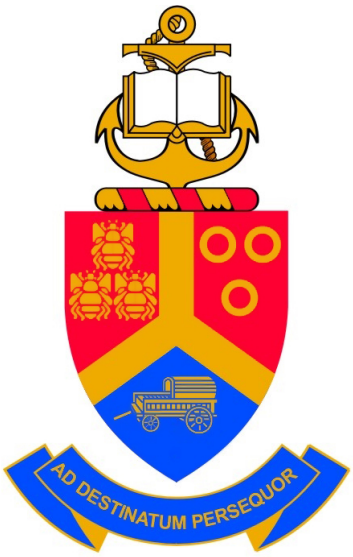

3D COMPUTER VISION CONTACT PATCH MEASUREMENTS INSIDE OFF-ROAD VEHICLE TYRES

by

Alan Glenn Guthrie

Submitted in partial fulfilment of the requirements for the degree

Master of Engineering (Mechanical Engineering)

in the

Department of Mechanical and Aeronautical Engineering
Faculty of Engineering, Built Environment and Information Technology
UNIVERSITY OF PRETORIA

February 2016

SUMMARY

3D COMPUTER VISION CONTACT PATCH MEASUREMENTS INSIDE OFF-ROAD VEHICLE TYRES

by

Alan Glenn Guthrie

Supervisor(s):	Prof P. S. Els
Department:	Mechanical and Aeronautical Engineering
University:	University of Pretoria
Degree:	Master of Engineering (Mechanical Engineering)

The interaction between tyres and terrain is one of the most studied areas in the vehicle dynamics and terramechanics research communities because it is the only region where excitation forces act on the vehicle if aerodynamics is not considered. The tyre area which deforms against the ground is called the contact patch. Measuring the contact patch has been accomplished statically in the past; however, measurement while a wheel rotates has proven difficult.

A number of attempts to measure carcass deformation from inside the tyre have succeeded in measuring small areas or single points but full field measurement has never been attempted. This study describes the design and testing of a system which uses stereo cameras and image correspondence to measure the deformation of the inside of the tyre carcass in the contact region completely. The system includes a stabilisation mechanism which prevents the cameras from rotating, ensuring that the cameras view the inside of the contact region at all times.

Software to capture and process the images captured is developed and tested to ensure measurement accuracy. The 3D results produced by the software are compared to one another where possible and any trends or problems are discussed. Results indicate that the full 3D displacement field in the contact region can be measured accurately. The information produced is expected to be extremely valuable for development and validation of tyre and vehicle dynamics models.

LIST OF ABBREVIATIONS

Abbreviation	Expansion
2D	Two dimensional
3D	Three dimensional
ABS	Anti-lock Braking System
ADC	Analog-to-digital converter
BM	Block Matching
CAN-bus	Controlled area network bus
CCD	Charge coupled device
CMOS	Complimentary metal-oxide-semiconductor
DAQ	Data acquisition system
HTD	High torque drive
JPEG	Joint Photographic Experts Group
LED	Light emitting diode
LHS	Left hand side
MOSFET	Metal-oxide-semiconductor field-effect transistor
OpenCV	Open Computer Vision Library
RHS	Right hand side
SAD	Sum of Absolute Differences
SAW	Surface Acoustic Wave
SGBM	Semi-global block matching
USB	Universal Serial Bus
WFT	Wheel force transducer

LIST OF SYMBOLS

Roman Symbols

Symbol	Description	Units
0	A zero displacement vector	m
b	The baseline distance between the optical centres of the left and right cameras	m
C	The camera matrix which is the product of the intrinsic and extrinsic matrices	
C_l	The camera matrix which is the product of the intrinsic and extrinsic matrices for the left camera	
C_r	The camera matrix which is the product of the intrinsic and extrinsic matrices for the right camera	
c	The capacitance in applied at the shift register output node	F
d	The disparity between x_{P_l} and x_{P_r}	pixel
d_x	The x dimensions of a single pixel	m
d_y	The y dimensions of a single pixel	m
e_l	The location of the epipole on the left image plane	pixel
e_r	The location of the epipole on the right image plane	pixel
f	The focal length of the camera lens	m
H	The projective transform called the homography matrix describing a plane pose	
H_i	The ith homography matrix describing a plane pose	
H_l	The projective transforms which makes the left image lie on the same plane as the left image	
H_r	The projective transforms which makes the right image lie on the same plane as the right image	
h^i_j	The jth column of the homography matrix R describing the ith plane pose	
I	An identity matrix for rotation	
I_l	The intensity at a pixel coordinate in the left image	

I_r	The intensity at a pixel coordinate in the right image	
K	The intrinsic matrix of the camera	
K_l	The intrinsic matrix for the left camera	
K_r	The intrinsic matrix for the right camera	
N	Number of electrons at the shift register output node	
N_s	The number of teeth on the sun gear	
N_p	The number of teeth on the planetary gear	
O_l	The optical centre of the left camera	m
O_r	The optical centre of the right camera	m
P	A point in 3d space	m
P_1	A second point in 3d space	m
P_2	A third point in 3d space	m
P_l	The projection of the 3D point P on to the left image plane	pixel
P_r	The projection of the 3D point P on to the right image plane	pixel
q	Electrical charge (1.6×10^{-19})	Coulombs
R	The rotation matrix of the right camera relative to the left camera	
R_c	The rotation matrix between the world and camera coordinate systems	
R_l	The rotation matrix of the left camera	
R_r	The rotation matrix of the right camera	
R_x	The rotation matrix about the x axis	
R_y	The rotation matrix about the y axis	
R_z	The rotation matrix about the z axis	
R_{data}	The combination of rotation matrices about R_x, R_y, and R_z	
r	Radius of distortion in the lens	m
r_j	The jth column of the rotation matrix R_c	
r_j^i	The jth column of the rotation matrix $R_{[c]}$ describing the ith plane pose	
s	A scale factor used to allow equivalence in the pinhole camera model	
t	The translation vector of the right camera relative to the left camera	m
t_c	The translation vector between the world and camera coordinate systems	m
t_l	The translation vector of the left camera	m
t_r	The translation vector of the right camera	m

t_c	The translation vector between the world and camera coordinate systems	m
t_c^j	The translation vector between the world and camera coordinate systems for the jth plane pose	m
V	Voltage at the shift register output node	V
X_c	The x position of a 3D point in the camera coordinate system	m
X_w	The x position of a 3D point in the world coordinate system	m
x	The x coordinates of a point in the image plane	pixel
x_0	The x position of the image plane centre	m
$\tilde{x}$	The x position of a projected point on the image plane	m
$\tilde{x}_{dist}$	The x distorted position of a projected point on the image plane	m
x_{P_l}	The x coordinate of the point P_l on to the left image plane	pixel
x_{P_r}	The x coordinate of the point P_r on to the right image plane	pixel
Y_c	The y position of a 3D point in the camera coordinate system	m
Y_w	The y position of a 3D point in the world coordinate system	m
y	The y coordinates of a point in the image plane	pixel
y_0	The y position of the image plane centre	m
$\tilde{y}$	The y position of a projected point on the image plane	m
$\tilde{y}_{dist}$	The y distorted position of a projected point on the image plane	m
y_{P_l}	The y coordinate of the point P_l on to the left image plane	pixel
y_{P_r}	The y coordinate of the point P_r on to the right image plane	pixel
Z_c	The distance from the aperture of a 3D point in the camera coordinate system	m
Z_w	The z position of a 3D point in the world coordinate system	m
κ_i	The ith distortion coefficient for modelling the lens	
λ_i	The scale factor for the ith plane pose	
ω	The matrix used to estimate the intrinsic coefficients of the camera	
ω_c	The angular velocity of the carrier ring	rad/s
ω_p	The angular velocity of the planetary gear	rad/s
ω_s	The angular velocity of the sun gear	rad/s

Greek Symbols

Symbol	Description	Units
α_x	The ratio between the focal length and the horizontal pixel dimension	
α_y	The ratio between the focal length and the vertical pixel dimension	
κ_i	The ith distortion coefficient for modelling the lens	
λ_i	The scale factor for the ith plane pose	
ω	The matrix used to estimate the intrinsic coefficients of the camera	
ω_c	The angular velocity of the carrier ring	rad/s
ω_p	The angular velocity of the planetary gear	rad/s
ω_s	The angular velocity of the sun gear	rad/s

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

1.1 CONTEXT OF THE PROBLEM

The interaction between tyres and terrain is one of the most studied areas in the vehicle dynamics and terramechanics research communities. This is because the areas where the tyres contact the road are the only regions where excitation forces act on the vehicle if aerodynamics is not considered. The area where the tyre deforms to contact the road is known as the contact patch or tyre footprint and is very important when attempting to understand the forces applied to a vehicle by the road surface.

The size and shape of the contact patch is difficult to measure because it cannot be directly observed under normal driving conditions. Measurements are therefore currently done statically in a lab with the tyre depressed onto a markable surface or onto a pane of glass with a camera on the other side. Valuable information on the contact patch and tyre deformation can be obtained by measuring displacement of the inside surface of the tyre. It could be argued that since the carcass stiffness is considerably higher than the global tyre stiffness a fairly accurate estimation of the outer tyre geometry can be obtained by measuring the inside surface and superimposing the carcass thickness onto this data. This may be very accurate in terramechanics applications where the shear forces between the tyre and terrain are relatively low.

As discussed in subsection 2.1.2 all attempts to measure dynamic tyre deformations are limited to either single point or a small region inside the tyre and are used to estimate other vehicle parameters (e.g. vertical load or friction) rather than a dedicated system for

observing the contact patch. Static measurements, while extremely valuable, are not an accurate representation of the contact patch that occurs when the tyre is rolling or deforming over an object (such as a cleat) and the existing dynamic systems do not capture enough information to represent the entire contact patch.

Due to the shortcomings of existing measurement techniques there exists an opportunity to develop an entirely new system which will allow measurement of the full contact patch while the vehicle is traversing any road surface.

1.2 RESEARCH OBJECTIVE

The aim of this study is to design, implement and test a new system which will solve some inherent problems of dynamic tyre deformation measurement by approaching the problem with a novel technique. The system should therefore be able to measure as large an area as possible while the rolling tyre deforms over both discrete objects, and varied terrain.

1.3 PROPOSED SOLUTION

This study proposes measuring a large deformed area from inside the tyre using a pair of inexpensive, off-the-shelf, high speed digital cameras. The two cameras will be mounted next to one another (facing the same direction) to capture stereo images which can later be used to produce sets of 3D data points through the application of stereo correspondence techniques. The system will require a mechanism which ensures that the cameras face the area of interest at all times, lights to illuminate the measured area and a data acquisition system manage the cameras and store the images.

1.4 RESEARCH CONTRIBUTION

The system will be a proof of concept to investigate the viability of using stereo cameras inside a tyre as a measurement technique for the dynamic deformation as well as to develop the experience and algorithms required to extract useful data from the raw camera images. The study aims to produce data which is currently unavailable to researchers in the vehicle dynamics and terramechanics communities and discover paths for future research. Techniques developed in this study could be extended to facilitate the capture not only the portion

of the tyre in contact with the road but the response of the entire tyre as obstacles are negotiated.

1.5 OVERVIEW OF STUDY

The study is divided into the following sections:

Literature Study: A thorough review of current contact patch measurement techniques and some important parameters when choosing a digital camera.

Computer Vision Techniques: An explanation of the computer vision techniques used in the study as well as their software implementation.

