The Effects of Task Intensity on Attention

A Major Qualifying Project Report

Submitted to the University of

WORCESTER POLYTECHNIC INSTITUTE

In partial fulfillment of the requirements for the

Degree of Bachelor of Science in

Computer Science

By:

Daisuke Abe

Corey Christous

Advisor, Professor Emmanuel Agu

Table of Contents

1. Introduction 5

1.1 Problem Statement 7

1.2 Previous Works 7

The Goal of this MQP 8

2. Background 9

2.1 Visual Attention 9

2.2 Attention Based Rendering 10

3. Requirements 11

3.1 Genre 11

3.2 Why Doom III? 12

3.2.1 The Engine 12

3.3 Experiment Needs 12

4. Design 13

4.1 Map Design 13

4.1.1 The Task 14

4.1.2 Layout 14

4.1.3 Objects 15

4.1.4 Data Collection 15

4.2 Difficulty Divisions 15

4.3 Experiment Logistics 16
4.3.1 Survey 16
4.3.2 Map selection 16
4.3.3 Script 16
4.3.4 Follow up 17
5. Analysis 18
5.1 Time vs. Pictures Seen 18
5.2 Relations to Skill Level 19
6. Conclusions and Problems 23
6.1 Conclusions 23
6.2 Possible Problems with the Experiment 23
7. Future works 24
7.1 Improvements 24
7.2 Additional Ideas 25
8. References 26
Appendix A: Pre-Implemented Objects Used In the Map 27
Appendix B: Images of Famous Figures Used in the Map 28
Appendix C: Pre-experimental Questionnaire 30
Appendix D: Map Blueprint 31
Appendix E: The Pre-Experimental Scripts 32

Easy Map Script 32

Medium/Hard Map Script 32

Appendix F: In Game Screenshots 33

Abstract

Visual attention characterizes what features humans notice most in a graphics scene and why. By exploiting patterns of visual attention, game developers can render scenes more efficiently by focusing more rendering effort on areas of the scene that is noticed most. We have studied what visible features people notice during gameplay while they perform a short-term task. By carrying out experiments on custom maps in Doom 3, we found that increasing task intensity caused a proportional decline in user attention.

Acknowledgements

We would like to thank our advisors Professor Emmanuel Agu and Cliff Lindsay for the advice and support they provided.

1. Introduction

1.1 Problem Statement

Advanced graphics rendering algorithms such as ray-tracing and global illumination have improved the quality of computer generated graphics, but these operations are expensive in terms of time. This is an important factor to consider in the video game industry, where the game can become somewhere between irritating and unplayable if the machine cannot keep up with the rendering. To play games with high system requirements, players need to invest a lot more money than for the games themselves. Many games have settings that the player can modify to improve performance at the price of having less than optimal graphics, but part of the intent of the game designers is to have players appreciate the graphics of the game. Ideally, the game should not require prohibitively expensive (for a casual gamer) machine specs while displaying appreciable graphics.

1.2 Previous Works

There have been attempts to solve this problem through an ingenious method of combining graphics rendering techniques with research results from visual attention. Visual attention is a topic in cognitive psychology dealing with how people notice objects or areas with certain characteristics in a given scene. For example, there have been technologies developed to concentrate global illumination calculations on areas of the scene that are more likely to be noticed and to reduce detail in areas that don't require it [5]. As a result, the images are perceived to be the same quality but are rendered much more quickly. To facilitate the development of these technologies, there has been research done to explore how visual attention in computer generated 3-dimentional environments works. Cater et al. concluded from their research that many people do not notice degradations in the quality of the image unrelated to their task. El-Nasr and Yan conducted an experiment to find what people's attention focused on while they were playing a game. In addition to contributing to more efficient algorithms for graphics rendering, these findings also aid game designers in creating environments such that players will be more engaged and less frustrated by the game. Novice gamers sometimes get lost in 3-dimentional environments, or miss pick-up items that they didn't notice [1]. These studies will aid the improvement of such environments.

The Goal of this MQP

The goal of our project is to explore the relationship between task intensity and visual attention in a computer generated 3-dimentional environment. If we are able to find this relationship, our results could be used for game designers to determine how much and where the detail would go unnoticed when users are under pressure to complete an in-game task. The unnoticed detail would not need to be rendered, saving computation time. Our results might also be used for determining what level of task intensity is appropriate if the game designer wants the player to appreciate certain scenes in the game.

2. Background

2.1 Visual Attention

When a human view a scene, he/she does not notice all objects in the scene at once. The fovea has a visual angle of about 2 degrees, and objects outside of this area are seen in peripheral vision. The eye focuses on one area at a time and jumps between areas in sudden motions called saccades. Studies in visual attention have shown that there are certain rules that govern these motions in a given scene.

