Inducing Cooperation Through Virtual Reality

Daniel W. Zhang

Submitted in partial fulfillment of
the requirements for the degree of
Master of Science

Department of Computer Science
Courant Institute of Mathematical Sciences
New York University

May 2017

Dr. Ken Perlin

Dr. Andy Nealen

Abstract

There has been a recent resurgence in Virtual Reality (VR) as a new medium for entertainment and communication. With these potentially exciting developments, I decided to create an experiment to test how people can potentially be influenced by virtual reality interfaces. I had hoped that I could induce people to cooperate better in a virtual reality version of a task compared to an un-augmented version of the task in the regular world. After conducting 16 separate trials, with half in VR and the other half in the regular world, there is no conclusive evidence to completely confirm or deny this hypothesis. I have found evidence to suggest that there can be such an influence, as there were more successes in the VR trials than the regular trials, but they can potentially be explained away by the sample size and the attitudes of participants before starting the experiment. This data suggests that further research in this field can lead to interesting discoveries regarding human behavior in virtual reality environments, and that the Holojam framework invented by the Future Reality Lab at New York University can be very helpful in designing experiments for this research.

Acknowledgements

I would like to thank Doctors Ken Perlin and Andy Nealen for agreeing to serve as my thesis readers . Their support has allowed me to delve into a potentially exciting and unexplored field of research. I would also like to thank the Computer Science department at New York University, especially Courtney Miller and Dr. Benjamin Goldberg, for allowing me the opportunity to conduct this research. I extend my thanks to the members of the Future Reality Lab and the Game Innovation Lab at New York University for helping to cultivate the ideas behind my thesis and for their technical support throughout the various stumbling blocks I encountered.

I would also like to acknowledge the help and support I received from Doctors Walker M. White, Erik Andersen, and David Mimno of the department of Computer Science at Cornell University. Without their support, I doubt I would have reached this stage of my academic career. In addition, I extend my thanks to Dr. Ronald L. Seeber and Professors Rocco M. Scanza, J.D. and Christina Homrighouse alongside the various professors I had the pleasure of interacting with at the School of Industrial and Labor Relations for instilling in me the necessary frameworks and knowledge for conducting this thesis.

Last, but not least, I would like to thank my family and friends for their support and encouragement over the years. Without their support and trust in me, I would not have made it this far in terms of pursuing my research interests.

Table of Contents

Introduction

Virtual Reality (VR) has taken off in recent years in consumer and enterprise markets, with Sony, Oculus, and Valve serving as the leaders in the market of high fidelity headsets and multitudes of other companies driving the proliferation of lower cost solutions. Though the concept of VR has been around for many years in academia since Jaron Lanier first proposed the concept in the 1980s, VR as a widely available product and service has only now seen a surge in the public eye, with financial news institutions such as Bloomberg proclaiming a booming industry that will produce billions, if not trillions, of dollars in revenue over the next decade. However, it seems shortsighted to only focus on the potential profits of the industry, as there are many interesting questions to ask about the significance of VR as a permanent fixture in our daily lives. In particular, *Hamlet on the Holodeck*, a book written in the 1980s by Janet Murray, poses fascinating applications of the technology in the space of theater and play, and science fiction media such as *Star Trek* depict speculative worlds where VR is used in the daily lives of people. There is one major difference in how VR is depicted in these media and what we currently have on the market: VR is assumed to be a tool that the average person can incorporate in daily routines.

To clarify, those media have as references holograms and other sorts of technologies that multiple people can interact with at the same time, similar to how multiple people can view a television set without sacrificing parts of the experience. The hardware solutions available on today's market are primarily for single person use, mostly in the areas of arts and entertainment. For example, Oculus and Sony are continually funding the development of exclusive digital games for their proprietary headsets, and Oculus and Valve have introduced software for creating digital art in new ways, with Oculus offering a sculpting tool and Valve, a three-dimensional drawing application. We can even look at Facebook's attempts to introduce virtual environments for long distance communication over the internet. Unfortunately, a grand majority of these efforts at popularizing and populating the VR marketplace are isolating experiences, generally made for the enjoyment and engagement of a single person in the same physical space. This is not an attempt at blaming the market, though, as there are currently no easy and cheap solutions for researching the area of shared space virtual reality.

My thesis is an attempt to kindle interest in the field of how VR can serve as an everyday tool and fixture to affect how people interact with each other on a social level. I am not the first to attempt research in this field, as Dr. Ken Perlin, the head of my research lab, has led the charge over the past few years to encourage different ways of thinking about VR. Dr. Perlin believes that this technology will enrich aspects of our lives beyond the arts and entertainment. For instance, with this technology, we have a way of interacting with four-dimensional objects in the third dimension, something that is impossible in the physical world. Such an application allows us to reason about properties of these objects beyond their standard mathematical equations, and may potentially serve as interesting educational toys for children to augment their thinking.

