Honours Project Report

School of Rock Art: Browser based cave exploration

Author:
Kaitlyn Crawford

Supervisor:
Dr. Hussein Suleman

	Category	Min	Max	Chosen
1	Requirements Analysis and Design	0	20	15
2	Theoretical Analysis	0	25	0
3	Experiment Design and Execution	0	20	15
4	System Development and Implementation	0	15	10
5	Results Findings and Conclusion	10	20	10
6	Aim Formulation and Background Work	10	15	10
7	Quality of Report Writing and Presentation	10		10
8	Adherence to Project Proposal and Quality of Deliverables	10		10
9	Overall General Project Evaluation	0	10	0
Total		**80**		**80**

University of Cape Town

Department of Computer Science

October 29, 2012

The financial assistance of the National Research Foundation (NRF) towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at, are those of the author and are not necessarily to be attributed to the NRF.

Abstract

The School of Rock Art website was designed to be an educational tool based on rock art in caves in the Western Cape. The website is made up of three components, namely the Guided Tours, StoryTelling and Cave Navigation components. This report describes the development and evaluation of the 3D, Web-based cave navigation component. This component of the website enables users to navigate through 3D virtual representations of the caves, where rock art can be found, and view images of that art. An iterative development process was followed to produce the component using cave models from the Zamani Project at UCT and images of rock art from the Archaeology Department. The results of the usability study carried out on the component show that it is both usable and useful.

Contents

List of Figures

Acknowledgements

The first person I would like to thank is my supervisor, Dr Hussein Suleman. Thank you for the effort and attentive guidance that you gave us at all hours of the day and night. The sacrifice you made to help us is greatly appreciated. I believe you might just be superhuman.

To Professor John Parkington, thank you for your time and assistance during this project. I hope I have produced a system that comes close to your vision. To Professor Heinz Rüther and Roshan Bhurtha, thank you for taking time out of your busy schedules to provide me with the data I needed when I needed it.

I would like to thank my project partners, Joanne Marston and Marco Lawrence. Working with both of you has been an absolute pleasure. Your attitude and perseverance throughout this process has been inspirational. Thank you for your support and friendship - I will always cherish it.

To the honours class of 2012, thank you for an exceptional year. The joy and tears shared between us will be always be close to my heart. Thank you for the late night sessions, the food runs and the coffee. To Michiel Baird specifically, thank you for the effort that you selflessly spent helping me improve my project and keeping me sane through the late nights.

To the National Research Foundation (NRF), thank you for the financial assistance provided towards my university tuition.

Lastly, but certainly not the least important, I would like to thank people from my personal life who contributed to the completion of this work. Thank you to Marc Pittermann and Emma Cecchini for your understanding and support through the late nights and long absences. To my parents, Dawn and Lindsay Crawford, I also say thank you for both the financial and emotional support throughout this year.

Chapter 1

Introduction

1.1 Problem Outline

South Africa has many historical sites containing ancient artefacts and artwork. These sites can often contain vast stores of knowledge and information but in many cases they can also be fragile and must therefore be protected (Pwiti & Ndoro 1999). One manner in which this can be achieved is through stewardship, where local residents are educated about nearby sites and how to protect them without preventing access and the opportunity to learn from them (Pwiti & Ndoro 1999). They are made stewards of the sites and the information about heritage stored there.

The Archaeology Department of the University of Cape Town (UCT) has been collecting data and information on many of these historical sites. Therefore, they have a large repository of images from the sites that is currently used for little more than cataloguing. Additionally, the Geomatics Department of UCT is in the process of recording laser scans of many of these heritage sites (Rüther, Chazan, Schroeder, Neeser, Held, Walker, Matmon & Horwitz 2009) and would like for this data to be put to use. Understandably, both departments felt that these resources could be used to build a system to educate potential stewards and the general public about heritage and the importance of heritage sites.

1.2 Proposed Solution

For this reason, the Archaeology Department requested the development of an interactive tool that could be used to educate the general public about the heritage sites in and around the Western Cape. A specific requirement of this system is that it made use of the electronic laser scans of some of these sites, provided by the Geomatics Department.

A solution was proposed to provide the two departments with an interactive, education-focused website that makes use of the data. The system was chosen to be a website instead of a device-specific application as this allowed for a larger audience. It also meant that the application would be accessible to users who lacked the storage capacity to save the data locally. The proposed website is called The School of Rock Art in order to highlight the educational aspect of the system.

The website designed and developed combines the three components outlined below.

Guided Tours: Developed by Marco Lawrence, this component enables users to follow predefined paths through images related to specific topics. It also allows users to create these paths for others.

Story Telling: Developed by Joanne Marston, this component enables users to read traditional stories that have been translated into English. Users can also design the appearance of the storybook in which the stories appear and add relevant images to the book.

Cave Navigation: Developed by Kaitlyn Crawford, this component allows users to explore some of the heritage sites in 3D and view images related to each specific site. This component is described and discussed in this report.

Each of these components made use of the data provided by the Archaeology Department to provide an educational website that could be used to educate users about the importance of heritage. Additionally, the Cave Navigation component also made use of the laser scans from the Geomatics Department to achieve this goal. This report outlines the design, implementation and evaluation of the Cave Navigation component of the website. This design focused on the following research question:

1.2.1 Research Question

Is it possible to build a usable and useful 3D cave navigation system to encourage learning about rock art?

1.2.2 System Outline

In order to answer the research question satisfactorily, the system developed needs to effectively allow users to navigate through a cave and view images of the art found there. This system would present the cave environment to a user in a browser using the models built from the laser scans. Artwork on the cave walls would be made identifiable to users and they would be able to view images of that art from the database.

1.2.3 Methodology

An iterative design process was followed to develop the cave navigation component. Each iteration was evaluated and the findings of the evaluation used to inform the next iteration.

First Iteration: The first iteration was used to determine the feasibility of the concept and design based on requirements set by the Archaeology Department. Evaluation on this iteration was carried out through a short evaluation by the second reader, Mr Stewart. The goal of this evaluation was to determine if the problem could be solved by the proposed solution.

Second Iteration: Feedback from the first iteration, as well as feedback from a focus group of users, was incorporated into the design and implementation of the second iteration. This iteration was intended to implement the core functionality of the component. It was evaluated by a small set of expert users in the form of a heuristic evaluation. Feedback gained in this evaluation was used to expose the most critical usability problems in the interface.

Final Iteration: The final iteration was designed based on feedback from the second round of evaluation. The goal of this iteration was to improve on the usability of the system. Evaluation of this iteration was conducted through formal user testing to determine its usability. The feedback gained from users was used to determine if the system successfully answered the research question posed.

1.3 Report outline

Chapter 2 of this report discusses previous work done in the field of visualising large information spaces similar to the database of images provided by the Archaeology Department. The design chapter, detailing the iteration cycles involved in the development of the cave navigation component, follows on from there. The evaluation chapter describes the evaluation process followed to evaluate the final iteration of the component and presents the findings of this evaluation. The report ends with a conclusion chapter that summarises the report and presents possible future work.

Chapter 2

Background

Education of residents of the community close to heritage sites is the key to stewardship programs, and digital libraries containing heritage and archaeological information that could be used for this task do exist (Rüther 2007). Unfortunately, many are not easily accessible to those who have little to no research experience (Isaacman, Lalu & Nygren 2005). In order to correct this, these libraries must be presented to users in a way that provides intuitive navigation of the information space. Therefore, visualisation techniques that could be applied to digital libraries must enable this.

Additionally, these libraries often contain information stored in a variety of different formats (Baldonado, Chang, Gravano & Paepcke 1997) or draw in information from multiple different sources (Rao, Pedersen, Hearst, Mackinlay, Card, Masinter, Halvorsen & Robertson 1995). For this reason, potential visualisation techniques for large information spaces must also be able to handle multiple sources or formats in order to be considered suitable for application to digital libraries. This chapter presents past work done in the presentation of archaeological information to the public and then separates visualisation techniques into those that display the data in a 2D format and those that display it in a 3D format. These visualisation formats are then analysed with respect to navigational ability and the diversity of information formats handled.

2.1 Archaeological Information

The traditional method of presenting heritage information to the public is to construct rock art centres and museums (Mazel 2008). These sites have the distinct disadvantage of not being accessible from other geographical locations. Additionally, large amounts of the artwork may be susceptible to damage from people or nature and are, therefore, not open to the public (Dorn, Whitley, Cerveny, Gordon, Allen & Gutbrod 2008) (Daz-Andreu, Brooke, Rainsbury & Rosser 2006).

For this reason, a number of initiatives have began to digitise heritage sites and archaeological data (Isaacman et al. 2005) (Bryan & Heritage 2009) (Gonzlez-Aguilera, Muoz-Nieto, Gmez-Lahoz, Herrero-Pascual & Gutierrez-Alonso 2009). This will enable research to be done on the digital copies of the artifacts, without endangering them, while making them accessible to researchers all over the world. Unfortunately, not much information can be found with regards to making this information easily accessible and usable to the public. With this aim in mind, a number of methods can be considered to display the digital information kept in large archaeological databases in a way that allows non-technical users to navigate through it.

2.2 2D Visualisations

2.2.1 Information Murals

Techniques for visualising and navigating large information spaces are not uncommon. Information Murals are one such technique. This technique focuses on condensing large amounts of information into a global 2D visualisation that is able to fit easily onto a computer screen (Jerding & Stasko 1998). The visualisations preserve information on the context of the data being viewed in a way that supports analytical and navigational tasks that a user may want to perform (Wan 2006). The goal of the software is to represent a large information space within the confines of a computer screen without losing information due to compression, while still allowing inspection of details without losing contextual information.

2.2.2 Macroscoping

Macroscoping tools employ a variation on the pan-and-zoom technique of visualising large information spaces that are organised or can be visualised hierarchically (Lieberman 1994). This technique uses transparency to preserve a visual representation of the original context from which the current detailed view comes (Harrison, Ishii, Vicente & Buxton 1995). The main goal of macroscoping is to allow the user to zoom-in to get details, while still being able to see the context that they occur in. This is achieved by changing the transparency of the original view and overlaying it on the zoomed-in view. This technique could be useful for navigating the GIS visualisations and maps contained in digital libraries (Wan 2006).

2.2.3 ActiveGraph

ActiveGraph uses scatter-plot graphs to depict datasets of digital library documents. The service effectively handles user queries by filtering objects using predefined attributes and metadata. It is most notably useful for researchers who wish to query citation data (Wan 2006). Built to enable collaborative projects, ActiveGraph also allows users to contribute to the digital library by editing the metadata of objects and storing the history of who edited the data and when (Marks, Hussell, McMahon & Luce 2005).

2.2.4 UC

The UC system for visualising and navigating digital libraries uses Treemap layouts to present collections of documents within the library (Good, Popat, Janssen & Bier 2005). Continuous and Quantum treemaps are used to provide an overview of documents in a set and then facilitate navigation amongst these documents. The system specialises in allowing the user to interact with the document instead of the tool and enables manipulation of documents as well as the ability to compare them. While it is designed for the visualisation of personal libraries (where the user has the right to use the data as well as local possession of it), it can be used for more public libraries (Wan 2006).

2.2.5 Context Preservation

To provide easier and more intuitive navigation, Information Murals and Macroscoping both preserve a visualisation of the context that details come from (Jerding & Stasko 1998), (Lieberman 1994). However, the transparent images used by Macroscoping to display this information could be misleading and distracting, obscuring the true information and making navigation difficult (Harrison et al. 1995). This idea of a preserved overview or global representation is not applied in the UC or ActiveGraph techniques and this could hamper a user's navigation of the information (Jerding & Stasko 1998). Interfaces designed to show subsets or sections of information amongst larger sets of data or in large spaces, where relation might be important, should strive to provide detailed views of the subsets without losing the contextual information.

2.2.6 Type Independence

A drawback of both the Information Murals technique and the UC system is that only one type of information in the space can be viewed at any time (Jerding & Stasko 1998) (Good et al. 2005). The Macroscoping technique suffers from the need for the information to be hierarchically connected or organised. These problems make these techniques inapplicable to visualisations of digital libraries that can contain many objects stored in different formats and related in different ways (Baldonado et al. 1997). In contrast, a dataset displayed by the ActiveGraph system can correspond to any objects in the library regardless of the format or medium in which they are stored and related (Marks et al. 2005).

2.3 3D Visualisations

2.3.1 LVis

LVis (Digital Library Visualizer) is a visualisation tool that extracts semantic relationships from data in a library and then uses a Boltzman algorithm to lay the data out in space (Borner, Dillon & Dolinsky 2000). The system has both a 2D and a 3D user interface, with the 3D interface making use of the CAVE (CAVE Automatic Virtual Environment) virtual reality tool to create an immersive environment through which to display the information. Users of the 3D interface enter a virtual reality version of Easter Island and can choose to walk through gates that separate the information categories into rooms where the information is displayed on objects modelled in the space (Wan 2006).

