Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

School of Engineering

Attention and Communication in Virtual Worlds: Interacting with Non-Player Characters in Virtual Reality

David Carroll

April 2022

A dissertation submitted in partial fulfilment
of the requirements for the degree of
MAI (Electronic and Computer Engineering)

Declaration

I hereby declare that this dissertation is entirely my own work and that it has not been submitted as an exercise for a degree at this or any other university.

I have read and I understand the plagiarism provisions in the General Regulations of the University Calendar for the current year, found at `http://www.tcd.ie/calendar`.

I have also completed the Online Tutorial on avoiding plagiarism 'Ready Steady Write', located at `http://tcd-ie.libguides.com/plagiarism/ready-steady-write`.

Signed: ______________________ Date: 19/04/2022

David Carroll

Abstract

Virtual Reality is fast becoming a useful tool for an increasing amount of applications. With the advancement of Virtual Reality technology, there is an increasing need for improved levels of believability and immersion in virtual environments. The inclusion of simulated humanoid crowds is one way in which the believability and immersion of a virtual world can be enhanced. Enabling a Virtual Reality user to interact with the Non-Player Character crowd members can increase their sense of presence within these environments. In this project, a method of creating simple, non-contact hand gesture interactions with these Non-Player Characters is investigated. Models of attention and communication are designed to provide the Virtual Reality user with the capability of believably interacting with members of a simulated crowd. The resulting design of these models succeed in providing a Virtual Reality user with a method of interacting with a simulated crowd in a way which allows for the variation of crowd responses. The design of these interactions created for this project allow for interactions to be tailored to suit the needs of different applications.

Acknowledgements

I would like to express my gratitude to my supervisor, Professor Carol O'Sullivan, for her continuous guidance and assistance provided throughout this project.

I also wish to extend my thanks to the staff of both the Trinity College Dublin School of Engineering, and School of Computer Science and Statistics for the support they provided me with this year.

Finally, I wish to acknowledge the support and inspiration provided by my parents, Bryan and Emer, and my girlfriend, Julia. I could not have completed this dissertation without them, and I am forever grateful for their understanding and the sacrifices they have made throughout my education.

Contents

List of Figures

Nomenclature

NPC	Non-Player Character
VR	Virtual Reality

1 Introduction

Virtual Reality (VR) is a technology that uses a head mounted display and computer graphics to simulate an immersive 3D environment. This VR environment can provide a greater sense of presence and immersion for a user when compared to environments displayed on traditional desktop user interfaces. This sense of immersion and presence provided by VR environments is contributing to the growing popularity and usage of VR technology.

VR has only been widely available to consumers for a relatively short period, but has seen rapid growth in popularity and investment, particularly in the entertainment industry. The market for VR technologies is expected to have a further major increase in growth in the near future. In addition to it's growth in the entertainment industry, VR has become an important topic of research. This research investigates the many possible applications for VR technology in areas such as education, medicine, science, socialising, building design and more. This research has led to the implementation of VR technology in applications such as medical training, and military and emergency response training.

Virtual Reality's major advantage over traditional user interfaces is it's ability to provide the user with a greater sense of immersion in a Virtual World. For many of these immersive experiences, a simulated crowd is used to improve the users sense of presence in the Virtual World. These simulated crowds can be used in applications such as games, teleconferences, chat rooms, and simulated training exercises. The simulated crowd is created by populating the Virtual World with Non-Player Characters (NPC) that can traverse the scene in order to further the users immersion and sense of presence in the virtual environment.

NPCs are a well researched component in traditional desktop interfaces, and have existed for quite some time. While the rest of the gaming industry has seen major growth in recent times, the development of NPCs has lagged behind. There has arguably been no need for further development of NPCs in traditional user interfaces, as users focus on the main character and are willing to ignore artefacts in fringe characters such as NPCs. However, there is a need to develop NPCs for the more immersive environments that VR interfaces provide. These NPCs can break a users immersion in a VR experience if they are not believable. There has been little research carried out in the area of creating interactive NPCs in a VR environemnt.

These VR environments offer new ways to interact with the Virtual World that are different from the traditional key press. The traditional key press interaction was limited to a set amount of predetermined actions. In VR, the user has the ability to interact with much more freedom of movement. This freedom of movement ensures the user is not restricted to a set of predetermined key press commands. The user now has the freedom to make hand gestures and use their natural gaze to interact with the environment. With users having more possible ways of interacting with the Virtual World, there is a need to research and develop how the user interacts with NPC members of a simulated crowd in VR.

This research project aims to investigate a means of interacting with NPC crowd members in Virtual Reality. This interaction can help to further develop and enhance the users sense of immersion and presence in a Virtual World. As the focus of this project is interacting with an NPC in VR, the methods of interaction selected are unique to VR, and not available to desktop interface users. The selected methods of interaction for this project are the users gaze and hand gestures. The users gaze and hand gestures are natural methods of communicating that the VR user should find intuitive. This selection of natural movements is aimed at furthering the users immersion and sense of presence in the Virtual World.

The proposed interaction with the NPC for this project is divided into two main steps. First, the user must gain the attention of the NPC using their gaze. Following this, the user communicates with the NPC using a hand gesture. This proposed interaction is investigated using an interactive VR environment created in a 3D engine. Following the investigation of this proposed interaction method, the results are used to define a simple framework for creating interactive NPC crowd members for a VR environment in a 3D engine.

The remainder of this paper is structured as follows: Chapter 2 'Background', provides background information for this project. The motivation for this project is outlined and relevant work in the domain of the project is reviewed.

In Chapter 3 'Design and Implementation', the design choices for the project are outlined and explained. The implementation of the research undertaken is also outlined. This implementation is broken down in to three main sections, 3.2.1: Virtual World, 3.2.2: Attention, and 3.2.3: Communication. The results of this research are used to define a simple framework for creating interactive NPC crowd members for a VR environment in a 3D engine.

In Chapter 4 'Results and Discussion', the results of the investigations and tests performed in this project are presented and discussed. A final Virtual World in which a VR user can interact with various crowd NPC's is also presented.

Chapter 5 'Conclusions and Future Work', consists of the conclusions drawn at the completion of this project, and possible areas of future work and improvements that can be made are discussed.

2 Background

2.1 Growth of Virtual Reality

Virtual Reality has seen a surge in popularity in recent years. Although the technology has been around for some time, recent advances in computing power and newly released content have helped to drive this technology's popularity with consumers and researchers. This new content, such as the groundbreaking video game Superhot VR [2], and the Massively multiplayer online game VRChat [3], demonstrate the new ways VR can allow users to interact with virtual environments that were not previously possible with desktop user interfaces. With affordable VR headsets being made more widely available to the public, this growth in popularity is expected to rise.

The huge growth seen in the investment of the development of VR technology has contributed to making VR more accessible to consumers. With Facebook's $2 billion purchase of Oculus in 2014 [4] leading to the commercial release of the Oculus Rift in 2016 [5], it is clear to see that large investors see a bright future for VR technology. Facebook further confirmed their belief in VR technology in 2021, when the company changed its name to "Meta Platforms" with the intention of developing a large scale metaverse [6] in which VR technology will play a vital role. This investment in VR technology is expected to further increase as the VR market is forecast to experience major growth in the near future [7].

The growth in popularity of VR technology will drive the need for further development of this technology.

2.2 User Immersion In Virtual Reality

With the expected need for further development of VR technology, it is important to note that VR is already a popular topic of research. Ongoing research has been carried out to investigate the useful applications for VR such as gaming, entertainment, and as an educational tool for medical, military and emergency response training.

Video games are a well established and rapidly growing industry with an estimated 3.24 billion gamer's across the globe [8]. As gaming is currently the most prominent use for VR, research comparing VR gaming to traditional computer monitor gaming can provide a very good indication of whether or not VR has an advantage over the traditional user interfaces. Shelstad et al. [9] investigated how VR affects the satisfaction levels of game users when compared to the traditional computer monitors. The results of this investigation showed that VR enhanced the users overall satisfaction and engrossment in the games they played. The participants in the study found that the games played were more "open to creative freedom" when playing the game using a VR headset. This may be because the user can be more actively involved in a VR environment and can interact in ways such as physically moving their head to direct the camera in the VR environment. The improved immersive capabilities of VR is further investigated in [10] where it is seen that VR technology can enhance the users spatial perception of 3D CAD models at different complexity levels when compared to desktop interfaces. These results confirm that user immersion in a Virtual World is increased when using a VR headset, when compared to a traditional computer monitor.

