ISSN 0280-5316
ISRN LUTFD2/TFRT--5705--SE

Kinematic Analysis of Rapid Eye Movements for Vestibular Disorders

Gerard Duffin

Department of Automatic Control
Lund Institute of Technology
June 2003

Department of Automatic Control
Lund Institute of Technology
Box 118
SE-221 00 Lund Sweden

Document name
MASTER THESIS

Date of issue
June 2003

Document Number
ISRN LUTFD2TFRT--5705--SE

Author(s)
Gerard Duffin

Supervisor
Rolf Johansson at LTH. Måns Magnusson, Mikael Karlberg and Per Anders Fransson at Lund University Hospital.

Sponsoring organization

Title and subtitle
Kinematic Analysis of Rapid Eye Movements for Vestibular Disorders (Kinematisk analys av snabba ögonrörelser vid vestibulära störningar)

Abstract
The system under development provides a means to assess the semi-circular canals of the human vestibular system. The Impulse Test is a simple method to detect disorders within the three sets of semi-circular canals, by stimulating each pair of canals in turn. This report describes the work carried out to develop a simple, non-intrusive system whereby the patient can be assessed in a matter of seconds. The system consists of a single high-speed monochrome camera connected to a computer with the developed software. The main area of work so far, has been the implementation of an accurate image processing technique to track both the head and the eyes. Pattern recognition was attempted first, but this was met with limited success. The method of image processing then shifted to thresholding performed upon the eye. Modelling the head and eye in three dimensions were also integral parts of the project. The eye's origin must be accurately represented as eye velocity is measured relative to this point. Therefore, inaccuracies in describing the eye's origin yield misleading results.

Keywords

Classification system and/or index terms (if any)

Supplementary bibliographical information

ISSN and key title
0280-5316

ISBN

Language
English

Number of pages
100

Recipient's notes

Security classification

LUND INSTITUTE
OF TECHNOLOGY
Lund University

Kinematic Analysis of Rapid Eye Movements for Vestibular Disorders.

Gerard Duffin

Master Thesis Project Report – 2002 / 2003

Lund Institute of Technology

Imperial College, London

Abstract

The system under development provides a means to assess the semi-circular canals of the human vestibular system. The Impulse Test is a simple method to detect disorders within the three sets of semi-circular canals, by stimulating each pair of canals in turn. This report describes the work carried out to develop a simple, non-intrusive system whereby the patient can be assessed in a matter of seconds. The system consists of a single high-speed monochrome camera connected to a computer with the developed software. The main area of work so far, has been the implementation of an accurate image processing technique to track both the head and the eyes. Pattern recognition was attempted first, but this was met with limited success. The method of image processing then shifted to thresholding performed upon the eye. Modelling the head and eye in three dimensions were also integral parts of the project. The eye's origin must be accurately represented as eye velocity is measured relative to this point. Therefore, inaccuracies in describing the eye's origin yield misleading results.

Table of Contents

1 INTRODUCTION 9

1.1 BACKGROUND 11
1.2 PROJECT AIMS 11
1.3 PROJECT RESULTS 12

2 BIOMEDICAL BACKGROUND 13

2.1 THE VESTIBULAR SYSTEM. 15
2.1.1 The Inner Ear 15
2.1.2 The Vestibular Labyrinth 16
2.1.3 The Otolith Organs : Saccule and Utricle 16
2.1.4 The Semi-Circular Canals 16
2.1.5 Vestibulo-Ocular Reflex (VOR) 20
2.1.6 Symptoms of Vestibular Disorders 20
2.2 LISTING'S LAW 21
2.3 CURRENT TECHNIQUES TO TEST THE VESTIBULAR SYSTEM 22

3 IMPULSE TEST SYSTEM 23

3.1 SYSTEM AIMS 25
3.2 SOFTWARE ENVIRONMENT AND THE HARDWARE. DETAILS AND SPECIFICATIONS 25
3.2.1 The Camera 25
3.2.2 The Lens 26
3.2.3 FRAME GRABBER NI 1422 26
3.2.4 The Computer 26
3.2.5 Software 26
3.3 OBTAINING DATA 27
3.3.1 Data Source 27
3.3.2 Recording New Data From The Camera 27
3.4 THE IMAGE ACQUISITION BUFFER 28
3.5 CALIBRATION 29
3.5.1 The Head Marker Template 30
3.5.2 Estimation of Rectangles 31
3.5.3 Estimation of Circles 33
3.5.4 Where to sit 33
3.6 THE CURRENT IMAGE PROCESSING METHODS 34
3.6.1 Setting the Parameters 34
3.6.2 The Procedure 35
3.6.3 The Results 36
3.6.4 The Loss of a Template 37
3.6.5 Difficulties with the Selection of Objects 39
3.6.6 Lighting 40
3.7 FURTHER IMAGE PROCESSING DEVELOPMENTS 41
3.7.1 Edge Detection 41
3.7.2 Area Detection 42
3.7.3 Thresholding 43
3.8 SUGGESTED TECHNIQUE 43
3.8.1 Light Intensity 45
3.8.2 Determining the Pupil's Location 46
3.8.3 Using the Reflection to Describe the Eye 47
3.8.4 Eyelashes and mascara 49
3.8.5 Handling Blinking 50
3.8.6 Advantages of this Image Processing Implementation. 51

4 HEAD MODEL 53

4.1 THE HUMAN HEAD 55
4.2 THE CURRENT LABVIEW MODEL 55
4.3 MATLAB MODEL 55

4.3.1 Head Design/Structure *55*
4.3.1.1 Graphical Implementation 56
4.3.1.2 Trajectories 57
4.3.1.3 Limitations 58
4.3.2 Locating the Eye's Origin *58*
4.3.3 Results *59*
4.3.3.1 Sensitivity to the Head's Horizontal Radius 59
4.3.3.2 Sensitivity to the Distance between the Eye and the Head Marker 60

5 SIGNAL ANALYSIS 63

5.1 COMPARISON OF RESULTS 65
5.1.1 Results from the 2D model with the Pattern Matching *65*
5.1.2 Results from new technique *67*
5.1.3 Results with 3D head model and pattern matching *70*
5.2 SYSTEM IDENTIFICATION 73
5.2.1 Experimental Aims *73*
5.2.2 Data *73*
5.2.3 Initial Spectral Tests *74*
5.2.4 Non Parametric Modelling *74*
5.2.5 Parametric Modelling *75*
5.2.5.1 Fitting ARMAX models 76
5.2.5.2 Fitting Box Jenkins models 78
5.2.5.3 Comparison Between ARMAX and Box Jenkins 79
5.2.6 Sub Space Based Identification *80*
5.2.7 Results *82*

6 DISCUSSION AND FURTHER WORK 85

6.1 IMAGE PROCESSING ALGORITHM 87
6.2 MODELLING 89
6.3 SYSTEM CONFIGURATION 90
6.4 CALCULATING THE VOR 90
6.5 SYSTEM IDENTIFICATION 91

7 CONCLUSIONS 93

8 ACKNOWLEDGEMENTS 99

9 BIBLIOGRAPHY 100

Table of Figures

Figure 1. An image of the Inner Ear 15
Figure 2. The Vestibular Organs; Three Semi-Circular Canals, Saccule and Utricle 17
Figure 3. Semi-circular canal plane orientations 17
Figure 4. The non-linear response from cilia within the lateral canals. Left, the normal case where the cilae function correctly in both canals. Right, the case where the left canal is affected. 18
Figure 5. Head and eye velocities. The left image shows the normal case and the right shows the case where the left posterior canal is affected [2] 19
Figure 6. How the eye rotates. Left, two pupil positions upon the eye. Right, the choice of rotation axes between the pupil locations [3] 21
Figure 7. An example of how Listings Law applies [3] 22
Figure 8. The Pulnix TM-6710 Camera 26
Figure 9. Positioning the patient relative to the camera 28
Figure 10. Head marker 1 30
Figure 11. Head marker 2 31
Figure 12. The current head marker 31
Figure 13. Searching for boxes. The green lines define the search regions, the yellow dots define the detected edges and the red lines show the result. 32
Figure 14. Searching for circles. The green lines define the search regions, the yellow dots define the detected edges and the red lines show the result. 33
Figure 15. The effects of Sub-pixel estimation : Top, results with sub-pixel estimation. Bottom, raw data. 34
Figure 16. Image Processing Properties 35
Figure 17. Image of a search 36
Figure 18. Image Processing Results 37
Figure 19. Pupil behind the eyelid 39
Figure 20. Effect of the white dot. Top, the initial image, with correct tracking. Middle and bottom, subsequent frames showing poor tracking qualities as a result of the reflection 41
Figure 21. Top. A simulation of the effects of the reflection upon edge detection. Bottom left, an image of the eye with the reflection within the pupil. Bottom right, the same image with circular edge detection performed. 42
Figure 22. An image using the centroid technique. 43
Figure 23. Thresholding. Left, the original image. Right, the image after thresholding. Two threshold levels were employed. (100 and 200) 43
Figure 24. Binary Thresholding. Left, the original image. Right, the binary image. (Threshold level 170) 44
Figure 25. Images of the same eye taken from the same test sequence. 2 threshold levels were applied. (100 and 240) 44
Figure 26. Images of the same eye taken from the same test sequence. 5 threshold levels were applied. (50, 100, 150, 200 and 240). Compare with the previous figure. 45
Figure 27. Left, the iris is selected. Right, it is split into four regions to determine the light intensity for the threshold levels 45
Figure 28. Images of the same eye taken from the same test sequence. 5 threshold levels were applied. (31, 100, 150, 200 and 240). Compare with the previous two figures. 46
Figure 29. An image of how the pupil becomes obscured by the reflection [1]. 46
Figure 30. Top left, how it was assumed to appear. Top right, how the eye appeared. Bottom, a real image showing this effect 47
Figure 31. A sequence of eye movements with the pupil and the reflection both tracked. Left, the normal position. Middle, looking down and left. Right, the reflection on the other side of the pupil, while in the resting position 49
Figure 32. A continuation of the same sequence. Left, looking right. Middle, looking up. Right, looking down. 49
Figure 33. Detecting the eye through the eyelashes 50
Figure 34. A sequence of images showing how a blink is recognised and tracked. 51
Figure 35. A diagram of the head model, with eyes elevated by 30 degrees. 56

FIGURE 36. THE TRAJECTORIES OF THE HEAD MARKER AND ROTATION POINTS. TOP LEFT IS FROM A SIDE VIEW. TOP RIGHT IS FROM A PLAN VIEW. BOTTOM SHOWS THE 3D RESULT 57
FIGURE 37. 3D MODELLING RESULTS. TOP LEFT, A FRONTAL VIEW. TOP RIGHT, FROM A SIDE VIEW. BOTTOM LEFT, A 3D VIEW. BOTTOM RIGHT, HEAD AND EYE VELOCITY. 59
FIGURE 38. SENSITIVITY OF THE EYE VELOCITY TO VARIATIONS IN HEAD'S HORIZONTAL RADIUS. RADIUS: TOP LEFT, 10CM. TOP RIGHT, 12CM. BOTTOM, 14CM. 60
FIGURE 39. SENSITIVITY OF THE EYE VELOCITY TO VARIATIONS IN DISTANCE BETWEEN THE EYE AND THE FOREHEAD. LEFT, 3CM. RIGHT, 5CM. 61
FIGURE 40. RESULTS FROM A NORMAL PATIENT. LEFT, RAW POSITIONAL DATA. RIGHT, PIXEL VELOCITY DATA WITH A SIMPLE 2D MODEL. 65
FIGURE 41. THE EFFECTS OF DIFFERENT EXPERIMENTAL SET UP CONDITIONS WITH THE PATTERN MATCHING TECHNIQUE 66
FIGURE 42. RESULTS FROM AN AFFECTED PATIENT. LEFT, RAW POSITIONAL DATA. RIGHT, PIXEL VELOCITY DATA WITH A SIMPLE 2D MODEL 67
FIGURE 43. DATA FROM A NORMAL SUBJECT. TOP LEFT, RAW POSITIONAL DATA. TOP RIGHT, PIXEL VELOCITY. BOTTOM, ANGULAR VELOCITY 68
FIGURE 44. DATA FROM AN AFFECTED SUBJECT. TOP LEFT, RAW POSITIONAL DATA. TOP RIGHT, PIXEL VELOCITY. BOTTOM, ANGULAR VELOCITY 69
FIGURE 45. A SACCADE WITH NO UNWANTED SPIKE. TOP LEFT, RAW POSITIONAL DATA. TOP RIGHT, PIXEL VELOCITY. BOTTOM, ANGULAR VELOCITY 70
FIGURE 46. RESULTS FROM A NORMAL PATIENT. TOP LEFT, POSITION DATA. TOP RIGHT, LINEAR VELOCITY. BOTTOM, ANGULAR VELOCITY 71
FIGURE 47. DATA FROM AN AFFECTED SUBJECT. TOP LEFT, RAW POSITIONAL DATA. TOP RIGHT, PIXEL VELOCITY. BOTTOM, ANGULAR VELOCITY 72
FIGURE 48. A SACCADE WITH NO UNWANTED SPIKE. TOP LEFT, RAW POSITIONAL DATA. TOP RIGHT, PIXEL VELOCITY. BOTTOM, ANGULAR VELOCITY 72
FIGURE 49. TOP, IDENTIFICATION DATA. BOTTOM, VALIDATION DATA 73
FIGURE 50. THE COHERENCE PLOT RESULTING FROM THE SPECTRUM FUNCTION. 74
FIGURE 51. CORRELATION ANALYSIS AND THE ESTIMATED STEP RESPONSE 75
FIGURE 52. THE BODE PLOT FROM THE INPUT / OUTPUT DATA. 75
FIGURE 53. ARMAX MODELLING RESULTS FOR POLYNOMIALS OF ORDER 10. TOP LEFT SHOWS THE POLE ZERO CONFIGURATION. TOP RIGHT SHOWS THE RESIDUAL ANALYSIS BASED UPON THE VALIDATION DATA. BOTTOM LEFT SHOWS THE ESTIMATED STEP RESPONSE. THE BOTTOM RIGHT IS THE ESTIMATED TRANSFER FUNCTION 76
FIGURE 54. ARMAX MODELLING RESULTS FOR POLYNOMIALS OF ORDER 1. TOP LEFT SHOWS THE RESIDUAL ANALYSIS PERFORMED UPON THE VALIDATION DATA. TOP RIGHT SHOWS THE ESTIMATED TRANSFER FUNCTION. BOTTOM SHOWS THE ESTIMATED STEP RESPONSE 77
FIGURE 55. POSSIBLE MODEL ORDERS RESULTING FROM A BALANCED REALISATION ON THE TENTH ORDER MODEL 77
FIGURE 56. COMPARISON OF TENTH AND FIRST ORDER MODELS. LEFT, THE IDENTIFICATION DATA. RIGHT, THE VALIDATION DATA. 78
FIGURE 57. THE RESULTS FROM BOX JENKINS MODELLING. NB=1, NC=0, ND=0, NF=1. TOP LEFT SHOWS THE RESIDUAL ANALYSIS PERFORMED UPON THE VALIDATION DATA. TOP RIGHT SHOWS THE ESTIMATED STEP RESPONSE. BOTTOM SHOWS THE ESTIMATED TRANSFER FUNCTION. 79
FIGURE 58. A COMPARISON OF THE BOX JENKINS (GREEN) AND THE ARMAX METHODS (PURPLE). TOP LEFT THE RESIDUAL ANALYSIS BASED UPON VALIDATION DATA. TOP RIGHT, THE STEP RESPONSE. BOTTOM THE TRANSFER FUNCTIONS. 80
FIGURE 59. SINGULAR VALUES OF THE HANKEL MATRICES OF THE IMPULSE RESPONSE 81
FIGURE 60. THE RESULTS OF A THIRD ORDER SUBSPACE MODEL. TOP LEFT IS THE ESTIMATED TRANSFER FUNCTION. TOP RIGHT IS THE ESTIMATED STEP RESPONSE. BOTTOM IS THE POLE ZERO PLOT 81
FIGURE 61. THE STEP RESPONSE FROM A SECOND ORDER SUBSPACE MODEL 82
FIGURE 62. TOP, THE RESULTS OF PROCESSING THE IDENTIFICATION DATA. BOTTOM, THE VALIDATION DATA ON THE THIRD ORDER SUBSPACE MODEL (LEFT) AND THE SECOND ORDER SUBSPACE MODEL (RIGHT) 82
FIGURE 63. A POSSIBLE INDIRECT LIGHTING ARRANGEMENT. 89
FIGURE 64. HOW THE REFLECTION MOVES RELATIVE TO THE EYE'S ORIGIN 90
FIGURE 65. HOW THE VOR IS EXPECTED TO DEPEND ON HEAD VELOCITY 91

1 Introduction

1.1 Background

Vestibular disorders are common amongst many spectra of people, regardless of age and sex, and certain disorders can even be hereditary. They can be difficult to diagnose correctly as many of the symptoms are common to several different disorders. Therefore the ability to distinguish between them can result in earlier diagnosis and better treatment and recovery for the patient. Within the inner ear are five organs, three of which act as angular acceleration transducers and two as linear acceleration and position transducers within a gravitational field. It is the three angular transducers, which are under investigation in this work. There is a great need to diagnose Vestibular disorders quickly and accurately, as they can prevent a subject leading a normal life.

Tests have been developed to stimulate these organs, provoke a neural response and assess their physical state. One of the methods for stimulating the three angular sensors is the 'Impulse Test', which has been employed in this work. For each of the angular transducers under investigation, an operator rotates the patient's head in the relevant direction by approximately 30 degrees, while the patient attempts to fixate their eyes on a target straight ahead. The resulting movement of their eyes gives an insight into the state of each of the transducers and quantify the value of the Vestibulo-Ocular Reflex (VOR).

The best method currently available to assess the physical state of each of these semi-circular canals is the Search Coil technique [2]. This is a highly accurate method, but the high running costs in conjunction with the highly specialised equipment makes it a difficult procedure to conduct on patients on a regular basis.

1.2 Project Aims

The eventual aim of this system is to have the means whereby a patient can be fully assessed by sitting in front of a video camera. Each of the six canals shall be assessed in turn to determine whether any of them display characteristic symptoms of the disorder. If so, the effect shall be quantified to define a VOR for each canal and it is this value which the system should eventually produce. None of the highly specialised techniques associated with the Search Coil method need be employed, and the whole test from beginning to end should take a matter of only a few minutes.

From the point of view of this project work, the aim was to establish a firm foundation from which data could be obtained and recorded accurately, so initial testing could be performed to assess the system. When this thesis work began, there already existed a basic version of the system and this version was used as a starting platform. Some of the original software elements still exist in the current version and they will be highlighted in the relevant places.

Image Processing is a vital part of the system and therefore a method to accurately track the eye up to and beyond the limits of the Impulse Test movements was a main target to provide accurate results. The ability to transform the two-dimensional data

into three-dimensional data was also another fundamental aim within the project, as this would enable the system to provide meaningful results compared to the two-dimensional model employed in the initial system.

1.3 Project Results

Although the system is not yet complete, each of the modules developed can currently be linked together to yield a more complete picture. Each stage was critically assessed to evaluate its performance and the results were used to decide whether development should be continued or reworked.

The overall system structure has been laid down, with further development work in mind. The Image Processing implementation was subject to constant scrutiny from the beginning, as its accuracy is perhaps the single most important element of the entire system. The initial implementation was removed in favour of IMAQ Imaging tools, but again these showed weaknesses during patient testing and hence, they were redeveloped again. Currently, the method under development eliminates many of the issues raised with the previous implementation, although it has not been integrated into the Impulse Test system.

The head modelling yielded moderately successful results, although it highlighted many areas in which the system required further development.

2 Biomedical Background

2.1 The Vestibular System.

The Vestibular System is a sensory system whose purpose is to detect position and acceleration of the head in the gravitational environment. The vestibular system's organs and detectors are located within the inner ears with neural connections to the brain. Albeit it is a sensory system and we are not consciously aware of its presence or function, we are dependent upon its operation. It is responsible for us being able to maintain our posture and balance in space. Therefore, disorders of the system accentuate our awareness of it, even in mild forms, as the symptoms will prevent the subject from leading a normal life.

An overall description of the vestibular system will not be presented here, although the focus of this report will be upon the inner ear, and specifically, the three semi-circular canals within each ear. The purpose of the Impulse Test System is to facilitate the assessment of each of the six canal plane's functioning, hence the attention to this particular area.

2.1.1 The Inner Ear

The ear is composed of three regions, the external ear, the middle ear and the inner ear. The external ear comprises of the parts visible on the outside of the head. The middle ear contains the auditory ossicles and it is separated from the external ear by the tympanic membrane, or ear drum, and from the inner ear by the oval window.

The internal ear consists of the membranous labyrinth, which contains organs for hearing and balancing, and is suspended within the bony labyrinth. The bony labyrinth is composed of two parts, the cochlea and the vestibular labyrinth. The cochlea is responsible for converting the sound waves travelling from the external environment, into neural signals.

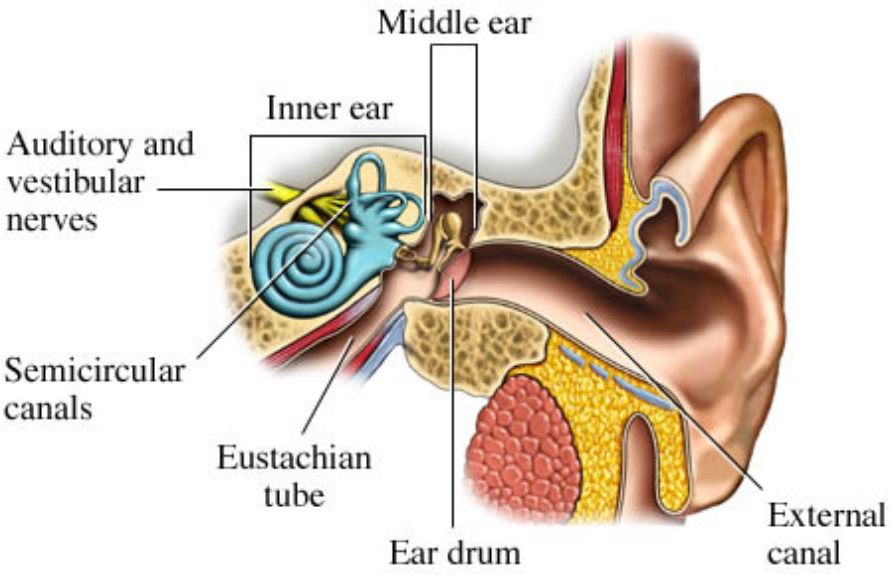

Figure 1. An image of the Inner Ear

2.1.2 The Vestibular Labyrinth

The vestibular labyrinth is an oval, bony chamber containing five organs for both the detection of linear and angular acceleration. The static and linear sensors relate mostly to the position and linear motion of the head. These are the otolith organs. The semi circular canals constitute the dynamic transducers sensors, which detect the angular acceleration of the head and for rotational movements.

- Linear Acceleration. The utricle is the otolith organ, which senses linear head acceleration in the horizontal plane. The saccule is the otolith organ, which senses linear head acceleration in the vertical plane.

- Angular Acceleration. The three semicircular canals combine to sense head rotation in the three directions. The canals are positioned at approximately 90 degrees to each other; one approximately in the lateral plane and two in vertical planes.

All of these organs are connected together in the membranous passage, containing an endolymphatic fluid.

Each of the five balance related receptor organs within the vestibular labyrinth are structured and positioned in a specific way to sense their respective motions. But several of the fundamental characteristics of the structures are common. Each organ contains a volume of space (either a canal or a cavity), where the endolymphatic fluid moves due to the head movement in certain directions. Along the boundaries of the canals or cavities are sensory epithelial cells coated with cilia, which become displaced with the fluid's motion. These receptor hair cells convert their motion into neural signals and convey these by the vestibular nerve.

2.1.3 The Otolith Organs : Saccule and Utricle

The two otolith organs depend on gravity to detect motion in three dimensional space. They each consist of a cavity in the bony labyrinth containing a small, dense mass, the otoconial layer. This mass is suspended in the cavity with a fluid and as the head moves, the otoconial layer lags behind the head due to inertial forces. This causes the cilia attached to the cavity wall to deviate, detecting the movement.