2.3.2 The 3D Vase Museum

3D spacial metaphors are used in other digital library visualisation tools as well, such as the 3D Vase Museum developed at Tufts University that allows a user to change their view of the information as they navigate through it (Shiaw, Jacob & Crane 2004)

Figure 2.1: A view of the 3D Vase Museum

(Wan 2006). At a high level, the information is displayed as a birds-eye view of a museum room; this forms a visualisation similar to a scatter-plot with the walls as axis against which the vases are plotted. As users move closer to the objects, the view changes to be a perspective, eye-level view that allows examination of the physical details of the vases as can be seen in Figure 2.1. As the user gets even closer, the view focuses on a vase and brings up metadata on it without the view leaving the room (Shiaw et al. 2004).

2.3.3 Augmented Reality

Projects such as LVis and the 3D Vase Museum raise the question of the use of virtual reality in information visualisation techniques. Guven & Feiner (2006) draw from projects such as FlyAbout (Kimber, Foote & Lertsithichai 2001) and Movie-Maps (Lippman 1980) to build a tool that uses augmented reality to display information. FlyAbout provides the spacial navigation through video, while Movie-Maps provides overlays of information on these graphics. Guven & Feiner (2006) focus on visualising information in its actual location and context by overlaying it on pictures or videos taken from a camera. A Virtual Field of View is used to ensure the scale and positioning is correct in relation to the information from the camera as it moves.

2.3.4 Context Preservation

The LVis system is an example of mapping a user's navigation of a 3D structure to their navigation of the information contained in the digital library being navigated (Borner et al. 2000). However, unlike the 3D Vase Museum and some of the 2D techniques described, it does not maintain a concept of the context that details come from (Shiaw et al. 2004). Guven & Feiner (2006) preserve context very well but suffer in that users must navigate to a different physical space in order to view information from the new space. However, it does allow the user to navigate the information at their own pace and leisure. Due to the addition of a dimension, context preservation in a 3D space is more complex than in 2D space. The addition of the third dimension provides more space to display contextual information (Shiaw et al. 2004), but it can also allow user to get lost in the space (Kimber et al. 2001).

2.3.5 Type Independence

LVis again falls short when compared to the 3D Vase Museum since only information stored in image format is displayed (Borner et al. 2000), whereas the 3D Vase technique can be extended to incorporate other sources of information (Wan 2006). Guven & Feiner (2006) benifit from the Movie-Map technique as it can display any information that can be viewed on a camera.

2.4 Summary

Digital libraries can be used to prepare potential heritage stewards and give citizens detailed information about their local heritage. Unfortunately, most digital libraries were built for the purpose of research (Isaacman et al. 2005) and users who do not have the research skills may find it difficult to navigate the large information spaces of these libraries and will therefore be unable to make full use of the resources contained in them.

With this in mind, this chapter looked at methods of visualising large information spaces, such as digital libraries, in both 2D and 3D environments before analysing the techniques with respect to navigation and variety of information that they can handle. An important factor noted throughout the research was the significance of allowing users to examine details without losing sight of the context from which the information is drawn.

Systems such as ActiveGraph and the 3D Vase Museum were found to be strong with respect to navigation and compatibility with multiple information formats or sources. Therefore, these systems can be drawn from in the design and implementation of the School of Rock Art cave navigation system to ensure that detailed, lower level information can be retrieved from the system without losing the context that it comes from.

Chapter 3

Design and Implementation

An iterative design process was followed for the development of the cave exploration section of the School of Rock Art system. This chapter discusses the considerations for the design and then goes on the describe the design, implementation and evaluation phases of each of the three iterations completed. Feedback gained from each evaluation step was used to inform the design of the following iteration, as indicated in Figure 3.1. The evaluation phase of the final iteration is discussed in the evaluation chapter of this report.

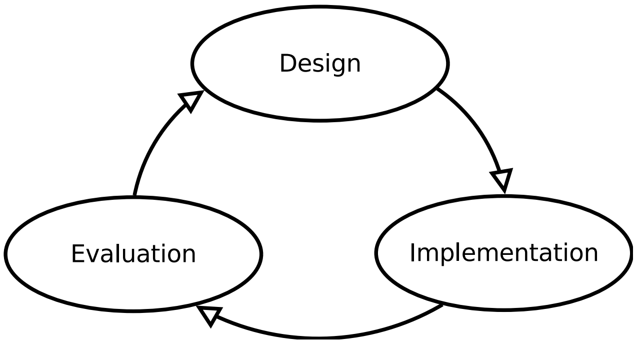

Figure 3.1: Iteration Cycle

3.1 Design Considerations

Requests made by the Archaeology Department imposed specific considerations on the system. These requests and the resulting consequences for the system are described below.

Cave Scans

The main specification for the interface was that it should make use of cave scans developed by the Geomatics Department. One such model can be seen in Figure 3.2. Considering this constraint and the photographs and information that the Archaeology department was willing to supply, it was decided that the system described in this report would suit the requirements.

Due to the nature of the laser scans, the models contain some holes where the data was not accurately recorded. An open project with the Computer Science Department focuses on developing an algorithm to fill these holes, so this was considered to be outside of the scope for this project. Therefore the models presented to users have some holes present. It was determined that maintaining the accuracy of the models was preferable to filling these holes.

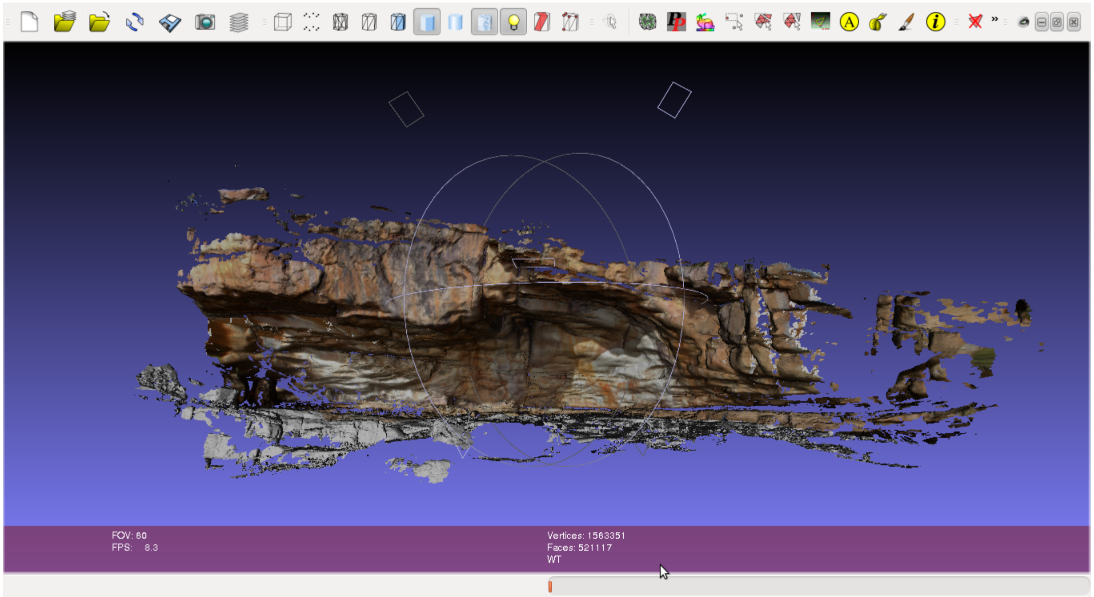

Figure 3.2: Model of a cave displayed in Meshlab

The formatting of the files had other consequences for the design of the component. The files are formatted according to the Stanford Triangle Format or PLY (This format describes 3D models in a way that is both flexible and portable, since the binary version available makes it platform independent) (McHenry & Bajcsy 2008). This had to be kept in mind when deciding on the tools and frameworks to use. Since the cave model was already textured, it was not necessary to be concerned with texturing it.

Available Resources

Since the main aim of the system was to educate users on rock art, it was determined that additional resources would be required to give information and to make the system novel and interesting. The Archaeology Department was willing to provide the necessary information or resources from existing data.

Due to the visual nature of the project, and the fact that the rock art in the caves is not clearly visible on the walls of the model (even when textures have been added), it was determined that photographs of the art would be an appropriate addition to the system. The photographs available from the Archaeology Department were in various formats but, since the most common image formats are supported by most software, it was felt that this was unlikely to cause a problem.

The images are archived in a file system store structured according to location. They are first grouped according to the map code of the 1:50000 map that they can be found on. The images in these folders are then subdivided according to the name of the property that the cave is found on. Each property may have a number of sites on it, so the images are then grouped in folders according to the site code. Finally, since any number of excursions may have been made to collect the photographs, they are grouped according to the excursion that they were taken on.

At the level where images are grouped according to site, an XML file specifies metadata for that site and what artwork can be found there. This file records information such as the dimensions of the site, its local name, the types of rock found there and the number of artworks of each animal that can be found there.

Accessibility

It was determined that a wide audience would be interested in the system. This audience would include, but not be limited to, students, school learners and teachers, researchers in a number of fields and the general public. For this reason, it was determined that the system must be highly accessible.

To ensure this accessibility, a Web based approach was chosen as a requirement for this system. This decision had to be kept in mind when determining what tools and frameworks to use and it influenced all choices regarding design and implementation.

Using this information, a system was envisioned where a user could explore a cave environment and examine the rock art found there. The models provided would be used to display the virtual environment in a Web browser. Users would be able to explore the cave itself or to focus on artwork and more closely examine this through different photographs of that item of art.

3.2 First Iteration (Feasibility Demonstration)

The main goal of this initial iteration was to show the feasibility of the concept. Focus was placed on rendering a 3D environment in a Web browser and allowing a user to navigate in this environment. Additionally, it was shown that images could be displayed over the rendering of the environment. At this time, a number of tools and frameworks were considered as options for implementation.

3.2.1 Design

The design of the initial prototype was partially informed by the constraints and comments from the Department of Archaeology at the University of Cape Town, as described above. Further inspiration for the design was taken from other software applications that model 3D environments. The most common of these applications are variations on *Computer Assisted Drawing* (CAD) systems or *First Person Shooter* (FPS) games. FPS systems use a perspective view to show a 3D environment in an attempt to make it as lifelike as possible. CAD systems use orthographic views to display the detail and measurements required for modelling 3D. The decision was made to develop an interface similar to that of FPS games since CAD systems require a greater amount of training and FPS systems are more intuitive. The features designed at this stage of development are discussed below.

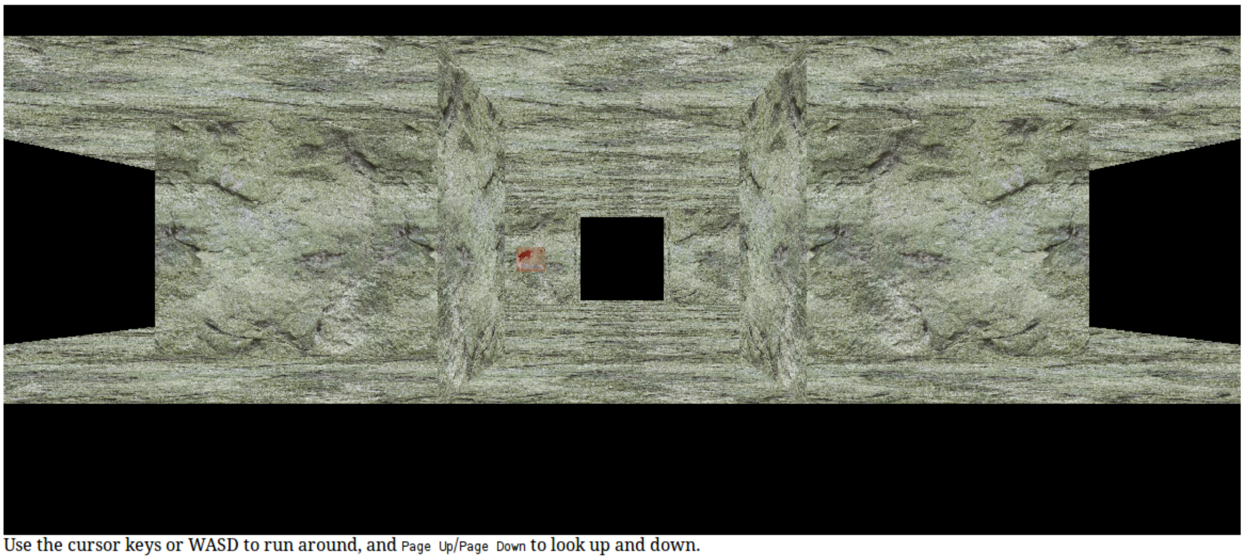

Figure 3.3: Basic environment for the Feasibility Demonstration

3D Environment

At this stage a simple 3D environment was built to be displayed in the browser. As can be seen in Figure 3.3, the environment was very basic, consisting of a single room with a roof, floor and four walls, each with an entrance. The planes forming this environment were all textured with a rock texture to give the feel of a cave. An image of rock art is affixed to one of the walls in the environment. This is to simulate a user seeing an article of art painted on the wall.