My research in this field is unique in terms of how VR can affect people's perception of each other and of their tasks in various social environments. If we envision a future where VR is portable and everywhere, perhaps in the form of lightweight glasses, we'd imagine that there are a multitude of applications to exploit the technology for conveniences we could not have had before. In fact, the ideal VR technology would encompass many aspects of Augmented Reality (AR) technology, enabling new innovations such as geophysical waypoints that we can follow instead of needing to refer to a map every few minutes, or input methods that allow us to draw or create virtual objects in midair to aid explanations and lectures. In that case, we are combining objects that are not real in the traditional sense of the word with what we perceive to be real. The release of Pokemon Go in 2016 demonstrated how AR can affect human behavior, with multiple people involved in trespassing incidents and physical accidents as the result of searching for "fake" Pokemon spread across the real world. Given these results, it is not too much of a stretch to imagine how the ideal VR technology proposed above can affect human behaviors.

With these thoughts in mind, I proposed an experiment to examine the differences in how humans would behave in simulations involving resource constraints when placed in a virtual environment compared to the regular world's. Part of my motivations lie in attempting to validate research performed in the fields of behavioral psychology with reference to collaborative and competitive behavior. The other major portion of my motivations is to potentially derive insights into how the design of virtual reality interfaces can influence and affect human behavior, as mentioned above. The rest of this report details the exact motivations and inspirations behind this experiment, going further into the design and implementation of this project. Afterwards, I compile my experiences administering my experiments with various human subjects and my thoughts and conclusions on the results. Due to the nature of this report, I refrain from making final conclusions about the results or statistical validations of the data, as I view this as an initial foray into an area of research that I and many others will be able to further develop in the future.

Motivations

This section details previous work that relates to my study. Due to the interdisciplinary nature of this topic, the sources comprise a mixture of psychology and computer science research papers published over the past few decades and up until as recently as last year.

One of the main motivators behind my thesis is the classic game theory scenario of prisoner's dilemma[1]. To summarize, two criminals in the same gang are placed in separate rooms and asked to betray the other person to commute a potential jail sentence. If both accept, both criminals will face a sentence of two years in prison, and if one accepts while the other declines, the one who accepts goes free while the one who remains silent gets three years in prison. The last case is where both criminals choose to remain silent: both will serve a sentence of one year. The typical game theoretic analysis suggests that it is in the prisoners' best interests to both betray one another, as there is no mode of communication between the two criminals, and hence no potential avenue to make sure that both will choose to remain loyal to one another and not betray the other, making mutual betrayal the dominant strategy to be chosen by rational agents, or the Nash equilibrium.

However, one of the mistaken lessons many people take away from this classic problem is to try and apply the proposed solution to negotiations in the physical world, or behave as if opposing parties in a negotiation will always wish to take advantage of one another. The key to why many make this fallacy lies in the fact that the payoff matrix (the representation of the potential outcomes of the participant's choices in the scenario) almost never corresponds to the actual outcomes of a negotiation; in other words, there is no guarantee that there cannot exist a mutually beneficial solution in a given negotiation or resource-limited scenario.

This belief that one's goal and another's in a negotiation are diametrically opposed to one another is known as the fixed-pie perception, the belief that there can only be a winner and a loser in a negotiation, with no room for mutual gain [2]. De Dreu et al. tackle this fallacy in their paper and present strategies for integrative negotiation, more commonly known as expanding the pie. The general idea is that by having participants in a negotiation exchange information about their motivations, the participants can focus on trying to help each other solve their problems rather than relying on a fixed mindset of what success means to them. A motivating example to illustrate this involves apples and bananas as resources. If person A wants apples from person B, and person B wants bananas from person A, the outcome of the negotiation could potentially be very different if person A reveals that he has a large supply of bananas and if person B reveals that he has a large supply of apples. Instead of needlessly arguing over how much of each fruit each party can give one another, the two people involved can engage in a fair trade of apples and bananas. In the case where both parties choose to refrain from sharing information with each other, both will most likely resort to preconceived notions of "needing to win" in the negotiation, and at least one person will come out empty-handed.

This next paper involves how people's appearances can affect the outcomes of negotiations. In 1991, Professor Ian Ayres of the Harvard Law School published a paper that detailed experiments involving car dealerships in the Chicago area [3]. In the paper, he attempts to determine if there exist potential biases against customers based on their appearances. By sending multiple participants of different genders and ethnicities to various car dealerships in the Chicago area and having each of the participants employ similar negotiation strategies, he discovers that there exists a dramatic difference in the average, final negotiated price among white males, white females, black males, and black females. He includes a table that estimates the average dealer profit, which I have reproduced below, with slight modifications:

Customer Identity	Average Dealer Profit For Final Offers ($)
White Male	362
White Female	504
Black Male	783
Black Female	1237

Though this study would be considered inappropriate for an academic institution to authorize today since the car dealers at the time were not aware of the experiment, it paints a picture of how a person's appearance can affect the outcome of a negotiation.