2.1.4 The Semi-Circular Canals

The three semi-circular canals are responsible for determining the angular acceleration of the head. They are each positioned at approximately 90 degrees to each other, lying in three spatial planes. The Lateral Canal mainly detects angular motion in the horizontal plane although it is actually inclined by several degrees to the horizontal, so it is also affected to a lesser extent by any vertical movements. The second and third canals are called the Anterior and Posterior Canals and are arranged in vertical planes, rotated left and right around the sagittal plane by approximately 45

degrees. These angles are not exact and they vary from one person to another, just as most other physiological features do.

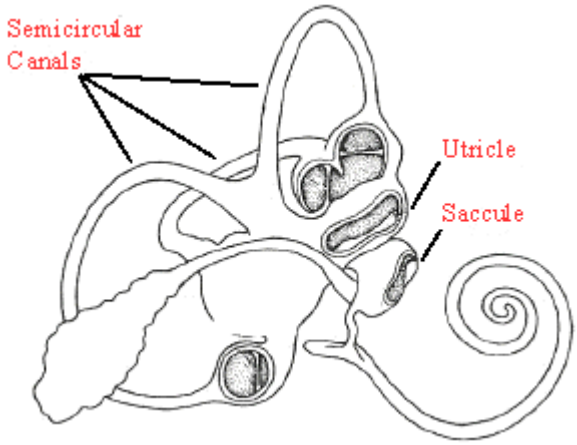

Figure 2. The Vestibular Organs; Three Semi-Circular Canals, Saccule and Utricle

The canals naturally function in pairs when motion is induced in their plane, one from each ear, to determine angular motion in each direction. They pair up as;

- Left Lateral and Right Lateral
- Left Anterior and Right Posterior (LARP)
- Right Anterior and Left Posterior (RALP)

The reasoning is more obvious when looking at their orientations and comparing the planes they lie in. So, for example, if the patient's head is moved in horizontal directions only, the fluid in the two lateral canals will become displaced, relative to the canals. In a patient with no vestibular disorders, the cilia will also become displaced within the fluid resulting in an electrical charge being induced and sent to the brain via the vestibular nerves. A similar mechanism drives the other two canal plane pairs (CPP).

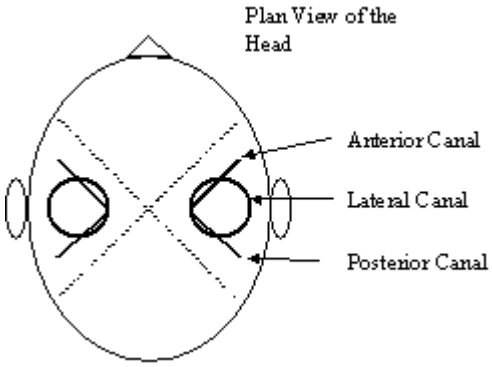

Figure 3. Semi-circular canal plane orientations

The reason for their orientations is not fully understood but it is the result of biological evolution and the arrangement is similar in all species, with only small

differences. From an engineering perspective, it may seem more logical to arrange the canals such that they correspond exactly to horizontal, vertical and torsion movements. In this fashion, there would be less interaction between canal planes in solely horizontal, vertical and torsional, making it easier to diagnose which, if any, of the canals are malfunctioning.

But as this is not the case, testing each canal plane pair is not so simple. To get the maximum response from the CPP under investigation, the effects from the other two CPP should be minimised. This can be achieved by rotating the head in the plane of the relevant canals, one plane at a time. The difficulties with this method are that the canals are not at exactly 90 degrees to each other and without a (for example) CT Scan it is not possible to know exactly what orientations the CPP lie in. So there is always likely to be interference to a small extent from the other canal planes. But with good choices for the planes of head rotation, these unwanted effects can be reduced. It was normal for the three separate tests to be performed in the lateral plane (to test the lateral canals), and in the vertical plane with the head rotated by 45 degrees left (to test RALP) and right (to test LARP) of the sagittal plane.

So to test each CPP, the head is orientated correctly and then moved vertically or horizontally, depending on the CPP in question. But for each individual canal to be investigated, the test must be performed in two directions (left/right for lateral tests and up/down for vertical tests). This is because the canal planes are non-linear systems. In the case of the lateral canals, one can consider the following;

Under stationary conditions, the head is at rest and so too is the fluid within each lateral canal. When the head moves to the right, the hairs in the right hand side canal induce an increased lining frequency than at the rest position. Simultaneously, the hairs in the left hand side canal induce a reduced lining frequency. This causes a slow eye deviation of opposite direction and the gaze remains stable in space. The frequency of the induced signals depends on the speed of head rotation. If the speed is high enough, the frequency in the left canal will drop to a minimum value, effectively "saturating" it and isolating it from the experiment. Meanwhile, the right hand side canal can still gain more charge, rendering it the only effective canal plane. The reverse is true for head movements in the left direction and a similar explanation is valid for LARP and RALP planes.

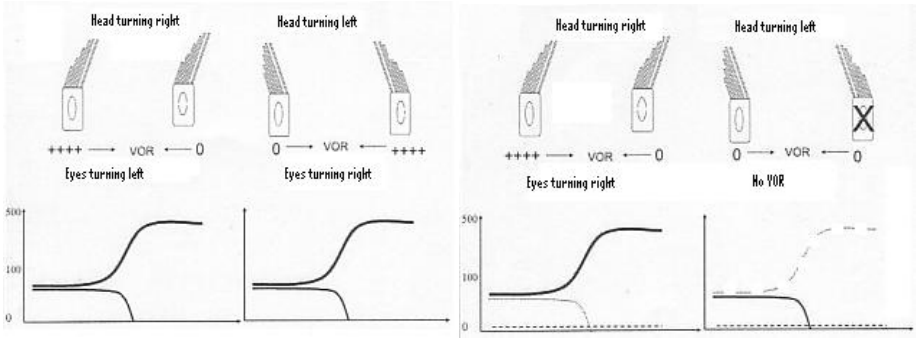

Figure 4. The non-linear response from cilia within the lateral canals. Left, the normal case where the cilae function correctly in both canals. Right, the case where the left canal is affected

It is the condition of these hairs and the fluid within the canal that determines the level of normality of the canal. In a fully functioning canal, all the hairs are able to become displaced and induce a charge upon the connecting nerve. But if some of the cilia are damaged or the canal becomes lesioned, then the response of the canal to rapid head movements becomes impaired. In the above right figure, the dashed line shows where the response is expected in the normal case, but due to the disorder, only a reduced neural signal is sent along the vestibular nerve

But how is it possible to determine whether an individual canal is normal or not? Does the electrical charge induced through the head rotations need to be measured? Obviously this is not a practical solution as the inner ear is not easily accessible and would be destroyed if exposed to such procedures. Thankfully, because of the close connection between the vestibular system and the ocular system, it is possible to observe the eye movements in relation to head movements to determine the condition of each canal to a certain degree.

If a subject stares at a fixed point in front of them and their head is rotated slowly, such that none of their canals become saturated, then the subject's eyes can track the target. But if the head velocity is increased such that one canal at a time becomes saturated (as described above), then the eye will respond through a combination of signals from the vestibular system and conscious decisions about the target's location. If the non-saturated canal is functioning, the eye will be able to track the target correctly with minimal time delay. This inherent short delay is due to the time between the canals sensing the motion, and the brain calculating the correcting eye response. But if the canal is damaged, then the eye will not initially rotate about its own centre to compensate for the head movement, and the subject will gaze in the wrong direction. After a longer time interval, approximately 100ms, the subject will consciously realise that they are staring in the wrong direction and correct their gaze position. It is this time delay, and conscious decision to rotate the eyes, that determines how badly the canal is damaged.

The following diagram shows two sets of results [2], one from a normal subject, left, and one from a subject who had a loss of functioning in the left posterior canal, right. It can be seen here, that the eye does not fully compensate for the head velocity until approximately 200ms, when the eye saccade occurs. This manifests itself as the short, rapid eye velocity lasting for approximately 50ms.

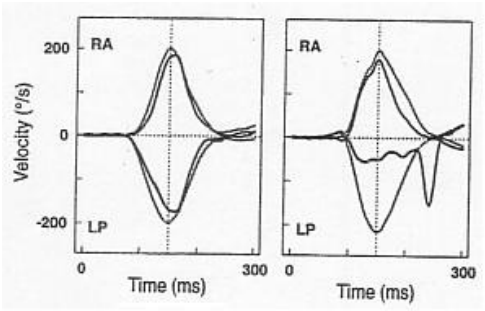

Figure 5. Head and eye velocities. The left image shows the normal case and the right shows the case where the left posterior canal is affected [2]

This situation provides a basis for the "Impulse Test", whereby the head is moved rapidly in each of the three CPP, isolating two CPP and one of the remaining canals per test.

2.1.5 Vestibulo-Ocular Reflex (VOR)

The Vestibular-Ocular Reflex (VOR) is a neurological response that starts with the stimulation of the vestibular organ by head movements and results in compensational eye movements [2], [4] and [5]. It forms the basis of our ability to move and see simultaneously As the head is spun in one direction, the VOR produces coordinated motion of the eyes in the opposite direction, in order that they fixed on a specified target. In subjects with a poor VOR, a delayed saccadic eye movement (fast changes of eye position between fixations) will prevail. A method to quantity the VOR for each vestibular organ is the overall aim of the Impulse Test system. Within this system, a VOR is defined for each individual SCC [4], as defining a single VOR for a patient makes little sense when investigating each SCC separately.

The method to test the lateral canals is well understood and relatively easy to perform. The patient's head is rotated rapidly in the horizontal plane in order to stimulate a response. The eye's response then comes mostly from the lateral canals. But in the case of the vertical CPP, the head must be rotated by approximately 45 degrees left and right from the sagittal plane. Then the rapid rotations are performed in the vertical plane.

2.1.6 Symptoms of Vestibular Disorders

There are many different symptoms of vestibular disorders, which can be interpreted to give varying diagnoses. A patient may experience a combination of symptoms and care must be taken to ensure a correct diagnosis is made.

Perhaps the most common symptom is Vertigo. The affected patient feels dizzy and the environment appears to be spinning around them making it difficult to stand or walk. The dizziness results from the apparent movement of the environment. Other symptoms, such as nausea and vomiting, partial or full loss of hearing and tinnitus (a ringing sensation in the ear) can all be used in the diagnosis. A serious problem that is especially common in elderly sufferers is that when they feel dizzy and disorientated, they fall down, potentially causing further injury to themselves. Vertigo can also affect the patient's vision. They may experience motion or swirling of static images, giving a sensation of being in a moving environment.

Nystagmus is a term used to express the undulating movements of the eyes. It is a common symptom of vestibular disorders and it results in the subject being unable to maintain steady fixation and clear vision. It may be defined as a periodic rhythmical ocular oscillation. Nystagmus usually involves both eyes moving in the same direction and in phase with each other, but there are many different forms of nystagmus and hence it is not possible to give a general, yet exact explanation of what happens.

2.2 Listing's Law

So far, all the details presented give no indication that the horizontal, vertical and torsional movements are in anyway related. It is easy to understand that to maintain a fixed view of a target, up/down head movements require down/up eye movements and left/right head movements require right/left eye movements. But this gives no indication of the torsional eye movement. There is however a relationship relating the gaze direction to the ocular torsion.

Listing's law [3] and [6] provides this relationship qualitatively and quantitatively. For a stationary head in the upright position with both eyes looking towards infinity, it says:

When the line of sight is moved from the primary position to another position, the amount of torsion in this second position is such as if the eye had rotated about a fixed axis, which is perpendicular to the line of sight in the two positions.

There are several interesting features of Listing's Law which result from three-dimensional rotational mathematics. Perhaps the most interesting is that during saccades, the eye follows a sub-optimal trajectory, where the shortest route is optimal. Consider the following case [3];

Let the eye be inclined at an angle α degrees from the horizontal. Then the red axis is also rotated by α degrees from the vertical (green) axis. The blue axis represents an inclination of $\alpha/2$ degrees.

The shortest route for the eye to travel from position A to position B is by rotating the eye around the red axis and the longest route, is by rotation around the green axis. But the eye rotates around the blue axis, at an inclination of $\alpha/2$ degrees. The reason for the eye following this sub-optimal trajectory can be interpreted from Listing's Law.

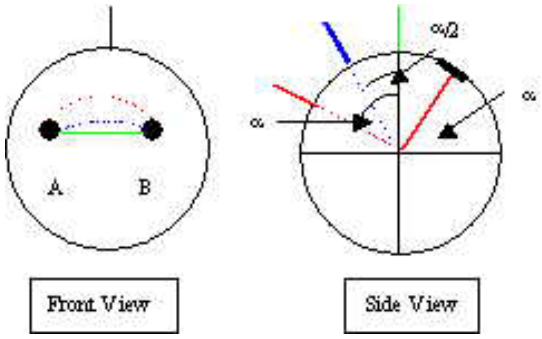

Figure 6. How the eye rotates. Left, two pupil positions upon the eye. Right, the choice of rotation axes between the pupil locations [3]

If the three dimensional rotation is represented in rotation vector form (a 3x1 vector, with each element representing an angle in one dimension), then for a given reference gaze position, all the rotation vectors describing eye movements, lie in a plane called the displacement plane. The orientation of this displacement plane is dependent on the reference position gaze position, so to describe the displacement plane for a particular set of eye rotations, the reference gaze position must be specified. In the special case

where the gaze position is perpendicular to the displacement plane, the displacement plane is called the Listing's Plane.

From the above discussion on the sub-optimal axis of rotation, it can be shown that describing one set of eye rotations with different reference positions leads to different orientations of the displacement plane. Moving the reference gaze position by 2α degrees moves the listing plane by α degrees as can be seen in the next figure. The first diagram shows the special case of the Listing's Plane and in each of the subsequent cases, the reference position is moved by 2α degrees and the corresponding displacement plane moves by α degrees [4].

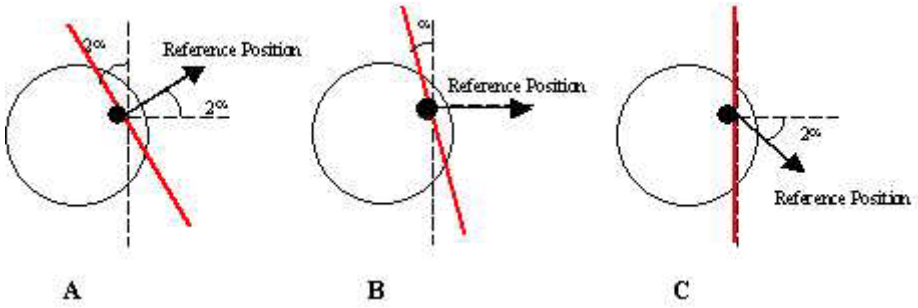

Figure 7. An example of how Listings Law applies [3]

So the movement of the eyes is a complex series of inter-related horizontal, vertical and torsional movements. The Impulse Test system does not need to account for Listing's Law because of the types of head rotations performed.

2.3 Current Techniques to Test the Vestibular System

The most accurate method for testing the canal planes and quantifying the VOR is with the Search Coil method [2], [4] and [5]. This method involves placing magnetic search coils on the pupils and performing the test in a magnetic field. This procedure is the most accurate currently available, but its highly specialised equipment and set up costs make it an impractical system for common use. The cornea must be anaesthetised properly before application and each coil can only be used five or six times. In addition, each coil costs in the region of $100 US and therefore, the running costs of the system can be very high.

The test is performed within a magnetic field, which requires calibration before each test can begin. Each coil must be tested after the calibration to ensure they are working correctly and there is no coil misalignment upon the eye.

3 Impulse Test System

3.1 System Aims

The Impulse Test System under development is intended to be a fully stand alone, automatic system for the detection of damage to any of the six semi-circular canals. The level of input to starting the system from the operator should be minimal for ease of use and so as not to alter any system properties set for specific environmental conditions.

The patient should sit approximately 50 centimetres in front of the camera, with the operator positioned behind. The first step is to perform a sequence of calibration tests required to describe the head model in detail. Once complete and the actual test has begun, the operator should not leave the patient to attend the system, as translational movement of the patient's head can result in moving outside of the camera's field of view. Therefore, the system was designed to be as fully automatic as possible. When the tests have been completed, the results are displayed on the screen and this gives an indication of the patient's state of health.

The results of the Image Processing can then be sent to a head model to describe the three-dimensional head and eye motion.

3.2 Software Environment and the Hardware. Details and Specifications

The current system is comprised of a single high-speed digital camera connected to a PC workstation with the associated software. The single camera poses specific challenges for the system in terms of two-dimensional to three-dimensional data conversions.

3.2.1 The Camera

The Pulnix TM-6710 is a high-speed monochrome device, capable of up to 600 frames per second (fps). The standard picture frame rate is 120fps, but the current Impulse Test software employs a frame rate of 240fps. Higher frequencies result in different scanning implementations onboard the camera, reducing the number of lines read from the Charged Coupled Display (CCD), which in turn reduces the maximum size of the image available to the system. This represents a conflict between high-speed imaging and the available screen size. Hence, a compromise must be found. (Refer to section 3.5.4). The camera produces images with 8-bit grey scale resolution, with a maximum Active Area of 648x484 pixels.

Figure 8. The Pulnix TM-6710 Camera

The images are digitised onboard the camera and communication with the Frame Grabber (Refer to section 3.2.3) is via the SCSI port. Some of the camera properties can be read from the device via the serial port, although certain properties, such as focus, cannot be set from the PC.

3.2.2 The Lens

The camera is attached to a 55mm "close up" lens for excellent picture quality in this application. At a distance of approximately 50cm, the resulting images are of high enough quality for the image processing to give excellent feature detail.

3.2.3 FRAME GRABBER NI 1422

National Instrument's IMAQ PCI-1422 frame grabber [9] allows up to 16 bit gray scale with a clock speed of 40MHz (total acquisition rate of 80MBytes/sec). It also contains 16MB of onboard memory for image storage and processing, although the current Impulse Test system transfers all the images directly to the PC for processing. It also provides extensive Region Of Interest (ROI) tools, for processing and analyzing image data sets. These facilities have been employed successfully throughout the developed system.

3.2.4 The Computer

The computer is a Dell PC with a PentiumIV 2Ghz processor. The system began with 512Mb RAM memory, but due to demanding requirements from the Impulse Test System software, the RAM was upgraded to 1024Mb. The PC runs under a Microsft Windows 2000 Service Pack 3 operating system.

3.2.5 Software

The programming environment used to develop the Impulse Test system is National Instrument's LabVIEW version 6.1. It provides all the functionality required to communicate with the camera for image acquisition and processing. LabVIEW is a graphical language, which can appear aesthetically pleasing to look at when used

properly. But if care is not exercised, the source code can quickly become unreadable and difficult to develop.

Math Works' Matlab version 6.5 was also used extensively for the physical modelling of the human head and examining signal properties.

3.3 Obtaining Data

3.3.1 Data Source

The system allows for data sets to be read from the camera, or from file locations. This allows for new patients to be examined and their results to be analysed, while the file option enables the operator to re-process results from earlier experiments if necessary. This re-processing may be necessary for improving the accuracy of the results due to different experimental conditions affecting the image processing algorithms. (Refer to section 3.6). Storing the data sets to file can also be useful for tracking a subject's condition, as comparisons can be made by processing image data sets recorded over moderate intervals of time.

3.3.2 Recording New Data From The Camera

Before the camera can start delivering images to the computer screen, the camera needs to be fully initialised with the properties required for recording the data. Upon starting the system, the user can select the camera option by clicking the camera button, which starts a series of sequences to initialise the camera and the associated memory management. Most of the low-level camera facilities are set in the Init.VI file, which is called first in event handler for the camera button. This file communicates with the camera to define parameters, such as the active area, and it also creates a buffer to store the data recorded from the camera. Most of this file was inherited from the original version of the system, although modifications have been made to the way in which the image window is created and initialised for more programming flexibility. The Window Setup.VI file incorporates these changes.

With the camera initialised and a buffer created for storing the experiment's image data, the next stage of the recording can begin. The system firstly reads in new images from the camera without storing them in the buffer. The purpose of this is to let the operator position the patient correctly relative to the camera, while observing the images on the screen. During this time, exercises can be performed to relax the patient's neck muscles, so they do not resist the head's motion during the tests. By not storing the images in the buffer, the system's memory is preserved for the test and is not used for storing data which contains no relevant information about the Vestibular system. As will be seen later, the memory storage capability of the system is limited due to the high frame rate and size of the images.

Once the initial, pre-test procedures are complete, the next stage within acquisition sequence is to obtain the calibration data. The first of these is to determine the

distance from the front of the eye to the forehead and the second is to determine the horizontal and vertical head radii. These quantities are required to accurately define the physical dimensions of the head model. For descriptions of the calibration processes, refer to section 3.5.

The next stage in the sequence is the collection of the data set, which will be used to perform the analysis upon. For this particular section of the test, it is of utmost importance that the patient sits at the correct distance from the camera. The reason for this is due to the image processing. The closer the patient can sit to the camera, the better the tracking will be, but the distance is limited by the camera's field of view. The image window has three regions defined upon it, one for the head marker and one for each of the eyes. These provide a rough guide for the operator as to where to position the patient. It is important that the head marker is within its designated region, as the calibration is based upon finding the marker within this region and determining its dimensions. The operator is then free to move the patient's head in any of the relevant directions to test the canal planes either until the testing is finished, or until the image buffer is full. Upon completion, the image buffer is sent through to the analysis stage.

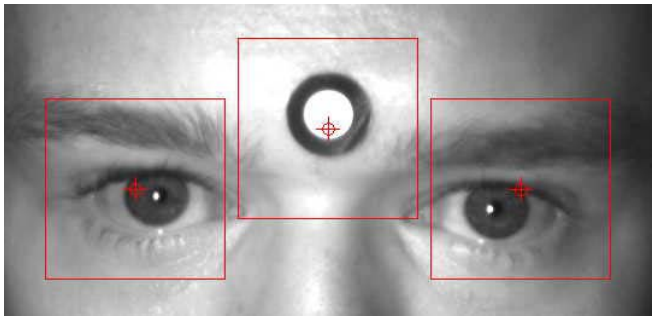

Figure 9. Positioning the patient relative to the camera

3.4 The Image Acquisition Buffer

The camera operates at a high frequency to capture as much information as possible about the head and eye movements. The frame rate is as high as 240 fps because the Impulse Test may only last for approximately 300ms and the distances, which the head and eyes move are relatively small. Therefore, as much detail as possible is required and 240 fps will result in approximately 60 positional head and eye values per test.

In order to examine all the canals, the experiments need to be performed in each of the relevant canal planes and there will also be time associated with moving the head back to its resting position at a much slower speed after each test. Therefore, many of the images in the buffer are not relevant to the calculations and can be neglected. But the fact remains that they are in the buffer and occupy computer memory. The buffer originally could contain up to 1000 full size (648x484 pixels) images but this corresponded to only four seconds of data recording time. This was deemed impractical, as each of the tests would need to be recorded separately and saved to file before the next test could begin. Therefore, the computer's RAM was upgraded from 512Mb to 1024Mb and the buffer has been successfully increased to hold 3000

images, corresponding to over 12 seconds of data acquisition, although the system performance during the analysis stage is greatly reduced. The speed could further be increased by closing all other non-essential system processes and with some modifications to the tracking algorithms. With a 2000 image capacity, the system appears to function at approximately normal speed. As the analysis is performed off-line, the speed of the system at this stage will not affect the data collection, or the quality of the results, and could therefore be tolerated.

3.5 Calibration

The system performs several calibration experiments in order to determine the physical size of the objects it sees. Several of these experiments are still in the development stage and they require extensive testing to ensure they perform as expected. Several of these calibration experiments have been devised after initial results from the head model showed that the model was sensitive to variations in the physical dimensions required to describe the head structure and in particular, the eye's location within the head.

The focus properties of the camera cannot be set automatically through the interface with the PC and therefore must be performed manually. As the system's performance is highly dependent upon the camera's focus, it was decided to try to perform the calibration tests without altering the focus settings. The results of the tests varied, depending on which property was under investigation.

The first calibration test is to determine the horizontal distance from the front of the eye to the front of the forehead, where the head marker is located. The patient should sit in a perpendicular direction to the camera, so the side of their face is in the camera's field of vision, with their head in the normal position. The operator is firstly prompted to select a region in the image window to describe the location of the head's calibration template. Only when the system recognises this template can it proceed. Once the template is recognised, the system can calibrate the image, extracting the relevant information needed to proceed with further tests. The user is then prompted to point out both the location of the front of the eye and the front of the forehead. With these positions and the calibration data, the distance can be computed in real world coordinates. This experiment worked well with proper attention paid to the objects in question.