Intuitive Controls

The controls to navigate the camera around the environment should be intuitive and easy to use. For this iteration, users were able to navigate through the environment on a fixed plane. The navigation used the WASD keys or the arrow keys to move about the environment, as is the case for most FPS systems. The difference is that, while the W and S keys (or up and down arrows respectively) are used to move backwards and forwards as usual, the A and D keys (or left and right arrows respectively) are used to rotate the camera to move the view left or right. This key mapping is shown in Figure 3.4. Most FPS systems use the movement of the mouse to control the rotations of the camera and the A and D keys to move left and right in the environment.

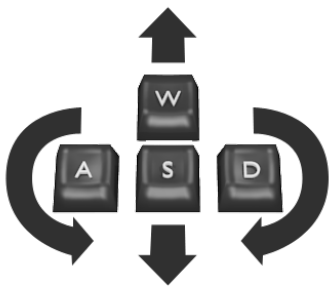

Figure 3.4: Mapping of keys to camera movements in the first iteration

Photograph pop-ups

In order for users to be able to explore the environment itself, the photographs that allow for closer inspection of the artwork should not obscure the cave itself. For this reason it was determined that the images should be displayed in an overlay that was rendered over the view of the environment when users' focus was on the relevant item of artwork.

This iteration made use of a user's distance from the artwork (indicated by an image fixed there) to determine if it was their focus. This meant that if they were close enough to the artwork then the image would "pop-up" and be displayed on a larger scale in an overlay in front of the environment. An example of this can be seen in Figure 3.5.

Figure 3.5: Overlay of artwork in the Feasibility Demonstration

3.2.2 Implementation

The first implementation set out to prove the feasibility of the concept. This involved creating a navigable 3D environment in a web browser. As described above, the features to implement were the 3D environment, photograph pop-ups and intuitive controls. The execution path of the system can be seen in Figure 3.6.

Basic HTML5 and JavaScript was used for this implementation since it supports hardware accelerated rendering of 3D graphics through WebGL. WebGL is an application programming interface (API), written in JavaScript, that enables the rendering of graphics in a web browser without the use of plug-ins. However, due to the varied browser support for HTML5, specifically WebGL, and the scope of the project, it was determined that only support for one browser would be implemented. Google Chrome was chosen as the main browser for development since it is the forerunner for HTML5 adoption.

3D Environment

An HTML5 *canvas* element is used to display the environment to the user. JavaScript is used to create and initialise the scene to be shown and handle all events, animations

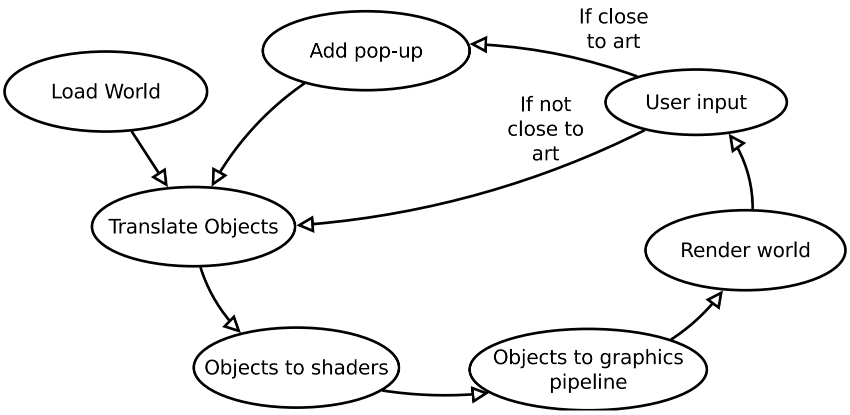

Figure 3.6: Execution path of the prototype

and rendering. First the dimensions of the canvas are set, then the simple vertex and fragment shaders (written in C++) are initialised. Following that, the textures and world are loaded. Lastly, the animation step that handles input and renders the environment is called.

Since only a basic 3D environment was needed for this iteration, a simple four-walled room was used to show the feasibility of the concept. The environment is stored in a text file that specifies the positional and texture coordinates of the corners of the planes that are used to make up the walls and floors of the room. In order to load the world, an asynchronous *GET* request is sent to the server to retrieve the contents of the file. When this request has successfully completed, the information is parsed and stored for later use.

On every animation step, when the scene is drawn, the rotations for the camera are computed and applied (these are dependent on the controls and are discussed in the section on controls below). The positional information for each plane in the environment is then computed and sent to the shaders, along with the corresponding texture information. The shaders determine what color each vertex in the environment should be and pass that information to the graphics pipeline, where the processing happens and the scene is rendered.

Intuitive Controls

As previously stated, the WASD keys control the view seen through the camera. An event handler is assigned to the HTML document to deal with all keyboard events. When a key is pressed the event is called and the action corresponding to that key is executed. In the case of the W and S keys, the action will set the speed of the camera to be either positive or negative (positive for W and negative for S), while in the case of A and D the action will set the change in the rotation about the y axis to be either positive or negative (positive for A and negative for D).

When the scene is drawn, the speed and rotation angle are used to calculate the position and rotation of the camera in the coordinates of the scene. Matrix multiplication is then used to translate and rotate the camera to show the correct view. The graphics pipeline then uses the view to render the scene.

Photograph pop-ups

The photograph pop-up was also created using JavaScript in the *canvas* element. When the page is loaded and the environment is created, another plane is created to hold the photograph that will be displayed to the user. On every render step, the distance from the camera to the object is calculated. If this distance is below a certain threshold then the plane for the pop-up is rendered, otherwise it is not. The transformation matrix used to change the view is not applied to this object; instead a different matrix is used to translate the plane's position to be directly in front of the camera. This means that it is always rendered in the same position relative to the camera.

3.2.3 Evaluation

Since the first iteration of the system had a small level of functionality and was intended to prove the feasibility of the concept, the evaluation carried out on it was very short. It was demonstrated to the second reader (Mr. Stewart), who provided the feedback below.

Jarring pop-up effect

It was noted that the appearance of the image when a user is close to the item of rock art is very jarring. This is due to the fact that there is no warning or indication that this is the effect walking closer to the image will produce. Some warning or indicator should be provided. Possible solutions are to make the image pop-up a user initiated event or to make the image turn to face the user as they move in the environment and grow larger as a function of the user's distance from the artwork

Not clear that the artwork is different

It was felt that the images used might be difficult to identify against the backdrop of the cave. This is especially true in the real cave models as the photographs of the artwork are similar to the photographs used to texture the models. For this reason, the artwork should be displayed on the wall in a manner that clearly indicates that it is different. Possible solutions are to highlight the artwork or to direct the user towards it with arrows.

Audio narration

The suggestion was made that audio narration should be provided for the environment. This narration would be tied to the cave and the artwork to provide functionality similar to a tour guide. This was considered to be beyond the scope of the project based on the availability of applicable content. However, it was deemed to be a feature that would very positively impact on the experience of the user and is considered as an important inclusion in future work.

3.3 Second Iteration (Functional Prototype)

The main goal of the second iteration of development was to implement the bulk of the functionality for the system. It was also necessary to consider the changes deemed necessary during the evaluation of the previous iteration.

3.3.1 Design

The design of this iteration of the system was informed by the findings from the evaluation of the previous iteration and a small focus group of users. The focus group was given screen shots of the previous iteration and asked to provide feedback on the design. Three students in their first year of study at the University of Cape Town volunteered to take part in the focus group. This section describes the changes made to the design of the system based on these two sets of information.

Environment using the cave models

Both the previous evaluation phase and the focus group highlighted the need to make use of the cave models provided, instead of the basic room, as soon as possible. Using the caves would provide more insight into the problem and help identify changes that might need to be made to the system in order to incorporate them. Similar to the previous prototype, the cave would be displayed on a black background, as it was felt that other colours might detract from or obscure the cave.

FPS Controls

Users felt that the navigation controls for this iteration should be made to more closely resemble those of a *First Person Shooter* (FPS) system. These controls are considered to be more intuitive and since many users are still familiar with them, they will require less training. The actions performed by the A and D keys and the usefulness of the mouse were to be changed.

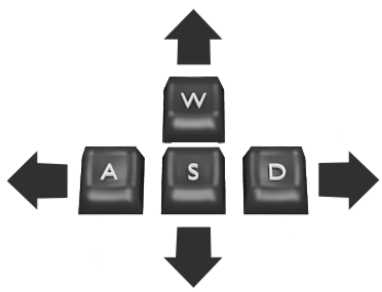

Figure 3.7: Mapping of keys to camera movements in the second iteration

As can be seen in Figure 3.7, the A and D keys move the camera left and right respectively. The rotation of the camera is controlled by the movement of the mouse and the view is rotated only when the mouse moves. Moving the mouse left will pan left and moving the mouse right will pan to the right, while moving the mouse up on the screen will pan the camera up and moving the mouse down will pan the camera down.

Disappearing Artwork Tags

Since items of rock art or other interesting aspects of the cave might not always be visible on the model, it is advisable to provide a way for users to find the art without obscuring the cave itself. The solution to this was to add tags to the environment that indicate where a user can find something that may be of interest to them. These tags are intended to give the appearance of frames around the artwork on the cave wall, as can be seen in Figure 3.8.

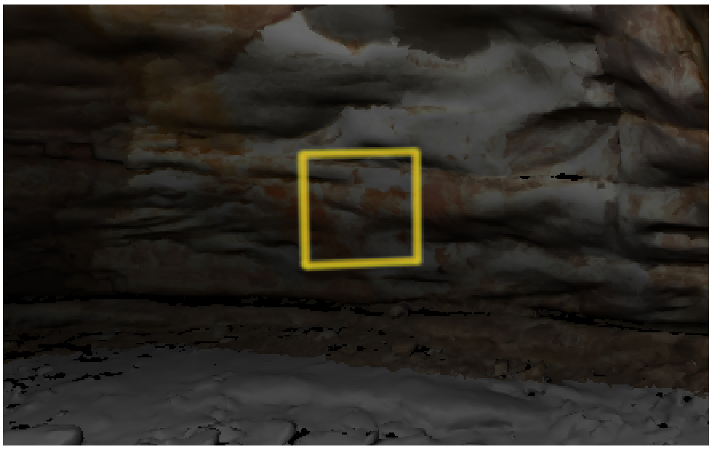

Figure 3.8: Tag indicating an item of rock art in the cave

In order to prevent these tags from obscuring the cave, they are invisible and only become visible when a user is focused in that area. Instead of determining a user's focus by their distance from the artwork, as was done in the previous iteration, it was decided that a user would be considered focused on an object if the mouse was over the object. This means that the tags only become visible when a user mouses over them. Users then click on the tags to view photographs of the cave.

Lightbox Overlay

One of the main concerns raised during the evaluation of the previous iteration was the suddenness with which the photographs of the artwork were displayed. A more gradual transition would be necessary for this iteration. Research suggested that a number of tools had been made in JavaScript to display an image in front of a Web page in an overlay. The most commonly used of these tools is Lightbox[1]. This popularity means that there is a large amount of online support for Lightbox. For this reason, it was the tool chosen to display the images in a more gradual manner.

When a user finds a tag indicating that there is some artwork in that area, they can click on it to view photographs of it. An overlay, like the one pictured in Figure 3.9, opens with the images of the artwork loaded. If there is more than one image associated with the artwork, mousing over the left or right edge of the overlay will bring up an arrow that will scroll to the next photograph when clicked. There is a caption at the bottom of the image and buttons at the top and bottom right of the image. The cross button at the bottom closes the overlay, the printer button next to that allows users to save the image and the button at the top right with the arrows at the corners enlarges the image to fill the screen.

[1]Lightbox: http://lokeshdhakar.com/projects/lightbox2/

Figure 3.9: An example of a Lightbox overlay showing some rock art

Edit Mode

A further feature that was decided would make the system more novel and useful was an edit mode. In this mode, users would be able to add new tags to the cave and identify images from the server as corresponding to an existing tag. This would allow information to be added to the system as opposed to it being static. Therefore a button was added in the top left corner that allows users to switch to edit mode and back to normal mode.

3.3.2 Implementation

The implementation of the second iteration was where a large amount of the development time for the project was spent. This is due to the fact that the aim of this iteration was to implement the major functionality of the system, as described in the section above.

After careful consideration, it was decided that the Three.js [2] library would be used for the implementation. Three.js is an open source JavaScript library specifically for developing 3D environments in a Web browser (Danchilla & Danchilla 2012). It is used to abstract away the complexity of the WebGL API.

