Well Played

a journal on video games, values, and meaning

A Special Issue on Post-Mortems of AR & VR Games

EDITED BY DREW DAVIDSON

VOLUME 12 NUMBER 2

Well Played (Vol. 12, No. 2)

WELL PLAYED (VOL. 12, NO. 2)

A Journal on Video Games, Values, and Meaning | A Special Issue on Post-Mortems on AR & VR Games

ETC PRESS

Carnegie Mellon University: ETC Press

Pittsburgh, PA

ISSN 2164-344X (Print)
ISSN 2164-3458 (Online)

All submissions and questions should be sent to: etcpress-info (at) lists (dot) andrew (dot) cmu (dot) edu For formatting guidelines, see: www.etc.cmu.edu/etcpress/files/WellPlayed-Guidelines.pdf

This book was produced with Pressbooks (https://pressbooks.com) and rendered with Prince.

CONTENTS

INTRODUCTION

This special issue was inspired by the interesting interactive experiences people are designing and having with immersive mixed reality. Augmented Reality and Virtual Reality (AR and VR) games are growing in popularity and access while also increasing the variety of experiences players can have that are specific to this nascent medium. They offer multiplayer embodied experiences, ranging from augmented reality games that tap into GPS and have you walking around like Pokemon Go, to virtual reality games in a headset like Beat Saber, to mixed reality experiences that are location-specific like Ghostbusters: Dimension. Interestingly, the open call for this special issue garnered post-mortems by designers and developers exploring the unique affordances of AR & VR.

The authors all share a sense of their design philosophy and lessons learned in their practice through the discussion of their design and development processes in relation to their goals for the gameplay. Throughout, they help define a sense of what it means to have a well played AR & VR experience.

CHAPTER 1.

The Impact of a Quadrupedal Avatar on Virtual Reality Locomotion

LAURELINE CHIAPELLO & CHARLES-RENÉ CHOUINARD

Keyword: Design process, virtual reality game, avatar, locomotion, controls, engagement, cybersickness, presence, arm-swinging, project-grounded research, natural mapping.

ABSTRACT

In virtual reality, the choice of locomotion can drastically change the perception of the game experience. Different locomotion schemes have already been explored and lead to various challenges for players, concerning mostly movement and cybersickness. However, there have been fewer investigations into the relationship between avatar design and locomotion.

In this paper, we will show how we explored different locomotion schemes during the creation of a game where the avatar is a wolf. The purpose of this experiment was to see how locomotion and the avatar are linked to each other and how it affects the experience we aim to create as game designers.

By utilizing a project-grounded research approach, we documented how we defined the best choice of locomotion for our project and how locomotion can create presence and participate in the design of a specific avatar. We will detail the making of a prototype for the virtual reality game *Howl's Adventures*. We found that the *arm-swinging* approach for locomotion, a bridge between indirect locomotion and gesture-based locomotion, was the most valuable way to create the

fantasy of being a wolf, as it provided a natural and intuitive form of mapping, which granted a form of physical immersion in the experience.

1 – INTRODUCTION

Would you like to know what it is like to be a wolf? While we might never fully be able to reach this fantasy, virtual reality offers a way to explore new points of view on the world. Over the last two years, we developed *Howl's Adventures*, a virtual reality prototype that allows the user to experience the fantasy of being a wolf. Following a "research *through* design" methodology as defined by Alain Findeli (2015), the project's general goal was to create an original experience with an immersive control scheme that allows players to embody wolf traits. Players can experience the joy of running free in the wild, jumping over obstacles, chasing creatures, grabbing objects with their mouths, and howling under the moon.

This paper delves into the investigation of the designer's own process to find the best locomotion for controlling a wolf in virtual reality. How does the avatar design impact the locomotion and game controls? This type of research emphasizes the design process and the choices made by the designers compared to evaluative research, which assesses the impact on players.

The first part of this paper will be dedicated to presenting the most popular forms of locomotion for virtual reality games. The second part will explain the research *through* design methodology and detail the project we developed: *Howl's Adventures*. Additionally, we will provide a description of the auto-ethnographic tools we employed to collect our data and explain how we utilized them to analyze our project. Finally, we will present our research findings regarding the impact of avatar design on locomotion and game controls.

2 – FROM AVATAR TO PRESENCE AND LOCOMOTION

In 2016, access to virtual reality (VR) equipment exploded in popularity as a result of becoming easily available and affordable for the mass market (Sherman and Craig, 2019). This created a demand for VR experiences, which resulted in a multiplication of titles on official game stores, as seen on websites like Virtual Reality Databases (VRDB) (Virtual Reality Databases, 2020), an online database that maintains records for all new virtual reality games released on each platform.

The virtual reality experience can vary greatly between each game. Some titles put players directly in control of their avatar in a first-person view, like *Beat Saber* (Beat Games, 2018) and *Richie's Plank Experience* (TOAST, 2016), while others grant players a third-person view, like *Astro Bot Rescue Mission* (Sony Interactive Entertainment, 2018) and *Tetris Effect* (Enhance, 2018). All of these games use VR to create an immersive and sensory experience. The illusion of interacting with the environment is done through a head-mounted display (HMD) and the system controllers. This immersive experience relies on different aspects, but a major one is the idea of presence. Sherman and Craig define presence as: "short for sense of presence; being mentally immersed" (Sherman and Craig, 2019, p.10). Being mentally immersed refers to the state where players are deeply engaged, involved in the action, and perceive a suspension of disbelief (Sherman and Craig, 2019).

"Being in the world" is thus a major concern for VR game designers, and it often starts with a fundamental action: moving around. Therefore, a primary challenge for VR designers is to discover the optimal method of navigating the environment, commonly referred to as 'locomotion' (Bozgeyikli et al., 2016). The link between the controls scheme, which button to press, and locomotion is one of the earliest challenges tackled by any game developer (McEntee, 2012). This challenge can be seen in various game design models and is commonly referred to as *the "3Cs"* which are *Camera, Character,* and *Controls* (McEntee, 2012; Rogers, 2010). Figuring out how to implement locomotion early in the game development process is not an easy task, and it becomes even more challenging when facing new technologies.

2.1 – Camera and Character/Avatar

In the game industry, the term *avatar* helps to distinguish the playable character from the other non-playable characters (Schell, 2008). Avatars are a tool used to navigate the game world and solve tasks at hand (Fullerton, 2008; Rogers, 2010). As such, the avatar can take multiple forms, be it human-like Mario in *Super Mario Bros* (Nintendo, 1985), an animal like the goose in *Untitled Goose Games* (House House, 2019), an anthropomorphic cartoon character like Sonic from *Sonic The Hedgehog* (Sega, 1991), or it can even be an abstract object like the pieces in *Tetris* (The Tetris Company, 1996) (Rogers, 2010). With these avatars, players project themselves into the game world (Bateman 2010; Schell, 2008).

In virtual reality, the experience is presented in a first-person perspective, and rather than featuring a full avatar, players observe elements that represent the human body, such as hands or feet (Sherman and Craig, 2019). This empty shell of an avatar eases players into the virtual reality game experience (Dufour et al, 2014).

Virtual reality offers an exclusive approach to the game experience, one that is centered around player body gestures. These gestures help by creating both mental immersion and physical immersion. Sherman and Craig describe physical immersion as "bodily entering into a medium, synthetic stimulus of the body's senses via the use of technology; this does not imply all senses or that the entire body is immersed/engulfed" (Sherman and Craig, 2019, p.10). Adding a physical dimension to the experience helps to push further the boundary of player immersion in order to create a greater feeling of presence in those virtual worlds.

Quadrupedal avatars, such as Amaterasu in *Okami* (Capcom, 2006), are less common in video games compared to their bipedal counterparts. They are often utilized to explore new gameplay styles that require unique characteristics not found in other avatars (Rogers, 2010). Particularly, animals are frequently used to create and present a contrast between the animal and human experiences (Bateman, 2010). Referring to our initial example with Amaterasu, she is the sun goddess of the Shinto religion, reincarnated as a white wolf who happens to be unable to speak directly to humans. However, she is able to understand them. In her game, she is initially feared by all humans for simply being a wolf, which sets the mood and challenges for our heroine. The human fear refers to the stereotype commonly attached to wolves in media and folklore, where they are scary and dangerous beasts. In those cases, animals observe and act upon humans. As for developing our project, we have yet to see any commercially released virtual reality games that utilize a quadrupedal avatar.

2.2 – Controls

The main framework for describing locomotion in virtual reality is the one elaborated by Frommel et al., who observed that with our current virtual reality gear set up, four main first-person controller-based locomotion can be found in games: free teleport locomotion, fixpoint teleport locomotion, indirect locomotion, and automatic locomotion (Frommel et al., 2017). For each locomotion method, these authors emphasize two components: movement and cybersickness. Movement refers to the way players navigate the virtual world with their avatars. Cybersickness encompasses any symptoms of exhaustion experienced during the VR experience, which can result from the activity itself or from wearing virtual reality equipment, or various forms of motion sickness, ranging from nausea to headaches (Davis et al., 2014). Furthermore, certain locomotion forms are known to cause more cybersickness than others (Frommel et al, 2017).

Free teleport locomotion allows players to move their avatar by using a pointer tied to their hand position (Frommel et al., 2017). With the controller, players are able to cast a pointer from their avatar's hand in the virtual world (Figure 2 and Figure 3). The pointer determines the end position where the avatar will be teleported. Because the player's point of view changes instantly, this creates little to no cybersickness, which allows players to explore the virtual world freely and relatively fast.

Fixpoint teleport locomotion is similar to free teleport locomotion, but instead of giving the option for players to freely teleport their avatar anywhere in the world based on an arc from their avatar's hand, fixpoint teleport locomotion locks the avatar to a predetermined destination point by selecting specific markers chosen by the game developers (Frommel et al., 2017). This restricts players' movement and guarantees they receive all the necessary information to solve the task. Locking avatars in a specific area prevents players from getting lost. As a result, this locomotion doesn't cause much, if any, cybersickness.

Indirect locomotion allows players to navigate the virtual world using tools like joysticks or buttons on controllers (Frommel et al., 2017). This form of locomotion shares many similarities to traditional first-person games or third-person games locomotion on a personal computer or home console. This locomotion grants freedom to players, allowing their avatars to navigate and move wherever they want. However, this locomotion is often seen as slow for virtual reality and leads players to experience more cybersickness because of the avatar speed (Frommel et al., 2017).

Automatic locomotion automates camera movement by putting it on rails (Frommel et al., 2017). This form of locomotion removes all freedom of movement from the players, leaving them only the ability to aim their camera to observe their surroundings, allowing game developers to create events at specific locations. Unfortunately, this form of locomotion generates a significant amount of cybersickness. Cybersickness primarily arises from the uncontrollable speed of the automation, the player's tolerance to speed, and the removal of the avatar's movement controls from players while they remain static.

Other forms of locomotion, beyond those presented by Frommel et al., are more complex or require specific equipment, such as omnidirectional treadmills. In particular, gesture-based locomotion usually requires players to wear trackers on their leg or elsewhere on their body to capture the player's walking motion, in order to make the avatar move (Slater et al., 1995). Gesture-based locomotion offers the ability for players to use a mimetic approach for controlling their avatar movement while traversing the virtual world. In recent years, this locomotion has evolved in order to utilize controllers as a form of motion capture.

In order to gain a deeper comprehension of diverse locomotion methods and the importance of presence in the virtual reality realms via avatars, we designed a game prototype to investigate the creation of an immersive experience featuring a wolf as our avatar.

3 – PROJECT-GROUNDED RESEARCH

This study follows a research *through* design approach, specifically inspired by "project-grounded research", as defined by the French design researcher Alain Findeli (2015). The concept of research *through* design originates from Christopher John Frayling, who introduced different forms of interaction between research and design: research *for* art and design, research *into* art and design, and research *through* art and design (Frayling, 1993; Godin and Zahedi, 2014). Research *through* design allows researchers-designers to explore and understand their own practice within a design project.

Findeli refined Frayling's initial approach and developed "project-grounded research", which relies on a project to serve as a research field where data is collected (Godin and Zahedi, 2014). The goal of this approach is to produce valuable knowledge while exploring the researcher/designer's own perception of design problems (Findeli, 2015; Chiapello, 2019). It is essential for the researcher to actively take part in the project in order to develop those new valuable insights (Findeli, 2015).

Findeli invites researchers to link design projects and research to conceive an environment that allows practitioners to make new discoveries: "If you want to understand a phenomenon or a concept, put it into a [design] project" (Findeli, 2015, p. 56). This type of research differs from user experience research, which focuses on evaluating design solutions through user testing. Research *through* design allows for the discovery of the knowledge and understanding what underlines the creation or design of an artifact by the researcher (Godin and Zahedi, 2014).

In project-grounded research, researchers adopt a reflective stance toward their own practice (Schön, 1983). The objective is to discover tacit knowledge emerging from practitioners' actions and to reveal the frames they adopted (Schön, 1983). The practitioner-researchers need to be aware that they're entering into a conversation with the situation in order to construct new knowledge.

Furthermore, Schön's goal with reflective practice was to address the issues observed within *Technical Rationality*. Technical Rationality is a positivist model of practice, where professionals are trained to apply preexisting knowledge and tools to solve problems without interrogating those tools and solutions, as the positivism movement separates research from practice (Schön, 1983). On the contrary, research *through* design grants that specific knowledge can emerge from practice. Practitioners are not just applying knowledge to a situation: they can craft new understandings and refine existing knowledge in the midst of action.

This is the originality of this research: as previous authors have already described many different challenges related to VR, we wanted to see how we could explore and refine their concepts during a complex project. Rather than constraining our exploration to evaluate isolated concepts, we embarked on an expansive design project centered around delivering a comprehensive player experience – embodying the role of a wolf – to explore emergent challenges.

Therefore, during the development of the *Howl's Adventures* project, we delved into various challenges. In addition to findings on the impact of avatars in VR experiences, parallel outcomes on game design were revealed and enhanced our understanding of the 3Cs model. However, these outcomes will not be discussed in this paper.

In this paper, our primary objective is to delve into virtual reality locomotion concerning non-human avatars. We aim to detail the design process pursued by a developer targeting this unique experience. By leveraging qualitative data derived from the developer's experiment, our intention is to shed light on the challenges encountered and the subsequent design decisions made within this particular context.

3.1 – The project: *Howl's Adventures*

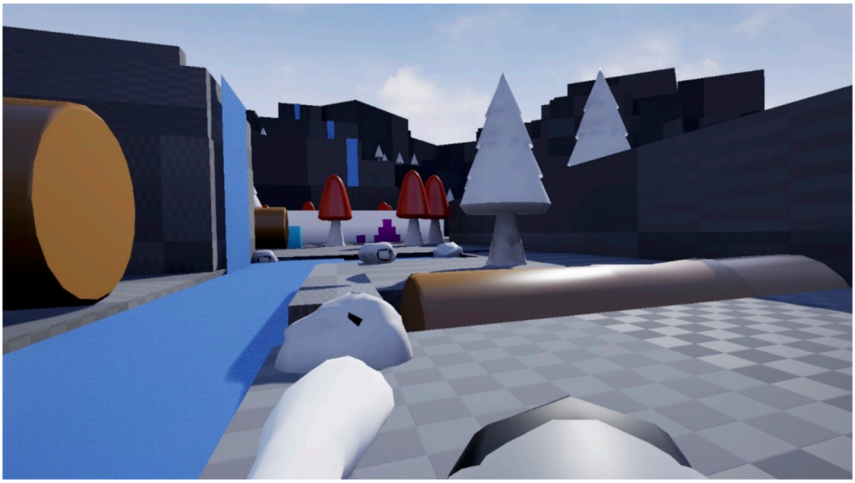

Figure 1: A screenshot of Howl's Adventures prototype level.

In this study, we developed a virtual reality game prototype for the *Oculus Quest*, now known as the *Meta Quest 1* (Facebook Technologies LLC, 2019), in which the player's avatar takes the form of a wolf. The game was developed as part of a Master's degree program in Art, with a specialization in digital design. The development team consisted primarily of a student who served as the main designer under the supervision of a director. The student was responsible for the majority of the game's creation, guided by their supervisor, particularly in research aspects. Additionally, a programmer provided assistance in specific code segments, and other research lab students tested the game informally. The prototype, as a game design project, did not incorporate artistic elements, such as detailed realistic 3D models. The entire project was developed with *Unreal Engine 4* (Epic Games, 2014).