Measurement System Design and Implementation: An explanation of how the mechanism , power system, and data acquisition system interact with one another.

Testing and Results: A description of the accuracy of the system, the test procedure, the results produced, and an interpretation of results explaining any issues in the data and the reasons for them.

Conclusion: A discussion of the findings made in the study, the performance of the system and avenues of future research.

CHAPTER 2

LITERATURE STUDY

The literature study presented below attempts to describe the existing technology and techniques used to measure both static and dynamic tyre deformations.

2.1 TYRE DEFORMATION MEASUREMENT

Traditional pneumatic tyres have three regions through which forces are transferred. The first being the area where the tyre contacts the road and the other two are the beads which contact the rim (Potenger, 2006). The tyre-road interface is extremely important since the high dynamic and static forces developed by the vehicle are applied at one of the road interfaces.

The area of the tyre which deforms to conform to the road is referred to as the contact patch. The contact patch size defines the pressure and dimensions of the compressed soil area underneath the tyre; therefore, the applied force can be directly linked to the area of the contact patch. Tyre flexibility, and the driving torque required to rotate the wheel can also be determined from the dimensions of the contact patch (Sharma and Pandey, 1996). Study of the contact patch is also extremely complex because, aside from being hidden under normal vehicle operating conditions, tyres have a doubly curved profile. A doubly-curved profile is one that cannot conform to a flat or round surface by simply bending, but rather requires tension and compression of the tyre compound in addition to bending to conform to the surface applying a load (Potenger, 2006).

The size and shape of the contact patch is becoming more important in traditional vehicle dynamics as new tyre models (such as Cosin's F-Tire (Cosin, 2015)) are developed and

need more information both for characterisation and validation. In the terramechanics field, knowing the size and shape of the contact patch is extremely important because it defines the pressure on the soil below the tyre. The contact patch can therefore describe both the deformation of the tyre due to the road and (on soft soil) the deformation of the soil due to the tyre.

2.1.1 Static Contact Patch Measurement

The measurement of tyre deformation has until recently been almost exclusively measured on stationary tyres due to the difficulty of measuring the deformations while the tyre is rotating. Some standard techniques used for capturing static measurements, as described in Potenger (2006) and Sharma and Pandey (1996), are:

- Placing carbon paper beneath the wheel on an extremely flat surface and imprinting the contact patch onto a piece of heavy paper (often referred to as ink-block printing).
- Glass plate photography is used to take contact patch measurements by directly viewing the tyre deformation from beneath a pane of glass.
- Pouring plaster of Paris into tyre imprints until the height was 6mm greater than that of the lugs when dealing with soft soil can provides some indication of tyre deformation.

These techniques have the potential to produce good results for static tests however once the tyre begins to rotate the shape of the contact patch changes considerably.

2.1.2 Dynamic Contact Patch Measurement

Some attempts have been made to modify static measurement techniques to capture the dynamic case:

- Glass plate photography has been used by allowing the camera to take multiple shots as the tyre rolled over the plate (Potenger, 2006).
- Wulfsohn and Upadhyaya (1992) ran a thin piece of shielded wire across a patch of soil in the tyre's path and connected the wire to a spring-loaded potentiometer. The extension of the wire was measured as the tyre rolled over it and used to fit a contact patch size estimation model to the extension data.

Until recently these attempts have been relatively primitive, with no studies attempting to instrument the tyre to capture the contact patch deformation continuously while the tyre is rolling. This has changed in recent years however with several authors developing sensor systems to dynamic deformation.

Magori et al. (1998) developed an ultrasonic sensor to measure the load on the tyre and the vertical deflections. The sensor was mounted to the rim, inside the tyre, and measured a small area beneath the sensor and as such only produced useful vertical deflection information once per revolution. Tests were performed by rolling the tyre over a 3m track with the wheel mounted to a rig which could change the vertical load, the camber and the slip angle. The data from the sensor was compared against a conventional measurement system consisting of an inductive sensor for deflection and a load cell for the vertical force. The system was able to accurately measure the tyre deflection at discrete rotation intervals and provide some indication of the deflection when the sensor was not directly over the contact patch. Consequently, the vertical force could also be estimated accurately for each interval if the tyre pressure was known.

Tuononen (2009) designed and implemented an optical sensor system utilising a position-sensitive detector mounted on the rim, inside the tyre, and a light emitting diode (LED) directly beneath it glued to the tyre carcass. The position of the LED relative to the sensor allowed the vertical and lateral deflection to be measured. Vertical and lateral forces were estimated based on these deflections. The sensor was mounted through a specially designed rim and was able to communicate measurements through a standard Controlled Area Network (CAN) bus with integrated error checking and wireless communication. Tests were performed on a Volvo FH12 truck and results proved promising for deflection and force estimation. This system ultimately suffers from the same drawback as Magori et al. (1998), where contact patch deformation can only be measured once per revolution because only a single point can be measured at a time.

Sergio et al. (2006) proposed a method of measuring tyre deformation continuously using the tyre as a sensor. The steel wires, which add stiffness to the tyre carcass, are used as electrodes in an array of passive impedances where the deformation of the tyre affects the magnitude of these impedances. A large amount of information is made available by the system, such as instantaneous pressure and deformation by direct measurement and provided good results in

the testing performed. The main drawback of the system is the need to electrify individual wires in the carcass meaning that the only experimental testing performed was on rectangular cut-outs of car tires and no vehicle tests were performed on rotating tyres. No indication was given whether methods to perform these tests on full tyres was ever investigated.

Pohl et al. (1999) used passive Surface Acoustic Wave (SAW) sensors to measure the lateral and longitudinal deformation of a single tread element from inside the tyre. These sensors are passive requiring no power and only capture measurements when a request radio signal is applied to them. The sensors were embedded in the tyre carcass with a small arm penetrating through to the tread element of interest. This arm would deflect when the element was loaded and the corresponding deflection can be measured by the sensors. The purpose of these measurements was to estimate the friction between the road surface and the tyre to overcome one of the problems of ABS systems that the friction coefficient is not known before the wheel reaches full lock and the pressure is reduced. The system was tested in a rolling tyre at speeds of up to 140km/h. While successful in its aim of measuring friction at the road surface, the system itself can only capture data at discrete points which are widely spaced. This is not ideal for contact patch measurement.

Sandu et al. (2012) developed a system for measuring the rolling radius from inside a tyre. This system is used to control the longitudinal slip by accurately measuring the rolling radius of a tyre being tested in a terramechanics rig. Twelve sensors each equipped with five infra-red LEDs and covering a 25° range within the tyre are mounted to the rim. Timestamped data from the sensors is stored on a memory card inside the tyre and the data is retrieved wirelessly. The sensors provide coverage over most of the inner surface of the tyre and only the sensors facing the contact patch are sampled (determined using an accelerometer). The system is very effective at finding the rolling radius of the tyre under controlled conditions however it is battery powered and can only be used for few test runs before it needs to be recharged.

Matsuzaki and Todoroki (2007) proposed using flexible capacitive sensor to measure the strain in the tyre and transmit the measurements wirelessly to a central processor. The intent of the system was to investigate the ability of such a sensor to produce reliable strains for friction estimation. The sensor comprised of two capacitive copper substrates (one each for the lateral and longitudinal directions) embedded in an ultra-flexible epoxy resin. Under strain

conditions the resin would take the majority of the strain and be measured by observing the capacitance across the substrates. A major drawback of this system was that the capacitance in the sensor was sensitive to temperature and when tested in field conditions the sensor required extensive compensation to produce usable results. The study concluded that the sensor in its current form was not suitable for tyre strain measurement.

A common theme in the papers described above is that of the "intelligent tyre". This refers to a tyre which is no longer just a piece of rubber meeting the road but rather it becomes an integral part of the sensor system used to aid the driver. This concept has led to most of the systems being useful only for determining vertical and lateral deflection of a single point and then estimating important vehicle forces and characteristics rather than measuring the entire contact area for tyre and terrain characterisation. Due to this concentration on force estimation rather than contact patch deformation there is a complete lack of literature exploring dynamic tyre deformation measurement for tyre and terrain modelling. There, therefore exists a need for a system which can measure the entire contact patch at all times during while the tyre is rotating.

2.1.3 Computer Vision Measurement of Contact patch

Hiraoka et al. (2009) was the first to propose using computer vision inside a tyre to measure deformations. A study was conducted into the use of a single camera with digital image correlation techniques to measure the deformation and strain in a fixed section of the tyre as it rotated. Deformation was measured using a method called spotlight projection. In spotlight projection a laser point is shone onto a patch of tyre a known location, if the camera detects that the laser point has moved relative to the known location it is possible to calculate the deformation of the surface. The system proved to be extremely accurate with a sub 1% error for the deformation at the laser points.

Green (2011) describes the implementation of a system similar to that of Hiraoka et al. (2009) but using a high resolution camera and a speckle pattern to determine the deformation with sub-pixel accuracy. The system proved to be accurate but suffered from lens defects which made measurement of the out-of-plane deformation impossible.

Botha et al. (2012) is the direct predecessor to the present study where the idea for using

stereovision cameras on a track to measure contact patch deformations was originally proposed. The paper describes the construction of such a system and outlines some required components but lacks detail in terms of the design of these components. Also described is the accuracy of such a system and the potential for strain estimation using the data captured by such a system. Specifically the paper does not include a working method of stabilising the cameras to prevent rotation, a method of delivering power inside the tyre or a means of controlling the cameras inside the tyre. The techniques described in this paper have great potential however the problems detailed above need to be resolved before the techniques can be applied.

2.2 CONCLUSION

The research discussed in the literature survey shows that the contact region of the tyre plays an important role in both vehicle dynamics and terramechanics. It also includes many studies aimed at measuring the contact region geometry both dynamically and statically. A few studies had considered the possibility of using digital cameras to measure tyre deformations however none of the literature showed any attempt to measure the contact region consistently as the tyre rotated. Using digital image techniques for tyre deformation measurement represents a new direction and encouraging opportunities for further investigation.

CHAPTER 3

COMPUTER VISION TECHNIQUES

3.1 DIGITAL CAMERA TECHNOLOGY

In order for the reader to better understand the techniques detailed later in this chapter a short description of the different types of digital sensors and the capture & storage of images.

3.1.1 Image Sensors

A digital camera is a device that captures light from scenes using an image sensor to convert light intensities into electronic signals which, after passing it through an analogue to digital converter, can be stored as digital data. Image sensors typically fall into one of two categories: Charge Coupled Devices (CCD) or Complimentary Metal-oxide Semiconductors (CMOS).

Both CCD and CMOS sensors use silicon detectors to convert light into electrical charges. These detectors are sensitive to photons at wavelengths where electrons can be released due to the photoelectric effect (a range which is generally slightly larger than the visible spectrum encompassing ultraviolet and infrared spectrums with bias to the infrared side (Nagappan, 2007)). The silicon detectors are typically arranged in a grid pattern with each detector element of the grid being referred to as a pixel (picture element). A lens is used to focus light onto the grid so that each pixel will capture a unique part of the scene.