Ethnicity and gender seem to have played a significant role in the final prices each participant bargained for in the process of purchasing a car, with white males receiving the best deals and black females, the worst. Ayres's work underscores the biases and attitudes that participants carry with them when entering a negotiation. In mentioning potential reasons for why certain combinations of gender and race yield worse prices, Ayres specifically points to the "racial conceptions" that salespeople may bring to the job, especially because of the uniformity of "disparate treatment in different neighborhoods." There are other potential reasons he mentions for this difference in treatment, including the typical purchasing patterns of consumers of specific genders and races, but the information most relevant to this thesis suggests that participants in competitive scenarios will use visual information in an attempt to gain an upper hand when there is a lack of information. As mentioned above, having both parties sharing information in negotiations can increase the possibilities for an integrative negotiation, and this example of a competitive situation demonstrates how the lack of information can cause participants to extrapolate based on external information.

This does not mean that visual information is the key to dominating in a competitive scenario. In the car negotiations, participants were taught to follow a specific script in order to gather data about the potential price they would have to pay in order to purchase the car; in a more realistic scenario, it is likely that the participants could have chosen to take advantage of other tactics to lower the price of the car. More importantly, the unwitting salespeople in the study were exploiting personal experiences in making sales that the participants most likely could not match, in addition to their preconceived notions of how customers of the same demographic would act to maximize their potential profits. How would an average person fare in such a negotiation where one would need to read visual cues to extract information from one another? Bond et al. suggest that relying on such visual cues to glean information from others is usually not that helpful; their analysis of 206 documents and 24,483 judges suggests that on average, people achieve “an average of 54% correct lie-truth judgments,” more specifically classifying “47% of lies as deceptive” [4]. In a proposed scenario where two people are placed in a simulation with little to no information about a task that involves constrained resources, this study suggests that if both untrained participants were to try and read each other’s thoughts via body language and facial expressions, it may not be that successful, especially if the participants are attempting to hide information from each other, such as their personal states of mind. An addendum to this train of thought involves the Kuleshov effect, a psychological phenomena proposed by Soviet filmmaker Lev Kuleshov suggesting that the context in which a person views another person’s face can have a dramatic effect on how the first person perceives the second person’s feelings [5]. Barratt et al. demonstrate in an experiment how they could induce participants to judge a neutral face in a scene by simply changing the clip that takes place the shot of the face. This research suggests that if humans are already somewhat flawed at trying to make judgments based on other people’s appearances, that the time and place can also have a potential effect.

Finally, Peck et al. conducted an experiment that placed light skinned participants into an immersive VR experience and gave a portion of those participants dark-skinned virtual bodies [6]. They found that this experience suggests a decrease in implicit racial bias when comparing their results on a racial Implicit Association Test before and after the experiment. This finding suggests that the virtual representation of a person can affect their thought processes subconsciously, though it is more conclusive regarding one's own appearance rather than when viewing another person's avatar.

With all of the previous work that I collated together, I decided that I wanted to test the following hypothesis: if two people are placed in a potentially competitive scenario involving limited resources and limited information, they may resort to relying on each other's body language and facial expressions to try and extract useful information to use against each other. However, if the psychological research is to be believed, such gestures cannot be reliably trusted, and may potentially be more noisy than helpful in the process. If that is the case, then perhaps it would be more helpful to place the participants in a virtual reality environment where they see each other as virtual avatars, occluding the potential noise and, therefore, potentially inducing the two to cooperate more to find alternative solutions to the scenario. To simplify, I wanted to test the hypothesis that we can induce people to cooperate better in a virtual reality environment compared to in the regular world when given similar tasks.

Design and Implementation

To create an appropriate experiment to test my hypothesis, I had to create a scenario that could be enacted in both the regular world and in virtual reality. This ruled out tasks that involved peripherals such as video game controllers and digital games, as there would be potential issues in replicating the actions performed in both types of experiments. As a result, I decided that the experiment should involve the movement and placement of objects, as I could easily create a VR experience with a one-to-one mapping of the movements as performed in the un-augmented world.

In order to provide some motivation for participants to move objects around, I came up with the idea of specified “work zones” into which the participants could place the objects. Each zone is associated with a participant and a progress bar that represents the amount of work needed to accomplish a task by the same participant. As long as any of the objects intersects a zone, that counts as doing work, and the associated progress bar will fill up to indicate this. Figuring out how to display the progress bar in the regular world and in VR was a massive roadblock, as I initially couldn’t figure out a good way to display it in both representations without massive differences; I decided that such a display would work well enough as long as I had a computer screen to display the progress bars in the regular world, and to have a floating screen in the VR version of the game.

I initially wanted to have a collection of objects that participants would place into the work zone into a specific order to accomplish a large assortment of tasks, the idea being that the activity would become some kind of competitive game where both participants would need to figure out the order to place the objects into their respective zones under a time limit. This initial proposal can be seen in Figure 1. Unfortunately, this proved to be rather unwieldy. I talked with several colleagues about this plan, all of whom required multiple clarifications before claiming to understand the premise. Upon attempting to simulate such experiments with these colleagues, it quickly became apparent that they did not understand the experiment as proposed, so I decided to try and simplify the experiment.