The game concept revolved around experiencing the fantasy of being a wolf. Titled *Howl's Adventures*, this action-adventure game tasked players with exploring a 3D environment and solving puzzles to unlock new areas. The tasks found in the prototype consisted of interacting with the environment, flipping switches, pushing objects, and grabbing a key item and using it to solve some small puzzles. The wolf avatar is able to run around, grab objects with its mouth, push objects with its head and paws, bark, and, of course, howl.

While developing this game experience, we encountered several issues that limited player gestures. During play testing, the HMD was connected to a PC with a USB-C cable. Users had to contend with this additional cable, which deviated from the cable-free experience expected with a *Meta Quest 1*. This added cable often became entangled around the player's body, leading to restricted movement. In such cases, we had to intervene to untangle players from their constricted state, which disrupted their immersion and impacted their sense of presence within the game world. Surprisingly, this single cord caused more inconvenience than anticipated.

As explained earlier, as the project progressed, locomotion models and their limitations surfaced as a focal point, prompting our team to delve deeper into their intricacies. We explored which form of locomotion would offer the most immersive experience for a virtual reality game and how our unique avatar would help shape this reasoning. The design process included determining the avatar proportion, the body collision size, the gameplay action, the camera position, and the avatar aspect, such as the paws and muzzle.

Given that this prototype aimed to offer an immersive and playful experience with a wolf avatar that closely mimicked the behavior of a real wolf, all gameplay tasks were tailored to suit a quadrupedal avatar. As an illustration of how object interaction diverged from conventional VR games, players use the wolf avatar's mouth to manipulate objects, simulating how canines naturally pick up objects instead of relying on hand controllers.

Constructing the game with the player assuming the role of a wolf as their avatar presented a novel and uncommon experience, particularly within the realm of video games and more so in the context of virtual reality. As we delve into this, our original concept introduced a host of design challenges, sparking insights related to game design, locomotion, and the concept of presence.

3.2 – Data collection and analysis

We collected qualitative data using auto-ethnographic tools. Data gathering was accomplished through the maintenance of a comprehensive logbook that housed all necessary technical information related to the project, as well as personal reflections on game development and design. We acquired data through various mediums during the researcher's exploration, including voice recordings, handwritten notes, hand-drawn diagrams on paper, and screen captures. While we conducted informal playtesting within the team, our research did not involve external participants.

This approach allowed us to document both the researcher's intentions and their emotional states, which can significantly influence our research practices (Rondeau, 2011). This strategy enabled us to incorporate the subjectivity of the designer-researchers, thereby preserving the human experience as an integral part of the project development (Rondeau, 2011) and enhancing our understanding of the design process. This subjectivity serves as both a strength and a limitation of the research *through* the design approach.

Analyzing the raw data from our logbook led to the emergence of valuable insights. The logbook itself may initially appear as an unorganized collection of information that is challenging for unfamiliar individuals to comprehend (Rondeau, 2011). To make sense of this apparent chaos, it became necessary to take a step back, observe the interconnected elements, and craft a more structured report. In order to explicitly convey the designer's perspective, this report is typically composed in the first person. The aim is to convey the richness and subjectivity inherent in the design process. This approach enabled us to gain a deeper understanding of how the avatar's locomotion influences the sense of presence for players.

4 – RESULTS

The design process does not follow a straightforward and linear path, as one might initially expect. In this section, we will explain how we created three different versions of *Howl's Adventures*, each incorporating a distinct locomotion system. More importantly, we will elucidate the design choices that guided this trajectory. Our objective was not to comprehensively test all the previously mentioned locomotion methods but rather to demonstrate how our prior knowledge of these methods, combined with the specific project context, informed our deliberations regarding locomotion, the avatar, and the sense of presence.

While developing the game experience centered around the wolf avatar in *Howl's Adventures*, the majority of our explorations steered us toward a choice of locomotion that effectively maintained the sense of presence while reducing the occurrence of cybersickness in players.

We implemented and tested two out of the four forms of first-person virtual reality locomotion listed by Frommel et al.: free teleport locomotion and indirect locomotion. However, these two methods did not convincingly capture the experience of embodying a wolf. Instead of the locomotion methods proposed by Frommel et al., we explored a gesture-based locomotion system that relied on arm-swinging motions.

4.1 – Free teleport locomotion

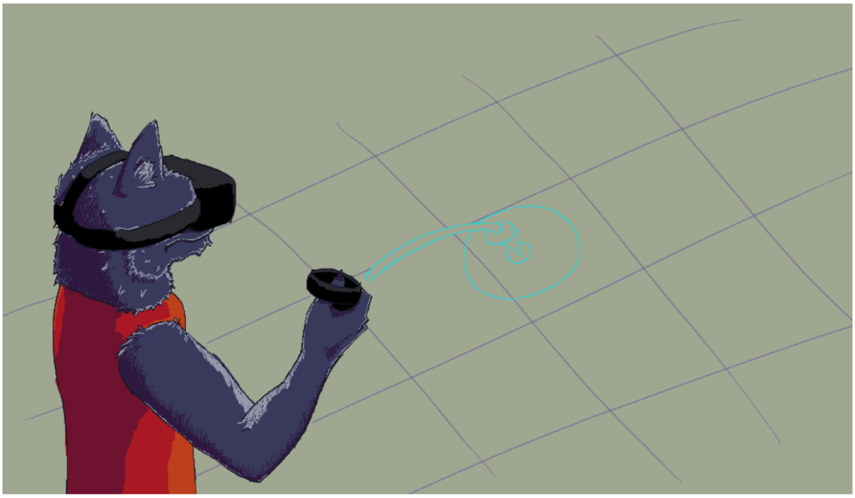

Figure 2: A player aiming while using free teleport locomotion to navigate in the virtual environment.

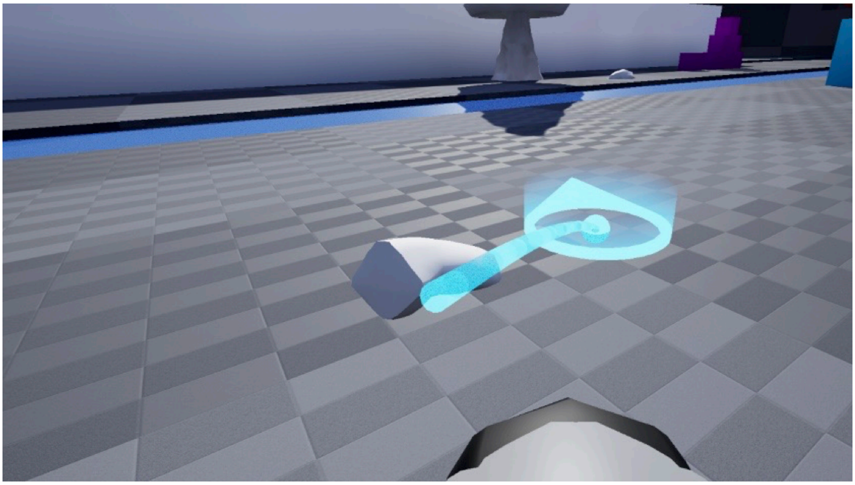

Figure 3: A player using their paw "hand" to point their end destination for free teleport locomotion.

As an initial exploration, the team opted for free teleport locomotion since it was provided as the default locomotion with *Unreal Engine 4*. In the game, players would use the avatar's paw

(Figure 2 and Figure 3) to point and spawn the teleportation arc showing where they would teleport to.

This locomotion method proved highly effective in granting us the freedom to move about and explore the game levels at our own discretion. To enhance the effectiveness of this locomotion, we decided to play while standing up. This not only provided us with a deeper awareness of our surroundings but also imbued the overall experience with a heightened sense of dynamism, closely resembling real-life movement. While standing up, we exhibited a greater tendency to actively survey our environment and engage in physical movement, in contrast to when we remained seated. When it came to navigating the environment, utilizing teleportation was simple to perform. The act of aiming and pointing with the controller was quick and easy to use.

The main problem with teleportation was that we could navigate through the environment too fast and easily get out of bounds because the game levels were designed as an open area. This locomotion made the traversal of levels so straightforward that it was possible to skip some sections. Implementing this locomotion was extremely simple, but it ended up requiring additional attention to find ways to prevent unintended teleportation outside the game levels without impacting the level design. It required additional development to come up with a solution that restrained player's freedom in the environment. One of our solutions was to add additional scenery objects that would naturally limit player movement, like trees, cliffs, or caverns. In specific scenarios, an invisible wall was added that disables the teleportation arc altogether.

Moreover, the ability to travel fast across game levels made us less aware of our surroundings, which lead us to be less immersed in the environment. Therefore, we tried some ways to slow us down. However, no matter the restrictions, nothing would adequately restrain us from speeding through the environment with teleportation. One other solution that was brought up, but not implemented as we thought it would make the overall navigation frustrating, was to add a cooldown to the teleportation system. All of these additional restrictions to the avatar's movement would contradict the freedom found with free teleportation locomotion.

In addition, for this particular locomotion version, we incorporated the ability to grab objects both with the avatar's paws and its mouth. However, we quickly noticed that we predominantly used our hands to interact with objects rather than emulating canine behavior by using the avatar's mouth. It became evident that grasping objects with the avatar's paws was more intuitive. While wearing a VR headset, attempting to grab objects with our actual mouths was initially perceived as amusing and inventive. Nevertheless, when compared to the precision and efficiency of using our hands, this approach was swiftly dismissed as challenging and inefficient.

With this locomotion form, we were not influenced by the quadrupedal avatar. In other words, we relied on our intuitive bipedal approach to complete the task at hand and explore the environment. This made the avatar more of a costume than a feature of the experience, which resulted in a lack of immersion and less attachment toward the character.

Moreover, seeing the wolf's avatar teleporting itself around felt absurd. This would require heavy lifting in the narrative department to make sense of this ability. Because of the low feeling of presence in the world, it was decided to implement another type of locomotion for this prototype.

4.2 – Indirect locomotion

Figure 4: A player using hand controllers in indirect locomotion, simulating the experience of using a traditional gamepad.

For our second exploration, we choose indirect locomotion. For this locomotion setup, our goal was to implement a navigation system similar to first-person games for the home console by building on top of our initial control setup for free teleport locomotion. To achieve a similar control scheme, we selected the left joystick of *Meta Quest 1* controllers to be used for movement. We left the camera aiming attached to the HMD, and we decided to play while sitting in a swivel chair. The swivel chair allowed us to easily change the direction of the avatar, thanks to the HMD. We decided to play while sitting in order to make the experience more closely to how one plays video games.

This form of locomotion allowed us to navigate the playable area more naturally. As a side effect, this greatly slowed our movements, which resulted in giving us more time to pay attention to our surroundings. We evaluated this form of locomotion positively, finding it easy to use and more in line with our previous gaming experiences. This, in turn, gave us more time to visualize and immerse ourselves in the environment.

Nevertheless, this form of locomotion also had its shortcomings. Cybersickness was the main issue plaguing this version, leading us to feel discomfort, which broke our presence in the virtual world. The rupture of presence led us to be disconnected and disinterested in the avatar. This discomfort mainly stemmed from the low tolerance thresholds for the avatar and camera speed in a first-person perspective experience. Moving slowly bored us while moving too fast would lead to cybersickness. Every time the in-game camera stopped adequately following real-life motion, we perceived a dissonance between our actions and those seen with the avatar, initiating nausea. We ended up using a more conservative movement speed to prevent the chance of gaining cybersickness.

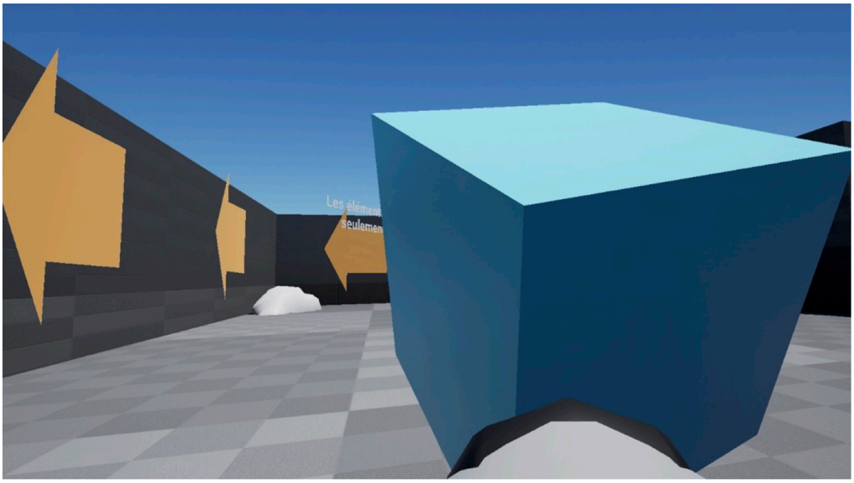

Figure 5: A player holding an object via the wolf avatar mouth.

Lastly, indirect locomotion did not distinguish between a quadruped and biped for movement. Moving with a joystick de-emphasized the use of our hands. This resulted in an experience that felt too static for a virtual reality game, as most VR experiences tend to aim for physical immersion. This means that both hands and the head should be involved in the action. However, in *Howl's Adventures*, as explained earlier, players are not supposed to use their hands to grab objects because their avatar is a wolf. Instead, they grab objects with the avatar's mouth, located slightly below the muzzle of the wolf avatar, as shown in Figure 5. However, indirect locomotion removes an opportunity for the player to connect with the avatar via the paws. Most of the time, we would leave the hand controllers lying on our lap, which removed the paws from the screen and out of our sight. This sitting position removed one of the few elements that reminded us that we were a wolf.

Moreover, animations in video games must go hand in hand with the avatar to help sell movement and make the characters feel alive (Rogers, 2010). In first-person games like *Mirror's Edge* (Electronic Arts, 2008) or *Doom Eternal* (Bethesda, 2020) where the avatar is not fully seen on screen, only their hands, weapon, or any other object being held by them are visible, it is possible to feel the weight of their avatar movement thanks to the animation. Those animations create a feeling of presence in the game world for players.

When it comes to virtual reality games, adding those types of animation to the avatar can lead to new problems. With indirect locomotion, the avatar position is controlled by a joystick. On the other hand, the camera and the controllers both follow their own devices for orientation and position. In virtual reality, if animations overlap the player's own motion, they end up decoupling the player's motion from their avatar, disrupting the player's presence and resulting in a disjointed experience. Furthermore, in first-person virtual reality games, it is not necessary to add animations on the avatar because the player's gesture will perform the required animation for exploring the environment. This results in players creating their own presence in the virtual world. Finally, we observed that indirect locomotion reduces the amount of motion performed, which resulted in feeling less presence from our action.

To summarize, indirect locomotion allowed us to navigate and immerse ourselves in the environment easily. However, in this specific scenario, indirect locomotion resulted in an unrefined experience where it was easy to forget that we were supposed to play as a wolf. The avatar had limited visibility on the screen, and since its paws had no direct purpose, this only left the muzzle as a reminder of the wolf shape. We then decided to try a form of locomotion that would utilize a mimetic approach to the controls in order to embody the wolf better.

4.3 – Gesture-based locomotion

After trying both free teleport locomotion and indirect locomotion, it was concluded that Frommel et al. framework might not provide a satisfying solution for our game. Frommel et al. evoke gesture-based locomotion, but they do not develop it thoroughly. At this time, we believed it to be a promising avenue to employ a mimetic approach for locomotion and chose to explore arm-swinging.

Figure 6: A picture presented in game to instruct players on how to perform the arm-swinging motion.

The arm-swinging motion seemed to be appropriate for the wolf avatar in *Howl's Adventures*. In this scenario, the player's hands are the front paws of the wolf avatar. Moving them around allowed us to simulate the walking cycle of a wolf, as shown in Figure 6. Changing the controls to this type of mapping provided the most natural form of controller mapping for the games. The goal of natural mapping for video games is to handle controls of the game in a way that eases players into a mental model similar to their real action counterparts (Skalski et al., 2011; Steuer, 1992; Tamborini and Bowman, 2010). Furthermore, this form of control allows us to change our speed from walking to running like a wolf simply based on how fast we swing our arms. This results in everyone navigating the environment at their own speed. These emerging results guided our interest in creating an experience that would cherish this form of control for the game's locomotion.

Arm-swinging motion allows us to remind ourselves that the avatar, which normally is humanoid, was actually a wolf. The positive outcome of forging intuitive game controls is that it makes the avatar act as an extension of our body, which results in developing presence (McDonald, 2012; Swink, 2009).