The charges developed at each pixel are handled differently depending on the type of sensor

used. A CCD sensor contains both a photo-detector and a shift register at each pixel. All the pixels are exposed at the same moment and for the same duration. The charges developed in the pixels are then moved to a shift register where they are stored as the charges are moved sequentially into the serial register and moved towards a single output node as can be seen in Figure 3.1(a). The charges are converted into voltages at the output node by using (3.1)

$$\mathrm{V} = \frac{Nq}{c} \tag{3.1}$$

where c is the capacitance of the node (in the order of 10^{-13}F), N is the number of electrons at the node and q is the electronic charge (1.6×10^{-19}C) (Durini, 2014). The voltage produced is then amplified by using a Metal-Oxide-Semiconductor Field Effect Transistor (MOSFET) to generate a measurable voltage across an external resistor and converted to a digital value by an analog-to-digital converter (ADC). The output node's charge must be reset to zero after each measurement so that the charge in the node does not accumulate. This reset is performed by a reset MOSFET before the next charge is moved into the ouput node.

Figure 3.1(b) shows the configuration of CMOS type sensor compared to a CCD sensor. The main difference between a CMOS sensor and a CCD sensor is the number and placement of the elctron-to-voltage converters. In a CCD sensor there is only one converter and one output node whereas in a CMOS sensor conversion is performed at each pixel location so there are an equal number of pixels and converters. Amplification and ADC is generally performed on each column separately so an entire row can be sampled simultaneously. This property of the CMOS sensor allows only single rows to be exposed at a time removing the need for a physical shutter if exposure is done sequentially. Sequential exposure to capture an image is referred to as a rolling shutter and is discussed further below.

While most CMOS sensors use rolling shutters, it is possible to build a sensor which exposes the entire photo-sensitive area simultaneously (a global shutter), however the cost generally significantly higher due to the increased number of transistors required. CMOS sensors with global shutters are becoming more common as technology advances and manufacturing costs are reduced.

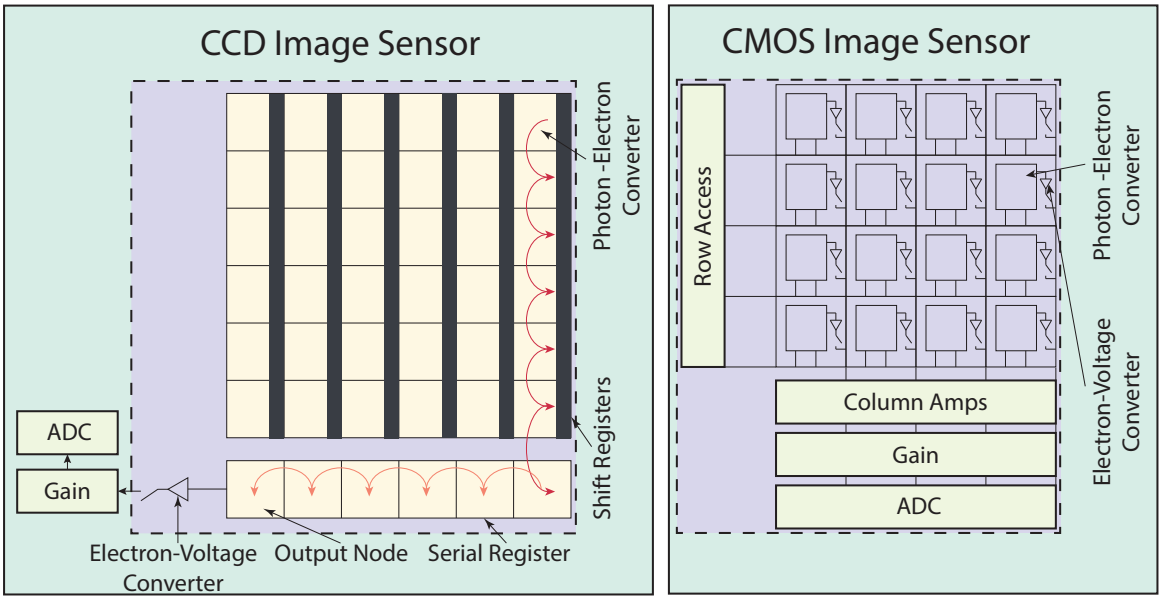

(a) CCD image sensor diagram **(b)** CMOS image sensor diagram

Figure 3.1: Image sensor diagrams

3.1.2 Rolling Shutters Vs Global Shutters

It is important to consider the type of sensor when choosing a camera for computer vision because cheaper CMOS cameras use a rolling shutter to lower the sensor costs. An object in motion will be distorted if captured by a rolling shutter and these distortions cannot be compensated for. Figure 3.2 shows a good example of the distortions caused by rolling shutters.

Rolling shutters are unsuitable for computer vision applications which assume that the captured image is a perfect projection of the 3D world on a 2D plane. A global shutter captures the entire image at the same instant and is more appropriate for computer vision. Global shutters are available in all CCD sensors and some higher quality CMOS sensors.

3.1.3 Digital Representation of Intensity Data

The intensity data captured by the cameras is stored as two dimensional matrices where each matrix element contains three bytes representing red, blue and green light intensities (for a colour image). Since a byte is able to store 256 discrete values it means that an image stored in this manner can have $256^3 = 16777216$ different colour values. Having 3 bytes represent

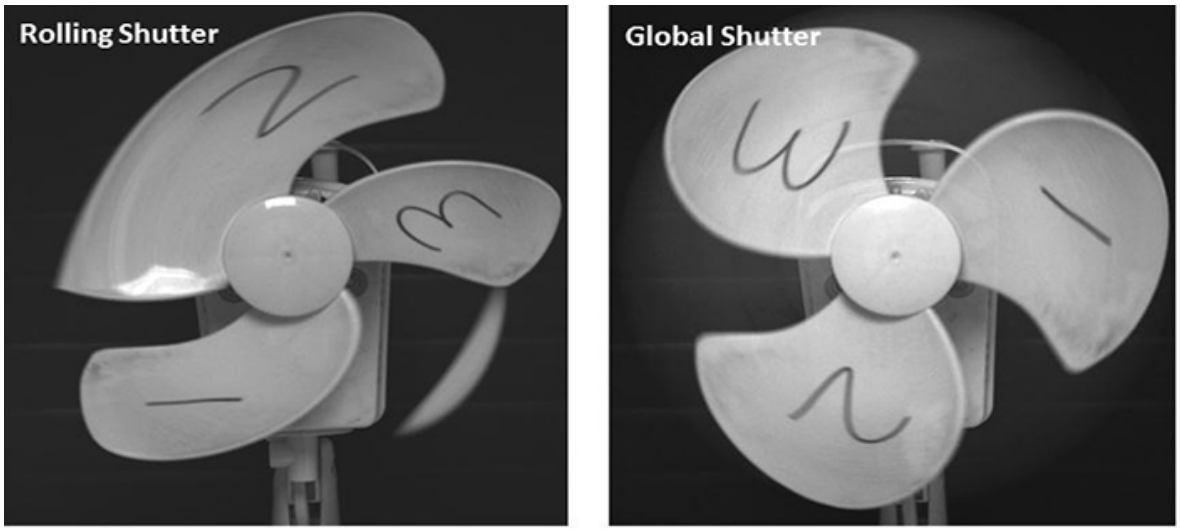

Figure 3.2: Example of a rolling and a global shutter capturing the same moving fan (Andor Technology Ltd, 2015)

every pixel means that image data is often very large with a simple 640 frame taking up 900KB of storage space. Due to these large sizes of this images are normally compressed using an algorithm such as JPEG.

3.2 CALIBRATING CAMERAS AND PROCESSING IMAGES

Images produced by digital cameras are not perfect representations of the 3D scene they are capturing, as the lenses introduce distortions. Fortunately the distortions are constant and their effects can be calculated and removed using well documented techniques. The process of quantifying and removing these distortions is known as camera calibration (discussed in subsection 3.2.1) and is the starting point for most image processing techniques. The use of stereo cameras to capture depth information from scenes requires an extra calibration step described in subsection 3.2.2 to determine the orientation of the cameras in relation to one another. Stereo images can be used to measure depth if the cameras are calibrated correctly and the process of progressing from two images to a 3D point cloud is discussed in subsection 3.2.3 - 3.2.6

The computer vision techniques explained in this section were implemented in software to produce useful results as described in section 3.3

3.2.1 Single Camera Calibration

3.2.1.1 Pinhole Camera Model and Intrinsic Parameters

The simplest means of modelling a camera is to use the pinhole assumption. This assumption states that a single ray of light emanating from an object in a scene enters the camera centre (modelled as a tiny hole) and is projected onto a 2D image plane. The image projected is assumed to be in perfect focus and all object scaling within the image is related to the distance of the object from the camera and the focal length of the lens.

Figure 3.3(a) shows a typical pinhole camera (though only a for a case with a 2D object location) where X_c is the height of the object relative to the camera coordinate system, Z_c is the distance from the aperture to the object relative to the camera coordinate system, f is the focal length of the camera lens and $\tilde{x}$ is the height of the object projected on the image plane. A relationship (3.2) between these parameters can be developed so that the projected height in the image plane can easily be determined (Cyganek and Siebert, 2011).

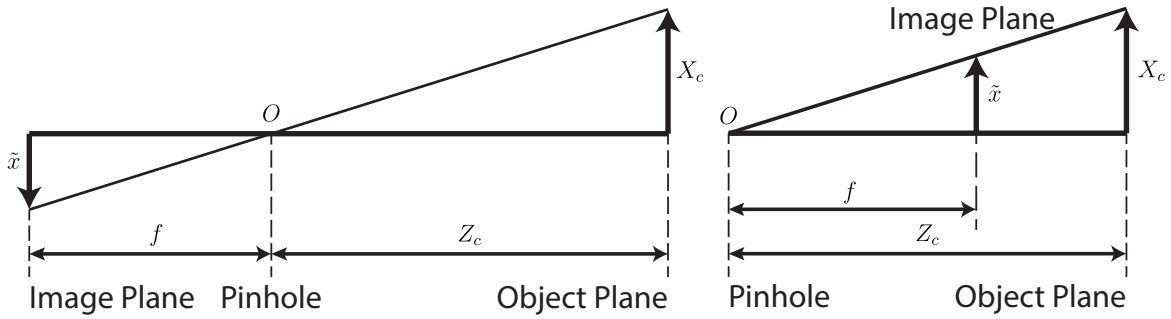

(a) Conventional 2D pinhole model **(b)** Modified 2D pinhole model

Figure 3.3: Pinhole camera models

$$\tilde{x} = -f\frac{X_c}{Z_c} \tag{3.2}$$

It is clear from this equation that the resultant image will be inverted as the ray of light crosses the optical axis as it passes through the pinhole (Radke, 2012). This inversion is often removed (for mathematical convenience) by simply moving the image plane to be in front of the pinhole with light rays intersecting it. This modification is shown in Figure 3.3(b) and results in (3.3).

$$\tilde{x} = f\frac{X_c}{Z_c} \tag{3.3}$$

The pinhole model holds for object points in three dimensions therefore (3.3) can be extended as shown in (3.4)

$$\tilde{x} = f\frac{X_c}{Z_c} \text{ and } \tilde{y} = f\frac{Y_c}{Z_c} \tag{3.4}$$

Equation 3.4 allows a 3D object to be mapped onto a 2D image plane. However, all coordinates on the plane are measured from the plane centre in meters, whereas the digital images are generally measured from the upper left hand corner of the image in pixel values. A transformation is therefore needed to account for these inconsistencies. Equation 3.4 is modified to become (3.5) where d_x and d_y are the dimensions of each pixel measured in meters and the point (x_0,y_0) is in the centre of the image plane (usually referred to as the principal point and x and y are the coordinates in pixel coordinates).