The next stage is to determine the radii of the horizontal and vertical head rotations. The patient is moved back to face the camera and their head rotated slowly in the horizontal and in the vertical directions, by at least 30 degrees, preferably 40 – 45 degrees. By observing how the head marker contracts / expands, rotation angles are computed. An experiment was conducted to rotate the head marker in purely horizontal and vertical directions and the results were of varying quality for the different directions. The radii for both directions were 150mm.

- In the horizontal direction, the radius was computed by using the angular motion from the head marker and the lateral distance in which it moved. This method worked well, estimating the radius to within 5% of the actual length, but the results were less accurate in the vertical plane. The system could only

accurately describe the radius of the vertical motion to within 20% and this would yield a much distorted head model and eye velocity.

- Another method was attempted to measure the vertical radius. By rotating the head into a perpendicular plane to the camera, the trajectory of another small marker placed on the head was followed during the vertical motion. This method worked slightly better under certain conditions, producing results to within 15%. But when tested on a subject, small head translations appeared to significantly affect the result. Therefore, because of the difficulty in estimating the vertical radii, the value of the horizontal radius was used for that of the vertical radius. This approximation can be considered valid, as the Impulse Test is only performed with moderately small head movements, which can be approximated to the surface of a sphere.

3.5.1 The Head Marker Template

The system has been designed with the aid of simple object shapes, which can be easily recognised and measured by the system. The two standard geometrical shapes employed by the system are circular and rectangular markers, although each yielded varying results.

The first type of marker used was simply to aid the image processing track the head movements. It was a small marker with concentric circles to be placed on the forehead between the eyes. The outer circle was white and the inner, black, giving as large a difference in contrast as possible. The image processing algorithm inherited from the original system was a simple attempt at "colour thresholding", which searched for areas of high and low contrast. (Black = low, white = high). The same algorithm was then be used to track the eye, as it too could be approximated to a grey circle (the iris) surrounding an inner black circle (the pupil). The system then looked for the centre of the low contrast areas to represent the centre of the marker and of the eye. No calibration was performed upon this marker, although it did provide a basis for further development.

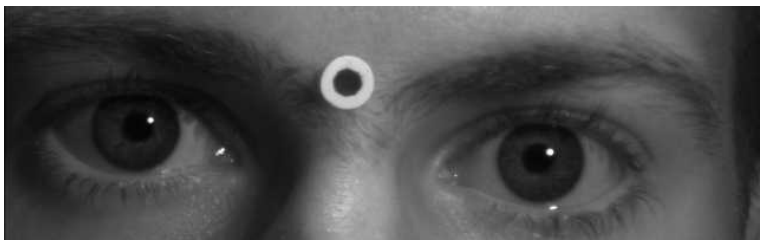

Figure 10. Head marker 1

The second type of marker was introduced after the image processing algorithms had fundamentally changed to the current version, based on pattern recognition, and therefore the marker was less constrained to the previous structure and colour. This too was a concentric circle, but with the colours reversed, black for the outer circle, and white for the inner. The reason for the switch of order was to enable the system to be able to estimate the size of the circles in addition to tracking them. The estimation was performed using a least squares algorithm on contrast levels, so large differences

in contrast were required. The forehead appears light grey when illuminated with a fluorescent lamp and therefore, a black outer region with a white inner region provided the greatest contrast differential.

Figure 11. Head marker 2

The third type of marker tested was a black square surrounding a white square for the same reasons as before. The decision to try a square came as a result of one of the calibration tests, which required knowledge of the head's rotation. When viewed in two dimensions, the marker on the forehead appears to expand and contract during head rotations and it was easier to determine the dimensions of a rectangle (a contracted square) than of an oval (a contracted circle).

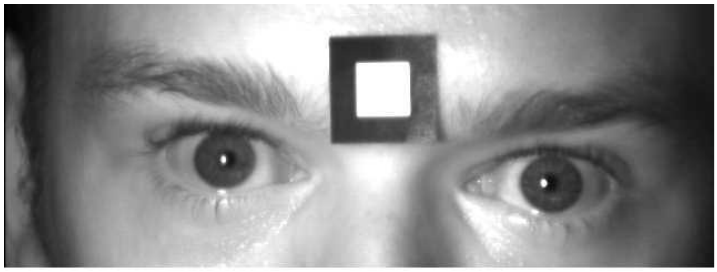

Figure 12. The current head marker

A facility exists to record an image of a new template from the camera and extract important information about the object's shape. In the case of a rectangular marker, the height and width, in pixels, of the inner and outer rectangles are calculated and stored to a text file. The physical dimensions in centimetres are also required and are also written to the file. In the case of a circular marker, the inner and outer radii in pixels are calculated and stored along with the radii measured in centimetres.

The image file and the associated text file can then be read and used in each of the calibration stages as required. The process of recording a new template should be performed as accurately as possible, in terms of camera focus, template position relative to the camera and minimising any rotation of the marker.

3.5.2 Estimation of Rectangles

The system can currently estimate a rectangle in the form of horizontal and vertical lines. This is quite restrictive as any torsion of the head cannot be accounted for and the estimated rectangle would be larger than in reality. The method to determine these horizontal and vertical marker dimensions came from the original calibration test

developed. This was to determine a scaling factor between pixels and real world coordinates assuming the patient was facing the camera, with their head resting in the normal position. The function to determine the size of the boxes was not needed to quantify the marker at any further stage in the experiment. But as subsequent calibration tests were considered necessary, the ability to recognise a rotated rectangular marker became more apparent. It was decided to continue with the current functionality and determine whether or not the calibration tests under development performed sufficiently well within "ideal" circumstances ie, during vertical and horizontal rotations only. Developing the functions to estimate the dimensions of a rectangular head marker, which can account for head torsion, is certainly an area for further development.

The function to detect and quantify the rectangles searches for a pair of horizontal and then vertical lines within a specified search area. Once found, the perimeter and the centre of the estimated rectangle are used to define a new search area for the inner rectangle. The search is performed around the centre of the outer box, within a rectangular perimeter twenty pixels smaller in height and width than the outer rectangle. This reduced area is especially important as it prevents an edge of the outer box being estimated again. This effect was noticed when the camera was not properly in focus and the result was that the boxes were estimated incorrectly or simply not recognised.

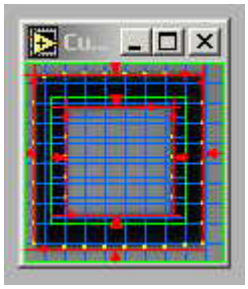

Figure 13. Searching for boxes. The green lines define the search regions, the yellow dots define the detected edges and the red lines show the result.

The vertical lines are estimated by defining a series of horizontal lines along which the contrast levels are detected. If the contrast levels change by a pre-defined threshold level within a pre-defined number of pixels, then the change was recorded as an edge. Then, with the result of each of the horizontal lines, least squares estimation was performed to define the vertical line. The same method was used for defining horizontal lines, although the search lines were vertical.

As the head rotates in the horizontal direction for example, the horizontal edges of the marker appear to contract and once this depth reduces close to twenty pixels, the estimation of the inner box begins to fail. Therefore a thick outer rectangle is required to allow the inner rectangle to be estimated. Testing showed that rotations of up to at least 45 degrees from the normal position could be achieved with a 21mm outer square size and 10.5mm inner square size. As the Impulse Test should never involve rotations of this magnitude, this marker size was assumed suitable.

The position of the marker on the patient's forehead is another major factor. When the head is in the normal position, the marker should be fully perpendicular to the camera

axis. (The “normal” position corresponds to the patient sitting in the upright position, facing the camera). If the marker is positioned on the slope between the forehead and the bridge of the nose, the system will estimate its size incorrectly, affecting all further calculations upon the resulting data. It was therefore required to position it as low as possible on the flat surface of the forehead above the eyebrows. The next issue resulted from the maximum available image size from the camera. During the vertical tests, the head marker disappeared from the camera’s field of view, especially when the head was moved in the upward direction. This must be prevented from happening, as the image processing defines search areas for the eyes based upon the head markers position.

3.5.3 Estimation of Circles

The estimation of the radii and centre of the circles in a head marker is performed in much the same manner as for estimating the dimensions of the rectangles. First the search area is specified to estimate the outer circle and then, the inner circle is estimated from within a region specified from the results of the outer circle. The same least squares method as before was used, but the search lines spanned the image radially from a defined centre point of the search area. During horizontal and vertical rotations, it proved difficult to estimate an oval with accuracy. It also would not yield any torsional information and hence the circle marker has become increasingly redundant.

The system can still track these concentric circles and perform some calibration experiments upon them, but due to the increasing trends towards the rectangular markers, development of new calibration tests has focused less upon the circular markers. It is expected that eventually they will become obsolete.

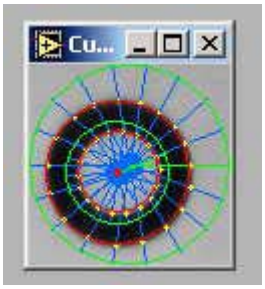

Figure 14. Searching for circles. The green lines define the search regions, the yellow dots define the detected edges and the red lines show the result.

3.5.4 Where to sit

As described in the section detailing the system hardware, there exists a trade-off between the image frame rate and the acquisition window size. This must be balanced in such a way that the patient’s head is always within the camera’s field of view, while the frame rate is high enough to precisely track the head and eye movements. With 240fps, the largest acquisition window available is 648x484 pixels and this determines how close to the camera the patient can sit. If they are positioned too close, the possibility exists of losing sight of the head marker or an eye. But too far away, and the quality of the image processing will deteriorate. The current camera

settings allow for a distance of approximately 50 centimetres, although this is subject to change depending on the final implementation.

3.6 The Current Image Processing Methods

The image processing techniques are performed with National Instruments' IMAQ Imaging and Machine Vision tools.

3.6.1 Setting the Parameters

Upon starting the system, several of the parameters used by the image processing can be set and changed. These are all given default values which appear to allow the system to function correctly, but this facility allows further improvements to be made with respect to the standard of the image processing and hence the results. They are as follows [10];

- Sub-pixel Estimation : Subpixel analysis estimates pixel values which would be available from a higher resolution image. This type of analysis is useful in this application because it can estimate positional head and eye values to a greater level of accuracy than would be available with only using the raw data.

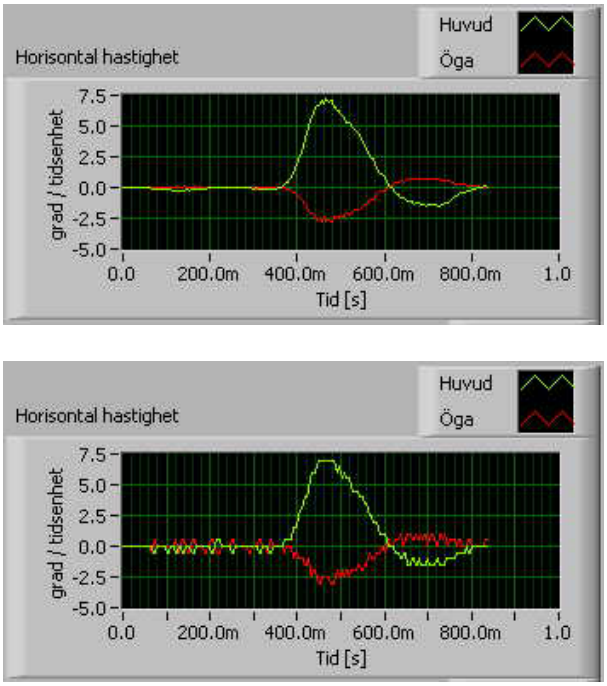

Figure 15. The effects of Sub-pixel estimation : Top, results with sub-pixel estimation. Bottom, raw data.

- Match Mode : The mode set here determines whether the pattern matching algorithm is restricted from finding a rotated version of the template or not. The Shift Invariant mode searches for the template pattern in the image within ± 4° (maximum) of rotation. The Rotation Invariant mode places no restriction

on the rotation of the template it is searching for, although it takes longer to perform.

- Minimum Score : This is the lowest acceptable match value which can be accepted from the image processing. Templates found with accuracy less than this value are classified as not found throughout the rest of the system.

- Minimum Contrast : This value defines the minimum range of contrast values expected in the image.

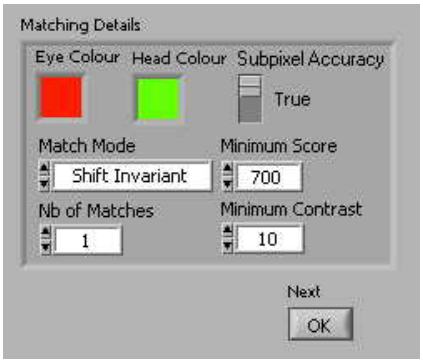

Figure 16. Image Processing Properties

3.6.2 The Procedure

After the Impulse Test has been completed, the system replays all the images stored in the buffer to extract the positional data of the eyes and the head. Before the replay begins, the operator is prompted to highlight both of the eyes and the head. This is performed by dragging the mouse over each of the objects in turn to form rectangular regions. Once each object has been highlighted on the screen, a button is pressed to initiate the first step of the image processing sequence. All the image data within each of the selected regions is "learnt" by the IMAQ Learn Pattern function, and placed into newly created image templates for storage.

Initially, the system was designed to track only one eye and the head template. From a medical point of view, both eyes move together in an almost identical manner and therefore no advantage would be gained from tracking both eyes. But for several reasons, the decision was taken to modify the software so as to track both eyes.

- The results from the Matlab head model show that significant errors can result when torsion is not accounted for. For small degrees of torsion, (less than a few degrees) it is possible to neglect torsional movements as noise introduced from the image processing can hide these effects.

- Particularly during lateral tests, it is possible that one eye could disappear from the camera's field of view. Instead of discarding the data from that particular test, the other eye could be used to provide the results.

- Also during lateral tests, as the head moves in the left/right directions, the left/right eye can become difficult to track as it rotates away from the camera and can disappear behind the bridge of the nose.

- Lighting could also become an issue. As the head rotates, the light intensity falling on different regions of the face varies, changing the contrast of the same areas over the rotation time. (Refer to section 3.6.6)

- Problems can also arise due to the shape of the region used to describe the eye. (Refer to section 3.6.5)

Although the capabilities now exist to track both eyes, extract the positional data and calculate several results from the tracking, the head model currently only accepts data from one eye and the head. It is only this stage of the system, which requires development to handle data from both eyes.

Once each of the three objects has been selected and their patterns learnt, the system iterates through the rest of the images in the buffer to search for the previously acquired image templates. All three searches occur together and the properties set at the beginning of the program are extracted and used at this stage. When each object has been found in the current image, it is highlighted on the screen to show the success of the pattern recognition software.

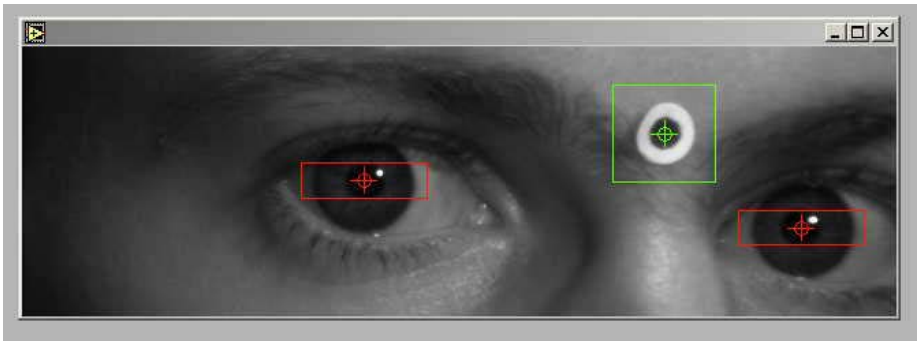

Figure 17. Image of a search

3.6.3 The Results

When a search has been performed on an object, the IMAQ Find Pattern function returns the number of matching templates found, the corresponding accuracy of the result and several other parameters. These values can be used to give the operator an indication of how well the experiment was conducted. They are used to calculate the following results for each object;

- The number of image frames where the template was not found.

- The mean accuracy from all image frames where the template was found. The mean accuracy is defined to be the mean of all the resultant template accuracies, which yielded a positive result. This deliberately excludes the templates not found, as it should give an indication of how accurately the algorithm worked, not how well the test was performed in the given

surroundings. This can be judged from the number of templates missed and the resulting graphs.

- The worst accuracy from all the images frames where the template was found.

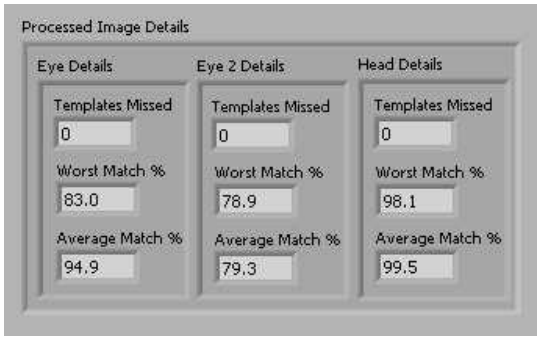

Figure 18. Image Processing Results

The exact details of the algorithms for the IMAQ imaging tools are not available in the associated help files, but the standard methods used for simple image recognition involve cross correlation algorithms. Assume a region w(x,y) has been selected from an image f(x,y). Let w(x,y) be of size KxL pixels and let f(x,y) be of size MxN pixels, where K ≤ M and L ≤ N. Then the correlation between the two images at a point (i,j) is given as follows [10];

$$C(i,j) = \sum_{x=0}^{L-1}\sum_{y=0}^{K-1} w(x,y) f(x+i, y+i)$$

But this basic method is very sensitive to the intensity of the data within the image. Therefore, it is common practice for these algorithms to calculate a normalised correlation function [10];

$$R(i,j) = \frac{\sum_{x=0}^{L-1}\sum_{y=0}^{K-1} (w(x,y) - w^*)(f(x+i, y+i) - f^*(i,j))}{\left[\sum_{x=0}^{L-1}\sum_{y=0}^{K-1} (w(x,y) - w^*)^2\right]^{\frac{1}{2}} \left[\sum_{x=0}^{L-1}\sum_{y=0}^{K-1} (f(x+i, y+j) - f^*(i,j))^2\right]^{\frac{1}{2}}}$$

where w* is the average intensity of the pixels in the selected region, w, and f* is the average value of f in the region from which w came. Therefore;

$$-1 \le R(i,j) \le 1$$

and is therefore independent of the intensity values in f and w. Obviously, this is not the actual method of implementing the pattern recognition tools, as the computational complexity is huge, although this forms the basis from which modern image recognition tools are formed.

3.6.4 The Loss of a Template

The result of the search for the head marker is critical as this marker is used to define the search area for both eyes in the subsequent frame. The current method of defining these search areas for the eyes is to split the screen into two sections, left and right. When a head marker is found, its centre coordinates are extracted and the screen divided into two vertical regions around this point. If the screen is not divided correctly, the detection of the eyes can become difficult because of the similarities between the two objects. The detection methods search the whole screen and the result is that one eye can be found in the current frame and the other eye in the next. This will obviously yield meaningless results.

The loss of an eye template is not critical to the system's performance as in the case of the head marker, but it is to the quality of the results from a particular test. There are several reasons why the eye templates would be not found and they are presented here;

- Blinking. On average, the human eye blinks every four or five seconds and each one lasts approximately 200ms to complete. There are different reasons for blinking and the two which are perhaps of most relevance here, are those which occur as part of a response to an unexpected stimulus and those which are protective, reflex blinks. If blinking occurs before or between tests, the effects upon the results are not so important. But if the blink occurs during an impulse test, the results should perhaps be discarded. As each Impulse Test lasts approximately 300ms and a blink is of the same order of time magnitude, the eye response could be hidden from the camera. During a blink, the system cannot detect the eye object, but once the eyelids open, the object is immediately found again. A simple technique to detect the onset of a blink has been developed, although it requires more testing and development. (Refer to section 3.8.5)

- As described before, during lateral tests it is possible that one eye could partially disappear from the camera's field of view, degrading the system's ability to track it. If this is the case, the results from the other eye could be used. This facility is not currently available because of the limitations of the head model.

- During vertical testing, it is possible for the patient's pupils to partially disappear behind the eyelid, hiding the region which the image processing is trying to track. The extent to which this occurs is dependent upon how far the patient's head is rotated and the features of their eye and eyelids. It was observed, especially in older patients, that drooping eyelids could cover portions of the eye required for tracking. It is therefore necessary for patient's to have their eyes open wide during the test.

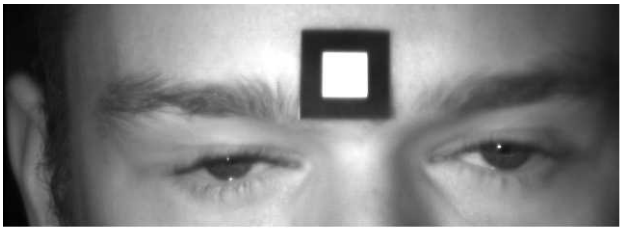

Figure 19. Pupil behind the eyelid

3.6.5 Difficulties with the Selection of Objects

The quality of the results obtained from the image processing algorithms is highly dependent upon the manner in which they are selected. As the objects appear to be described by their contrast properties, it was found necessary to provide as large a range of contrast with each selected area to provide accurate results.

The head marker was designed to be a simple, yet distinctive object, so that the possibility of the system recognising another object with a similar appearance was minimised. Selection of the best region containing the head marker was achieved by selecting the entire marker and a small portion of the forehead surrounding the marker. Therefore two large contrast differences result; the first is between the inner (white) and outer (black) rectangles and the second between the outer (black) rectangle and the forehead.

The selected region around the eye is rather more difficult to choose. The contrast levels available in the eye are less well defined but with the given structure of the eye, certain features can be used to the system's advantage. The three regions of the eye which provide greatest contrast levels are firstly the pupil, secondly the iris and thirdly, the sclera. The results and observations of several selection areas are described below;

- Selection of the whole pupil and a little of the surrounding iris was met with only limited success, as the search in subsequent image frames yielded a "noisy" velocity signal.

- The next attempt was to select the pupil, iris and part of the sclera for a larger range of contrast. This proved more successful during lateral tests, but upon vertical tests, as the eye moved behind the eyelid, the image processing struggled to track the template, degrading the analysis results.

- The height of the selected region was then reduced in an attempt to improve the system's performance during vertical tests. There was no noticeable deterioration in the signal quality from the lateral tests, and the vertical tracking worked until approximately 45 degrees of head rotation. As the Impulse Test should not require head rotations of more than 30 degrees, this was considered to be an acceptable result. But the noise introduced in the vertical direction was more than that in the horizontal direction. This could be as a result of having very little changing contrast in the vertical direction.

Hence, the best region found was a long, thin, horizontal rectangle which covered the pupil, the iris and the beginning of the sclera in the horizontal direction, and high enough to capture some of the circular properties of the iris and pupil.

3.6.6 Lighting

The first effect noticed from performing a test was that normal room lighting could not be switched on. The 240fps camera frame rate was much higher than the mains frequency of 50Hz and therefore the system produced images with varying contrast and intensity properties. A standard fluorescent desk lamp was used to illuminate the patient's head instead. The position of the lamp had to be chosen carefully as illuminating one side of the head more than the other also deteriorated the quality of the results.

It was noticed that in the pupil, there was a small white dot as a result of the fluorescent lamp illuminating the head. Initially, attempts were made to exploit this fact by including it in the selected eye region. A white dot surrounded by a black circle (the pupil) provided a large contrast difference over a very small area and this appeared to be a positive result of the lighting. But it was noticed that, especially in patients with a prominent disorder, this white dot in the eye was a serious drawback. When the head rotated and the eye had not yet compensated for the head movement, the white dot was still located on the front of the eye, not on its original position with respect to the pupil. This forced the image processing to compromise between the structure of the eye and the white reflection from the lamp.