Using a mimetic approach to game controls reintroduced the concept of physical immersion, which is unique to the virtual reality experience. This allowed us to be fully immersed, creating a more meaningful experience. Arm-swinging motion became an integral aspect of the overall experience and game design. Juxtaposing a control scheme that fits the avatar, in coordination with player motion and a mimetic approach, resulted in creating a presence for players.

Furthermore, using a mimetic approach helps by solving the issue with the static experience encountered with indirect locomotion. Adding an arm-swinging motion makes locomotion an active and dynamic experience. The light physical activity performed with a mimetic approach creates presence and gives relevance to the game world.

Contrary to what was originally expected, locomotion using arm-swinging does not create a lot of cybersickness. We were tiring ourselves from physical workouts, but we did not suffer cybersickness from screen motion or in-game motion. It's believed that synchronization of HMD and controllers to player's motion tricks the brain in a way that prevents them from gaining cybersickness.

In the end, making players act like a wolf, from walking to grabbing objects, helped to increase their presence. Adding mimetic motion like arm-swinging made players more aware of the avatar's nature, leading to greater engagement in performing gameplay actions more suited to a quadruped than a biped. This form of locomotion provided a deeper and more immersive experience, increasing player presence.

5 – DISCUSSION, LIMITS, AND CONCLUSION

The investigation of how a quadrupedal avatar can challenge game design was an unusual way to approach a topic that lacks exploration in conventional video games or in virtual reality. Building upon the insights gleaned from our results, it became evident that the avatar played a pivotal role in shaping the discourse of our research, exerting influence over both game design and locomotion choices. Ultimately, we opted for the arm-swinging locomotion method, as it proved to be the optimal choice for enhancing the sense of presence in our virtual reality experience. This approach engendered a fusion of mental and physical immersion.

Furthermore, the arm-swinging locomotion method offered the most natural mapping, effectively enabling us to embody the movements of a wolf, mirroring our avatar's actions. This locomotion method emerged as the most effective means for us, as developers, to establish a profound and meaningful sense of presence.

In contrast, the other two locomotion methods left us wanting for more. Teleportation, while simplifying navigation through the virtual environment, removed the challenges tied to exploration and appeared disassociated from our avatar. Indirect locomotion, on the other hand, restricted our gestures and interactions.

Through our testing, we observed that each type of locomotion induced varying levels of cybersickness. Free teleport locomotion had the lowest cybersickness, while indirect locomotion was the one where we experienced the most cybersickness. As for arm-swinging locomotion, it did not make us feel motion sickness. Instead, it resulted in us exhausting ourselves faster than any other locomotion thanks to the light activity done from swinging our arms around to move in the virtual world.

Returning to Frommel et al. four types of popular locomotion, the arm-swinging mimetic approach can be the starting point of a new type of locomotion, one that we call *indirect mimetic locomotion*. This form of locomotion refers to one that uses mimetic motion like arms-swinging utilized in tandem with existing motion controllers to control direction, speed, and player movement, in a desire to offer a similar movement experience seen with indirect locomotion. *Indirect mimetic locomotion* acts as a hybrid form between indirect locomotion and gesture-based locomotion to offer the best of both worlds.

Through this project, we discovered that when designing an experience featuring a non-bipedal avatar for VR, a mimetic approach can help foster the connection with the avatar. This approach resembles the imaginative roleplay commonly seen in childhood games of make-believe (Bateman, 2010). Bateman's posits that the use of representation in make-believe serves as a tool for individuals to shape an image and comprehend their fictional self. This mimetic projection of players into their avatars should facilitate presence. We thus hope that this mimetic approach to control will grant players a heightened sense of expression within the experience, aligning with our design goals.

Providing the necessary tools for players to freely roleplay significantly guided the direction of this project.

In conclusion, there are various forms of locomotion in the realm of virtual reality. In the scope of this project, we specifically focused on a select few. We chose to implement free-teleport locomotion, indirect locomotion, and a gesture-based method known as arm-swinging. These choices were guided by their accessibility, the availability of information, and their relevance to our desired VR experience using the *Meta Quest 1*. The quadrupedal avatar further accentuated the limitations of VR locomotion. One might even wonder whether these issues constitute one of the main challenges to the democratization of virtual reality games.

Through a project-grounded research approach and by adopting a reflective perspective, we determined that gesture-based locomotion involving arm-swinging was the most effective means to immerse ourselves and instill a sense of presence in the *Howl's Adventures* prototype. This choice was notably well-suited to our game's distinctive context, which featured a quadrupedal avatar. This locomotion method allowed us to establish a connection between ourselves and our wolf avatars, employing a mimetic approach and natural mapping to provide an immersive experience based on our avatar's actions.

Although we championed arm-swinging as a VR locomotion method for our project, it doesn't solve all VR movement challenges; it's not a one-size-fits-all solution. In contrast to our findings, another study favored 'walk in place'—moving by simulating walking—over arm-swinging (Tunnel Wilson et al., 2016). According to it, users performed better with 'walk in place,' accurately estimating distances, and making fewer turning errors compared to arm-swinging. Depending on the game's intended experience, each team might face different problems. Choosing the right locomotion method depends on the game's context and significantly influences research outcomes, as our study also showed.

Indeed, we hope to contribute to valorizing scientific research using qualitative data and promoting the designer's point of view, as it appears to be less common than quantitative approaches (Zielasko and Weissker, 2023). As stressed in Zielasko and Weissker research, most human-computer interaction (HCI) studies for VR locomotion are based upon quantitative data and adhere to a positivistic epistemology (Zielasko and Weissker, 2023). Rather than delving into the metrics of each locomotion technique, our goal was to comprehend and share the intricate design process involved in implementing different locomotion schemes within our distinctive context. While our study could benefit from an evaluative study utilizing extensive large-scale playtests, our primary objective was to demonstrate how decisions about the locomotion during the prototype stage are influenced by the designer's intentions concerning the avatar.

We hope this study will attract more attention and shed light on the game development design process. We aspire to encourage future experiences, take creative risks, and adopt distinctive approaches to video games, as we did with our canine companion, to expand the horizons of gameplay and game design.

ACKNOWLEDGMENTS

The authors would like to thank Rochelle LaForest for her English vocabulary and grammar corrections. We would also thank Soka, Charles-René's dog, who sparked the entire idea of this project.

BIBLIOGRAPHY

Bateman, Chris (2010), *Imaginary games*. Winchester, UK: Zero book.

Beat Games (2018), *Beat Saber*, Beat Games.

Bethesda Softworks (2020), *Doom Eternal*, id Software, Bethesda Softworks.

Bozgeyikli, Evren, Raji, Andrew, Katkooti, Srinivas and Dubey, Rajiv (2016), 'Point & Teleport Locomotion Technique for Virtual Reality', *Proceeding of the 2016 Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '16),* Austin, Texax, 16-19 October, New York, NY, pp. 205-216.

Capcom (2006), *Okami,* Capcom.

Chiapello, Laureline (2019), 'Le pragmatisme comme épistémologie pour le design de jeux, Enquête sur la créativité et le processus de design'. Ph.D. thesis, Montreal : Faculté d'aménagement, Université de Montréal.

Davis, Simon, Nesbitt, Keith and Nalivaiko, Eugene (2014), 'A systematic review of cybersickness', *Proceeding of the 2014 Conference on Interactive Entertainment (IE2014),* Newcastle, Australia, 2-3 December, ACM Press, New York, NY, pp. 1-9.

Dufour, Tristan, Torregrossa, Théo, Ibarra Cortes, Ricardo, Pellarrey, Vincent, Nachbaur, Vladimir, Clermont, Jonathan, Chagnon, Philippe, Li, Cheng, Dumas, Florent, and Majdoubi, Ahmed (2014), *ASCENT : A First Person Mountain Climbing Game on the Oculus Rift.* CHI PLAY '14, Oct 19-22, 2014, Toronto, ON, Canada.

Electronic Arts (2008), *Mirror's Edge,* DICE, Electronic Arts.

Enhance (2018), *Tetris Effect*, Enhance.

Epic Games (2014), 'Unreal Engine 4', Unreal Engine, https://www.unrealengine.com/en-US/. Accessed 21 September 2019.

Facebook Technologies LLC (2019), 'Oculus Quest', Meta Quest, https://www.oculus.com/quest/refurbished/. Accessed 9 January 2023.

Findeli, Alain (2015), 'La recherche-projet en design et la question de la question de recherche : essai de clarification conceptuelle', *Science du Design*, 1:1, pp. 45-47.

Frayling, Christopher John (1993), 'Research in art and design', *Royal College of Art Papers series*, 1(1).

Frommel, Julian, Sonntag, Sven, and Weber, Michael (2017), 'Effects of controller-based locomotion on player experience in a virtual reality exploration game', *Proceedings of the 12th International conference on the foundations of digital games (Article 30),* Cape Cod, MA, 14-17 August, Hyannis, MA: Association for Computing Machinery, pp. 1-6.

Fullerton, Tracy (2008), *Game design workshop, second edition: A playcentric approach to creating innovative games*. Burlington, MA: Morgan Kaufmann.

Godin, Danny, and Zahedi, Mithra (2014), 'Aspects of Research through Design: A Literature Review, in Lim, Youn-kyung, Niedderer, Kristina, Redström, Johan, Stolterman, Erik, and Valtonen, Anna (eds.), *Design's Big Debates – DRS Interna Conference 2014,* 16-19 June, Umeå, Sweden.

House House (2019), *Untitled Goose Games,* House House.

McDonald, Peter (2012), 'On couches and controllers: Identification in the video game apparatus', in M. Wysocki (ed), *Ctrl – alt – play: Essays on control in video gaming,* Jefferson, NC: McFarland, pp. 108-120.

McEntee, Chriss (2012), Remaking Rayman: Camera, character, controls. *Gamasutra,* October 2012, pp. 26-31.

Nintendo (1985), *Super Mario Bros,* Nintendo.

Rogers, Scott (2010), *Level-up, the guide to great video game design,* Chichester, UK: Wiley & Sons.

Rondeau, Karine (2011), 'L'autoethnographie: une quête de sens réflexive et conscientisée au coeur de la construction identitaire', *Recherches qualitatives,* 30:2, pp. 48-70.

Schell, Jesse. (2008), *The art of game design, a book of lenses*. Burlington, MA: Elsevier, Inc.

Schön, Donald A. (1983), *The reflective practitioner, how professionals think in action*. New York, NY: Basic Books.

Sega (1991), *Sonic The Hedgehog,* Sega.

Sherman, William R. and Craig, Alan B. (2019), *Understanding virtual reality: Interface, application, and design,* 2nd ed., Burlington: Morgan Kaufmann.

Skalski, Paul, Tamborini, Ron, Shelton, Ashleigh, Buncher, Michael, and Lindmark, Pete (2011), 'Mapping the road to fun: Natural video game controllers, presence, and game enjoyment', *New Media & Society,* 13:2, pp. 224-242.

Slater, Mel, Steed, Anthony, and Usoh, Martin (1995), 'The virtual treadmill: A naturalistic metaphor for navigation in immersive virtual environments', in M. Göbel (ed), *Virtual environments '95,* Springer, Vienna, pp.135-149.

Sony Interactive Entertainment (2018), *Astro bot Rescue Mission*, Team Asobi, Sony Interactive Entertainment.

Steuer, Jonathan (1992), 'Defining virtual reality: Dimensions determining telepresence', in F. Biocca and M. R. Levy (eds), *Communication in the Age of Virtual Reality*, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 73-93.

Swink, Steve (2009), *Game feel: A game designer's guide to virtual sensation*. Burlington, MA: Elsevier, Inc.

Tamborini, Ron, and Bowman, Nicholas (2010), 'Presence in video games', In C. C. Bracken and P. Skalski (eds), *Immersed in media: Telepresence in everyday life*. New York, NY: Routledge, pp. 87-109.

The Tetris Company (1996), *Tetris*, The Tetris Company.

TOAST (2016), *Richie's Plank Experience*, TOAST.

Tunnel Wilson, Preston, Kalescky, William, MacLaughlin, Ansel and Williams, Betsy (2016), 'VR locomotion: Walking > walking in place > arm swinging'. *Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry,* Anaheim, CA, 24-28 July, New-York: Association for Computing Machinery, pp. 243-249.

Virtual Reality Databases (2020), 'VR store statistics', VR Store Statistic, https://vrdb.app. Accessed 5 December 2022.

Zielasko, Daniel, and Weissker, Tim (2023), 'Stay vigilant: The threat of a replication crisis in VR locomotion research'. *Proceedings of the 29th ACM Symposium on Virtual Reality Software and Technology*, Christchurch, NZ, 9-11 October, New-York: Association for Computing Machinery, pp. 1-10.

CHAPTER 2.

EXPLORING REINFORCEMENT LEARNING AS AN APPROACH TO DEVELOPING VR GAMES

MITCHELL FOO, JINGYUAN KANG, & OLIVIA ZHANG

Tags: game development, reinforcement learning, virtual reality, human-AI interaction

INTRODUCTION

More questions need to be asked about the impact deep learning can have on the future of developing games for virtual reality (VR). The rapid development of deep learning technology has had an unquestionable effect on the perception of content creation and design workflows such that we are pressed to ask: how do these technologies change the way that we, as video game developers, artists, and players, reapproach the creation of novel experiences? Especially with the increasing investment in VR technology and adoption, what does deep learning pose to a platform that immerses us more deeply into a virtual context than ever before? The only way to begin answering these questions is not to shy away from the lesser explored; we instead embrace the novelty this opportunity presents by developing a VR game that pits users against AI trained with deep learning.

We explore the potential uses of deep reinforcement learning (RL) for developing game AI that users play against within an immersive VR environment. RL is a method in deep learning that pertains to training an agent on a task in a given environment, awarding positively for desirable actions taken, and punishing undesirable ones based on learning through trial and error. RL's more recent integration with Unity (Juliani et al., 2020) has made it both accessible and prime for experimentation in the video game development space. In this paper, we document our exploration into using RL as a tool for developing a game in VR. Being a team with design, art, and engineering experience, we together reflect on both the technical potential of deep learning and its artistic influence by developing a game experience that competitively engages users with RL-trained agents in VR.

CONTEXT

Video games and virtual environments are not solely means of entertainment for human users. Since Deep Mind's 2013 paper announcing the success of using RL to train an AI capable of outperforming a human expert in three Atari games (Mnih et al., 2013), video games such as Starcraft II (Vinalys et al., 2017), Dota 2 (OpenAI et al., 2019), and plenty of others have become the benchmark for the cutting edge of RL models. And with each model and each newly benchmarked game, it is always asked, how much better is the agent than a human player?

With the recent surge of AI relevance, the question of deep learning's capabilities has never been less questioned, and the same trajectory of development has also occurred in the field of RL. So much so that it can be expected that if an agent can learn how to play a game of reasonable complexity, it likely is going to be able to outperform a human player with enough training. So why is it essential that we can create a near-perfect AI? Would it surely not be fun to play Pong against a faultless opponent? Well, we do not need to have a perfect AI trained with RL it could be almost human-like to a fault — which is where we consider imitation learning.

Imitation learning is a technique used in RL to help supplement agent training by providing human demonstration to the agent on the task it is trying to learn. Simply, instead of an agent starting from scratch with no knowledge of the task, it can draw from prior human examples to learn what works in an environment far faster and potentially with greater performance. But what if there is more to just boosting agent performance? Would the agent, as a result, also then do the task similar to that particular human demonstration?

Debnath et al. (2017) showcase the use of Generative Adversarial Imitation Learning (GAIL) (Ho & Ermon, 2016), a more sophisticated imitation learning algorithm, to the problem of training a virtual humanoid to walk in a physics simulation. Using human walking motion capture data, they showed that the agent trained on those demonstrations performed the task with more qualitative similarity to the human than the agent trained without. Hence, the demonstration has a large effect on the outcome agent's behavior. Furthermore, we might begin prioritizing the appearance of the agent's behavior alongside the task itself.

We are now situated at an intersection between video game development and RL, and can only wonder when RL might start to give back to video game development. Currently, from discussions with peers and online development communities, there is a consensus that an ML-based approach overcomplicates the traditional development of game AI, not being friendly to the iterative design process of game development and often too unpredictable. That being said, we would argue that more work can be done in exploring RL's impact on developing game AI.