This increased sense of immersion for VR environments is leading to the technology being used as an educational tool. Izard, Santiago González et al. [11] investigated the use of VR as a training tool for medicine. The results of this research found that the VR technology used will raise the quality of the academic process and allow students to better manage their knowledge and practical skills. It was also discussed that VR technology can allow medical professionals to better understand the real world surgical processes by practicing in a virtual environment to become more familiar with the tools used. Simulations like these can provide a means of educating more students virtually, and can reduce the need for education and training in rooms with limited capacity.

This reduction in the need for education and training in real world environments is of particular interest when these real world environments have a degree of danger to them. Environments used to educate people such as soldiers and emergency response personnel present an inherent danger to the students. Much like flight simulation, it is desirable to create simulated environments for training students in potentially dangerous situations. A good quality simulation can allow a student to make potentially costly mistakes in a virtual world without there being any immediate safety concerns, or major financial repercussions. VR is beginning to open the door for many more simulated environments. The immersive nature of VR allows for a wide variety of training scenarios to be carried out in a safe environment. Research has been carried out investigating the potential that VR technology has for competency-oriented education and training of highly trained soldiers [12]. Alexander et al. [12] found that VR technology made it possible to add realism and interactivity to training environments. They concluded that with additional feedback on simulations, VR has a promising future as a tool in competency-oriented training for highly trained soldiers. This conclusion is evidence that VR is allowing users to interact with the virtual environment in a wholly new way when compared to traditional desktop user interfaces.

This new level of interaction that VR technology provides can be used in a wide range of applications and has been used as a tool to monitor human behaviour in crowds [13]. Simulations like this make researching topics such as human behaviour in crowds much more accessible and allows experiments to be run a number of times at a greatly reduced cost.

The results of all the research discussed in this section demonstrate how VR can provide a more immersive experience in Virtual Worlds and can help to simulate real life situations. This demonstrates a bright future for immersive simulated environments in VR and provides motivation to undertake research in the area.

2.3 Simulated Crowds in Virtual Worlds

A key component in many simulated Virtual Worlds is a simulated crowd. Crowd simulation plays a major role in many forms of visual media such as video games, feature films, architecture, and also in training simulations. While human crowd simulations are widely used in virtual worlds, they are the main element that is failing to enhance the users sense of presence [14]. Pelechano et al. [14] review the recent advances in immersive virtual crowd simulations and discuss areas that are in need of improvement to create a crowd that is believable to the VR user and will increase their sense of presence. One such area that would improve the sense of presence in the crowd simulation is the addition of models of attention that would provide more lifelike qualities to the crowd members [14]. Applying a model of attention to a character to make them look around the scene at interesting objects and at the user can increase the users sense of presence.

The varying density of a VR crowd can affect the user experience and behaviour [15]. As the density of the crowd increases, it has a significant negative impact on the VR users enjoyment. This is similar to responses of humans in real crowds where high crowd density causes an uncomfortable experience for the user. Dickinson et al. [15] found that crowd agent behavioural artefacts may have also contributed to this negative impact as the presence of agent behaviours that were perceived to be rude increased with crowd density. This research can inform the improved design of simulated crowds in virtual environments.

The addition of basic social interactions such as gaze and verbal salutations strengthened a users sense of presence in a simulated crowd in VR [16]. Kyriakou et al. [16] examined the attributes of an NPC crowd members behaviours that impact a user's experience in VR. Participants were presented with three scenarios with different crowd interaction levels.

These scenarios consisted of a crowd that ignored the participant, a crowd that had collision avoidance enabled, and a crowd with basic social interactions such as gaze and verbal salutations. The addition of collision avoidance contributed to the environment and the whole VR system appearing more lifelike and realistic to the participant. The addition of the basic social interactions further contributed to this realistic feeling for the participant.

It is clear that there is a need for the development of social NPCs in a simulated crowd in immersive VR to enhance the sense of presence that a user feels. This development of a social NPC requires the NPC to have a degree of believability.

2.4 Non-Player Characters (NPC)

2.4.1 NPC Believability

Recent development and growth in the video game industry has seen large improvements for both graphics and audio in virtual worlds. Unfortunately, the development of NPCs has not matched the growth of graphics and audio. While there have been improvements in the development of NPCs, the field as a whole has lagged behind the rest of the video game industry [17, 18]. Warpefelt et al. [17, 18] explores the believability of social NPCs in video games. It was found that NPCs are capable of being believable, particularly in social settings. The believability of an NPC in a social setting relies heavily on pre-scripted behaviours to allow the user to interact with the NPC [17].

An investigation into the extent to which an NPC appears as a person or if they're regarded as an object in virtual worlds was carried out [19]. This investigation discovered that players do not regard an NPC as a person in a virtual world. The NPCs were not viewed as social persons largely due to the lack of ability for emergent social behaviour. However, NPCs are afforded empathy by players. The empathy afforded to NPCs was not demonstrated by all human players as different players developed different relationships with NPCs.

This information confirms that it is possible to develop the social NPC needed to enhance the sense of presence a VR user feels in a virtual environment.

2.4.2 Uncanny Valley

As this project involves the use of humanoid NPC characters, the Uncanny Valley Effect [20] was researched to ensure the project is not negatively affected by unintentionally creating a sense of uneasiness in the VR user.

The design of a 3D character has an impact on the users level of bonding and interaction with the character [21]. Bouwer et al. [21] recorded the perceptions of audiences observing 3D character of various animation styles. It was found that realistically animated characters suffer negatively from the Uncanny Valley Effect if there are imperfections present in the 3D character or its animation. The Uncanny Valley Effect does not negatively affect the users bonding or interaction with a stylised character to the same extent.

Adamo et al. [22] investigated the perception of emotions in realistic and stylised characters. The findings of this investigation showed that people could better perceive the emotions of a stylised character compared to a realistic character.

The findings above inform us that hyper-realistic models are not needed to increase a users connection or level of bonding with an NPC.

VR technology does not provide much of an advantage for every day computing needs and is unlikely to replace traditional user interfaces for day to day work. However, it is evident from the research discussed in Section 2.2 that VR technology holds a major advantage over traditional desktop user interfaces for immersive simulated environments. It can also be seen that VR allows for a new level of user interaction with these virtual environments. This advantage and possibility for improved interaction, along with the projected growth of VR outlined in 2.1, provides motivation to research and develop user interactions in Virtual Worlds using VR.

Many immersive Virtual Worlds contain a simulated crowd. Section 2.3 describes how to design these simulated crowds to ensure the highest level of immersion for the user. The current failings of crowd simulations to fully immerse a user are also presented. Areas to help prevent these failings are identified, with one of these areas being the improvement of the social presence of crowd members. An improved social presence of crowd NPCs will enhance the VR users immersion in the virtual world.

For the users immersion to be enhanced by the social NPC, the NPC needs to have a level of believability. Section 2.4 presents research on how to ensure believablity of these social NPCs. This research will allow for the development of a social NPC that will create a more believable crowd simulation.

Overall, this research provides motivation for the creation of a social NPC that can be used to simulate a virtual crowd to enhance the users VR experience and immersion in the virtual environment.

3 Design and Implementation

3.1 Design

The aim of this research project is to create an NPC that will enhance a VR user's sense of immersion in a simulated crowd. As outlined in Chapter 2, there is clear motivation for this research.

The design of this project centers around the creation of a simple NPC that a VR user can interact with. This project will focus on enabling the VR user to preform a non-contact interaction with an NPC through the use of their gaze and hand gestures. This interaction can be broken down into two main steps, attention, and communication. The decision to focus on non-contact interactions was made due to the difficulty associated with creating a believable interaction involving collisions between the VR user and the crowd members. Unlike in traditional media, in VR, the user's character cannot be forced to change location or fall to the ground should this be required by a collision. Creating believable collisions in VR is a separate topic of research and beyond the scope of this project.

Upon completion of this project, a VR user should be able to gain the attention of an NPC and communicate with the NPC in some way that enhances their sense of immersion in the Virtual World. To create this interactive crowd NPC, there are a number of prerequisite elements that need to be implemented.

3.1.1 Prerequisites

A 3D virtual environment, in which a VR user can be placed amongst a crowd, must be created to investigate the interaction between the VR user and the crowd NPCs. For this project this virtual environment is to be created using a 3D Game Engine.

The simulated crowd created for this project is to be quite basic, but closely follow the requirements of immersive crowd simulation mentioned in Section 2.3. The crowd NPCs must be capable of dynamic pathfinding, and traverse the scene in a believable way. The crowd NPCs must also avoid collisions with the VR user in order to ensure a sense of immersion. As described in Section 2.4.2, these NPCs are not required to be hyper-realistic characters in order to create this sense of immersion for the user. The dynamic pathfinding capabilities of the crowd can be tested in a simple 3D environment with obstacles. This same scene can be used to test the crowd members' collision avoidance with the VR user.