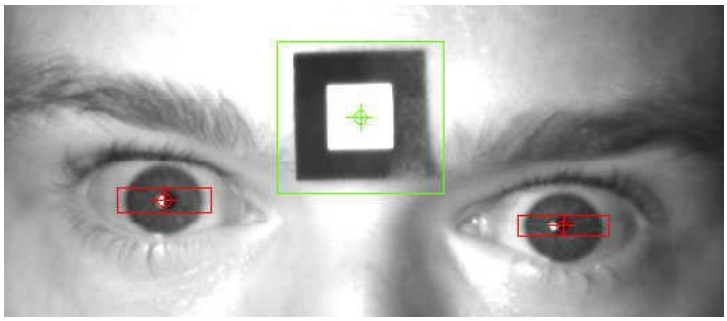

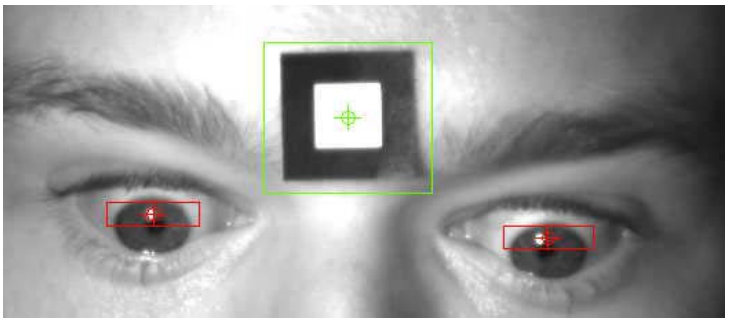

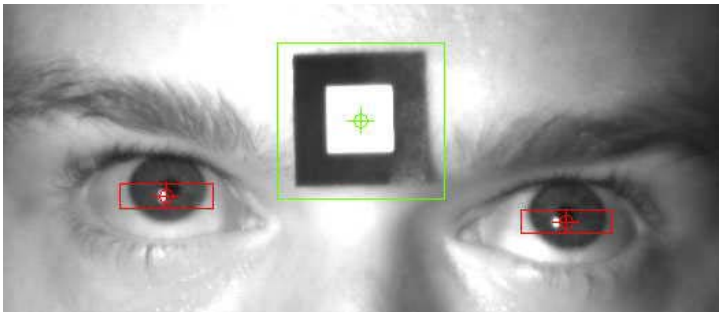

Figure 20. Effect of the white dot. Top, the initial image, with correct tracking. Middle and bottom, subsequent frames showing poor tracking qualities as a result of the reflection

As will be described later (Refer to section 3.8.3) this white dot was again turned into an advantageous side effect of the lighting. One other option to illuminate the head would be to use an indirect method of lighting. For example, shining a strong light upon a white umbrella positioned in front of the patient would illuminate the subject's head. This would remove the white reflection associated with the fluorescent lamp and it could also illuminate the head to a greater extent during rotation.

3.7 Further Image Processing Developments

The Image Processing algorithms employed in the current Impulse Test system perform accurately under certain experimental conditions. The pattern recognition tools allow for the head and eye objects to be tracked well during analysis of normal subjects, but experimental set up properties, such as lighting, pose challenges for the system. Other methods exist to track defined objects, and some of these were experimented with, although not employed in the current system due to time constraints.

The main difficulties with the current image processing methods which must be overcome are summarised here.

- The reflection from the fluorescent lamp.
- The eye becoming hidden behind the eyelid and partially hidden behind the bridge of the nose.
- Poor selection of the eye object for tracking.

Attempts were made to investigate different approaches in order to solve these issues. Two of the more common approaches in estimating the location and size of the pupil within the eye are the Edge Detection and the Area Detection methods.

3.7.1 Edge Detection

Edge detection processes usually involve locating the boundary between the iris and the pupil. This is a variation on a process which has already been employed in earlier stages of the system to obtain the circular head markers. In the case of identifying the pupil, it is theoretically a simple procedure to follow, as methods exist to locate the edges of common shapes, such as circles.

The first attempt at implementing an edge detection algorithm was met with poor results. The contrast levels between the iris and the pupil were generally relatively small and detecting a well-defined boundary was extremely difficult. The results were highly sensitive to the focus and even when boundary was located, the eyelid occasionally came into close proximity to the pupil, distorting the estimation.

When the reflection of the lamp was within the pupil, the pupil/iris boundary could be estimated moderately accurately, but when it fell on the boundary between the pupil and the iris, the resulting outliers distorted the estimated size of the pupil and hence its centre. As the pupil was normally small in diameter due to the lamp illuminating the face from close proximity, it quite often occurred that the reflection moved across the pupil-iris boundary. If this were implemented in the system, results could show that the pupil was moving while in fact it remained stationary. Therefore, this method on its own would not be practical for the final implementation.

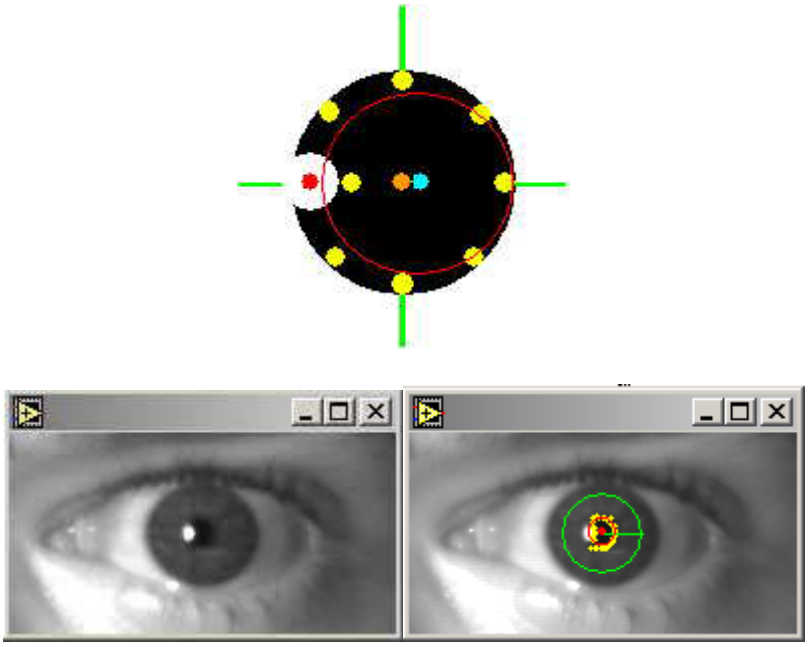

Figure 21. Top. A simulation of the effects of the reflection upon edge detection. Bottom left, an image of the eye with the reflection within the pupil. Bottom right, the same image with circular edge detection performed.

3.7.2 Area Detection

Area detection methods can, in general, be more precise than edge detection methods, depending on the application, as they attempt to determine the centre of energy within an area. This is commonly implemented with the Centroid technique, which assumes the pupil is totally black in order for the centre of energy to correspond with the actual centre of the pupil. The main problem with this assumption is again due to the lamp's reflection. With the reflection within the pupil, or upon its edge, the centre of energy was distorted, hence suggesting that the eye had moved to a different position, as was the case with the edge detection algorithm.

Figure 22. An image using the centroid technique.

3.7.3 Thresholding

Another commonly used method to assist image processing algorithms is Thresholding. This is a technique to simplify an image by reducing the number of image colours. This procedure removes a lot of unnecessary information, but retains enough to describe the original image in sufficient levels of detail for further processing.

Figure 23. Thresholding. Left, the original image. Right, the image after thresholding. Two threshold levels were employed. (100 and 200)

3.8 Suggested Technique

In order to determine accurately the position of the pupil, a series of procedures, combining all of these techniques was employed to describe the eye's image in a heavily reduced level of detail, while retaining enough information to extract the relevant features of the eye.

The first attempt involved obtaining a binary image (1 bit grey scale) by setting only one threshold value, splitting the image into black and white only. This separated the strong white reflection from the rest of the image, allowing for quantifying its size and intensity. But no other objects could be obtained without introducing unwanted objects into the final image. Therefore, the resulting image was required to contain more detail.

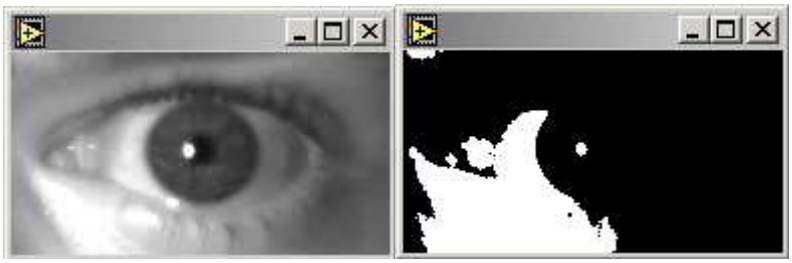

Figure 24. Binary Thresholding. Left, the original image. Right, the binary image. (Threshold level 170)

The next step was to develop a multi-threshold version of the previous method. The choice of threshold levels was non-trivial and results varied depending on the number of, and values assigned to, the thresholds.

By setting three levels, the eye was characterised in terms of the reflection as a white region, the iris, pupil and eyelashes as a black region and all other points as grey regions. The pupil could not be distinguished from the iris and hence difficulties were encountered with defining the centre of the pupil. The iris appeared as a single black object connected to the eyelashes and so attempting to describe the iris as an oval yielded poor results.

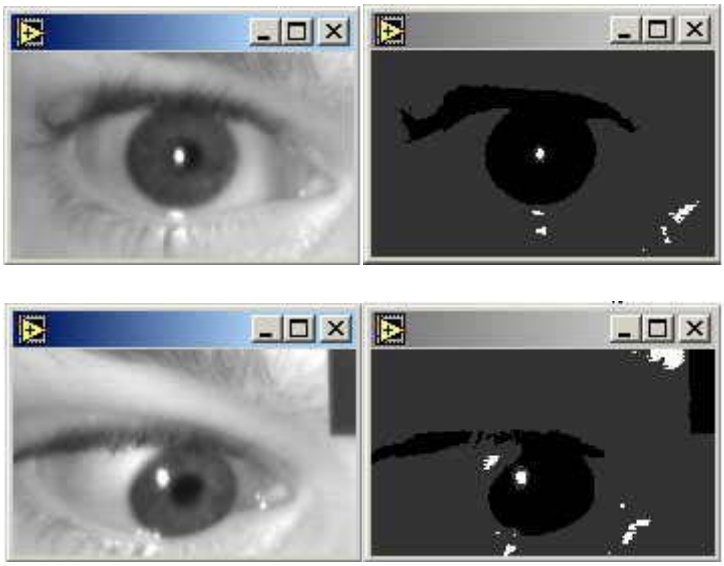

Figure 25. Images of the same eye taken from the same test sequence. 2 threshold levels were applied. (100 and 240)

The number of contrast levels were increased to six and the resulting eye structure was characterised in sufficient detail to begin extracting and defining features. The threshold levels were defined as follows;

Lower Threshold	Higher Threshold	Resulting Value
0	50	0
51	100	51
101	150	101
151	200	151
201	240	201
241	255	255

These values were manually obtained from a test set of data and it is necessary to use different threshold levels for different subjects. An automatic method to define the threshold levels was developed and is detailed below.

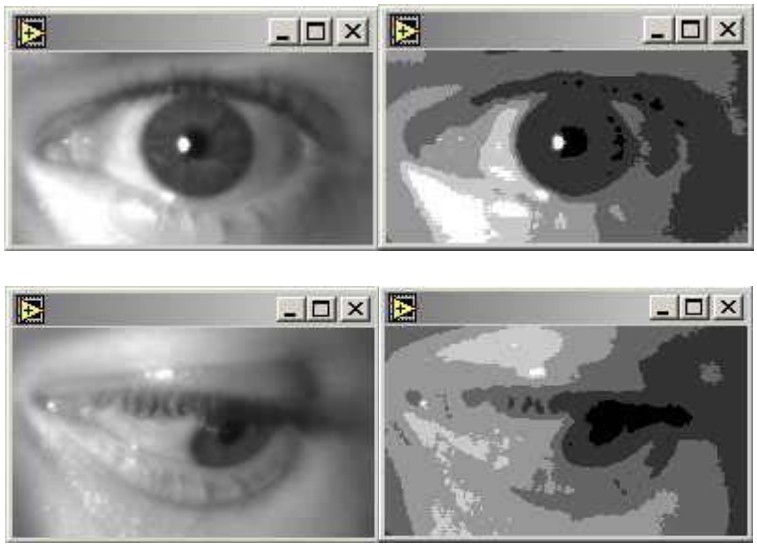

Figure 26. Images of the same eye taken from the same test sequence. 5 threshold levels were applied. (50, 100, 150, 200 and 240). Compare with the previous figure.

3.8.1 Light Intensity

The threshold values defined earlier were obtained manually form a test set of data. When applied to other data sets, the resulting threshold images were generally poor. Therefore, a facility to obtain the light intensity levels of the iris was developed. The user is prompted to select a region on the screen corresponding to the iris, without the pupil or the reflection. Within the resulting concentric circles, four rectangular regions are defined automatically to quantify the grey scale level. This facility returns the upper, lower and mean intensity values, as well as the standard deviation. The lowest value obtained from all four rectangles is used as the first threshold value, and all other values are defined as shown earlier. This technique could also be applied to other regions of the face in further work to define the remaining threshold values. Applied to the same set of data, the lower threshold value returned was 31 (compared to 50 earlier). The resulting sequence yielded better results, although issues remain with respect to making the system more robust to eyelashes.

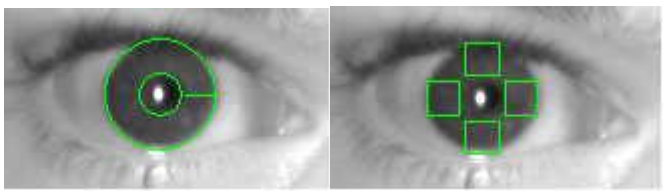

Figure 27. Left, the iris is selected. Right, it is split into four regions to determine the light intensity for the threshold levels

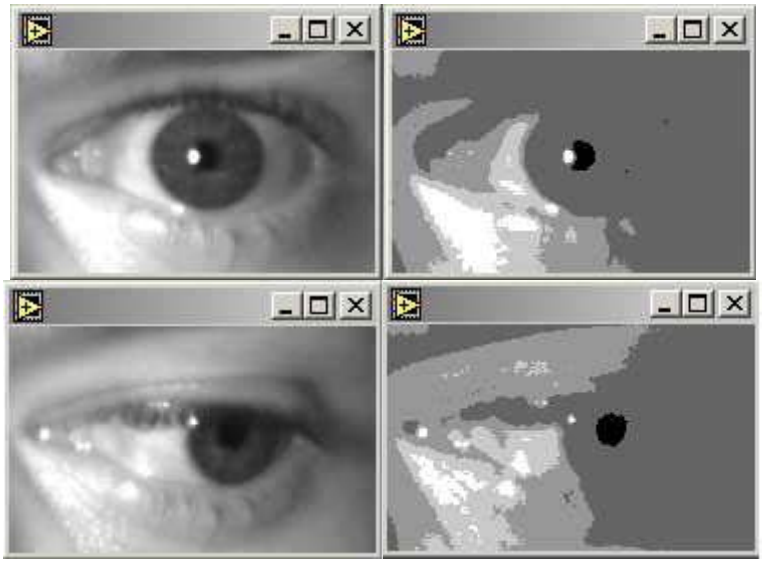

Figure 28. Images of the same eye taken from the same test sequence. 5 threshold levels were applied. (31, 100, 150, 200 and 240). Compare with the previous two figures.

3.8.2 Determining the Pupil's Location

These values allowed the pupil and the reflection to be distinguished from the rest of the image and therefore, both objects could be quantified. By considering the black pupil on its own as representing the entire pupil, an initial estimate of the centre position of the pupil could be obtained using the centroid method. But as the reflection was often located within the pupil, neglecting this region resulted in an offset in the central coordinates. Furthermore, this offset was not constant as the reflection did not always occur in the same location within the pupil and when it moved out of the pupil altogether, the offset disappeared. By using particle analysis upon the pupil region, it was possible to determine whether the reflection was firstly within the pupil region and secondly, apply the centroid technique to it. Therefore, knowledge of the black pupil region and the white reflection could be used to approximate the centre of the pupil more accurately.

When both the entire pupil and the dot were described as circles, the method of determining the correct coordinates of the centre of the pupil could be done as follows [1];

Let the pupil be of radius r, the reflection of radius d and the area of the reflection be S;

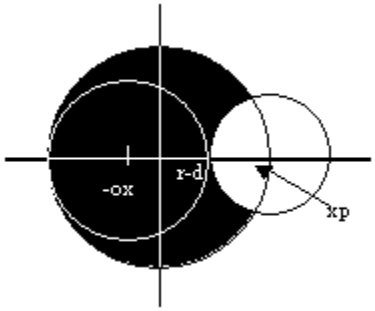

Figure 29. An image of how the pupil becomes obscured by the reflection [1].

$$S = \int_{r-d}^{xp} 2\sqrt{d^2 - (x-r)^2}\,dx + \int 2\sqrt{(r-ox)^2 - (x+ox)^2}\,dx$$

where xp is the x coordinate of the intersection between the pupil and reflection boundaries.

$$xp = \frac{r^2 - r \times ox - d^2/2}{ox + r}$$

Unfortunately, due to the nature of the thresholding and the experimental set up, it was not always possible to describe the detected region of the pupil as a complete or partial circle. The transition region between the intensive black and white regions of the pupil and the reflection resulted in a gray area which was not accounted for in the above method. This area merged into the surrounding iris, and hence developing this method further proved extremely difficult. As the pupil did not have a well defined circular outline, it was difficult to describe it fully as a crescent, and hence a circle. The figure below shows how the pupil actually appeared and how it was assumed to appear.

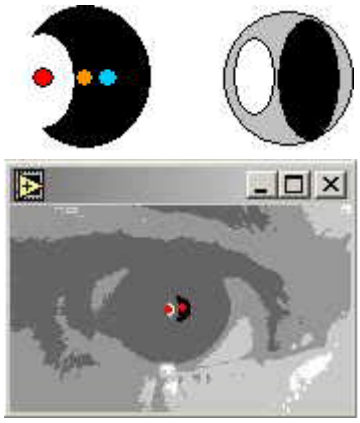

Figure 30. Top left, how it was assumed to appear. Top right, how the eye appeared. Bottom, a real image showing this effect

Therefore, a simpler method to account for the reflection was attempted, yielding positive results. Using the first frame, the size of the pupil was estimated in pixels and hence whenever the white dot was inside this area, it was accounted for in the following manner;

$$Pupil\ Centre = \frac{(Estimated\ Pupil\ Centre \times Area) + (\text{Re}\,flection\ Area \times Area)}{Total\ Area}$$

The increase in accuracy with only this simple correction method was significant. With a pupil of radius 10 pixels, the offset was corrected by 2 pixels on average.

3.8.3 Using the Reflection to Describe the Eye

As will be described later in the head modelling section, the biggest challenge to quantifying the ratio of eye to head movement is accurately describing the location of

the centre of the eye. Errors of only a few pixels can result in the system describing eye movements which may not have actually occurred. Therefore it is imperative to locate both the centre of the pupil and of the eye's origin accurately.

Locating the eye's origin within the head proved difficult, but using a technique to track the reflection, one can determine how the eye has actually moved. As the illuminating lamp is situated in a fixed position relative to the camera, its reflection on the cornea will always be at the closest point to the camera. Hence, a straight line from the camera to reflection projected further into the eye will intersect the eye's origin. Therefore, tracking eye movements relative to its own rotation point becomes much simpler than predicting its rotation point from the head model.

The anatomy of the eye requires a model to be implemented to account for the position of the reflection on the eye. Most of the eyeball is hidden within the head and it can be assumed spherical, but the structure is more complicated than that. The front of the eye protrudes out from the spherical region and forms another curved surface. This must be accounted for in the final model.

Several techniques exist to track the reflection by comparing its shape, but in this context, only its area, intensity and centre are exploited. By assuming it is a perfect circle, it could be quantified by using the above properties, although this was not implemented. Therefore with the centre of the pupil and of the eye described, one can project this two-dimensional data on to a three-dimensional model.

There are several difficulties and limitations with using the reflection to track the eye movements. The shape of the reflection distorts due to the shape of the eye during rotation, but as the centroid technique searches for the centre of energy, this effect can be minimised. More importantly, this technique will only work for relatively small head movements, as the tracking is not reliable once the reflection approaches the edge of the iris/sclera boundary. The white sclera appears to merge with the reflection and therefore the centroid technique becomes less accurate. Testing on a patient with a severe disorder in the right lateral canal showed that the head movements involved were not large enough to result in the reflection moving from the pupil to the iris sclera boundary. It remained well within the iris, justifying the method to use it as a way to determine to the eye's location.

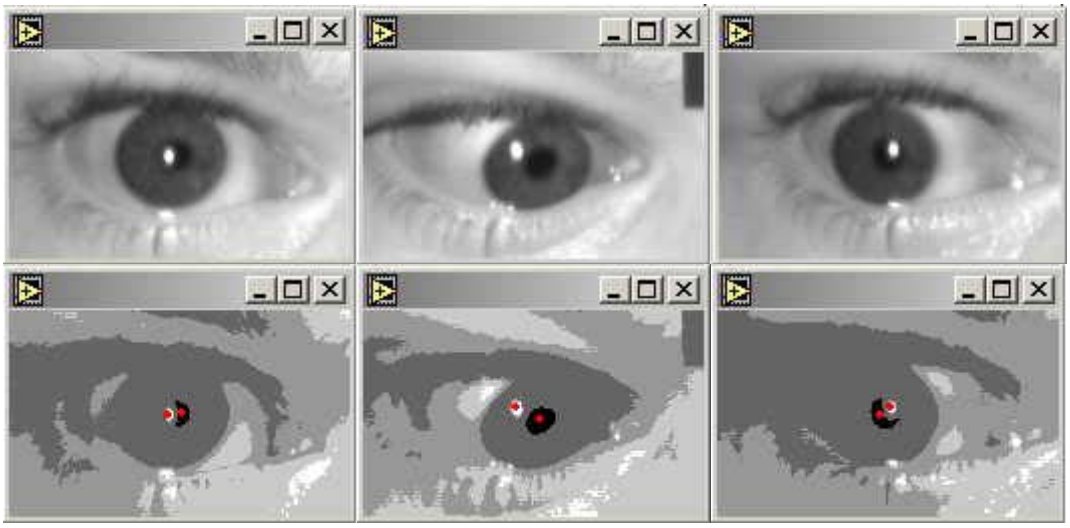

Figure 31. A sequence of eye movements with the pupil and the reflection both tracked. Left, the normal position. Middle, looking down and left. Right, the reflection on the other side of the pupil, while in the resting position

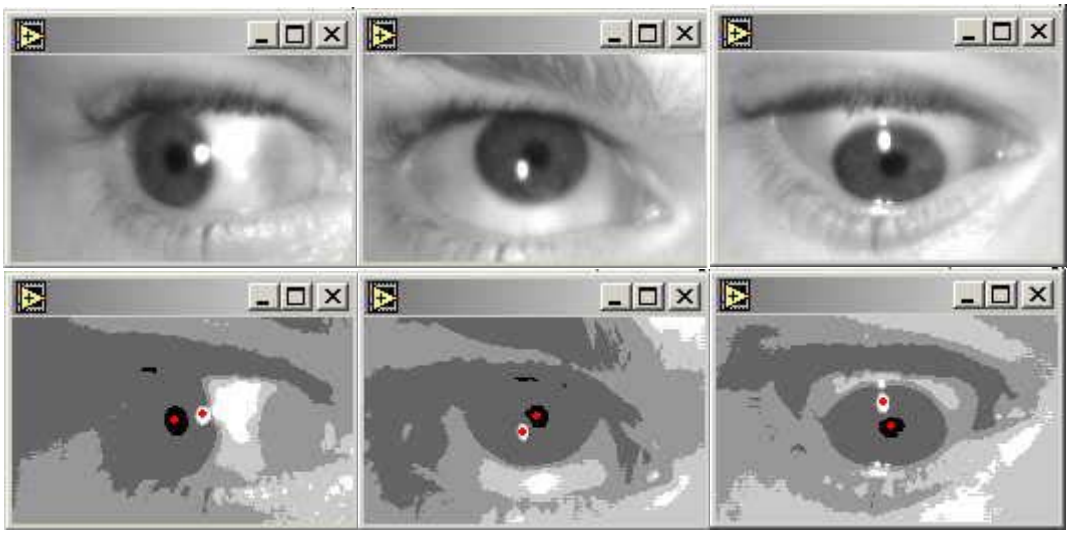

Figure 32. A continuation of the same sequence. Left, looking right. Middle, looking up. Right, looking down.

Another advantage of using this technique to determine the centre of the eye is that the physical size of the human eye varies much less between one patient and another than the required dimension of the human head. This should produce much less variation in estimating the rotation points, and hence better eye velocity signals. A comparison between the different implementations is described in section 5.1.