Specifically, imitation learning allows us to develop agents to learn implicitly human-like behaviors that otherwise would be hard to code explicitly. We consider that in the case of applications to VR, detailing agents with human demonstration might have a more significant impact on the immersive user experience than traditional screen-based experiences. Some helpful precedents begin to examine the potential for VR to affect the RL training process, namely Koganti et al.'s (2018) work on creating demonstration data with a VR headset, and Gao et al.'s (2019) *VRKitchen*, a task-oriented environment allowing agent and VR user to embody the same virtual avatar. Once again, we see VR explored for RL development, leaving a gap in the condition where RL enhances a user's VR experience. We assert that ultimately the agent's task is not primary, where instead, the user gameplay experience is, that what we strive for is an AI that contributes to the primary goal of building something that simply is — fun.

THE SCIENCE PROJECT MEETS VR GAME

Our goal for this project was to build a fun and immersive VR game that could also serve as an exploration into the potential for RL and imitation learning to augment the development of game AI. Our game, *Imitation Ball*, was designed to consider how a VR game, unique from 2D games, can help us create novel and more immersive experiences with a game AI. From this project, what we can observe is how RL can start to shape our game design process in unique and unexpected ways, both technically and visually.

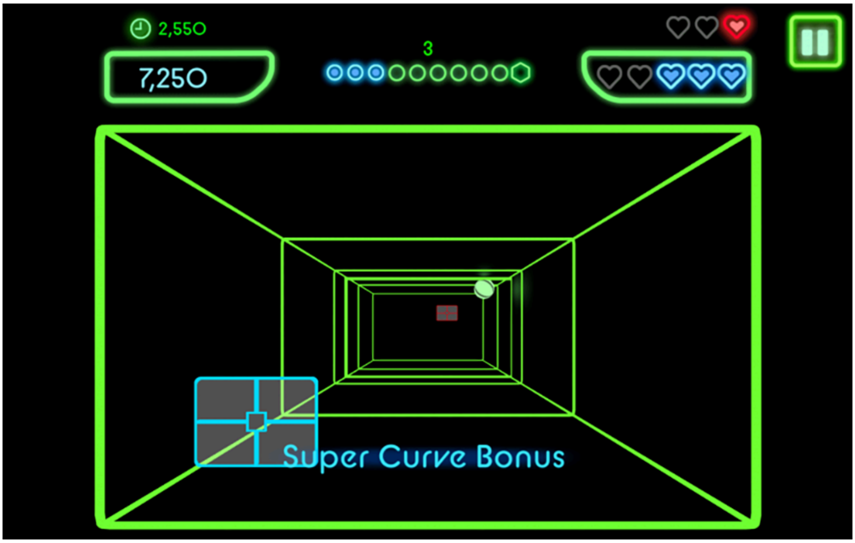

Image 3.1: In-game capture from the flash game Curve Ball 3D

The original idea for *Imitation Ball* was to develop a game that allowed humans and agents to interact with each other through a competitive setting while utilizing similar gameplay controls. The game we decided to model ours after to fit this criterion was a VR hallway-style Pong, most similar to the popular flash game *Curve Ball 3D*. In this game, the user controls a 2D paddle that deflects and curves a 3D ball, bouncing off of the hallway's walls and back to the opponent. Players play against an opponent's AI paddle to try not to concede any balls, similar to Pong. We propose that bringing this more 2D experience into VR space would create a gameplay experience that is both fun and immersive from a human perspective while also posing enough potential for the AI to express an understanding of gesture, depth, and movement.

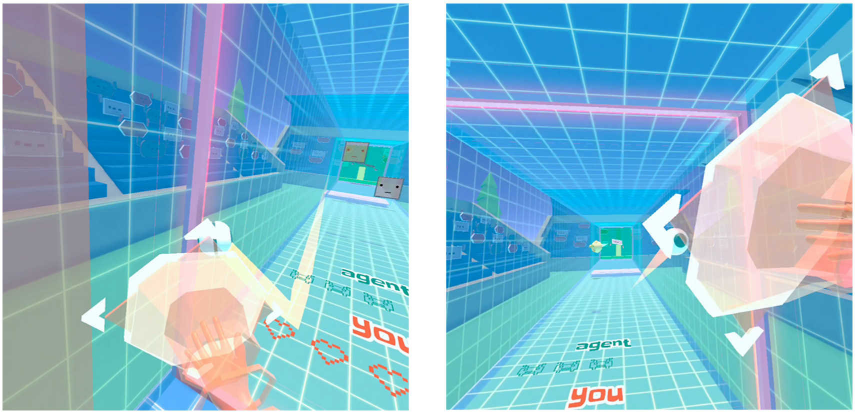

Image 3.2: In-game capture from the final build of Imitation Ball

The core VR gameplay of *Imitation Ball* sets players in the middle of a rectangular hallway facing off against an AI paddle opponent at the opposite end, rallying a bouncing ball back and forth. Players stand and move freely within a box at the end of their hallway and cannot move into the hallway. A flat paddle is affixed across the right hand of the player meaning that players do not need to grasp or press any buttons on their controller, but only watch the movement of the ball and place or swing their paddle to meet it as it approaches near. A successful impact between the paddle and the ball sends the ball in the direction back toward the AI opponent, and conditions like the angle of the paddle upon contact affect the return trajectory of the ball. If the ball passes by the player and reaches the surface behind them, then the ball is reset to the middle of the field, and the player loses a visible life. Players can expect the ball to come from all angles and heights, requiring a moderate amount of active movement, like reaching for corners and crouching for low balls.

In implementing the gameplay, it originated as a 'training gym' before an actual player-vs-enemy environment. The closest comparison would be to playing tennis off of a wall as a means to hone your instincts before hitting the court. A ball spawns in the middle of the hallway, launched randomly, deflecting its bounce against any wall collider. When the ball nears, the agent or player bounces the ball off of their paddle back against the wall at the opposite end of the hallway. If the ball passes, the ball simply resets to allow for continual practice. This training gym is used to teach the agents the core fundamentals of playing our game. It is also used for creating human demonstrations that will later augment agent training.

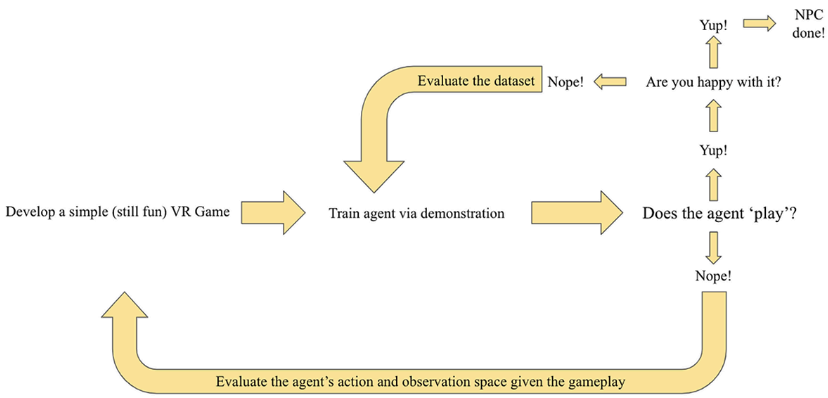

Image 3.3: Initially developed flowchart that frames the goals of this project

We wanted to use our training environment to develop distinctly different behaving agent opponents. In keeping with arcade traditions, we wanted to develop an easy, medium, and hard agent with a behavior associated with each. Through these agents, we wanted to explore the effect that a human demonstration would have on the trained behaviors using GAIL, in contrast to an agent trained without the inclusion of demonstration. As a result, the process of using RL as a design methodology has an element of unpredictability. We determine what goes in, but we will have to see what comes out.

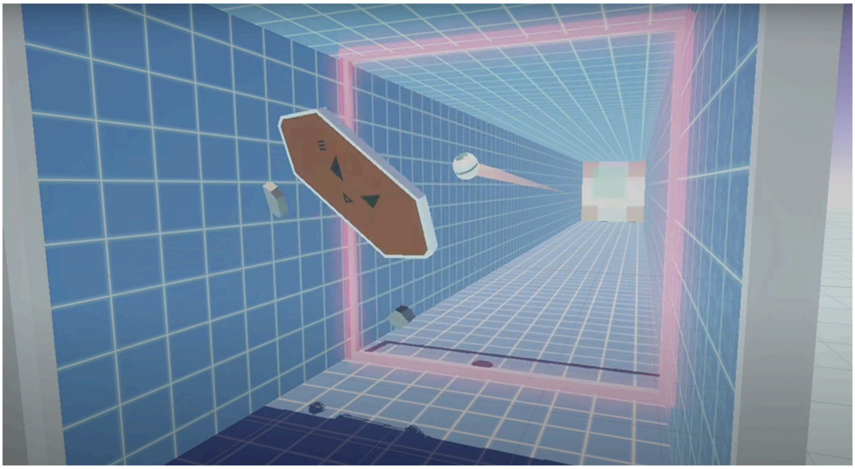

Image 3.4: Agent being trained with RL in the training environment

To clearly describe the objective for the agent, we need to determine what the agent can do, what it can observe, and what its rewards are within the training environment. In RL, we are training a feed-forward neural network that acts like the brain of the agent. It is fed a vector representing an observation within the environment. It then learns to output an optimized corresponding action vector, which is converted to a directional input for the agent. This is similar to how we press buttons on our gamepad to move our game character, given what we see on our screen.

What the agent observes is a 13-length vector that includes its quaternion rotation (4), its position (3), ball position (3), and ball velocity (3). In essence, the agent can see its orientation in the environment as well as its relationship to the ball and the ball's trajectory. The agent, in response to the observation, outputs a 6-length action vector, which includes directional input for position (3) and Euler rotation (3).

Condition	Reward Value
Contacting the ball	+0.5
Contacting the ball at the center region of the paddle	+2
Conceding, or letting the ball past	-2
Remaining (*x, y*) centered in the hallway	+0.01

Table 1: Agent Reward Values Given by Responses in the Training Environment

For the agent's reward, we can base it on a set of conditions during training as seen in Table 1. Notice that the nature of these rewards implicitly teaches the agent how to play our game, albeit defensively. Just try and hit the ball, and do not let it pass. You might also note the rewards for hitting the center region of the paddle and a light reward for remaining in the center of the hallway. During training, we noticed that we did want to suggest certain stylistic behaviors to the agents and were able to do so suggestively via its reward conditions.

With the ability to train an agent, what about our demonstration? A step is needed to translate our actions in VR to one the agent understands. We are not just recording our actions directly but instead controlling the agent in VR, like a puppet, showing them how they should be playing the game. The beauty of VR, from a human's perspective, is that we freely and intuitively use our natural hand motions in the virtual environment without considering directional inputs. Unfortunately, it would be a lot to expect the agent to have such complex sensibilities. So instead, demonstrating to an agent requires heuristics that compare the current agent's pose to the pose of your hand control. If your hand is rotated more right and is above the agent, the agent is given a corresponding input to realign itself according to your hand.

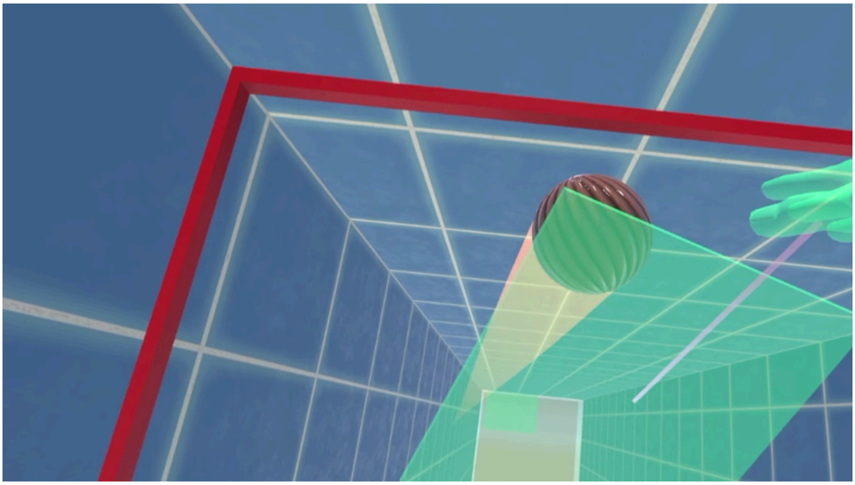

Image 3.5: Human user using VR controls to demonstrate to the agent how to play the game

Having recorded about 45 minutes of human practice in the training environments, we can begin experimenting with the three desired agents and crafting their behaviors. Table 2 helps describe the differences between how the agents were trained based on iterations of tuning that determine the weight of the demonstration and the agent's rotation and movement speeds.

Agent Name	**Aimee**	**Gail**	**Daimian**
Intended Difficulty	Easy	Medium	Hard
Demonstration Weight	0	0.1	0.02
Movement Speed Factor	2	3	25
Rotation Speed Factor	100	250	1000
Training Rounds	1	1	3

Table 2: Three Agent Variants and their Training Variations

The rotation and movement speed are used to add explicit parameters to help us characterize these agents. For instance, by making Aimee move slower, it would in theory, make it both harder for the agent to play the game and longer for it to learn an optimal strategy, like setting a deliberate handicap. Gail was our designated agent for seeing if a heavily weighted demonstration would reflect in its behavior. Lastly, we wanted Daimian to be incredibly tough to play against, such that we trained it three times longer than Aimee and Gail.

From our design method, we can qualitatively observe distinct behaviors and playstyles between the three agents as described:

- Aimee: Lethargic. Slow to move and most prone to letting the ball pass. Does not rotate and slowly drifts to the ball.
- Gail: Flippy. Often spins in full rotations as the ball approaches, actively moving across the hallways. Gets to most balls but is prone to error, especially when it flips and knocks the ball backward.
- Daimian: Efficient. After too much training, Daimian does not let any ball pass. It does not need to move quickly aside from the occasional quick dart to the ball due to expert anticipation skills. Does moderate swings when returning a ball.

What we can conclude from our developmental process is that approaching AI development through implicit training creates unique and desirable agent behaviors. Instead of explicitly determining states and animations, these agents learn to act according to the situation and appear like they are naturally reacting, especially the agents trained with GAIL. We consider it likely that Aimee does not flip around because there is no incentive to; it is not expressed anywhere in any reward. In contrast, in the demonstration, the human swung at balls and reacted with hand rotations, qualities we see in Gail and Daimian and find give the AI a natural-looking behavior.

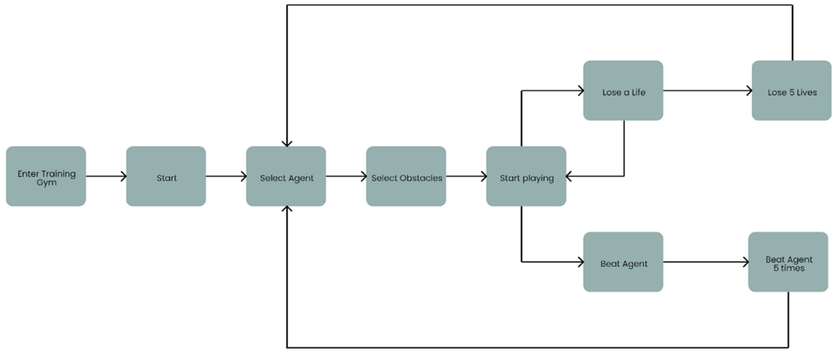

Image 3.6: Final gameplay loop for Imitation Ball

For the final build of our game, Imitation Ball focuses on celebrating, observing, and evaluating these novel agents, exploring their effect on our VR gameplay experience. We developed a simple gameplay loop around these agents allowing for playtesters to give us feedback on the behaviors they observe. Both agent and player are given five lives, which they must protect to beat their opponent.

We also included obstacle selection in our hallway to make it even more interesting, one of which includes a special fourth agent. Because we enjoyed our three agents so much, we had the idea of training an agent that disrupts and causes mayhem in the middle of the hallway; we call this agent the monkey. Aside from being an additional fun feature, by adding obstacles and other agent types, we want to suggest further potential interactions and complexities atop games developed with RL.

Having developed these agents with unique behaviors, we also need to do them justice by giving them a visual identity. Part of creating a game with RL allows us to tackle another question regarding how we want to extend and provide these agent personalities within the dimension of VR.