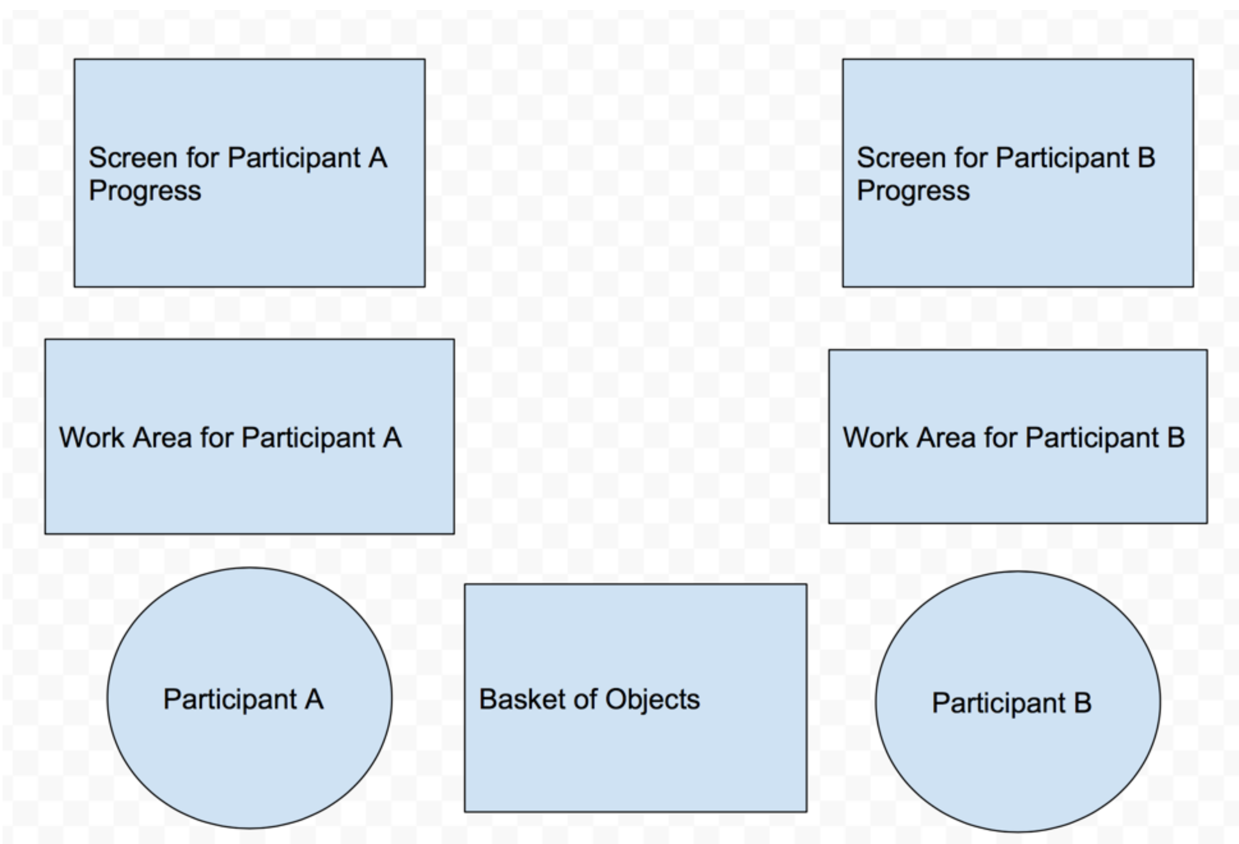

Figure 1. The first attempt at designing the tasks

My second attempt at designing the experiment still involved the ideas of objects colliding with work zones to accomplish tasks, but this time I reduced the number of objects to three and decided that the solution to this game would be to have the third object act as a “bridge.” This meant that in order for both participants to succeed in completing their task, they would need to place the third object such that it would intersect both work zones, where both work zones were located close enough in proximity. Then, if the participants could place the remaining objects into their respective work zones, the game would be considered a success, as the participants have decided to cooperate in order to complete their respective tasks in time, as multiple objects in a work area increase the rate at which work is being done.

This idea seemed to work rather well, as my colleagues were able to understand it very easily. However, as I tried simulating the scenario repeatedly, it became clear that there was no need for a third object, as it could easily be carried out with two objects instead. In this scenario, the participants would need to figure out that placing both objects such that they both intersect the two work zones would be enough to increase the rate of work such that both participants’ tasks would be complete. With this, the general idea for the game was complete, and I was ready to start the concrete implementation of the game.

One of the biggest problems in trying to conduct an experiment like this today is the lack of easy to access tools for hosting VR experiments in the same room, where participants can see each other in the VR space. As mentioned before, most VR equipment is intended to be used by one person at a time in one room, which prevents local multi-user experiences. Luckily, the lab I conduct most of my research in had the perfect setup for this. In the lab, we make use of twelve motion capture cameras to capture the movement of objects in the room. Taking advantage of this, we attached tracking markers to portable mobile phone VR headsets to get the positions of people wearing the headsets. This information is then transmitted over a wireless local area network to a server, where the locations of the people are sent wirelessly to the applications running on the mobile VR headsets. As a result, we can display models to the people wearing the headsets that match the location of where other people are in the same space. This was exactly what I needed for the purposes of my experiment, as now I had the ability for participants in the experiment to see each other in the VR version of the game. Due to the versatile nature of the cameras, I could create the objects necessary for the game simply by attaching markers to physical objects, such as boxes, and using that object's coordinates, can associate those objects with 3D models in the VR game.

