Title
Supporting Social Activities in VR

Author
Jhaveri, Samvid

Publication Date
2018

UNIVERSITY OF CALIFORNIA

SANTA CRUZ

Supporting Social Activities in VR

A thesis paper submitted in partial satisfaction of

the requirements for the degree of

MASTER OF SCIENCE

in

COMPUTATIONAL MEDIA

by

Samvid Jhaveri

December 2018

The Thesis of Samvid Jhaveri is approved:

Professor Katherine Isbister, Chair

Professor Angus Forbes

Lori Kletzer
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

Abstract

Social VR currently includes a diverse set of applications and competing models of what it means to be social in VR. This work tries to grasp the current situation of what industry is focusing upon and then demonstrates a few interventions in the social VR meeting spaces and analysing the usability with other developers and understanding the future research in Social VR.

Acknowledgment

This project was a collaborative effort and team members will be credited appropriately. My project advisors, Professor Katherine Isbister, Professor Angus Forbes always offered me unconditional support, inspiration, and expertise from varied areas of study. A very special thanks to Joshua-McVeigh Schultz for helping out and guiding on each and every step of development. A very special thanks to Max Kreminski on his collaboration on the Virtual Pet section. And finally, thank you to Anya Koleshnichenko and Erica Nguyen without them the visuals could have been super different and we never would have found the theme of Ba-Ke-Neko.

Introduction

There is currently a great deal of interest in social VR environments, with many commercial examples, such as Facebook Spaces, VR Chat, and RecRoom. My masters research examines the terrain of design for supporting social activities in VR. I have created a taxonomy of current Social VR activities, examining some current and popular applications. Drawing upon this taxonomy, I have created three prototypes of social activities in VR which together form my project which I titled ‘Ba-Ke-Neko’: Virtual Pet, Gaze Exploration, and Social Trace. Each activity has a unique set of interactions aimed at supporting social activity. Finally, I have conducted a preliminary user study of these prototypes, toward drawing some conclusions that might guide other developers interested in supporting social activity in VR.

My thesis draws upon my experience and learning in the field of Computational Media. This work made me use different skills--sometimes more than one--defined in the final report on Media System<need citation>. This work takes advantage of such as technical and creative (Designing & developing different activities) and, also collaborative (Worked with a few people along the way of development).

There have been previous works on Collaborative Virtual Environments(CVEs) and social activities in the virtual environments<need cites>. My research takes these projects and their findings into consideration, and also draws upon research examining gaze, social proxemics, and social navigation, using all of these influences toward creating social activities in VR.

This report is structure in five different sections. First, I describe my explorations of different current VR meeting spaces and the activities in those meeting spaces, and my reasoning of classification of those activities. Taking inspiration from both past work and this taxonomy, I then present the three different VR activities that I created in Ba-Ke-Neko. First is a virtual pet, where the goal of the prototype was to create an alien species with which users can interact. Second is Gaze Exploration: this prototype is inspired by work in social proxemics, and uses that knowledge to create an activity that makes use of a person's gaze. The last activity is Social Traces, built upon the knowledge of social navigation. Users can see all of their activity as well as interact with it. The final part of this report is the user study. I ran a user study with VR designers & developers, to get their feedback on the prototypes that I built.

A Taxonomy of Social VR Meeting Spaces

Introduction:

Before setting up Ba-Ke-Neko, I wanted to scope out existing applications and use cases in Social VR meeting spaces. As mentioned in the earlier section, the third generation of VR is on the rise, providing features at a lower price point that were previously only available to research institutions. Engaging in social media interactions within VR is something that has been a part of the concept for this generation of VR headsets, so creating a taxonomy of existing social applications and use cases can be beneficial to understanding the state of Social VR, as well as setting the stage for my project, Ba-Ke-Neko.

Related Work:

In this section, I cover other researchers' attempts to taxonomize virtual reality and related domains. There have been many attempts to survey to classify and subdivide the hardware side of Virtual Reality. In [1]the authors present different types of VR and CAVE environments, and generated a taxonomy of VR hardware. In [2], the authors present the evolution of the definitions of virtual reality, based on technology and immersion. They also outline a brief history of virtual reality technology. Furthermore, the authors discuss key image processing techniques needed in designing virtual reality systems and virtual environments.

Another related work was recently presented in (Liu et al., 2012)[3], in which the authors reviewed the current state of melding-related techniques in virtual worlds.

Melding is the core of creating the illusion of a shared reality. The authors introduced a taxonomy of consistency models that helps in providing users interacting within a shared virtual world with the illusion needed to improve their virtual experience. The taxonomy was applied to several shared virtual worlds in the paper, for example. Finally, the authors discussed the challenges and promising solutions of state melding in large-scale virtual worlds.

In another work, Zhao[4] has created a taxonomy of copresence. He also defined the meanings and subtypes of copresence and how this copresence refers to the physical conditions that structure the human interaction. Zhao defines copresence as consisting of two dimensions: copresence as *mode* of being with others, and copresence as *sense* of being with others. Mode of copresence refers to the physical conditions that structure human interaction. Six such conditions are classified by Corporeal copresence, Corporeal telepresence, Virtual copresence, Virtual telepresence, Hypervirtual copresence and Hypervirtual telepresence.

One of the very related work can be found in the classification of the Ubiquitous computing. As explained in [5], this classification covers some of the meeting spaces, and the current work is also inspired a little bit from the work explained in this paper as my work explores more in depth in meeting spaces and it explores the area of having the smart capabilities in the interactive environments and gives the detailed examples about that as well.

There also some taxonomy of Virtual Reality Environments in the software side as well. This work [6]classifies different types of network structures implemented in networked virtual environments. The authors classify network communication considering latency, distribution, schemes and reliability.

Objectives and Scope of the Current Taxonomy:

The main objective we had in creating my taxonomy was to provide a comprehensive review of the current status of different social VR meeting environments. We believe that experts from Virtual Reality fields would benefit from this review. Also it could help new designers as well, because this will help them classify their work, as well as find the areas where they can develop more.

Another objective we had was to use this work as the base structure in order to classify the work that has been done in Ba-Ke-Neko. This work is to show what inspired me and the team I worked with to create different social superpowers and different activities in Ba-Ke-Neko.

Thirdly, the accompanying visualization of the taxonomy was intended to give deep insight to non-designer users as well, who could then see and understand more easily by looking at this visualization. Users can also explore this visualization and find applications they would like to know more about, and some of the general activities they can do in those meeting spaces.

Finally, the companies that design this type of Social VR experience can use this research to better understand which domain they want to design for and how they can use their existing work from one domain and modify to use in other domains, by going through the examples provided in this taxonomy.

Variables of the Taxonomy:

Before classifying existing applications and use cases, we decided upon some basic variables for creating the taxonomy. As a starting point for our work, we had the following questions in mind:

- Why do people choose to use a particular Social VR application?
- What do people do inside a given Social VR application?
- What are the choices designers made within each application to increase and enhance social interaction between people?