There are two widely recognized processes for visual attention, both of which are addressed in attention based rendering. The first is the bottom up, or stimulus driven, process. This process examines inherent properties or features that make objects more salient. For example, the sun in the sky is very noticeable because of its brightness and color contrast against the rest of the sky. Other properties that affect bottom up visual attention include orientation, shape, size, lines, and motion. The other process for visual attention is the top down process. When a subject is pursuing a goal, his/her attention is attracted to objects related to the goal. For example, if the subject were given a task of opening a locked door, he/she would look for objects that look like a key, and these objects would be more noticeable to them. An important fact to note about these two processes is that when visual attention is focused on certain objects, the subject will fail to recognize other features of the scene.

In 1985, Koch and Ullman introduced the idea of a saliency map, a two-dimensional map encoding areas that are likely to attract people’s attention in an environment. However, Marmitt and Duchowski showed in 2002 that such bottom up models were not always completely reliable in predicting where people’s attention would go. In 2002, Cater et al. showed that objects that would attract attention when the viewers were given no task were ignored when they were given a task unrelated to those objects [4].

While there have been a significant amount of research on the top down process of visual attention, to our knowledge, there has not yet been a study on visual attention in relation to the intensity or difficulty of a task in video games. We believe that the intensity of a task will have an effect on the top down process to recognize task-related objects and that this will in turn impact the player's bottom up process to recognize unrelated objects.

2.2 Attention Based Rendering

Attention based rendering is the application of visual attention research to computer graphics. The main concept of attention based rendering is that given a certain computer generated scene, a viewer's attention is more likely to be attracted to certain areas of the scene than others, and the areas that are less likely to receive attention require less computation for the whole scene to appear convincing. This allows the reduction of time and processing power required to render the image while the quality of the image as perceived by a human remains at the same level. For example, in 2001, Haber et al. developed a method that took into account both bottom up and top down processes to determine the salience of objects. This method enforces additional computation of objects that are more likely to be seen. Another example of attention based rendering was developed by Yee and Pattanaik. Their algorithm uses a saliency map for acceleration of global illumination of a pre-rendered scene (2004). The method first creates a rough estimate of the scene, and uses the estimate to figure out the areas that need less computation. Their method takes into account both bottom up saliency and humans' loss of contrast sensitivity when viewing high-spatial frequencies and motion. They discovered that their algorithm was an order of magnitude faster than an algorithm without the saliency map.

3. Requirements

After completion of our background research we were able to begin outlining the goals of our experiment. We needed to design a 3-dimentional environment in which users could move around and observe objects while attending to a task. The subjects could then be questioned on what they saw in the environment. We also needed separate environments of varying task intensity to reach conclusions about task intensity and attention. This required us to use a platform in which we could create and edit maps and introduce new objects into these maps.

In addition to the map, we needed a way to collect data. We needed to record some statistics of each participant's run through the map, such as time. We also needed to find our participants' gaming background and experience, because we suspected a relation between skill levels and attention.

3.1 Genre

The first priority of the experiment was to choose a game genre which would provide us with an environment that addressed graphics rendering in a 3-dimentional environment as well as visual attention. Some choices for genre were action, sports, and first person shooters. We decided on first person shooters for maximum engagement. Our rationale was that if a player had the same perspective as his/her avatar, they would feel more engaged in the game and could concentrate on their task. This genre is also fairly straightforward, and novice players would not be confused with complex tasks. Out of the many games available in this genre, we chose Doom III, a game released in 2004. Specific details outlining our game selection are provided in the following section.

3.2 Why Doom III?

There were several reasons to choose Doom 3 over other platforms. We wanted a platform with a high level of realism and the latest graphical techniques. This would make our study more relevant to new developments in the gaming industry. In addition, it would help people who did not spend much time playing video games feel more immersed in the environment. Another advantage of Doom III is that the game is open source, and we would be able to modify the game for free, as well as study others' modifications. This gave us four levels of modification capabilities: code, scripting, variables, and maps.

3.2.1 The Engine

The name of the engine that was created specifically for Doom III is id Tech 4. This engine is popular for the development of first person shooters. Some other games that use this engine are *Quake 4*, *Enemy Territory: QUAKE Wars*, and *Return to Castle Wolfenstein II*. The main innovation in this version of id Tech 4 is its real-time dynamic lighting [3], meaning that the lighting can change as the game is being played. This produces moving shadows, not implemented previously, which adds realism to the game environment.

3.3 Experiment Needs

There are many tutorials available online on how to make modifications to the game. This facilitated our experiment design using previously unfamiliar software. The tutorials helped us learn to use the map editor to create a map suitable for our experiment. We also learned how to incorporate 3-dimentional objects created with 3D Studio Max into our map.

4. Design

After deciding on our platform, we were ready to begin the actual design of our experiment. The first step of designing a scientific experiment is having a hypothesis, so we came up with these hypotheses:

1. As a user's task becomes more difficult to perform, their top down attention begins to dominate their bottom up attention, and they will notice less objects unrelated to their task even if the objects are salient from a bottom up perspective.

2. Given a task-intensity level, users with more expertise will notice more task-unrelated objects because the task requires less attention for them.