Unfortunately, the camera system was ultimately not incorporated into my work, as the lab made significant changes to the software used for creating the VR experiences. The name of this framework is Holojam, a software development kit built on top of the game engine Unity. It enables people to create shared room experiences in VR, as long as they have the right equipment. Originally, Holojam used the data sent by the cameras in order to accurately reflect the positions of the participants in their respective headsets. The maintainers of the software decided to shift most projects in the lab to the HTC Vive system instead, a high fidelity VR headset that has its own methods for tracking position. I was initially disappointed by this change, but as I dove into using the Vive headset for development, I quickly realized how easy it would be to integrate the Vive game controllers as tracked, physical objects in my experiment since they were tracked by default due to the nature of the Vive system and how it relies on its controllers for any interesting input in its VR experiences.

As a result of this change, I now had to work with two Windows computers instead of two mobile phones, as the Vive only supports Windows computers at the time of writing this report. The computers I used for the experiment were designed to be worn as backpacks in order to have a wireless experience, as the Vive headset relies on external sources of light to calculate the user's current position rather than rely on camera data. However, for the purposes of the study, I ultimately treated the backpacks as computers to attach the headsets to, as the backpacks are not well suited for development purposes. Since they do not come with built-in monitors or input methods, that meant that that if I had participants wearing the backpacks, it would be a major inconvenience should I need to restart the experiment, as I would have to attach some kind of input method to restart the whole experiment.

Figure 2. A reference picture of the backpack in question

Setting up the experiment on the software side took some time, as I had to relearn how the Holojam software framework was setup, since the previous versions of the framework only required calibration of the cameras. With this version, I had to constantly configure the backpacks, as the Vive system needs to have accurate information about where the location of the floor is relative to itself along with the bounds of the current room. This proved to be a consistent struggle throughout the duration of my experiments, and I will go more into detail on this in the results section of the paper. Aside from the calibration issues, porting my game's design to the Holojam framework required a bit of time because of the necessity of synchronicity between the participants of the game. In a more traditional, isolating VR game, developers can implement features at a faster pace because there is no need to account for different users and how they might view different objects in the same space. With the multiplayer nature of Holojam, this becomes a bigger problem, especially with the rendered virtual elements. Since the positions of the headsets and controllers are coming from an external source, there is no potential danger of the lack of synchronicity there among the headsets, as all of that data is being taken care of by the networking stack in the Holojam code. That leaves the potential issues of synchronizing collisions of the controllers with the virtual work zones and displaying accurate information on the holographic screens to both users in the experiment.

The setup I used to solve these problems in the end made use of a third computer as a "master client," meaning that this computer also runs an instance of the game, specifically for synchronizing data between the users on the network. This meant that when creating and deploying builds of the game, that same computer does not serve as a traditional VR machine, as it only needs to be connected to the same local network as the other participants' in order to synchronize data. This pattern was already supported in the Holojam framework, but I needed to adapt specific parts as needed to suit my needs for this experiment. To make things easier, all physics code, such as the collisions between objects and zones, is run on the master client, and all physics code on the participant's machines is turned off to prevent any overlap in physics and potential duplication of actions, such as when work is being done in a work zone. In terms of the digital entities involved, I designated modified cylinders as the work zones, and added colliders on them so that whenever a Vive controller collided with a cylinder, I could keep a reference to it and process that as “doing work” as long as a controller continually intersects the cylinder. I also attached a capsule collider to the controllers, as the in-game meshes for the controllers would have been computationally expensive to calculate collisions for.