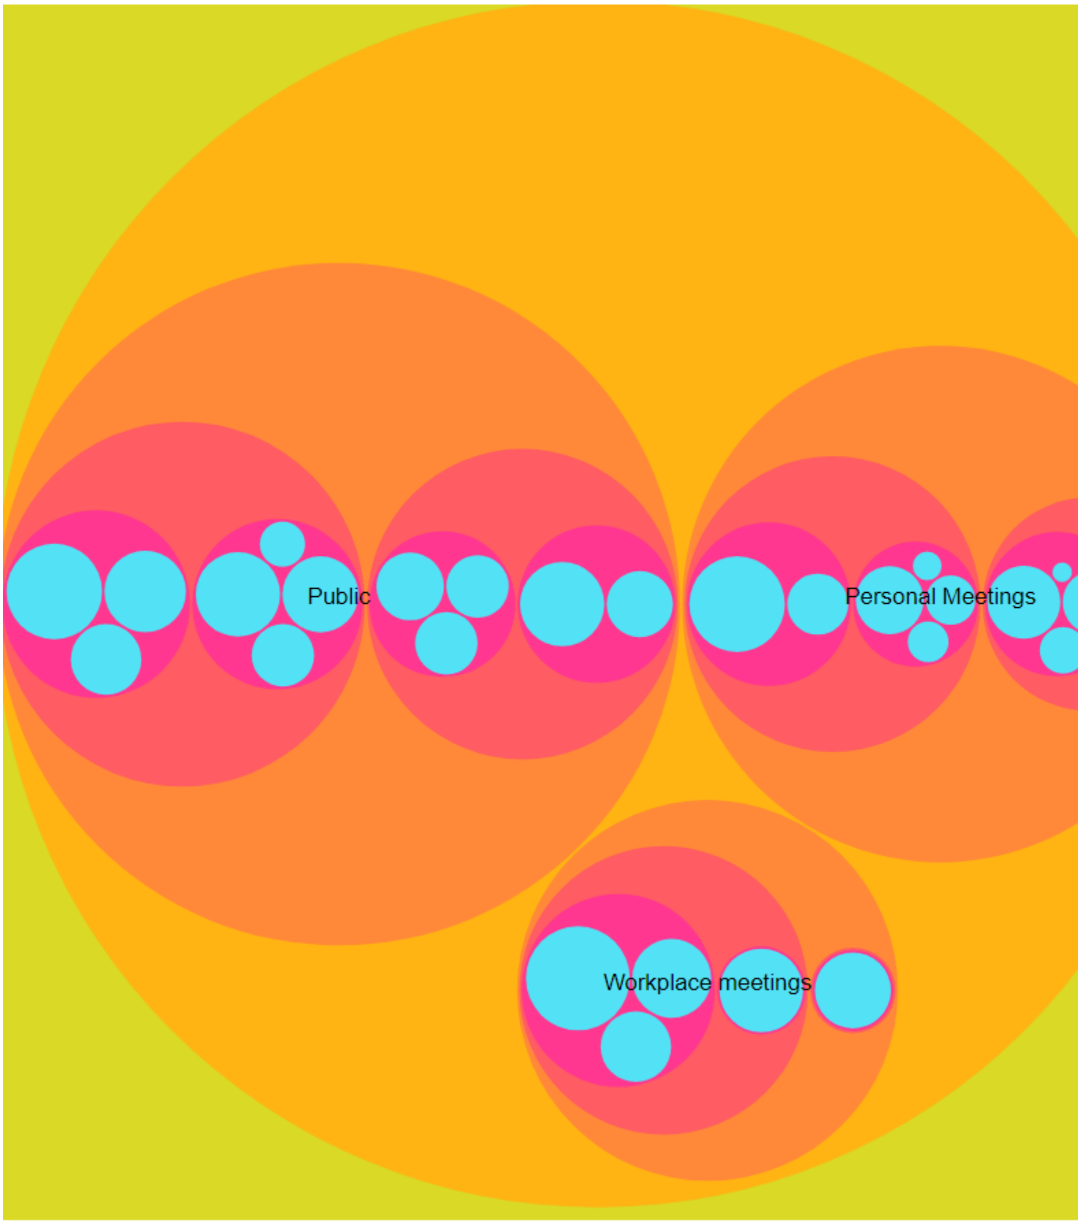

Fig.1: Taxonomy of meeting spaces

We spent time in several social VR applications, talking to people and spending time interacting, to get more clarity about these questions. For our work we have chosen applications such as Facebook Spaces, Rec Room, High Fidelity, VRChat, AltspceVR, and Cisco Spark VR. These are some of the most popular meeting spaces.[7]–[9] Some applications are marketed specifically for a single domain.

For example, Cisco Spark VR is marketed only as the extension of Cisco Spark Chat, which is mostly used for workplace meeting spaces in virtual reality. As you can see from the taxonomy visualization(Fig1), this space is divided into three main domains from the use cases we observed:

1. Public spaces
2. Workplace meeting spaces
3. Personal meeting spaces (with close friends and family).

But when we looked at the platforms using the designer perspective, we found we could divide activities into two main functional categories:

1. Ambient Activities
2. Coordinated Focused Activities

Ambient activities require less focus on the activity from users. Ambient activities are the type of things that you usually do in a group where It matters less upon the specific contribution of a person and it more depends upon the group of people and their interactions with the environment and experience as a group. For example, watching a SpaceX launch together in VR where you don’t interact with the environment much but you experience it as a group. On the other hand, coordinated focused activities are targeted for a small group of people. These are the activities that requires more attention from the group, and they are combined activities between a small group. For example, playing a game like paint ball in VR which requires active participation from each and every member of the group.

Upon further inspection, we were able to understand more clearly that each domain has both types of activities, and applications are divided among these activities. There are some applications that cover multiple domains, while some are just focused upon a single domain that can be easily distinguished by the activities offered in the application. For example, VR Chat focuses upon meeting new people, but they have some activities such as watching your own uploaded photographs and video on a large screen, which people usually prefers to share with friends and family members only. Whereas Facebook Spaces focuses only upon personal meeting spaces.

There is one more feature that few applications provide, which falls in the middle of all these domains, which is creation of private spaces. These private spaces can have the same sorts of activities as other platforms, but the distinguishing feature is that they are controlled by the user rather than the company—e.g. the user can decide how their room shall be used and who is allowed to enter it. We were trying to understand how designers think about creating social VR spaces, so we are excluding the option of private spaces for now. In end of this taxonomy, I will mention the platforms which give this option of creating private spaces, and whether they provide any different functionalities.

1. Public Spaces:

   I define public spaces as places in social VR where you can meet new people. The applications in this domain mainly have open environments where users can meet and greet other people. This domain has the greatest number of activities, and there are a lot of applications in this domain as well. In order to help people to make good connections with other users, these applications let you join in small activities such as games, which are more focused upon team play in their environments. Now let’s discuss the activities in these environments:

- Ambient Activities:

An example of an ambient activity could be watching a concert or watching some sport together, or dancing in a club-style environment. These are activities that are more focused upon the experience of the whole group. As we can see from the example activities, they do not require a person’s undivided focus. These activities are designed to enhance meeting and talking with new people.

Most social VR platforms include some kind of ambient activities. Platforms like High Fidelity and Alt Space also host some weekly Open Mic for music and comedy, which users can attend, and they also support big events like the launch of the SpaceX, providing people with ways to view such events together in social VR. RecRoom has a giant gym room where you can meet and talk to new people and interact with the objects around the room as part of the activities which are completely optional.

- Coordinated Focused Activities:

In the public domain, often these activities are combined with ambient activities, so that a user can meet new people and then to get to know more about the other person as they jump into these activities. These activities often require more focus from the user than ambient activities.

Playing a game is one of the biggest example of this type of activity, and most of the Social VR meeting spaces include some game. Another example of this kind of activity is visiting other peoples' virtual worlds—something that can be done in AnyLand. AnyLand has a mechanic where you can create your own world and define some basic physics rules. You can visit your friend's virtual world as well, or make it open for the public to explore. Another example is making a cult and following it. A few months ago, Ugandan Knuckles were pretty popular as a joke in the application VRChat. People used a specific avatar and a set of dialogues to communicate with other users. In case of Ugandan Knucles popular avatar is in the figure below and one of the dialogues was, "Do you know da way?" and this became an internet meme for a couple of months.

2. Workplace Meeting Spaces:

The core purpose of these types of meeting spaces is to create and use the environment for work-related tasks. They are more formal and created for workplace practices such as stand-up meetings. The types of activities in this area are pretty skewed towards most of the activities being of the coordinated focused category. Example platforms are very few, Teleporter, Cisco Spark VR, Telia VR in this section, as this is a small and targeted market compared to broad scale consumer-facing social VR, and that's why this domain is the least explored of all 3 of these domains.

- Ambient Activities:

Ambient activities are usually created with a goal that they don't require much focus from the user, but in the office environments these activities seem to be rare, perhaps because companies see them as distracting.

The only example activity that I was able to place in this section was some background music before the meeting starts in Teleporter.

- Coordinated Focused Activities:

The user in these activities is required to focus upon the activity very heavily. For this domain of office meeting spaces, attention is focused solely upon the work, and these are activities created by the designer for the sole purpose of increasing the efficiency of the employees and providing more tools for meeting environments. Because of that, the example activities are spread across a very large area from having different type of meetings to small features within meeting spaces.

For example, Sprint meetings is one core activity in this domain. This activity is one of the model use cases for this domain—this is when users can do their Agile-style daily sprint meetings within the virtual environment. Another example could be, Slide Presentation or Volumetric Drawing. As stated above, this is a type of activity that can be a part of another activity such as Sprint meetings, but the reason for listing both of them separate, is that Slide Presentation can be used as its own entity and can be its own activity. Example environments focused on workplace activities are Cisco Spark VR and Telia VR. This domain is not yet well explored, as this market is still pretty small.

4. Personal meeting spaces (with close friends and family)

The main goal of this type of meeting place is to meet your friends and family members, the people whom you already know. The activities in this domain may overlap with those in public spaces, but there are some activities which are unique to this domain. There are very few applications for example Facebook Spaces that clearly focus upon this domain, although most of the applications have private spaces, which essentially convert activities in public spaces to personal ones.

- Ambient Activities:

The activities in this domain are designed to enhance a person's experience in VR with their friends and family members. An example of this type of activity is having a virtual creature, or watching a whale from inside VR. Such activities tend to need less focus from the user, and provide a backdrop for the interaction between people rather than taking up all attention. Other ambient activities include exploring different environments, sharing a meal, fishing together with someone. As I have mentioned before, there are very few applications which focus upon this aspect of VR. Facebook spaces is the only application that facilitates this behaviour.

