

Quantifying Immersion in 21st Century Virtual Reality

Final Report

Submitted for the BSc in
Computer Science with Games Development

May 2015

by

Alex Kane

Table of Contents

Introduction	3
Aim and Objectives.....	4
Background.....	6
21 st Century Virtual Reality	6
The Pausch et al. Experiment.....	6
Subsequent Research.....	7
Technical Development	9
Virtual Environment.....	9
Testing Design and Procedure	11
Results and Analysis	13
Evaluation	17
Project Achievements	17
Results.....	17
Further Work	18
Conclusion	19
Appendix A: Initial Project Time Plan	20
Appendix B: Participant Testing Written Consent Form.....	21
Appendix C: Participant Pre-testing Questionnaire	22
Appendix D: Blueprint Scripting – Randomise Target Character.....	22
Appendix E: Blueprint Scripting – Set Target Character Above Door	23
Appendix F: Blueprint Scripting – Hide Target and Randomise all Characters.....	23
Appendix G: Blueprint Scripting – Randomise all Characters without Hiding Target	24
References.....	25

Introduction

The aim of this dissertation was to follow in the footsteps of a research paper by Pausch et al. entitled 'Quantifying Immersion in Virtual Reality', published in 1997 (Pausch, et al., 1997).

It was the goal of the study to ascertain whether the results obtained by Pausch et al. hold true with modern, 21st century virtual reality (VR) technology – namely the Oculus Rift Development Kit 2 (DK2).

This report outlines the aims and objectives of the dissertation and the tasks required to achieve them, presents detailed background information in the chosen topic of research, describes in detail the experimental design and methods used, provides insight into the testing software creation process, analyses and presents the findings of the research conducted, and critically evaluates the project as a whole.

This goal was achieved through the completion of a number of objectives, each comprised of a number of tasks, which are included within the 'Aims and Objectives' section of this report, and the 'Task List', Appendix A, respectively. In brief, these were the creation of a virtual environment for use in testing, the gathering of study participants, testing in a controlled environment using the Oculus Rift DK2, and the analysis and breakdown of the results into usable, useful data as presented in this report.

The subject research of this dissertation was a study conducted by pioneer virtual reality expert Randy Pausch and his team in 1997, designed to study the ability to measure 'immersion' – defined as 'complete involvement in some activity or interest' (Merriam-Webster, n.d.), in virtual reality, and commonly used as a yardstick in the field of virtual reality for its efficacy in defining a quality experience.

The experimental focus of the study asked participants to search for a 'target letter' amongst a set of similarly shaped letters distributed on the walls, ceiling and floor of a virtual environment, either using a head-tracking virtual reality head-mounted-display, or a handheld, custom-built input device and the same display.

The research concluded that the head-tracked display provided more immersion, ascertained by the results showing that the time taken for a subject to declare that the target letter was not present in the environment was significantly quicker using said display than with the non-head tracked condition. This conclusion was justified with the explanation that increased immersion led to less back-tracking when searching the environment (Pausch, et al., 1997).

It was the goal of this new research to determine whether the results of the original study still apply when a modern 21st century virtual reality head-mounted-display is used, compared with a traditional 'Desktop VR' setup, and to increase the reliability of the original research.

Aim and Objectives

To ascertain whether the findings of the 1997 Pausch et al. study 'Quantifying Immersion in Virtual Reality' still hold true in the 21st century, and to increase the reliability of its results.

This aim was achieved by completing the following series of tasks:

1. Creating a virtual environment like that used in the original research.
2. Obtaining participants for use in a scientific study.
3. Conducting a scientific study focused on immersion within the created virtual environment.
4. Collecting, analysing, and compiling the data obtained from testing.
5. Summarising and reporting on the results obtained and whether the aims of the project were achieved.

Objective 1 – Creating a virtual environment like that used in the original research.

To perform this research, a simulated virtual environment (VE) was created using a game engine – specifically Unreal Engine 4, which supports the Oculus Rift natively without the need for manual low-level integration of the hardware.

The VE produced was effectively identical to that created for the original research. Specific sizes of components of the VE such as the size of the walls and individual letters are supplied in the research paper, and were used in the creation of this VE.

This objective included tasks such as researching the ideal engine for use, gaining an understanding of how the Oculus Rift is integrated with the engine, the integration of the Rift, creation of the environment itself in terms of physical geometry and lighting, as well as underlying programming to allow for test conditions to be achieved, such as the display of different sets of letters shown in the environment.

Altogether, this objective was scheduled to take no more than ten weeks of work from start to finish.

Objective 2 – Obtaining participants for use in a scientific study.

Once the VE was built and all relevant testing-related programming completed, participants were found to take part in the experiment. 48 users took part in the original research (using a between-subjects design), and it was planned that at least 24 participants would be found to take part in this study (using a within-subjects design). User age, gender balance and general backgrounds were also planned to be very similar to the original research, with all participants being between 18 and 25 years old, most with no VR experience and mostly undergraduate students.

Participants were found through social media, departmental advertising and word of mouth.

Objective 3 – Conducting a scientific study focused on immersion within the created virtual environment.

After the VE had been created and study participants found, testing could then begin – all testing was carried out using the author's Oculus Rift Development Kit 2 and standard PC hardware, and conducted in the 'Green Room', within the University of Hull's Computer Science department.

This objective included tasks such as finding an appropriate testing location, preparing the room for the testing, briefing, testing, and debriefing the participants, and collecting experimental data.

This objective was scheduled to take at least six weeks to complete.

Objective 4 – Collecting, analysing, and compiling the data obtained from testing.

The data obtained from testing needed to be analysed and formatted into readable, useful results, so that a conclusion could be reached that possessed scientific worth. This was done by the production of graphs, tables and diagrams similar to those in the original Pausch et al. study, and allowed for easy comparison of final results between both pieces of research.

This objective was scheduled to take one week.

Objective 5 – Summarising and reporting on the results obtained and whether the aims of the project were achieved.

The production of this report was designed to produce a scientifically valuable summation and breakdown of the entire research study process from start to end, as well as fulfil the aim of the project; *'To ascertain whether the findings of the 1997 Pausch et al. study 'Quantifying Immersion in Virtual Reality' still hold true in the 21st century, and to increase the reliability of its results.'*

This objective and thus completion of the project was scheduled to take four weeks.

Background

21st Century Virtual Reality

Virtual Reality technology has improved dramatically since the original Pausch et al. study was published, particularly in the domain of head-mounted displays (HMD's) which are the focus of this research.

The explosive technological advancements in this area recently are largely due to the Oculus Rift. The very first Oculus Rift prototype, intended for developers (Development Kit 1 – DK1), was made available to the public and funded through the crowd-funding platform Kickstarter in late 2012, which raised \$2.2m over and above an original goal of \$250,000 (OculusVR, 2012).