There were three main ideas that needed to be designed for this experiment: the map, the different task-intensity levels, and a questionnaire to gather information. Our first task was the map design.

4.1 Map Design

The first step of map design was to determine what type of map to create. The main part of this was deciding on a task for the participants to carry out, since we needed a map that made sense for the task. The reasons for our task and map designs are explained in the following subsections. After we decided on the basic structure of the map, we modified it to meet our experiment's requirements such as adding task-related and unrelated objects for our subjects to see and data collection mechanisms into the map. These are also detailed in this section.

4.1.1 The Task

We had several ideas for tasks. The two ideas that we ended up in the end with were to collect a number of items or to kill a number of monsters. We decided that a collection task would be a better option because it did not rely too heavily on the participants' skills in first person shooters. We did not want to have novice players be stuck forever trying to hit monsters. An item collection task is straightforward and requires less skill, so all of our participants would be able to understand the task and collect the items in the room in a reasonable amount of time. We implemented the item collection task by placing ten blue balls for the participants to collect throughout the map. The blue balls were salient from the bottom up perspective, because they glowed and were easily noticeable against the relatively dark environment of the room. Of course, they were also salient from the top down perspective, since the participants were instructed to look for them. When a player finished collecting all of the maps, the game ended automatically.

4.1.2 Layout

We decided the layout of the map to be a large square floor surrounded by walls. There were columns evenly distributed throughout the map to block vision so that our participants needed to move around to find the balls. The columns did not have a fixed size or shape. These metrics, in addition to their location, were determined somewhat randomly. This resulted in some columns border the walls or creating dead ends, but navigation through the room was trivial.

4.1.3 Objects

We placed in the map 18 unique objects that were already implemented in the map editor (see Appendix A). These objects looked consistent with the environment. We also added images of five famous figures on the walls (see Appendix B). These pictures had a strong contrast against the rest of the room and were easily noticeable. We used famous figures to be sure that if they caught the attention of participants, they would notice the images. The room was well lit so that the participants could identify all of the objects easily. The objects were placed in such a way that the participant would see most of them when they collected the balls, regardless of what route they took.

4.1.4 Data Collection

As the participants played the map, we wanted to record the time they spent in the map and the route that they took. However, with our resources, we were unable to actually track the route of each participant, so we came up with an alternative method. By recording the order in which the participants picked up the balls, we would be able to guess roughly what route they took. We wrote a script to trigger a timestamp whenever the participant picked up a ball. This script logged which ball was picked up (numbered 1-10) and the time that it was picked up.

4.2 Difficulty Divisions

In order to create different levels of task intensity we divided the map into three different versions. We used the original map as the easy difficulty. We added slow monsters to the map for the medium difficulty. We changed these zombies to faster monsters for our hard difficulty. The monsters were set to be unable to attack, because if our participants were to die, it would make that run of the experiment inconsistent with the other runs. The participants were not informed of this, and they were encouraged to avoid the monsters.

4.3 Experiment Logistics

4.3.1 Survey

Before each participant played the map, they were asked to answer a few questions about themselves including their computer usage and gaming experience (see Appendix C). The purpose of this was to provide us with a basic idea of their background and to provide us data about their experience and expertise.

4.3.2 Map selection

Because participants would inevitably learn from playing the map multiple times, we could only have each participant play one difficulty. We decided to evenly distribute our participants to each difficulty. We were unable to distribute proportional numbers of novices and experts to each map, because it would have been impractical to try to find out everyone's level of expertise before they played the map. We would have had to see everyone twice, once for them to answer the questionnaire and again to play the map. Instead, we decided to assign difficulties to participants in order of participation, beginning the first with easy, second with medium, third with hard, back to easy with the fourth, and so on. Theoretically, this would distribute participants proportionally given a large enough sample.

4.3.3 Script

The participants were read a standard script (see Appendix E) corresponding to the difficulty they were to play on. Every participant was instructed to collect ten blue balls as fast as they could. In addition, participants for the medium and hard maps were told to avoid the monsters due to the fact that they cannot effectively fight them.

4.3.4 Follow up

After each participant finished the map, they were asked to list all of the objects in the room. We used these lists to figure out how many objects each person noticed. It was possible that participants would not write down the names of the famous figures because they did not know them rather than because they didn't notice them, so we also asked each participant to identify all of the figures printed out on a separate piece of paper.

5. Analysis

For each participant, we gathered data of their gaming experience and background, how they thought their gaming skills compared against their community, the order and time that they gathered the balls in while they played our map, and the number of objects and pictures that they noticed in the map. We graphed several categories of our data against each other. Here we present the graphs that we drew our conclusions from.

5.1 Time vs. Pictures Seen

We graphed the time players took in the map (recorded by the logging script) against the number of images they noticed (the number of images they wrote in the post-experimental list of objects). As you can see in Figure 1, the average number of images seen increased as the participants spent more time in the map. This makes sense, because the longer someone spent in the map, the more chances they would have had to recognize things in the map. This phenomenon is also seen in Yee's study, where he noticed that subjects who saw the same animation multiple times noticed an increasing number of objects (2001).

