

**Eye Tracking with State-of-the-Art
Radiography**

Michael Terzza
Computer Science
Session 2009 / 2010

The candidate confirms that the work submitted is their own and the appropriate credit has been given where reference has been made to the work of others.

I understand that failure to attribute material which is obtained from another source may be considered as plagiarism.

(Signature of student)_____

Contents

1	Introduction	1
1.1	Specialist Equipment	3
1.2	Aim	4
1.3	Objectives	4
1.4	Minimum Requirements	5
1.5	Schedule	6
1.5.1	Milestones	6
1.5.2	Deadlines	7
2	Background Research	8
2.1	Radiology Research	8
2.1.1	Pilot Study	8
2.1.2	Eye tracking in Radiology	8
2.1.3	Meetings	9
2.2	Eye Tracking Systems	9
2.2.1	Previous Projects	9
2.2.2	SR Research	9
2.2.3	General Purpose Eye Tracking	10
2.2.4	Eye Tracking Analysis	10
3	Data Collection Phase	12
3.1	EyeLink II System	12
3.2	Software	14
3.2.1	Design	14
3.2.2	EDF Data Files	16
3.2.3	Code Base	18
3.2.4	Testing	18
3.2.5	Re-adjustments	19
3.3	Experiment	19
3.3.1	Experiment Design	19

3.3.2	Deployment Preparation	20
3.3.3	Deployment	22
3.3.4	Data Collection	22
3.4	Evaluation	22
3.4.1	Communication with L.G.I.	23
3.4.2	Specification Issues	23
3.4.3	Unfamiliar Language	24
3.4.4	Cross Platform Issues	24
3.4.5	Eye Tracking Calibration	25
3.4.6	False Conditions	26
3.4.7	Uniqueness of Human Participants	27
3.4.8	Fatigue of Participants	28
3.4.9	Unforeseen Issues	28
3.4.10	Phantom Issues	29
4	Analysis Phase	30
4.1	Planning	30
4.1.1	Tools	30
4.1.2	Data Correction	31
4.1.3	Nodule Locations	32
4.2	Required Results	34
4.2.1	Coverage	34
4.2.2	Accuracy	35
4.2.3	Confidence Rating	37
4.2.4	Results	37
4.3	Further Results	41
4.3.1	Spatiotemporal Analyses	41
4.3.2	Observations	42
4.3.3	Analysis	43
4.4	Evaluation	44
4.4.1	Saccades and Fixations	44
4.4.2	Values Encoded in Circle Diameters	45
4.4.3	Animated Heat-Maps	45
4.4.4	Nodule Coordinates	45
5	Conclusion	47
5.1	Outcomes	47
5.2	Future Work	47
5.3	General Evaluation	48

SUMMARY

The main aim for this project was for it to be the technical compliment to a pre-planned radiology study by a 4th year medical student.

The aim of the radiology study was to use eye tracking technologies to help assess the way a pool of Radiologists' viewed and reported on images taken with a state-of-the-art X-Raying technique, Dual Energy Imaging. Dual Energy X-Rays are a fairly modern technology and among their many uses, can provide significant advantages in the correct detection of the early stages of lung cancer.

The project involved using expensive, specialist equipment including a head mounted eye tracker and a high quality medical grade monitor.

Equipment was set up in the Leeds General Infirmary for data collection, which lasted just over a week.

ACKNOWLEDGEMENTS

First and foremost my thanks go to my project supervisor Dr Roy Ruddie. He has provided excellent support and guidance throughout the project. He has regularly helped me to develop some of the raw ideas I've had, into worthwhile work, with his knowledgeable input, and has kept me motivated and challenged in my work.

He also provided the correct statistical analysis of results needed, from the data I have gathered and processed, for use in this project.

I would also like to thank Andrew West and David Brettie, the medical collaborators that I worked on this project with, for their hard work and help introducing me to the field of Radiology.

Thanks also to all the Radiologists who generously gave their time to volunteer in the study.

Last but not least I would like to thank Sam Johnson, who provided his time and excellent Matlab skills to help me to develop one of my ideas.

Chapter 1

Introduction

It should be noted that the reader is encouraged, if possible, to view the accompanying online materials referenced within this report for extra understanding.[23][25][24]

When reporting on Chest X-Rays, one of the many problems that Radiologists look for is small abnormal growths in the lungs called nodules. These nodules may be calcified (fairly dense) or non-calcified (a lesser density, similar to that of soft tissue). It is these non-calcified nodules that are the dangerous, early stages of lung cancer.

The most common form of diagnostic procedure for suspected lung cancer in a patient is, firstly, a traditional chest X-Ray. This may then lead on to more in-depth testing i.e. CT scans. This project will look at a similar, but more advanced technique, when compared to the traditional chest X-Ray, Dual Energy X-Raying.

Dual Energy X-Raying is a fairly new technique where by two traditional X-Rays are taken in close succession, but at differing levels of radiation energy, producing two similar X-Ray images. Due to the fact that more dense material absorbs radiation at a different rate to soft tissue, mathematical calculations can be performed on these two images to give two ‘subtracted’ images: a soft tissue image and a bony image [12].

This technique therefore provides an advance tool allowing Radiologists to better identify a nodule from normal anatomical features, and to better classify whether identified nodules are of a calcified or non-calcified type. In essence, a more accurate diagnosis. The uptake of the Dual Energy process has been a lot slower than expected in the field on radiology [12] and this research project aims to aid in the understanding of why this is, and test methods in which we may be able to better the use of the Dual Energy process.

For clarity it should be noted that, throughout this report, when referring to either the left or right hand side of an image, it is meant from the reader’s view point, not from the patient’s.

Figure 1.1 shows the two outputted images from a Dual Energy Chest X-Ray after the subtraction processing. There are two nodules present in this image, a non-calcified nodule in the upper left lung, and a calcified nodule in the upper right lung. The nodules can be seen as small light circles. This pair of images demonstrates some of the advantages of Dual Energy

X-Rays. The soft image (on the left) shows minimal bone structure, therefore not obscuring potential nodules. It also shows the non-calcified nodule more prominently than the calcified nodule due to the filtering effect of the more dense materials. The boney image (on the right) shows minimal soft tissue and the calcified nodule in the right lung is clearly visible, whereas the non-calcified nodule is hardly visible at all.

A Radiologist should be able to use this extra information to better diagnose whether the detected nodules are calcified or non-calcified, and hence potentially dangerous or not. It may also allow them detect nodules that may have been obscured by the rib cage.




Figure 1.1: Example of the output from a dual energy X-Ray. Soft tissue image on the left, boney image on the right.

This project has been carried out in conjunction with the Medical Physics Department at the Leeds General Infirmary (LGI), and in particular, with a forth year medical student.

One hypothesis proposed by the collaborators to explain the slow uptake of Dual Energy imaging, despite its clear advantages, is that the images produced are unfamiliar to the majority of Radiologists. It has been shown by Kundel [7] that there is a trend for a Radiologist's gaze to focus on areas of high contrast in the image (i.e. the borders of the ribs) when reporting on an X-Ray. It is thought that the removal of "anatomical clutter" in the soft tissue images not only leaves some Radiologists with an unfamiliar view, but also removes key land-mark features that would normally aid in their visual search strategy. This may reduce the visual coverage of an image and may reduce the accuracy of the diagnosis.

It is therefore proposed that the quality of visual search on the soft tissue images may be improved by overlaying a low intensity of the boney images. The characteristic land-marks would then be visible to aid in visual search, but not to the extent that bones may obscure

potential nodules.

A pilot study has previously been carried out by the collaborators to gain an insight into the correctness of this hypothesis.

1.1 Specialist Equipment

Successful execution of this project requires the use of many pieces of specialist equipment:

The School of Computing owns a state-of-the-art, head mounted, eye tracking piece of equipment, the EyeLink II from SR Research (as shown in Figure 1.2). This was to be used to for the acquisition of the participant's gaze direction for this project.

Due to the importance of the viewing quality of an image in radiology, a PC complete with a medical grade monitor was loaned from the medical physics department. This can be seen later in Figure 3.1.

A lung phantom (as shown in Figure 1.3) was used as the subject for all the X-Ray images. This allows for arbitrary nodule placement in the lungs, whilst providing realistic X-Ray images. The phantom, which belongs to the Medical Physics Department, was the LungMan N1 by Kyoto Kagaku, Japan. X-Rays of this lung phantom were then taken with a GE Healthcare Definium™ 6000 system.

Specific details of this equipment will be described later in this report.



Figure 1.2: The EyeLink II Headset used for eyetracking data collection.




Figure 1.3: The Lung Phantom. Disassembled in the left picture with a nodule inserted into the top of the right lung, and reassembled in the right picture, in front of the X-Ray Machine.

1.2 Aim

The aim of this project is to build upon the collaborators' pilot study and extend the research by using more advanced methods of both data collection and analysis. This requires the development of a new system to track the gaze of a Radiologist whilst reporting on Dual-Energy X-Rays and to research and develop methods of visualising the eye tracking data collected from the studies for analysis. This should aid in both the understanding, and improving the viewing, of this new form of medical imaging.

1.3 Objectives

The objectives of the project are to:

- Gain an understanding of the field of Radiology and more precisely the use of traditional Chest X-Rays and Dual Energy Chest X-Rays.
- Gain an understanding of the field of Eye Tracking and its use within medical imaging.
- Successfully co-ordinate the development of the proposed systems with the pre-planned medical study, to ensure that all systems function as expected and are able to collect the desired data as accurately as possible.
- Analyse and interpret the collected data effectively for conclusions to be drawn, based upon the needs of the medical collaborators.

- Expand the analysis of the collected data further, using visualisations in an attempt to gain insight into detail such as search strategies.

1.4 Minimum Requirements

The minimum requirements are:

- Develop a library to interface with the EyeLink II eye tracker. This will be implemented in C++ to ensure it is fast enough to deal with the possible 500Hz refresh rate of the eye tracker hardware.
- Develop an OpenGL Application to display medical images and incorporate the eye tracker data for future analysis. Functionality will include timing of the viewing trials and the ability to click on a suspected nodule and record the co-ordinates. This must be suitable for and executable on the PC belonging to the LGI, due to the need to use the medical grade monitor in the study.
- Develop applications to display / visualise results from eye tracking trials. These should have the ability to replay the subject's eye's scan-path and also generate heat-maps, coverage maps and mouse clicks from the collected data.

The possible extensions and enhancements to the project are:

- Visualisation applications to analyse and understand the search strategies of a Radiologists.
- Develop new ways to analyse the dataset collected during the study.
- Training applications for Radiographers (Radiography Graduates) and / or Registrars (Medics undergoing advanced training in a specific field), based on the patterns and strategies of the Radiologist's observed trials.
- Testing applications for Radiographers or Registrars, both for self-assessment and for assessment of suitability for certain roles.
- Viewing applications to aid in the viewing process of Dual Energy X-Rays. This could possibly include advanced image pre-processing of the X-Rays, allowing abnormalities to be more easily noticed and correctly diagnosed.
- Development of generic applications and / or software libraries that will be easily incorporated into future research projects or studies that make use of the eye tracker.

1.5 Schedule

The project was separated into two distinct phases, the Data Collection Phase and the Analysis Phase. The Data Collection involved design, development, and execution of the data collection experiments. The Analysis Phase involved studying the collected data to produce visualisations and analysis related to the experiment.