CREATING A VR AND AI IDENTITY THROUGH ART AND DESIGN

Alongside the implementation of RL into our gameplay design, we took this as an opportunity to visually define our own unique design philosophy and art direction within our VR environment. Central to our vision is a gaming experience where players engage with AI opponents who have undergone extensive training through thousands of playthroughs and observations of human gameplay. Early in our development process, we identified that emphasizing the thrill and novelty of "dueling an AI" should be our primary focus for game design and art direction.

Immediately, we wanted the player to be confronted with their AI opponent, understanding that it possesses equal abilities and limitations. That being said, some of the decisions in development to tune the speed parameters of the agents remain as invisible modifications to the player. We want the perception of the behaviors of the agent to be as if they were not explicitly predetermined. The agents are not just a level to beat but defend their own lives like the player.

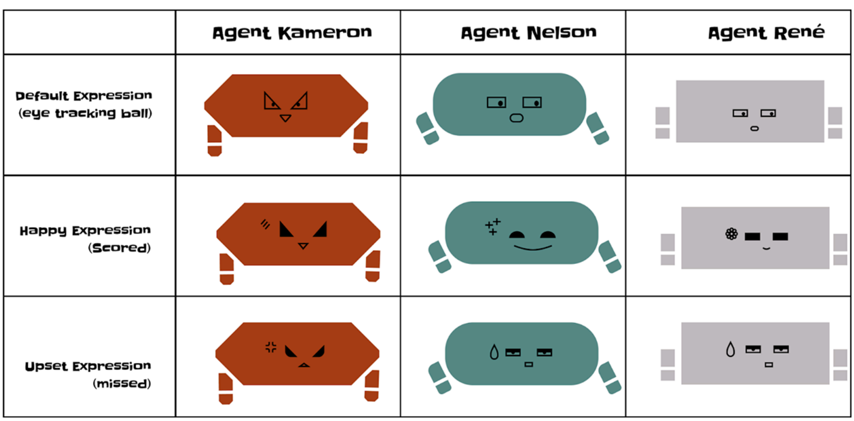

Image 4.1: Expression system of the AI agent variants

Regarding art direction, our objective was to imbue the AI opponents with distinctive characteristics through 3D modeling and sound design. On the one hand, AI opponents should be considered real opponents by the players. On the other hand, there should be a clear line to distinguish those AI agents from humans. To achieve this, we developed a playful identity for the agents as paddles featuring facial expressions and hands, which react dynamically to the game's progress. The hands of these AI opponents wave and swing as they move within the game environment. These features are a response to what we discovered from the behaviors, embellishing what we discovered from our training with RL. Our additional chaotic feature, the 'monkey in the middle,' makes the best example of how dangly hands are a great accompaniment to our RL agents.

Image 4.2: 3D model of the 'monkey in the middle'

For game environment design, we created a stadium populated by numerous AI characters that resemble the player's opponent. This helps establish a sillier and light-hearted theme in contrast to the technical aspect of our exploration. Highlighting to the player thematically that we are exploring the domain of professional AI paddles in their home field helps remind them that AI and deep learning can be contextualized in a user-friendly and game-oriented way.

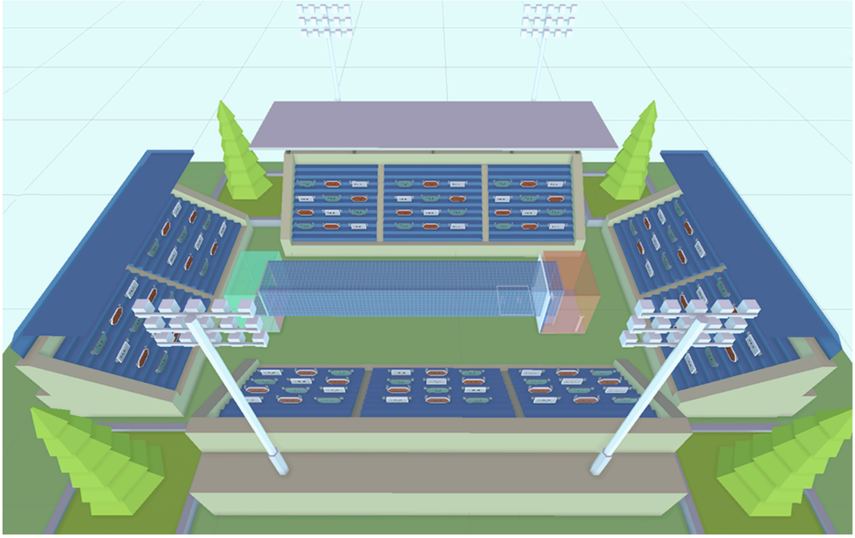

Image 4.3: 3D stadium environment contextualizing the core gameplay arena

For the sound design, we opted for an unconventional approach by eschewing the typical physical impact sounds when agents contact the ball. Instead, the agents produce peculiar and amusing noises upon hitting or missing the ball, as if they were living creatures. Our inspiration for creating these sounds was based on *Animal Crossing*, which strikes a balance between human and AI speech that distinctly gives each character a unique personality. Each of our three agents has its bank of sounds that is randomly played given an explicit situation, like returning the ball. Furthermore, the player can hear cheering and disappointment sounds from the crowd, but the crowd is cheering for your AI opponent, not you. This plays a role in further enhancing the player's immersion and engagement with the AI opponents in their world.

Our approach to level design diverges from traditional linear progression. Instead, we offer players the flexibility to select their opponents based on varying difficulty levels and choose their preferred arena, complete with customizable obstacles. This design choice deviates from the conventions of arcade games, enabling players to concentrate solely on the gameplay experience and the behaviors of their AI opponent. This level design philosophy fosters an environment in which players can focus on the core gameplay mechanics, fully immersing themselves in the challenge of overcoming the AI opponents without the distraction of unrelated game elements.

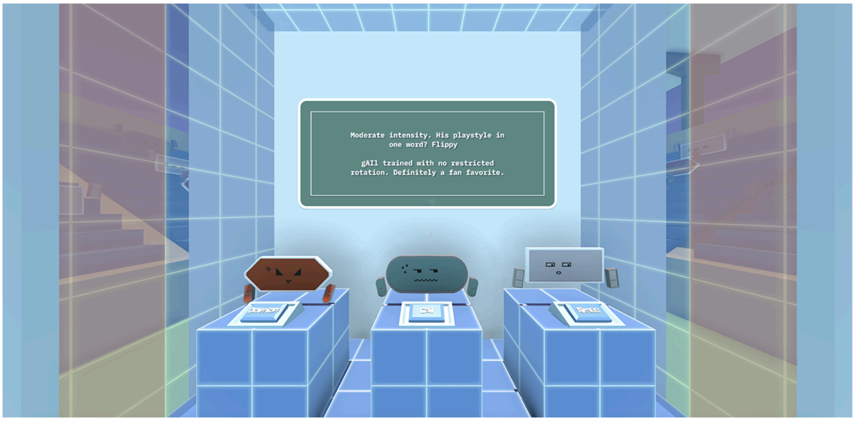

Image 4.4: Menu selection screen for picking agents

Image 4.5: Menu selection screen for picking hallways obstacles

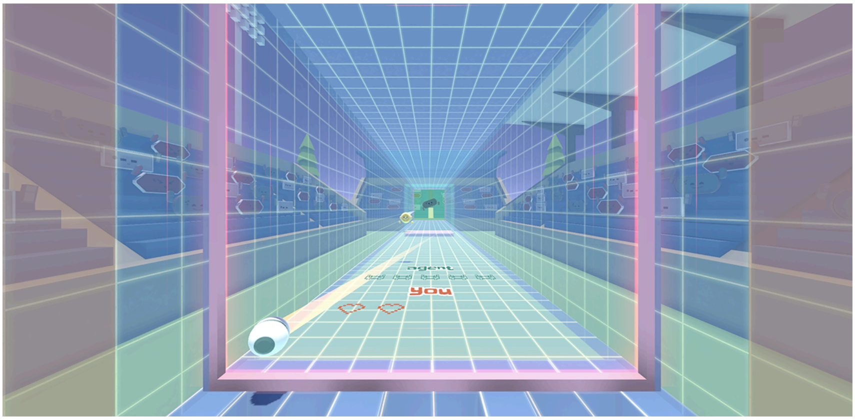

Image 4.6: First-person view of core gameplay

Overall, the art, design, and interactions are built on the same foundation: to create a mindset for players that they are immersed in a playful environment focused on exploring and meeting their new AI opponent.

RESULTS AND PLAYTESTING

Our initial prototype contained an empty hallway environment and three agents that users were able to select and play against. We used an Android build for Quest 2, allowing for a mobile VR experience. Before playing, users were given a summary of the purpose and design of the game. During gameplay, they were asked a series of questions while engaging and were encouraged to think out loud. Each user had varying levels of progress in learning the controls and "mastering" gameplay such that they were able to score and beat an agent. Thus, most of the initial feedback was focused on how to incorporate and modify UI elements such that the lives and scoring are more intuitive to the user. We had already incorporated elements to keep track of the ball through a trail and a motion-tracking light which worked well for users. In general, this is more traditional gameplay feedback.

After the controls were mastered, most users steered their attention toward the interaction between themselves and the agent. Those who had a deeper understanding of ML gave feedback in terms of how to train and design the agents based on that process, whereas general users gave feedback on how the agents and environment could be designed to be more engaging through suggestions such as music, more distinct playstyles between agents, and environmental assets. We observed that users enjoyed the minimal sound effects that were incorporated and saw the potential to elevate the experience by incorporating more audio effects.

We playtested with users with a variety of experiences in VR, from those who have played VR games avidly to first-time users. Users were generally engaged due to the variation in the ball's location and agent behavior. However, we did find that first-time users were more intrigued by being in VR and their abilities to bounce the ball rather than paying much attention to their agent opponent. Furthermore, users began to disengage when they couldn't score against the agents due to a lack of familiarity with the VR controls.

Evaluating the agent's dynamic facial expressions and animations we hypothesized that it would distinguish the agent 'personalities' or play styles and greatly enhance user experience. Unfortunately, technical limitations due to the Quest 2's limited resolution made it such that in playtesting, the playtesters could not observe the agents' expressions. Instead, feedback from more experienced VR users commented that the personality of agents was well conveyed by the motion and playstyle of the agents, describing them through characteristics such as "flippy" or "aggressive."

Based on the feedback and our observations, we redesigned our game and incorporated two more rooms with different assets to give users more choice and control over the difficulty. The most interesting observation was with the 'monkey in the middle' agent due to its unique design and playstyle in its attempt to intercept the ball, subverting a lot of the gameplay's predictability for both the user and AI agent. We further wanted to distinguish the agent's behaviors, increasing the difficulty of our hardest agent, as well as adding brief descriptions of the different agents and environments to better prompt the players in the menu selection phase.

For this iteration, user feedback focused on the enjoyability of gameplay in terms of the variation in difficulty. We observed that users began trying different techniques and motions to defeat the agents; there was more engagement overall due to the time and effort it took to either score or lose against an opponent. In essence, the stakes felt higher. Regardless of whether users read the descriptive text provided for each agent, there was more commentary on the distinct personalities of the agents.

Another piece of feedback from external playtesting highlighted the hallway's narrowness, which confined players. After assessing the issue, we realized that simply enlarging the hallway would be impractical from a human factor perspective, as it could create unreachable corners for the player's controller while AI agents could still cover the entire area, potentially resulting in unfair gameplay. To address this by designing a transparent hallway, allowing for the greater stadium environment to be seen. This solution maintained the original dimensions of the playable area while giving players the impression of being situated in a much larger space. Subsequent playtesting demonstrated that this approach was highly effective in enhancing the player experience.

DISCUSSION LOOKING BACK

The project began more focused on exploring training agent AI with RL; however, with the majority of our team being more design-oriented, our shifted focus also highlighted questions of enhancing the visual user experience in the context of VR, game design, and AI. Throughout the design process, research and multiple iterations of the UI and models were needed to address the challenges that VR brings compared to traditional interfaces. By keeping the mechanics simple with a focus on elements that encouraged user engagement and exploration with AI, we could constantly refine and improve our game while capitalizing on each team member's skills and interests.

Reflecting on the final gameplay and experience of our playtesters, we were able to create an engaging VR experience that many users enjoyed and felt took advantage of novel VR controls. With a simple and intuitive gameplay loop and artistic polish to the game, demoing the final build had many users treat the experience as a proper game rather than an 'ML science project' that was core to *Imitation Ball*'s development at the onset. In other words, the agents trained with RL visually and functionally fit well into our design ambitions for *Imitation Ball*.

From our point of view, the gameplay enhancements from the RL-trained agents were a success as the spatial VR experience felt more immersive and entertaining due to the dynamic and unique personalities expressed by our three AI opponents. Due to the dimensionality of VR and the expressive movements of the agents, observing the agents play the game alone was captivating. In a game that required fluid movements from the AI, we could not have imagined being able to animate the AI opponents deterministically better than the final trained agents.

In terms of future iterations on *Imitation Ball*, we would consider incorporating a greater variety of agents that had been trained on individual team members' demonstration data. This would help us see the capabilities of GAIL in suggesting a design process for creating unique dynamic behaviors through the demonstration itself. Additionally, we would add more mechanics to gameplay, such as power-ups, fast-balls, or curve balls, to further increase variation and gameplay complexity, similar to that seen in the *Mario Tennis* series. This would be both more interesting to the human player and provide a more complex situation for the agent, where in addition to learning continuous movement, it must learn to take discrete offensive actions to try and win the game.

LOOKING AHEAD

At the end of our project, the question becomes, would we make another game that used RL that was not just a science project? Our answer is, of course! Both functionally and exploratory, there are several concluding reasons we would continue developing VR experiences with RL. That being said, it is important also to note the limitations of this approach when it comes to game development.

The development of *Imitation Ball* showcased a shift in approach to developing game NPCs that point toward future design methodologies that leverage deep learning technology. Where typically animators and artists must meticulously handcraft deterministic animations to create immersive NPCs, RL, in this context, replaces keyframed animations with a more implicit approach. We, too, consider that in VR especially, realistic 3D NPC animations need to be more detailed and organic given the nature of its immersion. With further integration with VR human demonstrations, it is possible to imagine NPCs behaving more 'human-like' via a process that does not require much deterministic animation but rather through tuning environmental training parameters.

That being said, one of the main difficulties with using RL for our game was finding out what works to make the agent learn the task in the first place. One has to find an observation, action, and reward function for the agent that implicitly considers all the possible states the agent might find itself in. Furthermore, say we train several agents to a desirable level after a couple of hours, but we want to alter a significant aspect of the game environment. This means that we have to retrain the agents again and potentially modify their training environment as well. This makes implementing AI with RL in games less flexible than traditional methods of AI that rely on explicit heuristics that can be altered and tested instantaneously. So is there a right context for RL where the payoff is worth it?

Though RL might be challenging to fully adapt as a core solution for a complex and responsive AI system, we speculate that training movement behaviors can be worthwhile for specific applications in VR. Instead of thinking of an RL agent learning how to beat a complex game, the objective can be more directed toward smaller tasks that help increase a user's spatial immersion. For example, consider an open-world town that needs a population of pedestrian NPCs. One can imagine replacing a handcrafted path-finding AI system with an RL-trained agent tasked with moving toward a certain marker while avoiding key obstacles. The benefit is that the RL agent can be trained with different human behaviors to create more realistic appearing and naturally reactive NPCs. Furthermore, slight adjustments to different training parameters, similar to how we created our three AI opponents in *Imitation Ball,* can craft different NPCs that do the same task but in behaviorally different ways. Adding these subtle but essential details is a promising scope for RL to keep exploring.

At a wider level, what this work highlights is the need to pursue more explorations on the emerging relationship humans and deep learning are developing, given VR technologies. On one hand, the experience of interacting with AI models in VR is a question of representation and canniness while on the other, using VR headsets and controls creates data that, in turn, can be used to influence how we train deep learning models. Though we found that using our motion data captured in VR can be a compelling means to design and animate AI NPCs, it should also be noted that the data being captured about our bodies during VR experiences can also be used for any number of machine learning applications. As much as we see AI in VR, it also 'sees' more of us.

In conclusion, as VR is expanded into a platform for immersive experiences, the question of how interactable NPCs become extended by deep learning technology is pressing and relevant. It is necessary to start confronting how these experiences will be developed given these future developments in gaming and technology. We have only provided one small project that scratches the surface of creating immersive experiences with deep learning, but who knows what *crazier* things are yet to come?

Are we ready for it?