These activities need little focus and they keep moving from background to foreground and after a while back to background. This is the starting inspiration for the Virtual pet in Ba-Ke-Neko.

- Coordinated Focused Activities:

These activities are small activities which we engage in with our friends just to have fun. Most of these activities are physical activities recreated in VR, for example Pictionary and Karaoke. Also other small coop games like charades or Pictionary are implemented in VR. They are not created with intense and complex mechanics and game design in mind, but are instead small, fun, low stakes experiences.

Example applications of this type of activity occur in Facebook Spaces and Rec Room. Facebook Spaces gives users cards to play Pictionary, and lets users draw in VR to recreate the experience of Pictionary.

Another type of Social VR experience that is not listed here is the asymmetric VR experience. These are activities that are performed between a person in VR and a person in the real world. There are very few example activities of these type of experience. One example would be a dancing game Astaire which is being developed in SETLab, in which one person is holding one controller in the VR world, with the other person standing outside with the other controller, who guides the person in the VR from the obstacles.

Virtual Pet as an Intimate Ambient Social Affordance

Introduction

After exploring the taxonomy of the VR social interactions, we focused upon finding an area where we wanted to design an experience.

This section of the research contributes:

- A prototype of a virtual pet that occupies a shared (networked) VR environment.
- The first step towards a framework of how to design AI agents that play the role of companions to humans in virtual spaces.

The goal of this part of the research was to create a virtual pet in a shared VR environment to encourage prosocial interaction, both person-to-person and person-to-pet. Our eventual long-term goal is for this agent to be added as one component of a larger existing social VR environment, where we can test this AI agent in different social settings. Of particular interest to us is the role that such an agent might play in mediating 'intimate' ambient activity, such as social interaction between friends and family members. As such, our focus is on the presence of an AI agent during intimate ambient interaction and how this will affect the emergent social dynamics in such a setting.

There have been many prior attempts [10], [11]to create virtual pets, but most of these attempts have been focused on developing features that more closely resemble biological pets (feeding, petting, caring for), and all of these applications have been created with 2d or 3d screen-based interfaces in mind[12], [13]. Our agent instead resides within a shared (or "multiplayer") virtual reality environment, where users have 6 degrees of freedom and can interact with the environment freely. Our project adds to this environment an AI agent with an anthropomorphic (or, perhaps more precisely, a "zoomorphic"[14]) interface, intended primarily to provoke conversation or social interaction between multiple human participants in the virtual space through its behaviors. In addition, our agent functions as a companion to the space's human inhabitants, with interactions such as receiving petting being specifically designed into the agent's behavior: the virtual creature is aware of these actions and has the ability to appropriately respond.

We also used this project as an opportunity to explore alternative ways of conceptualizing the relationship between humans and AI. In particular, because the project explicitly casts an AI agent in a role that would normally be taken up by a biological pet (such as a cat or dog), it benefits significantly from engaging with Donna Haraway's notion of the "companion species[15]" and how AI might play the role of a companion species to humanity going forward.

Related work

This project meets the definition of "intimate technology" given by Marc Bohlen and Michael Mateas in "Office Plant #1":[16] it is a technology which is "best situated" not in a laboratory, but in an office or home space where it is "close to people". This agent is a mediator between machine precision and human emotions. It attempts to understand human emotions by analysing the movement of the players as well as comparing human action with its current state, and behaving accordingly, as an anthropomorphic or "zoomorphic" agent. For example, if our virtual pet moves towards the person with an expectation of spending time with it, and if the pet senses extreme hand movements, then it will change its behaviour and do a different activity, because this interaction will trigger the state of scared in the pet and it will in turn trigger another state depending upon it's hunger and sleep values.

Like Office Plant #1, many of the cues it takes from the environment in determining how to act are implicit cues that are not deliberately intended as commands or suggestions to the agent. Unlike Office Plant #1, however, the agent (in being modelled more after an animal than a plant) provokes more direct interaction between itself and human participants in the space, meaning that some of the cues it takes from the environment are explicit as well. Rather than a passive observer and commentator along the lines of Office Plant #1, our agent is cast more as an active participant in the social space.

Similarly, as an "ambient, non-human, embodied, intelligent agent", our agent also arguably meets the definition of "alien presence" given by Adam Smith et al. in "Tableau Machine: A Creative Alien Presence"[17]. Like Tableau Machine, our agent makes use of an AI-based perception/interpretation/output loop to determine how it will behave on a moment-to-moment basis. The input in our case is the behavior of the humans and human actions in the shared virtual environment; these actions are then interpreted by the AI and used to determine what actions the embodied agent should take in response.

Unlike Tableau Machine, however, our agent is intended to be recognizably less alien and more familiar in its embodiment and behavior, while remaining largely non-anthropomorphic (at least in the strict sense of "attempting to appear human"). We intend to achieve this by leaning away from anthropomorphism and towards zoomorphism instead. By modelling the behavior of our agent on behaviors we expect many people to be familiar with through the observation of pet dogs and cats, we hope to create an agent that is recognizably inhuman but also less wholly alien than something like Tableau Machine (which is totally non-zoomorphic in form).

Finally, in deliberately attempting to provoke social interaction between multiple human participants in the space (in addition to direct one-on-one interaction between the human and the AI agent itself), our agent goes somewhat beyond both Office Plant #1 and Tableau Machine in overall design intent. The added challenge of sensing and responding to social dynamics in the environment did not seem to be a particular priority for either of these earlier projects, whereas here it plays a relatively important role.

Implementation

The pet's behavior is managed by a state machine. Transitions between states are regulated by numerical "needs" ranging from 0-100. At any given moment, the pet is in one of four different states: Wary, Hungry, Curious and Affectionate.

Initially, the pet starts out in the Wary state. In this state, it will attempt to keep its distance from the player, moving away from the player to preserve this distance if the player attempts to approach it. "Aggressive" actions from the player (such as moving too quickly) will increase the pet's wariness, causing it to remain in the Wary state for a longer period of time. If the player behaves non-aggressively, however, wariness will gradually decrease. At the same time, the pet's other needs (hunger and curiosity) are constantly on the rise. Once one of these other needs becomes more pressing than wariness, the pet will then move from the Wary state into either the Hungry or Curious state as appropriate.

If the pet becomes Hungry, it will move towards the food dish and start eating its food. While eating, to signal this behavior to the player, the pet will continuously move its head up and down. Eating will both decrease the pet's hunger need and make the pet more curious to explore the virtual environment.

If the pet becomes Curious, it will move towards the player and start jumping around them in order to catch their attention. This is the cue for the player to pat the pet. This is where the virtual reality shows its own idea where the way to pat the animal makes a lot of difference. For example, if you pat gently around its body, it will move towards the Affectionate state. On the other hand, if you behave rudely with the pet by putting your hand inside it, it will become angry and scared, moving back to the Wary state.

In the Affectionate state, the pet will approach the player and sit at the player's legs to show affection. For a little while, the pet will remain at rest at the player's feet. At the same time, however, the pet's own biological clock is also in progress, so after a little while it will become Hungry and potentially shift states on its own. In addition, while the pet is resting, any sudden movement from the player will scare the pet. For example, if the player moves both of their hands very rapidly, the pet will become scared of the player's behavior and return to the Wary state.

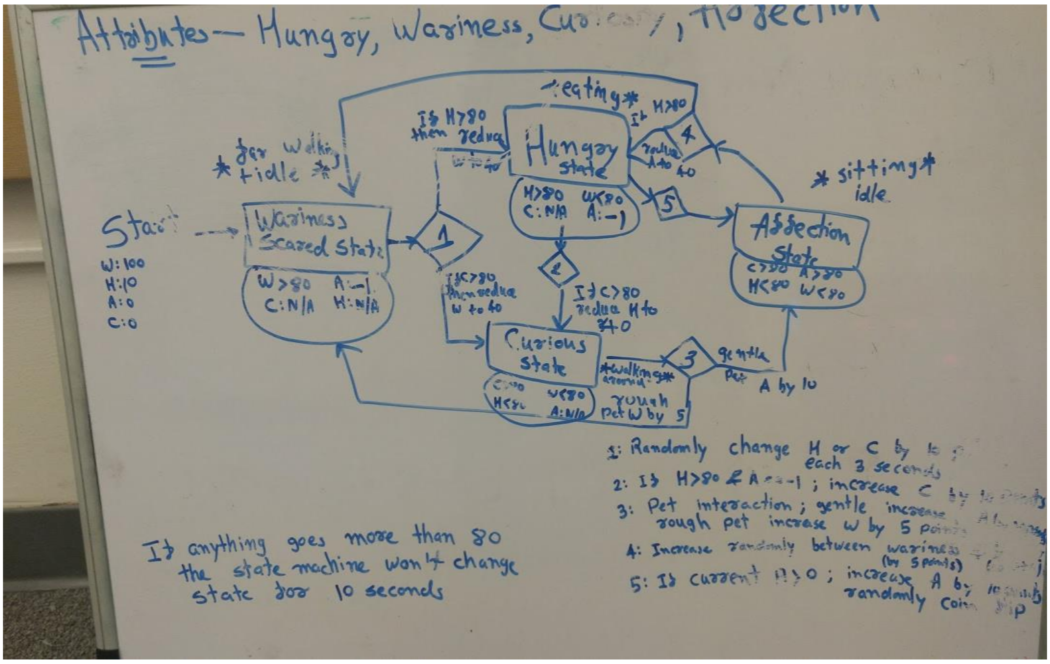

Fig.2: A diagram of the state machine governing the pet's behavior.