Simple animations are to be applied to the crowd NPCs to increase the user's sense of immersion. This crowd animation is to be basic, as the main focus of this project is to investigate the user's interaction with the NPCs. Crowd animations are a well established area of research, and complex crowd animations in this project will not contribute any meaningful advancements to the research topic of crowd animation. Therefore, basic crowd animations are deemed appropriate for this project and its timeline. The crowd animator created for this interactive crowd must include a state for scripted reactions to the VR user's communication. This pre-scripted animation can enhance the believability of the NPC in the crowd as discussed in Section 2.4.1.

With these prerequisite elements implemented, the attention model for this project's NPC can be developed.

3.1.2 Attention

The attention model for this NPC is designed so that VR users can interact believably with the NPC. For this project the VR user gains the attention of the NPC through the use of their gaze. This attention model design has constraints to ensure the VR user can only interact with NPCs who's attention can be realistically obtained.

The first constraint is that the user should not be able to gain the attention of an NPC that is not facing them. It is unrealistic that the VR user could use their gaze to interact with a crowd member that is facing away from them. This first constraint can be tested in a simple environment with stationary NPC crowd members.

Secondly, the VR user should not be able to gain the NPCs attention immediately. The VR user must first engage the NPC they are facing with their gaze. When a VR user first engages an NPC, the NPC should return the user's gaze. If the user holds this gaze with an NPC for a set period of time, they gain the NPC's attention. Once the NPC's attention is gained, they should stop traversing the environment and face the VR user. The VR user is then in a position to communicate with the NPC. This second constraint can be tested in a simple environment with a single dynamic NPC.

The NPC's attention is held until the VR user averts their gaze. This gives the user a sense of control over the virtual environment and increases their sense of immersion within it. For this project it was also decided that while a VR user holds the attention of an NPC, all other NPCs should be ignored. This is so that the VR user can complete their interaction with the NPC who's attention they have gained.

Once the VR user gains the attention of an NPC, they are in a position to begin the communication step of the interaction.

3.1.3 Communication

For this project, the method of communication with an NPC consists of a simple waving hand gesture performed by the VR user. This method of communication was chosen as it is an input method unique to VR and cannot be performed with a traditional key press. VR handheld controllers will be used to create this waving gesture. The design choice to use VR handheld controllers for this project was made to allow the user to make use of the controller's joysticks to explore the environment. While this may mean that the waving motion is less natural for the VR user, it will allow them to freely traverse the environment. This is deemed to be an appropriate trade-off for this project as the user's immersion in the Virtual World is not likely to be broken by the use of the VR handheld controllers.

In order to communicate with the NPC using hand gestures, methods of recognising these hand gestures as input values are required. For this project only one gesture is used as an input to communicate with the NPC. However, it is important to test whether multiple gestures can be recognised as different input values in order for this communication with the NPC to be diversified in the future. The recognition of hand gestures as inputs can be tested in a simple scene containing only the VR user.

Following the successful recognition of a VR user's hand gesture, the gesture can be used as an input to communicate with the NPC. The NPC responses to this communication are pre-scripted, but can be varied as needed to suit the application being developed. For this project the NPC will respond to the VR user interaction with a waving animation. This is a sufficient response to demonstrate a VR user's interaction with an NPC in a simulated crowd.

3.2 Implementation

3.2.1 Virtual World

The Virtual World in this project is created using the Unity Game Engine [23]. Unity was chosen for this project as it is considered an industry standard 3D Game Engine. The use of Unity ensures that this research project is relevant to the state of the art products being created today. Unity also provides a very good architecture with improved support for VR platforms, making it an ideal 3D Game Engine for this VR project.

Dynamic Pathfinding and Collision Avoidance

This project requires a basic crowd of NPCs with dynamic pathfinding and collision avoidance capabilities in order to create a believable crowd for the VR user [16]. To create a crowd with these capabilities, Unity's NavMesh Agent [24] component is attached to simple characters in a virtual environment. This NavMesh Agent component uses Artificial Intelligence to allow the characters to navigate the environment while avoiding other NavMesh Agents and environment objects.

To control the overall movement of the crowd, each agent is provided with a goal location. A number of these goal locations are present in the scene. When a NavMesh Agent reaches their goal location, they randomly select a new goal location from those available in the environment. Unity's NavMesh Agent component allows the character to move to this next location with dynamic pathfinding and collision avoidance capabilities.

This movement is controlled with a simple C# script that provides a new goal location to a NavMesh Agent when the character is a certain distance from its current goal location. This C# script is applied to all characters in the scene to create a constantly dynamic crowd.

This dynamic crowd implementation is tested using very simple 3D capsule models to represent NPC crowd members. These NPCs are used in a basic environment, with obstacles placed between the goal locations, to test the dynamic pathfinding and collision avoidance capabilities of the crowd.

Humanoid Crowd Characters

Using humanoid 3D models to represent the NPCs in the virtual crowd is an important part of this project. As indicated in Section 2.4.2, these 3D character models do not need to be hyper realistic to be believable. To create this humanoid crowd, the 3D capsule model from Section 3.2.1 is replaced with a 3D Woman character model. This creates a crowd of humanoid NPCs with dynamic pathfinding and collision avoidance capabilities. The humanoid Woman character used for this project was provided by Trinity College Dublin.

A simple animator controller component was added to the humanoid NPCs to add animation to the crowd. The animator controller for this project provides the characters with an Idle, Walking, and Waving state. Each state simply has one animation that is triggered when the animator controller transitions from one state to another.

The Idle state is entered when the NPC stops moving around the scene. The animator controller changes state from Idle to Walking when the NPCs begins to move towards its goal location. These two states make up the entirety of the crowd animations in this project.

The Waving state controls the NPC's pre-scripted response to the user's waving hand gesture. This state is entered under the conditions that the user currently has the NPC's attention, and that the user has performed a waving hand gesture. This state is exited when the user averts their gaze from the NPC, and the state machine returns to the Idle state. The Crowd Animator controller state machine used in this project is presented in Figure 3.1 below.

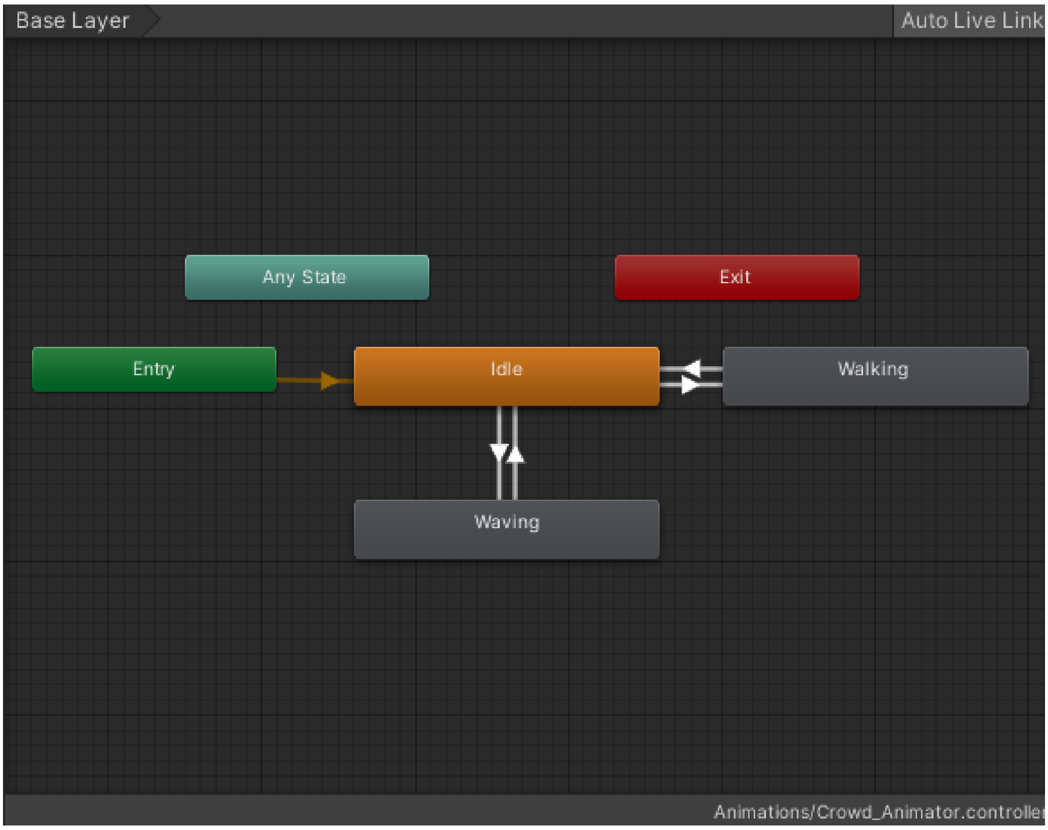

Figure 3.1: Crowd Animator Controller State Machine

The Walking animation for this Crowd Animation state machine was taken from the Standard Assets [25] package in the Unity Asset Store [26]. The Idle and Waving animations were obtained from Mixamo [27].