Palmer Luckey, founder of Oculus VR, has managed to launch a VR revolution, and the years since the launch of the DK1 have seen a huge leap in technology as well as enthusiasm and general mass-market appeal for the medium. Immersive virtual reality experiences are now within the grasp of the everyday consumer.

Today, the successor to the DK1, the Oculus Rift Development Kit 2, is available for purchase directly via the Oculus VR website for \$350, and was heavily backordered when sales started due to high demand. The DK2 uses an improved 1920x1080 resolution, low persistence OLED display combined with advanced gyroscopes, accelerometers and a positional tracking camera to create an extremely convincing 6-degrees of freedom (6-DOF) virtual reality experience (OculusVR, 2014), and is well suited for comparison to the HMD used in the original study, the Virtual Research Flight Helmet.

Several other VR head-mounted competitors are currently in development, as larger names in the industry attempt to capitalize on the excitement created by Oculus within the audio/visual community and the general public. For instance, Sony is developing 'Project Morpheus' for the PlayStation 4/PlayStation Vita (Sony, 2015), Microsoft announced the development of 'Microsoft HoloLens' (Microsoft, 2015) in January 2015, and Razer announced the 'OSVR Hacker Dev Kit' (Razer, 2015), an open-source take on an affordable, high fidelity display, in March 2015. Additionally, expensive HMDs designed for 3D viewing of movies like the Sony 'HMZ-T2 Personal 3D Viewer' (Sony, 2015) already exist, but are not suitable for consumer virtual reality due to their price and lower specifications. Mobile VR devices such as Samsung's 'Gear VR Innovator Edition' (Samsung, 2015), and Google's 'Cardboard' (Google, 2015), are not included in this discussion due to their numerous shortcomings and unsuitability for the experiment.

The Oculus Rift DK2 was chosen for this research thanks to its affordability, high specifications and availability.

The Pausch et al. Experiment

The original experiment that is the focus of this project was conducted in 1997 by Randy Pausch, Dennis Proffitt and George Williams, and was designed to test a real-world application of virtual reality to ascertain whether, in the opinion of the researchers, virtual reality was an area of research worthy of study and excitement.

The real-world application being tested was the concept of searching a heavily camouflaged environment for a target. Specifically, participants (half using a HMD and half using a stationary monitor) were 'placed' in the centre of a 4 metre by 4 metre virtual room (as shown in Figure 1 in the

'Aims and Objectives' section of this report), and asked to either locate a target letter amongst a large number of similar letters, or to claim it was not present in the scene.

170 letters in total were placed on all surfaces of the scene, including the floor, ceiling and walls. A door and two windows were present for orientation within the environment. Fifty percent of the time the target letter was present in the environment, masked by randomly assorted letters from one of two sets, AKMNVWXYZ (all of whose primary features are slanted lines) or EFHILT (straight lines) placed in the other 169 locations. The other fifty percent of the time, the target letter was not present. The time taken for participants to find the target letter or to claim it did not exist was measured, for both the VR and static display sets.

It was found that the VR users were not significantly any better than the stationary monitor users at finding the target when it was present, but they were 41% quicker at determining no target was present. The researchers believed this was thanks to less re-examination of areas already searched, due to the VR users being more immersed within the environment.

Subsequent Research

The Pausch et al. study has been heavily cited since its publication, and stands as an example of one of the very first pieces of research into immersion in virtual reality.

One important piece of subsequent research based on its findings is the 1997 study by Microsoft Research entitled 'Immersion in Desktop Virtual Reality' (Robertson, et al., 1997). The experimenters attempted to recreate part of the study performed by Pausch et al. but posited that the experimental procedure and equipment used (i.e. using the same HMD) was not indicative of 'Desktop VR' (that is, a 3D virtual environment presented on a desktop computer) as claimed by Pausch et al (Robertson, et al., 1997).

Robertson et al. believed that the Pausch et al. implementation was so different to 'Desktop VR' that they used a different term – 'Fixed HMD VR' (Robertson, et al., 1997). Robertson et al. performed a similar experiment to that of Pausch et al. by testing thirty-two participants in a hidden-character search task within a virtual rotational-study room identical to that of the original study. Participants used the experimenters' definition of 'Desktop VR' – a stationary desktop computer monitor, along with one of two input methods – a Microsoft two-button serial mouse, or the two-handed Polhemus tracker used in the original Pausch et al. study. Each participant performed two sets of 18 or 12 searches for a character camouflaged within the environment, with each character from the same two confusable letter-sets as in Pausch et al. appearing once in a target-present condition, and once in a target-absent condition. This led to three main independent variables being manipulated – input device, present vs absent, and letter set.

Robertson et al. concluded that the results of Pausch et al. did not apply to 'Desktop VR', as they observed predicted times for target absent conditions rather than the slower times that Pausch et al. observed (Robertson, et al., 1997). Additionally, the experimenters observed much faster search times overall than those observed by Pausch et al. This was believed to be due to the smaller viewing angle of Desktop VR compared to HMD VR making the environment faster to scan, suggesting an advantage of Desktop VR over HMD VR (Robertson, et al., 1997).

Robertson et al. theorised that by compensating for the difference in resolution between the Pausch et al. HMD and a traditional desktop monitor by using the HMD (fixed in place for 'Desktop VR') for both conditions in the original study, the experimenters skewed the results towards HMD VR (Robertson, et al., 1997). Additionally, they submit that the unfamiliar and cumbersome method of controlling the 'Fixed VR' avatar (a wired Polhemus tracker held in the hands) may have affected the results (Robertson, et al., 1997). The study also claimed at the time that '...current HMD-based VR

techniques suffer from poor display resolution, display jitter, and lag between head movement and the resulting change to the display. These problems tend to inhibit the illusion of immersion...' (Robertson, et al., 1997).

It is the author's belief that the HMD chosen for this new study was of such increased sophistication than those around at the time of the aforementioned experiments, that these issues are no longer relevant. The new research of this study attempts to find a middle ground between the Pausch et al. and Robertson et al. studies, using the increased resolution and response time of the Oculus Rift DK2 HMD, compared with a traditional desktop monitor and a familiar input device (an optical mouse), thus fulfilling the criteria of Robertson et al. for 'Desktop VR'.

By replicating both the testing of the original Pausch et al. study using a modern VR HMD, and the 'Desktop VR' conditions of the subsequent Robertson et al. study, these experiments are brought into the twenty-first century and are easily compared in a modern setting.