Design Considerations

The physical representation of the pet was created from primitive objects. We believe that, if we had used a higher-fidelity pet model (such as a more directly representational cat or dog model) that was more familiar to the player, it would encourage direct comparison between the virtual pet's behavior and that of the particular type of real-world animal being portrayed. If the virtual pet's behavior did not induce the correct emotions in the player (and thus "measure up" in this sort of direct comparison), this might in turn lead to a sort of uncanny valley effect[18], paradoxically making a higher-fidelity pet with imperfectly tuned behavior feel subjectively *more* alien rather than less. This assumption, however, remains untested.

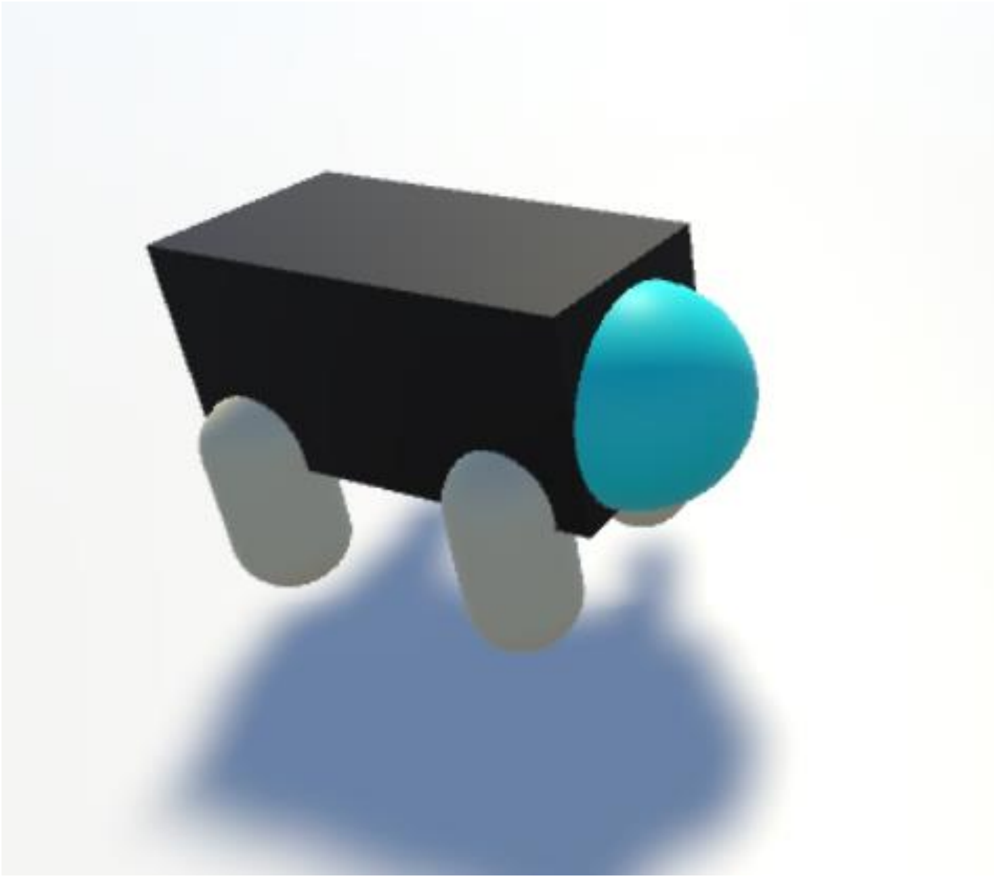

Fig.3: The virtual pet, created from primitive objects.

Another major obstacle we faced was the challenge of making the pet's internal state visible in a way that human players could understand. Creating detailed animations is not our forte, and would likely have taken far too much time for the scope of this project, so instead, we decided to use color and very basic animations to signal the pet's internal state to the player. In particular, the color of the pet's head is used to signal several of the pet's main "emotions" or emotional state categories. Yellow signals high curiosity, red signals high anger, and magenta signals high affection; in each of these cases, the shift in color is directly proportional to the current value of the pet's corresponding need.

One example of using basic animations to signal internal state can be seen in the pet's behavior while in the Curious state. In this state, the pet will approach the player and jump up and down in order to catch the player's attention. Scaling the pet is also used to signal internal state.

This idea was inspired by the book *The Illusion of Life: Disney Animation*, which presents twelve basic principles of animation; one of these principles specifically describes how scale may be used to convey emotion.

Future Work and Conclusions Concerning Virtual Pet

We focused our demo only upon a single human participant's interaction with the AI agent, demonstrating a few example behaviors of the virtual creature when it is placed in an environment with a human participant. Social interaction between multiple human participants in the VR environment is thus left for later, as this would add an additional degree of complexity to the challenge of demoing interaction with the AI agent without being necessary to show off the core of the project's functionality.

In our view, this project represents the first step towards the development of a framework for designing AI agents that are intended to serve in the role of companions to humans. Future work building on this project will also help determine the ways in which AI agents might impact social dynamics and social interaction in intimate ambient activities and social settings, as this AI agent is intended to encourage or afford specific kinds of social behavior among people interacting with it. In addition, introducing this type of AI agent into a virtual reality environment also enables us to explore the possibility space of deliberately anthropomorphized (or zoomorphized) objects in the ecosystem of virtual reality, drawing on the apparently natural human tendency to read anthropomorphic behavior into virtual characters.

Gaze Exploration as a Coordinating Social Affordance

Introduction

The research work described in this section uses gaze direction in social VR in a playful manner, in order to enhance communication between people in these environments. Our longer-term goal is to create a framework for designers to guide them in using gaze control to create more activities in Social VR environments. Our current focus is to use the gaze control mechanism to help people stay together and explore the environment while interacting. We are interested in how gaze control might enhance a sense of 'togetherness'.

There has been extensive prior work evaluating gaze direction and how it enhances face-to-face interaction in virtual environments, but all of this experiments have taken place using real world applications. Tracking gaze outside virtual environments requires special equipment, whereas it is easy in virtual reality to track gaze. This makes VR a convenient testbed for trying out some design ideas that rely on gaze tracking.

In the research prototype that we built, we used gaze as a way to explore alternatives for how mutual gaze can be used in social VR. We use gaze as a beacon. Knowing exactly where the other person is looking, and exploring that with the general place by using the position as a beacon, in order to show his location with lights.

Normal face-to-face communication is an extremely rich, multimodal form of expression. Aside from the verbal channels, non-verbal channels available during face-to face communication include gaze from head posture and eye direction, arm gestures, body posture and facial expressions. With the expansion of virtual reality platforms, hand tracking and head postures are visible to other users, but the gaze of the user is a very open subject to explore. There has been some work going in the use of gaze in collaborative virtual reality environment[19], [20].

In Ba-Ke-Neko we try to take this process one step further, using gaze information to heighten communication and interaction while exploring a dark environment. As noted in Benford's paper[21], one of the main challenges in Collaborative Virtual Reality environments is interest management. By arranging the virtual environments so that each participant is not overloaded and sees and hears "enough" of the world but no more, the problems of scale can be diminished. This was particularly useful to us while designing this activity. In the current activity users have to follow each other's gaze and work as a light source for each other in order to explore the hidden island within the social VR environment.