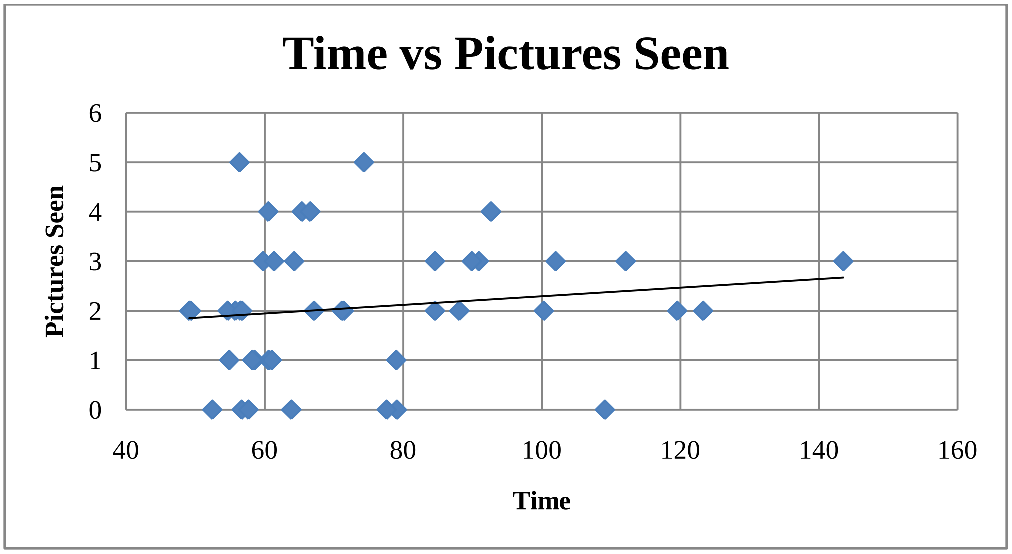

Figure 1: the graph of time vs pictures seen

5.2 Relations to Skill Level

Graphing the difficulty level against the number of images noticed showed us no trend. However, looking at our data, we noticed that the number of experts and novices that played each difficulty level was very unbalanced. Our next step was to check for our second hypothesis. We compared our participants' self-proclaimed skill levels to the number of pictures they noticed and the amount of time they spent in the map. Figure 2 is a graph of the former comparison. We saw that participants who rated themselves higher noticed more pictures. This is in line with our hypothesis that more experienced players would notice more task-unrelated objects than less experienced players. This also leads us to believe that many of our participants rated themselves rather accurately, contrary to our expectations.

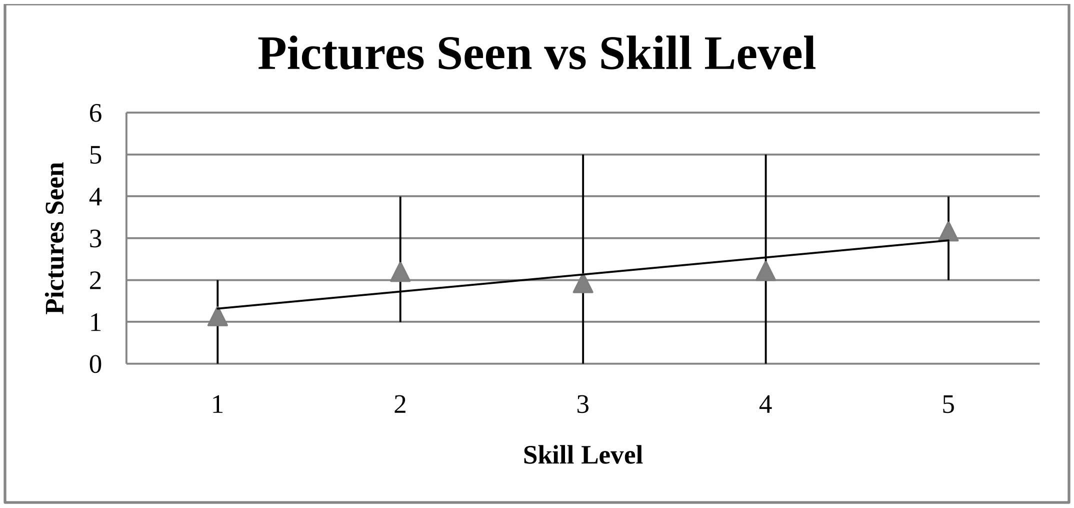

Figure 2: the graph of time vs pictures seen

We took this one step further to confirm our ideas. We graphed the same axes, but we separated the graph into three, one for each difficulty level. We saw in the graphs that the trend was the same for all of the difficulties, further reinforcing our hypothesis (Figures 3-5).