Virtual Reality User

As mentioned at the beginning of this section, Unity provides very good support for the development of VR projects. The Unity provided XR Origin [28] and XR Interaction Manager [29] were used in conjunction with the Oculus Quest 2 [30] VR system to enable a VR user to be added to the Virutal World in this project. A Locomotion [31] system was added to the project to enable the user to explore the virtual environment through the use of the handheld controllers.

3.2.2 Attention

Selecting NPC for Interaction

In order to create the attention model for this project, the VR user must be able to use their gaze to select an NPC to interact with. In order to make this selection possible, an XR Ray Interactor [32] component is added to the XR Origin's Camera Offset. There is also an XR Simple Interactable [33] component added to the crowd NPC. These components allow for collisions between the VR user's gaze direction and an NPC to be detected. This collision detection is used to select the crowd NPC the user wishes to interact with.

NPC Facing VR User

The first constraint that must be added to the NPC attention model in this project is that a player cannot interact with an NPC that is not facing them. The respective forward facing vectors of the VR user and the NPC they are trying to select can be used to implement this constraint.

We can use the Dot Product of these respective forward facing vectors to determine if the NPC is facing the VR user.

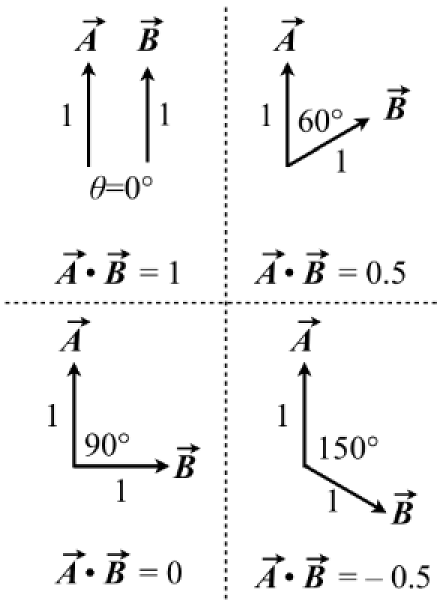

Figure 3.2: Dot Product Between Two Vectors [1]

Figure 3.2 illustrates that for vectors facing the same direction, the Dot Product returns a positive result, and a negative result is returned for vectors facing in different directions.

Using this information, it can be determined that the selected NPC is facing the VR user when the Dot Product of their respective forward facing vectors returns a negative result.

This result is used to set the value of a boolean to describe when an NPC is facing the VR user. This boolean enables the implementation of the first constraint for the NPC attention model created for this project. The NPC cannot be selected if the result of the Dot Product calculation returns a positive value.

Returning the User's Gaze

The second constraint for the NPC attention model is that a VR user should not immediately gain the NPC's full attention. To engage the NPC, the VR user makes use of their gaze. Once engaged, the NPC returns the VR user's gaze as a method of informing the user that the interaction process has begun. This use of the NPC's gaze can play a role in enhancing the VR users sense of immersion [14].

For the implementation of this, each NPC is assigned a target point in space which is the focus of their gaze. When the user first look's at an NPC that they are facing, the target point of the NPC's gaze is translated to the position of the VR users headset in the virtual world. This causes the NPC to meet the users gaze before the user has fully gained the attention of the NPC.

Gaining the NPC's Full Attention

Once the VR user has held the selected NPC's gaze for long enough, the NPC's attention will be on the VR user. At this point the NPC should come to a stop and face the VR user. To implement this, the NavMesh Agent's isStopped boolean is set to true. This disables the NPC's NavMesh Agent component and brings the NPC to a stop.

Once the NPC is stopped, the NPC must turn to face the user. To do this the NPC's forward vector ($\vec{a}$) is rotated an angle (θ) towards the vector between the NPC and the VR user ($\vec{b}$). To calculate $\vec{b}$ we subtract the world position of the NPC from the VR user's world position.

The formula for obtaining θ is

$$\theta = \arccos \frac{\vec{a} \cdot \vec{b}}{|\vec{a}||\vec{b}|} \tag{3.1}$$

The NPC is rotated about its up-axis by the angle θ to face the VR user. Once this is completed the user has gained the full attention of the NPC who then waits for a communication input from the VR user.

Maintaining Connection

For this project, the NPC's attention is held until the VR user averts their gaze. This means no other NPC is allowed to interrupt the connection between the VR user and the NPC they are interacting with. The implementation of this constraint was achieved by disabling the collider components of all NPC's that are not involved in the interaction. This disabling of the other crowd NPCs collider component's ensures that the VR user's XR Ray Interactor cannot detect other crowd members as they pass by.

Ending Interaction

The VR user ends their interaction with an NPC by simply averting their gaze. Once the VR user moves their gaze away from the NPC they are interacting with, the isStopped boolean of that NPC is reset to false, and the NPC continues along the path they were following before the interaction.

When the interaction is ended, the collider component's of all crowd NPC's are enabled and the VR user is free to interact with any viable NPC.

3.2.3 Communication

Enabling a VR user to communicate with an NPC in a Virtual World requires the recognition of the user's hand gestures as an input. This input must trigger a response from the NPC to complete the interaction.

Gesture Recognition

For this project, user gestures are recognised by tracking the position of the handheld VR controller in space. The user indicates that they are attempting to gesture to the NPC by pressing a button on the handheld VR controller. While this button is pressed, the controllers position in space is tracked, and this tracking is used to create a point cloud. The button press is used to ensure point clouds are only created when a user is attempting to gesture to an NPC.

This user created point cloud is compared to a pre-saved set of gesture point clouds through the use of an out of the box point cloud gesture recogniser [34]. This gesture recogniser returns the name of the pre-saved gesture that the user's input most closely resembles, along with a score for the recognition of the gesture. This returned value is used to ensure gestures are ignored if their recognition score is below a certain threshold.

When a user gesture is recognised and has an adequate recognition score, it is used as an input to communicate with the NPC.

NPC Response

Each of the VR user's gesture inputs has a pre-scripted NPC response. These pre-scripted responses are accessed by changing the state in the Animator Controller described in Figure 3.1. The NPC only responds to the VR user's hand gesture if the VR user has gained the NPC's attention.

The implementation of this gesture response is limited to a simple wave animation response in this project. These response animations could be chosen from a random pool of animations to create a sense of variety in the crowd and enhance the users sense of immersion in the Virtual World. The single waving hand gesture used in this project was deemed adequate to investigate the possibility of interacting with an NPC in VR.

4 Results and Discussion

4.1 Results

4.1.1 Virtual World

A Virtual World populated by a humanoid crowd, in which a VR user is free to explore, was successfully created for this project.

The Unity NavMesh Agent [24] successfully enabled a crowd of NPC's to navigate a simple virtual environment. Figure 4.1 shows the testing of the crowd's dynamic pathfinding and collision detection capabilities in this environment.

The resulting virtual environment populated with an animated Humanoid Crowd can be seen in Figure 4.2a. Figure 4.2b depicts the Virtual World from the VR user's perspective.