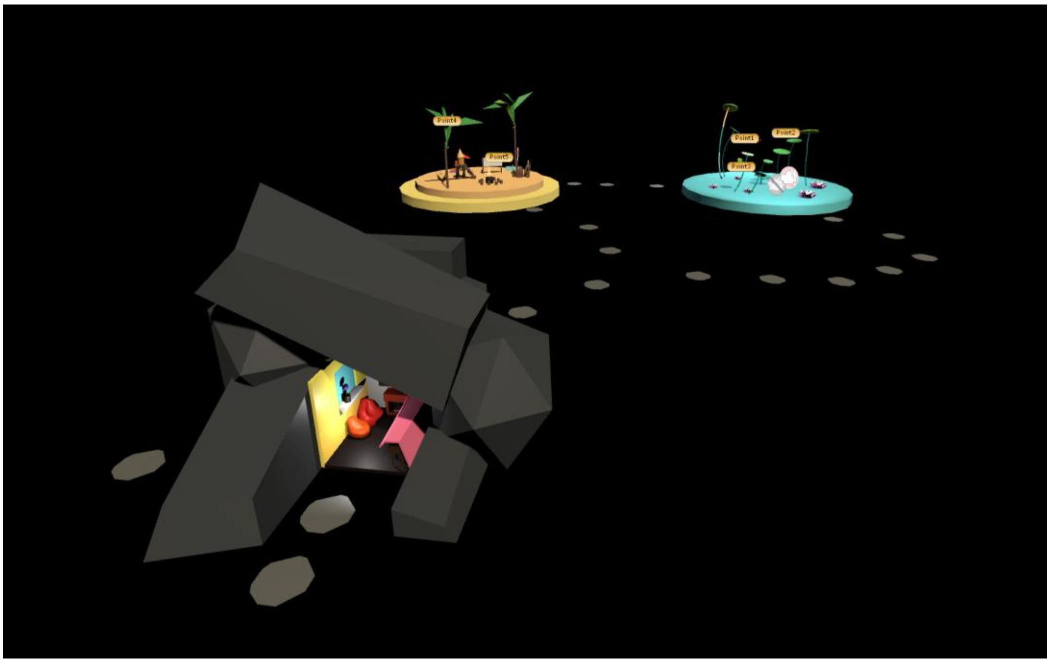

Fig.4: The hidden island

In this feature there is an island in the social VR world. This island is connected with the main island by small stepping stones. This whole island is invisible until it is turned on by the moderator or using the settings menu. After turning it on, in order to find this island, users have to work together. The stone bridge is only visible to you by staying with other users, as the player can use other users as their light source which will guide them to the next stepping stone, and this is same for all the users. Also you can only see one stone at a time, so you have to stay in close proximity with each other in order to find the other stepping stones. One more thing is, if a user's partner somehow falls down from the stepping stone, then they reappear at the origin point of the main island. Because of this, the other user has lost their guiding light, so it is impossible to spot the stepping stone. This forces the second user to jump down from the stepping stone and start the exploration all over again from the beginning. This forces both players to interact more with each other in order to find the island.

Related Work

The related work presented here explores different fields that are connected by VR; we have divided the overview into Proxemics, and Gaze and Attention. We do not cover gaze detection work outside VR [22].

Proxemics:

In an influential book[23], Hall categorized interpersonal space into four zones of intimate, personal, social, and public distances that people maintain towards each other in the real world. People adhere to similar norms in VR, for example keeping greater distances from humanoid avatars than cylinders [24]. Participants moved out of the way when approached by virtual avatars and also kept greater distances when there was mutual gaze. Other studies have also shown that real world proxemics behaviors work similarly in VR, for examples kin conductance response was shown to vary with the distance of one or many virtual characters as expected on the basis of proxemics theory[25].

Gaze and Attention:

Information about another individual's direction of attention is important in social interaction[22]. In this respect, gaze direction and head orientation provides the most accurate nonverbal information and one study shows that if though the gaze is averted it shows active attentional behaviour, and this shifts the other users attention in the same direction[22]. We have created two methods of using the user as a light source and by looking back to Peter et al. work we were able to understand the importance of having the information of a person's gaze direction and chose the most productive method. Another study analysed the head orientation's contribution in overall gaze direction and application potential in focus of attention estimation, which helped us calculate the estimated point of the person's gaze. [26]

Implementation

Exploration in this prototype is managed locally, by sending data concerning the other user's location and rotation. We then created a dark environment for users to find the hidden island. In order to find the island, users have to stay near each other as they can only navigate with the aid of the light coming from the other users. If they try to find it alone, it will be just a dark environment for them, and they won't able to find the path to the hidden island.

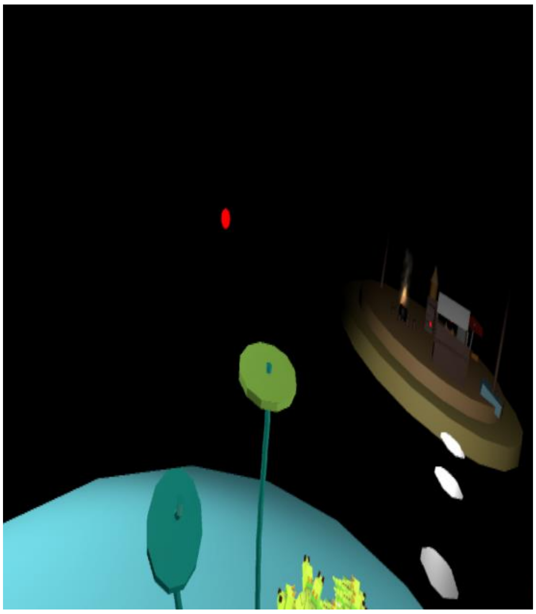

Fig.5: User's view to first island

Fig.6: User's view to hidden island

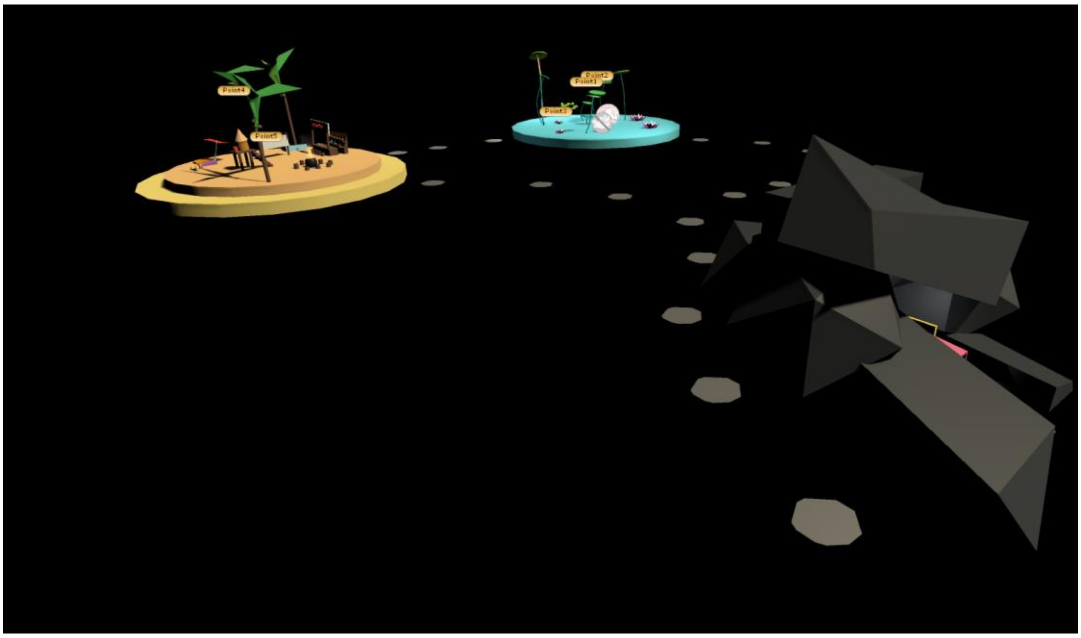

Fig.7: This is how the whole hidden island is mapped out

Before building this exploration environment, we needed to be able to detect gaze. In real world applications, ascertaining where a person is looking is a difficult task, and requires a lot of mathematics to exactly locate the gaze, but fortunately for VR environments that task is quite simple, as we can utilize gaze tracking that the headset approximates. In most VR headsets, in order to render the scene with a sufficiently high frame rate, engineers use a very neat optimization technique where they only render high quality images where the user is currently looking, and they slightly blur out the other areas. That is also a reason why VR headsets tend to have motion blur problems, and that’s why they apply local calculations in order to correct the motion blur. But for our work, we decided to utilize the approximated gaze from the headset calculations, which is in the middle of both the lenses, and by triangulation we can find the center spot that is in the focus of the headset. Achieving this gaze direction was the first step of creating this activity.

The next step was to communicate this information to all the other players connected in the environment. This was the tricky part, as it required some optimization around what data to be sent via the network. We use simple positional and rotational data from player position and send that to all the other clients.

Fig.8: End goal for user

The next step was to implement what we wanted users to find, and for that we have created a special cave which has a blanket fort in it. This is to give users an end goal. This cave was connected with the main environment with a series of stepping stones. All of these stepping stones are hidden because of the dark environment and very few lightings. In order to find each stepping stone, you have to be together with someone, because the other user's light is the only source of illumination that can be used to find the stepping stones which lead to the hidden cave. There are two technical approaches that we tried out to achieve the end effect we had in mind:

1. Directional light from the gaze.
2. Spotlight from the position.

Let's talk about spotlight from the position technique. In this method, we create a spotlight from all the clients' position except the main client. This will create an illusion that all the other players are illuminated and the user is not. This will in turn lead players to stay close to each other in order to explore the environment. This method takes the player position as the beacon in the environment and from the light emitted from the other person, making it easy to locate that person in the dark environment. This method makes exploration a little bit easier, as you don't have to coordinate much about where to go next, because the spotlight illuminates a big chunk of area. So in summary, spotlight method is just like having a light bulb in place of a person's head position, so because of it we cannot identify specifically where the person is looking which is one of the major disadvantage of this method.