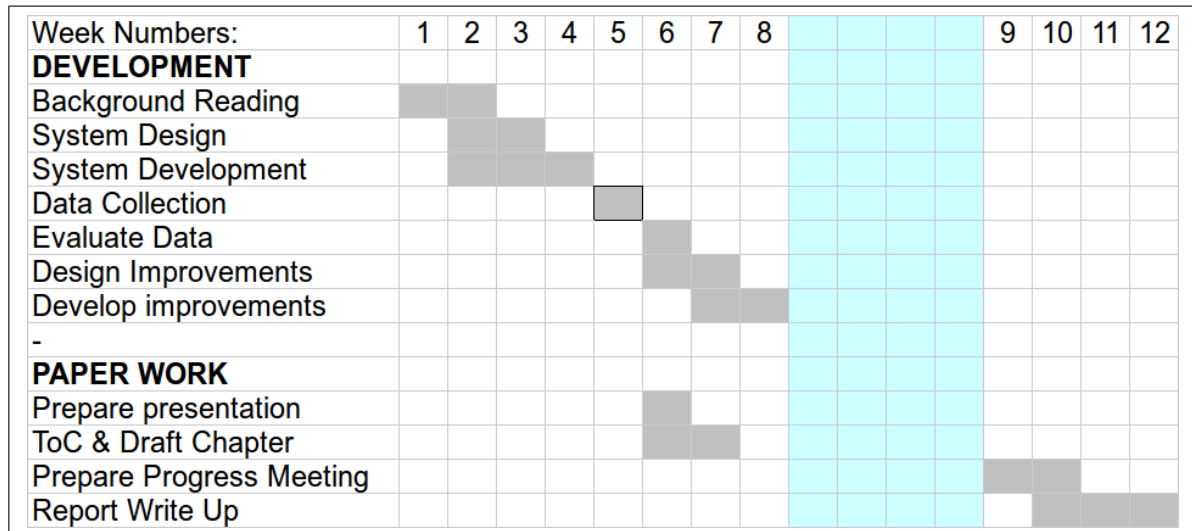


Figure 1.4: Schedule (including Easter Break between weeks 8 and 9)

The project schedule was designed in the form of the Gantt chart shown in Figure 1.4. Dates were based around deliverables set by the project itself and by the availability constraints of the Medical Staff and Student involved in the study.

1.5.1 Milestones

Milestones were set for the project based on key dates in the schedule. They were as follows:

- To have the eye tracking data collection application completely operational and tested by the end of week four.
- To have the presentation material prepared by the end of week six.
- To have a table of contents and Draft Chapter complete by the end of week seven.
- For all of the results needed by the medical collaborators to be complete by the end of week eight.
- To have all work ready for a demonstration by the end of week ten.
- To have a draft of the project report ready by the end of week eleven.

1.5.2 Deadlines

The initial schedule for the first five weeks of the project was extremely tight due to a number of fixed deadlines that needed to be adhered to.

The most prominent of these was for the data collection week, week five. This period was fixed and was the only possible time where both the medical collaborators were free of clinical commitments and The Academy in the LGI was available for use. This meant that the software to be built for interfacing with the eye tracker, displaying the X-Ray images and collecting all the required data, needed to be operational and fully tested by the end of week four. There was no opportunity for a re-run and the consequences of missing this deadline would render the rest of the project almost useless, as there would be no data to evaluate.

The eye tracker resides in the VR lab in the School of Computing and requires keyfob security access. Despite repeated attempts to arrange for this clearance to be given, it was not until the end of week two that the necessary security keyfob was made available. It was not possible to run any EyeLink II related software without access to this room and the eye tracking equipment, causing further development delays.

One of the requirements of the medical study is that the display monitor needed to be a medical grade monitor. This monitor and the accompanying PC was not available until the beginning of week three, leaving only two full weeks for development and testing on the designated hardware for the system.

To best deal with the lack of time available for development, the process was carefully planned including deciding on the bare minimum requirements for the data collection to proceed, and working on these first. To best compensate for this, it was decided that, for at least the first five weeks, project work was carried out seven days a week to try and make up for the lack of time.

Chapter 2

Background Research

The project involved the joining of two separate specialist fields. The radiology study and the eye tracking system. This has also been true for the background research in general, although there is some overlap.

2.1 Radiology Research

Without a medical background, it was found to be useful to have a medical dictionary [15] on hand whilst reading through medical research papers. Although many of the papers in this field do not go into advanced medical details, there has been a multitude of unfamiliar terms where it has been necessary to look up a more understandable definition. The medical dictionary used was recommended by the medical collaborators as a reputable source of definitions and information on medical terminology.

2.1.1 Pilot Study

As previously stated, this project is built upon a pilot study previously completed by the medical collaborators. The literature review and final report [26] for their pilot study was an obvious choice for the first piece of reading relating to radiology. The literature review examined academic research related to the study and the final report described the methodology of the pilot study in detail, allowing for further insight to the research to be done. The references from both of these documents became a hub for me to read through in a type of depth-first-search approach.

2.1.2 Eye tracking in Radiology

Although it seems that there has been no research applying eye tracking techniques to Dual Energy X-Rays, there has been numerous other studies using eye tracking on both Radiologists and Radiographers using medical images such as traditional chest X-Rays or CT scans.

Research has been conducted of the search strategies of radiologists based on their experience level by Manning [14]. This demonstrates the role that training plays in the search strategies of Radiologists. This training however has been shown to be specific to the context of the image viewed, as [17] shows, Radiologists do not seem to possess superior search strategies when searching for targets in other forms of images i.e. cartoons. Other noteworthy researchers in this field include Harold Kundel [6] [7].

Reading through their research has not only given a greater knowledge of the reporting process of X-Ray images, but also how and why eye tracking techniques can play a big part in understanding this process and how one may try to improve or train the Radiologists based on such data.

2.1.3 Meetings

Knowing that the medical collaborators have done previous research in this field, it was decided early on in the project to have a meeting with them to discuss the project, the background reasoning and possible future directions. A list of questions and discussion topics were thought through and prepared in advance. This meeting proved to be extremely helpful in the design and development decisions of the system. The complete study was discussed along with the necessary technicalities relating to what was wanted from such a system, and what was possible to be done with the equipment and the time frame available.

2.2 Eye Tracking Systems

2.2.1 Previous Projects

During 2004/05, Chris Rooney (a third year Computing student at the time) used the same EyeLink II eye tracker for his Final Year Project, where he studied the use of gaze direction as alternative interface tool. The report from his project [21] was an excellent piece of reading material for the technical aspects of using the EyeLink II, on both the software side, and the hardware and calibration side. After thoroughly reading this report, a meeting was arranged with Mr Rooney to demonstrate the work he had done, and to discuss how to approach my own project. Mr Rooney had previously written some C++ and OpenGL applications for the EyeLink II system and very kindly released the source code from his project, without restrictions, as a code base to learn from and build upon for this project.

2.2.2 SR Research

The company that made the EyeLink II eye tracker and accompanying software, SR Research, has an excellent catalogue of manuals and documentation ranging from setting up the software and hardware, developing applications with their supplied API and configuring and calibrating

the device on participants. All this is available on there dedicated support website [11] on which they also host a message forum where they seem to have been extremely diligent with regards to replies from other customer queries and problems. These resources have proved invaluable in correctly setting up and using the EyeLink II in this project.

2.2.3 General Purpose Eye Tracking

Eye tracking has been used to study participants in a broad range of fields, not only Radiology. Some of the current common uses for eye tracking include website usability studies [22], and vehicle driving, including driving simulators [18]. There is an abundance of fields where eye tracking technology has played a part in the understanding of human interaction.

The book ‘Eye Tracking Methodology: Theory and Practice’ by AT Duchowski [3] explains many of the key principles and uses of eye tracking and became a useful resource for general purpose information relating to eye tracking methods.

2.2.4 Eye Tracking Analysis

As mentioned, eye tracking techniques have been used in a massive variety of applications. Despite this variety of subject fields, analysis of eye tracking data is mainly done in two main styles.

The most commonly seen visual analysis for eye tracking is the heat-map. Heat-maps contain the original scene, but are overlaid with a 2D graphic that represents the duration the participant’s gaze was held in areas of the scene. The gaze duration is represented by colour, for example in the bottom image of Figure 2.1, the most intensely studied areas are red, followed by yellow then green. Areas not covered at all remain transparent, showing the underlying image.

Eye movements are not smooth motions, they consist of a series of fixations (points at which eye movement is still for a short time on an area of interest) and saccades (extremely rapid movements between fixations). A commonly used visual analysis to show these fixations and saccades is a scan-path map, an example of which can be seen in the upper part Figure 2.1. Like heat-maps, a scan-path map overlays the original scene. Fixations are plotted as a series of dots with sizes relative to the fixation durations. These fixation dots are then connected, in order, with lines representing the saccades.




Figure 2.1: Example of a standard eye tracking scan-path and heat-map visualisations.

Chapter 3

Data Collection Phase

This chapter describes the first phase of the project, the data collection. It begins with a description of the experiment and moves on to describe the design and development needed in order to collect the data and the experiment process its self. The chapter closes with an explanation of the actual data collected and a discussion covering issues from this phase as a whole.

3.1 EyeLink II System

Figure 3.1 is a photograph of the development set-up in the VR lab during development.

The EyeLink II system is a two PC system. It comprises of a host PC, a display PC, and the EyeLink II headset. The host PC has a special EyeLink PCI expansion card installed which has 3 ports. One port is connected directly to the EyeLink II headset via a long cable. Another is connected to four infra-red markers which are placed on the four corners of the display PC's screen. The final port is an Ethernet connection which connects to the display PC. The host PC runs a software package that controls and monitors the EyeLink II headset and allows the EyeLink system to accessed from the Display PC over a TCP/IP connection though its own C API. The screen and keyboard of the host PC are used for calibration and monitoring of the trials. The host PC is pictured in part [b] of Figure 3.1.

The display PC also has some EyeLink II software installed, mainly the EyeLink development libraries and calibration tools. As previously described, it connects to the host PC through an Ethernet connection and has infra-red markers attached to each corner of the monitor. The display PC is the target for the custom software built though out this project. The initial display PC was the School of Computing PC shown in part [c] of Figure 3.1.

The EyeLink II headset (as shown in Figure 1.2) contains three high speed infra-red cameras. Two of these are on adjustable mounts pointing towards the participant's eyes to track the position of the eyes relative to the participant's head. The other camera is mounted around in the forehead region, pointing forwards. This camera tracks the position of the infra-red

markers on the display screen and hence the position of the head relative to the display screen. Knowing the position of the eyes relative to the head, combined with the position of the head relative to the display screen enables the EyeLink II system to calculate the point of gaze of a participant, even with free head movement. In Figure 3.1, the EyeLink II is shown in-situ to the left of the host PC in part [b].

In order to be able to see all the detail when viewing medical images, high quality monitors are needed. Sufficiently capable monitors are expensive pieces of equipment and are commonly referred to as "medical grade". For accuracy in this study, a PC with medical grade monitor and graphics card from the LGI, was loaned for use as the display PC. The LGI PC, as shown in Figure 3.1 part [a], is a dual monitor system. It has a standard LCD monitor on the left used for general operations and a medical grade monitor on the right for viewing images. This PC was set up as the display PC for the EyeLink II system, including installing the infra red screen markers and installing the EyeLink II libraries and calibration software.

