Blascovich. What can virtual reality teach us about prosocial tendencies in real and virtual environments? Media Psychology, 11(2):259–282, 2008.

[3] Erdman Palmore. The ageism survey: First findings. The gerontologist, 41(5):572–575, 2001.

[4] Xueni Pan and Antonia F de C Hamilton. Why and how to use virtual reality to study human social interaction: The challenges of exploring a new research landscape. British Journal of Psychology, 109(3):395–417, 2018.

[5] Robin S Rosenberg, Shawnee L Baughman, and Jeremy N Bailenson. Virtual superheroes: Using superpowers in virtual reality to encourage prosocial behavior. PloS one, 8(1):e55003, 2013.