Setting up the scene with Three.js

Three.js simplifies the creation of the world that the environment will be drawn in. The HTML file to display the page only needs a *div* element with a specified *id* attribute. In the JavaScript file, that must also be included in the HTML, the *div* element is retrieved by its id. The library then appends the necessary HTML elements to the div and creates the scene in the *canvas* element that it added.

The library enables the creation of Scene and Camera objects as well as a Renderer object. A *PerspectiveCamera* object was used as it specifies the point of view from which a user will observe the world. A Scene is an easy way to handle all objects in an world. If something is to be rendered through the renderer, then it must be added to the scene.

[2] Three.js: http://mrdoob.github.com/three.js/

On every draw step, the renderer is called with the scene and the camera to calculate and redraw what the user can see. Three.js provides support for three different renderers; the WebGL renderer is used for this project due to its improved performance.

Inclusion of the cave models

Rendering the model in a browser was the major task of this implementation. It was found that the PLY format of the cave models did not easily convert into a JavaScript or json file that could be correctly loaded in a manner similar to that used to load the environment in the first iteration. This problem was the driving force behind the choice to use Three.js.

Even once the change to the 3D library had been decided upon, complications were faced due to the complexity of the models. The original model created by the Geomatics Department was comprised of more than half a million faces connecting over 1.5 million vertices and occupied 40MB of storage space. This proved to be too large for most of the hardware available for the project. For this reason MeshLab[3] was used to reduce the complexity and file size of the model by approximately half.

Once the complexity of the model had been reduced, it was loaded into Blender[4] 2.58. This was done so that it could be scaled down in size and rotated, ensuring that the minimum amount of processing had to be completed before it could be rendered at the correct scale and orientation in the client-side browser. Blender was also used to export the cave as an OBJ file. This allowed it to be converted to js format using a Python script packaged with Three.js.

After this preprocessing had been completed, the model was ready to be loaded into the browser. A *JSONLoader* object is used to read the file from the server when the document is ready. The callback function to be completed when this load is successful builds a mesh from the geometry returned by the loader. The camera is then pointed at the mesh and the mesh is added to the scene. In order for the model to be visible, lights are created and added to the scene as well. The finished product can be seen in Figure 3.10

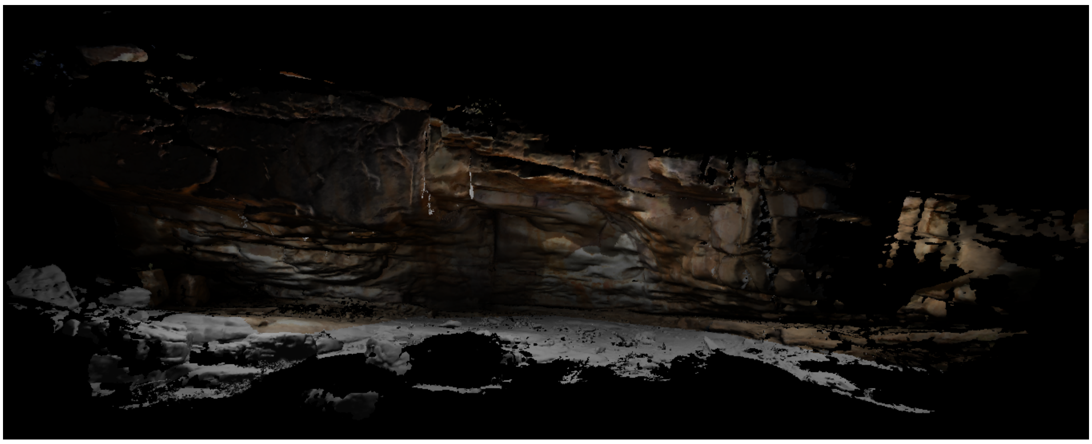

Figure 3.10: Cave model rendered in a browser (Using Three.js)

[3]Meshlab: http://meshlab.sourceforge.net/

[4]Blender: http://www.blender.org/

FPS controls

Three.js provides different types of control systems for navigating a 3D environment. The *FirstPersonControls* class enables navigation through the space using the WASD keys to move (A and D move the camera left and right respectively) and the mouse to control the rotations of the camera. Initially it was thought that this class would be enough to provide the controls sought for the system. However, upon closer inspection, it was found that these controls used the distance of the mouse from the centre of the screen to specify a rotational speed for the camera in the direction of the offset. This meant that it was difficult for the user to get the camera to be truly stationary.

The requirements of the system called for the mouse controls to be changed slightly. Instead of calculating the offset whenever the mouse was moved, the controls were made to calculate the difference between its current position and its last position. This amount is then added to the current offset from the position that it started at and the total is used to calculate the position that the camera should be pointing at.

These controls also provide additional functionality that had not been considered. Clicking with the mouse moves the camera forward if the left button is clicked and backwards if the right button is clicked. The R and F keys translate the camera up and down respectively and the Q key locks the controls so that no actions are taken when events occur until Q is pressed again.

Disappearing Artwork tags

PlaneGeometry objects textured with a PNG image are used to display the tags. All of the tags for a cave are stored in an XML file. Each tag has an *id*, a *position* and at lease one *file* element. The position specifies the x,y and z coordinates of the plane in the world, while each file element contains the relative path to an image of the rock art associated with that tag. When the page loads, a GET request is sent to a PHP script on the server that sends the XML file back in JSON format. The returned information is saved as an array of tags for future use and then it is used to populate an array of planes that are all added to the scene. All of these planes have their *visibility* attribute set to false, meaning that they will not be rendered.

Since objects in the canvas environment are not HTML elements, they cannot register events thrown by the document when the mouse moves over them. In order to register that the mouse is over a plane, a ray is cast from the camera to the mouse at every render step. If this ray intersects with one or more objects in the array of planes, then the first intersected plane has its *visibility* set to true, meaning that it will be rendered. If the ray does not intersect any of the planes then the one that it was previously intersecting has its *visibility* set back to false. Thus the tags appear when a user mouses over them and disappear when the mouse is no longer above them.

Lightbox Overlay

Generally, the Lightbox overlay is used to display larger versions of thumbnails already displayed on a Web page. In this case, the images are not visible at all except in the overlay. In order to display an image with Lightbox, the HTML element should have certain attributes set to specific values and when the user clicks on that image the overlay will open and display the image. Since the images used here are not specified by HTML elements or visible to be clicked on, another way to call the overlay with the correct images had to be found.

One option that was used as a starting point was to create invisible *image* elements in the HTML. When a user clicks on a plane the tag corresponding to that plane can be used to determine which images to display and simulate a click on one of those images. A more elegant solution followed on from this. Instead of specifying every image connected to a tag in the HTML file, a single *div* element was added to the file. When a user clicks on a plane, this *div* element is retrieved from the HTML and the jQuery html() method is used to add the necessary HTML markup to it. An *image* element is added for each image associated with that tag and a click event on the first one is simulated to call the overlay.

Edit Mode

In order to prevent the need to reload the model, this component of the School of Rock Art system does not open a new page to enter edit mode. Instead, a boolean value is set when the button to change modes is pressed. When events occur, the action taken is dependent on this boolean value.

Edit mode is intended to allow users to add new tags and images to the cave. If a user clicks anywhere on the cave when in edit mode, a ray is drawn from the camera to the mouse. If the ray intersects with the cave model then a new plane is created at that point of intersection and a PHP script is called. This script takes, as arguments, the position and id of the plane and appends new tag elements to the XML document for the cave. Due to time constraints, by the time this iteration moved to the evaluation phase, the capability to add images to tags (and thus file elements to the XML) had not yet been implemented.

3.3.3 Evaluation

The second implementation of the system was tested using heuristic evaluation (Nielsen & Molich 1990). Heuristic evaluation is a cheap and easy method of identifying the most visible usability problems so that they can be removed and the more formal user testing phase can focus on issues that are less visible and deeper (Kantner & Rosenbaum 1997). This section provides the details of how the evaluation was conducted and then lists the main problems and suggestions that arose and changes that could be considered for each of these.

A small set of three experienced users was asked to evaluate the system and provide their opinions on the design, what they felt was problematic about it and how it could be improved (Nielsen & Molich 1990, Muller, Matheson, Page & Gallup 1998, Nielsen 1994*b*). The evaluators were observed as they made use of the interface and observations and comments that they made were recorded. While this is an informal method of evaluation, it allowed the developer to assist the evaluator to understand and use the system. This assistance is acceptable since, unlike in traditional user testing, the developer relies on the evaluator's comments to evaluate the system and not on the observed actions or mistakes that they make. For this reason, more meaningful comments will be gained from expert evaluators or those who have had assistance in understanding how the component works (Nielsen 1994*a*). Evaluators were chosen from students and lecturers in the Computer Science department, since users experienced in user interface design or navigation in a three dimensional environment are favourable, considering the nature of the interface being evaluated. The findings of this evaluation are presented below.

No Loading Icon

A loading icon, or some form of indication to show when the system is ready for user input, is necessary. The large scale of the models used caused a significant delay in the time taken for the system to respond to initial user input. This was due to the time taken for both the loading and initialising of the Three.js *Mesh* object. The lack of information available about the status of the component may cause confusion for the users. Possible solutions to consider were to speed up the loading and initialisation steps, add a loading icon to indicate when loading has completed or to add a notification to users asking them to be patient as the loading may take some time. This last option can be seen implemented in Figure 3.11

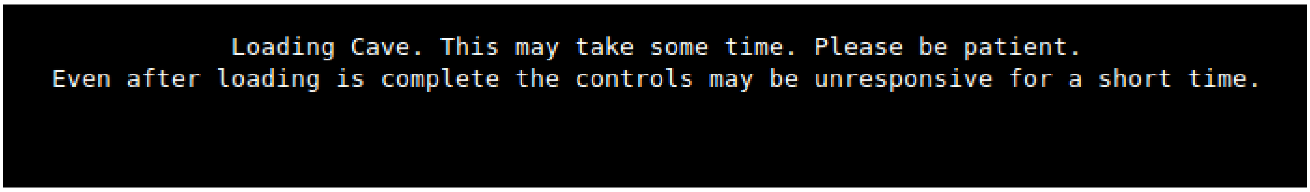

Figure 3.11: Possible solution to bad responsiveness: giving more information to the user.

Mouse controls too sensitive

The implementation of the camera controls at this stage involved using the mouse to determine the rotation of the camera and therefore, what a user could see. The speed with which the camera rotated was found to be confusing and disorientating and would sometimes lead to the user getting lost in the environment. This was easily corrected as the controls have a variable that determines the speed of the cameras rotation. This variable was reduced in value and the camera was moved slightly closer to the cave model so that there would be less chance of the user losing sight of it.

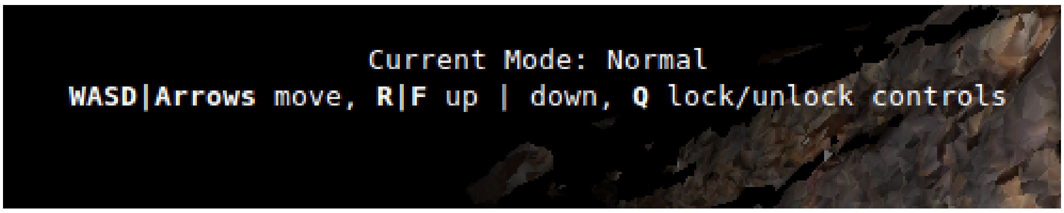

Figure 3.12: Poor control instructions in the second iteration

Poor control instructions

Information about the controls needs to be more explicit and obvious. At this point in development, the canvas element had not yet been incorporated into the HTML page that it was planned to be part of and this information was displayed as text above the canvas, as can be seen in Figure 3.12. A number of easily implemented solutions could be considered. This information needs to be provided to the user in some way and this can be done through text on the HTML page, as a user-initiated dialogue box or as a help page.

Arrows to rotate

Users should be given the option of how they want to control the camera rotation. Some users might not be familiar or comfortable with a control system similar to that used in games in the first person shooter genre. This requires the implementation of another control set or a mapping of different keys to the functions currently called when the mouse moves.

Only lock mouse

It was felt that the functionality to lock the controls should only affect the movements controlled by the mouse. When the camera rotations were being controlled by the mouse, users found it frustrating that they could not select a tag if their field of view was pointed in a different direction. It was suggested that users would seldom want to lock the keyboard controls but that they would probably often want to lock the mouse controls. The simple solution to this that was implemented was to ensure that the boolean that defined locked or unlocked controls was only checked for the mouse movements. This allows the user to use the keyboard controls even if the lock is set.