A laptop, as shown in Figure 3.1 part [d], was also used for general development purposes such as; software compilation, data backup, data integrity checks, and analysis.

Parts [b] and [c] of Figure 3.1 are the School of Computing's original EyeLink II set-up. Parts [a], [b], and [d] was the set-up deployed in the Academy at the LGI.




Figure 3.1: The equipment setup in the VR Lab. Showing [a] the LGI's PC and monitors (medical grade monitor on the right of the machine), [b] EyeLink II headset and Host PC, [c] School of Computing Display PC, [d] laptop for miscellaneous development.

3.2 Software

To collect the data reliably and efficiently, it was necessary to consider the details of the experiment process, and also, how to build the most robust and reliable software to collect and store the data. It was also essential to have a good understanding of the analysis that would be run later, to ensure that there is enough raw data of a high enough quality.

3.2.1 Design

The data collection software needed to be carefully designed around the needs of the above experiment. The following is a list of the different data that was decided as being essential to record:

- A unique identifier for each participant.
- The participant's level of expertise.
- The image set number and type being viewed.
- The Participant's eye movements (x, y coordinates in the screen's resolution and timestamps).
- The Participant's mouse clicks (x, y coordinates in the screen's resolution and timestamps).
- The Participant's confidence rating per image viewed.

It was decided to take note of some data that was slightly less essential to the results, but may be useful in later analysis. It was believed that as much data as possible should be recorded. Collecting excess data would not be a problem, but it would be near impossible to reproduce anything that was not recorded without re-running the whole study.

Examples of some such information is: does the participant wear eye-glasses, how many re-calibrations a participant underwent, and notes on any of the feedback or comments that participants gave during or after the studies.

The system would be required to perform several operations, the most essential of which are:

- To display the X-Ray images.
- To run on the LGI's PC and accompanying medical grade monitor.
- To monitor and record eye tracking data.
- To monitor and record mouse click data.

- To record certain participant meta-data.

Some of the recorded data was collected as follows. It was decided that as there would be no more than 26 participants involved in the study, therefore we were able to assign a letter to each participant as their unique identifier. At the beginning of each trail, participants were to be asked to state their level of expertise which was to be recorded as Registrar (Level 1 to 5), Consultant, or Senior Consultant. Following the viewing of each image, the participants were to be asked for a confidence rating in the range 1 to 5, where a rating of 1 meant that they do not know what they were looking at, and a rating of 5 being that they are sure that they have identified all possible nodules and have made no false detections.

To reduce the complexity of the participant's interactions within the trials, some of the data was collected verbally. Confidence ratings and level of expertise was spoken by the participant and recorded on a paper form by the experiment coordinator. The identifying letter of the participant would also be recorded on the same paper form.

Two ways of recording computerised data were identified. One being to design and use a custom file format whereby the display application would retrieve gaze direction data from the eye tracker, monitor mouse clicks, and store this data in the custom file format. Another option is that the EyeLink II software already has the ability to store eye tracking data in an easily accessible manner (in EDF files). After studying the manual for the EyeLink II system [9] [8], an option was identified for complete data recording. The EyeLink II API contains functions allowing one to pass a string from the display application to the EyeLink II library, to be stored with a time stamp, within the EDF files. This would therefore allow for storing extra data such as mouse click coordinates and image names etc.

The EDF file with messages was chosen as the best option for multiple reasons. Less new programming would be involved, which should mean a smaller chance of software bugs. The EDF files already contain all eye tracking data necessary and the message capability would allow for the storage of all other data including mouse click data, image number and type etc.

Another advantage of using the EDF files is that they also store indications of the start and end of both saccades and fixations, as well as the raw samples at up to 500Hz [20]. The measure of fixations and saccades is done automatically by the EyeLink II software, and is based on extensive research carried out by SR Research [10]. Fixations and saccades in this research are as defined by the EyeLink II software.

One of the ethical requirements of this experiment is that participants must be allowed to stop at any point for any reason, whether they wish to have a break, or terminate the experiment. It is therefore important to be able to pause and resume experiments with the software. The approach taken was to develop the software to only display one image at a time and record all data about that image in one EDF file. The software could then be called from a script prompting the user to start or pass running the software for each particular image in the specific set. This would allow for breaks or experiment termination without affecting previously

stored data. Some advantages of using a script to start each viewing would be that multiple scripts could easily be written the ordering and selection of images coded within. It would also allow for re-calibration of the eye tracker, mid-trial, if necessary.

During testing it was noticed that the calibration on the EyeLink II eye tracker can be easily upset for many reasons. This included scenarios such as slippage of the headset even from small head movements. Over a set of 32 images it was thought that the need to recalibrate would arise at some point. It was thought that an indication of calibration would be useful throughout the trials to effectively measure when recalibration was necessary. One approach would be to draw a form of marker on the screen at the point of gaze, before each image is viewed, the participant could then be asked to look at a specific point on screen i.e. the mouse cursor. It should then be fairly easy to gauge the accuracy of the eye tracker's calibration.

3.2.2 EDF Data Files

The EDF files mentioned are of a proprietary binary data format. Using the EDF2ASCII tool supplied with the system, it is possible to convert EDF files to an ASCII text file containing one line per piece of information. This meant that it would be possible to use all the data stored in the EDF files in a convenient way. Once EDF files were converted, all analysis applications could just parse the ASCII files for necessary data.

Each line in the EDF file is time-stamped. Time-stamps are incremented in milliseconds, hence using a sampling rate of 250Hz, data is sampled every 4 milliseconds. The following page shows examples from the converted ASCII files (lines in Mono Space style font are excerpts from the files).

EDF File Snippets

Lines beginning with a number (i.e. 988874) are details of the current gaze location in screen coordinates. The following line shows that at time 988874 the gaze was recorded at the screen coordinates (347.6, 733.5) with a pupil size of 490.0.

```
988874    347.6    733.5    490.0 .C...
```

A line beginning with "MSG" is a custom message inserted into the data by the display application. One of the first message to be stored was the file path of the image being used in the trial. In this example the file "data/image4/4-50.jpg" was used and 988870 indicates a time-stamp of the message. The image details are also recorded elsewhere, but it was thought that adding this detail may prove helpful for future analysis applications.

```
MSG 988870 data/image4/4-50.jpg
```

As stated, saccades and fixations are automatically detected by the EyeLink II software. Lines beginning with SFIX and EFIX denote the start and end, respectively, of a fixation. Lines beginning with SSACC and ESACC denote the start and end, respectively, of a saccade.

The following shows the end of a saccade (this line includes data such as the starting and ending time-stamps of the saccade, the duration, and the screen coordinates at the start and the end). This is immediately followed, as one would expect, by the start of a fixation.

```
988950    365.6    571.6    1354.0 .....
ESACC L  988794 988950 160    384.4    575.8    365.6    571.6    0.43    589
SFIX L   988954
988954    357.6    576.0    1354.0 .....
```

The following shows the end of a fixation (this line includes data such as the starting and ending time-stamps of the fixation, the duration, and the average screen coordinates for the fixation). This is immediately follow by the start of a saccade.

```
989078    376.4    583.8    1404.0 .....
EFIX L   988954 989078 128    367.2    583.7    1378
SSACC L  989082
989082    375.4    584.2    1404.0 .....
```

The mouse click data was also recorded in the the EDF files in the following format. This example shows a mouse click at time 1009874 with screen coordinates (421, 517)

```
MSG 1009874 mouse click. x: 421 y: 517 Ex:422.80 Ey:525.80
```

After each trial, all EDF files were converted to this ASCII format to ensure usability. From here on in the report, the ASCII text files generated by the ASCII2EDF tool shall be referred to as EDF files.

3.2.3 Code Base

It was decided early on in the project that the most appropriate programming tools for this piece of software were C++ with OpenGL graphics. This decision was based on several factors. The system would benefit from a fast performing language to reduce the possibility of latency issues considering the EyeLink II system can operate at data collection rates of up to 500Hz. Code to interface the eye tracker had already been written in these languages which was available to build upon, including libraries and APIs. There was also a number of people within the research group with vast experience and knowledge of C++ and OpenGL, so, should there be any software issues, help was readily available.

Interfacing with the EyeLink II's system is made possible through C API provided by SR Research. This API has already been used extensively by Chris Rooney in his Final Year Project [21]. Mr Rooney has kindly open sourced his software for use on this project. His code, along with the detailed project report, became a valuable asset in reducing the learning curve and development time of this software. Code from the NeHe OpenGL Tutorials [16] was used as a skeleton structure for the C++ and OpenGL code, due to the fact that it was fairly well documented and ran on the PC from the LGI.

3.2.4 Testing

For each milestone in the software writing process the software would be rigorously tested for both software bugs and compliance with the design specification. For example, the software would be given unexpected arguments and monitored to check that it responded appropriately. When possible, the source code was also inspected by others with expertise in C++ and OpenGL. When EDF files were generated, these were opened and examined to make certain that the data we expected to be stored would be available and correct.

Whilst developing the system, a conscious effort was also made to test the physical side of the system, mainly the eye tracker's calibration. Physical testing was initially carried out by testing and configuring the system on myself repeatedly. Later, volunteers from within the department were used to trial the system. This not only gave me practice in set-up and calibration on other people, but it also made me aware fairly early on that calibration issues may cause problems. Good calibration relies on many different factors, any one of which may easily have adverse effects on the data gathered.

Once the system was operational, the Medics involved in the project were invited to use the system. The system was explained and demonstrated to them and their feedback, which shall be discussed later, was able to be implemented into the final product.

3.2.5 Re-adjustments

It is important that the interface for the participants be as simple and natural to use as possible, in order to not distract from the main task of reporting on the X-rays. One observation during testing of the software was that the initial design required the participants to use both a keyboard (pressing space to start each trial) and a mouse (used to click on the suspected nodules). It was suggested that it may be more natural to discard the keyboard interaction and use only the mouse. Right click to start each trial and left click to identify suspected nodules.

This feature was quickly implemented and tested. Feedback was very positive and it seemed like the appropriate choice for this system. The user now only has to interact with the mouse and its two buttons. One unforeseen problem with this feature will be discussed later.

3.3 Experiment

3.3.1 Experiment Design

In the case of this project, the experiment was largely designed by the medical collaborators and based upon the pilot study that they had carried out in 2009 [26]. This design was discussed in person and the details were noted for use in the software design.

The basic overview of the experiment is for a pool of Radiologists to each look through a series of different images and search for lung nodules. For each image, the Radiologists should use the mouse to click on any suspected nodules. After each image they should give a rating indicating how confident they feel about their diagnosis. Their eye movements shall be tracked throughout the experiment for later analysis.

The initial experiment was based around a dataset of 16 Dual Energy Chest X-Rays each having 5 types of images associated with them: a soft tissue image, boney image overlays at 10%, 20% and 50% opacity, and a traditional Chest X-Ray style image. The Chest X-Rays vary in the type and count of the nodules placed in them from an image with no nodules, up to an image with 4 nodules. Each participant in the study was to look through 8 of each of the 5 image types, all presented in a random order.

The Chest X-Rays taken were of a lung phantom, meaning that the only difference between the 16 X-Rays sets, was the quantity and location of the nodules within the lungs. All other anatomical features were identical across all X-Rays.