3.8.4 Eyelashes and mascara

One common difficulty with other medical systems employed to track the eyes is distinguishing the pupil from eyelashes. This method has been able to track the pupil though the eyelashes, although with a small degree of error.

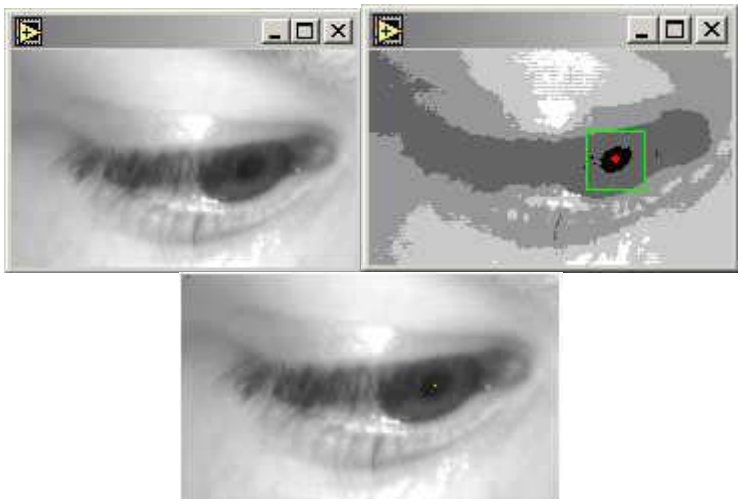

Figure 33. Detecting the eye through the eyelashes

A more challenging version of this problem presents itself amongst women wearing mascara. The eyelashes appear darker in colour and often the individual eyelashes clump together, creating more intensive areas than without wearing mascara. Nevertheless, this method still removes much of the associated dark regions, but still requires further development work to handle cases when it cannot fully distinguish between the pupil and other unwanted dark regions.

3.8.5 Handling Blinking

With the previous approach to the image processing, if an eye template is lost due to blinking, it can be re-established again when the eyelids open because of pattern matching. But this is not the case with the method described here. This works by looking at threshold contrast levels, and hence the search area cannot simply be widened to look for the pupil until the eye has re-opened. Therefore, the system must be able to detect an eye blink and prepare itself for the eyelids re-opening.

A simple method to detect an eye blink was created and implemented. The user is prompted to click a point on the upper and lower eyelids and on both corners of the eye, on the first image frame. The upper and lower eyelid templates are rectangular to encapsulate the eyelashes and the eyelids, while using as little of the underlying iris and sclera as possible as these will become hidden. The two side markers are square to capture as much information around them as possible, without using information from the neighbouring sclera, as this will also disappear from view during the blink. Using pattern recognition techniques alone yielded poor results as the patterns changed during the blink.

Thresholding was then employed alongside to this method and the results were better, although not perfect. As the top eyelid moves down, the shape of the eyelashes changes and they eventually hide the bottom eyelid. This results in the loss of the lower eyelid and the system struggles to recover it again. The top of the eyelid is also difficult to track due to the change in shape during the blink.

When a template is found, an area around its central coordinates is defined to search for the object in the next frame. If an object cannot be found in the current frame, the search area around the previous successfully found template is widened and applied to the next image. This technique normally worked well and the system could recover the original object within two or three frames. The exact details of this algorithm need refinement in order for it to be successfully and automatically applied to other subjects.

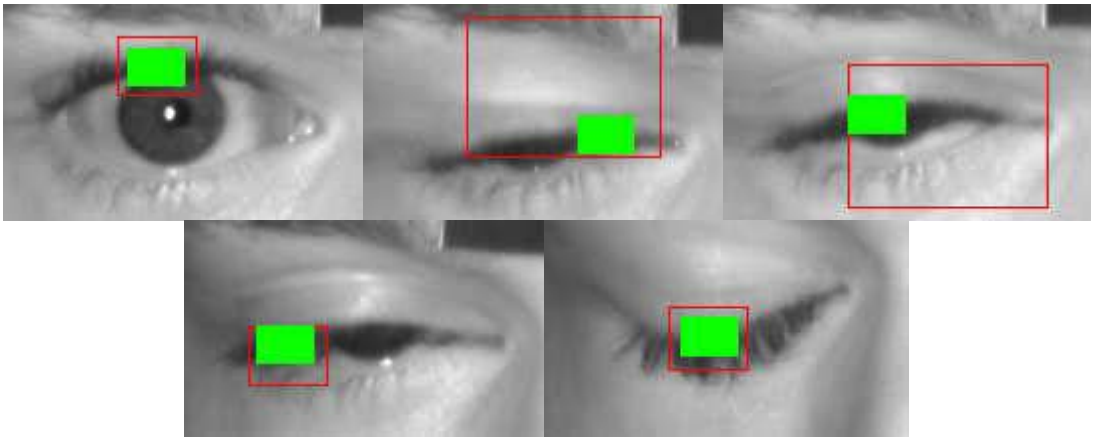

Figure 34. A sequence of images showing how a blink is recognised and tracked.

3.8.6 Advantages of this Image Processing Implementation.

This method of implementing the eye tracking has so far proved successful in testing and it also eliminates some of the problems associated with the current implementation in the Impulse Test system.

- The reflection from the fluorescent lamp has become beneficial to the image tracking implementation without the same restrictions as before. It works well for test data and more importantly, real experimental data.

- If the eye becomes partially hidden behind the nose, this method will still work, as it does not rely upon recognising patterns across the eye. It works only upon data centred on the pupil.

- The effects of poor selection of the eye object for tracking are no longer present, as no object structures are tracked.

Although this algorithm appears to work more accurately than the pattern matching, it has not been implemented into the system due to time constraints. This certainly provides the foundation for further work.

4 Head Model

4.1 The Human Head

The head is connected to the Cervical Vertebrae at the base of the skull and it is here where the head rotation is centred. The head's horizontal rotation point is quite well defined because the spine at this point is approximately vertical and if head torsion can be neglected, the rotation can be located accurately to be at a fixed position. The head's vertical rotation point is more difficult to define as one fixed point. The head can rotate vertically around approximately two points on the vertebrae, and this can be demonstrated by rotating one's head about the top of the neck and then the bottom of the neck. These points approximately correspond to the first and eighth cervical vertebrae segments.

4.2 The Current LabVIEW Model

The analysis of the data displayed in LabVIEW came mostly from the inherited system and therefore only a very brief overview of its operation will be given.

The two-dimensional coordinates extracted from the images are converted to angular position with the use of a simple two-dimensional model. The head movements were all measured relative to the first head position and the eye movements were all measured relative to the current head position. Then, by simply scaling the results with a constant value, the two dimensional position values were converted to angular position. This approximation was made as the Impulse Test only occurs over small angles and therefore results in relatively small errors. These angles were then differentiated to yield angular velocity and formatted to yield graphical results.

This two dimensional model will not yield any accurate results as eye position was measured relative to the current head position and not its own centre of rotation. Therefore, even if a patient was able to perfectly fixate their eyes on a target straight ahead, this model would predict that their eyes could not compensate correctly. Therefore a more detailed three-dimensional model was required to determine how the eye moved relative to its own origin.

4.3 Matlab Model

The head was modelled in detail with Matlab, using knowledge of the human anatomy. The data available to the model was in the form of two-dimensional position values and it was required to convert this into angular velocity signals. It was assumed initially that as eye torsion was not being measured and as the Impulse Tests were performed in the horizontal and vertical planes as much as possible, head torsion could be neglected.

4.3.1 Head Design/Structure

The idea behind the modelling was to represent head rotation points, the head marker and the eye within the head so all motion could be calculated within the model. The head was defined to rotate in only the horizontal and vertical directions about two rotation points on the cervical vertebrae. The first was the vertical rotation point, which was defined to be the stationary origin of the model and from this point, all other object locations were defined. The horizontal head rotation point was then defined to be a short distance from the vertical point and it rotated in a vertical plane around this point. As the model was being developed, approximate values were chosen to represent the physical dimensions of the head structure. The distances and angles from both rotation points to the head marker and the eye origin were required to map the two dimensional pixel data on to the model. The horizontal and vertical radii were also defined as two separate quantities, allowing for the fact that both are not necessarily the same. These are all variables which must be defined for each set of results.

The initial version on the software tracked only one eye and the model was developed based upon this information. Several stages were developed in an attempt to describe the head structure and trajectories for different input signals.

4.3.1.1 Graphical Implementation

The first stage was a visual structure of the head, detailing both eyes, and the connections to the horizontal and vertical rotation points. The file accepts inputs for both head and eye movements in the form of angles and it displays a short animation of how the structure moves over the course of rotating from its input angles, to the negative value of the angles in each plane. This was developed to show that the initial model was structurally and mathematically correct. The need to extend this model to accept pixel data was not considered relevant, as results could be assessed more accurately through the velocity signal graphs.

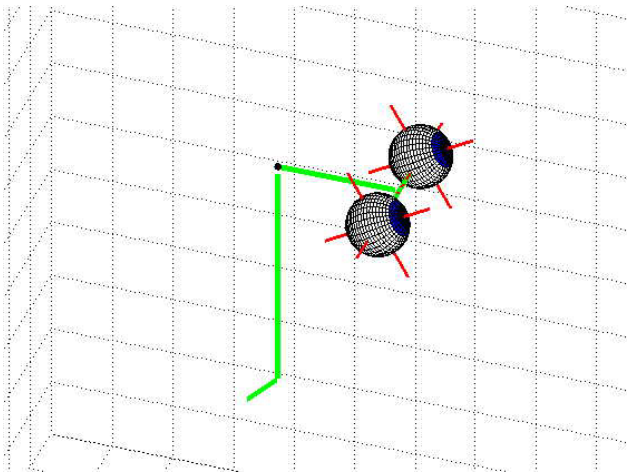

Figure 35. A diagram of the head model, with eyes elevated by 30 degrees.

The end points of the small green line at the bottom of the figure represent the two rotation points, while the two larger green lines attempt to give some perspective to the image, but represent no physical dimensions.

4.3.1.2 Trajectories

Due to defining the horizontal rotation point relative to the vertical point, the horizontal rotation point rotates in a circular trajectory around the vertical rotation point as the head rotates in the vertical plane. The following figures show how the horizontal rotation point moves, relative to the stationary vertical rotation point, and how the head marker rotates. The smaller inner arcs represent the horizontal rotation point's trajectory, while the outer lines represent how the head marker rotates horizontally for several vertical elevation angles. The colours simply define when the head has rotated above, blue, and below, red, its normal resting position.

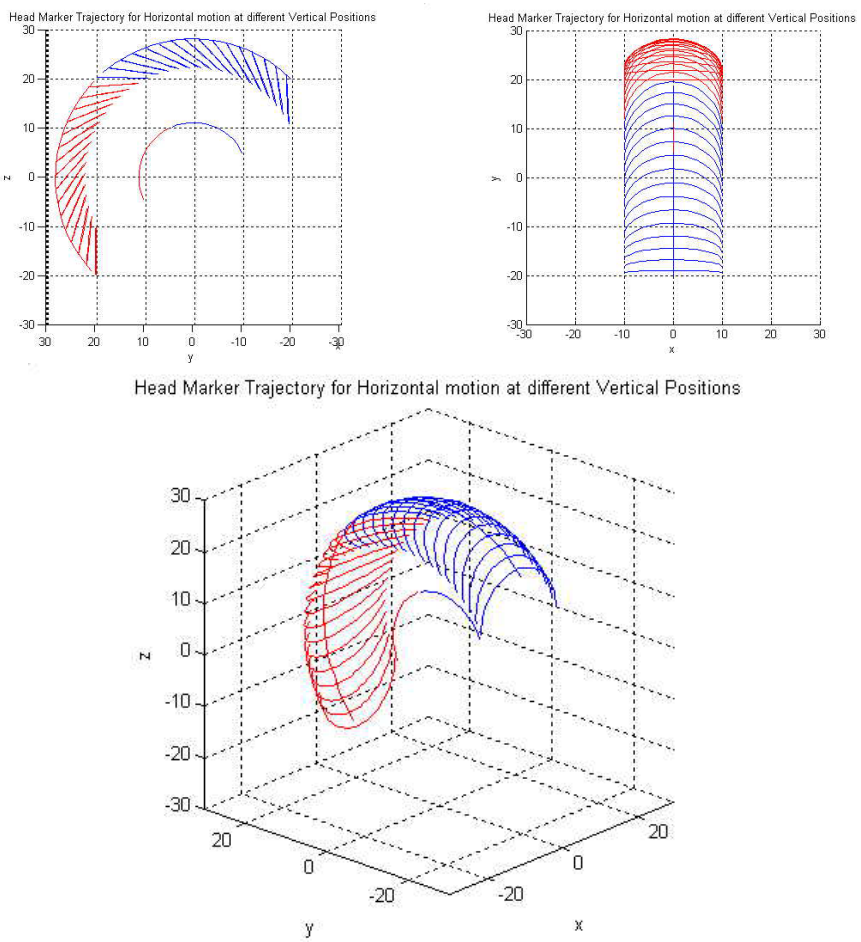

Figure 36. The trajectories of the head marker and rotation points. Top left is from a side view. Top right is from a plan view. Bottom shows the 3D result

It is obvious that motion of the head around the horizontal rotation point does not yield purely horizontal head motion. The actual trajectory is dependent on the level of vertical rotation. It is easy to visualise in the following way. By rotating one's head vertically so as to face the ground and then rotate the head about the top of one's neck (the horizontal rotation point), the trajectory of a point on the forehead moves as a swinging pendulum. This corresponds to the bottom most trajectory in the top left graph above.

4.3.1.3 Limitations

This head model described here has several limitations which should be overcome if the model is to accurately describe the complete head motion. Firstly, torsion of the head must be modelled. This could be performed by determining the torsional rotation angles of the head marker and sending them to the head model. Secondly, head translation is not modelled either. By tracking two points on the face, such as the outer corners of each eye and measuring the distance between them, the system should be able to determine whether any motion of the head involves rotation or translation. In the case of rotation, the distance between the points will decrease and with translation, the distance will remain constant. Currently, all motion of the head is defined as rotation and hence, any translation will result in errors. Thirdly, as described in section 4.1, the vertical rotation point will not necessarily always be constant. Although only small rotations are involved here, this difference in possible rotation points is likely to result in large eye velocity errors. As will be described in section 4.3.3, the computed eye velocity is sensitive to modelling parameters such as the head's radii. Overcoming this issue would be extremely problematic as determining the values for the radii is performed before the test begins and if the head was rotated differently during the actual test, then the radius would be incorrect.

4.3.2 Locating the Eye's Origin

The major problem with the simple two-dimensional model employed by the initial system was that the eye velocity was not measured relative to its own rotation point. It calculated the two dimensional velocity of the centre of the pupil relative to the head marker. This yields a result with no direct physical interpretation. The eye motion must be defined about its own origin and the head model attempted to describe the eye's origin based on angular motion from the head and predefined physical dimensions. By using the difference between the recorded eye movements and the predicted location of the eye's origin, eye velocity could be quantified.

The figure below shows the result of projecting the two-dimensional data upon a three-dimensional model. In the top left graph, one can see a small discrepancy between the predicted eye location (blue) and the recorded data (black). This is partly due to a small degree of head torsion, but also inaccurate dimensions have yielded this result. This represents a frontal view, so both of the eye signals should be as close together as possible. In the top right graph, the eye signals are further apart as the view is from the side and the difference corresponds to the eye's radius. The bottom left graph shows a three-dimensional view and the bottom right shows the resulting head and eye velocities.

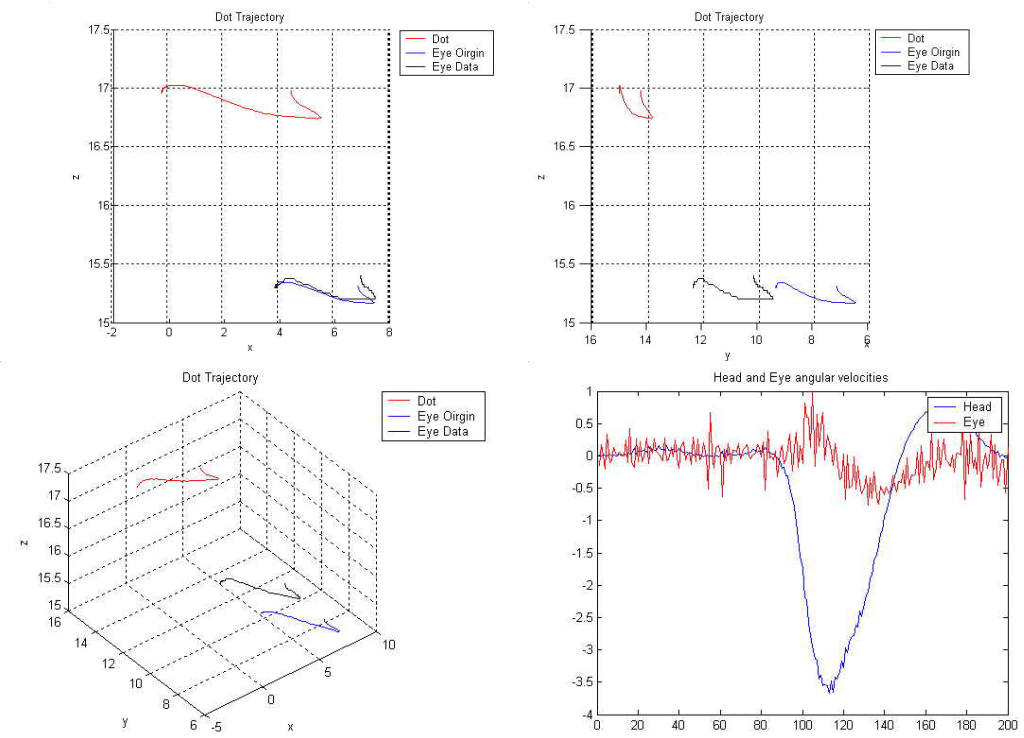

Figure 37. 3D modelling results. Top left, a frontal view. Top right, from a side view. Bottom left, a 3D view. Bottom right, head and eye velocity.

The difficulty with this method is that the radius of the eye is approximately 2.5 centimetres and as the impulse test only rotates the eye by 30 degrees, the actual movement of the pupil is very small. Therefore, the origin had to be described extremely accurately otherwise, the pupil's velocity would be computed incorrectly. This was a big problem which the calibration methods struggled to solve adequately. As it was initially decided to perform the calibration tests without altering the camera properties, the head's rotation points were not visible to the camera and they had to be estimated relative to the head marker. Errors introduced here will significantly alter the location of the eye's origin, immediately introducing errors into the resulting eye velocity.

4.3.3 Results

The model yielded both head and eye velocity results, which were used primarily to assess the model's validity and accuracy, as opposed to the patient's condition. The main conclusion from the modelling was that the location of the eye was not being described accurately enough. The results were sensitive to all the estimated and predefined dimensions, although some more so than others. The description of the head velocity was not as significantly affected as the eye velocity by varying dimensions, but the eye appeared highly sensitive to certain values.

4.3.3.1 Sensitivity to the Head's Horizontal Radius

The following results were obtained from a normal subject and reprocessed to display the effect on eye velocity of variations in the head's horizontal radius. The difference between the radius in each case is 2cm, yet the effects upon the eye velocity from inaccurately predicting the head's horizontal radius are dramatic. As the patient was normal, the results should show that the head and eye velocities are almost the same, indicating a VOR close to unity. The first graph is closest to the assumed result, but variations of only 2cm yields progressively worse results. Decreasing the radius from 12cm to 10cm almost yields a doubling in eye velocity.

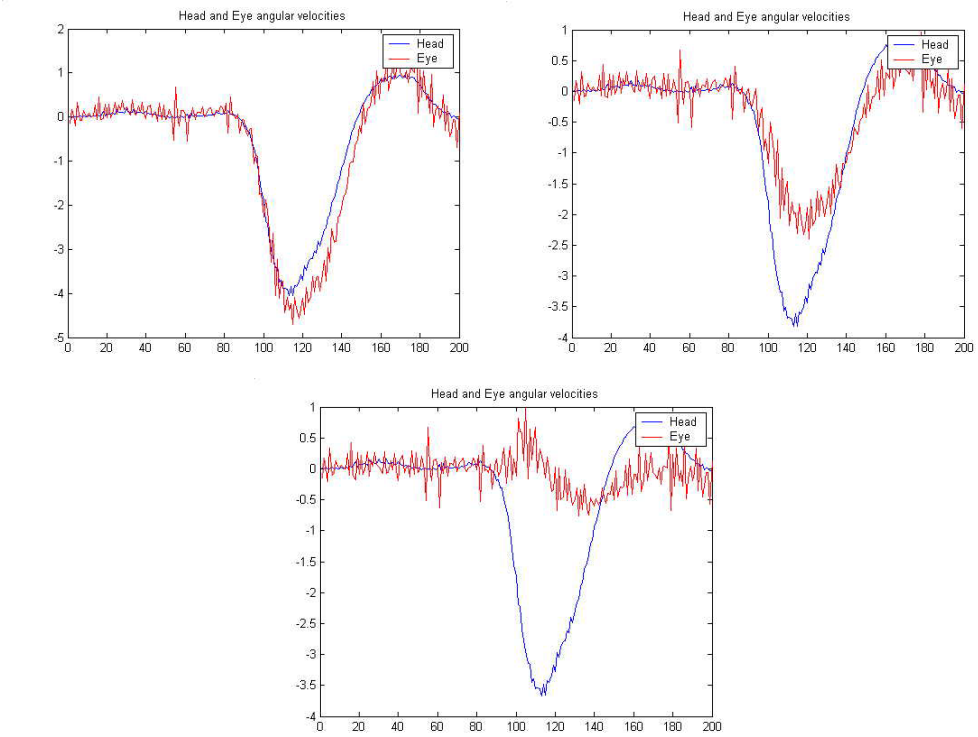

Figure 38. Sensitivity of the eye velocity to variations in head's horizontal radius. Radius: Top left, 10cm. Top right, 12cm. Bottom, 14cm.

4.3.3.2 Sensitivity to the Distance between the Eye and the Head Marker

The following results show the sensitivity of the eye velocity to the distance between the front of the eyeball and the forehead. A difference of only 2cm gives drastic differences in eye velocity. The velocity almost doubles when the distance is reduced from 5cm to 3cm.

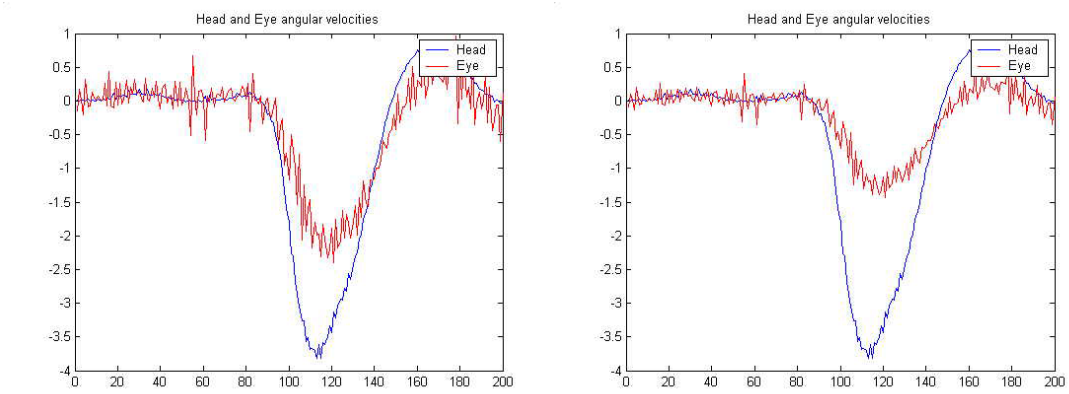

Figure 39. Sensitivity of the eye velocity to variations in distance between the eye and the forehead. Left, 3cm. Right, 5cm.

The head's vertical radius yields the same level of sensitivity, but the results are less sensitive to variations in the eye's radius. The eye, while assumed spherical, varies very little from one subject to another and the radius was defined to be a constant 2.5cm. As all of the parameters must be described accurately to obtain the correct results, this model will currently not yield good quality signals unless the calibration tests can be modified to define the values more accurately.

5 Signal Analysis

5.1 Comparison of Results

Each of the tracking methods with their respective models was critically evaluated based on several sample data sets. The data sets comprise of tests only in the lateral plane, as they are easier to conduct and the effects of disorders are easier to see visually.

The results are assessed in terms of the quality of the image processing and of the Vestibular information extracted from the head and eye movements.