Shortcut key for tags

At times some users expressed frustration at the difficulty of finding tags since a tag is only visible when the mouse is hovering over it. The solution to this was the inclusion of a shortcut key that sets all of the tags to be visible even if the mouse is not over them. When this key is pressed again, the tags are once again hidden so that they do not obscure the environment. The result of this change can be seen in Figure 3.13

Status indicators

When users activate or deactivate various controls or states, such as locking the controls so that they no longer receive input or switching to edit mode in the system, a persistent indicator must be activated to show that they are in this mode. The lack of indicators meant that the controls did not always respond as the user expected if the state of the system was different to the one the user believed it was in. Small icons were added at the top left of the screen to show the state. These icons were small enough not to obscure the environment.

Figure 3.13: Section of the cave with all tags set to visible

Smaller Edit button

Edit mode is intended to allow users to create new tags and add images to the tags. Concern was expressed at general users having easy access to this functionality. As a result, it was deemed necessary to change the button used to enter and exit edit mode to be less obvious. A smaller button was incorporated into the final iteration in place of the larger one in order to fulfill this criteria.

Findings related to Lightbox

The mouse controls to rotate the camera need to be locked when a user views images. When a user clicked on a tag to view an item of rock art in the *Lightbox* overlay, the controls for the camera continued to receive input from the mouse. This led to the camera continuing to move while the cave environment was partially obscured by the overlay. This was partially solved by connecting the click event handler to the canvas element instead of the document. However, due to the way that Three.js handles controls for the camera, this method was unsuccessful in preventing the movement of the mouse rotating the camera. Time constraints meant that at the time of user testing this problem had not yet been suitably solved. In the third iteration of the interface, the controls stop receiving input when a user clicks on a tag and begins to receive input again when a user exits the Lightbox. While this was a large problem, the changes implemented in the control and camera system meant that it did not have a significant impact on user testing.

Users should be able to zoom into images in the Lightbox. If users want to examine a section of an image more closely, they should be able to enlarge the picture and zoom into it. Due to time constraints, this was determined to be beyond the scope of this project and was not implemented. A zoom feature has been considered for future work.

The Lightbox obscures the cave. When the user clicks a tag to view an image, the Lightbox overlay opens in the centre of the screen, obscuring the environment. To solve this problem, the overlay should be moved to the side of the screen so that it does not cover the active portion of the cave.

3.4 Final Prototype

The final iteration aims to improve on the previous iteration by addressing issues that were highlighted during the evaluation phase of the previous iteration. The user evaluation of this iteration is then evaluated through a user experience study. The description and results of this evaluation phase are presented in Chapter 4.

3.4.1 Design

The design of this iteration aims to solve issues raised during the evaluation phase of the previous iteration. The goal is to provide a positive user experience through the system.

Smaller scale models

During the design phase it was noted that the size of the model, at 22MB, was too large to feasibly send over the internet to a user's browser. This, coupled with the time taken for the model to be processed and rendered in the browser, highlighted the fact that the complexity of the models needed to be reduced further.

The larger models also caused the system to be unresponsive for sometime while it loaded and initialised them. The confusion caused by this necessitated a loading icon that would indicate when the loading was completed and the system became responsive. When the models with a reduced complexity are used, the system becomes responsive sooner. This means that users can now interact with the environment almost immediately. Since the problem was almost entirely alleviated, no status indicator was included by the time the component went to the initial round of user testing. While this solved the problem on the equipment used for testing, the consideration of slower networks and hardware on the user side of the application may still necessitate the addition of this feature.

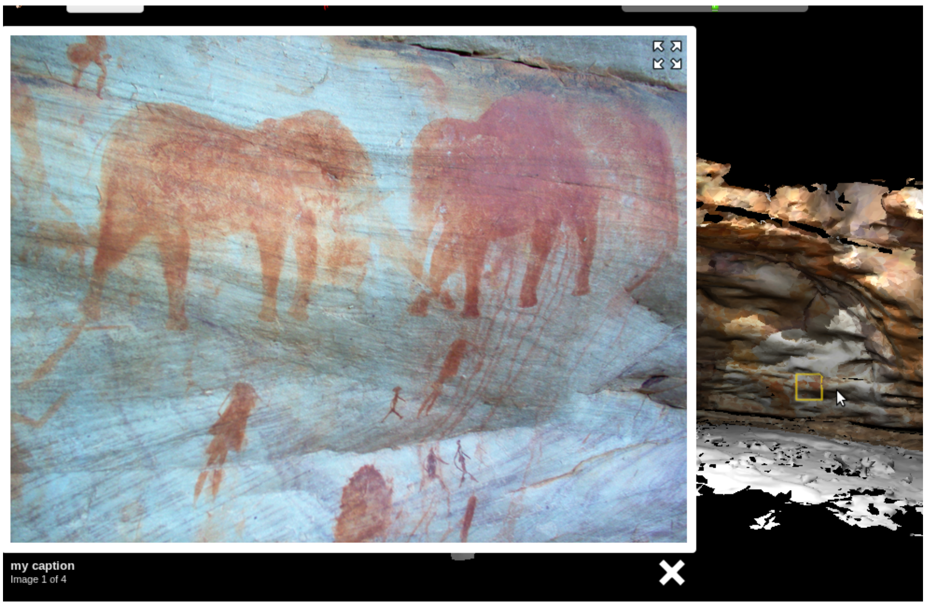

Figure 3.14: Lightbox to the side of active area

Lightbox moved to the side

As indicated in the previous evaluation, it was felt that the Lightbox overlay obscured too much of the environment. The solution to this was to edit the lightbos.css file packaged with Lightbox to move the overlay to the side of the screen so that the cave could still be seen on the other half of the screen. Despite this solution, wider images can still obscure the cave. However this was deemed acceptable as the viewing of the large images was considered one of the main features of the component. The background of the image overlay was also made to be entirely transparent so that the environment behind was still clearly visible. This also meant that a user could continue to interact with the cave while the Lightbox was active. The results of these changes can be seen in Figure 3.14

Controls Dialogue

A fully detailed controls dialogue should be added to enable users to quickly familiarize themselves with the system. The dialogue would be easily accessible through a button on the title bar and, as can be seen in Figure 3.15, would describe the controls for both the mouse-controlled camera rotations and the keyboard-controlled camera rotations.

This was considered to be the best of the proposed options since taking a user to a separate help page was undesirable and the inclusion of this text in the HTML would reduce the space available to display the environment. It was also considered favourable that the controls information be displayed at the request of the user, since experienced users are unlikely to need it.

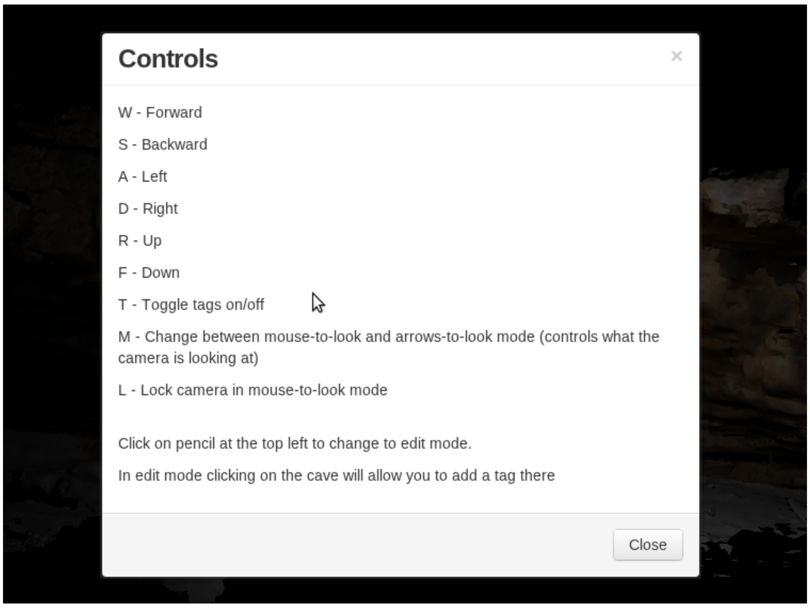

Figure 3.15: Dialogue box with detailed controls

Secondary controls

Evaluation of the second implementation indicated that users who are not familiar with FPS style controls should have the option to change to a different control system. This was added to the final iteration.

The M button on the keyboard was made to map to entering or leaving Mouse mode. Mouse mode works in the same way as the controls in the previous iteration. When not in Mouse mode, a user can control the camera entirely with the keyboard. The Arrow keys map to the rotations of the camera that are handled by the mouse in Mouse mode and the spacebar can be used to select a tag to view, thus calling up the Lightbox for that tag. However, the position of the mouse is still used to determine which tag the user is interested in.

Addition of new images

A function not yet added to the Edit mode of the previous iteration was the ability to link new images to tags when they are created. This functionality was added. When a user tries to create a new tag in edit mode, a dialogue appears containing the images of that cave that are available on the server. Users are able to select the image they wish to attach to the tag and confirm their selection before the tag is created. That image is then displayed as usual when they click on the newly created tag.

Incorporation into School of Rock Art website

Up to this point in the development process, the cave exploration component of the School of Rock Art project had not been connected to the website. The connection to the website entailed adding a title bar to the HTML page. This title bar contains links to the other components of the project and is where the status indicators are displayed.

3.4.2 Implementation

The implementation of the changes made for the final iteration are discussed in this section. This iteration aims to provide a fully functional system that can be evaluated by users in a user experience evaluation.

At this point an additional framework was incorporated to enable some of the changes to the system to be implemented. Bootstrap[5] is a front-end framework built to enable faster and easier Web development. In this implementation, Bootstrap is used to handle the appearance of the buttons, title bar and dialogue boxes.

Smaller scale models

Models of a lower complexity were needed to improve the responsiveness of the system and to reduce the bandwidth required by a user to start the system. Since the Archaeology Department had rated authenticity and accuracy of the system as important concerns, the further reduction of the models was handled by the Geomatics Department. The resulting models contained 130 000 faces and were approximately 6MB in size.

Unfortunately, this was still considered to be too complex. Further reduction was carried out using Meshlab and the resulting models contained 65 000 faces and were approximately 3MB in size. Due to time constraints, this final reduction step could not

[5]Bootstrap: http://twitter.github.com/bootstrap/

be carried out by the Geomatics Department. Regardless, these reduced models were considered acceptable and were used in the final implementation.

Secondary controls system

Three.js comes with a number of different controls classes. Inspection revealed that the *FlyControls* class would produce the functionality required for the secondary controls of the system. These controls work in the same way as the controls already implemented; the only difference is that the arrow keys are used to control the rotation of the camera instead of the mouse movements. Instances of both control classes are created when the scene is initialised. The change between controls is handled through a boolean variable that determines which set of controls is updated at every draw step.

It was found that when the new control system was set as the default (the system being used when the page is initialised), a bug inherent in the FPS controls system of Three.js was avoided. This bug causes the camera to point to the right of the scene when the page is loaded and overwrites any attempt to set a fixed direction for the camera to point in. With the new controls set to be the default the camera now points at the cave, as specified.

Addition of new images

When a user wishes to create a new tag in edit mode they need to be able to add a new image to that tag. To enable this, a PHP script is run when the page is loaded. This script loads all the images corresponding to the cave being viewed into an unordered list. This list is added to a *div* element in the HTML; this element is a member of the Bootstrap *modal* class.

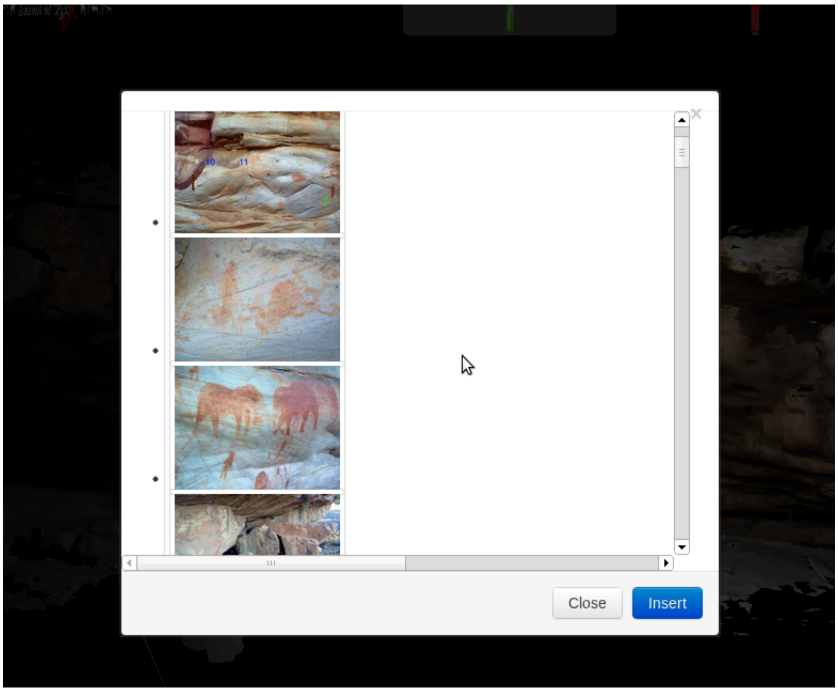

Figure 3.16: Dialogue for selecting images from the server

When a user attempts to create a new tag, a dialogue is displayed with the images in the list forming the body, as can be seen in Figure 3.16. Bootstrap handles the selection of the images and when a user confirms their selection, the path for the chosen image is retrieved. The PHP script to save tags saves this path name, along with the rest of the information about the new tag, to the XML file that specifies all of the tags in that cave.