It was suggested that the number of types of images from each X-Ray could be reduced. The visual difference between the 10% and 20% overlaid images was so small that it was almost impossible to notice. Discarding this 10% images could remove an extra variable from the analysis stage that added extra complication and little gain.

Originally participants would view the same image more than once, which although they would be of different types, would have the potential to warp results due to a memory of the

nodule locations. If a participant recognised that there is a nodule in the same position as there was in a previous image, there is a potential that they would recall the locations of other nodules in that previous image which may gain them an unfair advantage in the discovery of further nodules in the current image.

This could be overcome by reducing the number of images viewed by each participant to 16, or increasing the size of the dataset by adding extra Chest X-Rays. Fortunately the use of the lung phantom meant that it would be possible to split the images into a left side and right side and easily piece them together, mixing the different X-Ray image sets to produce new image sets. For example the left side of soft tissue image 3 could be added to the right side of soft tissue image 12 to give a new image, soft tissue image 17. This could easily be done with one of many image manipulation tools available.

After careful analysis and much discussion of these matters, it was decided that the following two amendments to the experiment should be made. The 10% image type was discarded from the dataset, slightly reducing the unnecessary complexity of the results.

Reducing the number of images viewed by each participant would mean that there would be a need for many more participants than previously planned for. Knowing that Radiologists are extremely busy people, and being unsure exactly how many Radiologist may be available to volunteer for the experiments, it was thought that it would be better to require fewer participants for statistical significance and be sure that we would have enough participants, than to reduce the image sets and require more participants than we knew to be available.

It was therefore decided that more X-Ray images were needed. Using a random selection process, a total of 16 new images were created with the method described above. Each participant would now be asked to look through 8 of each of the 4 types of images, totalling 32 images. Each of these images would be unique in terms of nodule quantity and distribution.

Figure 3.2 show an example of each of the four types of images used in the study. They range from a soft tissue only image through to a normal image, which is equivalent to a traditional chest X-Ray. Although the differences may be hard to notice in print, when these images were displayed on the medical grade monitor, the difference visually is a steady increase in the bone overlay from the soft tissue image through to the normal image.

3.3.2 Deployment Preparation

During development, a School of Computing PC was used to compile all the software. This machine would not be easily accessible during the data collection in the LGI, therefore the decision was made to set-up the same build environment on an available laptop and compile the software from there. This was merely a precaution against the possibility of needing to alter any software once the system was in place at the LGI. Fortunately this was not necessary to use.

Once the system had been developed and tested to what was thought of as an operational




Figure 3.2: Examples of the four types of X-Ray images used in the study, ranging from a soft tissue image (top left), to a normal chest X-Ray image (bottom right).

standard, the hardware was disassembled, ready to be transported to the LGI. During disassembly, careful notes were taken of connections and set-up details between all computers and the EyeLink II hardware to aid in reassembly later.

3.3.3 Deployment

Week five of the project had been designated as the data collection week. No trials were run on the first day for the following reasons. The morning consisted of setting up the hardware in the location provided (the Training Room, in The Academy, at the LGI). This was then tested to ensure that it still functioned as expected. Data recorded in these tests was also examined to confirm its correctness.

Once we were happy with the installation of the system, we started advertising our presence to radiologists in The Academy. Colour flyers describing the study were distributed throughout the department and wherever possible, we spoke with radiologists to explain our project. We also spoke with some of the more senior members of staff in the radiology department who were extremely helpful, for example, by emailing the staff mailing lists, describing our study and requesting participation.

3.3.4 Data Collection

During the week, a total of 14 radiologist kindly volunteered their time for the study. Data collected was validated and backed up after each participant in the study to ensure integrity.

At times when no radiologist were available, details of the analysis that was to be performed on the collected data was discussed. Each evening was then spent prototyping these ideas ready for the following morning so that the ideas could then be refined even more. Another benefit these prototypes provided was allowing a quick means of confirming the integrity of the data collected, shortly after a trial was completed. It was also useful to be able to show radiologists an analysis and visualisation of their data within minutes of them completing the study. Their feedback was helpful in creating further analysis.

3.4 Evaluation

Although the Data Collection Phase was successfully execute, there were small issues that needed to be dealt with. There were many issues in the development stage of this system that could easily have lead to over-running of the development and hence jeopardising the whole project. Small problems when at the LGI also arose and had to be dealt with quickly.

This section will discuss some of these issues and describe how the project was run to best deal with them.

3.4.1 Communication with L.G.I.

In general, communication with the medical collaborators was very good. A meeting was arranged very early on to discuss the details of the project, and plan out the requirements of the data collection system. A list of questions and discussion topics were prepared beforehand and the meeting was very productive.

Unfortunately, a few small unforeseen decisions arose during development (as discussed later), and there was a short delay in communication from the collaborators. This may have been avoided, had there been pre-arranged weekly meetings with the medical collaborators during the development of the software. Initially it was thought that the first meeting would provide almost all of the details of the specification and that any other minor details would be easily resolved. In hindsight, regular progress meetings and defined lines of communication would have been beneficial and allowed for a lightly smoother development phase.

3.4.2 Specification Issues

One issue that arose during development was the fact that the medical grade monitor has a rectangular screen, whereas the X-Ray images that were provided for testing purposes were of a much larger resolution (2021 x 2021) and square. This brought about questions of whether to crop the X-Ray images, to distort them or to keep their aspect ratio and shrink them. This issue was believed to be too critical to guess the ideal feature without consulting the medical professionals. Unable to get a response on this issue immediately, it was decided to write the software to display each image with a predefined scaling and transformation factor. This meant that as soon as the correct display process was able to be determined, all that was needed was to set the appropriate scaling factors in the code.

Another issue that had been overlooked was to do with the displaying of sets of images. It became apparent that it was unclear how many images each participant would be viewing and therefore the system had been built initially to display one X-Ray at a time and collect the relevant data. This program is started from the command line with the image name given as arguments. If there were to be a lot of images for each trial this would result in a lot of typing of file names. A solution was built with Windows batch scripts where by the whole set of images was pre-programmed to be executed as arguments to the program individually and the only interactions needed were to press enter to display the image, or "n" to decline the image. Another advantage of this approach was that the system could be stopped and resumed easily. This would prove extremely helpful for re-calibration and situations where a participant may want to stop halfway through, and potentially resume at a later time.

3.4.3 Unfamiliar Language

The most appropriate programming tools for this project were decided on as C++ using OpenGL for the graphics. Almost all of my programming experience has been with the Python programming language, along with some projects in Java and C.

This presented a great programming learning curve before development could even begin, although, considering myself as a proficient python programmer, it was believed this task would be more about learning the syntax and nuances of the language than having to re-learn how to program.

There are many resources available to aid in learning programming languages, after some research into what was available, including speaking with peers who already knew the language, the online C++ tutorials at Cprogramming.com [1] were chosen as the most appropriate. They present a logical flow of tutorials, complete with clear and concise descriptions and source code examples.

I was already a little familiar with OpenGL using Python from a previous university module in computer graphics I had taken. Following some recommendations, some of the NeHe OpenGL Tutorials [16] were studied, to help with familiarisation of the use of OpenGL in C++.

Another invaluable resource was the FYP Report and accompanying software from Mr Rooney's FYP [21]. This software provided 3 working examples of OpenGL programs written in C++, that also interface with the EyeLink II's API. With permission, one of these programs was stripped down and that became the initial base code for my system. It was thought that this would help compensate for the lack of time available for development in the project.

3.4.4 Cross Platform Issues

Due to unavailability of equipment, the initial system was developed on a standard School of Computing workstation. As soon as access was granted to use the VR lab, the system was moved to the lab and the eye tracking interface software was developed and implemented. This transition went reasonably smoothly.

The target machine for this system was the PC loaned from the LGI, due to its medical grade monitor. This PC ran the Windows 2000 OS, which was not thought to be an issue. The development PCs used where all running Windows XP, and it was thought that most software written for Windows XP would be backwards compatible with Windows 2000.

Once the LGI PC was installed in the VR lab, it was discovered that this was not the case. The software did not run at all and instead, only displayed a rather vague error message. After multiple failed attempts to fix this issue, it was decided that other graphics applications should be tested on this platform to find a combination of C++ and OpenGL code that would execute properly on this machine. A collection of programs were downloaded that used OpenGL with C++ and had freely available source code. After some testing and research, it became apparent

that OpenGL's GLUT Utility Toolkit was the problem.

Of the programs tested, the NeHe IPicture example [16], which used standard Windows system calls instead of GLUT, seemed to both work the smoothest and have the most usable code. The source code to this example was downloaded and compiled locally to ensure that it was possible to modify and use the code. This executed successfully on the LGI PC, therefore the logic from the previous application was added to a modified version of the IPicture code, to create an application executable on the Windows 2000 platform.

3.4.5 Eye Tracking Calibration

Before each trial with the EyeLink II, the system needs to be calibrated to that user. The first part of this process involves adjusting the headset, positioning and focusing the headset's cameras. The participant is then asked to follow, with their eyes, a series of nine dots around the screen to acquire their gaze direction. This process is repeated with another nine dots and the EyeLink II software compares the two sets of data to see how well they correlate. The software then returns a result of either GOOD, AVERAGE or BAD, relating to how well two calibrations match. Calibration was considered successful and accepted only if the system reported a result of GOOD.

As elementary as this process seems, it requires extreme precision and once a user has been calibrated to a level reported as GOOD, there can still be discrepancies between the user's actual gaze direction and where the EyeLink II reports the user's gaze direction. Also, even if the calibration is of a good accuracy, it can be easily disturbed by one of many factors.

The realisation of the calibration difficulties lead to an in depth study of the documentation provided by SR Research for calibration of the EyeLink II [9] [8]. SR Research also run a members only online support forum [11]. This online community was joined very early on in the project and the archived material available from other researcher's problems and solutions proved extremely useful in gaining a full understanding of the usage and calibration methods of the EyeLink II system.

One of the main factors that was found to introduce calibration errors, is any movement of the positioning of the headset. Once the headset is placed on a user it is tightened to the user's head with an adjustable head band. This presented a trade-off between a secure fit on the users head and a comfortable tension on the headset for the user. During trials it was found that the tension needed to secure the headset against slippage was uncomfortably tight for some users. In an effort to ease the pressure whilst maintaining a secure fitment, extra foam padding was sourced and added to the parts of the headband that seemed to cause discomfort. This seemed to aid in the comfort of the headset, which was considered especially important due to the fact that participants would need to wear the equipment for periods exceeding 30 minutes.

Even with a secure fitment of the headset, calibration could still be easily affected. The cable between the headset and the PC hangs from the back of the headset. Even just the weight

of this cable moving around, seemed to be enough to drastically affect calibration accuracy. To best resolve this issue, the slack in the cable was held in an elevated position by one of the experimenters during each trial. This seemed to minimise the effect of its weight and also stopped it getting trapped in the chair or on clothing which would again pull on the headset.

It was also noted that the calibration may be compromised in a similar way due to things such as rapid head movements, or even facial expressions of the user (This was particularly a problem with participants wearing glasses). Issues like these are hard to avoid and the best that was thought to be done was to ask participants to try not to use rapid head movements. Careful attention was paid to calibration during the trials and if it was thought that there was a problem, the experiments were paused, the initial calibration process was re-run, and once GOOD calibration was gained, the experiment would continue.