REFERENCES

Debnath, S., Devos, A., Heiden, E., Julian, R., & Khatana, F. (n.d.). Human Behavior Imitation Learning. Distill. Retrieved April 30, 2023, from https://uscresl.github.io/humanoid-gail/

Gao, X., Gong, R., Shu, T., Xie, X., Wang, S., & Zhu, S.-C. (2019). VRKitchen: An Interactive 3D Virtual Environment for Task-oriented Learning (arXiv:1903.05757). arXiv. http://arxiv.org/abs/1903.05757

Ho, J., & Ermon, S. (2016, June 10). Generative Adversarial Imitation Learning. ArXiv.Org. https://arxiv.org/abs/1606.03476v1

Juliani, A., Berges, V.-P., Teng, E., Cohen, A., Harper, J., Elion, C., Goy, C., Gao, Y., Henry, H., Mattar, M., & Lange, D. (2020). Unity: A General Platform for Intelligent Agents (arXiv:1809.02627). arXiv. https://doi.org/10.48550/arXiv.1809.02627

Koganti, N., Rahman H. A. G., A., Iwasawa, Y., Nakayama, K., & Matsuo, Y. (2018). Virtual Reality as a User-friendly Interface for Learning from Demonstrations. Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems, 1–4. https://doi.org/10.1145/3170427.3186500

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing Atari with Deep Reinforcement Learning (arXiv:1312.5602). arXiv. https://doi.org/10.48550/arXiv.1312.5602

OpenAI, Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., Pinto, H. P. d O., Raiman, J., Salimans, T., … Zhang, S. (2019, December 13). Dota 2 with Large Scale Deep Reinforcement Learning. ArXiv.Org. https://arxiv.org/abs/1912.06680v1

Play Curve Ball 3D: Curve and Score | Coolmath Games. (n.d.). https://www.coolmathgames.com/0-curve-ball-3d

Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K., & Hassabis, D. (2017, December 5). Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. ArXiv.Org. https://arxiv.org/abs/1712.01815v1

Vinyals, O., Ewalds, T., Bartunov, S., Georgiev, P., Vezhnevets, A. S., Yeo, M., Makhzani, A., Küttler, H., Agapiou, J., Schrittwieser, J., Quan, J., Gaffney, S., Petersen, S., Simonyan, K., Schaul, T., van Hasselt, H., Silver, D., Lillicrap, T., Calderone, K., … Tsing, R. (2017, August 16). StarCraft II: A New Challenge for Reinforcement Learning. ArXiv.Org. https://arxiv.org/abs/1708.04782v1

ACKNOWLEDGMENTS

A massive thank you to Thomas Corbett from CMU's Entertainment Technology Center for instructing us through his class, Research Issues in Game Development: Designing for XR. His guidance and insightful feedback played a critical role in shaping our project and learning experience.

CHAPTER 3.

MAKING COPIES

The Creation of a Virtual Reality Sculpting Game

JONAH WARREN

ABSTRACT

The following is a personal narrative describing the of development of Choppy Copies (Warren, 2021), a game about recreating abstract versions of famous classical sculptures made for virtual reality headsets. While I took a few creative liberties in this retelling, I did my best to accurately depict the game's development by referring to notes, interviews, screenshots, code, and GitHub comments from the project. At the time of this writing, Choppy Copies has not been released. It has been publicly exhibited several times, however, and was nominated for Best Live Action Game at IndieCade in 2022.

Additional note: In my experience developing games, I have found that a lot of important realizations happen during moments of private reflection, informed by personal experience. In post-mortems, these things often get overlooked in favor of more dramatic discoveries made during playtesting or in discussions between team members. My hope is that by retelling the story of Choppy Copies from a first-person perspective, I can highlight some of the decisions that happened quietly and might be challenging to fully articulate or appreciate if told from another perspective.

INTRODUCTION

It is the summer of 2021. I flip through a notebook trying to decipher sketches and notes I hastily scribbled months or years earlier. Phrases like "Two-Handed Simultaneous Grid-Based 3D VR Blek" and "Collaborative Wario Ware, 4-Person Multi Mouse Games" written in all caps pop from the pages. As I scan for promising ideas, I repeatedly say to myself, "What was I thinking?" Sometimes I genuinely have no idea what I was envisioning. Sometimes I do.

It is one of my favorite times of the year. Classes are done, grades are submitted, and there are a few independent game festivals with deadlines a month or two away. I am trying to decide which game idea to prototype. In my capacity as a professor of Game Design and Development at Quinnipiac, I recently finished teaching a course in Augmented and Virtual Reality for Games where I realized how much more I have to learn about designing for these spaces. I also just got an Oculus Quest 2. I decided I should make something in VR.

After being hesitant to work in 3D for most of my career with VR, I finally feel excited to do so. No projecting 3D space on a 2D screen. No awkward mappings of discrete button presses to 3D interactions (e.g., "Press F to Pay Respects" [Wikipedia, 2023]). No overly complicated camera systems trying to anticipate where players will look next. Players can see something in three dimensions, reach out, and interact with it. They can control the camera by turning their head. Six-degree-of-freedom controllers and head tracking naturally enable full-bodied game experiences.

As I think about what to make, I ask myself a few questions. What does this mean for me as a designer? What unique experiences can I, only now, create? What areas are least explored?

One path would be to embrace real-world interactions. Like Job Simulator (Owlchemy Labs, 2016), I could throw the player into an environment with a wide variety of realistic tools and familiar objects. Embrace the fact that players have a natural, intuitive way to interact with these things. No lengthy tutorials, no "getting the hang of" the control scheme, just go do stuff. Figure it out! Play!

Another path is to create something not of this world. Something that could only happen in a virtual environment. Like Beat Saber (Beat Games, 2019), I could place the player in an abstract world filled with blocks and patterns. Give them tools designed specifically for interacting with and navigating the challenges I create. No need to create a world that feels real. Design every object in the game and every piece of feedback so that it supports the VR-specific gameplay.

I decided to start by thinking about the interactions I want to explore and go from there. In this case, that means those unique to VR. I should create an experience that somehow *combines the visual perceptual skills and dexterity required to effectively move through 3D space and affect it.*

I realize I should be more precise. Considering the current state of VR's tracking technology, "moving through 3D space" is limited to head and hand movement a few feet in any direction. What experiences would work well within these constraints? I think of Tai chi, the martial art practiced for health and self-defense. I think of William Forsythe's Improvisation Technologies (2012), a vocabulary of dance movements I bookmarked a few months prior. While these seem like interesting perspectives to explore, most of the related game ideas I come up with feel very prescribed, without much decision-making or meaningful choice.

I continue flipping pages in my notebook. I see "Sloppy Sculptures." It is a riff on Sloppy Forgeries (Warren, 2018). Sloppy Forgeries is a fast-paced local multiplayer painting game I made a few years prior where players are given 90 seconds to recreate famous paintings from art history, armed with a mouse and a few simple paint tools. I wrote the idea quickly, without a sketch. I remember the idea floating around my head for a while before I even decided to write it down. It was more of a joke than a possibility. It felt too obvious, too close to something I already made.

Image 1: A screenshot of Sloppy Forgeries.

Upon seeing it again, I immediately envision a voxelized classical Greek sculpture, Muppet-like frantically flailing arms, and debris flying everywhere. I start getting excited. Thinking back, I now realize that one of my favorite student projects from that AR/VR class, a voxel-based prototype entitled Barbershop Simulator VR (Umelo, 2020), helped inspire this thought. The experience is what you might imagine: a head with a lot of hair and tools to help you artfully get rid of it—kind of like sculpting.

Although I have a degree in visual arts and took a course in the subject as an undergraduate, sculpture is not a medium I feel very comfortable with, especially in stone or marble. As a kid with shaky hands and an interest in art and computers, I spent hours in front of paint programs with one hand on the mouse and the other resting on Command-Z, drawing and undoing, drawing and undoing, until I got the exact line I was looking for. The idea that one misplaced cut could ruin hours of work and an expensive piece of marble sounds horrible.

Perhaps it is the combination of the context in which they are often viewed, the translucent marble, or the scale, but there is also something other-worldly about classical sculpture that interests me while thinking about a "Sloppy Sculptures" game. One of my favorite things about Sloppy Forgeries is the dramatic contrast between the formal, pristine masterpieces and the ridiculous, impossible task of recreating them in a minute and a half (not to mention the resulting crude reproductions). The contrast in "Sloppy Sculptures" seems even more pronounced: masterpieces are even more immaculate, a task even more impossible.

I realize that a VR game about copying sculptures also fits my interactive goals. It certainly *combines the visual perceptual skills and dexterity required to effectively move through 3D space and affect it.* It seems almost perfectly suited to take advantage of the affordances of VR. Sculpture is one of the most primal, basic, and intuitive ways of interacting with and creating three-dimensional forms.

Decision made. I get to work.

DEVELOPMENT

My first challenge is to get an abstract version of a famous classical sculpture into Unity. After a quick search, including some combination of "classical sculpture," "3D," and "model," I stumble upon Scan the World (MyMiniFactory, n.d.), a website that describes itself as the "world's largest ecosystem of free to download, 3D printable objects of cultural significance." Looking through the models, I am attracted to iconic works like Michelangelo's David and Auguste Rodin's The Thinker. I also find myself drawn to Myron's Discobolus. The dynamic form, outstretched arms, and negative space seem like the right kind of challenge. I start with these three.

I abstract the models using voxels. Voxels, which many may associate with the game Minecraft (2009), is a method of representing a 3D form through a collection of cubes aligned to a regular grid in three-dimensional space. It is essentially a 3d version of pixels and bitmaps. The higher the resolution of the grid, the more realistic the form becomes. The lower the resolution the grid, the more abstract.

Creating voxel versions of the models is surprisingly easy—there are a lot of free tools that do this. The tougher question is, what resolution? In Sloppy Forgeries, the paintings the player copies are abstracted to make things easier. Each painting has been posterized, reducing the color palette to four or five colors. I quickly learned that this abstraction suited some paintings better than others. Not surprisingly, Henri Matisse's The Dance works much better than Seurat's A Sunday on La Grande Jatte. I also discovered that I often needed to posterize an image over and over again and occasionally even rework it by hand to find a result that felt right.

I play around. Voxels too big and there is little resemblance to the original. Voxels too small and I worry it may start affecting performance or feel too overwhelming. I end up realizing that, like Sloppy Forgeries, it is a challenging thing to standardize. What feels and looks right for one sculpture may not work for another. I spend quite a bit of time finding the right balance, ensuring the abstracted sculptures are still recognizable. The process also cements my hunch that resolution could be a good way to adjust the game's difficulty, although keeping sculptures recognizable also seems important. I make a mental bookmark.

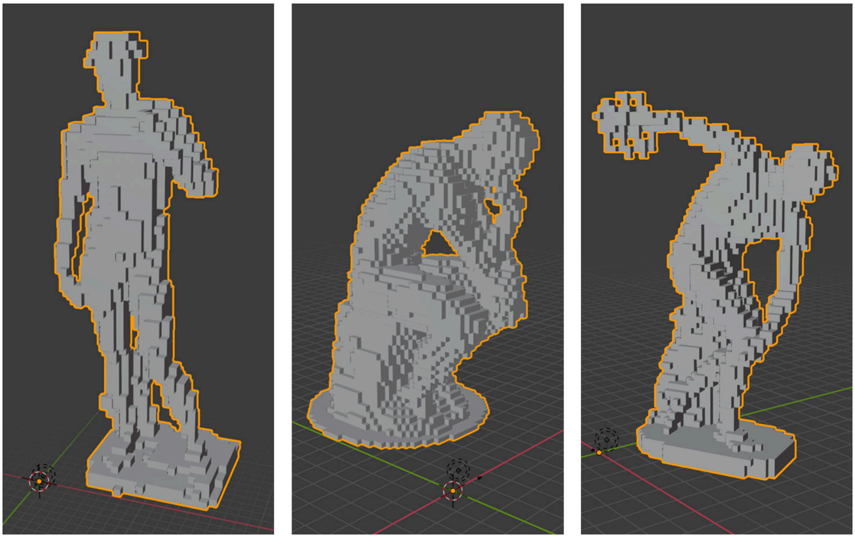

Image 2: Voxelized versions of Michelangelo's David, Auguste Rodin's The Thinker, and Myron's Discobolus.

Over the next few days, I concentrate on carving. I start with the simplest interaction I can think of. Give the player a sphere for a hand that destroys any voxel it hits. I try it out on an array of cubes and as expected, performance is an issue. Some forum posts about generating Minecraft-like landscapes made me realize that the solution is to generate the sculpture's mesh dynamically and update it as the player carves. Drawing only the sculpture's external faces speeds up things considerably.

Next, I give the player the ability to rotate their sculpture. When I first envisioned Sloppy Forgeries for sculpture and a flailing Muppet spraying voxels everywhere, the block they were sculpting was rotating. This was likely because of Barbershop Simulator VR. In it, the player could teleport or move around the head of hair to cut from different sides. In a critique of it, I mentioned swiveling barber chairs. I rarely get motion sickness from VR, but why include movement or unnecessary teleporting if you can avoid it?

This decision gives me pause since allowing the player to rotate the sculpture sacrifices realism in favor of ease of use, which goes against a common design philosophy in VR development. It favors diegetic interfaces over non-diegetic ones, where the more decisions and actions you can embed in-game (versus HUDs or abstracted UI), the more immersive the experience (Salomoni et al., 2016). However, this approach can be taken too far if it leads to an experience contrary to a designer's goals or negatively affects player experience. Plus, the more I think about it, magically teleporting or floating through space via a thumbstick seems just as unrealistic.

If I allow the player to rotate their sculpture, rotating the original sculpture simultaneously seems only natural. It should make it easier to keep the two in sync and emphasize the importance of their orientation. In Sloppy Forgeries, accuracy is determined simply by comparing each pixel in the original painting with the copy. If the color is the same, count it as correct. The percentage of correct pixels versus the total number of pixels determines the score. My plan is to do the same with voxels, but just counting on/off instead of color. The only problem is, with an additional dimension, it is not as obvious which side is front and which is back. I decided to also add guide arrows to indicate the forward and up vectors for both the original and copy sculptures to help. I add a bounding box as well.

I map rotating the sculptures to left and right on the joystick of the controller in the player's sculpting hand. I also allow the player to reposition their sculpture's location on the floor by mapping this movement to the joystick in the player's other hand. I do this because, while playtesting, I often find myself out of position, unable to properly carve. The sculpture would feel too close, too far away, or I'd creep out of the defined play space and hit a table with my hand. A few hours of repeatedly recentering my headset, adjusting the guardian, and fiddling with the starting position has me yelling, "Just let me move it!" which is shortly followed by, "Oh wait, I can do that."

I try the new controls. Movement and rotation feel great; however, I notice my head repeatedly moving back and forth, right to left, left to right. I look at the original sculpture, then at my copy, then my hand carving it, and back to the original. It is awkward. Line-of-sight is important. I realize that there are three important elements to consider, not two: of course, the copy and original sculptures, but the player's carving hand as well. After some fiddling, the best orientation seems to put the copy in front of the player and a touch to the right, and the original sculpture offset back and to the left. Although my arm is often across my body when carving, my copy, the original sculpture, and my hand are now simultaneously in my field of view. The copy also does not obstruct the original sculpture. It feels right.

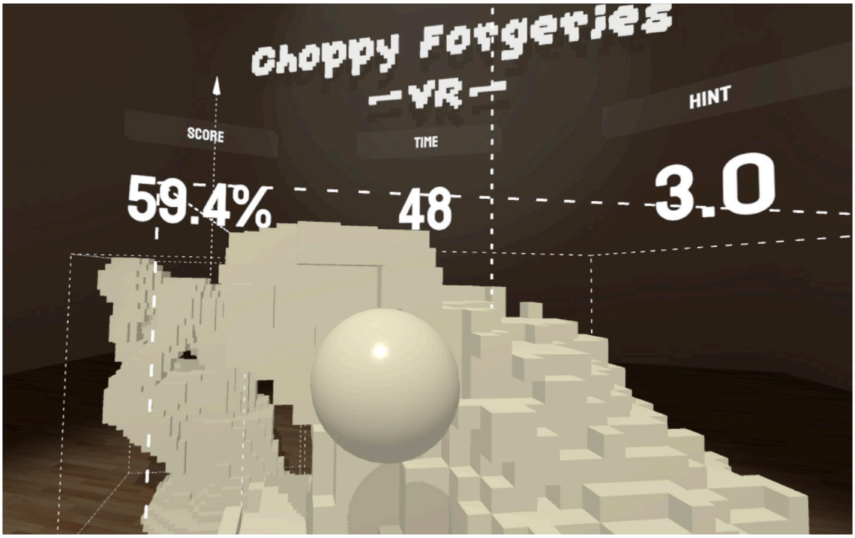

Image 3: A screenshot of Choppy Forgeries during play, demonstrating the player's line-of-site.