5.1.1 Results from the 2D model with the Pattern Matching

The current implementation of the Impulse Test system employs the pattern matching techniques with an inbuilt two-dimensional head model. When the experiment was performed on a normal subject, the results image processing results appeared positive under the correct conditions. Both positional and velocity data could be extracted with high Signal to Noise Ratios (SNRs). The positional data below exhibits very little noise in any of the signals and even when differentiated to yield velocity, the resulting data still has a high SNR.

From a normal patient, one would expect the eye to be able to compensate for the head movements almost instantly. This can be seen in the following graphs as the eye velocity increases at the same time as the head velocity.

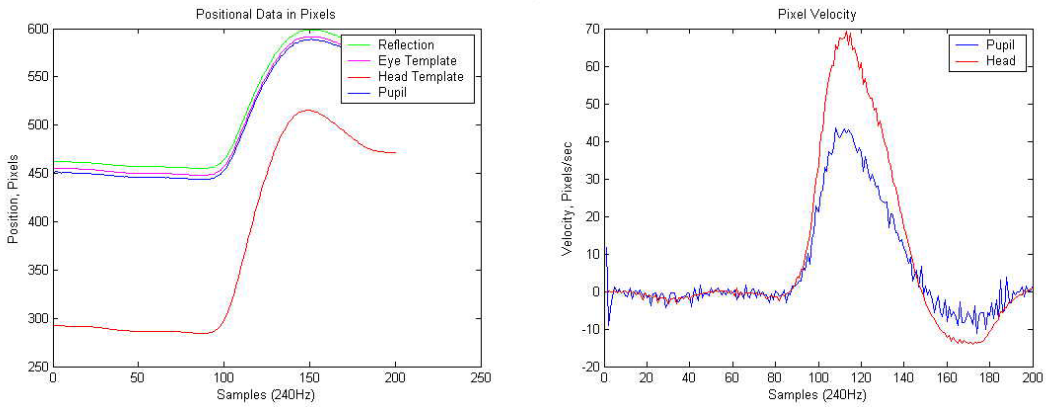

Figure 40. Results from a normal patient. Left, raw positional data. Right, pixel velocity data with a simple 2D model.

But depending on the set-up conditions, the results could quickly deteriorate. The following data resulted from choosing the head and eye objects incorrectly. Although there appears to be much noise present, it was possible to perform System Identification to obtain a model for the system using both sets of data.

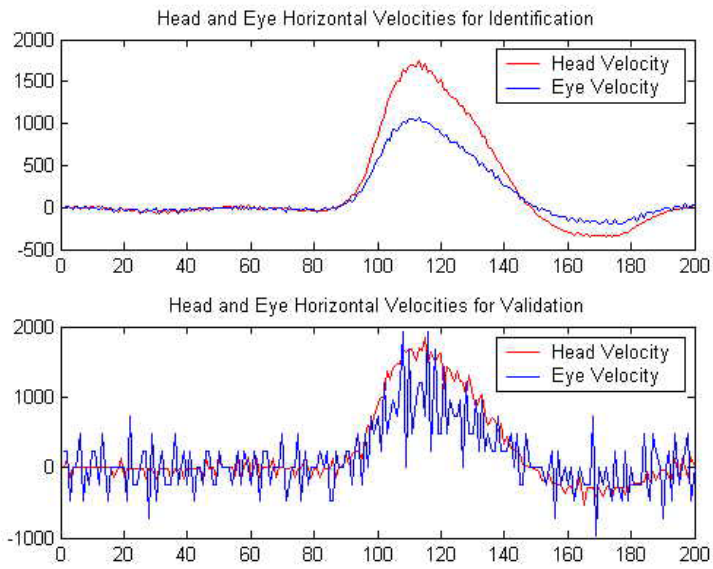

Figure 41. The effects of different experimental set up conditions with the Pattern Matching technique

When the analysis was performed with data from an affected patient, the results were much more difficult to obtain and interpret. The first problem was with the selection of the eye objects. When the reflection was within the selected region, the template could not be found when the eye had rotated from its original position. Therefore, care was exercised to exclude the reflection and the following results were obtained. One can see from the positional data that the magenta line, representing the eye template, has several interesting properties.

- The large spikes represent a failure in the pattern matching techniques to track the eye at those samples. These points correspond to the high head velocity rotations, so the poor results could be attributed to the rapidly changing environment. The result is also prominent in the velocity graph (large blue spikes).

- Towards the end of the test (samples 230 – 250 approximately) the eye position reduced, corresponding to the eye saccade. This was the result of the patient realising they were looking in the wrong direction. It is difficult to see in the velocity graph because of the red, head velocity line.

- Although the patient's eye did initially rotate with the head (hence not rotating with respect to its own origin), the eye velocity shows that it did. This is not as a result of the tracking algorithm, but it is as a result of the two-dimensional model. Therefore, this model is not precise enough to accurately describe the whole eye motion.

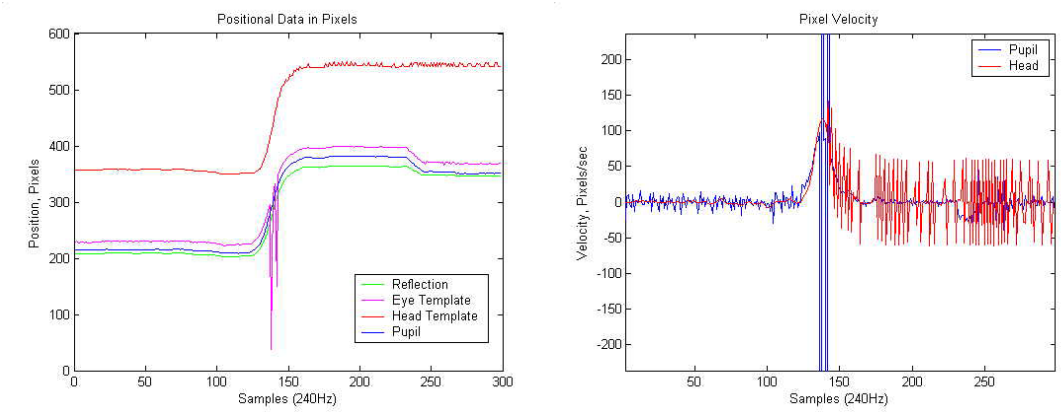

Figure 42. Results from an affected patient. Left, raw positional data. Right, pixel velocity data with a simple 2D model.

The head template results are also quite poor, as can be seen in the velocity graph. When the test was performed, the head was also rotated torsionally, which was not accounted for in the system. Therefore, this is a feature, which requires development.

Obtaining results of this quality for the affected case was extremely difficult as many of the image processing properties had to be tuned manually and selecting the object templates was extremely problematic. Therefore, this pattern matching approach to the head and eye tracking issue is not sufficiently reliable on its own.

5.1.2 Results from new technique

With the new image processing method developed, data was processed to show the effects of the algorithm and of the Vestibular disorders. The two-dimensional data was projected upon simple three-dimensional models to obtain results. The eye is represented as a sphere, which rotates at the same rate as the head. Then using the position of the pupil and the reflection from the image detection, the eye's velocity was estimated.

The results from the normal subject were almost as expected. The appearance of the positional and pixel velocity results below support the assumption that the subject was normal with respect to the VOR.

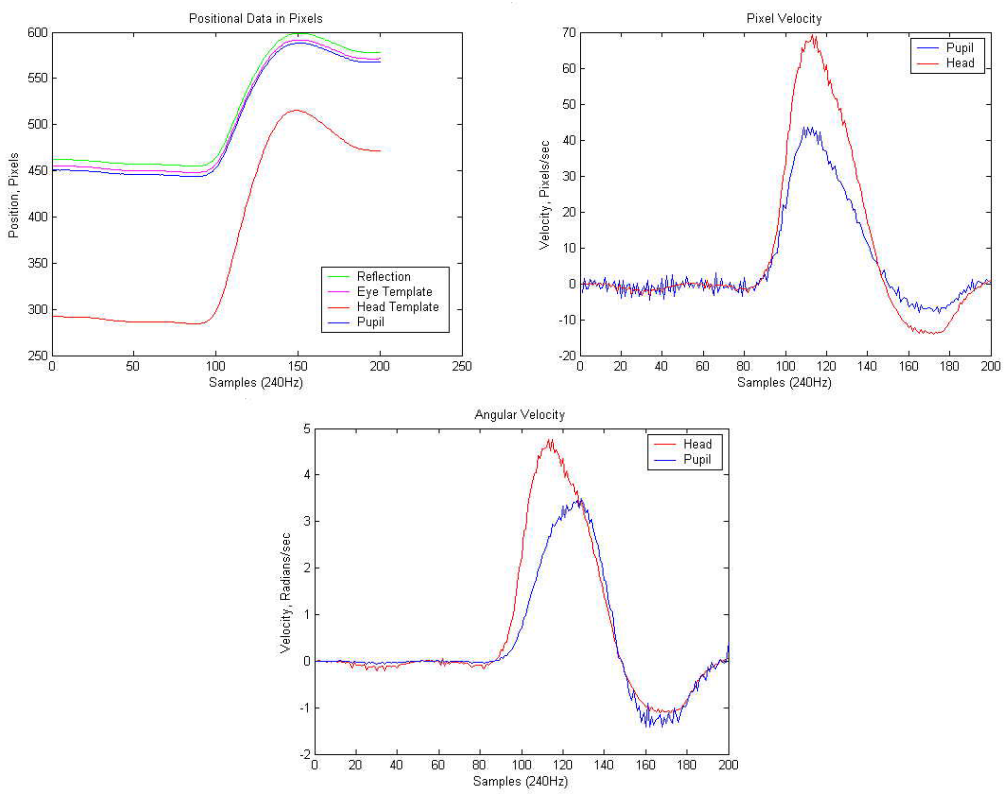

Figure 43. Data from a normal subject. Top left, raw positional data. Top right, pixel velocity. Bottom, angular velocity

The results from the tests performed upon an affected subject were more pleasing. From the positional data, one can see several interesting features.

- The eye template was tracked with the pattern matching methods and its line shows two large positional jumps and a large spike. The first jump corresponds to the system losing track of the left eye and finding the right eye in the next frame because of the changing position of the reflection during the rotation. The final step back to the left eye corresponds to the time when the saccade occurred. Therefore, the reflection moved back to its original position upon the pupil, making it appear more accurate than the right eye. The spike appears to be as a result of the tracking failing altogether for a single frame.

- One can see that the pupil and the reflection position lines cross each other twice in the figure below. As the eye rotated with the head, the reflection moved across the pupil and due to the small size of the pupil, the pupil appeared to jump. This effect can clearly be seen in the velocity graphs.

- From the pixel velocity graph, it is clear that the eye moves with the head, as the velocities appear very similar. Then, as the head stops moving, the eye saccade can be seen, although it appears distorted due to the spike corresponding to the pupil "jumping" across the reflection.

- From the angular velocity graph, the eye appears to move relative to its own origin as the head rotates. Upon closer inspection of the images, the eye did move a little, although not enough to explain the result in the graph.

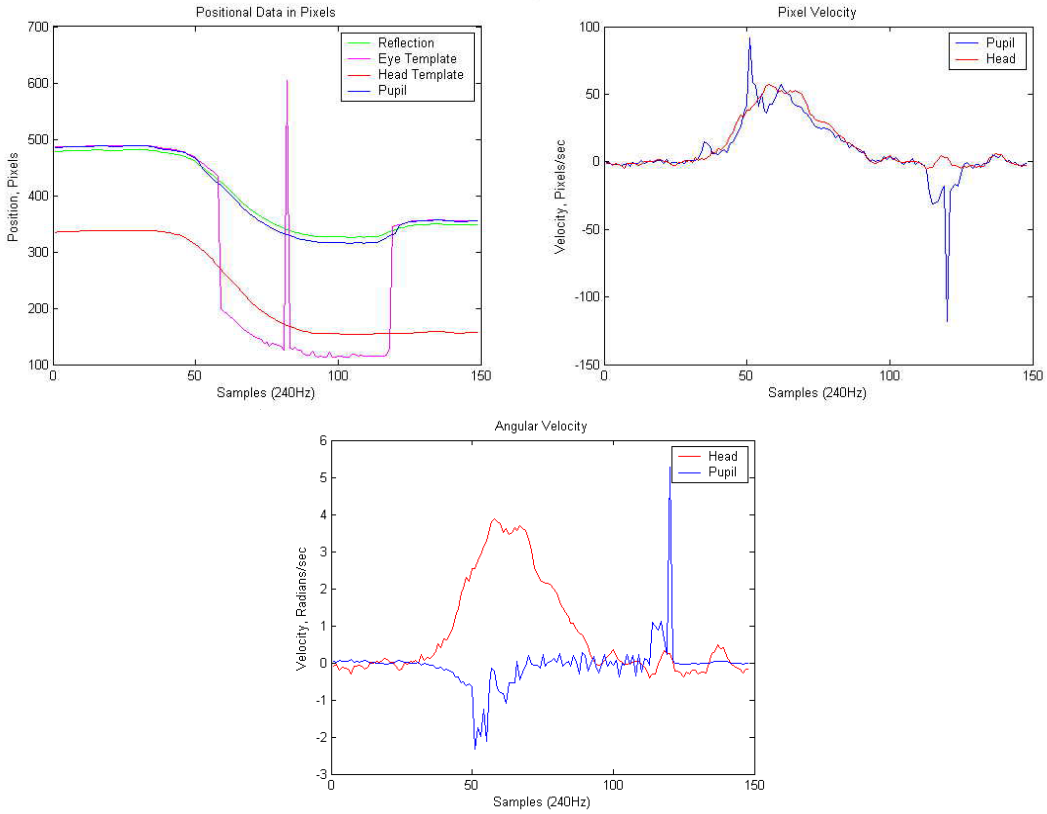

Figure 44. Data from an affected subject. Top left, raw positional data. Top right, pixel velocity. Bottom, angular velocity

The problems associated with the reflection moving across the pupil did not pose a problem when a simulation was performed. The test was organised in such a maner that the reflection left the pupil once the head rotated.

One can see that there was a distinct saccade from all three graphs, beginning with the 15th sample. But in the angular velocity graph, the eye is described as moving in one direction and then in the other while the head rotates. This suggests that the eye model is required to be accurate in its description and system calibration.

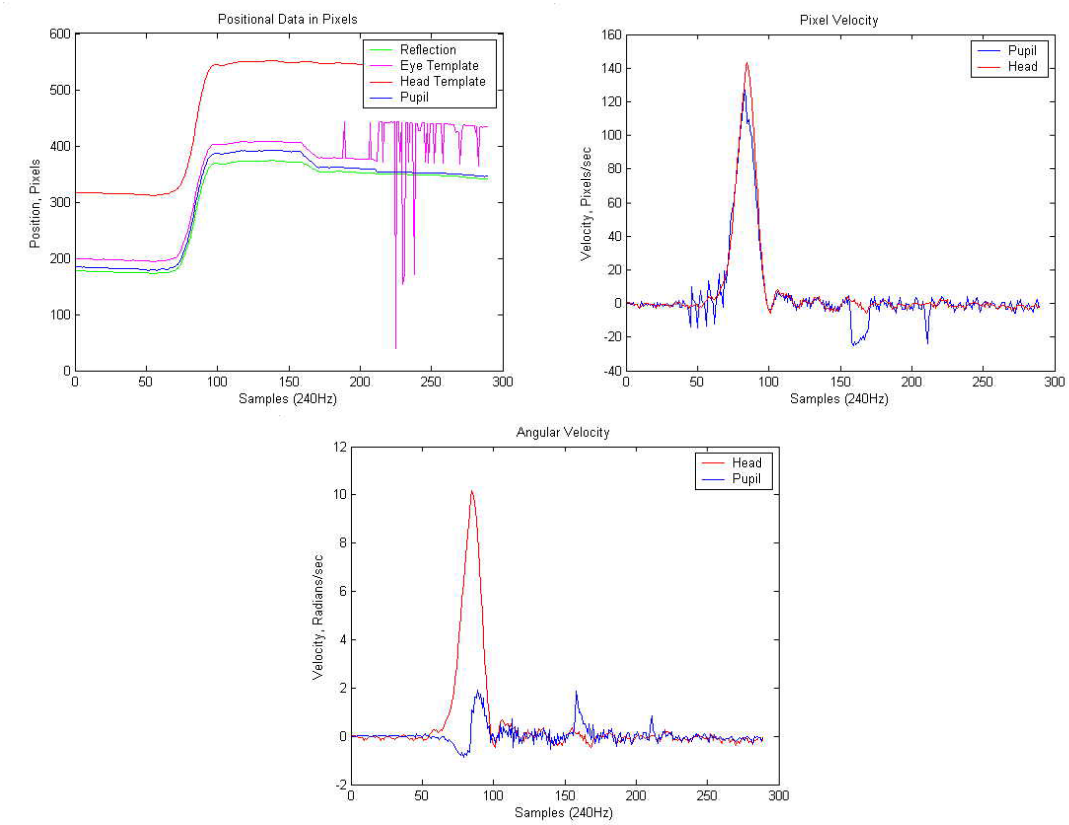

Figure 45. A saccade with no unwanted spike. Top left, raw positional data. Top right, pixel velocity. Bottom, angular velocity

The results from this method appear to be more positive than those from the pattern matching techniques above. The problems associated with the selection of each object disappear, although a more accurate model of the eye is required to accurately describe its motion.

5.1.3 Results with 3D head model and pattern matching

Due to the problems the pattern matching method experiences with respect to the reflection, the head results presented in this section will come from the pattern matching and the eye details will come from the new method.

When evaluating a normal patient, the results obtained were similar to those obtained previously. The angular velocity again was lower than expected as too was the linear velocity. Refer to section 4.3 for head modelling details.

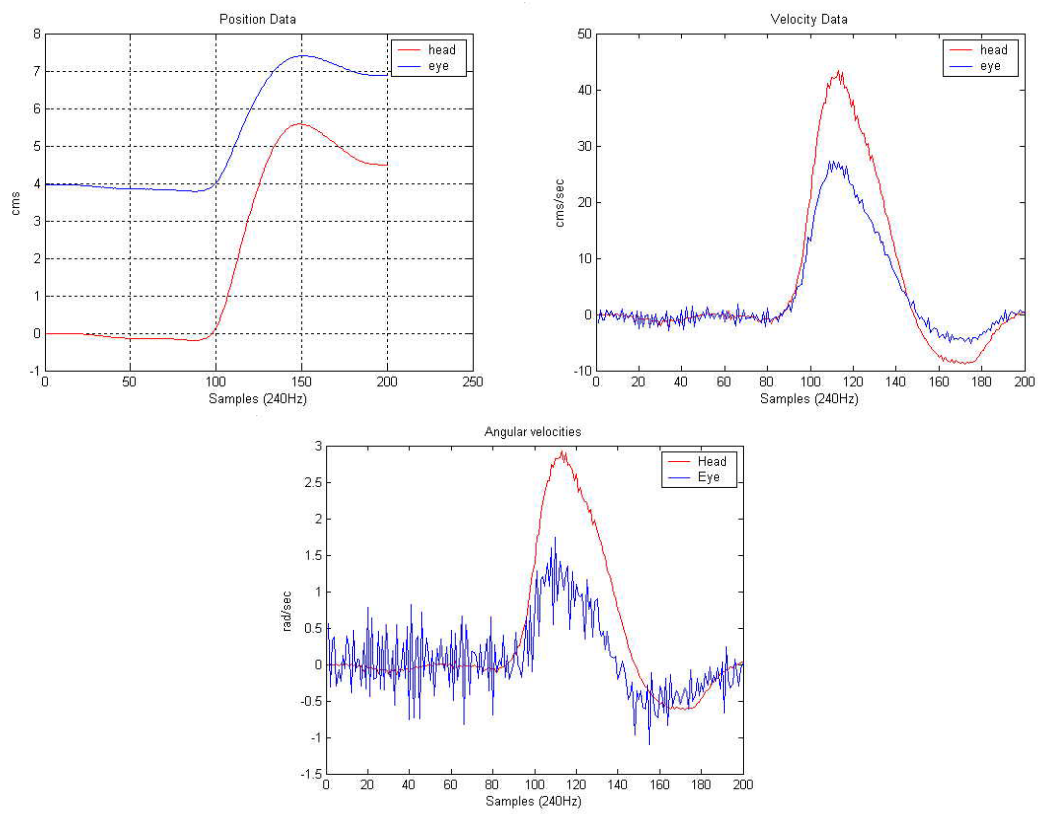

Figure 46. Results from a normal patient. Top left, position data. Top right, linear velocity. Bottom, angular velocity

When the data obtained from an affected patient are processed, the results appear similar appearance, but scaled differently from earlier results. Compare Figure 47. and Figure 48 with Figure 44 and Figure 45 respectively. An explanation for the difference is due to the different methods in obtaining angular data, but a more precise investigation must be undertaken to discover the root cause of the discrepancies.

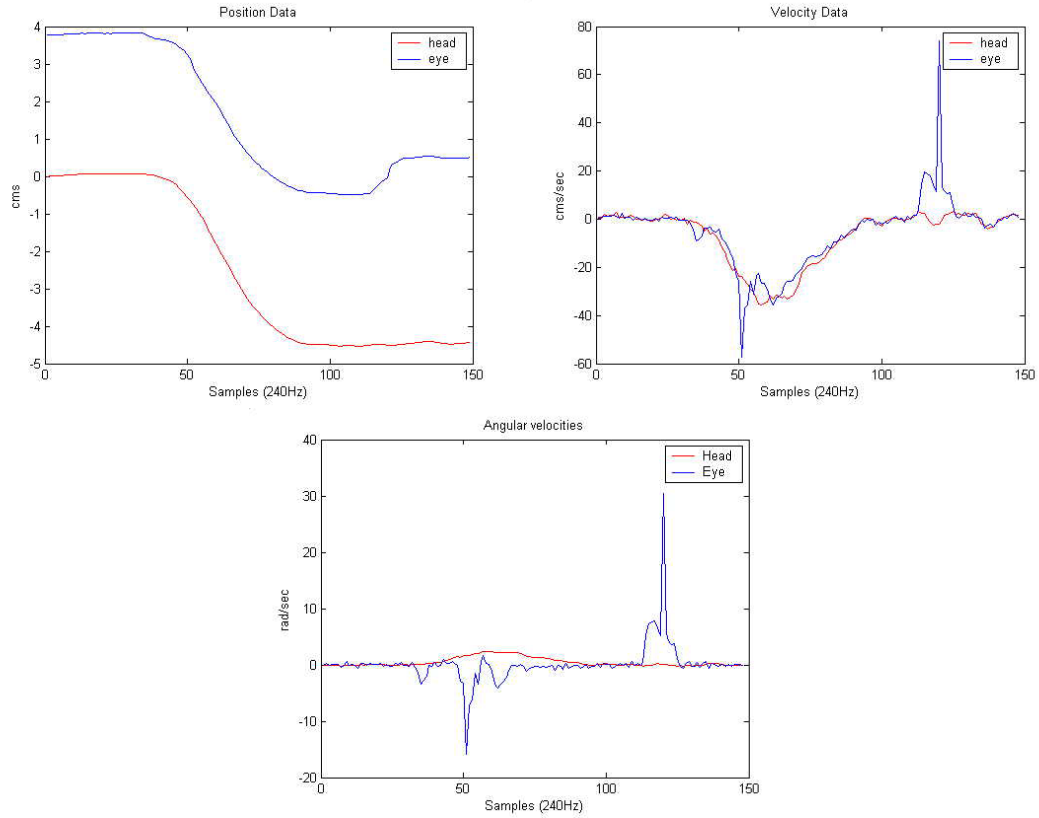

Figure 47. Data from an affected subject. Top left, raw positional data. Top right, pixel velocity. Bottom, angular velocity

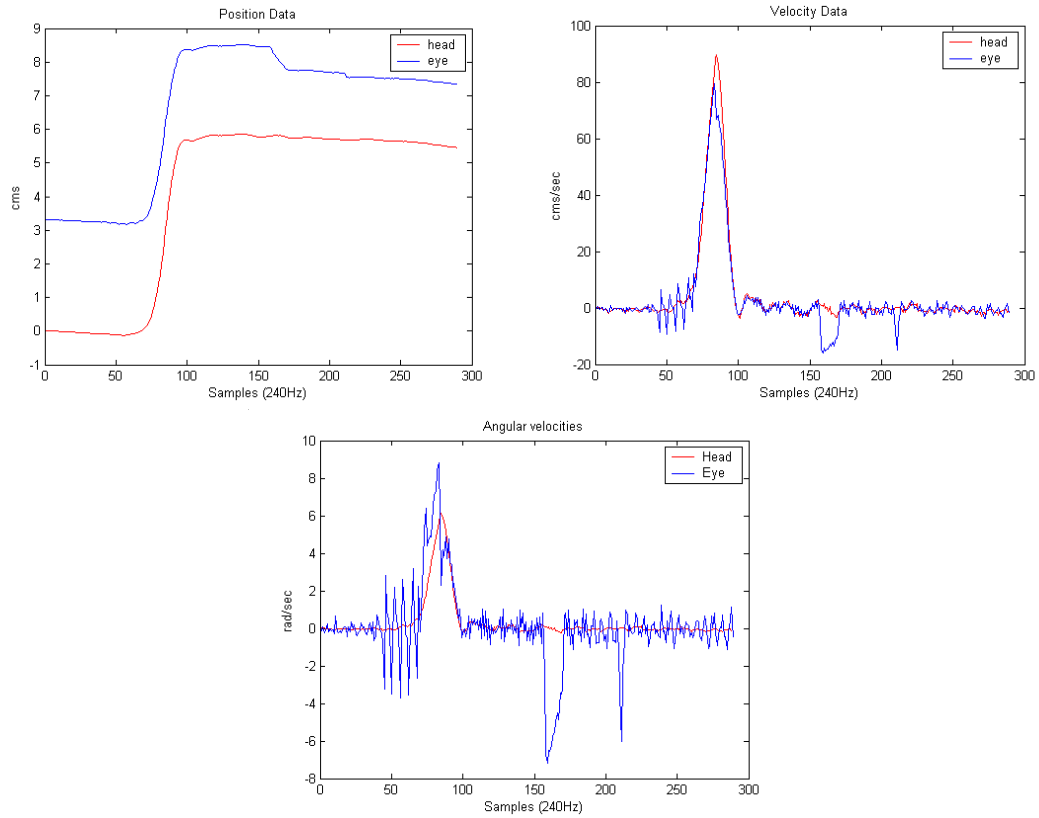

Figure 48. A saccade with no unwanted spike. Top left, raw positional data. Top right, pixel velocity. Bottom, angular velocity

5.2 System Identification

This section of the report details the system identification experiments performed. It is assumed that the reader is familiar with the methods and techniques employed, and therefore the background material will not be presented here. If further information is required, the reader is referred to [11] and [12].