3.5 Summary

This chapter described the design process followed to develop the Cave Navigation component of the School of Rock Art website. A three part iteration cycle was followed where each implementation was designed, implemented and then evaluated to determine the usability of the component. Three iterations of this cycle took place; the feasibility demonstration, the functional prototype and the final prototype. The evaluation of the final prototype is described in Chapter 4. The features of the final system are listed below.

3D Cave environment Reduced resolution models of the caves acquired from the Geomatics Department are used to provide the environment that the user wishes to explore.

Multiple navigation configurations Intuitive navigation through the environment is implemented using both the keyboard and mouse.

Controls information Displayed through a dialogue when the user clicks on the button to request it.

Disappearing tags Frames, that appear when the mouse moves over them and disappear when it moves off of them, indicate where in the cave artwork can be found.

Image overlay An interface that allows the user to see and scroll through photographs of a particular piece of rock art when they click on the corresponding tag.

Edit mode Provides the ability to add tags and images to the cave.

Chapter 4

Evaluation

The main goal of the Cave Navigation component of the School of Rock Art system is to provide users with a virtual, three dimensional representation of the cave that is easy and intuitive to navigate. This environment must provide an interface that is conducive to learning and allows them to gain information about the rock art through the use of the images from the Archaeology Department, without these images obscuring the environment itself. User testing was an integral stage of development and was used to inform the design process as well as to evaluate the effectiveness of the final iteration.

Two previous iterations assessed the feasibility of the concept and implemented the main features of the system. Changes made in the final implementation included the use of smaller resolution models, the addition of a secondary set of controls, the functionality to link images to a new tag when creating it and the adjustment of the Lightbox overlay to appear at the side of the screen. See Section 3.4.1 for a detailed list of the changes. This chapter describes and discusses the design, procedures and results of the evaluation phase of the final development iteration.

4.1 Experimental Design

The user testing conducted was used to test the three components of the School of Rock Art website. These being the Guided Tours, Storytelling and Cave Navigation components. This section describes the participant group, the task set and questionnaire used, the procedure followed and any bias that the experiment is open to and how that was mitigated.

4.1.1 Sample group

Since no specific user skills are needed to use features of the system, participants in the user testing phase were drawn from a large pool of students. This pool was limited to include only current students at the University of Cape Town. Participants needed to have basic competence with a desktop computer and the use of Web browsing software. Ethical clearance was obtained from both the Science Faculty Ethics in Research Committee and the Department of Student Affairs before testing on human subjects and testing on students of the University of Cape Town commenced. Refer to Appendix B for written confirmation.

A sample size of 24 was chosen and these students were recruited by way of email and posters placed in the northern section of UCT Upper Campus. All students were studying towards undergraduate degrees at the time of the study, with the majority of students being in their first or second year of study. This is estimated to put them between the ages of 18 to 25 years of age and the gender distribution in the sample was 10 females and 14 males. The range of cultural backgrounds meant that not all participants spoke English as their first language. However, it is unlikely that language barriers might have led to inaccuracies since all participants have undergone at least seven months of higher level education at an English language university and none of the language used in the interface was of a technical nature.

4.1.2 Experimental Procedure

Participants were requested to sign up for a time slot during which they would take part in the evaluation. Each time slot comprised of two or three users and was overseen by at least two group members at all times. Each group member was present to directly observe the test of their component of the system and monitor the interaction between the user and the component.

At the start of the test, the users were directed to workstations in the lab where testing would take place. Each participant was informed about the system being tested and the format and procedure of the evaluation. A group member then explained that all answers from participants were entirely anonymous and that users could choose not to answer any of the questions. They were informed that their actions would be observed and monitored but the examiner stressed that this was to determine how intuitive the system was and not to evaluate the user. The examiner also made it clear that they would be on hand to answer any questions or provide any assistance that the user might need. Once this was completed the testing commenced.

For each component of the system, the user was given a set of tasks to complete, using the component, and then a questionnaire related to that section. The tasks and questionnaires are described in the section below.

Once the questionnaires were completed for all components of the system, the participant was asked to sign a participation form and presented with a R30 participation fee. The participation form held the only information that could be used to identify an individual and it was never linked to the responses of any participant in any way, thus maintaining the anonymity promised to the participants.

4.1.3 Tasks and Questionnaire

Tasks

The tasks presented to the participants were designed to expose them to the most important features of the system to ensure that these were evaluated. The goal for the Cave Navigation component was for them to navigate through the environment, view art by selecting tags and add new tags to the interface. The tasks can be found as part of the questionnaire in Appendix A.

No strict time restrictions were imposed on the tasks but each testing time slot was loosely limited to an hour. Each task was accompanied by a short set of instructions to guide the user in the completion of that task.

The first task requested the user to experiment with the controls of the system and set it to a configuration with which they were comfortable and then to find a specific painting on the walls of the cave. The first section of this task was intended to enable the user to become more comfortable with the controls of the interface, as well as to test the controls and the instructional interface related to them. The intention of the second section of this task was to evaluate how easily a user could find the tags in a cave, use them to identify a specific item of art and how they felt about the image overlay.

The second task required the user to enter edit mode and add a new tag that would open a specific image. This task evaluated the conversion from the normal, exploratory mode to the edit mode, as well as the addition of new tags to the environment. It required users to scroll through the images that could be added, thus allowing this feature of the component to be tested as well.

Questionnaire

The evaluation of each component of the system was concluded with the completion of a questionnaire relating to that component. The questionnaire provided quantitative feedback on the user experience of the component as a whole. The usability, ease of use and satisfaction that the user experiences while using the interface formed sections of questions for the evaluation. Users responded to each question by ranking, on a five point scale, how strongly they agreed with the statement posed. A copy of the entire questionnaire can be found in Appendix A.

4.1.4 Bias

Various forms of bias can influence the results of an experiment of this kind. Forms of bias affecting this study are described below, along with where this bias arises from and how it was dealt with.

Acquiescence Response Bias

Acquiescence response bias is the tendency that participants in a survey have of answering questions in a positive (or sometimes negative) way, thus skewing results (Cloud & Vaughan 1970). This bias can arise in questionnaires such as the one used in this study where positive to negative scales are used to gage a participant's response. Acquiescence is reduced by having balanced questions in surveys. Balancing questions are a repetition of other questions in the survey that have been negated. This means that a positive answer in the first appearance of the question will correspond to a negative answer in its second appearance. The questionnaire used for this study clearly contains balanced questions to combat this bias.

Aesthetics Bias

Studies have shown that aesthetics that are perceived to be pleasing by users can lead to positive responses in terms of usability (Tractinsky 1997). The inherent nature of the interface proposed, with the three dimensional models and highly detailed photographs, creates a susceptibility to this form of bias. Since the models and photographs are integral to the interface being evaluated, no steps have been taken to mitigate this bias.

4.2 Results and Findings

The following section discusses the quantitative results gained from the survey completed by the participants of the user evaluation. Participants responded to statements posed to them on a five point scale ranging from "1, Strongly Agree", to "5, Strongly Disagree". Due to the small number of options available to users, only aggregated responses with very low variance can be considered to be suitably significant.

The tasks and questionnaire participants were requested to complete were used to test a number of attributes relevant to usability, namely: Learnability, Error, Satisfaction and Simplicity (Zhang & Adipat 2005). Only the results for the Cave Navigation component are discussed in this section. The responses as well as the observations made while the participants completed the tasks are discussed and analysed under their respective sections below.

4.2.1 Results

It was observed that a significant number of the participants chose to use the keyboard-controlled camera instead of the mouse-controlled camera. A few of these participants did not use the mouse-controlled camera for even a brief period. This may be due to a lack of clarity in the control instructions, since the control to switch between these modes was rather low in the dialogue and the difference between the two was not described. Care should be taken in improving the layout and content of this information.

In general, 13 of the 24 users agreed that the system as easy to use, and 7 others strongly agreed with this, while all but 1 user agreed that the system is usable. With 8 of those users strongly agreeing. These results indicate that users believe the system exhibits good usability and that they had a positive experience while using the system.

Learnability

Learnability is considered to be the time users take to accomplish a task the first time they use an interface (Zhang & Adipat 2005). An interface with good learnability is one that requires less time and effort to train users in the use of it. Usable and useful interfaces generally have good learnability.

Observations: Observations showed that, while the majority of users did not take a long time to determine how to navigate through the environment and complete the tasks necessary, *all* users needed to consult the controls information more than once. This indicates that training is needed to use the system.

Survey Responses: The 24 participants could not reach a general consensus on the learnability of the system. 6 participants agreed with the statement that they needed help with the system, while 11 disagreed or strongly disagreed with this statement. Users could also not agree on whether or not they would find the system easier to use with written instructions. 12 users either agreed or strongly agreed that they did not need instructions, while half (6) of that amount disagreed with the statement. 10 users felt that they did not need training before they could start with the system and 6 more felt strongly about this.

The large variance in these responses can be attributed to the degree of familiarity of individual participants with FPS games. When grouped according to how regularly user's play games of this genre, 6 of the 9 users who regularly play FPS games felt that they did not need training before they became skillful with the system.

Error

Mistakes made by users, such as incomplete tasks, button clicks that deviate from the correct path and detours, are used to evaluate usability of an interface under the *Error* attribute (Zhang & Adipat 2005). If these mistakes are not mitigated and easy to recover from, then the usability of the interface will be negatively impacted.

Observations: During the user testing it was noted that users generally made few mistakes. They were aware of where the cave was at all times and did not get lost navigating in the environment. When changing the state of the interface from the standard view mode to edit mode, it was observed that approximately half of the participants had some trouble and would sometimes press the button too many times. This resulted in them being in the wrong mode for the task they were attempting to complete. Three users also misunderstood how the art was represented in the interface and spent time examining the walls of the cave trying to find it. Another observation made was that a large number of participants clicked the "fullscreen" button on the image overlay when they were attempting to close it. Once the overlay was in this state they had difficulty closing it and returning to the cave environment.

Survey Responses: Despite these observations suggesting that errors were prevalent, when asked if they felt confident using the system, 12 users agreed that they did and 5 strongly agreed with it. Indicating that they felt errors were not overly prevalent and that they were not afraid or confused by how the interface responded to their actions. 9 users agreed and another 9 strongly agreed, amounting to 75% of participants who agreed, that they felt they could recover easily from errors that they did make. Only 3 participants disagreed with this statement. This again indicates that, when they did make a mistake, the actions needed to rectify the mistake were not surprising or unusual.

Satisfaction

The *Satisfaction* attribute of usability describes the users' attitude towards the application after using it (Zhang & Adipat 2005). This attribute is closely linked to a user's experience of the interface. A good experience will leave them satisfied whereas a bad experience will cause them to be unsatisfied. Satisfaction responses are generally prompted by emotive language.

Observations: Once users became familiar with the system, they began to show clear signs of enjoyment. No users seemed to be frustrated with the controls once they had learned how to use them.

When prompted for general comments on the interface, nine participants gave purely positive comments on their experience and two gave conditionally positive comments. The other participants chose not to comment on the system.

Survey Responses: Analysis of the survey results from the satisfaction section of the questionnaire indicates that 95.6% of users enjoyed using the interface. 1 user felt that they did not enjoy using the interface and another user refrained from answering. The remaining participants either agreed or strongly agreed, in equal numbers, that they liked using the interface. When asked if they thought the system was fun to use, *all* 24 participants agreed; additionally 50% chose to answer that they strongly agreed with this statement. This means that users were pleased with the overall experience.

Simplicity

Simplicity is a measure of how complex users find the system and the amount of effort required to complete a task (Zhang & Adipat 2005). An interface that requires a large number of clicks to complete a task or that has buttons that are difficult to find is considered to be more complex than one that does not exhibit these characteristics. Overly complex systems tend to have reduced usability.

Observations: Frustration or confusion in participants is a sign that the system is overly complex. Users of the Cave Navigation component expressed these emotions when they encountered errors but, since errors were not encountered often, this was not common and did not persist once they had recovered from the error. However, one aspect that did cause participants to display frustration and confusion was the small edit button used to enter Edit mode. While the small size was purposefully implemented, it may have had a negative effect on the user experience. Despite this, the majority of users completed the assigned tasks relatively quickly. These observations suggest that the interface exhibits simplicity and usefulness.