$$x = \frac{\tilde{x}}{d_x} + x_0 \text{ and } y = \frac{\tilde{y}}{d_y} + y_0 \tag{3.5}$$

Equation 3.5 can be written as a matrix of internal camera parameters, usually referred to as the intrinsic matrix K (3.6),

$$K = \begin{bmatrix} \alpha_x & 0 & x_0 \\ 0 & \alpha_y & y_0 \\ 0 & 0 & 1 \end{bmatrix} \tag{3.6}$$

where

$$\alpha_x = \frac{f}{d_x} \text{ and } \alpha_y = \frac{f}{d_y}.$$

We can therefore write a simple pinhole camera model as

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \propto K \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix}$$

where the RHS of the equation is equal up to a scale factor to the LHS. This proportionality may be used to create an equality by adding a scale factor s to the LHS as in (3.7)

$$s \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = K \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix} \tag{3.7}$$

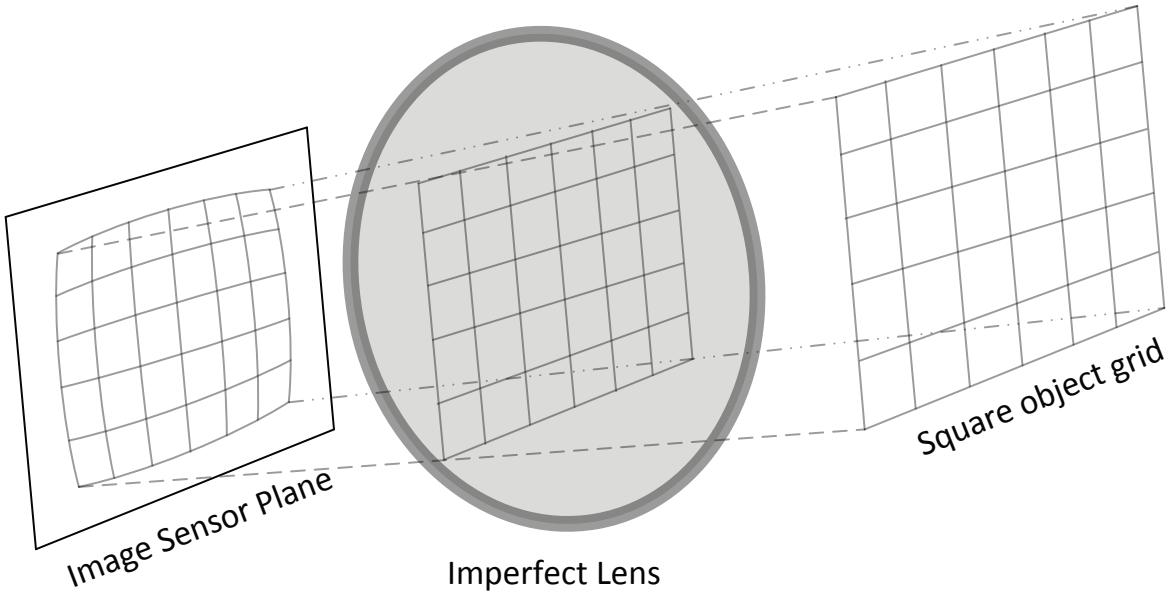

Figure 3.4: Radial Distortion

3.2.1.2 Distortion and Non-ideal Lenses

Real cameras are extremely complicated compared to a simple pinhole camera, often containing multiple lenses and even mirrors which introduce distortion into the images captured. The most common types of distortion are radial and tangential distortion.

Radial distortion is where the magnification of an image either increases or decreases based on distance away from the optical axis of the lens. It is most often caused by cheaply produced lenses or by specially manufactured wide angle lenses (which increase the field of view by introducing radial distortion). Radial distortion can be modelled mathematically using (3.8). This is known as the Brown-Conrady model (Brown, 1966), where κ_1 and κ_2 are distortion coefficients and $r = \tilde{x}^2 + \tilde{y}^2$. Figure 3.4 illustrates the radial distortion effect that an imperfect lens can have on a square object, it is clear that the edges of the image on the sensor curve outward as the distance from the lens centre increases.

$$\begin{bmatrix} \tilde{x}_{dist} \\ \tilde{y}_{dist} \end{bmatrix} = \left(1 + \kappa_1 r^2 + \kappa_2 r^4\right) \begin{bmatrix} \tilde{x} \\ \tilde{y} \end{bmatrix} \tag{3.8}$$

Tangential distortion is distortion caused by misalignment between the image plane and the lens. Figure 3.5 gives an example of how an image is distorted when the image plane is not perpendicular to the optical axis of the lens. These types of lens misalignment are commonly found in cheaper camera equipment as manufacturing tolerances and quality control are less

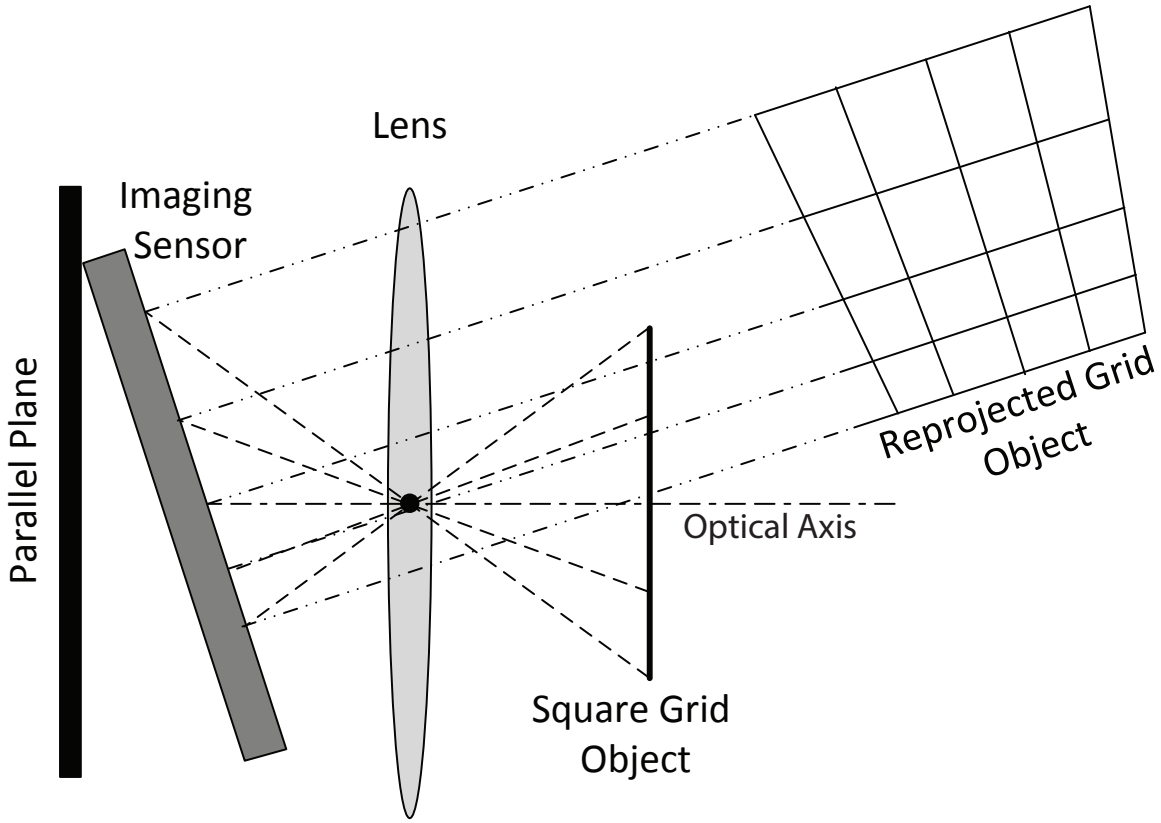

Figure 3.5: Tangential Distortion

stringent. Tangential distortion is generally less apparent than radial distortion however it should still be compensated for if accurate computer vision is to be achieved. The complete Brown-Conrady model (Brown, 1966) was designed to account for tangential as well as radial distortion by simply adding more distortion coefficients to give (3.9).

$$\begin{bmatrix} \tilde{x}_{dist} \\ \tilde{y}_{dist} \end{bmatrix} = \left(1 + \kappa_1 r^2 + \kappa_2 r^4\right) \begin{bmatrix} \tilde{x} \\ \tilde{y} \end{bmatrix} + \begin{pmatrix} 2\kappa_3 \tilde{x}\tilde{y} + \kappa_4 \left(r^2 + 2\tilde{x}^2\right) \\ \kappa_3 \left(r^2 + 2\tilde{y}^2\right) + 2\kappa_4 \tilde{x}\tilde{y} \end{pmatrix} \tag{3.9}$$

3.2.1.3 Extrinsic Camera Parameters

In the pinhole projection model discussed in subsubsection 3.2.1.1, the coordinates in 3D and the coordinates on the 2D plane are assumed to have the same coordinate system. This assumption is not valid because the the object being viewed will have coordinates relative to the universal "world" coordinate system. Figure 3.6 shows a camera with its own coordinate system in space in relation to the world coordinate system. A point in the world coordinate system can be mapped to the camera coordinate system if the rotation and translation of the systems relative to one another are known. These rotations are expressed as a rotation

matrix R_c containing sin and cos functions of the rotation angles about each of the three axes. The translation of the coordinate systems relative to one another are expressed in a vector t_c (shown in (3.10)).

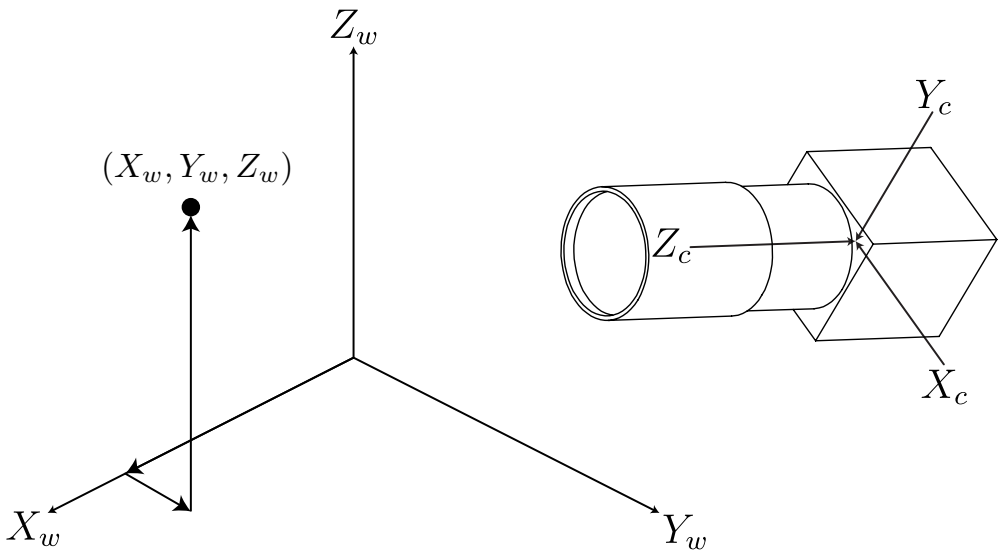

Figure 3.6: Rotation and translation between the world and camera axes

$$\begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix} = \underbrace{R_c}_{\text{Rotation Matrix}} \begin{bmatrix} X_w \\ Y_w \\ Z_w \end{bmatrix} + \underbrace{t_c}_{\text{Translation Vector}} \tag{3.10}$$