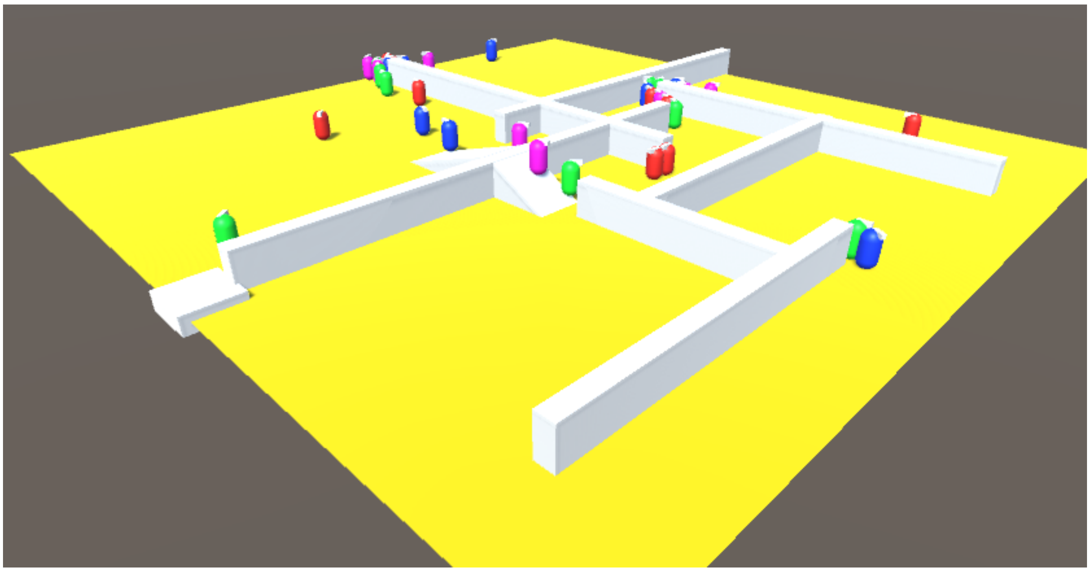

Figure 4.1: 3D Capsule NPC NavMesh Crowd

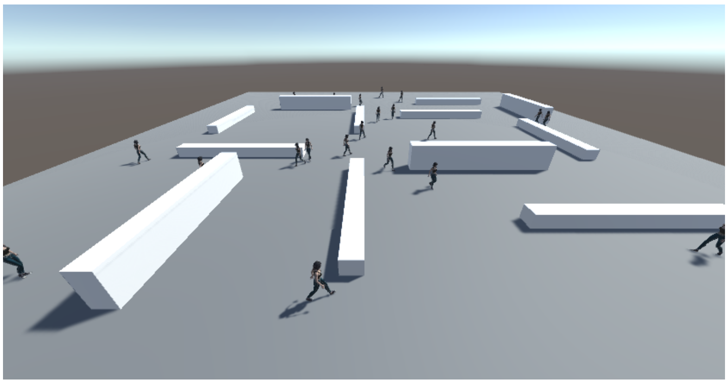

(a) Humanoid NPC NavMesh Crowd

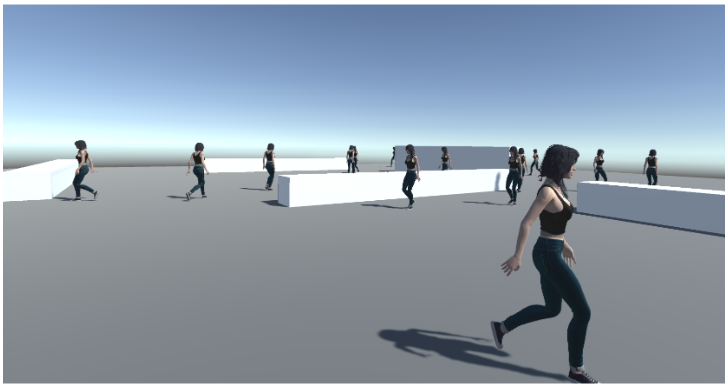

(b) VR User Exploring Virtual World

Figure 4.2: Humanoid Virtual World

4.1.2 Attention

The attention models developed for the crowd NPC in this project successfully allows the VR user to gain the attention of an NPC using their gaze. The two constraints for this attention model outlined in Section 3.2.2 were successfully implemented. The first constraint ensures that the VR user cannot interact with an NPC that is not facing them.

The results of the implementation of this constraint are displayed in Figure 4.3. The XR Ray Interactor component is displayed in Figures 4.3b and 4.3c to illustrate where the VR user's gaze is focused in the environment when testing this implementation. The boolean values returned by this test are displayed in Figures 4.3d and 4.3e.

The second constraint outlined in Section 3.2.2 ensures that the NPC's attention is not gained immediately by the VR user. Figure 4.4 illustrates the process of gaining an NPC's attention in this project. When the VR user first engages with the NPC, the NPC returns the user's gaze as displayed in Figure 4.4a. After the NPC's gaze has been held for long enough, the NPC stops to face the user who is then afforded the NPC's uninterrupted attention as seen in Figures 4.4b and 4.4c. When the VR user averts their gaze from the NPC, the NPC resumes its traversal of the environment.

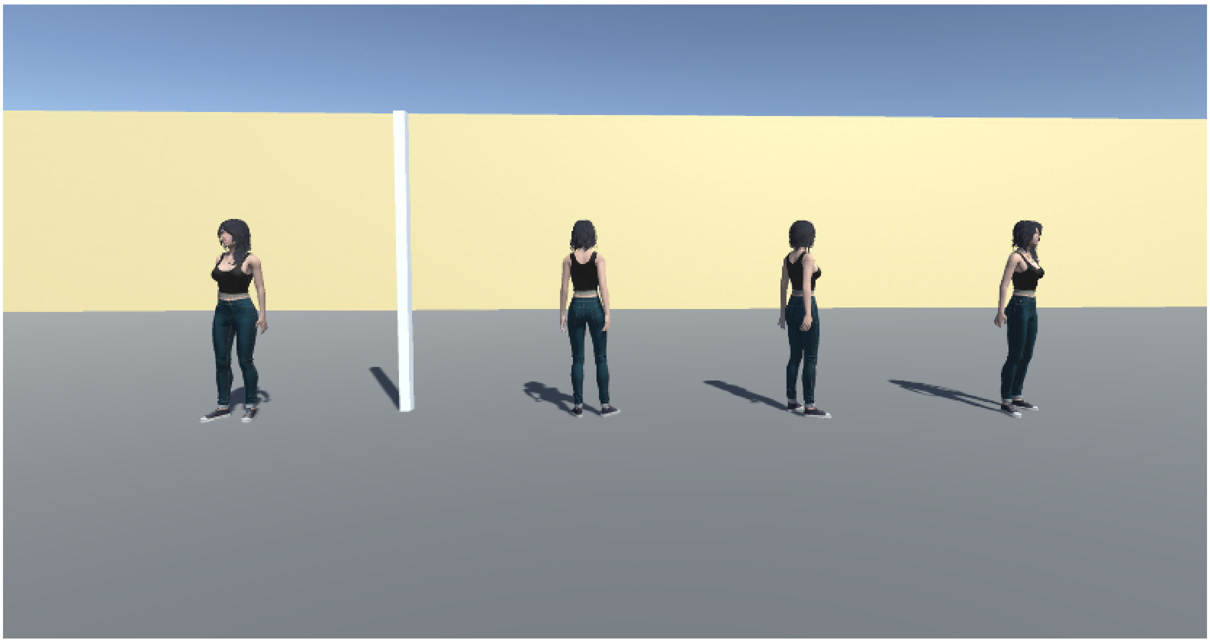

(a) NPC Facing User Test Environment

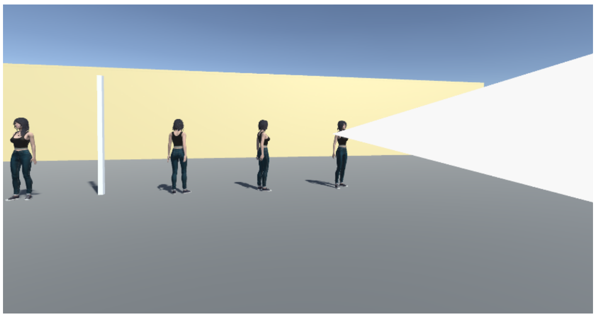

(b) Is Facing User

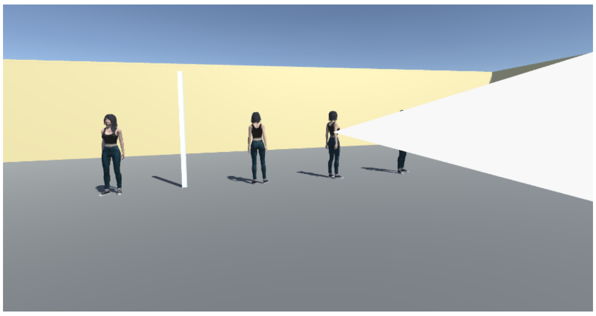

(c) Is Not Facing User

(d) Results of (b)

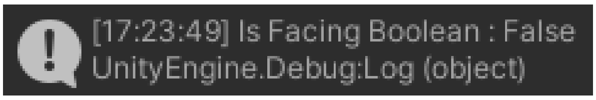

(e) Results of (c)