The other method of achieving the same end effect is sending in the gaze coordinate, and using a directional light towards those coordinates, which will be a little heavy in terms of processing. With this method we can acquire where the person is looking, and we can create a directional light over that coordinate values, which will give us a pretty good approximation of the gaze from the users. We can use that to create this mechanism of exploring the environment.

We have created a hidden environment connected to the main island with a series of stepping stones in the dark. The users can use the gaze light to explore the environment. They have to work with each other and look towards each next stepping stone for one another in order to progress in the environment. They have to go through a series of stones in order to find the hidden cave, which is designed as a blanket fort, in order to give the users a sense of satisfaction in their discovery.

Design Considerations

To get the better valuation of where the person is looking was one of the important facets of this whole activity and as described earlier, we created two techniques to support this exploration, but finally we have decided to go with the gaze control method, because the use of the spotlight makes it easier to explore the environment. But the major disadvantage is that we don't know the particular direction. That became our major consideration of using the direction light method where we use the raycast from the user's eye cameras and sending the data of the mean of two output points of that collision of the raycast which gives us the approximate point of where the person is currently looking.

Another major consideration is resetting the whole environment. The stepping stones are slightly disconnected and there is nothing between them, so if you don't travel through the stones normally there is good chance that you can fall down from those stones. For this situation we wanted to let the users show how the hidden island should appear only to those who want to explore the environment together with someone else, so if a person falls down from the stone they appear back at the starting point and the other person who is safe and sound has to jump behind the first person and also respawn at the starting point and has to go through the whole exploration from the beginning.

Communication between two people is also really important in the activity, because if you can't communicate with each other, it will be almost impossible to find the hidden cave, so if you sometime without communicating use brute force in order to find the cave, there might be a chance that you get stuck in the dark, where you can't see anything and your friend is far away, so they can't help you either. In those situation both participants have to jump off the stones and they have to start at the beginning.

Future Work and Conclusions about the Gaze Exploration Prototype

We have just focused this scenario on demonstrating how to use gaze data from the VR headset to enhance collaborative interaction. In this activity, users are forced to stay together in order to explore the environment, and need to communicate with each other in order to find the path pieces. We want to explore more around how we can use gaze in other social VR activities, for example how we can utilize gaze control in the office/meeting space activities to increase efficiency in the work environment.

In our view this work is a beginning step in this generation of VR environments. We demonstrated gaze control as part of intimate coordinated activity, which we believe could also be interesting in public spaces. The experiment affords a specific type of social behaviour amongst the people interacting through it. Also how the gaze detection in Virtual Reality can help the users in the other fields test their own system before implementing a very complicated setup of gaze detection.

Social Traces

Introduction and Previous Works

This ambient activity is focused upon how a user can look and interact with the places all the other connected users have visited and spent time with. We have created this activity in order to see what users can interpret with such social traces, and how this facilitates both designers and users in exploring the environment. In this activity, the users can see traces over the places they have visited in the environment as a sphere, and then they can also interact with them in order to identify the sequence of their movement from one place to another.

While exploring the High Fidelity and VRChat platforms, we identified an interesting phenomenon: they have a very large number of different environments and people are scattered all around those environments. Although you can see where your friends are currently, you don’t know where they have spent a lot of their time, and also where you have started from, and which other environments you have explored from the start of your session.

This led us to design a new ambient activity, and for that I started looking for concepts of analytics and I came across Unity analytics[27], which lets you keep track of what users are doing. For example, which level the user has spent the most time in, which helps the designer to identify where they need the improvement in their game. They are a lot of analytics tool out there, but I was thinking about what I could do to show these analytics to the users, how I could represent this in an interactive way.

Further exploring analytics led me to the concept of Social Navigation[28]. We have seen examples of Social Navigation a lot. For example, you are in the library and going through the book catalogue and trying to search for a book about designing in VR, and you find a book that is more worn out and dog-eared then the other books, suggesting that a lot of people have read it. You may decide to give that book a try to start learning, rather than going to the pristine books beside it on the shelf. In this case, instead of relying upon the description or other things in the book, you used information from other people to help you decide. This is a form of Social Navigation. Other example of Social Navigation can be asking for help at an airport helpdesk, who helps you understand how to find the baggage claim—this is also a type of Social Navigation. Social Navigation is a closer reflection of what people actually do, than it is a result of what designers think people should be doing.

Social Navigation doesn’t only take place in the real world. There are so many example of enabling Social Navigation with technology in the digital world. For example, the view count and trending section in YouTube. They help users in a way to decide which videos other people are watching that may interest them too. The view count example[29] can be same as the library, videos on a similar topic are selected between based upon a higher view count, and more popular channels can have a different impact on your choice than a new and smaller channel.

There have been many previous work of Social Navigation with digital media. Peter Brusilovsky[30] integrated social navigation with web lectures. He used implicit and explicit feedback from the system and created a new web-based filtering system that helped users identify the lectures and give more flexibility to share bookmarks as well. CityFlocks[29] helps visitor and new residents of a city to identify popular urban places that are tagged and commented upon by the local community. Their user tests have shown that mobile users really like the gathering of this information from the previous navigators. Another Social Navigation project supported a course recommendation system.[30]A tangent topic to the Social Navigation is the Recommender system, where this collaborative filtering approaches builds a model from a user's past behaviour and similar decisions by other users. This model is than used to predict items that a user may find of interest. YouTube uses a similar matrix to create their model and recommend videos to their users. By their statistics, recommendations account for about 60% of all video clicks from the home page[31].

Methodology

The whole concept is developed upon the creation of virtual spheres spawning in the environment with a timestamp stored in each of them so when they are pressed, a specific action can be played.

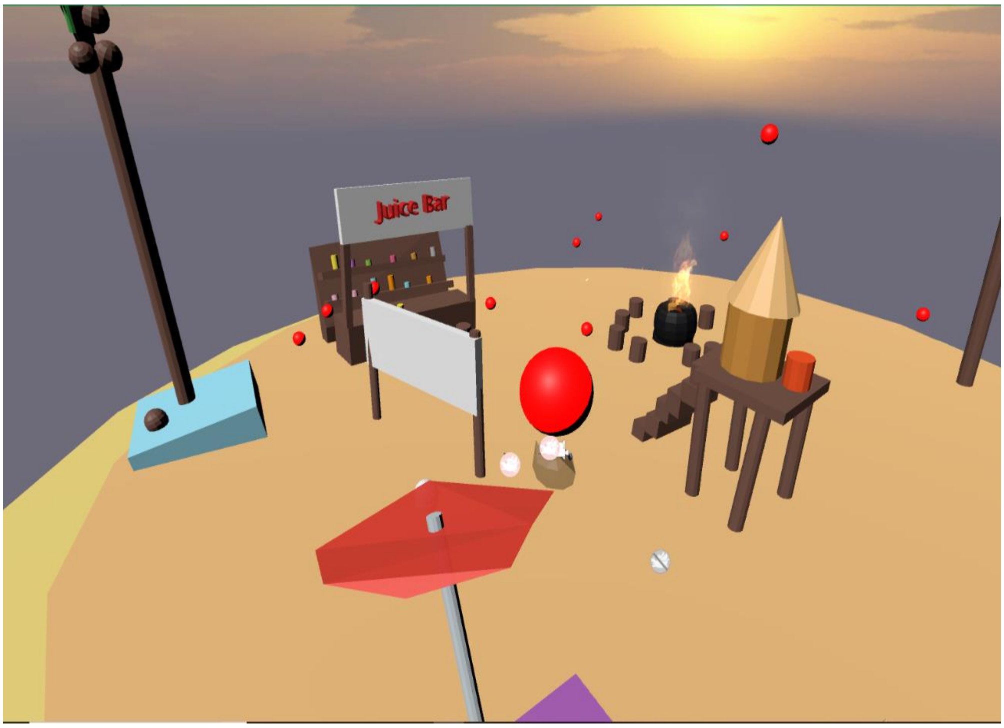

Fig.9: Social Traces and Color of the sphere for the main player

This intervention is created in the existing multi-user experience Ba-Ke-Neko, so one of the biggest design challenges was to make this available for multiple users without adding more data transfer between different users. For this we decided to use local data stored from the person's movement, instead of sending real time data of the user to all the connected users. Because of this choice, we were able to achieve this locally, and not use network bandwidth at all for this intervention. The next part was to create spheres at a user's position. We spawn a sphere with the timestamp that starts from 0, at the player's starting point, and we keep spawning a sphere every 3 seconds. We don't spawn spheres if we see that a person is not moving, then we skip the spawning of the next sphere. In that way we don't overdo the sphere generation, and keep the spheres properly aligned. To identify one user's trace from another's, we color coded the spheres. Local user is allocated with the white color and all the other users spawn with a random color that will be continuous throughout the duration of the experience.