Technical Development

Virtual Environment

Before the creation of any content for the project could begin, a decision had to be made on which engine to use - the choice being between Unity and Unreal Engine 4 (UE4). At the time, virtual reality content could only be created in Unity with a Unity Pro license, costing \$1,500. Conversely, UE4 utilised a very inexpensive subscription model (and has since turned free to use) (Epic Games, 2015) and was on the forefront of virtual reality support and features.

After UE4 was chosen as the engine base for the project, the next step was Oculus Rift integration. UE4 is yet to implement an official Project Template for the creation of content for the Rift, but a user on the Unreal Engine forums by the name of 'mitchemmc' created an unofficial template, which after research and testing, was settled on to create the project. Setting up UE4 to work with the Rift was still a task requiring significant effort, however the unofficial template makes the process much simpler and worked well after some initial problems were dealt with. The template and source are available on GitHub (Mitchemmc, 2014) and it has become a popular choice for developers creating VR content.

An explicit set of behaviours and features required to achieve testing conditions akin to those conducted by Pausch et al were drawn up:

- Ability to display 170 letters, each 0.6m long, on the walls/floor/ceiling of the environment.
- Ability to randomly choose a target letter from one of two possible sets, slanted characters and straight-line characters.
- Ability to randomly place this target letter in one of the 170 predefined positions.
- Ability to randomly distribute all the other letters from the target set to the other 169 positions, to mask the target letter.
- Ability to also place the target letter above the door to make it clear to the participant which letter is the target.
- Ability to either place the target letter in the scene, or to randomly distribute target set characters to all 170 positions.

These behaviours/features were achieved through three steps; construction of the environment itself, placement of characters on the surfaces of the environment, and scripting the testing behaviour.

To construct the basic virtual environment of the 'rotational study room' (Pausch, et al., 1997), UE4's built-in brush feature was used to create a 4m³ room. A solid door and window frames were added from UE4's example content, subtractive brushes were used to create holes in the walls for the windows (which were then filled with an opaque glass surface), and a spotlight was placed in the centre of the room near the ceiling to provide adequate lighting for the environment.

The next step was investigating how to display text on the surfaces of the environment. UE4 has a built in '*TextRenderActor*' which is used for showing text in the environment, and each object's properties can be modified dynamically in-game. Many *TextRenderActors* were needed to display the required amount of characters randomly within the environment, so an actor blueprint (akin to a class) was created to define behaviours common to each *TextRenderActor*. Next, the character objects needed to be arranged upon the walls, floor and ceiling of the virtual environment. Each character was set to have a scale of 1.0, with a 'world size' of 0.6, creating 60cm tall characters. The characters were then given a centre formatting to line them up correctly. Firstly a *TextRenderActor* containing six lines of six characters was used as a template for placement on the walls to ensure equal spacing.

Individual *TextRenderActor* objects containing single characters were then positioned over the template to maintain correct placement. Once the wall characters had been placed, more *TextRenderActors* were used to place letters on the floor and ceiling to match those in the original study.

Figure 1 shows the finished virtual environment in pre-test mode, compared with Figure 2, the original Pausch et al. environment. It was decided to mark the example target letter above the door in a different, distinct colour, to make it clear to the participant the target they were searching for. This change did not affect testing in any way, it simply aided the participant in confirming the target character.



Figure 1 - The author's recreation of the 'rotational study room' with un-randomised characters.



Figure 2 - The 'rotational study room' as created by Pausch et al. (Pausch, et al., 1997)

To achieve the testing behaviours required by the experimenter, a function was created within the extended *TextRenderActor* blueprint named 'SetTargetChar', to randomly set the target character. This was done by first randomly choosing one of the two possible character sets, 'Slanted' (A K M N V W X Y Z) or 'Horizontal/Vertical' (E F H I L T), by generating a random integer between 0 and 1, using the 'Random Integer' node. The resulting integer was then used with a 'Switch on Int' (integer switch) node to set a Boolean variable 'IsSlanted' (which is used in other functions) to either true or false depending on the outcome. From this, another random integer was generated (either in the range 0-8 for the slanted set, or 0-5 for the straight set) and compared with an integer switch to choose which of the actual characters in the set was made the target. The 'TargetChar' variable was then updated with the result and passed out of the function.

With the target character set, the next step was to distribute it, along with all the other characters in the target set, on all the surfaces of the environment. Firstly, two key-presses were chosen to decide whether the target was actually going to be hidden in the environment each time, setting the Boolean 'AreWeHidingTargetChar' to either true or false. If set to true, the 'Get All Actors of Class' node was used to retrieve a list of every *TextRenderActor* in the environment, and store it in an array. The order of contents of this array was then randomised using the 'Shuffle' node, to allow for proper randomisation of the target character's placement in the environment. Subsequently, a 'ForEachLoop' node iterated through every element in the array and for each a 1% chance was given using the Random Integer node in a range 0-100, to execute the 'HideTargetChar' function and place the target letter at this element. Once this function had executed, a Boolean *TargetHidden* was set to stop the function being executed again. For all other elements in the array, the 'DistributeChars' function was executed, randomly choosing a letter from the target set and assigning it to the character in the

environment. If the entire array of 170 characters had been iterated through and the target still had not been set (due to the 1% chance never occurring), the set of functions and checks executes again, until the target character has been set and the *TargetHidden* Boolean is set to true.

If *AreWeHidingTargetChar* was set to false by the experimenter's key-press, all the characters in the environment are again retrieved and stored in an array, but this time every one is set randomly to a letter that is *not* the target.

Operation of the testing conditions and scripting by the experimenter was done through simple key-presses connected to the relevant blueprint behaviour – one to choose a random target character (Appendix D), one to assign the target character to the distinctly coloured *TextRenderActor* above the door (Appendix E), one to randomise the 170 black characters within the chosen letter set and hide the target in the environment (Appendix F), and one to simply randomise the black characters without hiding the target (Appendix G).

Testing Design and Procedure

Subjects:

Twenty-eight (nineteen males) subjects participated in the study. All had normal or corrected-to-normal vision. Participant ages ranged from 18-25 years old. All participants were undergraduate students.

Stimuli and Design:

In all testing, the principal dependent variables were response time from onset of display, and response accuracy. All subjects were instructed to maintain a high degree of accuracy.

A 2 X 2 repeated measures design was used. Two independent variables were manipulated. Firstly, the display and input device used was either an Oculus Rift Development Kit 2 HMD (with 6 degree-of-freedom movement input), or a 24 inch LCD monitor and Microsoft optical mouse. All subjects used both input/display conditions. Secondly, half of the searches conducted by subjects contained the target item, half did not. The specific character used as the target item (chosen from either an angular set; AKMNVWXYZ or a vertical set; EFHILT) and its position in the environment were chosen randomly for each search trial.