Equation 3.10 can be extended by integrating the pinhole camera model (3.7), resulting in (3.11). In this equation the rotation matrix and translation vector are combined in a 3×4 matrix of extrinsic parameters which when multiplied by the intrinsic matrix results in a 3×3 matrix C, usually referred to as the camera matrix.

$$s \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = K \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix} = K \left(R_c \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} + \begin{bmatrix} t_{c_x} \\ t_{c_y} \\ t_{c_z} \end{bmatrix} \right) = \underbrace{K \begin{bmatrix} R_c & t_c \end{bmatrix}}_{\text{Camera Matrix } C} \begin{bmatrix} X_w \\ Y_w \\ Z_w \\ 1 \end{bmatrix} \tag{3.11}$$

3.2.1.4 Plane Based Calibration

The most popular calibration technique currently available was developed by Zhang (2000) who proposed capturing a series of unique images containing a pattern printed on a planar surface. The geometry of this pattern should have distinct features which can be tracked

between images (such as the corners of a checkerboard pattern or the centers of circles in a grid pattern) and be accurately measured in world coordinates. An assumption is made that the planar surface is held fixed in the Z_w plane so $Z_w = 0$, therefore

$$s\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = K\begin{bmatrix} R_c & t_c \end{bmatrix}\begin{bmatrix} X_w \\ Y_w \\ 0 \\ 1 \end{bmatrix} = K\begin{bmatrix} \overbrace{r_1\ r_2\ r_3}^{\text{Columns of } R_c} & t_c \end{bmatrix}\begin{bmatrix} X_w \\ Y_w \\ 0 \\ 1 \end{bmatrix} = K\begin{bmatrix} r_1 & r_2 & t_c \end{bmatrix}\begin{bmatrix} X_w \\ Y_w \\ 1 \end{bmatrix} \tag{3.12}$$

where r_1, r_2 and r_3 are columns of the rotation matrix R_c. When matrix multiplication is performed using the $Z_w = 0$ assumption, r_3 no longer has any effect on the model. The K matrix and rotation-translation matrix are multiplied together to form projective transformation matrix called the Homography matrix H, (3.13). This transform relates a point on the plane in X_w and Y_w coordiantes to a point on the image plane in x and y. Since both of these sets of coordinates are known at all times during calibration the homography for the plane can be estimated.

$$s\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \underbrace{H}_{3\times3\text{ Homography}}\begin{bmatrix} X \\ Y \\ 1 \end{bmatrix} \tag{3.13}$$

If i plane posistions are used we can calculate (3.14)

$$H_i = \lambda_i K\begin{bmatrix} r_1^i & r_2^i & t_c^i \end{bmatrix} \tag{3.14}$$

where λ_i is a scale factor applied at each plane position. With some manipulation it can be shown that

$$\begin{bmatrix} r_1^i & r_2^i & t_c^i \end{bmatrix} = \frac{1}{\lambda_i}K^{-1}\begin{bmatrix} h_1^i & h_2^i & h_3^i \end{bmatrix} \tag{3.15}$$

with h_1^i, h_2^i and h_3^i being columns of H. All rotation matrices have some common properties which prove useful as constraints in this problem. Firstly the all the column vectors are unit norms and secondly that the columns are at 90° to one another. These assumptions are shown mathematically in (3.16) and (3.17).

$$\left\|r_j^i\right\| = 1 \text{ and } \left(r_1^i\right)^T\left(r_2^i\right) = 0 \tag{3.16}$$

$$\left(r_1^i\right)^T\left(r_1^i\right) = \left(r_2^i\right)^T\left(r_2^i\right) \text{ and } \left(r_1^i\right)^T\left(r_2^i\right) = 0 \tag{3.17}$$

When these constraints are applied to the columns of H in (3.18) and (3.19)

$$h_1^{iT}\left(KK^T\right)^{-1}h_1^i = h_2^{iT}\left(KK^T\right)^{-1}h_2^i \tag{3.18}$$

$$h_1^{i\,T}\left(KK^T\right)^{-1}h_2^i = 0 \tag{3.19}$$

it is found that the only unknowns are in the intrinsic matrix, K, which is multiplied by its own inverse and inverted to form a matrix referred to as ω (3.20). The ω matrix is given in full in (3.21) and contains five independent variables. The constraints in (3.18) and (3.19) put two constraints on this system. Therefore, to estimate the value of ω accurately a minimum of three plane positions need to be considered.

$$\omega = \left(KK^T\right)^{-1} \tag{3.20}$$

$$\omega = \begin{bmatrix} \frac{1}{\alpha_x^2} & 0 & -\frac{x_0}{\alpha_x^2} \\ 0 & \frac{1}{\alpha_y^2} & -\frac{y_0}{\alpha_y^2} \\ -\frac{x_0}{\alpha_x^2} & -\frac{y_0}{\alpha_y^2} & \frac{x_0^2}{\alpha_x^2}+\frac{y_0^2}{\alpha_y^2}+1 \end{bmatrix} \tag{3.21}$$

Once ω has been estimated it is a matter of substitution using (3.22) to determine the intrinsic parameters of the camera and therefore the intrinsic matrix K.

$$\begin{aligned} x_0 &= -\frac{\omega_{13}}{\omega_{11}} & y_0 &= -\frac{\omega_{23}}{\omega_{22}} \\ \alpha_x &= \alpha_y\left(\frac{\omega_{22}}{\omega_{11}}\right)^{0.5} & \alpha_y &= \left(\frac{\omega_{11}\omega_{22}\omega_{33}-\omega_{22}\omega_{13}^2-\omega_{11}\omega_{23}^2}{\omega_{11}\omega_{22}^2}\right)^{0.5} \end{aligned} \tag{3.22}$$

With K and H_i known it is possible to determine the rotation matrix R_{c_i} for each frame position. The estimation of the camera matrices using the algorithm given above and estimating the distortion coefficients is often performed simultaneously because the distortion calculation also requires known world coordinates as input.

3.2.2 Stereo Camera Calibration

If two cameras (fixed rigidly with relation to one another) were calibrated individually using the same calibration points, then it is possible to estimate the relative translations and rotation between the two cameras. This is possible because if the cameras do not move relative to one another and view the same planar points it is possible to recover the relative rotation matrix R and translation vector t of second camera with respect to first camera using (3.23)

$$R_r = RR_lt_1 = RT_l + t \tag{3.23}$$

where R_l,T_l, R_r and t_r are the rotation and translation matrices of the left camera and right camera respectively. These rotations and translations are calculated in the individual camera

calibrations so it is possible to use the same poses from the individual camera calibrations to calculate R and T. To do this the reprojection error for all points in the poses is used as a cost function and minimised.

3.2.3 Definition and Importance of Epipolar Geometry

In any stereo camera setup that has its cameras viewing the same 3D point from slightly different points of view it is possible to calculate relations between the projection of the point on the one image plane and the projection of the same point on the other image plane. This relationship is called the epipolar geometry of the two views and is independent of the nature of the 3D world. The epipolar geometry rather depends on the internal parameters of the cameras and their orientation relative to one another (Hartley and Zisserman, 2003). This means that any two cameras which are described by pinhole models will be constrained by their epipolar geometry.

Figure 3.7 shows a left and right image plain containing a projection of the 3D point P at image points P_l and P_r. In this figure the baseline is defined as $O_l - O_r$, the line connecting the projective centre of each camera. In each image there is a point where the centre of projection of the other camera is projected onto the image plane. These points are called epipoles and in Figure 3.7 lie on the baseline at e_l and e_r.

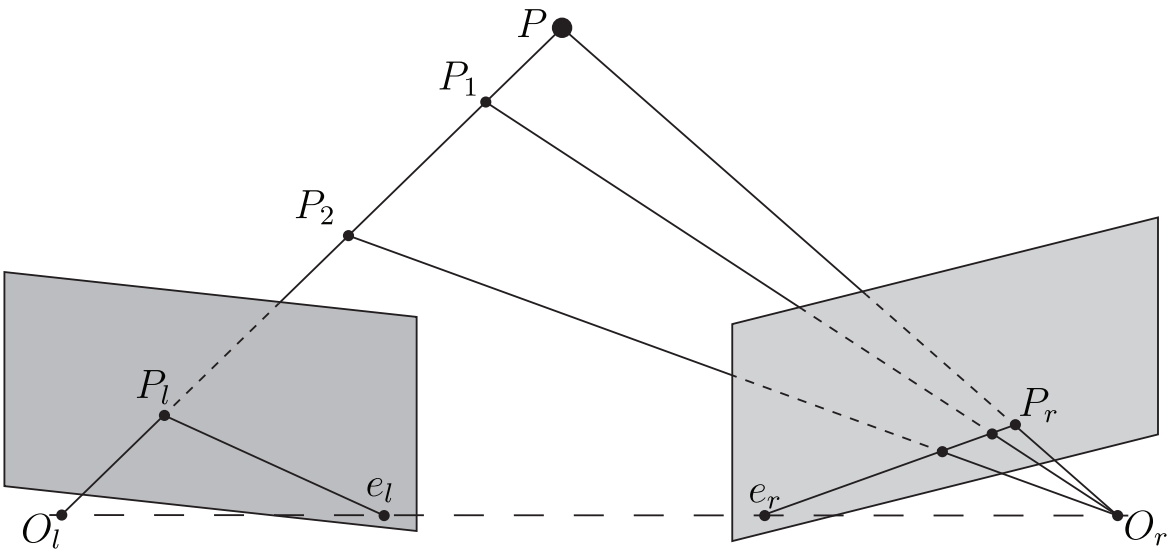

Figure 3.7: A 3D point with projections on the left and right image pair is used to illustrate how epipolar lines are found from a stereo pair

If the lines $P_l - e_l$ and $P_r - e_r$ line are drawn between the projected points and the epipoles

then epipolar lines are formed in both images. These lines form a constraint which maintains that any projected point P_l which is visible in the left image must be present in the right image and lie on the epipolar line $P_r - e_r$ which can be easily calculated. To illustrate this effect P_1 and P_2 are projected onto the left image at P_l and onto the right image along the epipolar line $P_r - e_r$.

The epipolar lines define search regions for matching points in image pairs which is useful for stereography. Figure 3.8 shows the special case where the image planes are transformed to be coplanar. In this case points e_l and e_r where the other camera's centre of projection appears on the opposite camera's image plane are so far away as to be considered infinite. Having epipoles placed at infinity means that the epipolar lines are infinitely long; therefore the lines are considered parallel to the baseline and to one another.

If the relative camera rotation matrix R and translation vector T are known from stereo calibration then it is possible to translate the image planes by T_y (the y component of the translation vector) so that the epipolar lines are co-linear.