Another thing is to add is control the number of sphere spawning and display clearly on how much time a user spent on a location. The way this is represent in the Ba-Ke-Neko is by checking if the user is in the close proximity of any already generated sphere and increasing the size of that sphere (Fig1, Fig2, Fig3, Fig4) and also adding that timestamp value in a separate array which is used to display the interaction, which is described in the next section.

Fig.10: Change of Scale 1/4

Fig.11: Change of Scale 2/4

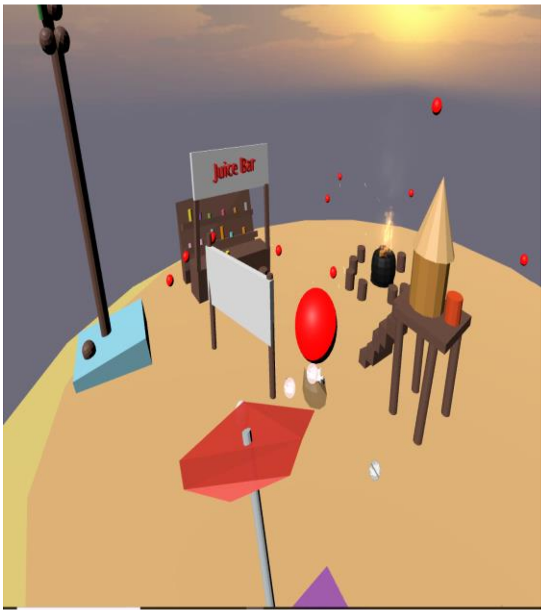

Fig.12: Change of Scale 3/4

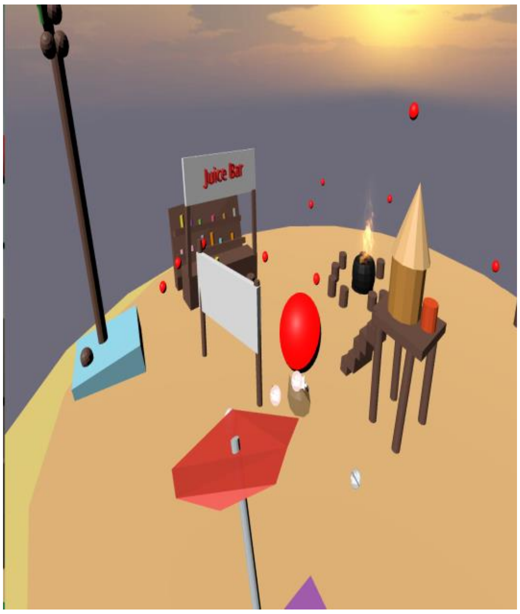

Fig.13: Change of Scale 4/4

Change of the scale overtime when a person is standing near an existing marker

The next step is to create the interactions with the spheres. As the time goes by, the spheres accumulates really fast, and it makes the whole trace very hard to understand, as going to the same place again and again creates more spheres. For that reason, having certain measures that help a user see where they have moved was really necessary, and for that I have created a simple button on the controller. Pressing the spheres will light up and show where a person has moved throughout their time in the virtual environment. This interaction was played for all the users connected in the environment, so a user could see that where all the other users travelled through their time in the environment.

Design Considerations

One of the important things we had to decide was how we wanted to display the traces in the virtual world, and how we wanted to distinguish one person's trace from another's. For that we decided to not make any complicated things and chose spheres as our method to display the trace. One of the main assumption for choosing spheres as markers was to make it seem to users that they are just spawned from their movements, and they not a part of the environments so user won't be able to interact with them. Color coding the spheres was also extremely helpful as users understood in an instant which markers belonged to them and which to other users.

Having an interaction which makes it easy to understand movement was a tricky thing as we have fiddled around with multiple ideas, such as changing the shape of the sphere to square, or changing the scale of the spheres. The problem that persisted among these iterations was they were unable to grab the attention of the users about the starting point of their trace, so we decided to go with particle effects around sphere as it became easier to manage and they grabbed the users' attention.

Future Work and Conclusions about Social Traces

The current version is only focused upon seeing everyone's trace at once, but having more control over the interaction to just see trace of single user, for example the power to only see a friend's social trace, could be actually helpful. Having control over whom to show the movement could be also helpful in this activity.

In this prototype, we have focused upon an ambient activity which has a use for both user and the designers of the virtual environments. A user can identify where all the users have spent the most time and where they can find new people in the environment, and designers can also identify the most popular environments and activities amongst the users as well. This activity can be a great use for both Social Navigation, and also helps create a matrix for recommendation system for the designers.

User Study

After defining domains of meeting spaces and creating three interventions for those domains, we ran a user study. We were interested in general feedback about the prototypes, and had some specific questions about each particular prototyped feature as well.

The study was conducted with 8 participants. (7 Male, 1 Female). All of these participants have been working/developing in VR. The average experience of their development in VR was about 1.2 years. All of them have used at least one Social VR platform. Some of the uses cases they have used the other platforms are as follows:

- Playing games with Friend
- Playing a game that re-enacted based on a real world game
- Meet new people
- Different experience of meeting a friend
- Attend some event in VR

Having users familiar with the concepts of development in VR & experienced in using Social VR application was extremely useful as they helped give a great critic on the controls and interactions and assert their opinion with more experience.

Study Participants and procedure:

As stated above we had a total of 8 participants and 6 of them were individual participants. In those cases, I was the other person in the virtual environment, whereas the one pair entered the world together without me present. My goal was to test out each individual intervention and compare it to a simple blank environment, in order to identify changes in participants' behaviour and receive feedback appropriate to each environment. Some questions that guided my observations and that were used as part of the participant interviews included:

- What domain did these interventions belong to from the participant's perspective?
- Are these activities helpful for a user to connect with other users?
- What other applications could be created by extending these interventions?

The procedure started with exploring the first blank island environment with nothing in it, and the other person started the conversation about the environment and usability of VR in social spaces. Then after a while we slowly introduced Virtual Pet, Gaze Tracking and Social Trace one after the other, by spending around 7-10 minutes in each of them, while having general conversations about all of those interventions.

After testing all of them, we have finished up with a couple of open-ended questions about their preference of activity amongst all three of them, and also we have recorded audio of all the test, so that we can use it to further analyse participants' reactions to each of them.

Results:

As discussed in the earlier section, we have introduced participants to each activity one at a time, so I would first like to discuss the results about each section separately, and then in the last sections, we have laid out some of the general thought about how these activities compare with having the simple environments and also with each other itself.

Virtual Pet:

Overall methodology for this part was to take the participant within the environment with the Virtual Pet and the Dragons, and let them explore more, and after that engage with them in a casual conversation about the design, interaction and their impact in the overall activity from the pet.

Fun and surprising behaviour: Five participants pointed out that the interactions with the pet dog was really fun and identifying the aggressiveness of the user and the behaviour of running away from the participant was actually surprising. As one user responded, “It took me a while but it is actually funny that if I put my hand inside of the pet he just runs away that made me very careful on interacting with this pet.” The dragons were just flying around the environment and which made one user respond to it as, “The dragons swirling around me is really astonishing and it gives me a sense of height difference from the lily pad environment.”

Interactions with pet: Seven participants pointed out to the interaction with the pet being really tricky as they crave user’s attention every time they arrive towards the user. “Not giving any attention to the pet when he is at my feet and seeing the pet really angry made me realize that I should focus upon playing with my pet for a second.”

Alien design of dragons and pet: Two participant referenced the body of the pet being something weird as their factor of being surprised the design choice and they also mentioned that “Having a different design than conventional cat or dog actually gave us time to concentrate and try to interpret the behaviour of the pet.” Other six participants did not have any preference about the shape and looks of either pet or dragon.

Overall people (87.5%) preferred a pet in the environment over the dragon, as the pet has more behavioural traits and it was also more interactive. With one exception of participant leaning more towards the dragon with the reason of, “I just loved the way their flying is animated with the flocking behaviour connected.”

Gaze Tracking and Island Exploration:

For this activity, I just informed participants that there was a hidden island in the environment, and they are tasked to find it. In case of a single participant, I always followed them on their every move so as to let them figure out that they could see the other user’s gaze as their directional light.

Working Together: All eight participants pointed out that this activity has forced them to work together with the other person. Some of the responses were, “It was funny that I left you behind in the dark and then I was also lost in the dark as well, so we both have to jump out, communicate with each other and start working together.”