Figure 4.3: Selecting an NPC for Interaction

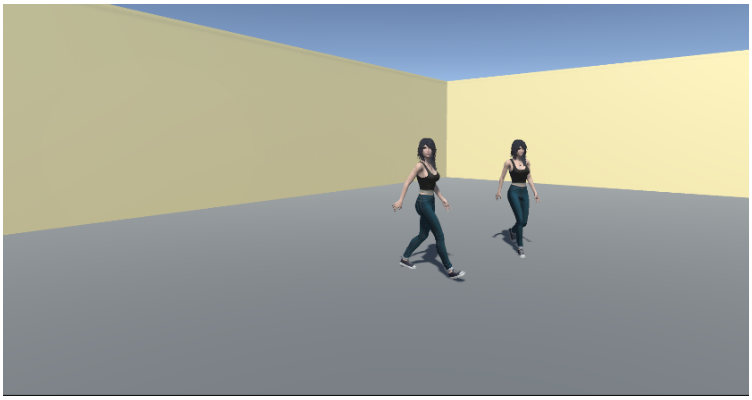

(a) NPC Meets User Gaze

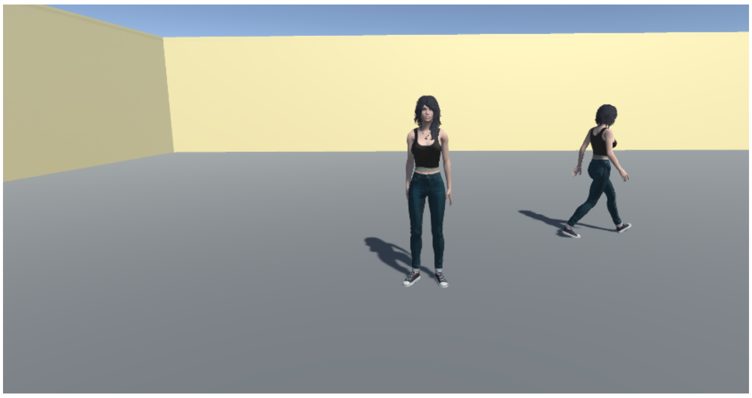

(b) NPC Stops to Face VR User

(c) Uninterrupted Attention

(d) NPC Resumes Traversal of Scene

Figure 4.4: Gaining attention of NPC

4.1.3 Communication

The results of the gesture recognition method employed in this project are displayed in Figure 4.5. These results demonstrate the use of point cloud gesture recognition for four different VR user hand gestures. The set of four hand gestures were pre-saved for the testing of this gesture recognition. These pre-saved gestures consisted of a V gesture, an R gesture, a Waving gesture, and a beckoning gesture labelled ComeHere. The recognition of these gestures was tested in a simple environment where the name of the gesture most closely resembling the point cloud created by the VR user was displayed. For each pre-saved gesture, the point clouds created by the VR user were correctly classified by the point cloud gesture recogniser [34].

The result of the correctly functioning gesture recognition allows the VR user to communicate with a Humanoid NPC in a very basic crowd. Figure 4.6 shows the VR user communicating with an NPC through the use of the Waving hand gesture.

All of the results presented above were used to develop a more advanced virtual environment with a wider variety of NPC's. This virtual environment can be seen in Figure 4.7. This scene was developed using the successful implementation of the design set out for this project, and allows the VR user to interact with a selection of six different NPC models. The environment was created using the Low-Poly Simple Nature Pack [35] available in the Unity Asset Store. The character models used for this environment were taken from Mixamo [27]. The resulting environment displayed in Figure 4.7 can also be seen in a real time video demonstration [36].