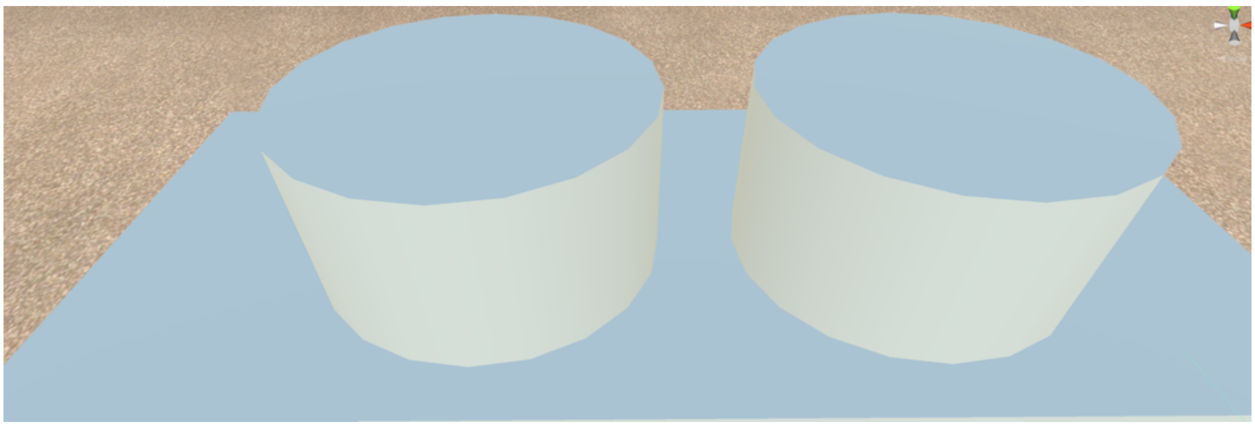

Figure 3. The cylinders acting as work zones

The key here is that the cylinder may not be visible in the regular world, but since it exists in the virtual space, as long as I have a physical object in its corresponding location in the physical space, I can have non-VR participants place their controllers on top of that physical object, and if they have a view of the progress bars, they will understand that this action has significance. To achieve this, I took a table in the lab and wore the VR headset to see the approximate location of where the cylinders would be in the space before placing the table there. I placed the cylinders high enough in the virtual environment and added a rectangular object beneath them to represent the table in the regular world. This positioning guaranteed that they would be located on top of the table, and I demarcated the division between the two work zones on the table in the regular world with a cord. This provided the setup necessary for replicating the experiment without VR headsets.

In addition to all of that work to ensure that the collision code worked well, I needed to ensure that data flowed properly from the VR clients to the master clients. This meant diving into the network stack in order to take advantage of Synchronizables, a component of Holojam that allows for developers to transmit various primitives and basic data structures over the local network. Since the locations of the controllers and the headsets were already being transmitted across the network, I needed to make sure that other variables such as the current time remaining for the experiment as well as the work being performed by each participant would be transmitted to the clients and rendered properly on each of their displays. This meant that in addition to performing the physics on the master client, I also updated the data on the master client with regards to work done, how long the controllers had been in collision with the work spots, and if the controllers were in position for the optimal solution. To make sure that participants did not accidentally stumble upon the solution by flailing around with the controllers, I made sure to also submit a timestamp of when both controllers collided with both work zones, and continually compared the current time with that timestamp. If the difference exceeded three seconds, then I could declare the experiment complete, as that duration is long enough for me to be convinced that the two participants made a conscious decision to collaborate and try out that configuration. All of this data is funneled from the master client to the server running on the same computer, which is then transmitted to the other clients in the scenario. Such clients then use this data to update their current displays, ensuring that participants are synchronized in terms of what they see.

The last remaining technical piece for my experiment was the display to inform participants of their progress as the experiment continued. I combined 3D text meshes and planes in order to give the illusion of there being a fully functional user display. This display, I placed in front of the cylinders such that when participants in headsets stand in front of the table, they will see the display floating in front of them as they perform the experiment. However, this would only serve well for the VR version of the game, as I did not possess the technology to display holograms in mid-air in the regular world. Fortunately, since I already had a third computer acting as a master client, I could utilize it as a display into the virtual world. Figure 5 shows what participants in the regular world experiment saw on a screen as they performed the experiment. It is a simple display with a progress bar and a timer counting down how many seconds the participants have before the experiment terminates.

Figure 4. An example of what the participants in the non-VR version of the experiment would see.

This representation of the work done for each task is, unknowingly to those who participated, misleading. Officially, each participant is to complete their task by making the word “Complete” appear on their side of the screen. However, the word “Complete” only appears when the participants achieve the cooperative goal of placing both controllers such that they intersect the cylinders. I purposely tuned the parameters of the experiment such that there was no way for participants to fill their progress bars in the time limit indicated if they choose to only use their own controllers. I did this hoping that, upon realizing that there was no way to completely fill the progress bar, the participants would search for other methods of filling the bar and ultimately come up with the out-of-the box solution necessary for the experiment to succeed.

Finally, I required a virtual environment for participants in the VR setting to perform the experiment in. Given the research regarding the Kuleshov effect and how people can choose to perceive each other based on their appearances and their environment, I decided to make use of digital deer masks and a model of the lab for this experiment. Figure 5 illustrates the results of this:

Figure 5. The deer mask and the lab environment

As there were potentially many different factors in the real world and virtual reality experiments, I decided that I wanted to, at the very least, keep the environments somewhat consistent. The model of the lab provided a rather accurate depiction of the actual lab, and many participants after the experiment commented on how it was nearly indistinguishable from the appearance of the lab. The deer masks were an attempt to induce people to cooperate better, as they provide an accurate measure of where each participant is and serve as a potential conduit onto which each participant can project emotions. By having floating deer heads in the space, I had hoped that the participants could potentially sympathize with these digital avatars, especially since by looking at each other this way, they could also realize that they themselves were being rendered the same way.