Procedure:

Subjects were welcomed upon arrival at the testing location (shown in Figure 3), made comfortable and asked to read and complete two identical informed consent forms (Appendix B), and a short testing questionnaire (Appendix C). One consent form was given to the subject to keep, and one was kept by the experimenter. The subject was reminded of their right to withdraw from the study at any point. No subjects answered affirmatively to the question regarding epilepsy/visual problems on the questionnaire, therefore no subjects were discounted from the testing for this reason.




Figure 3 – The testing location, and the testing equipment used by participants.

The testing procedure was explained to each subject to ensure understanding of the task required of them. At this point, subjects who were using the Rift condition first were introduced to it, the risks of simulation sickness were explained, and the subject's right to withdraw was once again mentioned. Rift-condition subjects put on the HMD after an explanation on how to adjust its fit, and were advised to keep their eyes closed until the experimenter transferred the display to the Rift's screens to avoid the potentially jarring transition. Subjects in both conditions were then able to look around the environment and become familiar with its features, and once again the search task was explained to ensure understanding.

Subjects were asked to complete a minimum of two un-timed practice searches for a present target, and given the opportunity for as many more practices as required – no subjects requested more than one more practice search. At this point the first portion of testing commenced; five timed searches for a present target. Subjects began each search facing the target character, displayed in blue above the door. The target character was then randomised, and the subject was asked to confirm the character they were searching for. Subsequently, all the black characters in the environment were randomised amongst the randomly chosen target character set (angular or vertical), and the target character was placed into the environment in a random location. Subjects were instructed to begin their search, and the timer was started, as soon as the black characters randomised. Subjects were instructed to say 'Found it' out loud as soon as they saw the target character in the scene, at this point the timer for each search task was stopped and the time taken recorded.

Following the initial five 'target present' searches, subjects moved on to the main testing section – ten searches for a target that may or may not be present in the scene, using either the Rift or Desktop test condition. Subjects were instructed that the target character may or may not be present in the scene, and were told to search until they either found the target and verbally acknowledge this (as before), or until they were comfortable that it was not present. Subjects were instructed to maintain a good degree of accuracy. On conclusion of these ten search tasks, subjects were asked to rest for five minutes before performing another ten randomised searches using the inverse testing condition.

Upon finishing all twenty 'present-not-present' search tasks, subjects were thanked for their time and participation in the study, debriefed, and any questions were answered. Subjects were provided with the experimenters' contact information in case of further questions or concerns, and departed the testing location.




Figure 4 – A subject about to perform the Desktop searching portion of testing.



Figure 5 – A subject performing the Rift searching portion of testing.

Results and Analysis

The results obtained from participants during testing were collated in a spreadsheet and each individuals' set of results was averaged within each test condition to form a single data point, as in the Pausch et al. study (Pausch, et al., 1997).

A total of sixteen errors were made by participants on the Rift configuration (a 6.67% error rate), all of which were made by declaring a target as not being present in the scene, incorrectly. A total of seventeen errors were made by participants on the Desktop configuration (a 7.08% error rate) - fifteen false-negatives. Two of the errors however, were due to participants claiming to have seen the target character when it was in fact not present in the scene. There is no significant difference in the error rates of the two testing configurations. Incorrect responses were omitted from the mean of each participant's results.

Table 1 shows the mean results obtained from twenty-eight participants, testing each participant against both 'Target Present' and 'Target Not Present' conditions, on both the Rift and Desktop testing configurations.

The results were analysed to determine if they fit into a Gaussian (normal) distribution, and hence could be analysed using parametric statistical tests. If the Kurtosis value of each set of results falls within the range 2-4 it could be said to be 'mesokurtic', within a normal distribution, and thus suitable for parametric testing. Kurtosis values for each set of data can be found at the bottom of Table 1. The Kurtosis values of the 'Target Present' condition for the Rift configuration made these results ineligible for accurate parametric testing. The Kurtosis values of the 'Target Not Present' condition for the Rift and Desktop configurations allowed for accurate parametric testing.

A Mann-Whitney-U test was conducted on the 'Target Present' data, and the difference was not found to be significant (Mann-Whitney $U = 358$, $n_1 = n_2 = 28$, $Z = 0.659$, $P = 0.582$ two-tailed).

From this result, we can state that there is no significant effect on search times (when a target is present in the scene) between the use of the Rift and the use of Desktop virtual reality. These results are presented in Figure 6. The results of the same experiment as conducted by Pausch et al. are shown for comparison in Figure 7. While still statistically insignificant, the difference in search times between the VR and Desktop conditions in the Pausch et al. experiment in this test case are visibly much more pronounced than those observed in our results.

Participant	Rift Configuration		Desktop Configuration	
	Present	Not Present	Present	Not Present
1	8.53	20.08	10.27	20.35
2	22.99	44.37	27.58	44.16
3	10.73	19.69	7.76	21.8
4	21.59	42.63	13.29	51.6
5	15.83	51.49	19.63	44.86
6	11.43	37.05	22.07	52.48
7	5.81	20.02	12.12	20.74
8	12.03	28.99	9.7	24.12
9	9.92	17.82	10.48	29.68
10	14.75	32.81	24.13	36.51
11	17.98	34.96	8.04	33.18
12	10.97	26.93	6.04	31.14
13	7.12	15.71	12.94	20.01
14	12.32	25.48	15.19	29.51
15	15.57	22.41	12.61	25.58
16	10.66	21.93	5.62	25.96
17	39.37	79.88	39.41	93.94
18	18.82	34.25	11.9	24.73
19	19.83	39.79	14.58	49.69
20	16.57	35.13	19.6	46.65
21	9.35	21.14	10.89	19.17
22	10.88	47.02	14.38	21.38
23	9.56	21.13	14.62	25.33
24	13.95	32.43	20.57	36.3
25	10.42	44.29	25.01	46.39
26	14.57	22.28	16.21	29.42
27	25.81	69.53	18.37	75.04
28	12.69	28.83	11.72	26.56
Mean	14.64	33.50	15.53	35.94
Std Dev	6.85	15.26	7.37	17.45
Kurtosis	5.28	2.78	2.48	3.76

Table 1 - Average time (seconds) taken for participants to find a camouflaged target or to declare it was not present.




Figure 6 - Average time for users in this study to locate a target present in the scene. Error bars show standard error (in parentheses).




Figure 7 - Average time for users in the original study (Pausch, et al., 1997) to locate a target present in the scene. Error bars show standard error (in parentheses).