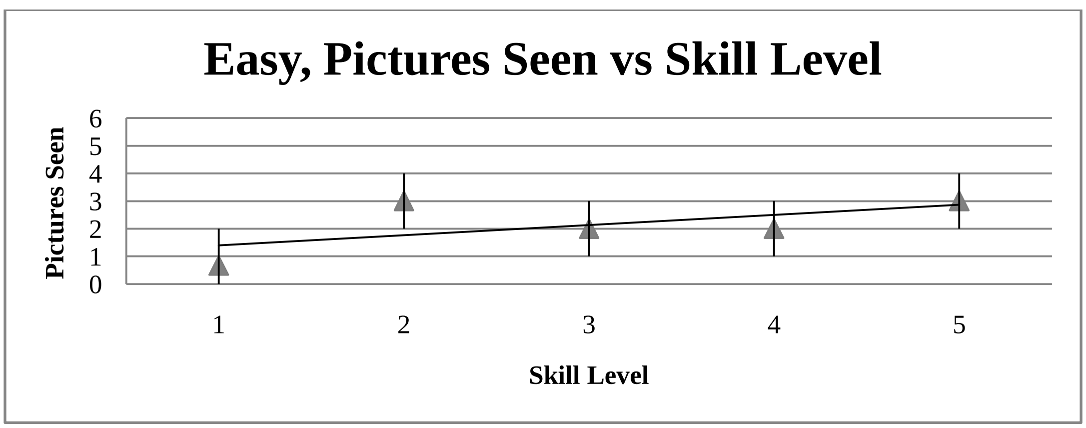

Figure 3: the graph of pictures seen vs skill level on the easy difficulty map

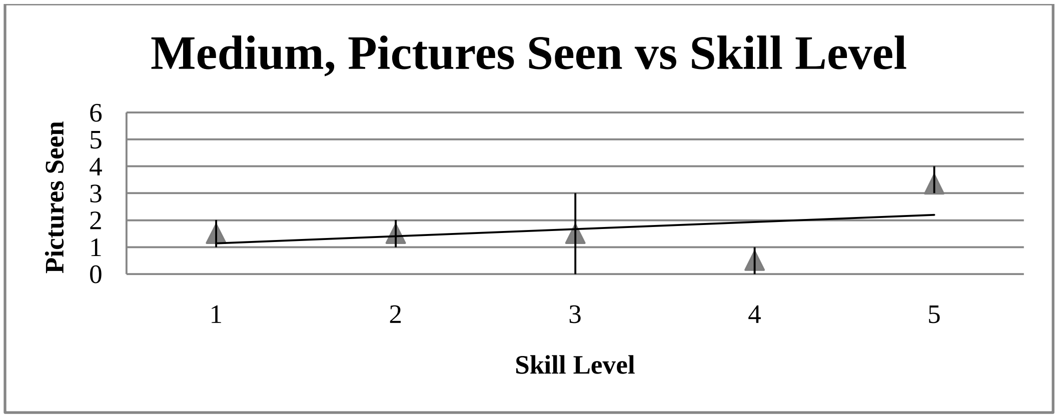

Figure 4: the graph of pictures seen vs skill level on the medium difficulty map

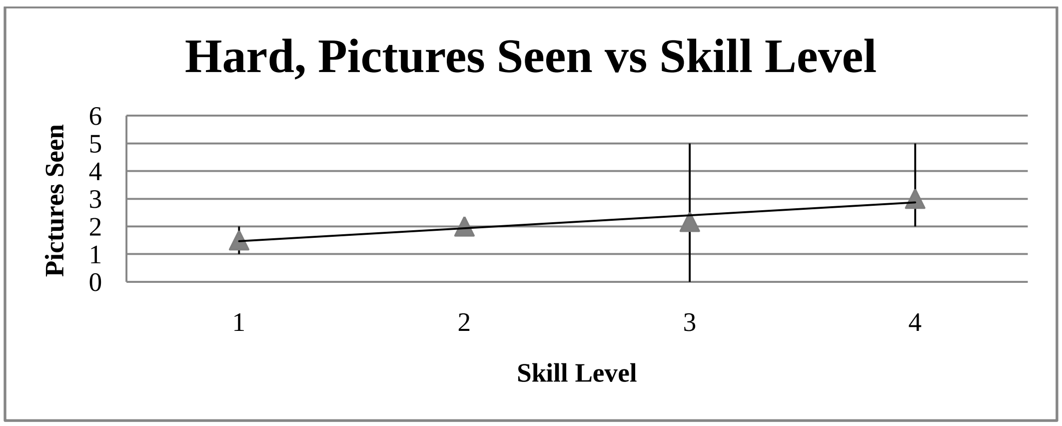

Figure 5: the graph of pictures seen vs skill level on the hard difficulty map

We also graphed the number of pre-implemented objects (listed in Appendix A) noticed against skill level, but found no trend. We suspect one reason for this was that experienced players of first person shooters subconsciously ignored many of these objects. To them, these objects did not mean anything because they knew from their previous experience in similar games that most of these objects do not have significance related to game play. On the other hand, the less experienced participants could have examined the same objects because they did not have the same background with first person shooters and thought every object had some significance.