The positioning of the cameras on the headset is extremely important for proper calibration. The full eye must be in the camera's view and properly focused, whilst also considering that it is not obstructing the field of view of the user. With practice and careful adjustment, this was not too difficult to achieve, although it proved considerably harder in participants wearing glasses due to the optics of the lenses and the obstruction of the frames of the glasses. Volunteers wearing glasses were used during development and testing to practice the set up of such participants. In general, camera positions were only affected if they were knocked accidentally, i.e. by a user rubbing their eyes or face. In this situation, it was decided best to reposition the cameras and repeat a full calibration before continuing, to ensure accuracy.

Another slightly more subtle issue that affects calibration of the EyeLink II is the size of a participant's pupil. SR Research have reported that fluctuations in the size of the pupil may introduce errors in the calibration of the system [9]. This may be caused by changes in ambient lighting, or large changes in the overall luminosity of the images appearing on screen.

When reporting on X-Ray images, Radiologists do so in a very dimly lit room. To best try and avoid variations in pupil size during the experiments, the lights were consistently dimmed for both the calibration and for the viewing periods.

The background colour of the display screen during calibration was found to be adjustable. A shade of grey that was very similar in luminosity to the X-Ray images being viewed was chosen as the new calibration background in an attempt to minimise the variation in pupil sizes of the participants.

3.4.6 False Conditions

It should be noted that although every effort was made for the conditions of the study for the Radiologists to be as close to their usual working environment as possible, there were some inevitable false conditions. This was confirmed by feedback given by the Radiologists after their trials. Some of these false conditions are outlined in this section.

When viewing images in a normal working environment, Radiologists use tools referred to

as "Windowing" tools. These are features such as zooming, panning, brightness and contrast adjustments. Windowing tools may allow for better detection of abnormalities within images, but the inclusion of such features would greatly over complicate the analysis of results in the study and therefore, although these tools were initially developed for the software, it was decided to disable them in the final implementation, to ensure a coherent set of results.

It is also unnatural for radiologist to scan through so many sets of images in one sitting. During a normal work flow they may view a handful of images in a row and then move on to another task, to return later to view more.

The addition of the EyeLink II headset presented more false conditions. Radiologists do not normally wear uncomfortable headsets, and many tend to prefer to move around whilst viewing images i.e. moving physically closer to screen upon noticing a potential nodule for a clearer view. Although the EyeLink II allows for head movement in participants, it is still somewhat restrictive itself, coupled with the fact that participants were made aware that sudden head movements may affect calibration adversely.

The decisions made by Radiologists whilst viewing images in their normal working environment may have great consequences on real peoples lives i.e. whether to diagnose a patient with cancer or not. The participants in this study were made aware of the fact that the images were of a lung phantom, not a real person, and that all data in the study was completely anonymised and not used in any way to judge their individual performance. This may have had a psychological effect (either consciously or not) of reducing a participant's effort in reporting on the images.

During the study, each trial (including its participants) were carefully observed by at least 2 experimenters. For some participants, this may have added an extra feeling of pressure, and possibly affected their performance.

Although any of the above may have adversely affected the results, they were conditions that were unpreventable, and considered to have not impacted on the legitimacy of the results gained. It is, however, worth noting that these conditions were present.

3.4.7 Uniqueness of Human Participants

The ideal participant for use with the EyeLink II system would; sit extremely still, not mind their head being clamped tightly, not wear glasses, have large eyes, have a perfectly symmetrical head and face, and have perfect vision. Unfortunately people are not quite so uniform and this made some participants easier and quicker to calibrate than others.

Most calibration attempts were successful, although one out of the fourteen participants who volunteered for the study proved extremely difficult to get an accurate calibration. After multiple failed attempts at recalibration and readjustments of the headset, it was decided to thank the participant for volunteering and terminate the trial. This was decided as we did not want to waste the Radiologist's time.

3.4.8 Fatigue of Participants

As mentioned, the comfort of the headset was an issue with some participants. Also, the total time participants were involved in the study varied from 30 minutes to over 90 minutes. These long sittings, coupled with the discomfort of the headset eventually became unbearable for some of the participants, leading to more than half of the participants in the study needing to take at least one break. Fortunately the software had been designed to allow for a trial to be paused and potentially continued at a later time.

Participant comfort was considered an important factor during the study. Before they began, participants were informed that they may stop at any time, they did not need to give a reason and they did not have to continue if they did not want to. Participants were also asked at regular occasions throughout the trials if they were comfortably and if they would like to pause for a break. Although the majority took at least one break, all 13 participants completed their trials and non asked for the experiment to be terminated.

3.4.9 Unforeseen Issues

As mentioned, to aid in the ease of interaction, the right mouse button was used by the participant to begin each trial. According to participant feedback, this proved to be an intuitive interaction tool.

Between trials, there was a short delay of approximately 5 seconds from the time that a trial finish, and the time that the next trial was ready to be started. This delay was due to a combination of the previous trial's data being saved to disk and the following trial loading.

Participants were given the instruction "When you see the black screen, right click to begin the trial.". Once used to this routine, one of the participants would often prematurely right click on the Windows desktop before the application had fully loaded causing the desktop properties menu to be display.

The problem being that if this menu was still open when the X-Ray application was fully loaded, the menu would persist on the screen, but the mouse focus would switch to the X-Ray application. Therefore this obscured the image and there was no way to remove the desktop properties menu, without exiting the X-Ray application.

Each time this happened, rapid action was taken by the experiment coordinator to exit the X-Ray application, cancel the offending menu and restart the X-Ray application at the appropriate trial.

This scenario only happened with one of the participants. Had it happened during testing, it would have been marked as a design flaw of the system and rectified.

After this participant's trials were completed, ideas to prevent this from happening again were discussed. It was decided that, as there seemed to be no way to disable the offending menu, it could potentially be remedied by opening a fullscreen application in the background that did

not respond to right click input. This idea was tested and seemed to be a good, temporary solution. Subsequent trials were then run in this way, although it proved to be unnecessary.

Other possible solutions to this problem might include; researching possible ways to disable this desktop menu, developing a dedicated cover application for the desktop or choosing another means of starting the trials.

3.4.10 Phantom Issues

Although the use of the lung phantom provided many benefits to this study, it also may have presented a few issues.

Despite being extremely realistic, X-Rays of the phantom contain a thin black horizontal line just below the heart region. Feedback from participants suggested that this abnormality may have confused, disoriented, or disrupted their search strategies.

Research by Krupinski [5] has shown that it is common for healthy features in the lungs, such as vessels, to be mistake for small lung nodules, and vice-versa. This is backed-up further by research, showing that only nodules above 3mm in diameter are visible, but are not accurately detected until they are over 8 - 10mm in diameter [6].

Feedback from Radiologist in this study was consistent with these findings. Many participants reported that they identified some features of the lungs as nodules, particularly around the hilar regions of the lungs. Only once they had seen a few images, did they realise that these were just normal vessels in the lungs of that phantom and proceeded to ignore those areas. Essentially participants were becoming accustomed to the physical anatomy of the phantom, which in turn, aided their search as they progressed through the trials.

Chapter 4

Analysis Phase

This chapter describes the second phase of the project, the analysis phase. Once all the experiments at the LGI were complete, work began on building software to analyse the data collected. The initial analysis performed was based on the results needed by the medical collaborators including, image coverage and nodule detection accuracy. Once these results were completed, software for real time replays of the data was built in order to gain insight into other more advanced types of analysis that could be performed. From these, further results and visualisations were done.

More than 40 different variations of data analysis were developed and explored during this phase of the project including both statistics and visualisations. This chapter describes the most notable of the completed analyses.

4.1 Planning

4.1.1 Tools

The majority of the software built for analysis was written in the Python programming language. Python was chosen for many reasons of which the following are just a few.

- Python is widely regarded for its short development times for application prototyping.
- There are a wealth of libraries freely available. These range from data and image manipulation, through to graphing and visualisations.
- It is the programming language that I am most familiar and comfortable with.

Using Python meant that time would not be wasted unnecessarily on application development. This is especially important when prototyping many ideas and concepts.

As discussed later, some other tools were used. Matlab was used for one of the visualisation prototypes and the SPSS software package was used to statistically analyse some of the results gained.

It became clear that, for almost all analysis applications, there would be a need to parse the EDF files and also to generate X-Ray images for each trial with certain features overlaid. For this reason it was decided that one of the first development task should be to create generic libraries of code that could quickly and easily be re-used in each application.

This was split into two libraries:

- Xreye, a parsing library that, when given the ASCII converted version of the EDF files, would allow access to all the data in an easy to use Object Oriented manner. Using this library now eases the process of accessing; mouse click data, fixation data, the name of the image used etc.
- Eyemage, basically provides a layer of abstraction on top of PIL (the Python Imaging Library). It contains methods to do many of the common tasks that were needed including; appropriate resizing of images, drawing of coloured circles and crosses at given screen coordinates etc.

4.1.2 Data Correction

Testing of volunteers during the development phase showed an average accuracy of up to two inches on screen, between a user's actual gaze direction and that of which the system reports even immediately after a GOOD calibration had been attained.

Figure 4.1 demonstrates two occasions taken from the study which show the effect of the calibration issues. These images have blue circles overlaid at the point of gaze for each fixation during the trial. Considering that the participants were asked to only examine the lungs, it follows that the shape of all the fixations for a trial should be roughly the shape of the lungs.

For the left image, it can be seen that there is clear a lungs shape in the fixations, although they are offset slightly to the left and up a little.

For the right image, the offset is more severe and the lung shape is slightly less apparent. It was decided that certain trials should be discarded due to excessive calibration errors. The criteria for discarding an image was; if the centre point of any of the fixations was outside of the space of the screen coordinates, that trial was considered too erroneous and discarded from the results.

With the assumption that the offset for these errors in calibration are the same within each trial, methods of data correction may be possible. Firstly, and rather crudely, it may be possible to set the X and Y offset variables by hand, moving the fixation data until it looks like it is in the correct location. This approach was discarded as it may be inconsistent over all trials, and is also highly unscientific.

Another method devised for data correction was to take the gaze locations at the time of a mouse click. Mouse click data is completely accurate, and does not suffer from the same errors. Assuming that a participant will be looking at the mouse cursor at the point of a click, and




Figure 4.1: Two examples of data with calibration errors.

possibly just preceding a click, the difference between the gaze screen coordinates and the click screen coordinates should be a very good indication of the offset error of the calibration. If these offsets for each click in a trial are averaged, the result could potentially be used to correct all the fixations within that image. One limitation to this method would be the fact there are a small percentage of the trials where a participant did not click at all, believing the image to have no nodules within it.

After careful thought relating to the need for correcting the data compared with the time available in the project, it was decided that this should be looked at later. More important analysis could be done without the need for this data correction and if time was available at the end of the project, this could be developed. Unfortunately, there was not enough time to completely develop this idea.

4.1.3 Nodule Locations

In order to calculate the accuracy of a participant two sets of data needed to be analysed, mouse click coordinates and nodule locations. Mouse click coordinates in screen resolution was available in the EDF files previously discussed. Nodule locations were available but this was in the form of a rough sketch of the chest, made at the time the X-Rays of the phantom were

taken. Details of the size of the nodules was also a vague measure. Nodules were categorised as either small medium or large. The nodule type was clearly defined as either calcified or non-calcified.