Survey Responses: 20 of the 24 participants felt that the controls were easy and simple to use (14 agreed, 6 strongly agreed). A smaller number of participants also felt that they were intuitive. 11 agreed and 6 strongly agreed with this statement. 7 participants strongly disagreed when asked if the system was unnecessarily complex and 14 disagreed with this statement. This amounts to 87, 5% who feel that the system is not unnecessarily complex. It can be seen that these responses gathered from the survey support the observations made.

4.2.2 Findings

In general, all users were able to complete the tasks put forward to them in a reasonable amount of time and a large percentage of them felt that the system provided a positive experience. This indicates that the system is both usable and useful. The findings with regard to the different usability aspects assessed are provided below.

Learnability

On the whole, it was found that users familiar with the type of controls used in the interface did not require as much training to use the interface as those who did not have that background. This means that those users who were familiar with the interface could have been considered to have recieved prior training for the interface. This encourages the finding that the interface requires a fairly large amount of training and is not particularly learnable.

Error

In terms of Error, the interface was not without this and users did experience errors as they navigated in and interacted with the environment. While these could potentially be significant errors and the developers need to confirm that this is the case in order to rectify the problem, they seemed to be semi-transparent to the user and did not have a major effect on the usability of the system.

Satisfaction

Responses and observations of the users seemed to convey that they were very satisfied with the system, despite the presence of errors and learnability concerns. As mentioned above, this result may be subject to aesthetics bias.

Simplicity

Users found that the system was very simple and straightforward. The original intended nature of the component was for users to be able to view the cave and artwork without having to leave the page at any time. It has come a great deal towards reaching that goal.

4.2.3 Summary

The user experience evaluation described found that the system developed displayed the following usability attributes:

- Error, at an acceptable level that does not impact significantly on the usability
- Satisfaction
- Simplicity

The system does not exhibit the learnability attribute of usability since some users need training to become familiar with the controls.

4.3 Discussion

Overall, the results of the user experience evaluation conducted show that users are very satisfied with the Cave Navigation component. Users particularly enjoyed the navigation system implemented, both the mouse-controlled and the keyboard-controlled configurations. Most users also enjoyed the ability to create tags. The aim of the Cave Navigation component of the School of Rock Art system is to provide users with a virtual representation of the cave that is satisfying and easy to navigate and that encourages learning about rock art. The results of the usability study show the satisfaction of participants, that they feel the component is easy to use and that the navigation controls are intuitive.

During the evaluation a number of participants showed an interest in the rock art as well as in the cave. These participants took care to examine all of the tags in the cave and to scroll through some of the other pictures available to add to the tags. They also explored the small corners of the cave. This shows interest in the artwork and heritage site. However, a more in depth study should be conducted to determine the level of interest and whether it was indeed a product of the Cave Navigation component. Despite this, it is plausible that the positive response from users is an indication that the system does encourage users to further their education in rock art. Thus, the component has come very close to fulfilling its goal to provide a usable and useful system to encourage learning about rock art.

Chapter 5

Conclusion

South Africa has a wealth of cultural heritage sites that contain fragile links to the rich past of the land. Some of these sites are in danger of degradation due to time, erosion and vandalism (Rüther et al. 2009). The key to protecting these sites lies in educating local citizens and the general public about the heritage and importance of the information contained there and teaching them to become stewards of the local heritage.

The Archaeology and Geomatics departments at the University of Cape Town have information and data that can be used to educate the public about local heritage. The School of Rock Art website was designed to use this information in a manner that allows scholars and researchers alike to use these resources for learning. The website was composed of three components and this report details the development of the Cave Navigation component. This component makes use of laser scans of the heritage sites from the Geomatics Department and images from the database from the Archaeology Department. These resources are used to build a Web-based system that allows users to explore the heritage sites and examine photographs of the artwork taken there. A user experience evaluation, using students of the University of Cape Town as voluntary participants, was used to determine if the final iterations of each of the components of the website satisfied this question.

The cave navigation component was built in order to answer the research question of: Is it possible to build a usable and useful 3D cave navigation system to encourage learning about rock art? The evaluation of this component revealed that users enjoyed using the cave navigation system and thought that it was a useful component. While the responses did indicate that the controls system used requires training and therefore the system does not exhibit the Learnability attribute of usability, the system does exhibit the usability attributes related to Simplicity, Satisfaction and Error. Therefore the system can be considered to be usable and useful. It wass also shown to be plausible that the system encourages learning about rock art. Thus positively answering the research question.

5.1 Future Work

At various stages during the development process of the cave navigation component, a number of possible extensions were considered or suggested. It is likely that these features would improve the usability and user experience of the system. However, for various reasons, these features could not be implemented during this project so they are detailed below for future work.

Audio

Users felt that the inclusion of audio in the system would have a large, positive impact on the experience of using the system. Ambient sounds such as soft wind or chirping crickets could be added to improve a user's immersion in the environment. Voice over sounds could be used to impart relevant information to the user when they view specific items of artwork. An important aspect to note, though, is that users felt these audio elements should be optional and that there should be a button to mute the sounds if the user so desires.

Zoom on Lightbox

Since the system is intended to be used by both scholars and researchers it is conceivable that users may wish to study the finer details of the artwork. Therefore, they should be able to zoom into the images of the art in order to better examine them.

Edit Mode Improvements

The results of the user testing showed that a number of users were frustrated by the small edit button. This was purposefully implemented in to prevent general users from adding erroneous information to the site. However, considering the impact that it had on the user experience, this button should be made bigger and easier to find and access to the Edit mode should be controlled in a different way. The Edit button should prompt a user to login to the server before they are able to make any changes to the site. This will ensure that all of the information on the site is accurate.

A major feature that needs to be added to the system, but was not implemented due to time constraints, is the ability to add images to existing tags when in edit mode. This feature will allow users to select a tag and link images from the server to that tag in much the same way that adding an image to a newly created tag works.

During user testing, users noted that they felt the functionality to delete a created tag or to remove images from a tag should be included in the Edit mode. The inclusion of this feature would make the need for improved access control to the Edit mode a greater priority.

A feature that was overlooked in implementation was that of allowing users to add captions to images when they are added to tags. This can be implemented through an editable dialogue box that accompanies the box to select an image to add.

Other caves

More heritage sites need to be added to the system to show that it is extensible. Users will be able to choose which cave to view through a selection page. This was not incorporated in this project due to time constraints.

Improved Controls Information

The controls information needs to be displayed in a way that allows new users to easily refer back to it and experienced users to remove it so that it does not take up space on the page. This can be done by placing the information on a sliding pane that users can choose to have open or closed. The important difference between this method and the currently employed method of using a dialogue box is that the new method would allow novice users to continue interacting with the cave while the controls are open.

Interaction with other components

Functionality to interact with the Storytelling and Guided Tours components of the website would be expected to greatly improve the experience of the website in its entirety.

Features that would enable this include a workspace that maintains its state between the pages. This workspace would enable users to collect information found on any of the components on the site in one place and export it in their chosen format.

Another possible way of integrating the Cave Navigation and StoryTelling components would be to have links present in both components if an image is used in both. This means that if a user is reading a story containing an image that is also displayed in a cave, they will see a link stating that they can see where this image is found in the cave and that link will take them to the cave, where the tag linked to the image in question will be highlighted. If a user is looking at an image in a cave that is also used in a story, a link will take them to that particular story.

Integration with the Guided Tours component can be achieved by allowing users to take a tour of the images in a cave. Another way of integrating the two components would be allowing a user creating a tour to record their navigation through a cave and then to have this replayed in the tour.

Information with images

To enhance the learning experience provided by the component, the image overlay can include information about the artwork in the image. This content would have to be created or added by users when they create new tags and so improved access control to the Edit mode would help assure the accuracy of the information.

Bibliography

Baldonado, M., Chang, C.-C. K., Gravano, L. & Paepcke, A. (1997), 'The stanford digital library metadata architecture', *International Journal on Digital Libraries* **1**, 108–121. [Online] Accessed at /url = http://dx.doi.org/10.1007/s007990050008.

Borner, K., Dillon, A. & Dolinsky, M. (2000), Lvis-digital library visualizer, *in* 'Proceedings. IEEE International Conference on Information Visualization', pp. 77 –81. [Online] Accessed at /url =http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=859740.

Bryan, P. & Heritage, E. (2009), 'Recording prehistoric rock art: A three dimensional approach', *Proceedings of Electronic Visualisation and the Arts (EVA 2009)* pp. 1–10.

Cloud, J. & Vaughan, G. (1970), 'Using balanced scales to control acquiescence', *Sociometry* pp. 193–202.

Danchilla, B. & Danchilla, B. (2012), Three.js framework, *in* 'Beginning WebGL for HTML5', Apress, pp. 173–203. [Online] Accessed at /url = http://dx.doi.org/10.1007/978-1-4302-3997-0_7.

Dorn, R., Whitley, D., Cerveny, N., Gordon, S., Allen, C. & Gutbrod, E. (2008), 'The rock art stability index: A new strategy for maximizing the sustainability of rock art', *Heritage Management* **1**(1), 37–70.

Daz-Andreu, M., Brooke, C., Rainsbury, M. & Rosser, N. (2006), 'The spiral that vanished: the application of non-contact recording techniques to an elusive rock art motif at castlerigg stone circle in cumbria', *Journal of Archaeological Science* **33**(11), 1580 – 1587. [Online] Accessed at /url = "http://www.sciencedirect.com/science/article/pii/S030544030600046X".

Gonzlez-Aguilera, D., Muoz-Nieto, A., Gmez-Lahoz, J., Herrero-Pascual, J. & Gutierrez-Alonso, G. (2009), '3d digital surveying and modelling of cave geometry: Application to paleolithic rock art', *Sensors* **9**(2), 1108–1127. [Online] Accessed at /URL = http://www.mdpi.com/1424-8220/9/2/1108.

Good, L., Popat, A., Janssen, W. & Bier, E. (2005), A fluid interface for personal digital libraries, *in* A. Rauber, S. Christodoulakis & A. Tjoa, eds, 'Research and Advanced Technology for Digital Libraries', Vol. 3652 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, pp. 162–173. [Online] Accessed at /url = http://www.springerlink.com/content/ bymr5tl90tl0rd1r/fulltext.pdf.

Guven, S. & Feiner, S. (2006), Interaction techniques for exploring historic sites through situated media, *in* '3DUI 2006. IEEE Symposium on 3D User Interfaces, 2006.', pp. 111 – 118. [Online] Accessed at /url = http://ieeexplore.ieee.org/stamp/stamp.jsp? tp=&arnumber=1647516.

Harrison, B. L., Ishii, H., Vicente, K. J. & Buxton, W. A. S. (1995), Transparent layered user interfaces: an evaluation of a display design to enhance focused and divided attention, *in* 'Proceedings of the SIGCHI conference on Human factors in computing systems', CHI '95, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, pp. 317–324. [Online] Accessed at /url = http://dx.doi.org/10.1145/223904.223945.

Isaacman, A., Lalu, P. & Nygren, T. (2005), 'Digitization, history, and the making of a postcolonial archive of southern african liberation struggles: The aluka project', *Africa Today* **52**(2), pp. 55–77. [Online] Accessed at /url = http://www.jstor.org/stable/4187703.

Jerding, D. & Stasko, J. (1998), 'The information mural: a technique for displaying and navigating large information spaces', *Visualization and Computer Graphics, IEEE Transactions on* **4**(3), 257 –271. [Online] Accessed at /url = http://ieeexplore.ieee.org/stamp/stamp.jsp? tp=&arnumber=722299.

Kantner, L. & Rosenbaum, S. (1997), Usability studies of www sites: heuristic evaluation vs. laboratory testing, *in* 'Proceedings of the 15th annual international conference on Computer documentation', SIGDOC '97, ACM, New York, NY, USA, pp. 153–160. [Online] Accessed at /url = http://doi.acm.org/10.1145/263367.263388.

Kimber, D., Foote, J. & Lertsithichai, S. (2001), Flyabout: spatially indexed panoramic video, *in* 'Proceedings of the ninth ACM international conference on Multimedia', MULTIMEDIA '01, ACM, New York, NY, USA, pp. 339–347. [Online] Accessed at /url = http://doi.acm.org/10.1145/500141.500192.

Lieberman, H. (1994), Powers of ten thousand: navigating in large information spaces, *in* 'Proceedings of the 7th annual ACM symposium on User interface software and technology', UIST '94, ACM, New York, NY, USA, pp. 15–16. [Online] Accessed at /url = http://doi.acm.org/10.1145/192426.192434.

Lippman, A. (1980), 'Movie-maps: An application of the optical videodisc to computer graphics', *SIGGRAPH Comput. Graph.* **14**(3), 32–42. [Online] Accessed at /url = http://doi.acm.org/10.1145/965105.807465.