Personal Space and privacy: As this activity forced people to stay near each other in order to find the next stepping stone, one of the concerns was people felt intrusion into their privacy, but seven out of the eight participants said that the activity did not feel like it's any intrusion in the privacy, and the distance was acceptable. Although I must note that three participants commented that the hidden island was really small in area, so they felt intrusion into personal space on the island. One response was," The island table is just too small even for two persons, and I don't like being this close to other person."

Communication: Seven participants told us that communication was actually a really important part of this activity, as without having any form of communication, it is hard to find the hidden island. Having a common goal encouraged participants to communicate with the other person: "I knew I had to find the hidden island and I wanted him to look for the next rock, that is why I was forced to communicate with the him."

Good difficulty setting: Two participants felt that this activity is like a game, and they have evaluated the activity in terms of difficulty level as well. One response is, "The difficulty was managed properly as it starts with Easy to Hard to Medium to again Hard. It's well balanced level of difficulty of searching the island."

Social Traces:

We wanted to get feedback from the participants about the Social Traces, and how they might envision using this feature in their application,as well as what other interactions they wanted, so we just informed them in the beginning to ignore the spheres currently, and after completing both virtual pet & gaze exploration activities we discussed the Social Traces that accumulated throughout the duration of their experience.

Controlled Interaction: Two participants asserted that being with an unknown person, they didn’t want to see that person’s trace, and did not want to show theirs, too, so having more control over which traces are visible to the participant is important. One response was, “If I am not friends with someone I don’t want them to see what I am doing in the environment.” It’s important to notice here that when they were asked that, if this activity is focused only for friends and family do they still feel the same need of control which they replied with, “No, if I know everyone in the environment that I wouldn’t care about it.”

Application in terms of Data Analysis: While asking about a possible application for this Social Trace, one of most popular replies was to understand their environments and activities, and which of them are more popular with their users and where their users spend their time. This conclusion was super obvious as all eight of the participants were developers in VR, and this information gives knowledge to the developer about the interests of their users. Some responses were,” This data will tell me where my users spend the most time and what they like.”, “I can identify more popular environments and recommend them to the other users.”, “I can find out user’s interest and create similar activities like those to increase the player retention.”

Surprisingly, I did not get much reply about how as a user they perceive this feature, as it was heavily focused upon their perspective as developers and how developers wanted to use that.

Conclusions from the User Study

After running the study, we were able to identify designer and developer thoughts about this intervention. These interventions can be a starting point for VR developers and designers in order to create new activities in the Social VR meeting spaces using the concepts like Gaze Tracking and Social Traces.

Conclusion

I think Ba-Ke-Neko falls into a late stage prototype which is currently intended to other developers in order to understand different possibilities in the Social VR meeting spaces. This work is also useful for all the aspiring social VR developers as this work helps understand the current situation of the meeting spaces in VR. Working on Ba-Ke-Neko taught me on how I can understand the social VR landscape through the lens of activities that people do and how I can interject into them and enhance those activities.

Creating a VR application which everyone can enjoy was my goal and along the way I got to understand what the developers need and hopefully this work will continue to be successful within the developer community and keep creating innovative activities in social VR meeting space.

References

[1] M. A. Muhanna, “Virtual reality and the CAVE: Taxonomy, interaction challenges and research directions,” *Journal of King Saud University - Computer and Information Sciences*, vol. 27, no. 3, pp. 344–361, Jul. 2015.

[2] N.-N. Zhou and Y.-L. Deng, “Virtual reality: A state-of-the-art survey,” *Int. J. Autom. Comput.*, vol. 6, no. 4, pp. 319–325, Nov. 2009.

[3] H. Liu, M. Bowman, and F. Chang, “Survey of state melding in virtual worlds,” *ACM Computing Surveys*, vol. 44, no. 4, pp. 1–25, Aug. 2012.

[4] S. Zhao, “Toward a Taxonomy of Copresence,” *Presence: Teleoperators and Virtual Environments*, vol. 12, no. 5, pp. 445–455, Oct. 2003.

[5] D. Lupiana, C. O’Driscoll, and F. Mtenzi, “Taxonomy for ubiquitous computing environments,” 2009, pp. 469–475.

[6] M. R. Macedonia and M. J. Zyda, “A taxonomy for networked virtual environments,” *IEEE Multimedia*, vol. 4, no. 1, pp. 48–56, Mar. 1997.

[7] “Top Five VR Social Experiences,” *VRFocus*. .

[8] “List Of The Most Popular Social VR Platforms,” *VR News, Games, And Reviews*, 13-Jan-2018. .

[9] J. McVeigh-Schultz, E. Márquez Segura, N. Merrill, and K. Isbister, “What’s It Mean to ‘Be Social’ in VR?: Mapping the Social VR Design Ecology,”

in *Proceedings of the 2018 ACM Conference Companion Publication on Designing Interactive Systems*, New York, NY, USA, 2018, pp. 289–294.

[10] T. Chesney and S. Lawson, “The illusion of love: Does a virtual pet provide the same companionship as a real one?,” *Interaction Studies*, vol. 8, no. 2, pp. 337–342, 2007.

[11] N. R. Aguiar and M. Taylor, “Children’s concepts of the social affordances of a virtual dog and a stuffed dog,” *Cognitive Development*, vol. 34, pp. 16–27, Apr. 2015.

[12] “*Dog’s Life*,” *Wikipedia*. 22-Jul-2018.

[13] “*Pet Society*,” *Wikipedia*. 10-Aug-2018.

[14] “Zoomorphism,” *Wikipedia*. 14-Apr-2018.

[15] “Companion Species - Cyborg Anthropology.” [Online]. Available: http://cyborganthropology.com/Companion_Species. [Accessed: 13-Aug-2018].

[16] M. Böhlen and M. Mateas, “Office Plant #1: Intimate Space and Contemplative Entertainment,” *Leonardo*, vol. 31, no. 5, pp. 345–348, 1998.

[17] M. Romero, Z. Pousman, and M. Mateas, “Alien Presence in the Home: The Design of Tableau Machine,” *Personal Ubiquitous Comput.*, vol. 12, no. 5, pp. 373–382, Jun. 2008.

[18] J. Seyama and R. S. Nagayama, “The Uncanny Valley: Effect of Realism on the Impression of Artificial Human Faces,” *Presence: Teleoperators and Virtual Environments*, vol. 16, no. 4, pp. 337–351, Aug. 2007.

[19] J. N. Bailenson, A. C. Beall, and J. Blascovich, “Gaze and task performance in shared virtual environments,” *The Journal of Visualization and Computer Animation*, vol. 13, no. 5, pp. 313–320, Dec. 2002.

[20] J. N. Bailenson, A. C. Beall, J. Loomis, J. Blascovich, and M. Turk, “Transformed Social Interaction, Augmented Gaze, and Social Influence in Immersive Virtual Environments,” *Human Communication Research*, vol. 31, no. 4, pp. 511–537, Oct. 2005.

[21] S. Benford, C. Greenhalgh, T. Rodden, and J. Pycock, “Collaborative virtual environments,” *Communications of the ACM*, vol. 44, no. 7, pp. 79–85, Jul. 2001.

[22] C. Peters, C. Pelachaud, E. Bevacqua, M. Mancini, and I. Poggi, “A Model of Attention and Interest Using Gaze Behavior,” in *Intelligent Virtual Agents*, Berlin, Heidelberg, 2005, pp. 229–240.

[23] E. T. Hall, *The Hidden Dimension*. Doubleday, 1966.

[24] R. J. Peters and L. Itti, *direction in*. .

[25] M. Sra, A. Mottelson, and P. Maes, “Your Place and Mine: Designing a Shared VR Experience for Remotely Located Users,” 2018, pp. 85–97.

[26] J. K. Hietanen, “Does your gaze direction and head orientation shift my visual attention?:,” *NeuroReport*, vol. 10, no. 16, pp. 3443–3447, Nov. 1999.

[27] U. Technologies, “Analytics,” *Unity*. [Online]. Available: https://unity.com/solutions/analytics. [Accessed: 26-Aug-2018].

[28] A. Dieberger, P. Dourish, K. Höök, P. Resnick, and A. Wexelblat, “Social navigation: techniques for building more usable systems,” *interactions*, vol. 7, no. 6, pp. 36–45, Nov. 2000.

[29] M. Bilandzic, M. Foth, and A. De Luca, “CityFlocks: designing social navigation for urban mobile information systems,” in *Proceedings of the 7th ACM conference on Designing interactive systems - DIS ’08*, Cape Town, South Africa, 2008, pp. 174–183.

[30] R. Mertens, R. Farzan, and P. Brusilovsky, “Social navigation in web lectures,” in *Proceedings of the seventeenth conference on Hypertext and hypermedia - HYPERTEXT ’06*, Odense, Denmark, 2006, p. 41.

[31] J. Davidson *et al.*, “The YouTube video recommendation system,” in *Proceedings of the fourth ACM conference on Recommender systems - RecSys ’10*, Barcelona, Spain, 2010, p. 293.