Figure 4.2: Example of the problems encountered whilst trying to attain accurate coordinates of the nodules from the given data.

Precise locations of the nodules were needed in order to perform proper analysis. The left hand image in Figure 4.2 shows an example of one of the sketches provided for the X-Ray shown in the center of Figure 4.2. Using the relevant sketch and the original 2021x2021 X-Ray images, exact coordinates of the center nodules were found by hand. The process was done using the GIMP image manipulation software and was aided by using zooming, brightness and contrast tools.

The most important aid in find these nodules was that fact that there was a "ground truth" X-Ray image. One of the X-Ray sets (set 16) contained no nodules. As all images were taken of the same phantom in roughly the same position, it proved invaluable for the more inconspicuous nodules, to compare the current image with an image from set 16.

Six nodules were used in the creation of the images. This comprised of a small, a medium and a large for both calcified and non-calcified nodules. From the X-Ray images, the diameters of these nodules in pixels were noted for use later.

An application using the Eyemage library was built to clarify the nodule locations, types and size were correct for every image. The output for one of the images can be seen in Figure 4.3. The left image is the original X-Ray, and the right is the same X-Ray, but with the nodule locations from the data previously describe, clearly marked by circles. The red circles are non-calcified nodules and the white circles are calcified nodules.




Figure 4.3: Showing the overlay of the nodule locations for an image. White dots are calcified nodules and red dots are non-calcified nodules.

4.2 Required Results

Although many different ways to analyse the data were identified, there were specific results required by the medical collaborators in order to to assess their part of the study. It was decided that it was important to work on these sets of analysis first and use the remaining project time to expand into other, more novel analysis methods. These initial requirements are outlined in this section.

4.2.1 Coverage

For the medical collaborators, the most desired analysis of the data was the coverage of the image being viewed. Coverage of an image describes how much of that image was scanned by the participant during a trial. A metric was devised for this as follows.

Firstly, it was decided that only fixations were to be used when calculating coverage as saccades were considered too fast a movement for anyone to be able to take in information from them. For each trail, fixations were plotted on the image viewed in the form of blue circles (as was seen in Figure 4.1). The diameter of these fixation circles as follows. The most sensitive part of the eye's field of view, known as the fovea, has been shown by many including Kundel [7]

to be approximately two degrees. Following information from Radiological viewing guidelines [19], stating a viewing distance for optimal comfort of between 51cm and 75cm, participants were asked to sit a distance of approximately 60cm from the screen. From this we identified the region of the screen where the eyes were most sensitive for each fixation as 2cm in diameter. This equated to 80 pixels in diameter.

The coverage maps produced were useful visually but a quantitative measure was needed for proper analysis. It was decided that the quantity of blue pixels in each image was an appropriate measure. A program was written to count the number of blue pixels in each image. This now gave a quantitative measure for each trial. It was also ensured that the coverage maps were saved as PNG images, due to the fact that this is a lossless format, therefore no blue pixels were missed on the basis of them not being exactly pure blue (0, 0, 255).

As mentioned previously, it was decided that some of the trials needed to be discarded due to excessive calibration errors. From a total of 414 trials with 13 Radiologist, 376 trials were considered usable, meaning less than 10% were discarded and deemed unusable for coverage analysis.

Coverage values were generated in CSV (Comma Separated Values) format for all trials to allow for importation into analysis software including Microsoft Excel/ Open Office Calc and SPSS.

4.2.2 Accuracy

As mentioned, accompanying the eye tracking data, Radiologists were asked to click, with the mouse, on where they thought nodules were located within the image. No distinction between calcified and non calcified nodules were made for this task. Reasonable competency with the use of a conventional computer mouse was assumed here.

The main point of interest from the mouse click data was to show the accuracy of Radiologists in detecting nodules across the four image types. There are a number of ways to analyse this accuracy data.

- The percentage of the nodules correctly detected from the nodules that were in the image.
- The number of False Positive detections in the trial (False Positives are classed as a point that a participant has clicked on the image where there is no actual nodule).
- Combinations of the above.

As displayed in Figure 4.3, nodule locations and diameters are known. This software was expanded to display the location of mouse clicks. Mouse clicks were added as green crosses as can be seen in Figure 4.4. These images now clearly show where a participant has correctly identified a nodule, where they have missed a nodule, and where they have made false detections in areas containing no nodules.




Figure 4.4: Two examples of images overlaid with nodule locations and type, and mouse click locations.

As with the coverage analysis, these images provide an excellent visual representation of the data in the individual trials, but for proper analysis, quantitative measures were needed.

In order to identify if a click was on target for a nodule or not, it was first necessary to identify a criteria for this. As previously noted, reasonable competency with the use of a conventional computer mouse was assumed. It was believed likely there may be instances where the mouse coordinates of a click lay just outside of a nodule's circumference, therefore if the analysis dictates that only clicks within a nodule are classed as on target, such clicks would falsely be regarded as false positives, and the nodule as undetected.

In order to clarify the accuracy of these mouse clicks a program was written to iterate over every mouse click in all the trials and report the distance to the centre of nearest nodule. This was normalised for the size of the nodules by calculating the distance divided by the radius of the nodule. When plotted on a graph, it became clear that the overwhelming majority of the clicks were within 1 radius from the centre of their nearest nodule i.e. the nodule had been clicked on. There were a few outlying clicks slightly further away, and then everything above a $1.5 \times$ radius distance from the centre of the nodule appeared to be uniformly random, hence false positives.

From this analysis it was decided that clicks within a distance of $(1.5 \times \text{nodule radius})$

from their closest nodule were to be considered a correctly identified nodule and clicks with a distance of (1.5 x nodule radius) or above from their closest nodule were to be considered false positives.

Using this metric, a program was written to process all clicks in all trials and report the results in CSV format for later analysis.

4.2.3 Confidence Rating

No processing was needed for the confidence ratings of each trial. This was recorded as a number from one to five on a paper form during the trials, therefore all that was needed was for it to be typed up into a spreadsheet for analysis.

4.2.4 Results

The data generated from the previously mentioned processing programs was collated into a spreadsheet ready for statistical analysis. A screenshot of this is shown in Figure 4.5

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
159	E	32	5	n	14.012	5	160384	1	1	2	1	0	0	
160	F	1	6	20	35.328	4	333871	3	2	0	2	0	1	
161	F	2	18	20	53.368	3	302430	0	0	4	0	0	0	
162	F	3	23	50	38.612	4	246224	2	1	2	1	1	0	
163	F	4	22	20	87.304	3	391080	2	1	2	0	2	0	
164	F	5	10	20	57.08	4	329323	3	2	2	2	1	0	
165	F	6	2	20	61.008	3	313227	3	0	3	0	2	1	
166	F	7	24	n	40.32	5	302030	2	1	2	1	1	0	
167	F	8	32	n	49.92	4	226659	5	2	2	2	2	1	
168	F	9	3	50	38.26	4	190939	1	0	3	0	1	0	
169	F	10	4	n	38.852	4	214984	2	1	2	1	1	0	
170	F	11	9	s	50.024	4	246246	0	1	2	0	0	0	
171	F	12	7	50	65.34	4	305392	3	1	2	1	2	0	
172	F	13	13	s	38.956	4	244322	1	0	1	0	1	0	
173	F	14	19	50	38.368	4	219357	3	2	2	2	1	0	
174	F	15	14	20	49.444	3	219592	1	1	2	1	0	0	
175	F	16	16	n	39.312	4	230804	0	0	0	0	0	0	
176	F	17	5	s	47.308	3	231600	1	1	2	1	0	0	
177	F	18	26	20	59.928	4	255543	2	1	2	1	1	0	
178	F	19	30	20	50.524	3	193018	2	0	2	0	1	1	
179	F	20	8	n	47.64	4	257971	5	2	1	2	1	2	

Figure 4.5: A screenshot of the data after being imported into Open Office. Columns are; Participant ID, Image order, Image number, Image type, Duration, Confidence Rating, Coverage (1920000 Max), Total Clicks, Number of Calcified Nodules, Number of Non-calcified Nodules, Correctly Identified Calcified Nodules, Correctly Identified Non-calcified Nodules, False Positives.

The data generated from the previously mentioned processing programs, relating to image coverage and trial duration, was analysed with the SPSS statistical analysis package as follows.

As mentioned, data from some of the trials where calibration problems had caused errors determined as unusable, was discarded. After this, from three to eight trials (average 7.2 trials) that each participant performed with a given bone intensity remained. The data for each participants non-discarded trials with each bone intensity were averaged, and analysed using repeated measures analyses of variance (ANOVAs) that treated bone intensity as a within participants factor. In the following analyses [†] indicates that the Greenhouse-Geisser sphericity correction was applied.

ANOVAs showed that participants covered significantly more of the image as the bone intensity increased ($F(2, 21) = 3.58^{\dagger}$, $p = .05$) (see Figure 4.6), but bone intensity did not affect the time that participants took to make diagnoses ($F(2, 21) = 0.81^{\dagger}$, $p = .44$) (see Figure 4.6). It follows that the speed at which participants covered the images was also significantly higher with increased bone intensities ($F(2, 23) = 4.15^{\dagger}$, $p = .03$) (see Figure 4.6). There was not a significant difference in the confidence of participants diagnoses with different bone intensities ($F(2, 21) = 2.74^{\dagger}$, $p = .09$) (see Figure 4.6).

For the following four graphs, boneIntensity values relates to the the image types. 1, 2, 3 and 4 relate respectively to Soft Tissue, 20% Overlay, 50% Overlay and Normal Image.




Figure 4.6: Image coverage percentages of the 4 image types.

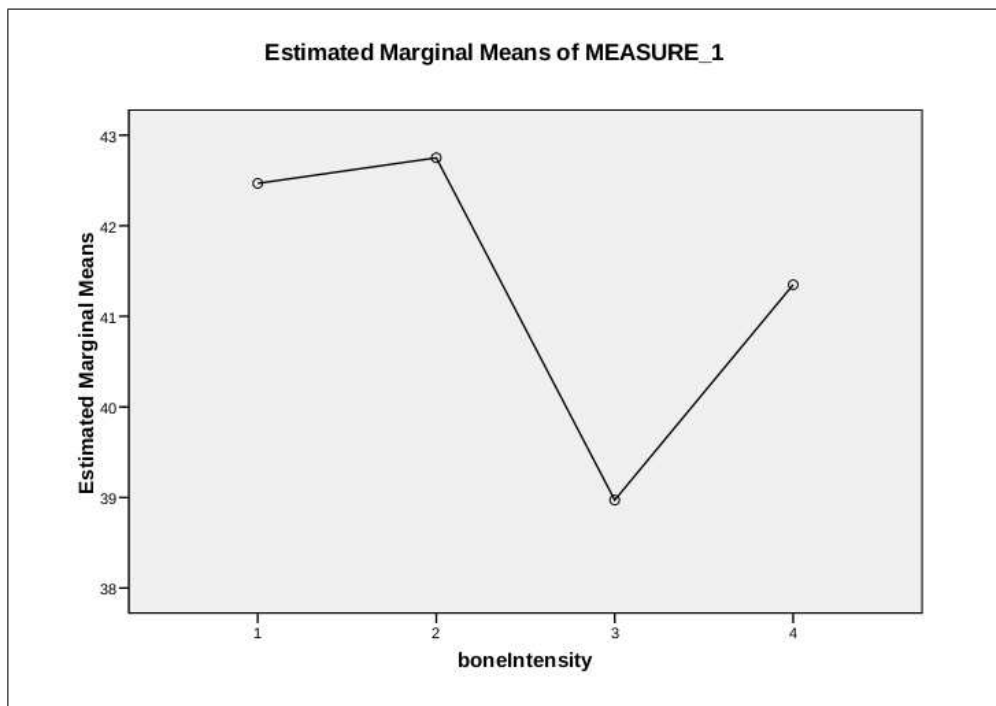


Figure 4.7: Trial duration in seconds for the 4 image types.