When considering search times for a target that was not present in the scene, we can predict the time taken for a perfectly efficient search as described by Pausch et al.:

If the targets are dense, and the users are efficient in their searching, we can predict how long this will take. Working backwards, consider an efficient user who takes 40 seconds to completely search a scene, with no wasted effort. On average, when a target is present, that user should find it in 20 seconds. Random placement may make the letter appear earlier or later in the search process, but on average the user will find the target halfway through the search. We know how long it takes users to find targets when they are present. If the users searched perfectly, it should take twice that long to search the entire room and confidently conclude the target is not there. Any time over that would imply that the users were re-examining portions of the room that they had already searched. (Pausch, et al., 1997)

Using this method, we can predict perfect search times for non-present targets for both the Rift and Desktop test conditions, as in Figure 8. The predictions made by Pausch et al. are shown in Figure 9.




Figure 8 - Predicted times for users in this study to perform a complete search of the environment.

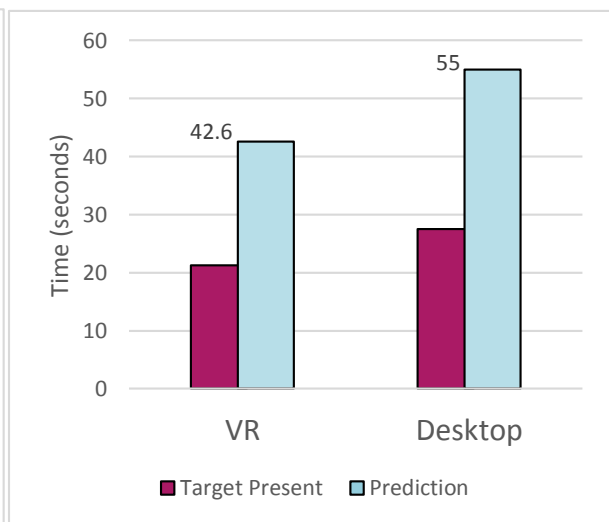


Figure 9 - Predicted times for users in the original study (Pausch, et al., 1997) to perform a complete search of the environment.

A paired two sample means T-test was conducted on the 'Target Not Present' data, and the difference was not found to be significant (conditions; $t(27) = -1.59, P = 0.122$). These results are displayed in Figure 10, compared with the average time taken to locate a target present in the scene (Figure 6 results) and the predicted time to search the entire scene (Figure 7 results). Figure 11 shows the same set of results, as published by Pausch et al.




Figure 10 - Average time for users in this study to search the entire environment and decide that no target was present. Error bars show standard error (in parentheses).




Figure 11 - Average time for users in the original study (Pausch, et al., 1997) to search the entire environment and decide that no target was present. Error bars show standard error (in parentheses).

The difference between the predicted and observed times (to search the environment and decide the target was not present) in the VR condition was 12.6% in this study, compared to 1.4% observed by Pausch et al. Under the Desktop condition, the differences were 13.58% and 41% respectively. These findings support informal observations that participants often rescanned areas of the environment before concluding that no target was present, under both testing conditions. Figure 12 shows all experimental results on one graph, for ease of comparison.




Figure 12 – Combination of Figures 6-11, showing average times for users in both studies to find a target present in the scene, predicted times for complete environment search, and average times for users to search the entire environment and decide that no target was present.

Robertson et al. conducted the same experiment, asking subjects to search for a target that was either present or not present within the rotational study room. Instead of directly comparing HMD-VR and (their interpretation of) Desktop VR as previously done by Pausch et al., the researchers compared input methods – a traditional computer mouse, and the same two-handed Polhemus tracker used in the original study. The experiment also controlled the Target Character’s letter set. The results of the Target Present and Target Not Present searches across both letter sets using the mouse input are each averaged to compare with our results, as presented in Figure 13.

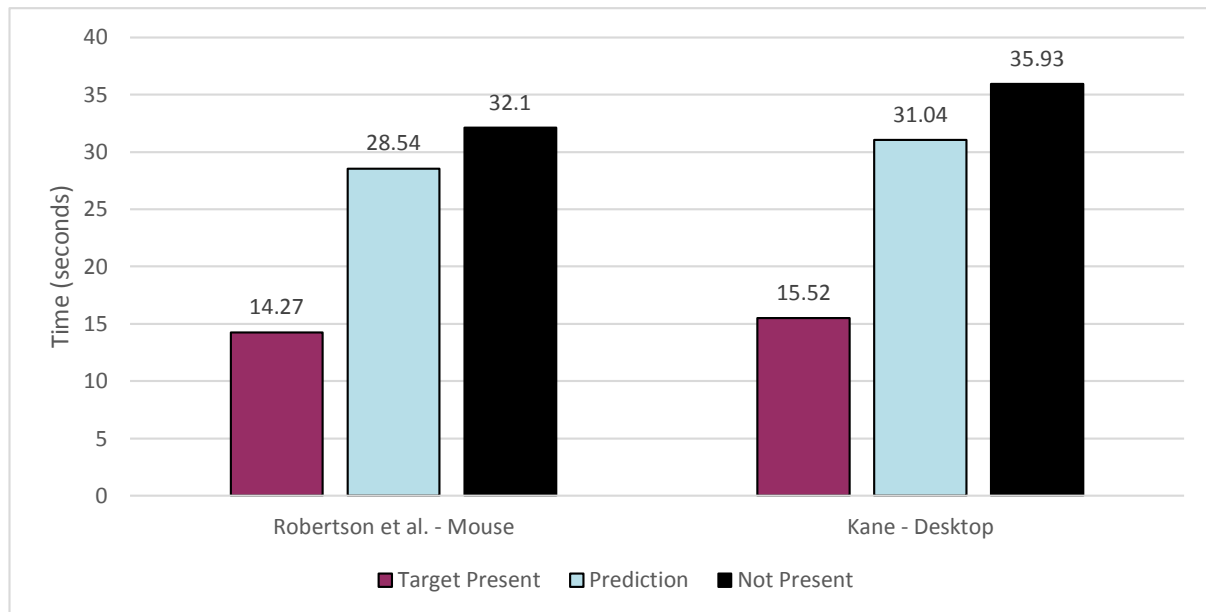


Figure 13 – Average search times for users in this study and in the mouse input condition of the Robertson et al. study (Robertson, et al., 1997) to locate a target in the environment or determine that it was not present.

The average Target Not Present search time in the Robertson et al. experiment was only 11.1% over the predicted ‘perfect’ search time, very close to the 13.58% observed in our testing. Robertson et al. claimed that this suggests that the findings of Pausch et al. do not apply to Desktop VR, a claim that is also supported by our results.