Marks, L., Hussell, J. A., McMahon, T. M. & Luce, R. E. (2005), 'Activegraph: A digital library visualization tool', *International Journal on Digital Libraries* **5**, 57–69. [Online] Accessed at /url = http://dx.doi.org/10.1007/s00799-004-0110-z.

Mazel, A. (2008), 'Presenting the san hunter-gatherer past to the public: a view from the ukhahlamba-drakensberg park, south africa', *Conservation and Management of Archaeological Sites* **10**(1), 41–51.

McHenry, K. & Bajcsy, P. (2008), An overview of 3d data content, file formats and viewers, Technical report, Technical Report ISDA08-002.

Muller, M. J., Matheson, L., Page, C. & Gallup, R. (1998), 'Methods & tools: participatory heuristic evaluation', *interactions* **5**(5), 13–18. [Online] Accessed at /url = http://doi.acm.org/10.1145/285213.285219.

Nielsen, J. (1994*a*), 'How to conduct a heuristic evaluation', *On the World Wide Web* .

Nielsen, J. (1994*b*), Usability inspection methods, *in* 'Conference Companion on Human Factors in Computing Systems', CHI '94, ACM, New York, NY, USA, pp. 413–414. [Online] Accessed at /url = http://doi.acm.org/10.1145/259963.260531.

Nielsen, J. & Molich, R. (1990), Heuristic evaluation of user interfaces, *in* 'Proceedings of the SIGCHI conference on Human factors in computing systems: Empowering people', ACM, pp. 249–256.

Pwiti, G. & Ndoro, W. (1999), 'The legacy of colonialism: Perceptions of the cultural heritage in southern africa, with special reference to zimbabwe', *African Archaeological Review* **16**, 143–153. [Online] Accessed at /url = http://dx.doi.org/10.1023/A:1021624632398.

Rao, R., Pedersen, J. O., Hearst, M. A., Mackinlay, J. D., Card, S. K., Masinter, L., Halvorsen, P.-K. & Robertson, G. C. (1995), 'Rich interaction in the digital library', *Commun. ACM* **38**(4), 29–39. [Online] Accessed at /url = http://doi.acm.org/10.1145/205323.205326.

Rüther, H. (2007), 'Building a digital library of scholarly resources from the developing world: An introduction to aluka', *African Arts* **40**, 1 – 7. [Online] Accessed at /url="http://dx.doi.org/10.1162/afar.2007.40.2.1".

Rüther, H., Chazan, M., Schroeder, R., Neeser, R., Held, C., Walker, S. J., Matmon, A. & Horwitz, L. K. (2009), 'Laser scanning for conservation and research of african cultural heritage sites: the case study of wonderwerk cave, south africa', *Journal of Archaeological Science* **36**(9), 1847 – 1856. [Online] Accessed at /url = "http://www.sciencedirect.com/science/ article/pii/S0305440309001344".

Shiaw, H., Jacob, R. & Crane, G. (2004), The 3d vase museum: a new approach to context in a digital library, *in* 'Proceedings of the 2004 Joint ACM/IEEE Conference on Digital Libraries.', pp. 125 – 134. [Online] Accessed at /url = http://ieeexplore.ieee.org/stamp/stamp.jsp? tp=&arnumber=1336109.

Tractinsky, N. (1997), Aesthetics and apparent usability: empirically assessing cultural and methodological issues, *in* 'Proceedings of the ACM SIGCHI Conference on Human factors in computing systems', CHI '97, ACM, New York, NY, USA, pp. 115–122. [Online] Accessed at /url = http://doi.acm.org/10.1145/258549.258626.

Wan, G. (2006), 'Visualizations for digital libraries', *Information Technology and Libraries* **25**(2), 88–94. [Online] Accessed at /url=http://0-web.ebscohost.com.innopac.
library.unr.edu/ehost/pdfviewer/pdfviewer?
vid=4&hid=112&sid=c09055c3-c6dc-4280-8ab9-2a00aaecd706
%40sessionmgr113.

Zhang, D. & Adipat, B. (2005), 'Challenges, methodologies, and issues in the usability testing of mobile applications', *International Journal of Human-Computer Interaction* **18**(3), 293–308. [Online] Accessed at /URL = http://www.tandfonline.com/doi/abs/10.1207/s15327590ijhc1803_3.

Appendix A

Questionnaire

School of Rock Art user evaluation Survey

Confidentiality and Privacy

If you consent to participate in this evaluation, your personal information will be kept confidential. Any information you choose to provide will be kept private between you and the researcher.

The record kept of your survey responses does not contain any identifying information about you unless a specific question asks for it. Answering any of the questions is optional, but please answer as many questions as possible to guarantee the validity of the results of the study.

Statement of Consent

"I acknowledge that I have read the above explanation of this evaluation. I understand that the collected data from this survey will be analyzed and used to evaluate the mentioned website. I also understand that the researcher will not disclose my personal information. By selecting the option 'Next' below I agree to participate in this evaluation"

There are 25 questions in this survey

About you

1 [1.1]

What is your main field of study (Faculty)?

Please write your answer here:

2 [1.2]

What academic qualification are you currently registered for?

Please choose only one of the following:

- Bachelors
- Honours
- Masters
- PhD
- Other

3 [1.3]

What is your year of study of the current degree?

Please choose only one of the following:

- First
- Second
- Third
- Fourth
- Fifth
- Other

Cave Exploration

Please complete the tasks below and then return to answer the questions

Task 1:
1 Play with the controls for a moment and choose the configuration you prefer.
2 Find the painting of the elephant in the cave.
Task 2:
1 Go into edit mode in the Keerbos cave (the pencil in the top left corner).
2 Add a new tag anywhere on the cave and add an image of an elephant to this tag.

18 [4.1]

How often do you play first person shooter style computer games?

Please choose only one of the following:

- ◯ A few hours a day
- ◯ A few hours a week
- ◯ A few hours a month
- ◯ Seldom
- ◯ Never

19 [4.2]Usability

Please choose the appropriate response for each item:

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree
I thought the system was easy to use	◯	◯	◯	◯	◯
I found the system unnecessarily complex	◯	◯	◯	◯	◯
I think that I would need the support of a technical person to be able to use this system	◯	◯	◯	◯	◯
I found the various functions in this system were well integrated	◯	◯	◯	◯	◯
I needed to learn a lot of things before I could get going with this website	◯	◯	◯	◯	◯
I felt very confident using the system	◯	◯	◯	◯	◯

20 [4.3]Ease of Use

Please choose the appropriate response for each item:

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree
It is usable	◯	◯	◯	◯	◯
I can use it without written instructions	◯	◯	◯	◯	◯
I don't notice any inconsistencies as I use it	◯	◯	◯	◯	◯
I can recover from mistakes quickly and easily	◯	◯	◯	◯	◯
I did not find it easy to navigate through the site	◯	◯	◯	◯	◯
The controls were easy to use	◯	◯	◯	◯	◯
The controls were intuitive	◯	◯	◯	◯	◯

21 [4.4]Satisfaction

Please choose the appropriate response for each item:

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree
It is fun to use	◯	◯	◯	◯	◯
I like using the interface of this system	◯	◯	◯	◯	◯
I found the system boring	◯	◯	◯	◯	◯
I would suggest this to my friends	◯	◯	◯	◯	◯
I want to spend more time using it	◯	◯	◯	◯	◯

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree
It is wonderful	◯	◯	◯	◯	◯
I feel I need to have it	◯	◯	◯	◯	◯
It is pleasant to use	◯	◯	◯	◯	◯

22 [4.5]

In your opinion, which feature(s) was the most useful?

Please write your answer here:

23 [4.6]

In your opinion, which features(s) was the least useful?

Please write your answer here:

24 [4.7]

What features do you think could be improved on or added?

Please write your answer here:

25 [4.8]

Any other comments?

Please write your answer here:

Thank you very much for completing this survey and for helping us with our project.

Please remember to collect your remuneration from us.

Please submit by 24.10.2012 – 00:00
Submit your survey.
Thank you for completing this survey.

Appendix B

Ethical Clearance

Department of Environmental and Geographical Science
University of Cape Town
RONDEBOSCH 7701
South Africa

e-mail: Michael.meadows@uct.ac.za
phone : + 27 21 650 2873
fax : +27 21 650 3791

27th September 2012

Ms Joanne Marston
Department of Computer Science
MRSJOA005@myuct.ac.za

Dear Ms Marston

Heritage and Learning

I am pleased to inform you that, having scrutinized the details of your above-named applications for research ethics clearance, the Faculty of Science Research Ethics Committee has approved your proposal in terms of its attention to ethical principles.

Your approval code for the project is: SFREC 43_2012

I wish you success in the work involved.

Yours sincerely

Michael E Meadows
Professor and Head of Department
Chair: Science Faculty Ethics in Research Committee

	RESEARCH ACCESS TO STUDENTS	**DSA 100**

NOTES

1. This form must be FULLY completed by applicants that want to access UCT students for the purpose of research.
2. Return the completed application form together with your research proposal to: Moonira.Khan@uct.ac.za; or deliver to: Attention: Executive Director, Department of Student Affairs, North Lane, Steve Biko Students' Union, Room 7.22, Upper Campus, UCT.
3. The turnaround time for a reply is approximately 10 working days.
4. NB: It is the responsibility of the researcher/s to apply for and to obtain ethical clearance and access to staff and/or students, respectively to the (a) Faculty's 'Ethics in Research Committee' (EiRC) for ethics approval, and (b) Executive Director, HR for approval to access staff for research purposes and the (c) Executive Director, Student Affairs for approval to access students for research purposes.
5. For noting, a requirement of UCT (according to Senate policy) is that items (1) and (4) apply even if prior clearance has been obtained by the researcher/s from any other institution.

SECTION A: RESEARCH APPLICANT/S DETAILS

Position	Staff / Student No	Title and Name	Contact Details (Email / Cell / land line)
A.1 Student Number	a)MRSJOA005 b)CRWKAI001 c)LWRMAR004	a)Miss Joanne Marston b) Miss Kaitlyn Crawford c) Mr Marco Lawrence	Mrsjoa005@myuct.ac.za Crwkai001@myuct.ac.za marco.lawrence85@gmail.com
A.2 Academic / PASS Staff No.	1331888 1331046 1330784		
A.3 Visiting Researcher ID No.	N/A		
A.4 University at which a student or employee	UCT	Address if *not* UCT:	
A.5 Faculty/ Department/School	Computer Science		
A.6 APPLICANTS DETAILS If different from above	Title and Name	Tel.	Email

SECTION B: RESEARCHER/S SUPERVISOR/S DETAILS

Position	Title and Name	Tel.	Email
B.1 Supervisor	Hussein Suleman	+27 21 650 5106	Hussein@cs.uct.ac.za

SECTION C: APPLICANT'S RESEARCH STUDY FIELD AND APPROVAL STATUS

C.1 Degree (if a student)	BSc (Honours) Computer Science
C.2 Research Project Title	Heritage and Learning
C.3 Research Proposal	Attached: Yes ■ No ☐
C.4 Target population	Undergraduate Students
C.5 Lead Researcher details	If different from applicant:
C6. Will use research assistant/s	Yes ☐ No ■
C.7 Research Methodology and Informed consent:	User Study #1 and #2: Participants will be voluntarily recruited from campus and asked to review the system. Participants will then be asked to fill out a questionnaire concerning the usability and design of the system. Users will be required to sign a consent form prior to the evaluation and may at any time decide to opt out of the review.
C.8 Ethics clearance status from UCT's Ethics in Research Committee (EiRC)	Approved by the EiRC: Yes ☐ No ☐ Awaiting response : ■ If yes, attach copy and state the date and ref. no of EiRC approval: Date of application if awaiting response: 5 September 2012

SECTION D: APPLICANT/S APPROVAL STATUS FOR ACCESS TO STUDENTS FOR RESEARCH PURPOSE

(To be completed by the ED, DSA or Nominee)

	Approved / * With Terms	* Conditional approval Terms	Applicant/s Ref. No.:	
D.1 APPROVAL STATUS	****Yes – Approved* With Terms** ■	**(a)** Access to students for this research study must only be undertaken after written ethics approval has been obtained. ***(b)*** . **(b)** In event any ethics conditions are attached, these must be complied with before access to students.	a) MRSJOA005/Joanne Marston b) CRWKAI001/Kaitlyn Crawford **c)** LWRMAR004/Marco Lawrence	
D.2 APPROVED BY:	Designation	Name	Signature	Date
	Executive Director Department of Student Affairs	***Ms Moonira Khan***	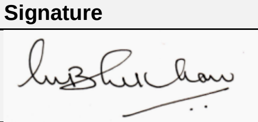	***17/09/2012***