Then it hits me—this only works if you're right-handed! I put my head in my hands. The arrangement I just arrived at is important. I need to support left-handed use, which means being able to change the sculpture arrangement and controller mappings dynamically, and some UI or a menu to facilitate switching back and forth.

An important lesson I learned when preparing for the AR/VR class is that UI elements in VR work best when anchored to a plane in space. This means that before I make menus, I need the space they will exist in. Where does this all take place? I create a tall, empty gallery space, placing the player in the center of the room facing three floating panels, which allow the player to choose a sculpture, start a game, or configure play. I also include a (currently non-functional) resolution menu and tack on a fourth panel with instructions since I'm not sure where else to put it. I place "Choppy Forgeries," the game's new title, in an extruded pixel font above the panels.

Image 4: A screenshot of the Choppy Forgeries menu.

I reuse this same three-panel arrangement in both configure mode and during play. In configure mode, the panels allow players to switch hands and provide instructions for how to move and rotate the sculptures via the controllers. It also indicates the location of the sculptures on the floor during play. When switching hands, the orientation of the sculptures and the instructions change, so players can see the effect of their choice. They can also practice using the controls described and watch the indicators move and rotate.

During play, I have the left panel show the player's accuracy score as a percentage, the middle show the amount of time left in the round, and the right the amount of "hint" left, displayed in seconds (similar to Sloppy Forgeries). "Hint" allows the player to quickly hold down the trigger button to see the original sculpture in place of their copy for a limited amount of time, making it easy to spot differences. I chose this arrangement since it puts the most important information directly in the player's field of view as they carve. As this is flipped in lefty mode, I decided I should flip the panel order too. It all works, except for the "hint" time. It is often obscured by the player's sculpture. Unfortunately, a submission date is approaching fast, so I just need to live with it for now.

During the last few days before submission, I finish things up. I adjust the time given for each round (two and a half minutes feels right). I add some fast-paced piano music (the third movement from Beethoven's Moonlight Sonata) from the fantastic Classical Piano Midi Page (Krueger, n.d.). I change the "time left" display to a "reset" button when a round is done and slowly rotate the sculptures, emphasizing the state change and encouraging reflection. Finally, I take screenshots, record a trailer, make a website, upload a build, complete the submission forms, and take a break.

A few months later, I have the opportunity to test the prototype informally at an assembly meeting for game design students and faculty at Quinnipiac. The timing works out well. Although the game was not accepted to any festivals, I was invited to exhibit it at an academic conference, now a month away. Hopefully, I can get some good feedback, fix some bugs, and iron out any kinks before the show.

Since the assembly is during the middle of the semester and playtesting a last-minute affair, there is little time to prepare. I boot up the game and watch students and faculty flail around. Overall, I am disappointed. Players are overwhelmed. There is too much to look at, too many controls. Once they do get the hang of it, recreating the sculpture seems like a monumental task. Most are frustrated and not in a good way.

I am not sure why this felt unexpected. It was a rookie game design mistake: underestimating difficulty. I had spent months slowly and incrementally adding new features, practicing my VR sculpting skills along the way. Of course, they were overwhelmed. One student gets into it though. He masters the controls after some practice, plays several rounds, and makes some great observations. He mentions the game's difficulty, his frustration at hitting the floor when sculpting, and acknowledges some awkwardness when changing the size of the carving tool. I make a note that he has a lot of experience with VR.

This leads to another somewhat obvious observation, albeit more nuanced than I anticipate: a player's previous experience with the technology will dramatically affect a VR playtest. What surprises me is how much this variation would affect my experience conducting the playtests. Since it was so relaxed and informal, I sometimes forget to ask the playtesters, "How much experience do you have with VR?" and adjust my expectations. Some players need a lot of help with the basics. Some just get it. Some even realize they dislike VR. I forget to keep track. I realize that quickly noting past experience is important, even in informal playtesting. Not only is it helpful in analyzing feedback, but also in being aware of how it affects your perception of a given playtest and the game you're developing.

Regardless, I learn a lot and the experience kickstarts development again. I make a list of issues and solutions. I start small. I add pedestals to both the original and copy sculptures, lifting them up so players avoid whacking their controllers on the ground. I make a 3D grid marked by white lines on both sculptures, giving players more points of reference when comparing. I add sound effects when carving for additional feedback, to help players understand how many voxels they are removing. I revise the written in-game instructions. I rework the overly complicated method for resizing the carving tool.

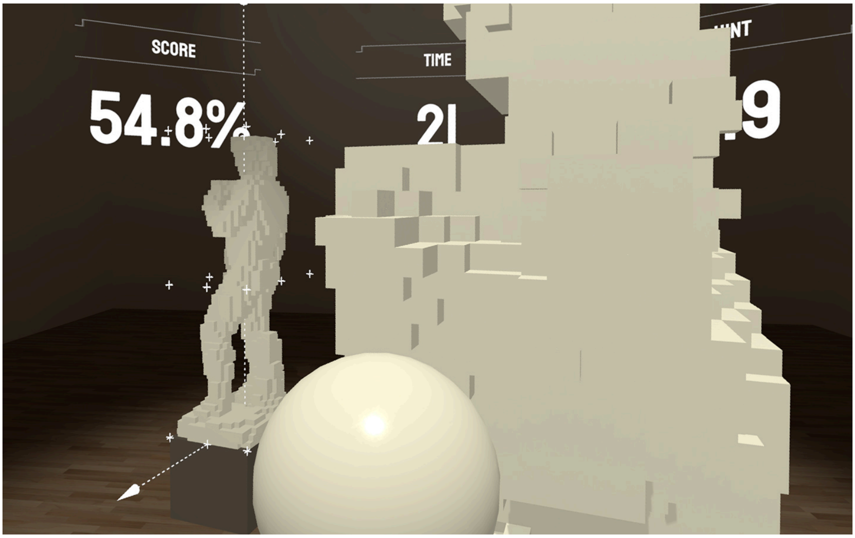

Image 5: Choppy Forgeries gameplay with pedestals and a grid added.

After trying it out, the game is still too hard. I need to explore lower resolutions. I realize that I should let go of the idea of the arbitrary requirement that sculptures must be recognizable. In fact, the mystery might even be compelling. I make a very blocky Discobulous at a resolution of 7 x 10 x 4 voxels, load it in, and see how it feels to copy.

It makes things easier. I get 100% after several tries, which is encouraging. However, surprisingly, it changes my experience of the game. At the higher resolution, where the default carving tool size is larger than any individual voxel making up the sculpture, it is difficult to remove individual voxels one by one. It would also be a bad strategy. To be successful, you need to remove large swaths of sculpture quickly. Fluidly waving your tool through those areas is the much more effective. However, at low resolutions, removing individual voxels works well. The interaction becomes less continuous, more discrete. I find myself poking at voxels one at a time.

I am initially taken aback by this. This is not the game I wanted to make. It feels obsessive and fiddly, like you need to get each voxel right before moving on to the next. One of the things I like about Sloppy Forgeries is that the scoring system incentivizes players to paint large swaths of the background color over the details of the foreground. It encourages the player to consider the entire canvas first, to not sweat the small stuff. This feels like the opposite.

But the more I play at low resolution, the more I see its value. When a sculpture is this abstracted, any association with what it represents is tenuous. It feels more like copying an arrangement of cubes in space than a human form. When working from life, this kind of dissociation can be beneficial to a beginning artist. There is a common saying in art education when drawing and painting from life, which also happens to be the title of an introductory book on drawing (Shirley, 2012): "Draw what you see, not what you think you see." Unconscious ideas of what an artist thinks a face or a body looks like have a way of seeping into an artwork, often with negative consequences.

Starting players off with abstract sculptures could help them practice this disassociation. After mastering the controls and recreating simple forms, gradually increasing the resolution would slowly require players to grapple with more detail, more recognizable forms, stronger associations, and a more challenging, fluid interaction. The progression makes sense in both game difficulty and art-making instruction contexts.

Most importantly, it resolves the difficulty issue, which was significantly limiting the game's audience. With these realizations, I let go of my aesthetic urge to keep all the sculptures at a consistent resolution and my pre-existing conception of the game's initial experience. I create low and medium resolution versions for all three sculptures and add them to the game. With the upcoming exhibition only a few days away, I create a build, put it on the headset, and halt development.

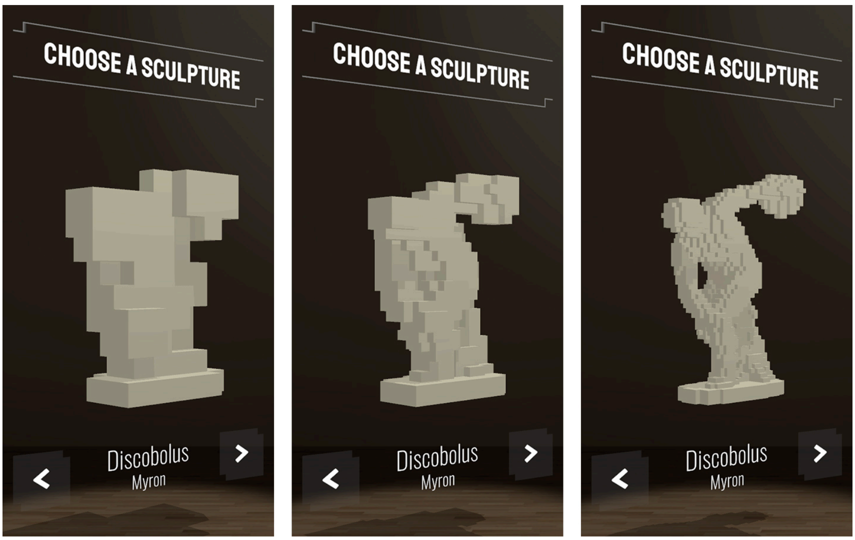

Image 6: Discobolus at low, medium, and high resolutions.

I figure the exhibition could be another good playtest opportunity, however, it does not go according to plan. First, only a handful of people play it, and second, I forget the link cable in my rush to get there. I figure this is fine, since I have the most recent version loaded on the headset, ready to go. What I forget, is that without the cable, I am unable to connect to another monitor and see the player's perspective as they play. I have no idea what players are experiencing.

I do learn a lot, though. Without the ability to see what the player is doing; I find it impossible to casually help them as they play. While there are instructions explaining the controls before the game, it is a wall of text few people read. I find myself answering a lot of questions about controller mappings. Through conversations afterward, I also discover that people have vastly different experiences of the game. I slowly gather this is because of the different resolutions. Some try low resolution first, some high.

These two issues have straightforward solutions, though they'll take some time to implement. First, to teach the controls, I need a tutorial. Although I was able to avoid one in Sloppy Forgeries, there are just too things to do in this game, too many controls. Second, to ensure players complete lower resolution sculptures before moving on to higher ones, I need some sort of progression system.

I do some research and find another competition with a submission date a few months away. I decide to use this as a deadline for completing these last two items.

During spring break, I tackle the tutorial. I want something quick, but clear. I place the player in front of the sculptures as they would be during play with a panel close to them, on the left. Sequentially, with concise language, I add text to the panel prompting the player to try each action. Upon completion, I check off the prompt and gray it out, leaving it there in case the player completes it by accident or needs to practice it again. When they finish, players can continue in this mode to practice carving, or return to the main menu.

Image 7: A screenshot of Choppy Copies' tutorial.

A week or two before the submission date, I make the progression system. Although the abundance of thoughtless achievements in games makes me hesitant to just add a list of tasks to complete, recently helping my daughter finish Untitled Goose Game makes me realize how effective a thoughtful list of tasks can be in incentivizing the exploration of a game space.

I decide to lock all the sculptures and resolutions when any player first starts up the game. I put a "Current Objective" section in the main menu screen right above the play button, with the first objective below it: "Complete the tutorial." I land on the following sequence, which progressively unlocks more sculptures and higher resolutions as objectives are completed.

- Completing the tutorial unlocks the play button and low-resolution Discobolus
- Scoring 70% on low resolution Discobolus unlocks all other low-resolution sculptures
- Scoring 80% on any other low-resolution sculpture unlocks all medium resolution sculptures

- Scoring 80% on any medium-resolution sculpture unlocks all high resolution sculptures

The last thing I do before submitting is rename the game. Although I like Choppy Forgeries, the more I work on it, the more it feels like its own game, deserving of its own name. I also find that people hear the name, expect a multiplayer experience like Sloppy Forgeries, and come away disappointed. I land on Choppy Copies. I redo the screenshots, trailer, gameplay footage, and the website.

A few months later, I hear back that the game is nominated for an award, which is a pleasant surprise. It feels good after all the effort. Though I wish things were different, due to a busy academic schedule, the game has not changed much since then. At some point, I would like to add a gallery, a high score list, and publish it. However, whatever happens, it has been a valuable, worthwhile journey.

Like playing Choppy Copies, making games can be messy, intimidating, and hard. It can be tough to know where to start, where to look, and what to prioritize. Often, the best way to overcome the paralysis all of this can cause, is to jump in, start making decisions, and trust that with some persistence and close observation, things will get better. Hopefully, you have some fun along the way, and, in the end, you can be proud of what you made, including all its wonderful imperfections.

REFERENCES

Job Simulator (VR version) [Video game]. (2016). Owlchemy Labs. https://store.steampowered.com/app/448280/Job_Simulator/

Beat Saber (VR version) [Video game]. (2019). Beat Games. https://store.steampowered.com/app/620980/Beat_Saber/

Forsythe, W. and Haffner, N. (2012) William Forsythe: Improvisation Technologies: A Tool for the Analytical Dance Eye. United States: Distributed Art Pub Incorporated.

Krueger, B. Classical Piano Midi Page. (n.d.) Retrieved June 15th, 2023, from http://www.piano-midi.de/

Minecraft [Video game]. (2009). Mojang.

Salomoni, P., Prandi, C., Roccetti, M., Casanova L., and Marchetti L. (2016). Assessing the efficacy of a diegetic game interface with Oculus Rift. 2016 13th IEEE Annual Consumer Communications & Networking Conference (CCNC), 387-392. https://doi.org/10.1109/CCNC.2016.7444811

Scan the World. MyMiniFactory. (n.d.) Retrieved June 15th, 2023, from https://www.myminifactory.com/scantheworld/full-collection

Shirley, R. (2012). Draw What You See Not What You Think You See. United States: CreateSpace Independent Publishing Platform.

Umelo, K. (2020) Barbershop Simulator VR. (VR version) [Video game] https://kene991.itch.io/barbershop-simulator-vr

Warren, J. (2018) Sloppy Forgeries. (PC version) [Video game] https://playfulsystems.com/sloppy-forgeries/

Warren, J. (2021) Choppy Copies. (PC version) [Video game] https://playfulsystems.com/choppy-copies/

Wikipedia. (n.d.) Retrieved June 15th, 2023, from https://en.wikipedia.org/wiki/Press_F_to_pay_respects

CHAPTER 4.

LESSONS LEARNED FROM DESIGNING IN MIXED REALITY

MARIE LEUNG

Mixed reality (MR) is an umbrella term for everything on the reality-virtuality continuum that lies between the real world and virtual worlds. On the far end of the continuum, virtual reality (VR) is a fully virtual representation of a 3D environment that a user can interact with immersively. VR is quite well-known – my grandparents can explain it to me. As for the closer end of the continuum, I daresay reality is quite popular as well. Falling within the range of mixed reality, augmented reality (AR) can be seen in many places, with virtual elements used to augment the real world. Augmented virtuality (AV) is less commonly pointed out, but involves putting real elements in a virtual space. To enhance mixed reality experiences, camera passthrough and other optical solutions on head-mounted displays are becoming more common, and depth sensing is improving on these devices as well.