During post-testing questioning, many participants verbally stated that they felt more immersed during the Rift testing condition than when using the traditional desktop display, as well as a sense of being ‘more in the environment’ – no participants expressed a sense of greater immersion in the Desktop condition.

Many participants also personally concluded that searching using one particular testing condition was ‘easier’ or ‘quicker’, with an equal number of participants expressing preference for each condition. Search behaviour was extremely varied between participants, from slow methodical searching in specific, repeated movements, to much faster, jerky movements in patterns that changed between searches. These varying search behaviours were exhibited by participants under both testing conditions.

Significant re-scanning behaviour under both conditions combined with such close timings between conditions, as well as the informal findings mentioned, suggest that comparing target-not-present search times to predicted ‘perfect’ search times is highly variable between participants, and not an effective measure of immersion in 21st century virtual reality. However, using this method, the results of this study suggest that there is no significant difference in immersion between HMD virtual reality and Desktop virtual reality, disputing the findings of Pausch et al. (Pausch, et al., 1997) and supporting the claims of Robertson et al. (Robertson, et al., 1997).

Evaluation

Project Achievements

New game-engine technology was investigated and learned in the construction of a custom virtual environment for testing, including the use of a bespoke visual scripting language to obtain the required software behaviour (achieving Objective 1). A successful full-scale scientific research project was undertaken involving human participants, including dealing with all the ethical concerns it entails (achieving Objectives 2 and 3). Techniques for statistical data analysis were learned and implemented with the data obtained from testing. These results were interpreted against the results of two previous virtual reality studies, and findings of scientific worth were obtained, furthering research into the fast-evolving field of modern virtual reality (achieving Objectives 4 and 5).

Results

While the results gathered from testing were previously analysed and compared against those reported by both Pausch et al. and Robertson et al., there is room for further insight and speculation on why these results were obtained. The significant improvements in search times for subjects using the Oculus Rift DK2 in comparison to the HMD used by Pausch et al. (for both target present and not present conditions) can be explained by comparing the technical specifications of the two HMDs.

The head-mounted display used in the original Pausch et al. testing, the Virtual Research Flight Helmet (VRFH), was at the time a very full-featured product, featuring a large field of view, comfortable ergonomics compared to other products at the time, and excellent lenses (VRifacts, 2010). However, the HMD was let down by a very low resolution LCD display of just 240x120 pixels – “the image becomes a collection of colored blocks rather than a smooth, continuous picture.” (Hezel & Veron, 1993). This issue of visible individual pixels has become known as the ‘screen-door effect’ in modern VR, and is a problem that is still being tackled today. The advent of small, affordable, high-resolution displays for use in smartphones is a driving force behind the abolition of this phenomenon. While the Oculus Rift DK2 does not achieve a resolution of 2048x2048 pixels - offered by Hezel & Veron as a “reasonable goal for HMD resolution” (Hezel & Veron, 1993), the 8/9x more pixels on display using the DK2 in comparison to the VRFH provide a far clearer image to the user.

	Virtual Research Flight Helmet	Oculus Rift Development Kit 2
Resolution and Display	240x120 pixels, 2x LCD displays	1920x1080 pixels, 1x OLED display.
Input	6 DOF tracking	6 DOF tracking
Cost at Time	\$6000+	\$350
Field of View	Approx. 100 degrees	100 degrees (nominal)
Weight	1.67kg	0.44kg
Frames per Second	60	75
Latency	100ms	20ms

Table 2 – Comparison of technical specifications of the HMD used by Pausch et al. against the Oculus Rift DK2.

Additionally, several other factors combine to make the Rift DK2 a more pleasant and immersive experience for the user, as shown in Table 2. Firstly, the cost of the VRFH was completely prohibitive to any non-research or military applications, compared with the still expensive, but much more affordable DK2. Secondly, the DK2 is almost four times as light as the VRFH, providing a much more comfortable experience. Finally and very importantly, the DK2 has one third of the VRFH’s ‘motion-to-photon’ latency. This is the time taken for user input to be fully reconciled to a change in the user’s

view of the environment. Low latency is an extremely important part of an immersive virtual reality experience, as a delay like that present when using the VRFH can lead to simulation sickness and easily disconnects the user from the virtual environment.

The improvement in technical specification between the VRFH and the Oculus Rift DK2 explains the improved search times observed in our results – subjects can more clearly see the virtual environment and the characters within it so are able to spend less time processing each individual character. Additionally, movement of the subject's head is more quickly processed into movement of the camera within the virtual environment, avoiding a jarring 'lag' which could slow subjects down.

Improvements

While the aims and objectives of the project were successfully achieved, there are areas in which the project could have been improved.

Firstly, while the original time plan and schedule produced for the project was realistic, due to unforeseen circumstances experimental testing could not commence until several weeks after the scheduled date. Fortunately, the large amount of time allotted for testing was sufficient to obtain the results necessary for the study. In future, even more time would be allotted for pre-testing requirements such as location booking, equipment sourcing, etc.

Additionally, a statistical power analysis was not carried out prior to testing due to a number of factors. This would have been useful to ascertain the required amount of participants needed for testing to achieve the results required beforehand, reducing the amount of time required for testing.

Finally, a potentially more representative sample of the population could be obtained for testing – all participants were students aged 18-25, with a good degree of familiarity with technology. It is possible that a sample consisting of participants in a larger age range may be more representative of the population as a whole. Although 68% of participants were male, there was no significant difference in results obtained for males compared to females, thus it is not believed that gender disparity in the participant sample was a problem.

Further Work

While the project itself was completed in full, it is but a small study in an extremely large area of ongoing research. There is also room for further study in this specific area - particularly in the conducting of modern equivalents for the additional experiments conducted by Robertson et al., such as the Navigation Hallway (Robertson, et al., 1997). Additionally, strong developments in virtual reality technology and the widespread development of head mounted displays by large companies such as Sony open the door to important questions in the field of virtual reality experiences such as; what experiences work well in virtual reality? How can video games be adapted to provide interesting and immersive experiences in virtual reality? How can simulation sickness be best avoided? These questions, and many more, are the focus of ongoing research.

Conclusion

The overall aim of the project was to ascertain whether the findings of the 1997 Pausch et al. study into virtual reality hold true when using modern 21st century technology. Pausch et al. claimed that virtual reality as experienced on a head-mounted display was more immersive than virtual reality as experienced in a desktop configuration, because experimental subjects were able to more quickly search a virtual environment and determine that a target character (which may or may not have been present) was absent from the scene, on the former (Pausch, et al., 1997). These findings were disputed in a subsequent study conducted by Robertson et al., which claimed that by using the same HMD for their desktop condition, Pausch et al. had not accurately portrayed desktop virtual reality. Robertson et al. repeated the experimental trial performed by Pausch et al. using a more traditional monitor and mouse setup, and found that the conclusions of the original study no longer applied.