We also compared the participants' self-ratings against the time they spent in the map. We saw that players who rated themselves higher were able to complete their task faster than players with low self-ratings. We draw from this result that the perceived difficulty of the maps was higher overall for novice players than expert players.

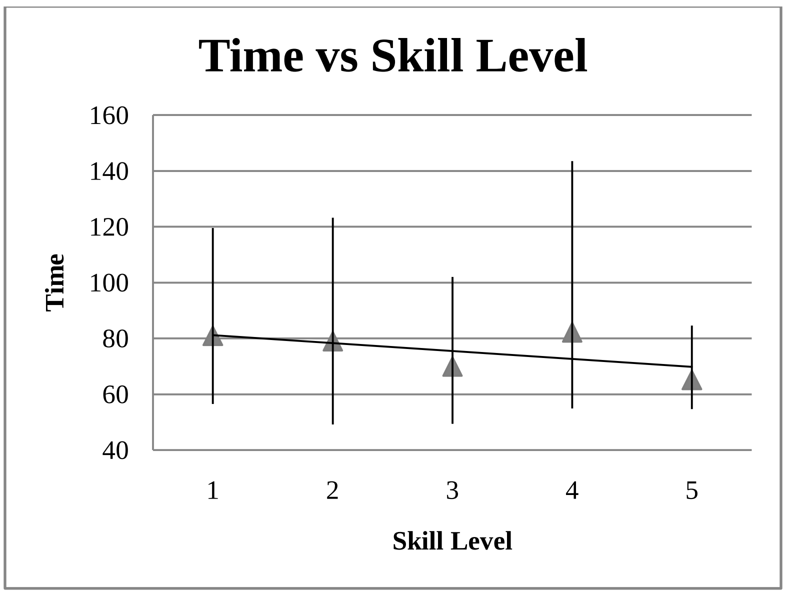

Figure 6: the graph of time vs skill level

6. Conclusions and Problems

6.1 Conclusions

Although we did not have statistically significant results, we were able to draw conclusions from some of the trends that we saw in our data. First, we conclude that a video game player's perceived task intensity given a particular task will vary depending on their level of expertise. We can see this in the graphs of pictures noticed and time to complete the task against each participant's self-rating. Participants with high self-ratings performed better overall (more pictures noticed/less time taken). Although we could not directly show with our data that absolute task intensity had an effect on visual attention, we have reason to believe that perceived task intensity did affect visual attention, confirming our hypotheses.

6.2 Possible Problems with the Experiment

The largest problem we had was the sample size. Because we needed to divide our participants in three different groups for difficulty levels, the sample sizes ended up being small (less than 20) for each of the groups. Because we distributed our participants randomly to the maps, the distribution of experts and novices in each map ended up being uneven.

The question on the questionnaire that asked about hours of games played may have had inaccurate responses for our purposes. The purpose of the question was to use it with the self-rating results to more accurately judge each participant's skill level. Our question did not address the fact that there were some people who played more often before but did not play as much currently. Some participants who rated their skill level at four or five spent less than five hours a week, while others played over twenty hours. We also received comments from people as they answered the survey that they used to play a lot but they do not play as much recently.

7. Future works

After performing this study, we realized there were a few things which could be improved upon. Some of these are a result of limitations of our experiment design, while other improvements could be made by the introduction of additional resources such as eye-tracking hardware. This section discusses such suggestions for future works.

7.1 Improvements

The medium difficulty map produced a lot of data that contradicted our expectations. One of the reasons for this is that the increase in task intensity from easy to medium to hard was not linear. It would be easier to draw conclusions from the results if the medium difficulty were discarded altogether. Eliminating the medium difficulty would also allow for a larger sample size for each of the other two difficulty levels. Another solution would be to devise a system in which the difficulty levels clearly progress linearly. Although this would be trivial in a simpler game, it would be difficult to show a linear progression of difficulty in a 3-dimentional first person shooter.

Although some of the trends seen in our results agreed with our hypothesis, our results were not statistically significant. We predict that if we were to increase the difference in difficultly levels, results would become statistically significant. Eliminating the medium difficulty would help this problem, but actually making the harder task more difficult is also necessary. Another idea is to add a visual timer in the hard difficulty to put more pressure on participants.

It was possible that some participants did not see all of the objects that we assumed they would see on their path. If we were to introduce eye tracking hardware into the experiment as well as saving videos of each run, we could verify that participants actually saw and failed to notice objects.

7.2 Additional Ideas

One of the most obvious suggestions for future work may be to try other game genres. Although the first person shooter genre was the most logical choice for us, results in genres such as action and other third person perspective games could be useful.