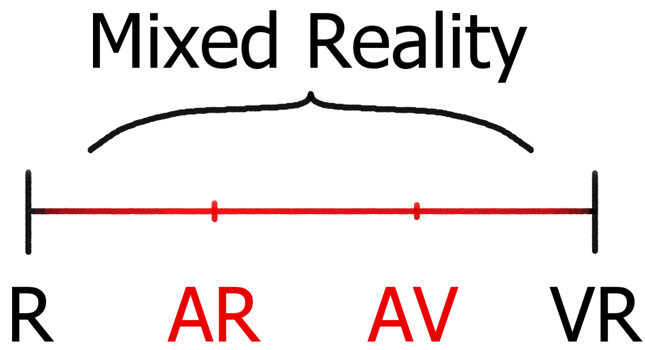

The RV continuum, introduced in 1994 by Milgram and Kishino, is a scale used to frame virtual and augmented reality (R=Reality, AR=Augmented Reality, AV=Augmented Virtuality, VR=Virtual Reality). Though this taxonomy has since evolved with the technologies, it still provides sufficient explanation for the experiences we create today (Skarbez, Richard et al.).

Why did I choose to design for mixed reality? It's a hard question to answer.

"Presence," I say. The feeling of being immersed in the space you're experiencing. Mixed reality allows for a unique juxtaposition of virtual interactions while remaining grounded in a simulation of what you perceive to be the real world.

"Taking advantage of physical space," I add. If your real-world environment is large enough, a wider variety of interactions can be woven into the real space around you.

"And honestly," I admit, "I just want to play with new technology."

On top of mixed reality's benefits, I believe there is much more to be discovered. To me, mixed reality extends beyond the equipment we put on our heads. It's a design space that necessitates new ways of thinking, one that holds potential in its nebulous definition.

DISCOVERING THE MIXED REALITY DESIGN SPACE

In my first year studying at Carnegie Mellon's Entertainment Technology Center, I pitched a mixed reality experience intended to help people overcome their social anxiety. I wanted two socially anxious strangers, each in the safety of their own rooms, to have virtual interactions with each other's avatars. Given the unique combination of presence and grounding provided by mixed reality, I hoped to create positive social interactions in a low-stakes space that felt real enough to help people gain the courage to interact with others in the real world.

To kick off our project, my team brainstormed what to include in our experience. Through many conversations with industry professionals and shared reflection as a team, we learned a lot about what kinds of ideas work in mixed reality.

I want to share questions you may consider when brainstorming for mixed reality:

- What do people normally do together in a shared space?
- What does mixed reality enable that could not otherwise occur in a real or purely virtual space?
- How can a virtual layer help people interact with each other?
- How do my virtual elements integrate with reality, and vice versa?
- How can virtuality help people share one or more physical spaces?
- Throwing away all existing notions of what happens in VR, AR, and the real world, what do I want to experience?
- For my ideas that have been conflicting with constraints, how can I work around the problem, solve it, or use parts of the idea in other solutions?

And for refining the ideas you have brainstormed:

- Why does this idea need to be implemented in mixed reality?
- How can I improve this idea's connection to mixed reality?

Our team also played *I Expect You To Die: Home Sweet Home* on the Quest Pro headset to learn about an existing use case of mixed reality. Part of a puzzle game series where you play as a secret agent trying to complete and survive your missions, *Home Sweet Home* includes one level that makes use of camera passthrough and the headset's ability to designate basic room features in your play space. The game demonstrates some ways to play with mixed reality. The player starts out in a box, only able to peer outside at their real world space through the small windows on the box. There is also a fire sprinkler virtually mounted on the player's ceiling, which involves using the real world ceiling designated during the device's room space setup.

After playing and speaking with the game's design director, we learned that part of the reason *Home Sweet Home* is simpler compared with other games in the *I Expect You To Die* series is that working in mixed reality means the developers are no longer fully in control of the environment the player sees. When we design for mixed reality, we want interactions to encompass reality and virtuality to make better use of the medium. However, we also have no idea where players will be when they launch our game, making well-integrated interactions difficult. This complication that we discovered by playing *Home Sweet Home* also affected our own mixed reality project and was the first of many challenges we encountered in the mixed reality design space.

OUR DESIGN'S CHALLENGES AND SUCCESSES

Once my team decided on an idea for our mixed reality project – a party simulator based on exposure therapy techniques – we began to encounter greater design and development challenges. Almost all of our challenges can be traced back to us having an unclear vision of what is effective for our target audience in a mixed reality space. Interestingly enough, a lot of our challenges can also be thought of along the reality-virtuality continuum because the level of realism needs to be decided for each element in our design.

We built a party simulator experience in mixed reality to help those with social anxiety.

As mentioned before, adjusting our experience to the players' environments became a huge challenge, especially for our developers. Because our players' real play spaces are all anticipated to be different, our experience needs to automatically configure itself to fit varied room layouts. We attempted to solve this in the first iteration of our project, which had two players in separate rooms. We proposed a synced room space system based on key furniture landmarks. For example, if both players have a couch and a chair, and Player A walks from their couch to their chair, their path might not make sense when seen in Player B's room configuration. Player A might appear to walk through

Player B's coffee table or even a wall. However, if each player marks the location of their respective couches and chairs in their rooms, as well as their other furniture and walls during room setup, we can create a logically walkable path for players' avatars to traverse from the couch to the chair. Though the concept never made it to our final prototype, we believe that this is a strong example of combining reality and virtuality.

Another issue we had with accommodating the player's environment was appropriately populating the interior of any room with non-player characters (NPCs). Since our experience simulates a party with lots of people, we have to ensure all NPCs are procedurally placed within the bounds of the real-world room and not in furniture or each other. This also applies to all other objects placed in the mixed reality environment, not just NPCs. They are placed procedurally to fit the room space designated during individual players' room setups. We learned that ideally, anything that is placed in the experience should take advantage of the mixed reality environment but not depend on static configurations, both programmatically and narratively. Thus the objects are virtual, but their locations are dependent on reality.

A magical moment we successfully achieve by actually using our player's real environment is the crossing of the threshold from the outside to the inside of our room. As someone with social anxiety, I can anecdotally attest to the unease felt when needing to enter a room where the situation holds a lot of uncertainty. Many people who tried our party experience were surprised to see the real room full of virtual NPCs the moment they passed through the doorway. This moment would not work so well in VR because the player would be immersed and expecting a filled virtual room to be behind a virtual door, but in MR, some part of the player is grounded in reality and does not always expect crazy things to happen. Of course, this depends on many factors like player psychology, but for now, we consider this a fun little trick. Remember that you may not be limited to the confines of your room when designing your experience.

As you can see, the environment that a mixed reality experience takes place in is very important to the experience itself. When we tried to fill that physical space with character models and user interfaces, we also had to ensure that the virtual elements we created fit the surrounding realistic environment.

Character design in a socially focused experience turned out to be much more nuanced than we expected. Our goal was to have characters that were human enough to induce social anxiety and not any other stray feeling, like aversion of an unpleasant 3D model. During design and playtesting, we alternated between needing to make our NPCs more realistic and more stylized. We originally thought that people would feel more comfortable interacting with cartoonish characters, but playtesters expressed the desire to speak to more human-like NPCs and thought that would improve their association of our experience with real-life situations. Next we tested realistic NPCs, but playtesters thought they were too uncanny and shouldn't be so realistic. Bouncing back and forth between realism and stylization, we finally found a balance in characters stylized similar to those found in popular animated movies. Though the NPCs were still overlaid on a passthrough feed of the real world, our final level of stylization didn't feel too out of place.

Our user interface (UI) mostly involves dialogue boxes that react to voice and text boxes that can be tapped in 3D space. This type of 3D digital guidance doesn't appear in reality and is typically seen as a benefit of MR and AR, but it has to be used carefully. Poor UI designs, such as unintuitive interfaces that have not been tested with users of different backgrounds, can easily be confusing instead of helpful. UI design is also more complex in 3D, as it has to be visible when needed, but also physically placed in spots that make sense for users. For example, our dialogue options should appear when it is the player's turn to speak in a conversation, but if the player is behind the NPC, it would not make sense for the UI to be visible, as that would imply the player should respond at the back of the NPC's head. Similarly, the player needs an indication of who in the crowd of people to go to next, but blatantly showing dialogue options in front of the correct NPC is immersion-breaking and could potentially imply that the player could speak to the NPC from 20 feet away. In our case, our 3D UI and even the NPC heights adjust to the player's height to achieve our desired impressions. We also keep the UI objects in certain places so as to not jeopardize the players' safety.

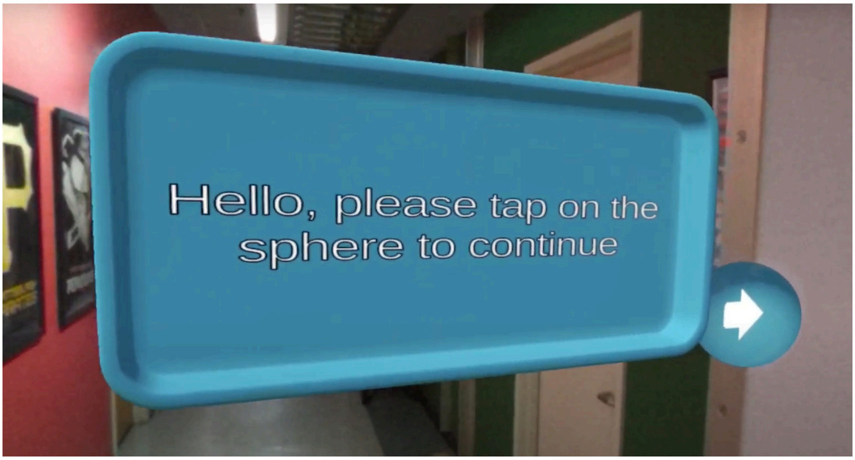

Our UI imitates traditional buttons in 3D and is designed to appear only where it makes sense in the 3D environment.

To best recreate a social situation in mixed reality, we opted to use vocal recognition and hand tracking as means of player input. Speaking out loud in simulated conversations and using one's hands to interact with the environment are much more grounded in reality and fit our goal of being able to connect our experience with real-world situations. These features add to mixed reality's more natural presence that differs from the presence felt in VR.

Vocal recognition identifies words that the player is saying so that we can match it to our dialogue options. Hand tracking identifies the player's hands so they can interact with the virtual world using their hands instead of controllers.

MY THOUGHTS ON MIXED REALITY

By the end of our project semester, we had a working prototype of a party simulator where a socially anxious user would have to walk into a real room, pick up a lost wallet, and talk to virtual individuals in the crowd of partygoers to find out who it belongs to. We spent a lot of time and effort designing appropriately for users with social anxiety and creatively wrangling our new hardware in order to achieve the results we wanted. We playtested to make sure the experience evoked the correct feelings in users, but did not have the time or resources to test the experience's long-term effects on social anxiety. Since our prototype was created in a short timeframe, we ended up developing a system that could be used to create any social experience that fits any room space so long as the dialogue lines and avatar assets are replaced to fit the new needs.

As a team, we learned how challenging it was to create an experience that both took advantage of mixed reality's new features and also made sense within our project's constraints. Because our social anxiety project was focused on helping people, we had self-imposed constraints related to safety, designing according to existing therapy techniques, and validating the effectiveness of our design. Mixed reality was intuitive to people who had never used head-mounted displays or played video games before, which made the medium a good choice for our target users, but the tradeoff was needing to be flexible with people's real environments as a design constraint. We learned that it was difficult to create a novel and intuitive experience due to technological and design constraints unique to mixed reality.

Achieving intuitive understanding of an experience is easier when creating a slightly augmented version of real-world activities that already exist, as we did. What's harder to create is a mixed reality experience where the real and virtual are thoroughly incorporated together, not just coexisting, in a way that still makes sense. The most difficult features for us to innovate on have roots in both reality and virtuality and require integrating the two worlds in a compelling manner. I believe that, even with the technological and design challenges, it is worth iterating on these difficult features. It is infinitely rewarding when players praise them as charming moments of discovery or compliment the intuitiveness of their experience.

Our project files are now archived with local academic communities and we are not planning on further development, but hope that our findings can help others venturing into the space. If we were to restart our project, I would choose to work with mixed reality again and attempt to create new core mechanics that cleverly integrate reality and virtuality in ways that people intuitively comprehend.

Why am I working with mixed reality now? Having experimented with the space, I can confirm that design potential lies in many aspects of a mixed reality experience, especially in visuals and interactions. I believe there are compelling mechanics and more uses of mixed reality beyond games and experiences to be invented. For example, what could you make if a headset captures a user's real environment and procedurally generates a virtual environment based on the sample capture? Technology and ideas are constantly improving, and mixed reality devices will continue to better record and create realities. As a game designer, I want to shape the experiences people have in new ways, and the mixed reality design space provides plentiful opportunities for me to challenge my own thinking and others' perspectives.

All paths have their challenges, and I believe mixed reality is a rocky path full of gold to be discovered. Mixed reality is accessible for most people and opens up worlds of possibilities to developers, designers, and users. If you have a chance to play with mixed reality, give it a try!

REFERENCES

Girdhar, Varun, et al. "Overcoming Social Anxiety with Mixed Reality." TerrARium, 4 May 2023, https://projects.etc.cmu.edu/terr-ar-ium/media/. Accessed 27 Nov 2023.

Skarbez, Richard, et al. "Revisiting Milgram and Kishino's Reality-Virtuality Continuum." Frontiers in Virtual Reality, vol. 2, 2021, https://doi.org/10.3389/frvir.2021.647997. Accessed 24 Nov 2023.

ABOUT THE AUTHORS

Laureline Chiapello
University of Quebec, Game Design
LChiapello@nad.ca

Charles Chouinard
University of Quebec, Game Design
cwchouinard@hotmail.com

Mitchell Foo
Carnegie Mellon University, Computational Design
mfoo@andrew.cmu.edu

Jingyuan Kang
Carnegie Mellon University, Entertainment Technology Center
jingyuak@andrew.cmu.edu

Marie Leung
Carnegie Mellon University, Entertainment Technology Center
mcleung@andrew.cmu.edu

Jonah Warren
Quinnipiac University, Game Design & Development
Jonah.Warren@quinnipiac.edu

Olivia Zhang
Tepper School of Business
osz@andrew.cmu.edu

ABOUT ETC PRESS

The ETC Press was founded in 2005 under the direction of Dr. Drew Davidson, the Director of Carnegie Mellon University's Entertainment Technology Center (ETC), as an open-access, digital-first publishing house.

What does all that mean?

The ETC Press publishes three types of work:peer-reviewed work (research-based books, textbooks, academic journals, conference proceedings), general audience work (trade nonfiction, singles, Well Played singles), and research and white papers.

The common tie for all of these is a focus on issues related to entertainment technologies as they are applied across a variety of fields.

Our authors come from a range of backgrounds. Some are traditional academics. Some are practitioners. And some work in between. What ties them all together is their ability to write about the impact of emerging technologies and its significance in society.

To distinguish our books, the ETC Press has five imprints:

- **ETC Press:** our traditional academic and peer-reviewed publications;
- **ETC Press: Single:** our short "why it matters" books that are roughly 8,000-25,000 words;
- **ETC Press: Signature:** our special projects, trade books,

and other curated works that exemplify the best work being done;

- **ETC Press: Report:** our white papers and reports produced by practitioners or academic researchers working in conjunction with partners; and
- **ETC Press: Student:** our work with undergraduate and graduate students

In keeping with that mission, the ETC Press uses emerging technologies to design all of our books and Lulu, an on-demand publisher, to distribute our e-books and print books through all the major retail chains, such as Amazon, Barnes & Noble, Kobo, and Apple, and we work with The Game Crafter to produce tabletop games.

We don't carry an inventory ourselves. Instead, each print book is created when somebody buys a copy.

Since the ETC Press is an open-access publisher, every book, journal, and proceeding is available as a free download. We're most interested in the sharing and spreading of ideas. We also have an agreement with the Association for Computing Machinery (ACM) to list ETC Press publications in the ACM Digital Library.

Authors retain ownership of their intellectual property. We release all of our books, journals, and proceedings under one of two Creative Commons licenses:

- **Attribution-NoDerivativeWorks-NonCommercial:** This license allows for published works to remain intact, but versions can be created; or
- **Attribution-NonCommercial-ShareAlike:** This license allows for authors to retain editorial control of their creations while also encouraging readers to

collaboratively rewrite content.

This is definitely an experiment in the notion of publishing, and we invite people to participate. We are exploring what it means to "publish" across multiple media and multiple versions. We believe this is the future of publication, bridging virtual and physical media with fluid versions of publications as well as enabling the creative blurring of what constitutes reading and writing.