By repeating the original experimental testing in a modern setting - comparing the Oculus Rift head mounted display against a contemporary monitor and mouse setup (a more representative form of 'Desktop VR' as suggested by Robertson et al.), we have obtained results that suggest that the findings of Pausch et al. do not apply when using 21st century virtual reality technology.

We also propose that the basis by which the level of immersion within a virtual environment is measured, as put forward by Pausch et al., is not an effective one. The experimenters suggest that, compared to the predicted, 'perfect' time to completely search the environment and determine the target is not present, any additional time spent is time used re-searching the environment, implying less immersion within it, due to not remembering where has been searched. According to the results of our quantitative testing using this criterion, subjects were no more immersed in one condition than the other. However, many subjects verbally expressed a feeling of more immersion when using the head-mounted display, while none felt that Desktop VR was more immersive. Re-searching appeared to be a highly individual factor during testing, largely down to the subject's specific searching behaviour – quick, flustered searchers tended to re-search areas more than slower, methodical searchers, regardless of testing condition.

Advances in computer graphics, VR and HMD technology, and numerous upcoming large-scale products in the virtual reality space, mean that fully immersive experiences will likely be within the grasp of consumers in less than 24 months. The Oculus Rift Developer Kit 2 is one of the first steps towards truly immersive virtual reality, and while not there yet, is a tangible vision of what's to come.

Appendix A: Initial Project Time Plan

		University Calendar Weeks																																		
#	Task Name	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36						
1	Topic Research	█																																		
2	Finalise Spec	█	█																																	
3	Choose Engine		█	█																																
4	Research Rift Integration			█	█																															
5	Integrate Rift				█	█																														
6	Create VE (Basics)					█	█	█																												
7	Create VE (Scripting)						█	█	█	█																										
8	Finalise Testing Data									█	█																									
9	Prepare Testing Equipment										█	█																								
10	Test Run											█	█																							
11	Interim report											█	█	█	D																					
12	Find Participants															█	█	█																		
#	Task Name	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36						
13	Organise Testing Room																		█	█																
14	Organise + book Participants																			█	█															
15	Perform Testing																				█	█	█	█	█	█										
16	Interpret Results																									█	█									
17	Final Report																																	D		

Calendar weeks correspond to those shown on the University of Hull Online Timetables web-page (Scientia, 2014).

Appendix B: Participant Testing Written Consent Form

Appendix III (F) - Consent Form: For Surveys & Questionnaires

Alex Kane, BSc Computer Science with Games Development, Department of Computer Science, University of Hull

Introduction/Purpose

You are invited to voluntarily participate in a study on quantifying immersion in 21st century virtual reality. This research is being conducted by **Alex Kane** at the Department of Computer Science. In total, we will have approximately **20-40** participants in this study.

Procedure

It will take approximately **20-40 minutes** to complete the testing.

The purpose of this study is to ascertain whether the use of virtual reality equipment in the twenty-first century leads to greater immersion within an environment than conventional display and input technology. You will be asked to locate letters within a virtual environment, and timed on how long you take to complete this task. A full description of each task and the procedure will be given before each task.

Voluntary Participation

You are free to choose whether or not to complete the study. You may stop the procedure at any time without loss of any benefits of participation and any information obtained from you will not be used.

Anonymity/Confidentiality

Any information concerning you and your participation in this study will be kept private and confidential. Your results will not be associated with your identity. If information about you is published it will be in a form such that you cannot be recognized. Data for the study will be used in scientific reports, but no names or identifying information will be included in these reports.

Contacts

The researcher will be happy to answer any questions that you might have about taking part in this study. If complaints or problems concerning this research project should arise, they should be reported in the first instance to **Alex Kane** at Department of Computer Science, The University of Hull on **07985540100** or to the Chair of the Ethics Committee, Department of Computer Science, The University of Hull on **01482 466948**

Thank you for your effort and honesty in your participation today.

Alex Kane

WRITTEN CONSENT

By signing here I consent to voluntary participation in this research study. I understand the procedures to be followed and the guarantees and limits of confidentiality. I understand that I will also receive a signed copy of this consent form.

Name: _____

Signature: _____

Date: _____

Appendix C: Participant Pre-testing Questionnaire

Oculus Rift – Dissertation Testing Questionnaire

Participant Number (to be filled in by experimenter):

Please circle the relevant answer.

Question 1: Do you have any previous experience with Virtual Reality?

YES

NO

Question 2: Do you suffer from motion sickness?

YES

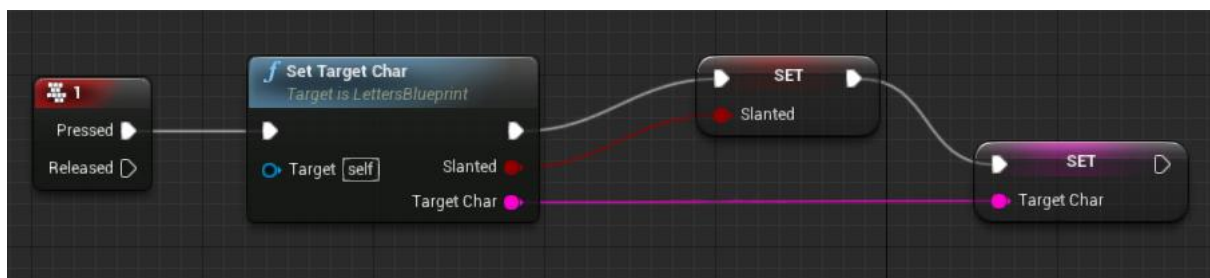
NO

Question 3: Do you have any history of epilepsy, seizures or any other visual impairments?