Figure 4.8: Coverage rate (percentage per second) for the 4 image types.




Figure 4.9: Confidence ratings of the 4 image types.

Due to time restrictions, accuracy data has not yet been statistically analysed, although initial tests seem to show no significance relating to the accuracy of Radiologists across the 4 image types.

4.3 Further Results

Eye tracking research has proven that people tend to use strategic patterns rather than random coverage when trying to visually locate something in a scene [27]. During their training, Radiologist are taught standardised methods for medical imaging search strategies. The remainder of the research in the project attempts to find insight into the search strategies of the Radiologist involved and also how their strategies may be affected by the different image types.

Some work has been done relating fixation duration to visual interest of the participant [2] including the field of Radiology [4].

Manning's work also states that the duration of a fixation on a feature in an image may prove useful in analysis of the likely correctness of the detection of that feature [13].

4.3.1 Spatiotemporal Analyses

As mentioned, visual search patterns vary among individuals, even in Radiologists who are taught specific search techniques. The following research aims to try and categorise different search strategies from the data collected and to try and devise quantitative measures relating to them.

Some of the previous results are insightful, but do not show much about visual search strategies, due to the fact that they disregard time. It was therefore deemed necessary to develop spatiotemporal analyses from the data.

An OpenGL application was written to display real-time replays of the trials from the data collected. Based on the same principles as the scan-path maps already mention, this application displays fixations as circles connected by lines, the saccades. The circle representing the current fixation grew in diameter with time until the next saccade. A trail of the 4 previous fixations and saccades was displayed on screen to aid in visual recognition of any perceived search strategies. A screenshot of this can be seen in the left part of Figure 4.10 and a video is available at [25].

The application accepted an EDF filename as its command line argument. This allowed for individual viewing of trial replays. Batch replays where easily constructed using shell scripts.

Another method devised for visualising the trial data mixed the concepts of the OpenGL application with heat-maps. The principle here was to created animated heat-maps at an accelerated rate of replay. This was developed in Matlab. For each trial a video was generated with a refresh rate of 12 frames per second. Each frame displayed a standard heat-map overlaid onto the image used in the trial. Fixation points were added one frame at a time and the

heat-map in each frame was generated from that fixation and all preceding fixations in the trial.

The resulting video was a an animated heat-map showing a clear representation of the searching strategies of the participant whilst maintaining historical data in the heat-map. A screenshot of this can be seen in the right hand part of Figure 4.10 and a video is available at [24].




Figure 4.10: Screenshots from two of the animated visualisations. Realtime scan-paths on the left and Accelerated animated heatmaps on the right.

4.3.2 Observations

It was not considered feasible to perform observations of all trials in this manner, due to the vast number of total trials. For this reason, a small subset of the trials were chosen to be analysed and observations were noted.

Upon viewing the OpenGL replays, some distinct patterns became apparent. From this, three key categories of searching strategies where characterised.

- Localised Search : When a participant's gaze was concentrated on a small area that was being studied in detail. Potential
- Strategic Search : When a participant is systematically searching through the image, looking for potential nodules.
- Overview Search : When a participant's gaze is moving rapidly around the image. This seems to be a sign of using peripheral vision.

Looking through different participants it also became clear that some have differing techniques in general, but these three categories still seemed to hold true.

These animated replays proved to be an excellent way to quickly and clearly understand the eye tracking data collected in the trials.

Once the three categories had been established, further observations were performed in an attempt to quantify them. It was hypothesised that localised search consisted of long fixations and short saccades, strategic search consisted of mid length saccades and fixations, and overview search consisted of longer saccades with much shorter fixations.

4.3.3 Analysis

In an attempt to show proof of the above hypothesis, an application was developed to iterate over the data in all trials and return a list all fixation durations and their preceding saccade lengths. The output of this application was plotted on a scatter graph as shown in Figure 4.11. The Y axis is duration of a fixation in milliseconds and the X axis is the length of the preceding saccade in screen pixels.




Figure 4.11: A graph showing Length of saccades from all trials, plotted against the time of the following fixation.

In the lower left hand region of this graph, the points plotted become so dense that the markers overlap and it becomes impossible to observe differences in the densities. For this reason, the same data was plotted in a heat-map style graph. Figure 4.12 shows a zoomed in version of this to accentuate the dense area in question. In this heat-map, dark blue represents areas with a low density of data points through to red representing areas with a high density of data points.

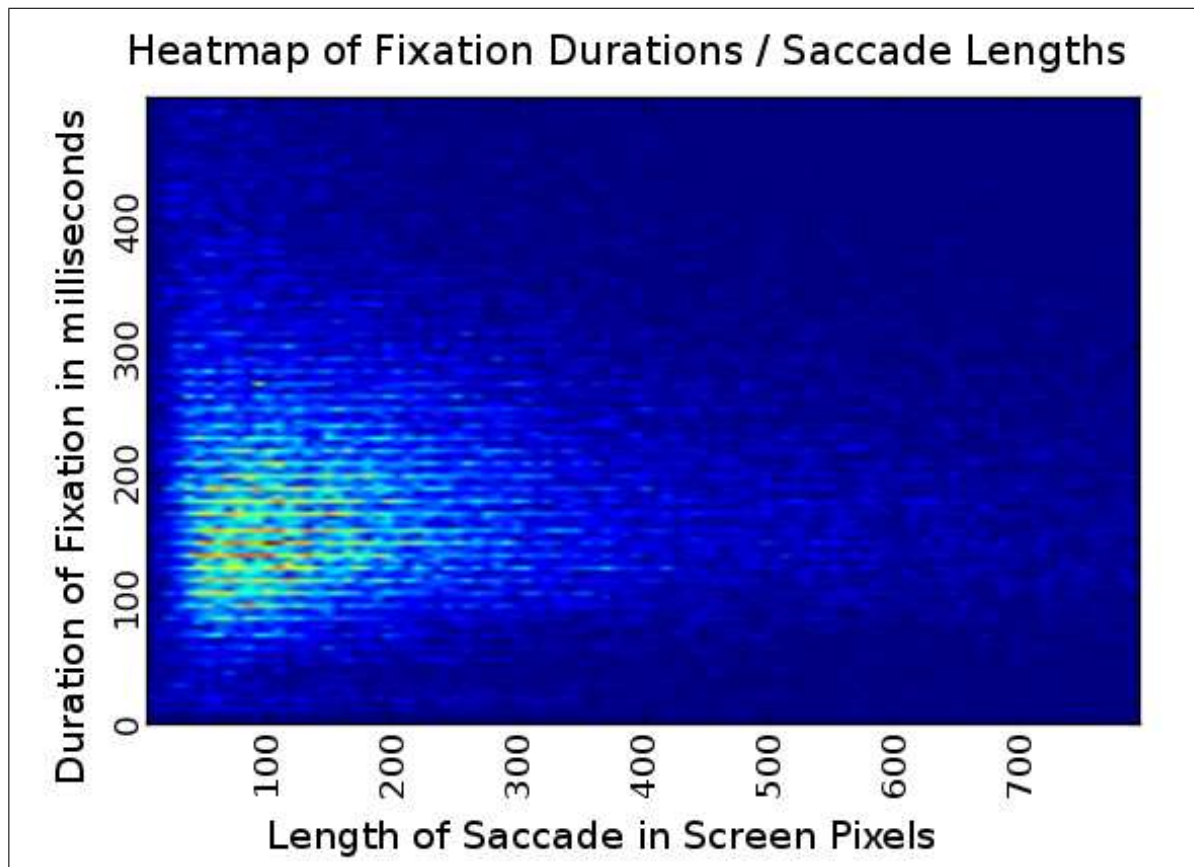


Figure 4.12: A heat-map visualisation of the graph in Figure 4.11, show only a zoomed in area.

4.4 Evaluation

4.4.1 Saccades and Fixations

The scatter plot and associated heat-map produced seem to show clear signs of a relationship between fixation duration and saccade length. They also seem to demonstrate, somewhat, the proposed search strategy taxonomy.

The densely populated region in the lower left of the graph relates to Strategic Search. The data points spread horizontally across the bottom of the graph relate to Overview Search.

The data points spread vertically across the left hand side of the graph relate to the Localised Search Strategy.

Another metric noted during observations, that would be worthy of exploring would be the angle between two saccades. Analysis of this may provide more insight into the search strategies. It would be an achievement to be able to define a mathematical model relating to the sometimes seemingly random search strategies.

4.4.2 Values Encoded in Circle Diameters

It was noted that, when viewing the OpenGL replays, as the circle representing the current fixation increases in diameter linearly with time, the area of the circle will be increasing exponentially.

Area of a 2D shape has been proven more likely to be interpreted as the encoded value than the diameter is, and therefore has the potential to be misinterpreted. This did not pose a problem for these observational tasks as circle size was used as a rough visual indicator, but the issue should be addressed if the software is used for future projects, to avoid misrepresentation of values.

4.4.3 Animated Heat-Maps

One of the problems with using standard, static heat-maps to visualise eye tracking data is, if an area on the heat-map is of a high colour value (e.g. red in the Matlab examples), it cannot be differentiated as to whether that area has been covered multiple times for short periods, or just once, but for a longer period. They also give no indication relating to the search patterns used.

The animated heat-maps created here address these issues.

A point worth noting about the animated heat-maps is thus. As mentioned, there is a steady frame rate of 12 frames per second and a new fixation is added with every frame.

Individual fixation durations vary massively, therefore, these replays are not only accelerated, but the original timings are warped in a non-uniform fashion. Without further research, it is unclear whether this is an issue, but it should be made apparent to viewers of these visualisations.

4.4.4 Nodule Coordinates

As mentioned previously, initial data relating to the position of the nodules in each of the images was approximate. Obtaining accurate locations for all nodules in all images was seen as critical to the integrity of any results based upon them.

The task had to be done manually and it soon became evident that, although it was obvious where some of the nodules were, others were extremely hard to locate. These less

obvious nodules required a massively time consuming process of searching in detail through the images, using many image processing tools and comparisons against other images in order to be sure of their proper location.

Figure 4.2 demonstrates a good example of some of the difficulties involved. The sketch on the left is the map given to represent the nodule locations and types in the image in the middle.

The image on the right shows the actual locations of the nodules in the centre image. It can clearly be seen that the locations provide in the sketch map are very different from the actual locations.

The actual location of the non-calcified nodule (red circle) on the right can be seen a lot lower than is demonstrated on the sketch map.

This is thought to be due to the nodule position being recorded at the time of placement in the lungs, but as the phantom was positioned in front of the X-Ray machine, the nodule became dislodged and fell down the lung region. This scenario occurred for many of the image sets. It took nearly a full week to produce a full set of accurate coordinates for all nodule in all images.

Chapter 5

Conclusion

5.1 Outcomes

There are many outcomes to this project.