5.2.1 Experimental Aims

The system under investigation is the human head, with head velocity as the input and eye velocity as the output. Although the real relationship is highly non-linear, under certain conditions it may be possible to approximate the input / output relationship as a linear system. The aim of these experiments was to determine whether or not this non linear system could, under the correct circumstances, be identified and described with linear methods. The saccade is one of the major non-linearities which experiments on normal subjects should render unnecessary to model. Attempts were made to fit different types of models to the data to try and determine characteristics about the system.

The 2D model was employed to model the head and eye relationship as the aim of these experiments was not to accurately identify the over simplified head to eye relationship but rather to test the hypothesis that linear identification methods could be applied to this data. The noise characteristics were also of great interest, as the image processing employed in the project was affected significantly by the surrounding environment. Lighting, camera focus and the contrast levels of the tracked objects all affect the quality of the image processing. Therefore, the noise must also be investigated.

5.2.2 Data

The data recorded was processed twice, once with the correct method to yield lower noise levels, and secondly with a set up designed to deliberately include more noise. The data is shown below.

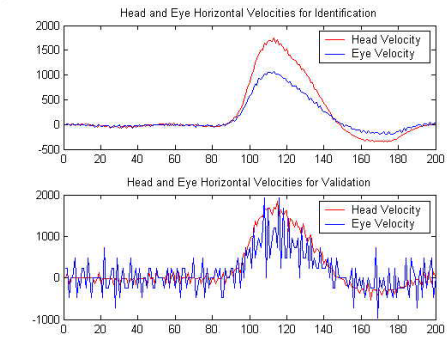

Figure 49. Top, identification data. Bottom, validation data

5.2.3 Initial Spectral Tests

The data was read in from text files and converted to the appropriate format for velocity data. Using the Matlab command Spectrum, the power spectral densities of the input, output and the cross spectrum were obtained. The input and output spectrums both showed low pass characteristics and appear to have very similar characteristics, as expected. With this simple model, when the head moves, so too will the eye although with lower velocity because the eye tries to remain fixated straight ahead, while the head continues to rotate. When the coherence spectrum was obtained using Matlab's Spectrum, it yielded numerically strange results;

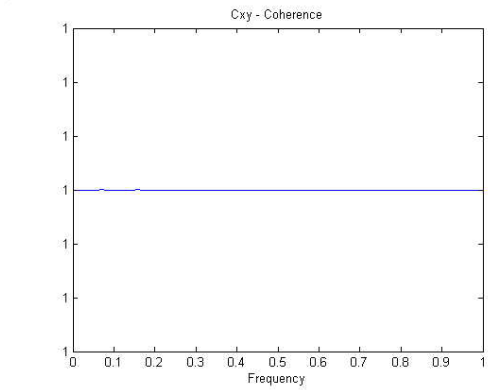

Figure 50. The Coherence plot resulting from the Spectrum function.

Numerically, it appears correct from the input, output and cross spectrums, although I cannot explain it. The same result comes from both the identification and validation data. The expectation was that the coherence spectrum would be very close to one for most of the lower end of the spectrum, but would decrease at higher frequencies, mainly due to the differentiation of the positional data and other noise sources. Therefore no results can be drawn from this.

5.2.4 Non Parametric Modelling

Correlation analysis was performed on the velocity data to estimate its impulse response and step response. The Matlab command CRA was used to estimate the impulse response of the data and then CUMSUM to derive the step response from the impulse response. Both of these analysis tools gave disappointing results. They appear to show that the data yields a marginally stable, or an unstable, system although we know from everyday life that this is not the case.

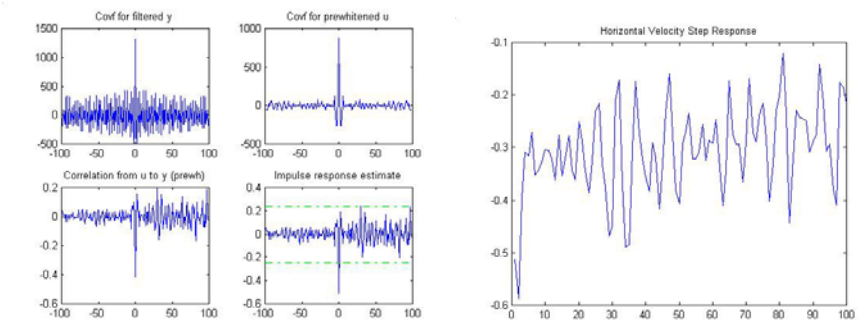

Figure 51. Correlation Analysis and the estimated Step Response

At this point, the method of de-trending the data was investigated. Performing a linear de-trend of the data yielded equally unstable results and processing the data with no de-trending appeared to shift the mean level of the step response, but still the result appeared unstable.

The next stage in the Non Parametric Modelling was to examine the data displayed in a Bode Plot. The theoretical transfer function should be approximately flat for a perfect system, representing the ratio between the eye and head velocities. The phase should be close to zero as the eye moves in the same direction as the head, measured from the head's rotation point. The results were almost as expected. At very low frequencies, the magnitude spectrum was flat, but at higher frequencies, the spectrum did not show any significant low or high pass characteristics, for example. This cannot be immediately interpreted as failure of the linear approximations to data, as the velocity data was obtained from differentiating position data and differentiation will tend to amplify high frequency noise. With a little optimism, one could argue that all is not yet lost.

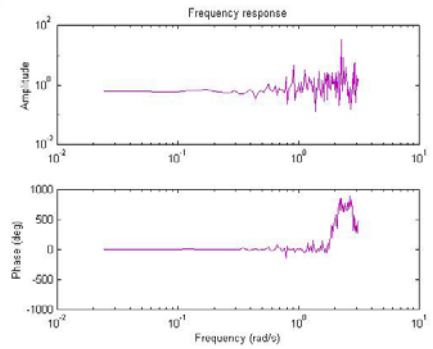

Figure 52. The Bode plot from the input / output data.

5.2.5 Parametric Modelling

Two attempts were made to fit different model structures and orders to the data; ARMAX and Box Jenkins models. The results varied from one method to another and this will be explained here.

5.2.5.1 Fitting ARMAX models

The first attempt to fit a parametric model to the data was with an ARMAX model. The general equation for an ARMAX model is as follows [11];

$$A(q)y(t) = B(q)u(t-nk) + C(q)e(t)$$

The initial order of the polynomials was the first choice and the orders of all three polynomials were set to 10 in Matlab's Ident facility. The time delay was set to zero. By observing the pole zero plot, there appeared to be much room for model order reduction, as several poles and zeros were very close together.

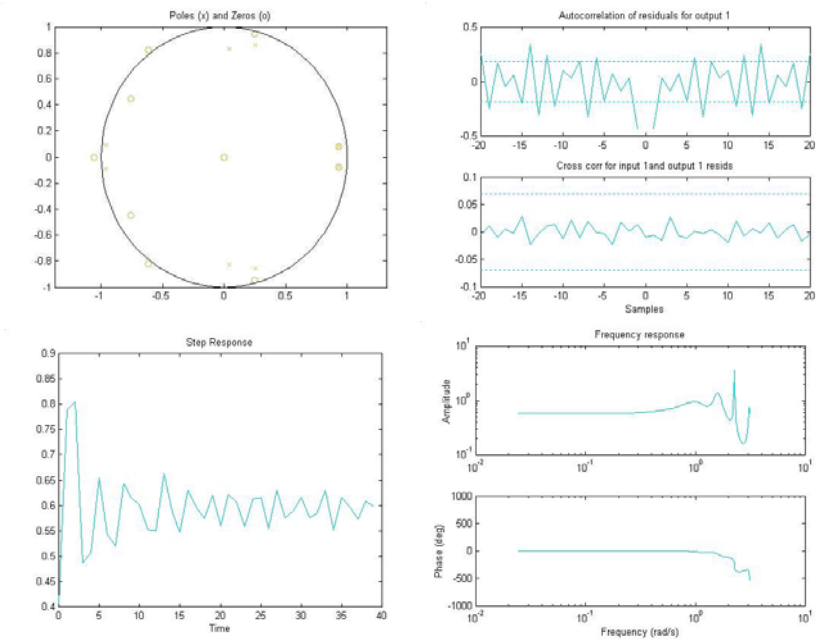

Figure 53. ARMAX modelling results for polynomials of order 10. Top left shows the pole zero configuration. Top right shows the residual analysis based upon the validation data. Bottom left shows the estimated step response. The bottom right is the estimated transfer function

The orders of the polynomials were continually reduced until the poles and zeros close together disappeared. The orders were reduced to one and the results appeared extremely positive. The step response rose to 0.6, which is approximately the ratio of the eye to head velocities, with no overshoot as expected, and the magnitude of the transfer function was flat over significant portions of the spectrum and dipped only slightly at higher frequencies.

The residual analysis performed upon the validation data yielded results which fell within the 99% confidence limits for both the tenth order and the first order models.

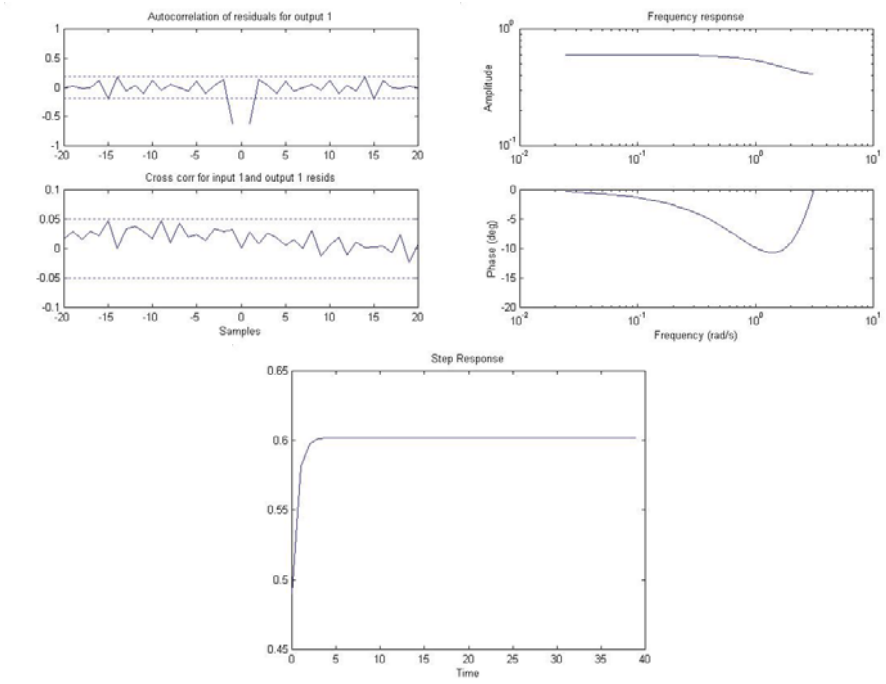

Figure 54. ARMAX modelling results for polynomials of order 1. Top left shows the residual analysis performed upon the validation data. Top right shows the estimated transfer function. Bottom shows the estimated step response

The results from the initial tenth order and the final first order ARMAX model were exported to the Matlab workspace and their information was displayed.

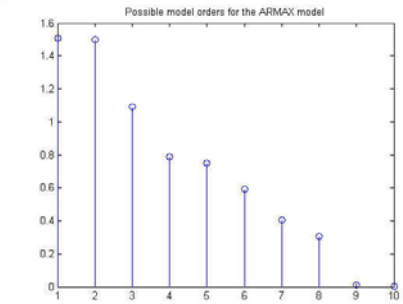

Figure 55. Possible model orders resulting from a balanced realisation on the tenth order model

The tenth order model was subject to a Balanced Realisation and the resulting diagonal elements of the gramian were displayed [11]. This was done to compare against the results of the model order reduction by manually removing pole zero pairs. It did not appear to suggest that the model order could be reduced as far as one, despite the results from earlier experiments suggesting that it could. One could argue that order eight would be a reasonable order for system reduction without losing valuable information. But an eighth order model also contained pole zero pairs which could be cancelled.

Therefore for several reasons, the lower order models appeared to be at least as good as the higher order models. Firstly, only when the model order was reduced as far as one were all the pole zero cancellations removed. Secondly, the lower order models produced superior results from the residual tests and finally, the transfer functions

from the lower order models appeared to agree with the prior knowledge of the system.

To compare the results from both the tenth and the first order models, the Matlab function dlsim was used to display the results.

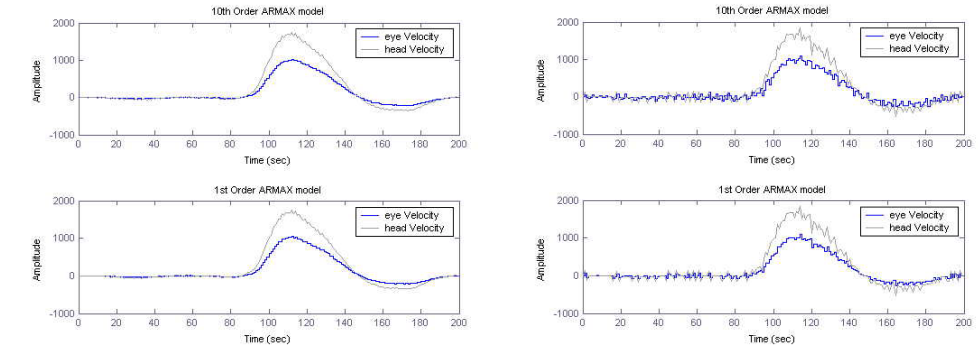

Figure 56. Comparison of tenth and first order models. Left, the identification data. Right, the validation data.

The results from the first order model with the identification and the validation data are as good as those from the tenth order, further justifying this low order of model. Compare with Figure 49.

5.2.5.2 Fitting Box Jenkins models

The data was also fitted to the Box Jenkins model as it contains a separate model for the noise from the system. The general formula for a box Jenkins model is as follows [11];

$$y(t) = \frac{B(q)}{F(q)}u(t-nk) + \frac{C(q)}{D(q)}e(t)$$

It allows both the system and the noise to be modelled separately, and the hope was that this would agree with, and further improve upon, the results from the ARMAX modelling. The initial polynomial orders were set to make the Box Jenkins model the same as the ARMAX model, ie; Nb=1, Nc=1, Nd=1, Nf=1 and the results verified this.

The noise model was then reduced in order and the best results were given by reducing the C(q) and D(q) polynomials to order zero. This suggests that only the current noise sample affects the output and hence no structured model is required to describe the noise characteristics. This was surprising, as part of the image processing works recursively. When the marker defining the head is found in a frame, the Impulse Test system uses its position to define a new area to search for the eye markers in the next frame. Although the expected interference was small, it was expected to affect the result more so than appeared in these results.

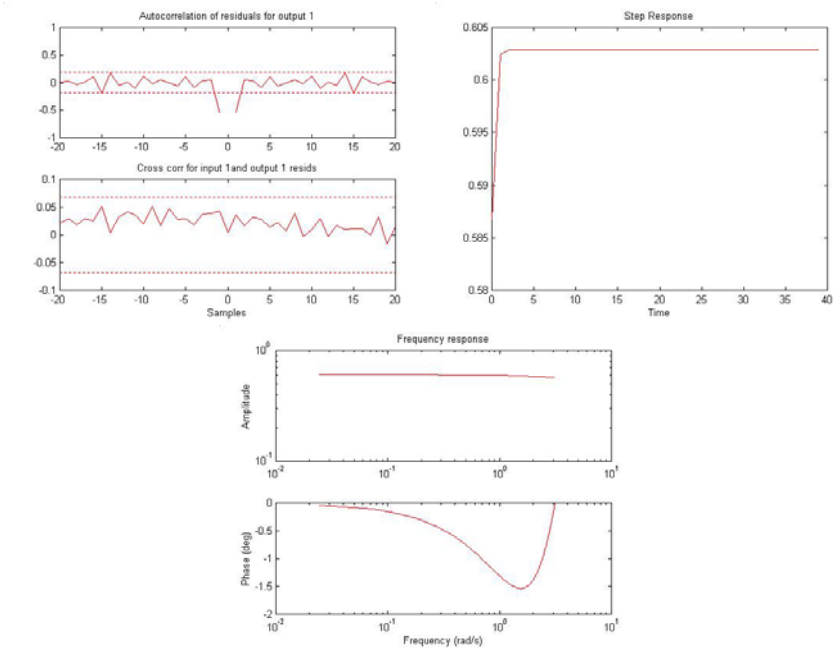

Figure 57. The results from Box Jenkins modelling. Nb=1, Nc=0, Nd=0, Nf=1. Top left shows the residual analysis performed upon the validation data. Top right shows the estimated step response. Bottom shows the estimated transfer function.

The results here are similar in appearance to those from the ARMAX modelling, and hence appear to further support the modelling methods.

5.2.5.3 Comparison between ARMAX and Box Jenkins

The big improvement in the results from the Box Jenkins modelling over the results from the ARMAX modelling is that of the transfer function. Its magnitude appeared much flatter, around 0.6, over a larger range of frequencies. Its phase also appeared much flatter, deviating significantly less from 0 degrees than in the ARMAX case.

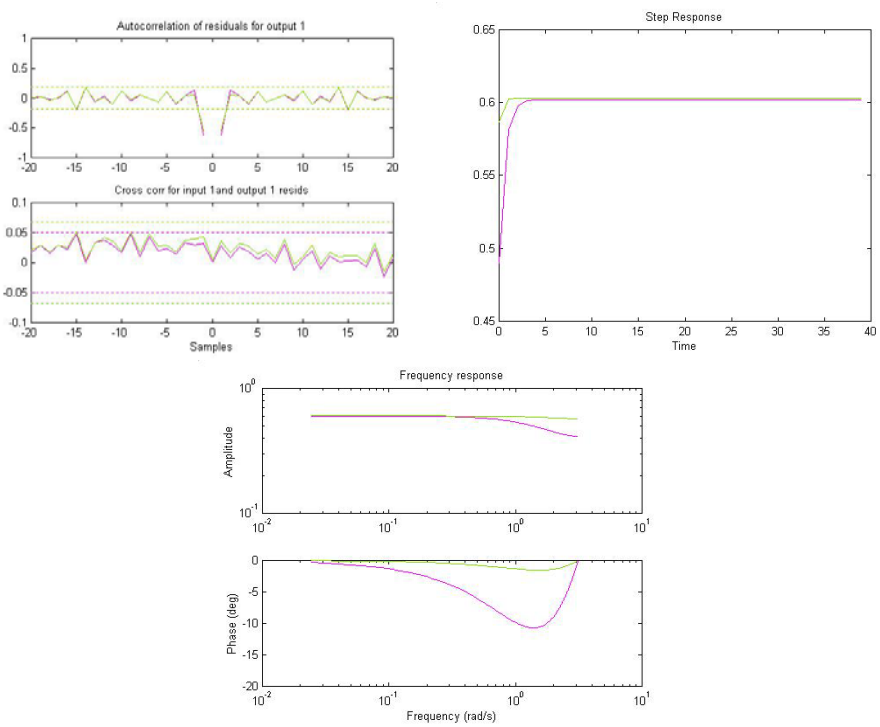

Figure 58. A comparison of the Box Jenkins (green) and the ARMAX methods (purple). Top left the residual analysis based upon validation data. Top right, the step response. Bottom the transfer functions.

The step response also rose faster with the Box Jenkins method and settled to approximately the same value as before. Again, there was no overshoot and the final value settled to approximately 0.6.

The residual analysis for the Box Jenkins modelling and the ARMAX modelling both yielded good results. The autocorrelation of the residuals were virtually identical, while the cross-correlation of the residuals and the input were also very similar, but with a slight offset. All the residuals appeared well within the 99% confidence levels.

5.2.6 Sub Space Based Identification

The final method used to identify the system form the data was using Subspace based methods. The Matlab command n4sid was used and tests were performed in an attempt to validate, or otherwise, the results obtained from the earlier experiments.

A range of model orders from one to six were displayed by computing the singular values of the Hankel matrices of the impulse response. The resulting models were investigated in terms of their transfer functions, step responses and stability characteristics.

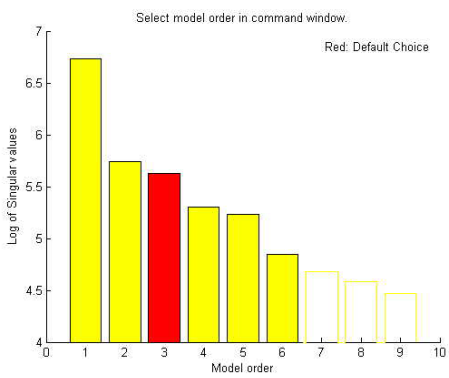

Figure 59. Singular values of the Hankel Matrices of the Impulse Response

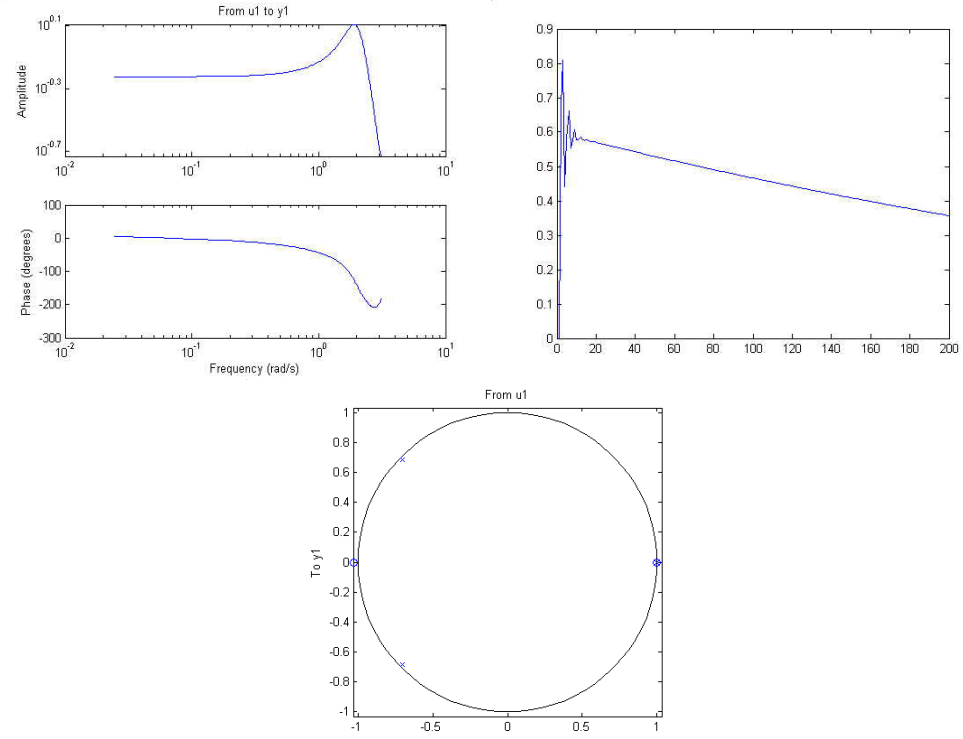

Figure 60. The results of a third order Subspace model. Top left is the estimated transfer function. Top right is the estimated step response. Bottom is the pole zero plot

The subspace algorithm suggested that a third order model was best, and it gave moderate results. The transfer function appeared to begin with approximately the expected values for both magnitude and phase, but as the frequencies increased, so too did the phase and the magnitude.