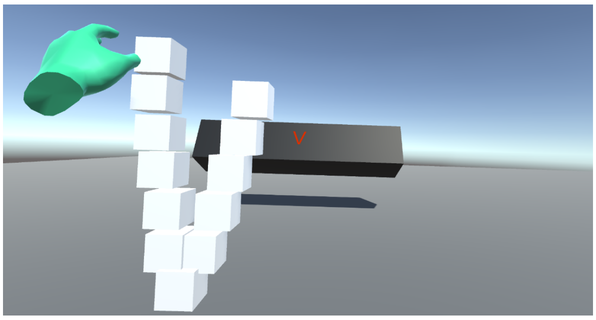

(a) V Gesture

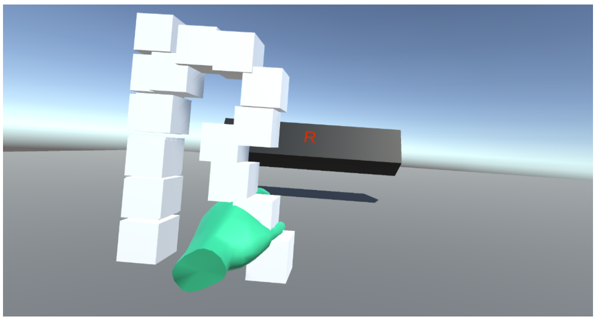

(b) R Gesture

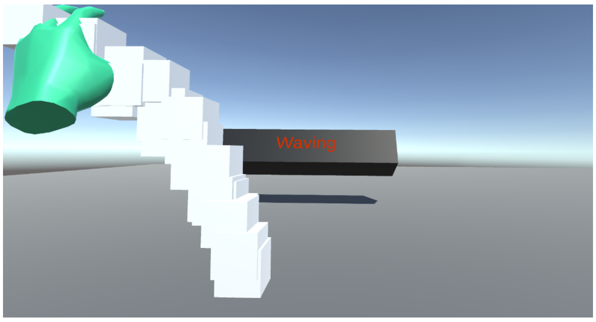

(c) Waving

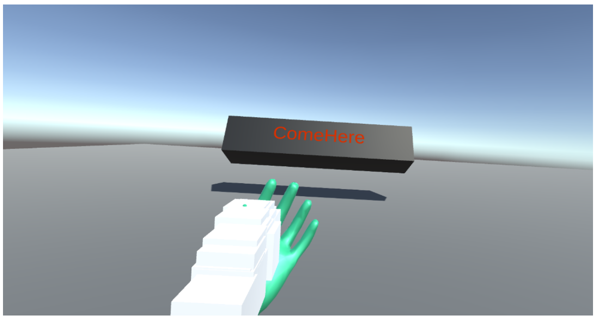

(d) Come Here

Figure 4.5: Point Cloud Gesture Recognition Test

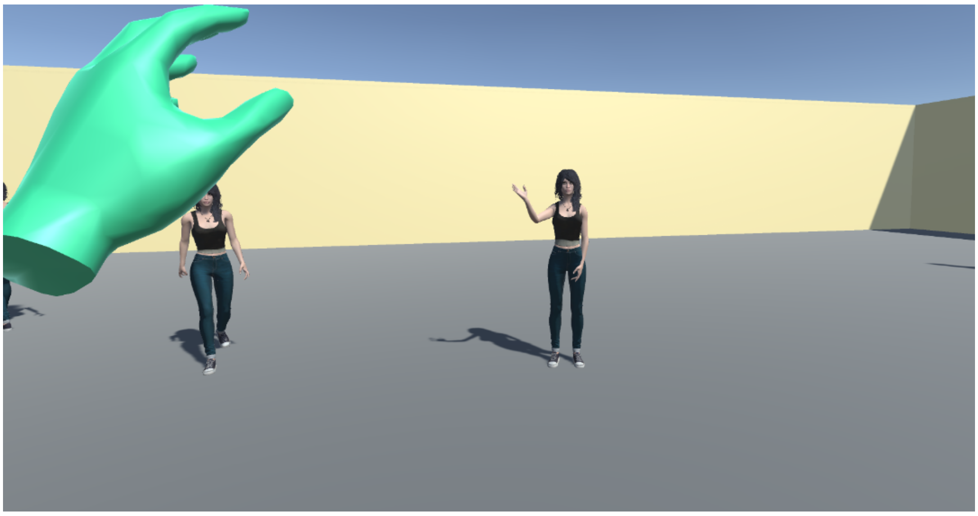

Figure 4.6: NPC Communicating With VR User

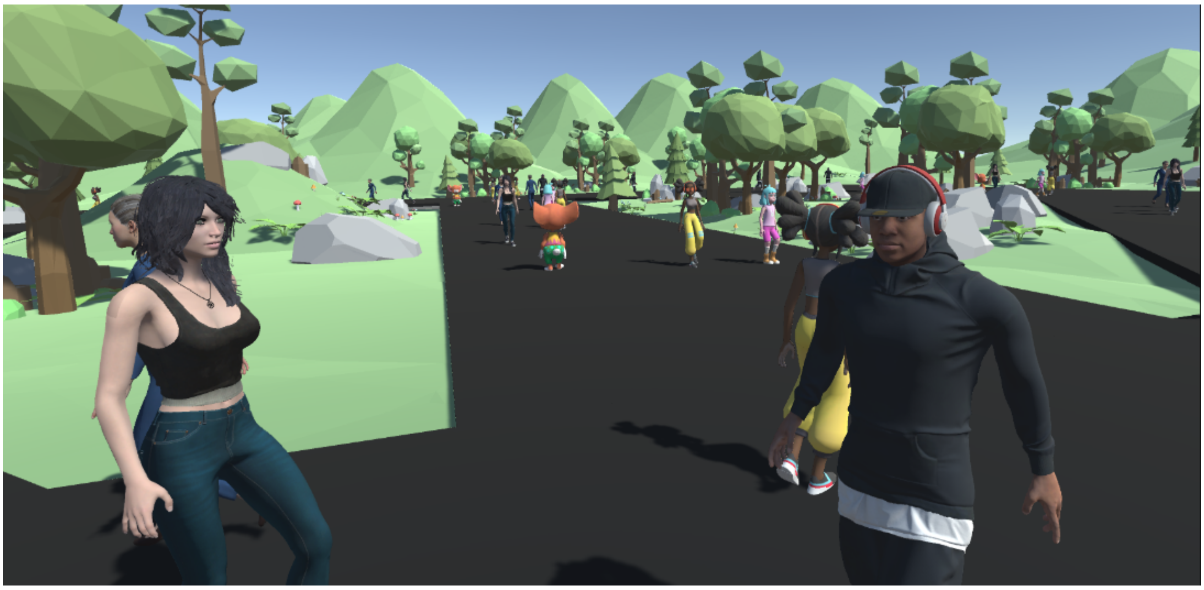

(a) Virtual Environment

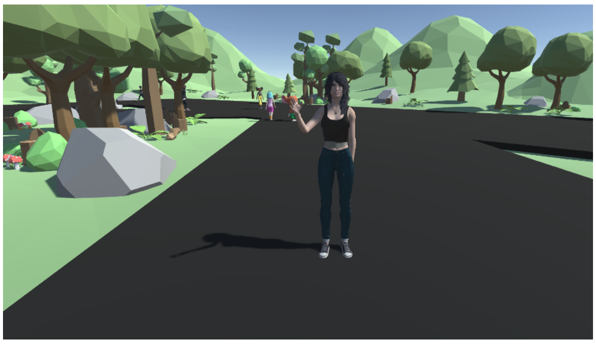

(b) Woman NPC Interaction

(c) Michelle NPC Interaction

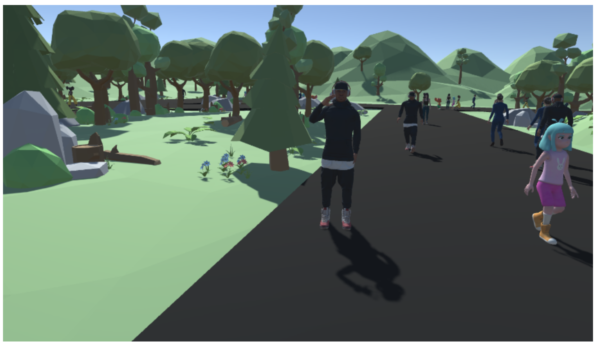

(d) James NPC Interaction

(e) Mousey NPC Interaction

(f) Amy NPC Interaction

(g) Louise NPC Interaction

Figure 4.7: Advanced Virtual World With Variety of NPC Models

4.2 Discussion

The simple virtual environment displayed in Figure 4.2 provides an adequate crowd simulation for the VR user to be immersed in. This simple crowd design can be scaled up or down quite easily, but consideration must be given to the effects that the density of the crowd will have on the VR users experience when creating the simulation. While there is room for improvement of this simulated crowd, it allows the VR user to test interactions with individual crowd members. The crowd could be improved with simple tweaks such as the addition of a strafing animation, or the reduction of footsliding present in some of the models. However, these improvements are not vital to this project, where the focus is on the VR user's ability to interact with crowd members.

The results of the attention model implemented in this project are in line with the desired outcome of the designed attention model. The results in Figure 4.3 are significant as they demonstrate that the VR user cannot gain the attention of an NPC who cannot see the VR user. This is very important in maintaining a realistic experience for the VR user, and enhancing their sense of immersion in the virtual environment.

The full functionality of the attention model implementation is shown in Figure 4.4. This result is important as it demonstrates that it is possible for a VR user to gain an NPC's attention in a Virtual World. Ensuring that the NPC comes to a stop and faces the player as seen in Figure 4.4b, provides a very good method of informing the user as to when they have gained the full attention of an NPC.

Providing the VR user with control over when this interaction is ended aids in the avoidance of creating an uncomfortable feeling for the VR user. If the NPC's were afforded control of the interaction, the VR user may become uncomfortable in the crowd. This could possibly have a negative impact on the VR user's experience in the virtual environment. In this project, the VR user is provided with control of these situations to give them a feeling of power and safety in the virtual environment. This limits the degree to which a VR user's experience and sense of immersion can be negatively impacted by unfriendly NPC behaviours.

This resulting attention model is limited to a single NPC reaction type as shown in Figure 4.4. Varying the NPC's reaction to a VR user's gaze could further enhance the VR user's experience in the Virtual World.

Providing the VR user with an uninterrupted connection with an NPC is a design choice that was revealed to be beneficial to the attention model. This design choice was made with consideration given to real life interactions in a crowd, where a passing crowd member does not usually prevent two people from interacting. Without this uninterrupted attention in the virtual world, it proves very difficult to communicate with an NPC in the simulated crowd.

This uninterrupted connection also serves to further the VR user's sense of control in the virtual environment.

The results for communication with an NPC provide evidence that it is possible for a VR user to interact with an NPC in a Virtual World. The resulting communication in this project is limited to that of a simple waving interaction. This communication can be improved by varying the NPC's reaction to the user input. Creating a pool of pre-saved reaction animations, and applying them to different NPCs in the environment could provide a greater sense of varying personality types amongst the crowd NPCs. This variation of personalities could enhance the VR user's experience.

The gesture recognition method implemented in this project functioned correctly and allowed the user to provide multiple different gestures as inputs to the communication with an NPC as illustrated in Figure 4.5. While successful, there are limitations to the method of gesture recognition implemented in this project. In this project the user is limited to creating gestures with one hand at a time. This limits the types of gestures available to the VR user. This limitation may make this implementation insufficient for virtual applications that require additional complex VR user gestures.

This gesture recognition method is further limited by the need for separate pre-saved gestures for left handed and right handed versions of the same gesture. For example, there is a need for separate pre-saved Waving gestures for the left hand and right hand in this project. This is due to the differences in the point cloud shapes created by these gestures.

Despite these limitations, the gesture recognition method implemented for this project provides the VR user with an acceptable method of interacting with an NPC that is not available to traditional desktop interface users. Figure 4.6 illustrates the communication between the VR user and an NPC. The waving gesture used to create this interaction provides evidence that a VR user's gesture input can elicit a response from an NPC in a virtual crowd.

The design of this project allows for the VR user's gesture input and the NPC's reaction animation to be tailored to suit different applications, such as video games or teleconferences. The possibility for the variation of this interaction is a positive result for the design of this project as it provides evidence that the design is not limited to one type of application.

5 Conclusions and Future Work

5.1 Conclusions

Overall, this project was successful in its aim of creating a means to allow a Virtual Reality user to interact with NPC crowd members in a virtual environment. The resulting design of this project allows for customisation of the interaction to suit specific applications. When following this design, there is the possibility for variation of the VR user interaction inputs and personalities of the crowd NPC's.

The design of this interaction complies with many of the guidelines for creating believable and immersive crowd simulations for VR outlined in Chapter 2. This compliance with previous research ensured the success of this project in creating a customisable and immersive VR experience. Although the results of this project are limited to one gesture, the implementation appropriately demonstrates the possibility of a customisable VR user interaction with an NPC that is not available to traditional desktop users. This customisation can lead to the creation of a crowd populated with NPC's of varying personality types. This will aid in enhancement of the VR user's immersion in a virtual world.

The design of this interactive NPC can be easily implemented in existing VR applications requiring crowds such as entertainment, teleconferencing, training simulations and more. This would further enhance the VR users experience in these applications and is very unlikely to reduce their level of immersion. This ease of implementation makes the addition of VR user interactions with crowd members useful to todays state of the art VR applications. Although this project was executed with limitations for the user interaction, the design allows for seamless further development of this interaction.

A survey to examine the difference, if any, in the levels of VR user satisfaction when placed in a crowd environment where interactions are enabled, when compared to user satisfaction in a non-interactive crowd would be a good addition to this project. It was not possible to carry out this survey as the required survey conditions could not be provided due to the ever changing status of the COVID-19 pandemic. The need for use of in person testing and the sharing of VR headsets meant that it was not possible to undertake this survey. This survey can be carried out in the future and can provide valuable information on the effect the results of this project have on a VR users sense of immersion in a virtual environment.

5.2 Future Work

Future work in this area could examine the use of verbal interactions with NPC's to gain their attention. This would provide the user with the ability to gain the attention of an NPC that is facing away from them as well as create variation in the attention model. The built in microphone in many VR headsets could be used to undertake this future work.

Additionally, future work could focus on extending the attention model so that NPC's in close proximity to the user's selected NPC are affected by the user's attempt to begin an interaction. This could be quite an important extension to examine if carrying out work in the area of the use of verbal interactions with NPC's. This work would allow for a more realistic sense of presence for the VR user in a virtual crowd, and would also be relevant to crowd simulations on traditional desktop interfaces.