Figure 6. An illustrative example of the experiment in progress

Results

To recruit participants for the experiments, I sent out announcements to students in a class I was acting as a teaching assistant for and promised each of them extra credit upon successful completion of the experiment, with the permission of the professor, who was my thesis advisor. I also contacted personal acquaintances to ask for their participation. Luckily, the mere concept of VR was enough to entice people to come to the lab to perform the experiment. In total, I had 16 pairs participate in the experiments, 8 in the regular world space and 8 in the virtual reality space. Before beginning the experiment, I told each participant the rules of the experiment:

1. You are expected to complete the task that appears in front of you. Any work done on the task is represented by the progress bar in front of you.
2. To accomplish the task, you need to place any game controller such that it intersects the zone in front of you. As long as there is an intersection, work will be done. I have given each of you a game controller for this purpose.
3. Your task is complete when the word "Complete" appears on your side of the screen or when time runs out.

These instructions were intentionally left unclear in order to see how participants would behave while undergoing the experiment. Some participants asked if this was a collaborative or cooperative task, to which I explained that it was up to them to decide. For each experiment, I set the timer to 5 minutes and set up the devices depending on the type of experiment. For the VR experiments, this meant readying the VR backpacks and the master client for the participants before beginning, and for the regular world experiments, I only needed the master client and one backpack to track controllers.

To break down the results:

1. Of the eight regular world experiments, only one pair succeeded in figuring out the “out-of-the-box” collaborative solution.
2. Out of the eight virtual reality experiments, two pairs were able to figure out the solution.
3. For all of the experiments, about three quarters of the pairs were willing to cooperate, as at least one person in each pair willing gave up their controller to the other in the hopes of getting at least one task done, regardless of whose it was.
4. The remaining quarter (I’ll refer to this group as the passive participants) chose to not interact with each other during the course of the experiment. Instead, a good amount of these participants thought that moving the controllers in some way would potentially cause the score to increase. Several of these participants did acknowledge the deer masks in the VR trials, but did not go further than talking with their partner about unrelated matters as they continued to experiment on their own.

In summary, it did appear that there were more people who were able to figure out the non-intuitive, collaborative solution to the problem when participating in the VR trials as opposed to the regular world experiments. However, this is not conclusive given the sample size of participants as well as other potential factors that I noticed while administering these experiments. I found that a stronger predictor than anything regarding potential success was how willing and open participants were to collaborate before beginning the experiments. In all instances of successful pairs, there existed at least one participant who was open to trying out different combinations of controllers in different zones.

This was not sufficient enough to reach the expected solution, as around three quarters of the tested pool had some subject who tried to experiment with the given tools. What was also needed was some level of curiosity and willingness to collaborate from the other participant; in the unsuccessful trials, there was often one person who took the initiative as a leader to try different approaches for increasing score, while the other person remained passive, often yielding to the leader's suggestions. The successful trials had engaged participation from both parties in the pair. This result echoes the findings of integrative negotiation as mentioned earlier, wherein participants in a potentially competitive scenario have a better chance of succeeding if both have some degree of active collaboration.

This raises questions about why the passive group did not choose to collaborate. When asked about the lack of collaboration, the passive participants claimed that they did not think of that option at all. One participant in particular was especially curious about why he himself did not think of this option at all. Other pairs in the passive participants immediately assumed that this was a competitive game, but did not actively try to coerce each other into giving up the game controllers. Part of this stems from the fact that since they were only open to experimenting within their own zones and with their own controllers, they never made the realization that more controllers would help increase the rate at which work was performed in a zone.

Admittedly, there are some key areas that can be improved upon in this experiment for rigorous analysis. For instance, I could have had a third type of experiment where participants could put on physical deer masks to mimic the virtual deer masks seen in the VR trials to see if that was sufficient enough to mimic any potential effects of VR had on participants. I also could have put more effort into diversifying the pool of participants, as the people who participated were mostly split between white and Asian populations, with one black participant as the sole exception. There could have also been more attempts at trying to make the experiment as unbiased as possible in the VR experiments; I had attempted to make sure that participants would not see each other before attempting the VR trials in the hopes that those participants would only work off of the sound of each person's voice and the appearance of the deer mask, but due to the fact that I had to share the backpacks with many other people and the various technical difficulties I encountered, this was often the case not possible.

In fact, technical difficulties were a major obstacle in completing these experiments. For instance, there were at least two occasions where I had to work together with other members of the lab to fix issues where the Vive headsets were not seeing each other in the correct positions. This was often the result of the external light sources not being calibrated properly since other members of the lab needed to use a different mode of those devices for their own projects. In at least one case, I had to turn away potential participants due to the hardware issues I was facing. An idealized version of this experiment would have a space solely dedicated to these trials, where there is no interference from outside parties. We can take this further by administering formalized pre-test and post-test questions on a Likert scale to potentially classify test subjects' general attitudes towards cooperation before pairing together participants based on how cooperative they are projected to be. Given that this experiment suggests that having two participants who are generally cooperative will result in success, it may be best to verify this finding and to see if we can tweak the parameters of the experiment to induce other personality types to be cooperative.