YES

NO

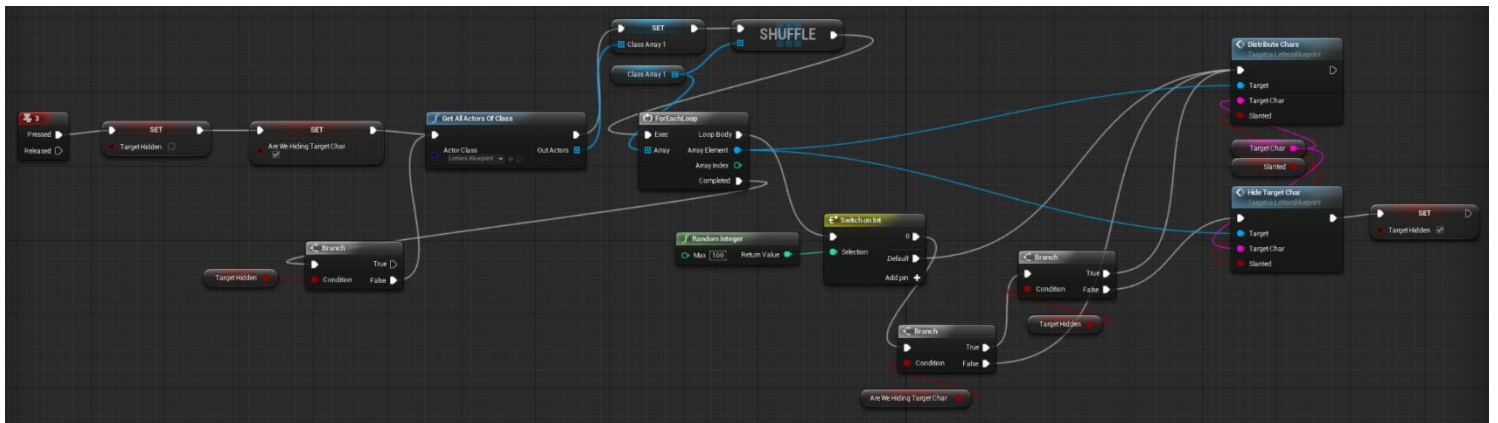
Appendix D: Blueprint Scripting – Randomise Target Character



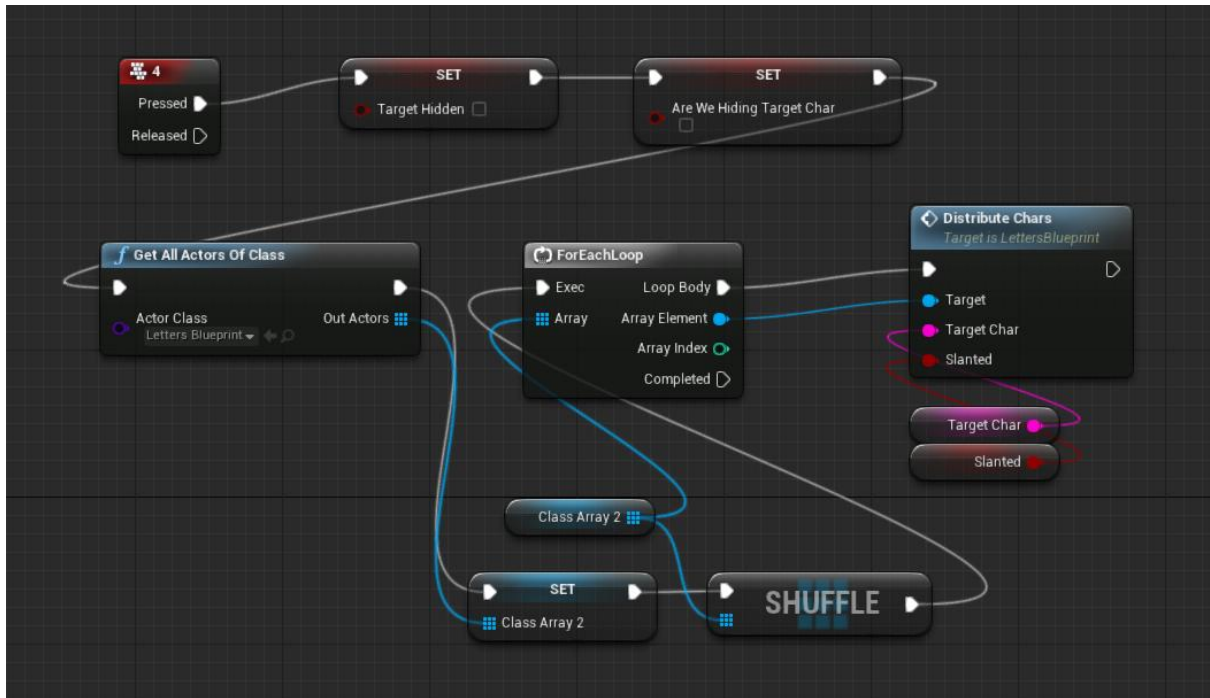
Appendix E: Blueprint Scripting – Set Target Character Above Door



Appendix F: Blueprint Scripting – Hide Target and Randomise all Characters



Appendix G: Blueprint Scripting – Randomise all Characters without Hiding Target



References

- Epic Games, 2015. *'What is Unreal Engine 4?'*. [Online]
Available at: <https://www.unrealengine.com/what-is-unreal-engine-4>
[Accessed 2014].
- Google, 2015. *Google Cardboard*. [Online]
Available at: <https://www.google.com/get/cardboard/>
[Accessed 2015].
- Hezel, P. J. & Veron, H., 1993. Head Mounted Displays for Virtual Reality. *DTIC*.
- Merriam-Webster, n.d. *'Immersion'*. [Online]
Available at: <http://www.merriam-webster.com/dictionary/immersion>
[Accessed 14 October 2014].
- Microsoft, 2015. *Microsoft HoloLens*. [Online]
Available at: <http://www.microsoft.com/microsoft-hololens/en-us>
[Accessed 2015].
- Mitchemmc, 2014. *UE4 First Person VR Template*. [Online]
Available at: <https://github.com/mitchemmc/UE4FirstPersonVRTemplate>
[Accessed December 2014].
- Pausch, R., Proffitt, D. & Williams, G., 1997. Quantifying Immersion in Virtual Reality. *SIGGRAPH*.
- Razer, 2015. *Open Source Virtual Reality Hacker Dev Kit*. [Online]
Available at: <http://www.razerzone.com/osvr-hacker-dev-kit>
[Accessed 2015].
- Robertson, G., Czerwinski, M. & van Dantzich, M., 1997. *Immersion in Desktop Virtual Reality*, s.l.: Microsoft Research.
- Samsung, 2015. *GearVR Innovator Edition*. [Online]
Available at: http://www.samsung.com/global/microsite/gearvr/gearvr_features.html
[Accessed 2015].
- Scientia, 2014. *Online Timetables*. [Online]
Available at: <https://sws.hull.ac.uk/default.aspx>
[Accessed 15 October 2014].
- Sony, 2015. *HMZ-T2*. [Online]
Available at: <http://www.sony.co.uk/electronics/head-mounted-display-products/hmz-t2>
[Accessed 2015].
- Sony, 2015. *Project Morpheus*. [Online]
Available at: <https://www.playstation.com/en-gb/explore/ps4/features/project-morpheus/>
[Accessed 2015].
- VRTifacts, 2010. *Flight Helmet - Redux*. [Online]
Available at: <http://vrtifacts.com/flight-helmet-redux/>
[Accessed 2015].