The python libraries that were developed to help deal with the EDF files were written in a generic object oriented manner, in the hope that they may also be useful for future projects. These libraries and the applications that were written with them will be made freely available to the School of Computing.

The results of the research taken place during this project will shortly be submitted in a paper to a reputable Radiology Journal.

Based on the ideas of one of the Radiologists involved in the study a prototype training application was built. The Radiologist was looking for an web based application for both training and self assessment of trainee Radiologists.

The basic idea was, to have a web page that displayed X-Ray images. The user would be able to click on the images where they had identified a nodule. The web page could then return results and statistics of their data. Essentially a web based version of the study in this project.

The Radiologists were not sure if such an application was possible, so the opportunity was taken to build them a working proof of concept for the idea. A screen shot of this can be seen in Figure 5.1, and a working prototype is available online here [23].

After the prototype was built, the website was decided as being too far outside of the scope of the project, but plans are being put in place to build this into a fully operational system.

5.2 Future Work

As has been mentioned within the body of this report, there are many areas of this project that could be continued or expanded. The dataset collected contains a wealth of information that may be analysed in many different way. A number of possibilities have been mentioned for

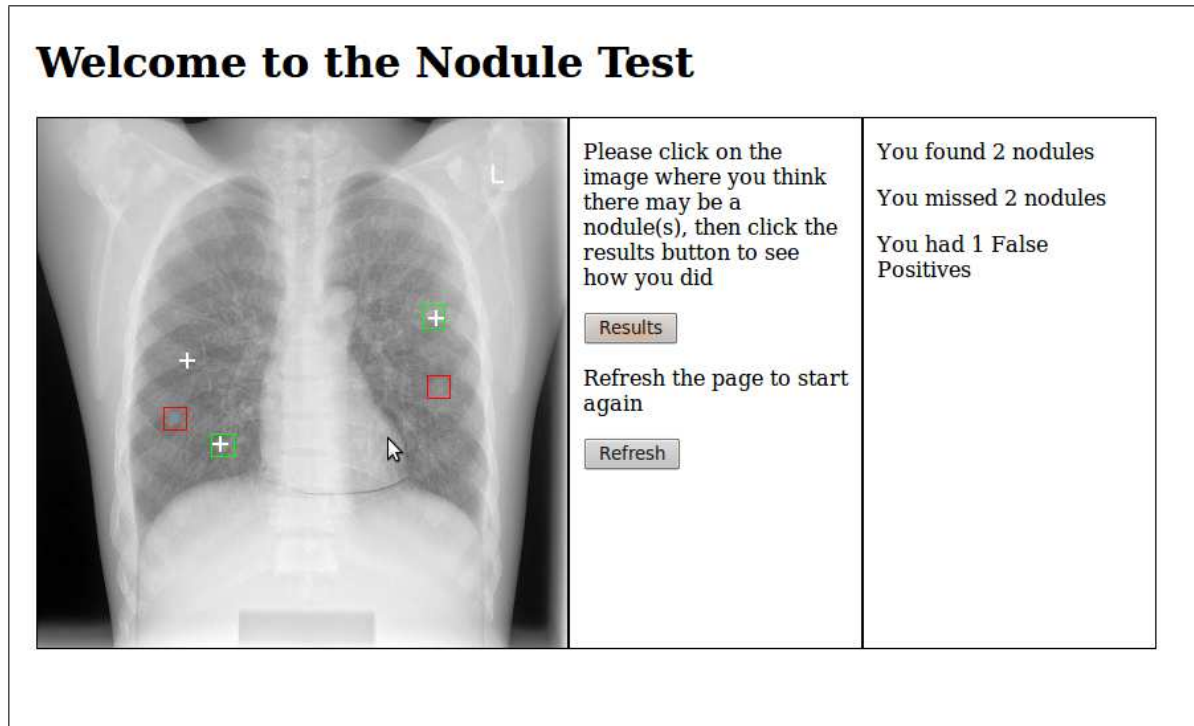


Figure 5.1: Proof of Concept for X-Ray self assessment test.

improved analysis and more advanced techniques of data processing. Using results from this study could potentially lead to many different training applications that could be developed.

5.3 General Evaluation

I believe this project to have been executed successfully. The aims set out in the project plan have been reached, and all minimum requirements have been met. The extra requirements have also been explored in as much depth as was possible in the time allotted, and have provided some interesting results.

The collaboration with an external party was properly planned and the project ran fairly smoothly. The few small set-backs in the project were dealt with swiftly and appropriately and did not cause concern.

As discussed previously, strict deadlines were set from day one. All deadlines were met with deliverables being produced within the time frame.

The hypothesis set by the medical collaborators, of which this project was to explore, has gained some significant results, and this shall hopefully lead on to further research in the area.

Bibliography

- [1] Alex Allain. CProgramming.com. <http://www.cprogramming.com/>. Date of Access 25/01/2010.
- [2] Chia-Fen Chi and Fang-Tsan Lin. A new method for describing search patterns and quantifying visual load using eye movement data. *International Journal of Industrial Ergonomics*, 19:249–257, 1997.
- [3] AT Duchowski. *Eye Tracking Methodology: Theory and Practice*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2007.
- [4] Krupinski EA. Visual scanning patterns of radiologists searching mammograms. *Acad Radiol*, 3(2):137–144, 1996.
- [5] Elizabeth A. Krupinski, William G. Berger, William J. Dallas, and Hans Roehrig. Searching for nodules: What features attract attention and influence detection? *Academic Radiology*, 10(8):861–868, 2003.
- [6] HL Kundel. Predictive value and threshold detectability of lung tumors. *Radiology* 1981, 139:25–29, 1981.
- [7] HL Kundel, CF Nodine, D Thickman, and L Toto. Searching for lung nodules. a comparison of human performance with random and systematic scanning models. *Investigative Radiology*, 5:417–422, 1987.
- [8] SR Research Ltd. EyeLink II Installation Manual. <https://www.sr-support.com/forums/attachment.php?attachmentid=779>. Date of Access 25/01/2010.
- [9] SR Research Ltd. EyeLink II User Manual. <https://www.sr-support.com/forums/attachment.php?attachmentid=778>. Date of Access 25/01/2010.
- [10] SR Research Ltd. SR Research Publications. http://www.sr-research.com/pub_full.html. Date of Access 25/01/2010.
- [11] SR Research Ltd. SR Research Support Site and Forums. <http://sr-support.com>. Date of Access 25/01/2010.

- [12] A Mackenzie and RE Trimble. Evaluation report: Dual energy subtraction for chest radiography. Technical report, NHS, 2007. <http://www.pasa.nhs.uk/pasa/Doc.aspx?Path=%5BMN%5D%5BSP%5D/NHSprocurement/CEP/X-ray/CEP07023S.pdf>.
- [13] D Manning, S C Barker-Mill, T Donovan, and T Crawford. Time-dependent observer errors in pulmonary nodule detection. *The British Journal of Radiology*, 79:342–346, 2006.
- [14] D Manning, S Ethell, and T Crawford. An eye tracking afroc study of the influence of experience and training on chest x-ray interpretation. *Medical Imaging 2003*, 5034:257–266, 2003.
- [15] EA Martin, editor. *Concise colour medical dictionary*. Oxford University Press, 2002.
- [16] NeHe. NeHe OpenGL Tutorials. <http://nehe.gamedev.net/>. Date of Access 25/01/2010.
- [17] C F Nodine and E A Krupinski. Perceptual skill, radiology expertise, and visual test performance with nina and waldo. *Academic Radiology*, 5(9):603–612, 1998.
- [18] Oskar Palinko, Andrew L. Kun, Alexander Shyrovkov, and Peter Heeman. Estimating cognitive load using remote eye tracking in a driving simulator. In *ETRA '10: Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications*, pages 141–144, New York, NY, USA, 2010. ACM.
- [19] S P Prabhu, S Gandhi, and P R Goddard. Ergonomics of digital imaging. *Br J Radiol*, 78(931):582–6, 2005.
- [20] SR Research. EyeLink II Technical Specification. Technical report, SR Research, 2010. Date of Access 25/01/2010.
- [21] C Rooney. Using Gaze Direction as an interface tool. Technical report, School of Computing, 2005. Date of Access 25/01/2010.
- [22] S. Shrestha and K. Lenz. Eye gaze patterns while searching vs. browsing a website. *Usability News*, 9(1), 2007.
- [23] M Terzza. A Self Test Web Application for Radiologists. <http://terzza.com/xray-test/> . Date of Access 10/05/2010.
- [24] M Terzza. Accelerated Replay of Eye Tracking Data using Animated Heat Maps. <http://www.youtube.com/watch?v=L319pLmzHVc> also available through <http://youtube.com/terzza> . Date of Access 15/03/2010.

- [25] M Terzza. Real Time Replay of Eye Tracking Data.
<http://www.youtube.com/watch?v=e7x5NjPMxDc> also available through
<http://youtube.com/terzza> . Date of Access 10/03/2010.
- [26] A West. The Effect of low intensity bony image superimposition on search strategies of experienced and non-experienced viewers. Manuscript in preparation, 2010.
- [27] Guohong Zhou, Koji Ezumia, and Lawrence W. Stark. Efficiency of search patterns.
Computers in Biology and Medicine, 23:511–524, 1993.

Appendix A

Project Reflection

This project has by far been my most enjoyable piece of work at university. It presented challenges that have really stretched my ability and also provided enjoyable ...

It has been a great experience to work with outside departments and learn about new fields. I have also enjoyed the fact (however small a cog in the machine I am), of being able to work on trying to battle a real world problem.

It has been hard work at times, the deadlines that I had to meet meant that I really pull out all the stops, working 7 days a week for a while just in an attempt to get the minimum requirements completed.

With the guidance of my supervisor I was able to break down the project from the humongous task I seemed to have ahead of me, into a properly organised schedule and work on the pieces bit by bit. I found my weekly supervisor meetings to be extremely useful, especially when having to re-arrange plans within the project and effectively manage my time.

I do have a tendency to be a little unorganised, so I decided to set up a wiki for my project notes. I used this for general notes on the project including brainstorming ideas, info about papers I read or wanted to read, road maps for project work and "to do" lists. I dedicated a page on the wiki to what I called my "Journal" and made a point to write at least a few lines in it every day, things like what I'd done that day, what went good/bad, any plans meetings that I had made. I didn't realise how valuable a resource that would quickly become. It helped to keep me organised and up to date, but most, when I felt that the project was moving along slowly for a few days I could have a quick look back in my journal and realise that I had actually been making progress, in the bigger picture. It was also helpful during the write up of the project report. I was able to look back through the wealth of notes of the day to day goings on that I might otherwise have forgotten.

To anyone about to embark on an FYP, I would highly recommend setting up a wiki, or even just a dedicated project note book, and write some of you progress in it every day. It may not seem worth while at the beginning, but it soon will be. If you don't have your own server, there are plenty of freely available wiki services on the internet that you can sign up for.

I think that I have handled my time well during this project and am please the outcomes.

The one thing that I feel I slipped up on is the time that I allocated to report writing. I had read previous student's reports stating how they underestimated the time needed complete the write up and thought that I would not make the same mistake. I had allocated nearly three weeks solid for the report and I still don't think it was enough. If there was anything I would do differently in this project, I would have put more time towards the report and started it properly a lot earlier in the project.