One major flaw may lie in the experiment itself. I was successful in explaining and demonstrating the exercise to my colleagues in terms of getting them to understand how to perform the test and also how to reach the non-obvious solution of placing the controllers such that they both intersect the work zones. The participants, upon seeing the proposed solution, all understood the simplicity of the solution, with several displaying embarrassment that they did not reach this conclusion. The problem lies in potential that there may be too much of an element of puzzle solving and creative thinking in an experiment that was meant to test cooperation. Though my original hypothesis was to encourage original thinking by having participants cooperate with each other in order to eliminate any fixed-pie perceptions that the participants may have upon starting the experiment, the exercise itself may need to be play tested further in order to further determine that it is of sufficient difficulty to be solved, yet not too difficult to the point where only a small portion of the subjects can reach the conclusion.

Regardless, I am convinced that the results of these experiments demonstrate great potential for further research in this field. I would not consider my analysis statistically rigorous enough to make finalized predictions on the statistical significance of my results, but I believe that the trends observed in these experiments are encouraging enough to lead to a class of experiments that are not prevalent yet, both due to the novelty of combining the fields of technology and behavioral psychology and due to the rarity of the technical setup afforded to me by the research lab.

Conclusions

Unfortunately, the results of my experiments are not enough to confirm nor deny my hypothesis. Instead, there is ample evidence that there exists an undiscovered wealth of information regarding human behaviors in virtual environments which can be uncovered through further research in the field. In particular, the data I collected from my experiments strengthens the arguments for inducing integrative mindsets in competitive situations and that individuals often interpret such situations based on their prior conceptions about competition and collaboration. There is not enough evidence to conclusively show if blocking out the appearance of participants from each can effectively induce them to cooperate with each other; though the data suggests that it is possible since we have more people succeeding in the VR trials, this could be explained by the sample size and attitudes of those involved. Further research needs to be conducted to see if this hypothesis holds promise.

In particular, I believe that we can design even more effective tests that can be held in virtual reality and the regular world. Researchers in the field of psychology are well versed in designing experiments like these, and my work in this paper has proven that we can replicate such experiments well enough in the virtual world. Participants in the real world experiments and in the VR trials were encouraged to examine the version that they did not participate in, and they were convinced that the tasks that they would have been asked to perform would be one-to-one analogues for what they did. In other words, I believe that the Holojam framework can be utilized to great effect by psychologists when designing experiments to the point where it may not require the skills of computer scientists like myself in the design of the experiment. Ideally, we could place psychologists and computer scientists/programmers on the same team to help realize proper implementations of these experiments.

The world of VR holds much promise in terms of potential future social interactions between different people, and yet the current industry is solely focused on projects in the arts and entertainment sectors. Even with companies like Facebook trying to bring in new communication platforms, they seem to be shortsighted when it comes to imagining a world where VR can be integrated into people's everyday social lives without needing to be physically isolating. Though I did not prove or disprove my hypothesis, I am satisfied that I have made interesting discoveries with regards to the potential of VR as a social medium, and hope to have the opportunity to hold more formal experiments and to inspire others to conduct research in this area.

I will conclude with a word of warning. If we have conclusive evidence that virtual reality interfaces can be designed to influence participants' thinking, the designers of such interfaces must be held accountable for their actions if they choose to psychologically manipulate the thoughts of those wearing VR headsets. Much like how psychological experiment designers have to undergo a brief exercise in ethics, such interface designers should also be educated on these aspects, as even with the best intentions, these experiments can potentially go wrong. There needs to be accountability in this area of research, and I believe that going forward, we in the technology field need to make this aspect of VR more prominent.

References

[1] William Poundstone. 1992. *Prisoner's Dilemma: John Von Neumann, Game Theory and the Puzzle of the Bomb* (1st ed.). Doubleday, New York, NY, USA.

[2] Carsten K.W. De Dreu, Sander L. Koole, and Wolfgang Steinel. 2000. Unfixing the fixed pie: A motivated information-processing approach to integrative negotiation. *Journal of Personality and Social Psychology* 79, 6 (2000), 975–987. DOI:http://dx.doi.org/10.1037//0022-3514.79.6.975

[3] Ian Ayres. 1991. Fair Driving: Gender and Race Discrimination in Retail Car Negotiations. *Harvard Law Review* 104, 4 (1991), 817. DOI:http://dx.doi.org/10.2307/1341506

[4] Charles F. Bond and Bella M. Depaulo. 2006. Accuracy of Deception Judgments. *Personality and Social Psychology Review* 10, 3 (2006), 214–234. DOI:http://dx.doi.org/10.1207/s15327957pspr1003_2

[5] Daniel Barratt, Anna Cabak Rédei, Åse Innes-Ker, and Joost Van De Weijer. 2016. Does the Kuleshov Effect Really Exist? Revisiting a Classic Film Experiment on Facial Expressions and Emotional Contexts. *Perception* 45, 8 (2016), 847–874. DOI:http://dx.doi.org/10.1177/0301006616638595

[6] Tabitha C. Peck, Sofia Seinfeld, Salvatore M. Aglioti, and Mel Slater. 2013. Putting yourself in the skin of a black avatar reduces implicit racial bias. *Consciousness and Cognition* 22, 3 (2013), 779–787. DOI:http://dx.doi.org/10.1016/j.concog.2013.04.016