We see from our results a correlation between the skill levels of the participants and their attention. A thorough pre-experimental test to accurately rate the skill of each test subject could produce clearer results. For example, the participants could be asked to play a different map and be grouped according to their score. If the skill levels of all of the participants were determined before the actual experiment, a proportional number of high and low skill level participants could be assigned to different difficulties.

Related to the previous suggestion, if only expert level participants were gathered, each of them could be asked to play on multiple difficulties. Because they are all expert level, they would not learn as much from the first run as novices, and the differences in the amount that each participant learns would be small. This way, each individual's run on different difficulties could be compared and the results would say for certain whether task intensity had an effect on their attention.

8. References

[1] El-Nasr, Magy Seif and Su Yan. Visual Attention in 3D Video Games. Hollywood, California, USA: ACE, 2006.

[2] Haber, Jörg, et al. Perceptually Guided Corrective Splatting. Saarbrücken, Germany: Max-Planck-Institut für Informatik, 2001.

[3] id Software. id Software: Technology Licensing. 19 April 2008 <http://www.idsoftware.com/business/idtech4/>.

[4] Sundstedt, V., et al. Visual attention for efficient high-fidelity graphics. New York, NY, USA: ACM, 2005.

[5] Turner, Bryan. "Real-Time Dynamic Level of Detail Terrain Rending with ROAM." 3 4 2000. Gamasutra. 18 4 2008 <http://www.gamasutra.com/features/20000403/turner_01.htm>.

[5] Yee, Hector and Sumanta Pattanaik. "Attention for Computer Graphics Rendering." 2004. 18 4 2008 <graphics.cs.ucf.edu/papers/nba2004.pdf >.

[6] Yee, Hector. Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments. New York, NY: ACM, 2001.

Appendix A: Pre-Implemented Objects Used In the Map

Barrel

Bin

Bottle of pills

Bucket

Can of soda

Chairs

Desks

Fire extinguisher

Hamburger

Lamp

Laptop

Locker

Microscope

Mop

Phone

Traffic cones

Wad of paper

Waste basket

Appendix B: Images of Famous Figures Used in the Map

Homer Simpson

Jackie Chan

Osama bin Laden

Hillary Clinton

Mario

Appendix C: Pre-experimental Questionnaire

Name: ________________________________

Gender: ________

Age: _______

Major: ________________________________

Do you play console or PC games?

Console PC Both

On average, how many hours do you spend using a computer in a week? _______

On average, how many hours of PC games do you play in a week? _______

On average, how many hours of Console games do you play in a week? _______

Rate your gaming level amongst the WPI community.

Poor Very Good

1 2 3 4 5

What is your favorite game genre?

Action First Person Shooter Music Puzzle Real Time Strategy

Role Playing Turn Based Strategy Other _____________

List your favorite game titles. ____

On average, how many hours do you play First Person Shooter games in a week? _______