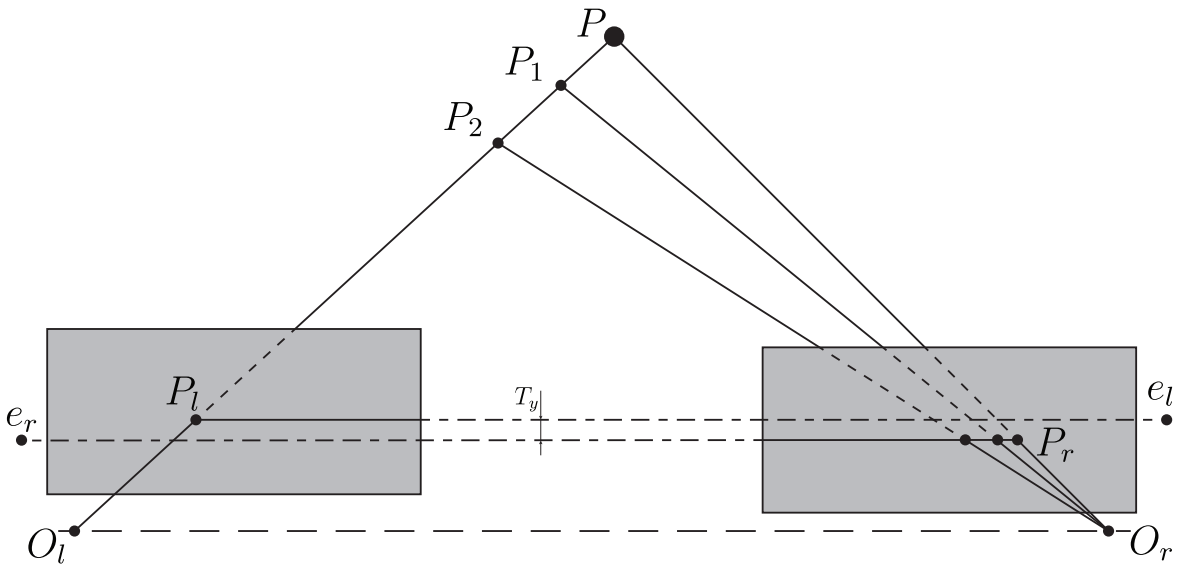

Figure 3.8: A 3D point and with projections on the left and right image pair is used to illustrate how epipolar lines are reach infinity when parallel to one another

The epipoles in the stereo images can be computed using only the intrinsic matrix, the extrinsic matrix of each camera as well as the rotation matrix and translation vector from the stereo calibration. This is accomplished by assuming that the coordinate system of the left camera is aligned with the origin of the world coordinate system and then using the

rotation and translation information to project the optical centres onto each plane. The first step is to compute camera matrices C_l and C_r which describe the projections from the world coordinates onto the image planes as given in (3.24) where I is a 3×3 identity matrix and 0 is a 3×1 zero vector.

$$C_l = K_l \, [I \; 0] \;\text{ and }\; C_r = K_r \, [R \; T] \tag{3.24}$$

These camera matrices are then used substituted into (3.25) and (3.26) to find the position of the epipoles (e_l & e_r) on the image planes using the world coordinates of optical centres O_l and O_r.

$$e_l = C_l \begin{bmatrix} O_r \\ 1 \end{bmatrix} = C_l \begin{bmatrix} -R^T T \\ 1 \end{bmatrix} = K_l [I \; 0] \begin{bmatrix} -R^T T \\ 1 \end{bmatrix} = -K_l R^T T \tag{3.25}$$

$$e_r = C_r \begin{bmatrix} O \\ 1 \end{bmatrix} = C_r \begin{bmatrix} 0 \\ 1 \end{bmatrix} = K_r [R \; T] \begin{bmatrix} 0 \\ 1 \end{bmatrix} = K_r T \tag{3.26}$$

3.2.4 Stereo Rectification

Rectification is the name given to the process of finding projective transforms which project a pair of stereo images to lie on the same plane and consequently have epipolar lines that are parallel to one another. The projective transformations called H_l and H_r are each composed of two distinct rotations. After the images have been reprojected onto a common plane the epipolar lines are shifted vertically to be co-linear with one another.

The first component rotation matrices (H_{l1}, and H_{r1}) of H_l, and H_r are created by rotating the optical axis (at the centre and normal to the image plane) of each image plane so that it is perpendicular to the baseline (as defined in subsection 3.2.3). This is accomplished by computing the cross product between the current optical axis and the desired axis to find an appropriate rotation matrix. This rotation will also ensure that the optical centres of the two image planes are also parallel to one another if the same desired axis is used to compute rotations for both image planes.

The epipolar lines in the images are computed by projecting artificial 3D points onto the image planes and connecting these points to the epipole locations from (3.25) and (3.26). The epipolar lines are then rotated to be parallel to the baseline (and each other) by taking

the cross product between the epipolar line and the baseline to find the rotation matrices H_{l2}, and H_{r2} (components of H_l, and H_r).

The projective transforms H_l and H_r are therefore defined by (3.27).

$$H_l = H_{l1}H_{l2} \text{ and } H_r = H_{r1}H_{r2} \tag{3.27}$$

The translation step is performed by using the translation vector T (known from stereo calibration) to shift the images vertically until their parallel epipolar lines are co-linear. With co-linear epipolar lines the process of finding a match for a point P_l (at pixel value (x_{P_l}, y_{P_l}) in the left image) is a matter of searching along the x axis of the right image at a height of $y_{P_r} = y_{P_l}$ until the matching P_r is found at (x_{P_r}, y_{P_l}) as shown in Figure 3.9.

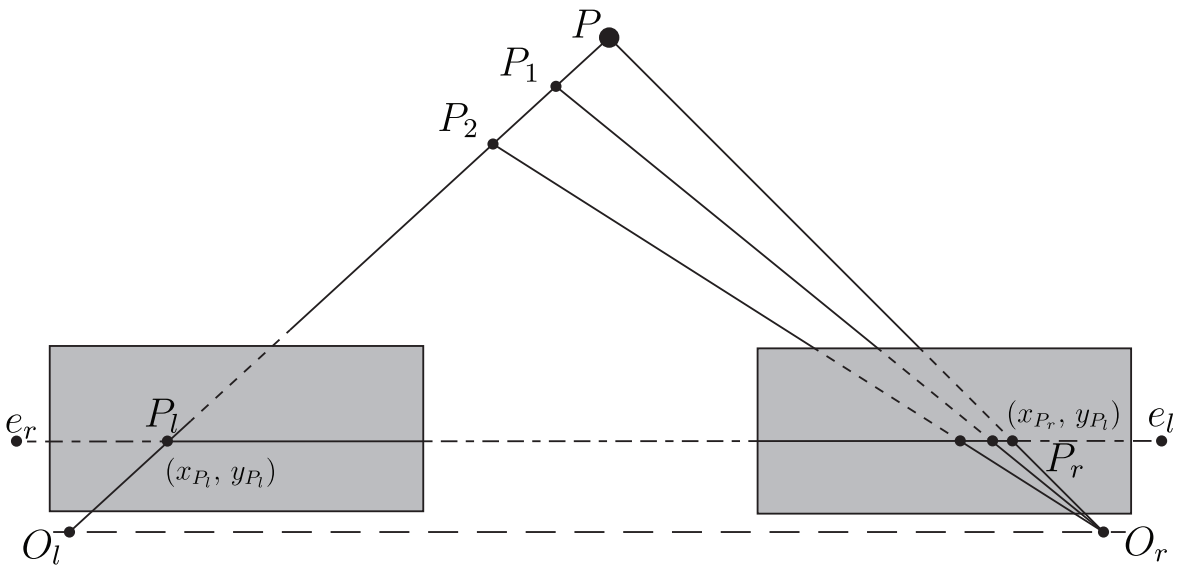

Figure 3.9: Co-linear epipolar lines allow for line searches at the same height in the opposite images

3.2.5 Stereo Correspondence Calculation

The most common use of stereo camera arrays is to extract depth information from a scene using triangulation. To find the depth in a scene it is necessary to find out the distance (in pixel coordinates) between the point on a left image P_l and the point on the right image P_r where the 3D point P is projected. If we once again use Figure 3.9 we see that the distance between the projected points is only in the x direction due to the rectification that has already been performed. This distance is called the horizontal disparity of the images and is defined

by (3.28)

$$d = x_{P_l} - x_{P_r} \tag{3.28}$$

A great deal of work has been performed on the subject of finding the disparity between two images and the problem itself is not a simple one with many extremely complicated algorithms producing excellent results as is evidenced by the stereo algorithms benchmarks at the Middlebury stereo vision page (Scharstein and Szeliski, 2014) having more than one hundred and fifty algorithms vying for top position. The simplest and fastest of these algorithms are window-based, which try to locate matching points based only on the intensity of the pixel in question and a surrounding rectangular window of reference points to the intensity of an equally sized sliding window in the second image as shown in Figure 3.10.

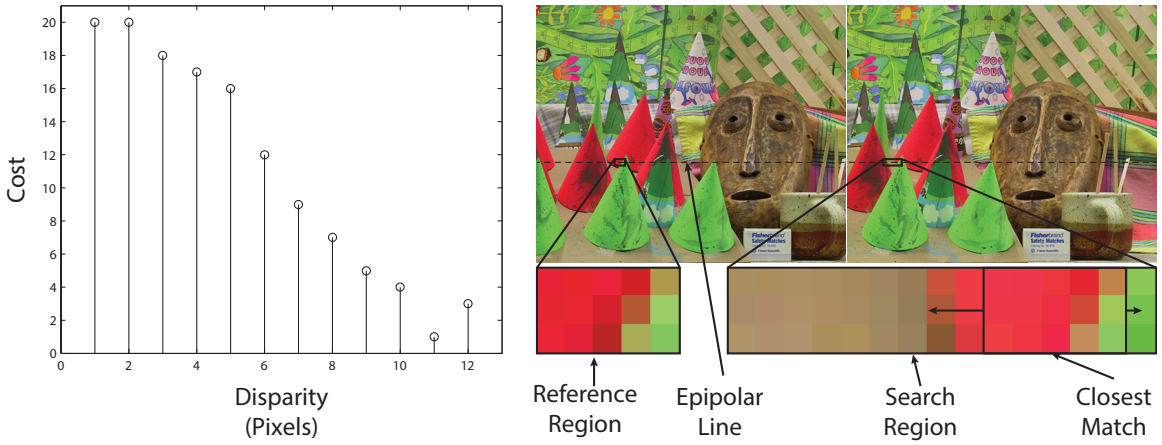

Figure 3.10: Local disparity calculation using a reference window in the left image and a sliding window in the right image (left image pair originally from (Scharstein and Szeliski, 2003))

These methods are termed local or block matching algorithms and compare regions by applying a cost function to obtain a single cost value for each pixel in the search region (Scharstein, 2002). There are many different cost functions that can be applied to the regions, however the simple local matching algorithm, referred to as Block Matching (BM) in the Open Computer Vision Library (discussed more completely in section 3.3 (OpenCV, 2014)), uses the sum of absolute differences (SAD) (3.29) between the pixels as a cost measure.

$$SAD(x, y, d) = \sum_{j=(y-w)}^{y+w} \sum_{i=(x-w)}^{x+w} |I_{ij} - I_{ij+d}| \tag{3.29}$$

In (3.29) $w = 0.5 \times$ width of window, (x, y) is the location of the point of interest in the left image and d is the disparity between the point of interest and the compared point in the right image. The SAD is a one of the fastest cost functions to compute because it uses simple subtraction rather than squaring the intensities as would be the case if a sum of squared

differences cost function were used. Once all the cost values have been computed along the epipolar line, the algorithm will find the minimum cost and corresponding disparity, then store it in a matrix comprising the x, y pixel coordinates and their corresponding d values.

The drawback of local matching algorithms is that they tend to produce less disparity matches (and hence much noisier disparity maps) than "global" optimisation methods such as graph cuts or belief propagation, however their speed is generally such that they are the only choice if real rime processing is required.