This variation in the attention model could be further developed by introducing different ways of indicating that the user has gained the attention of the NPC. One such variation could be to use the NPC's facial animation to indicate that they are paying attention to the user. This can ensure that not every interaction is the same and can provide a sense of variation in NPC personalities in the crowd. Another very simple way to vary this attention model is to vary how difficult it is to gain different NPCs' attention. By varying the time that the VR user must hold an NPC's gaze to gain their attention, the illusion of friendly and rude NPC's can be created.

Investigations could also be carried out in to the possibility of events in the environment catching the attention of NPC's. These could be events such as falling objects or loud noises. NPC's ignoring these events could break a users immersion in the environment. Preventing this break in a VR users immersion gives motivation for this investigation.

More advanced future work could involve an investigation into creating believable VR physical contact interactions with NPC's. This research area faces the difficulty of convincing the VR user that they have been involved in a collision with an NPC. Moving a VR user's camera around the virtual environment to simulate a collision without actually effecting the real world user themselves can break their immersion in the environment. This can also cause motion sickness for the VR user which will have a negative impact on their experience. This future work is an extremely difficult area of research, but successful solutions would unlock new levels of interaction and immersion for VR virtual environments.

Bibliography

[1] S. A. Hill, "An introduction to electricity & magnetism, chapter 1: Vectors, dot products," 2022. [Online]. Available: https://physics.nfshost.com/textbook/01-Vectors/04-Dot.php

[2] S. Team, "Superhot vr," 2016. [Online]. Available: https://www.oculus.com/experiences/quest/1921533091289407/

[3] VRChat, "Vrchat," 2017. [Online]. Available: https://hello.vrchat.com/

[4] V. Luckerson, "Facebook buying oculus virtual-reality company for $2 billion," 2014. [Online]. Available: http://time.com/37842/facebook-oculus-rift

[5] Wikipedia contributors, "Oculus rift — Wikipedia, the free encyclopedia," 2022. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Oculus_Rift&oldid=1083202164

[6] ——, "Metaverse — Wikipedia, the free encyclopedia," 2022. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Metaverse&oldid=1083622170

[7] G. V. R. (2021), "Virtual reality market size, share & trends analysis report by technology (semi & fully immersive, non-immersive), by device (hmd, gtd), by component (hardware, software), by application, and segment forecasts, 2021 – 2028," 2021. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/virtual-reality-vr-market

[8] J. Clement, "Number of video gamers worldwide 2021, by region," 2021. [Online]. Available: https://www.statista.com/statistics/293304/number-video-gamers/#:~:text=Number%20of%20video%20gamers%20worldwide%202021%2C%20by%20region&text=In%20total%2C%20there%20were%20an,billion%20gamers%20across%20the%20globe.

[9] W. J. Shelstad, D. C. Smith, and B. S. Chaparro, "Gaming on the rift: How virtual reality affects game user satisfaction," in *Proceedings of the Human Factors and*

Ergonomics Society Annual Meeting, vol. 61, no. 1. SAGE Publications Sage CA: Los Angeles, CA, 2017, pp. 2072–2076.

[10] N. Horvat, S. Škec, T. Martinec, F. Lukačević, and M. M. Perišić, "Comparing virtual reality and desktop interface for reviewing 3d cad models," in *Proceedings of the Design Society: International Conference on Engineering Design*, vol. 1, no. 1. Cambridge University Press, 2019, pp. 1923–1932.

[11] S. G. Izard, J. A. Juanes, F. J. García Peñalvo, J. M. Estella, M. Ledesma, and P. Ruisoto, "Virtual reality as an educational and training tool for medicine," *Journal of medical systems*, vol. 42, no. 3, pp. 1–5, 2018.

[12] T. Alexander, M. Westhoven, and J. Conradi, "Virtual environments for competency-oriented education and training," in *Advances in Human Factors, Business Management, Training and Education*. Springer, 2017, pp. 23–29.

[13] A. Nicás Miquel, "Development of crowd simulation models using unity for immerssive vr applications," Master's thesis, Universitat Politècnica de Catalunya, 2019.

[14] N. Pelechano and J. M. Allbecky, "Feeling crowded yet?: crowd simulations for vr," in *2016 IEEE Virtual Humans and Crowds for Immersive Environments (VHCIE)*. IEEE, 2016, pp. 17–21.

[15] P. Dickinson, K. Gerling, K. Hicks, J. Murray, J. Shearer, and J. Greenwood, "Virtual reality crowd simulation: effects of agent density on user experience and behaviour," *Virtual Reality*, vol. 23, no. 1, pp. 19–32, 2019.

[16] M. Kyriakou, X. Pan, and Y. Chrysanthou, "Interaction with virtual crowd in immersive and semi-immersive virtual reality systems," *Computer Animation and Virtual Worlds*, vol. 28, no. 5, p. e1729, 2017.

[17] H. Warpefelt, "The non-player character: Exploring the believability of npc presentation and behavior," Ph.D. dissertation, Department of Computer and Systems Sciences, Stockholm University, 2016.

[18] H. Warpefelt and H. Verhagen, "A model of non-player character believability," *Journal of Gaming & Virtual Worlds*, vol. 9, no. 1, pp. 39–53, 2017.

[19] J. Harth, "Empathy with non-player characters? an empirical approach to the foundations of human/non-human relationships," *Journal For Virtual Worlds Research*, vol. 10, no. 2, 2017.

[20] Wikipedia contributors, "Uncanny valley — Wikipedia, the free encyclopedia," 2022. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Uncanny_valley&oldid=1079899425

[21] W. Bouwer and F. Human, "The impact of the uncanny valley effect on the perception of animated three-dimensional humanlike characters," *The Computer Games Journal*, vol. 6, no. 3, pp. 185–203, 2017.

[22] N. Adamo, H. N. Dib, and N. J. Villani, "Animated agents' facial emotions: Does the agent design make a difference?" in *International Conference on Augmented Reality, Virtual Reality and Computer Graphics*. Springer, 2019, pp. 10–25.

[23] Unity, "Unity game engine," 2022. [Online]. Available: https://unity.com/

[24] ——, "Navmesh agent," 2022. [Online]. Available: https://docs.unity3d.com/Manual/class-NavMeshAgent.html

[25] U. Technologies, "Standard assets (for unity 2018.4)," 2022. [Online]. Available: https://assetstore.unity.com/packages/essentials/asset-packs/standard-assets-for-unity-2018-4-32351

[26] Unity, "Unity asset store," 2022. [Online]. Available: https://assetstore.unity.com/

[27] Wikipedia contributors, "Mixamo — Wikipedia, the free encyclopedia," 2022. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Mixamo&oldid=1072113984

[28] U. Technologies, "Xr origin," 2022. [Online]. Available: https://docs.unity3d.com/Packages/com.unity.xr.core-utils@2.0/manual/xr-origin.html#:~:text=The%20XR%20Origin%20represents%20the,Offset%20Object%2C%20and%20a%20Camera.

[29] ——, "Xr interaction toolkit," 2022. [Online]. Available: https://docs.unity3d.com/Packages/com.unity.xr.interaction.toolkit@2.0/manual/index.html

[30] M. Platforms, "Oculus quest 2," 2022. [Online]. Available: https://store.facebook.com/gb/quest/products/quest-2/?utm_source=https%3A%2F%2Fwww.google.com%2F&utm_medium=organicsearch

[31] U. Technologies, "Locomotion," 2022. [Online]. Available: https://docs.unity3d.com/Packages/com.unity.xr.interaction.toolkit@1.0/manual/locomotion.html

[32] ——, "Class xrrayinteractor," 2022. [Online]. Available: https://docs.unity3d.com/Packages/com.unity.xr.interaction.toolkit@0.9/api/UnityEngine.XR.Interaction.Toolkit.XRRayInteractor.html

[33] ——, "Class xrsimpleinteractable," 2022. [Online]. Available: https://docs.unity3d.com/Packages/com.unity.xr.interaction.toolkit@0.9/api/UnityEngine.XR.Interaction.Toolkit.XRSimpleInteractable.html

[34] D. V. Code, "Pdollar point-cloud gesture recognizer," 2015. [Online]. Available: https://assetstore.unity.com/packages/tools/input-management/pdollar-point-cloud-gesture-recognizer-21660

[35] JustCreate, "Low-poly simple nature pack," 2022. [Online]. Available: https://assetstore.unity.com/packages/3d/environments/landscapes/low-poly-simple-nature-pack-162153

[36] D. Carroll, "Attention and communication in virtual worlds: Interacting with non-player characters in vr," 2022. [Online]. Available: https://youtu.be/efNk7Sgddec