The step response had a very short rise time with significant oscillations, and it did not appear to settle to a constant value, as seen previously in the parametric modelling section. Although, the general appearance of the response was close to that expected.

From the pole zero plot, there appeared to be a cancellation at low frequencies and two poles very close to the unit circle. Therefore, a model of one order less was investigated and the results were found to be rather similar.

With a second order model selected, the step response improved slightly. The oscillations remained, and the response again did not appear to settle to a static value although, the decrease was less than with the third order model. The appearance of the oscillations indicate that the model is not as accurate as seen with the parametric modelling.

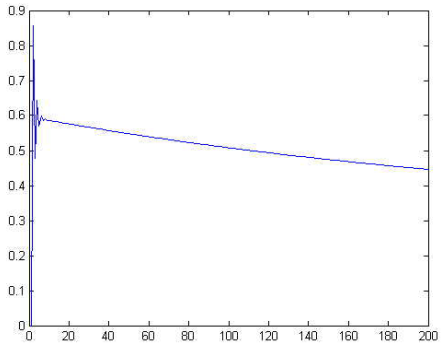

Figure 61. The step response from a second order subspace model

Reducing the order further yielded a very poor step response and therefore a second order model appears to give the best model approximation with subspace methods.

To test the subspace results, the same test was performed as with the reduced ARMAX model. The identification and validation data were both processed through the third and second order subspace models.

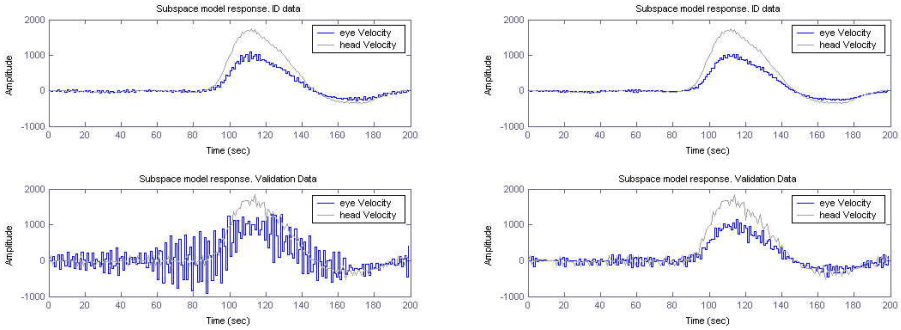

Figure 62. Top, the results of processing the identification data. Bottom, the validation data on the third order subspace model (left) and the second order subspace model (right)

The third order model appears to perform less well with the validation data when compared to the second order model. This goes some way to confirm that the system can be modelled with low order linear modelling techniques.

5.2.7 Results

The initial hypothesis of being able to approximate a non-linear system to a linear system under certain conditions was tested and gave varying results. The raw data was positional data and it was converted to velocity data by differentiation. This process

naturally amplifies noise, possibly presenting a more challenging identification process.

The initial spectral tests did not yield any results which could be used to support, or otherwise, the hypothesis as the coherence spectrum gave numerically strange results.

The first stage of the non-parametric modelling was to perform a correlation analysis. This suggested, from the step response, that the system was marginally stable or even unstable and hence very little information could be extracted from this to form a basis for further experiments. The Bode diagram did though give a transfer function which did, at least partially, agree with prior knowledge of the system. It exhibited a flat lower portion of the spectrum, but due to higher frequency noise, less system information could be extracted from the upper regions of the spectrum.

The ARMAX modelling showed quite convincingly that the initial polynomial orders of ten could be reduced to one, while still retaining all the properties that the system was assumed to possess. But as the noise is a major obstacle to the overall system performing accurately, a Box Jenkins model was fitted to the data. The initial values were chosen to match the results with those from the ARMAX modelling for comparison purposes. Only the noise polynomials, C(q) and D(q), were adjusted in an attempt to accurately describe the noise characteristics and the best polynomials were found to have order zero. This implies that only the current noise sample affects the system output. With the low (first) order system polynomials and only the current noise sample affecting the output, the results were good. The transfer function appeared flat over most of the frequency range and also the phase remained low. The step response was rapid, with a final value corresponding approximately to the ratio of the eye to head velocities. So a model with A(q) and B(q) polynomials of order one and simply the current noise sample produced best results.

The Subspace based methods also appeared to confirm that the data was more accurately modelled with a low order system. It yielded best results for a system of second order and when tested upon the validation data, the result was that the output contained much less noise than the original validation data. But when compared to the parametric modelling results, it can be concluded that parametric modelling describes the system more accurately.

6 Discussion and Further Work

The current version of the system is not in a state whereby patients can easily be assessed and hence further development is required in all areas. It does, however, provide a good foundation for the development of both the image processing and the modelling. After discussions with all parties involved, several approaches have been formulated for the next stage in the system's development. They will be elaborated upon here.

6.1 Image Processing Algorithm

The most accurate and reliable method so far is that which involves thresholding and tracking the reflection. This method has been shown to be robust to eye movements beyond those encountered within the Impulse Test, as well as becoming obscured by other obstacles such as the eyelashes. There exists the potential for accurately describing the eye, although some modifications are required to the basic algorithm in order to overcome the previously described limitations. The main issue to overcome is determining the centre of the pupil accurately without introducing errors due to the reflection. The reflection distorts the centre of the pupil after thresholding, and compensating accurately for the reflection has proved to be dependent on the direction of head rotation, and relative sizes of the pupil and the reflection. Three methods are discussed here to further develop the system.

1. Indirect Lighting Method

By altering the lighting conditions, as described in section 3.6.6, to an indirect lighting method, the pupil would be much more defined in terms of its shape and area, because of the absence of the reflection and all the associated problems. But the same fundamental problem of locating the origin of the eye arises again. This method would require a form of head model, which has already been shown to be sensitive to physical variations. Therefore, this appears to offer little, if any, advantage over the current implementation.

2. Characterising the Pupil

The changing appearance of the pupil as the reflection moves across the eye must be addressed to continue with this implementation. Although methods have been employed in an attempt to overcome this problem, the effects can still be prominent under certain conditions. The current correction algorithm attempts to quantify the pupil's area by compensating for the area of the reflection, while it is within the pupil's boundary. As the reflection moves across the centre of the pupil, the pupil can appear segmented, proving difficult to measure and yielding inaccurate eye velocity results. Therefore, by defining the pupil as a circle in the first image frame and extracting parameters from this, it could be described as an oval in subsequent frames with the same methods of compensation as before. The results from this method are subject to the lighting conditions as the size of the reflection becomes large with respect to the pupil, as has already been encountered. Therefore, the benefits gained from this method are not sufficient to fully quantify the pupil without compensating for the reflection.

3. Strobe Lighting

The main issue with describing the pupil accurately in the presence of the reflection is its changing appearance during eye motion. The ideal solution would be to be able to fully quantify the pupil without the reflection and still retain the benefits of locating the eye's origin through the reflection. This could be provided with a strobe light synchronised with the camera.

By altering the method of lighting to an indirect source, the head could be fully illuminated to the required level while a separate direct source could be used to introduce the reflection. The implementation is theoretically simple and will yield the benefits of locating the eye without the disadvantages of the reflection while quantifying the pupil. Consider the following;

- Assume the camera produces images at 240fps as normal. The images are then split into two classes; with and without the reflection.

- On every even numbered image frame, the direct lighting source can be switched on, lasting for approximately the length of the frame. This will allow the centre of the eye to be located using the method currently employed.

- On the odd numbered image frames, the direct lighting source can be switched off, allowing for the pupil to be fully quantified in the absence of the reflection.

- The camera and the strobe light must be fully synchronised in order for only the even numbered images to contain the reflection, while the odd numbered images contain no traces of it.

- The strobing can be preformed from one of the available trigger signals onboard the camera. The exact timing requirements have not been defined, but the idea in principle is valid.

- The choice of strobe light is important. As the head shall be illuminated with an indirect source, the strobe light does not need to be high power. In fact, it should be low power because at 120fps, as the light will appear as a constant source to the patient. A high power source could result in the overall intensity being too high, forcing the patient's pupil to shrink further and they may also squint. Both of these effects could be detrimental to the overall performance. Therefore, an LED strobe light could be used.

- The synchronisation of the camera and the lamp should not be performed through the computer, as time delays in reading and writing the signals through the com ports may result in strobing the LED at the incorrect time. Also, if the timing signal was read through the computer, careful real time process handling would be required. In addition to read and write delays, other delays resulting from higher priority computer processes blocking the system could result in the LED signal becoming badly delayed. Therefore, driving the LED signal directly from the camera's output would remove the real time computing issues.

- Due to the specific requirements, building the timing circuit and choosing a standard LED should allow more design freedom compared to using an industrial standard strobe light. The light would be triggered from the camera using a Monostable Vibrator in conjunction with other relevant timing and switching circuitry.

- The LED should be low power, but sufficiently powerful to produce a significant reflection upon the cornea. The intensity of the LED could easily be varied by controlling the current flow through it and hence, the circuit could be tuned to give the desired reflection properties. The LED should perhaps be green, as it is a natural colour, which the eye can easily adjust to. A white light, as is currently being used, can result in effects similar to "Snow Blindness". The light currently employed leaves a residual image on the retina for several seconds, due to its intensity. But the effect from a green light is of lesser extent.

- The indirect light source must also be chosen with care. It should not result in any reflection upon the eye. Shining the light upon a dull, mat-finished surface, a photographer's umbrella for example, should satisfy the requirement. Positioning two umbrellas around the patient would illuminate both sides of the head without casting a shadow from the camera.

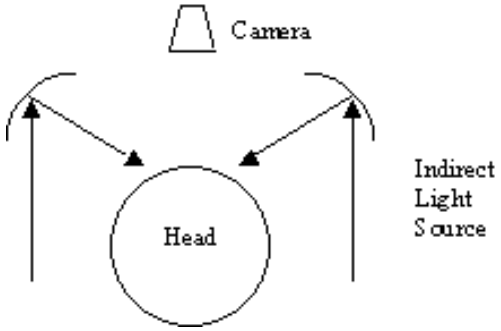

Figure 63. A possible indirect lighting arrangement.

6.2 Modelling

With the results of the head modelling in section 4.3.3, the decision to proceed with defining the eye with the reflection is further confirmed. The sensitivity of the eye velocity to physical dimensions in the model is too great to be able to assess the patients from the results. Therefore, the emphasis on modelling should shift towards detailing the eye.

In the simple model currently used to describe the eye, there are several limitations, which must be removed by increasing the accuracy of the eye model. Firstly, the eye is not perfectly spherical. It contains a perturbation at the front of the eye ball which should be modelled. This will enable a more accurate prediction of the eye's origin.

As described earlier in section 3.8.3, the reflection will be used to estimate the centre of the eye. By projecting a line from the camera to the reflection, the centre of the eye

can be found. But as the head rotates, compensation must be performed to allow for the fact that the centre of the eye is no longer directly behind the reflection, from a perpendicular view point to the patient. It will lie on the line extended back into the eye, which can be calculated from the head translation and the distance the patient is sitting from the camera.

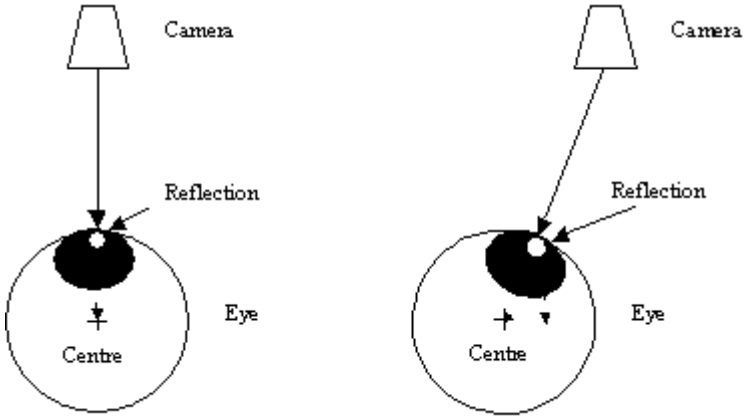

Figure 64. How the reflection moves relative to the eye's origin

6.3 System Configuration

The system was developed in such a manner that further sections could easily be added into it and connected to the relevant modules. Therefore, the overall structure of the underlying system is in place and once the improved image processing algorithm is fully developed, it should be implemented into the full system. This will allow for results to be instantly displayed, as opposed to running different stages of the system from various program files.

6.4 Calculating the VOR

There are several methods with which to define the VOR, and although they may not all give the same result, a combination of them could be calculated and displayed. They can each give an indication as to the state of the canal plane, with slightly different physical interpretations. Some of the ways in which the VOR can be calculated are as follows;

- The ratio of the average eye to head velocity during the head rotation.
- The ratio of the maximum eye to head velocity during the head rotation.
- A combination of the above with a measure of the time delay from the impulse in head velocity to the eye saccade.

All of these results could be calculated for each data set, allowing the operator to make the final assessment along with the graphs.

Once the system is operational, it will be possible to assess a patient's eye response and quantify the VOR as a function of head velocity. Little work has been conducted in this field, but the result could appear as follows;

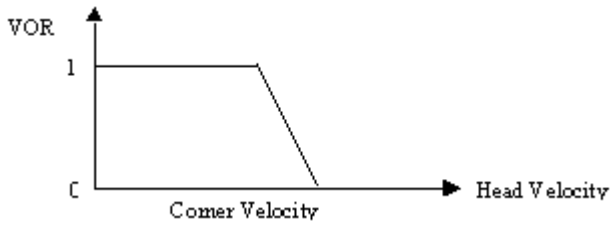

Figure 65. How the VOR is expected to depend on head velocity

This "corner velocity" would give an indication into the level of the canal's loss of functioning in addition to the above methods.

6.5 System Identification

The system identification performed upon data obtained from normal subjects could be used as a foundation for more advanced identification procedures, involving a non-linear representation and allowing data from affected patients to be used for identification and validation purposes. As the number of non-linearities in the relationship between the input head velocity to output eye velocity is not fully known, they could be approximated to be a simple, well defined non linearity such as a relay, representing the eye saccade. The results could then be assessed to determine the validity of such an approximation, and the model could be altered accordingly.

6.6 System Accuracy Assessment

Once the system is at the stage where subjects can be easily assessed, a comparison of results from this system with the search coil technique could be performed. It would perhaps be expected that the search coil technique would yield the more accurate results, as it can fully account for all three degrees of head and eye motion. There is no obvious reason though, to suggest that the imaging based system under development could not yield results with a comparable level of accuracy. One could suggest that the signal characteristics would differ between the two methods, as the potential sources of noise and other signal inaccuracies would be different. But the results shown in this work tend to suggest that this method does have a strong claim to being able to assess a subject and define the VOR for each canal.

7 Conclusions

The diagnosis and treatment of vestibular disorders have been highlighted as important parts of life for many patients who suffer from a range of vestibular problems. This work has been solely concerned with the diagnosis of problems within the semi-circular canals of the vestibular system. The main technique currently available to assess patients, the Search Coil Technique, is rather limited in its flexibility and practicality as it involves a high degree of system calibration, frequent replacement of the search coils and the cornea must be anaesthetised before each test.

The Impulse Test System under development has several advantages over the search coil method. It only requires a series of simple calibration tests to be performed and with the new method for tracking the head and eyes, some of the original calibration methods will become redundant. Integrating modern technology with medical practices makes the assessment easier both for the operator and the patient. The operator no longer needs to set up the search coil arrangement and the patient does not need the coils placed upon their cornea. Other than placing a paper marker on the patient's forehead and the operator performing the test, no other contact with the patient is necessary. Hence, no difficult and time consuming procedures are required with this system, and although it is not yet complete, many of the fundamental building blocks are in place to proceed with more work, as discussed in section 6.1.

With this digital imaging based algorithm, one of, if not the most important area of the system is the image processing algorithm. Several approaches were made to define an accurate algorithm, each with its own advantages and disadvantages.

The first attempt at an image processing algorithm employed Pattern Matching in each of the images recorded from the camera. This involved selecting a region around each of the pupils and the head marker in a predefined manner within the first image. The advantages of this method are that the system can extract features of the eye from the first frame and use them to determine the eye's position in subsequent frames. With tests upon normal patients, the algorithm worked well, allowing the pupil and the head to be tracked with a high degree of accuracy, depending on the experimental conditions. The method was simple, easy to develop within LabVIEW and the properties available to optimise the pattern matching yielded excellent signal quality.

The disadvantages somewhat outweighed the benefits though. The lighting conditions meant that the reflection on the eye was a prominent feature, even to the extent that it became the more important to the pattern matching technique than the pupil. The result of this was that the system was forced to compromise between the pupil and the reflection, yielding inaccurate eye velocity results. When data from affected patients was analysed, the reflection appeared to move relative to the pupil, making it very difficult to track the pupil or reflection accurately. As the eye initially rotated with the head, the reflection appeared to move with respect to the pupil. The pattern matching normally failed at this point, resulting in another object in the frame, closest in contrast terms, being located. This obviously yielded meaningless results and was most visible in the results of tests performed on affected patients. There were two other issues with this method of image processing, which were considered detrimental. Firstly the shape of the selected region of each of the objects in question appeared important because the region had to contain enough contrast information for the image processing algorithm to use accurately in subsequent image frames. The selection of the pupil object proved most difficult as the contrast levels between the

pupil and the iris are generally small. So, the sclera had to be included in the selected region, but again, this had to be done with care. As the eye rotated, the sclera disappeared behind the bridge of the nose or behind the eyelids, reducing the quality of any matching template found. Secondly the values assigned to each of the pattern matching properties affected how the function performed. The most important parameter appeared to be the lower bound upon the acceptable accuracy of the pattern matching function. If the value was set too high, changes in the object, during rotation for example, would result in image frames where the object could not be matched accurately enough. If the value was set too low and the object, again, changed in appearance, another feature in the image could be matched instead. Finding a different object yields large outliers in the resulting velocity data. The most accurate results were obtained when the head was in close proximity to the normal, resting position because this was the position where the pattern templates were defined. As the head rotated, the accuracy of the resulting templates depended upon the eye response. If the eye compensated for the head movement sufficiently well, then the eye appeared very similar to the original template. But if the eye did not compensate, then its appearance was significantly different to that from the first frame and the resultant accuracy decreased.

The second, and much improved, image-processing algorithm involves more complicated procedures to extract the relevant features of the eye, but the results are much better than those obtained from the first method and equally as importantly, many of the problems associated with the previous method are no longer applicable. It involves a combination of thresholding, area and edge detection, particle analysis and automatic contrast measurement on a selected region around the eye. This algorithm does suffer from a different set of problems, but as described in section 6.1, extending this method to incorporate a strobe light should remove many of them.

The three dimensional head model was initially intended to represent head movements and describe the location of the eye within the head. The model was designed to have two separate rotation points, representing the vertical and horizontal movements. Relative to these points, the locations of the eye and the head marker were described. Initially, approximate values were used to represent the physical dimensions of the head, but they required manual tuning in order to obtain the best results. Therefore a series of calibration tests were devised and implemented in an attempt to automatically determine the dimensions required, with varying degrees of success. But with the advent of the new image processing method, the head model is no longer required to describe the eye's location as this is done using the reflection. One would intuitively expect the eye's origin to be more accurately predicted within a small search area around the eye compared to predicting it over a volume the size of the head, using values for the head radii with as much as 20% uncertainty. The head model is perhaps over complex, given the level of accuracy resulting from the calibration methods. But as the head model no longer needs to predict the eye's origin, it could be simplified to account for head rotations only.

Therefore, the best combination of methods developed so far involves the newer image processing algorithm for tracking the eyes, a head model which will only be used to represent three-dimensional angular head velocity and a separate model for the eye. More development work, as described in section 6.2 is required to obtain the

next stage of system, whereby initial testing can be conducted with the aim of assessing patients, in parallel with the system's performance and accuracy.

The results of the system identification performed upon the initial system set-up showed that the proposed hypothesis was correct. Although the oversimplified two-dimensional model was employed, the experiments confirmed that under the correct conditions, the non-linear system could be approximated with linear techniques. A Box Jenkins model with first order system polynomials and zero order noise polynomials yielded the best results from the parametric modelling section. The subspace based approach yielded best results with a second order model, although they were not as accurate as those obtained with the Box Jenkins modelling. The system identification highlighted issues with respect to obtaining different results with different methods and therefore the methods which were successfully used in this experiment should be employed for further work. As described in section 5.2.3, the unexpected results could be due to the numerical representations employed or some other, as yet unknown, reason.

8 Acknowledgements

The author of this project would like to thank the following people and institutions for their contributions to this project and the exchange placement.

Professor Rolf Johansson, Department of Automatic Control, Lund University.
Project supervisor and technical advisor.

Professor Måns Magnusson, ENT Department, Lund University Hospital.
Project supervisor and medical advisor.

Doctor Mikael Karlberg, ENT Department, Lund University Hospital.
Project supervisor and medical advisor.

Per Anders Fransson, ENT Department, Lund University Hospital.
Project supervisor and technical advisor.

Doctor Philippe De Wilde, Imperial College, London.
Exchange coordinator at home institution.

Christina Grossman, International Office, LTH.
Exchange coordinator at host institution.

9 Bibliography

[1] K. Sung and M. Reschke, "*A Model Based Approach for the Measurement of Eye Movements Using Image Processing,*" Nasa Technical Paper 3680, April 1997

[2] P. Cremer, G. Halmagyi, S. Aw, I. Curthoys, L. McGarvie, M. Todd, R. Black and I. Hannigan, "*Semicircular canal plane head impulses detect absent function of individual semicircular canals,*" Oxford university Press 1998

[3] T. Haslwanter, "*Mathematics of Three-dimensional Eye Rotations,*" Department of Psychology, University of Sydney, Sydney

[4] T. Haslwanter, G. Halmagyi, I. Curthoys, R. Yavor and M. Todd, "*Three – Dimensional Vector Analysis of the Human Vestibuloocular Reflex in Response to High-Acceleration Head Rotations. I. Responses in Normal Subjects,*" The American Physiological Society 1996

[5] T. Haslwanter, G. Halmagyi, I. Curthoys, R. Yavor and M. Todd, "*Three – Dimensional Vector Analysis of the Human Vestibuloocular Reflex in Response to High-Acceleration Head Rotations. II. Responses in Subjects with Unilateral Loss and Selective Semicircular Canal Occlusion,*" The American Physiological Society 1996

[6] M Fetter and T Haslwanter, "*3D eye movements-basics and clinical applications,*" Journal of Vestibular Research 9 (1999) 181-187

[7] Y. Huang and T Huang, "*Facial Tracking with Head Pose Estimation in Stereo Vision,*" IEEE ICIP 2002

[8] M. Malciu and F Preteux, "*A Robust model-Based Approach for 3D Head Tracking in Video Sequences,*" Journal of Electronic Imaging

[9] National Instruments IMAQ PCI/PXI-1422 User Manual. April 2001 Edition, Part No 322158C-01

[10] National Instruments IMAQ Vision Concepts Manual, October 2000 Edition, Part No 322916A-01

[11] R. Johansson, "*System Modelling and Identification,*" Prentice Hall, Englewood Cliffs, New Jersey 1993

[12] L Andersson, U Jönsson, K H Johanson, "*A Manual For System Identification*"