Are you familiar with WASD key movement? _______

If so, are you proficient with it? _______

Appendix D: Map Blueprint

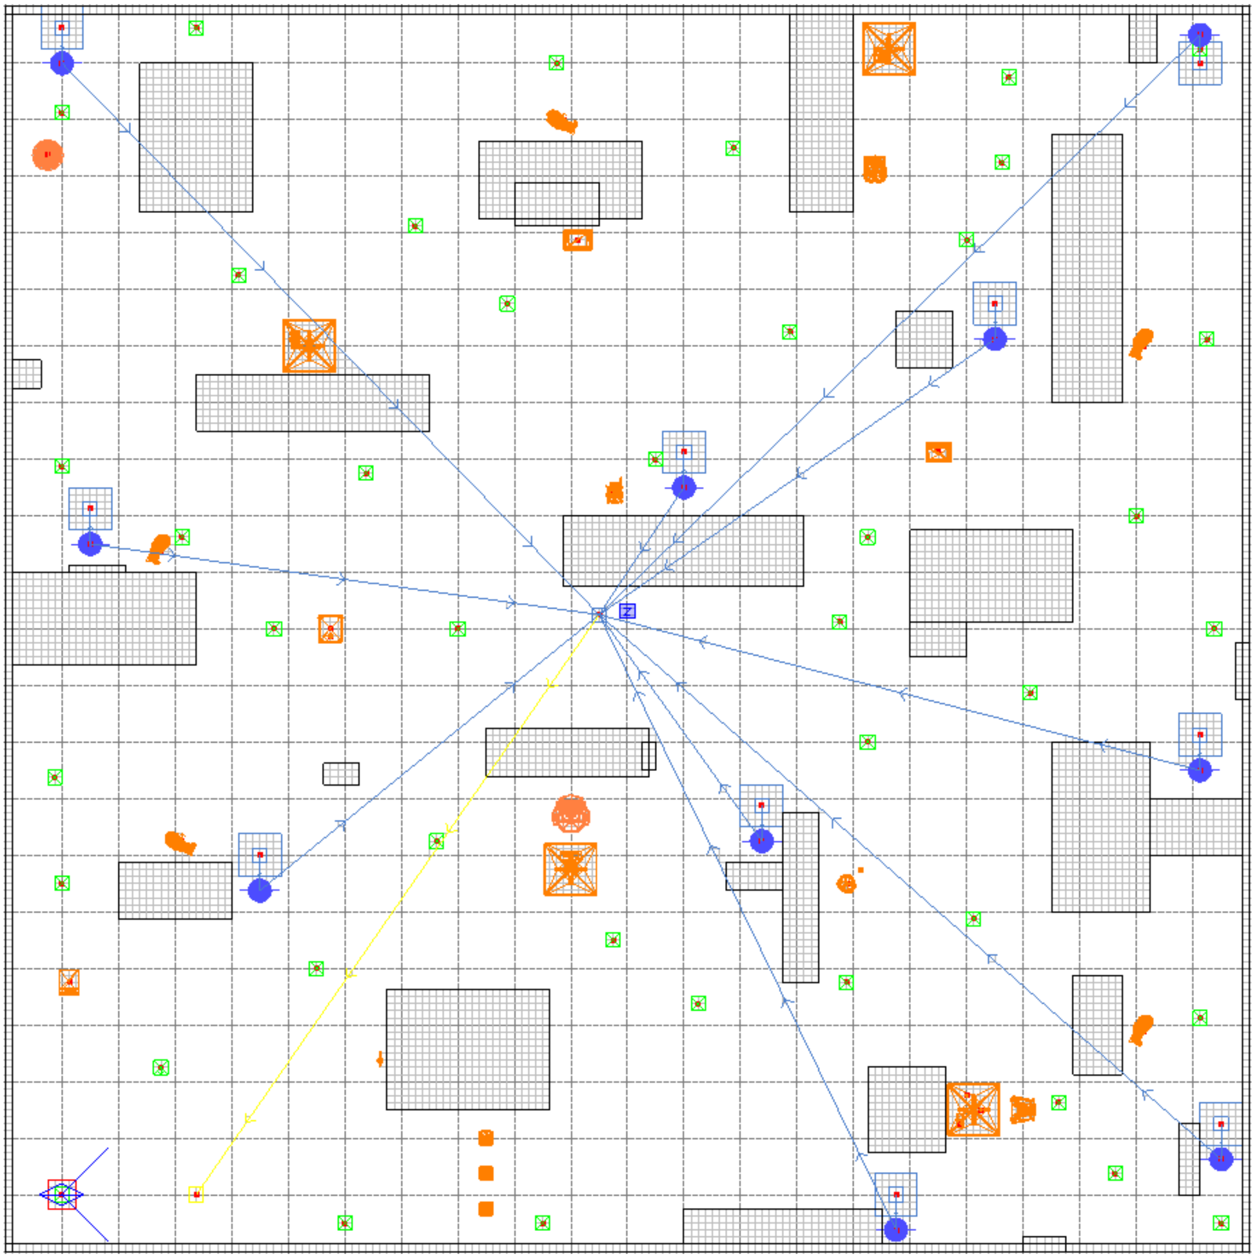

<u>Key</u>

Blue – Orbs

Orange – Objects and Monsters

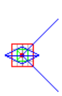 - Starting location

Appendix E: The Pre-Experimental Scripts

Easy Map Script

You are trapped in a room, and you need to collect 50 armor to get out. Armor shards are placed in blue glowing balls, and each one gives you 5 armor. Collect the armor shards in the map as fast as you can. Are you ready?

Medium/Hard Map Script

You are trapped in a room, and you need to collect 50 armor to get out. Armor shards are placed in blue glowing balls, and each one gives you 5 armor. Collect the armor shards in the map as fast as you can. There will be enemies in the room trying to stop you. You can try to fight them, but we're not giving you a gun so that's probably not a good idea. Try not to die. Are you ready?

Appendix F: In Game Screenshots

Screenshot 1: Attention Based Saliency Example, Colors and Illumination

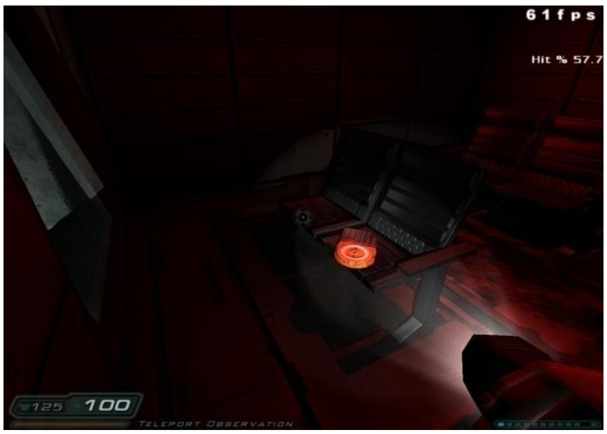

Screenshot 2: Task Based Saliency Example, Game Objective

Screenshot 3: Blue orb containing an armor shard

Screenshot 4: Medium level zombie

Screenshot 5: Hard level zombie