The second matching algorithm available in OpenCV is the Semi-Global Block Matching method (SGBM) developed by Hirschmuller (2005) which uses a hybrid approach to attain accuracy close to that of global methods while having a significantly lower processing time comparable to local methods. SGBM takes into account depth discontinuities which cause significant problems for other block matching algorithms by using the simple cost function (3.30) to directly compare the absolute difference of single pixel intensities.

$$C(x, y, d) = |I_l(x, y) - I_r(x - d, y)| \tag{3.30}$$

Equation 3.30 is then extended to include neighbourhood information as disparities are compared to other disparities in their neighbourhood (comparison is in the vertical and horizontal direction) to determine if the disparity is at a depth discontinuity and what the most likely disparity value is. If the pixel adapts its disparity value to its neighbours then it is given a lower cost value and is therefore more likely to be a match (Michael et al., 2013).

The difference in accuracy between BM and SGBM is illustrated in Figure 3.11(c) and Figure 3.11(d) where disparity maps are generated from the same image pair. Figure 3.11(a) is the left image of the pair and is provided by Scharstein and Szeliski (2003) and comes with a corresponding true disparity (Figure 3.11(b)) for comparison with the calculated disparities. It is clear from these figures that SGBM provides a far smoother disparity with much improved handling of disparities at depth discontinuities for only a small loss in terms of speed.

(a) The left hand image of the standard cones dataset from Scharstein and Szeliski (2003).

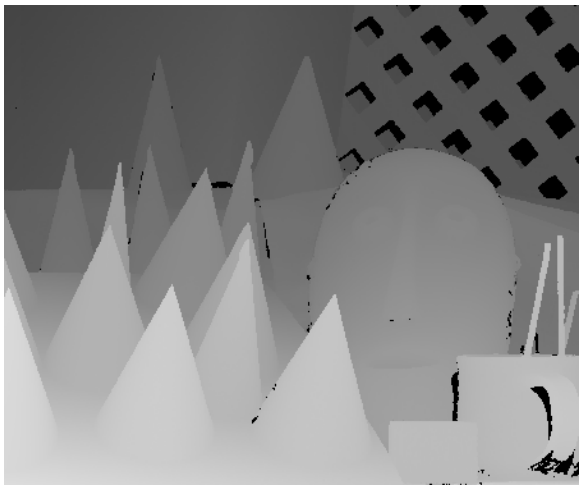

(b) The ground true disparity map provided with the Scharstein and Szeliski (2003) dataset.

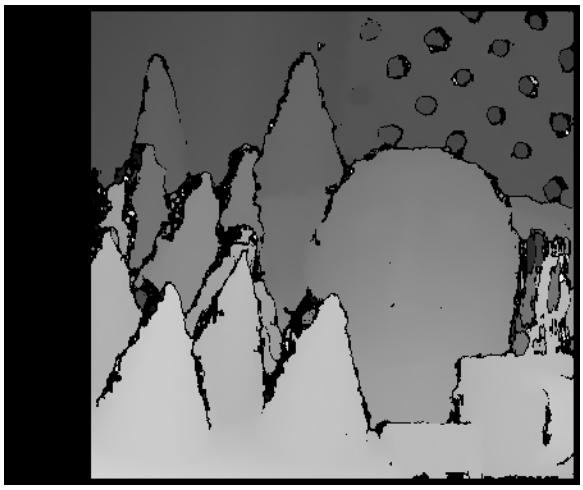

(c) The disparity map computed using the block matching algorithm.

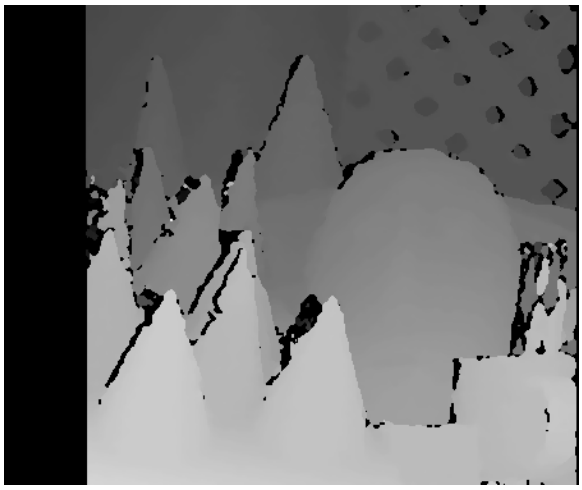

(d) The disparity map computed using the semi-global block matching algorithm.

Figure 3.11: Cones dataset

3.2.6 Triangulation to Determine 3D Points

Having generated a depth measure from the images, in the form of a disparity map, it is possible to recover the 3D coordinates of points in the world with respect to an origin at the optical centre of the left camera O_l. This operation is performed using the assumption that the images are in the same plane and the optical axes are parallel. This assumption is valid for rectified images and allows a system such as the one shown in Figure 3.12 to be constructed.

The figure shows a 3D point P projected onto the left and right image planes at points P_l and P_r respectively. Measured from the top left corner of each image the distance to the projected points is defined as x_{P_l}, y_{P_l} and x_{P_r}, y_{P_l} if the y direction is perpendicular to the page. The baseline b is defined as the distance between the optical centres O_l and O_r, the focal length is f and the distance from the baseline to the point P is Z_w. Using these values

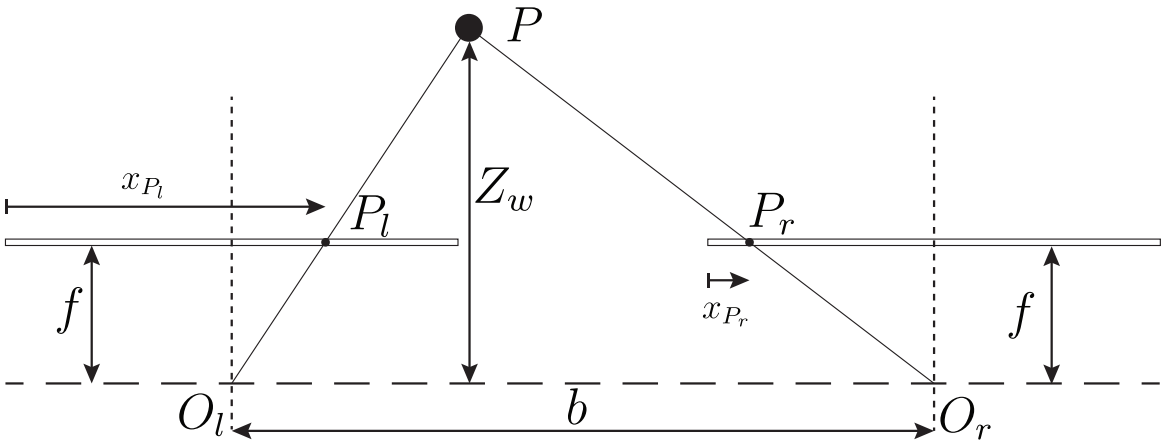

Figure 3.12: Triangulation can be used to determine the distance from the cameras' optical centres to a 3D point

and the fact that the dimensions can be related by similar triangles it is possible to develop (3.31) which links the z displacement in world coordinates (relative to the origin at O_l) to the disparity calculated earlier, the focal length and the baseline length.

$$Z_w = \frac{fb}{d} \tag{3.31}$$

Similarly it is possible to develop equations (3.32) and (3.33) for calculating x, y in world coordinates relative to the origin by using the calculated Z_w value.

$$X_w = \frac{x_{P_l} Z_w}{f} = \frac{x_{P_l} b}{d} \tag{3.32}$$

$$Y_w = \frac{y_{P_l} Z_w}{f} = \frac{y_{P_l} b}{d} \tag{3.33}$$

These equations state that the distance Z_w is inversely proportional to the disparity, so a larger disparity means the point being viewed is much closer to the camera. Another effect is that as Z_w increases, the difference in disparity between objects decreases amplifying errors in disparity calculation. This effect can manipulated to increase accuracy by changing the camera's relative positions which will in turn change the baseline value b.

3.3 SOFTWARE IMPLEMENTATION

All the software written as part of this study utilises an open source computer vision and machine learning library called OpenCV (OpenCV, 2014). OpenCV is written in C++ (although there are interfaces available for other languages such as Python and C#) and allows vision applications to be built quickly using over 2500 functions from all areas of computer vision (Bradski and Kaehler, 2008). Originally started at Intel, OpenCV was designed with computational efficiency in mind and is used heavily in real time industry applications.

OpenCV was chosen for this application over other computer vision packages (such as MATLAB's computer vision toolbox) because OpenCV is open source, free, very well supported by the computer vision community, and capable of extremely fast execution due to it's C++ base. Execution speed is an important parameter when designing computer vision software because due to the large volumes of data contained in each image processing time is often high. Even though the videos generated in this study do not require real-time processing (images processed at 30 frames per second (FPS) or greater) faster execution allows results to be analysed more quickly using multiple processing approaches.

The calibration, disparity calculation and triangulation methods described in section 3.2 are available in OpenCV and are implemented as functions that can be used freely in one's own C++ code. A thorough description of the use of these OpenCV functions for calibration and 3D point recovery is given below in subsection 3.3.1.

3.3.1 Calibration

The camera calibration process described in section 3.2 was implemented in software using the OpenCV functions. Figure 3.13 shows a flow diagram of the process.

Calibration is initiated by capturing a series of images (shown in Figure 3.14) containing calibration surface poses captured at the same instant by both cameras. Each pose is searched for a dot pattern with a specific shape and size, in this case a 4×11 asymmetric dot pattern was used, using the function `findCirclesGrid()`. If the pattern is found in both images simultaneously the position of the dot centres is recorded as a vector of 2D locations (one

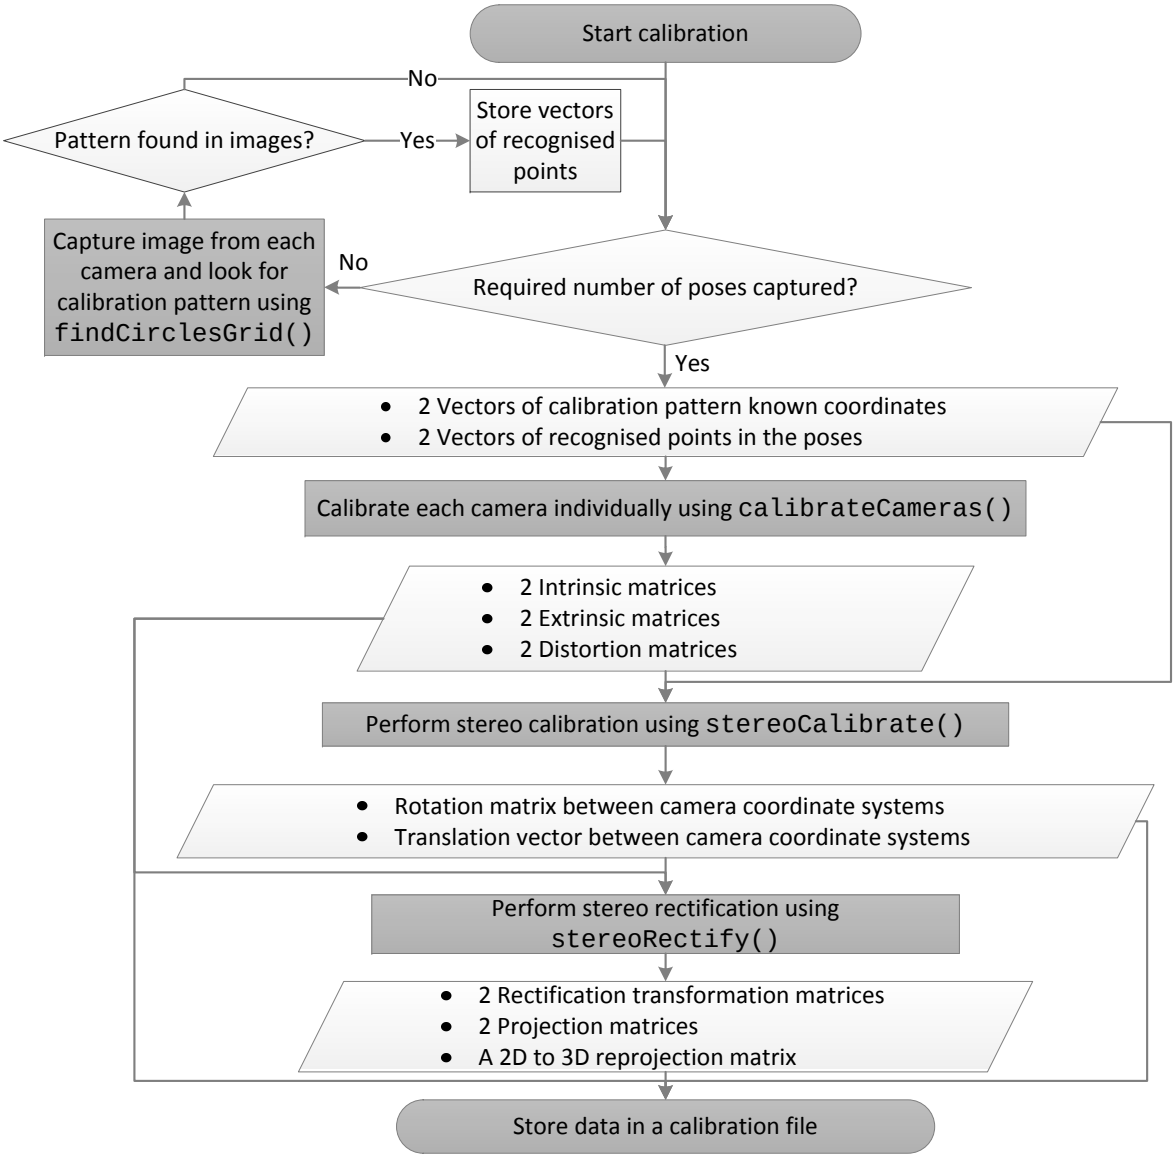

Figure 3.13: Flow Diagram of the calibration process

vector for each camera). This process is repeated a specified number of times adding to the point vectors accordingly. When the required number of poses has been reached the point vector along with a vector containing an idealised calibration surface geometry are passed to the `calibrateCamera()` function.

The `calibrateCamera()` function is run for each camera individually and implements the plane based calibration created by Zhang (2000) to calculate the intrinsic, extrinsic, and distortion matrices. These matrices are then passed to the `stereoCalibrate()` function which calculates the rotation matrix and the translation vector of the second camera relative to the first camera.

With the rotation matrix and translation vector between the cameras known it is possible to calculate the epipoles in the images and estimate the transformation matrices required to make the images lie on the same plane. It is also possible to generate a 2D to 3D reprojection matrix which will allow the 3D coordinates of a point to be calculated if the disparity and image coordinates are known.

The intrinsic, extrinsic, and distortion matrices as well as the rotation matrix, translation vector, rectification, and reprojection matrices are saved into a calibration file so that they can be used at a later time. A set of cameras need only be calibrated once unless their relative rotation and translation is modified, in which case a new calibration file is required.

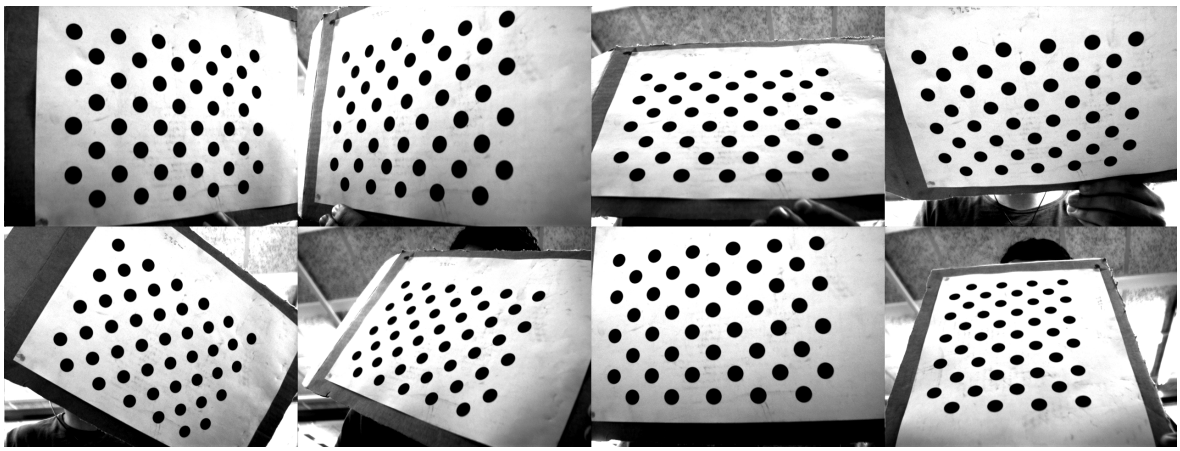

Figure 3.14: Calibration surface in eight different poses

3.3.2 Retrieving 3D Points from Images

Once the calibration parameters are available it is possible to retrieve 3D coordinates from any pair of images captured by the stereo cameras. However, first each set of raw images needs to be processed. The process which is applied to every pair of images is shown as a flow diagram in Figure 3.15.

The calibration file is initially read from the hard disk and the important parameters are loaded into memory. Intrinsic, extrinsic, distortion, and rectification matrices are passed to the `initUndistirtRectifyMap()` function which generates maps which relate old pixel locations to new pixel locations in new, rectified images. The maps in this case represent the transformation required to rectify and undistort the images. The `remap()` function is used to apply the maps to the distorted images and produce fully rectified, undistorted

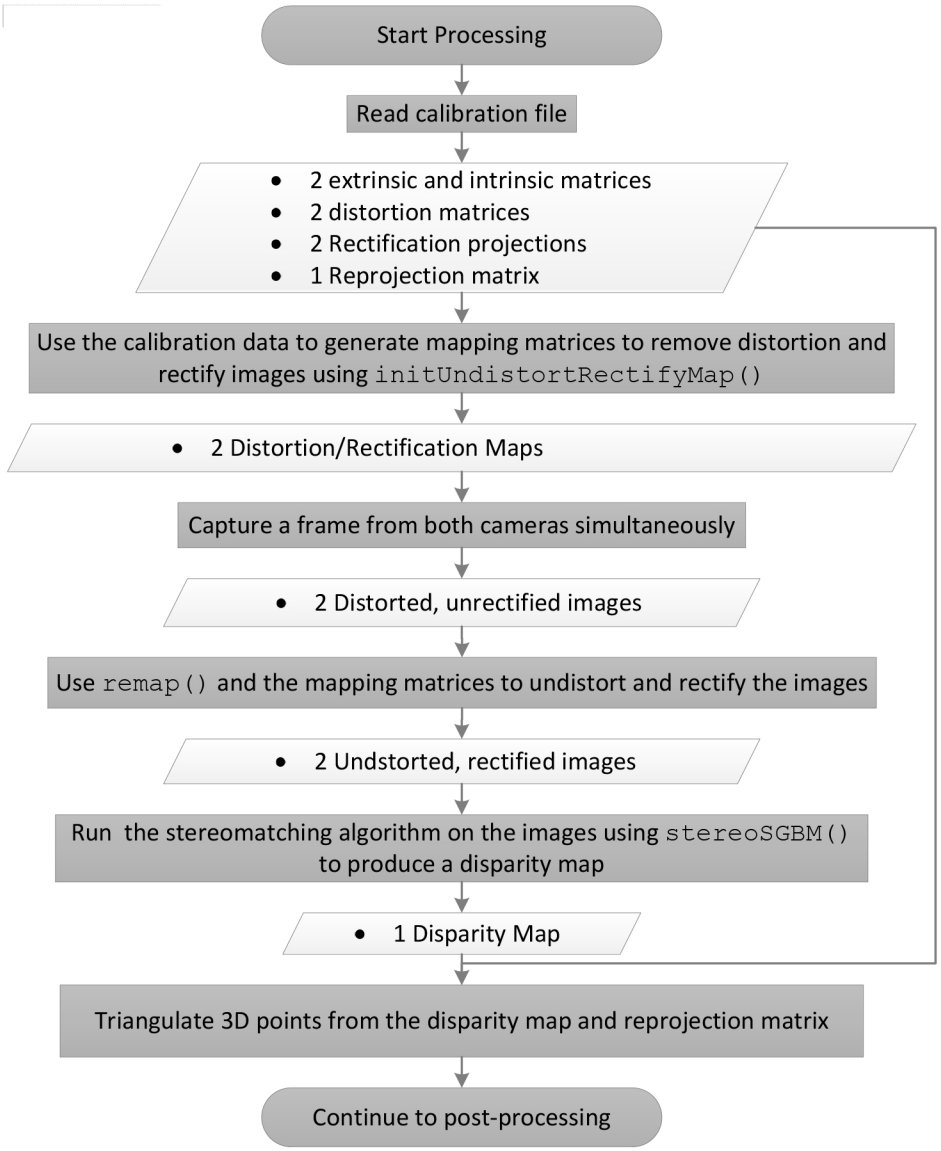

Figure 3.15: Flow Diagram of the stereo measurement process

output images. Maps are used in this case to reduce the processing time when applying the rectification and distortion transforms.

The rectified, undistorted images are passed to the `stereoSGBM()` function which produces a disparity map by implementing semi-global block matching algorithm developed by Hirschmuller (2005). Triangulation is then performed using the disparity map and the reprojection matrix recovered from the calibration file. The reprojection matrix is a matrix of equations (3.32), (3.33), and (3.31) so that each x, y, d point can be reprojected with a simple multiplication procedure. After reprojection the 3D data is made available for further post processing such as the rotation procedure described in section 4.5.

3.3.3 Video Capture Application

The details of the hardware used in the present study will be described in chapter 4. In short, images from two Point Grey Cameras (PointGrey, 2015) and captured using a Hyper-KBN single chip computer system (IEI, 2015). The software used to capture the images on the Hyper-KBN was custom written to be used with PointGrey cameras on a Linux platform. Designed to be a simple command line application, the program initially checks that both cameras are connected and working properly before prompting the user to enter a frame rate at which to capture images. The application waits for the user to start sampling by pressing the "s" key or "esc" to quit. While sampling, images from each camera are stored in a pair of memory buffers. When sampling is completed (by pressing the "q" key) the images are written to the hard drive as ".avi".

3.3.4 Stereovision Calibration and Tuning Application

Software was developed as part of this project that aimed to make the calibration and visualisation of stereo camera systems easier. The software was built for the Windows operating system using the OpenCV for computer vision and the Qt library (Digia, 2015) for user interface. The following features are available in this stereo vision application:

- Generates calibration data for stereo cameras using any asymmetric dot calibration surface and save the data to an ".xml" file.
- Rectifies and show images by applying calibration data (from file or newly generated) to input data.
- Allows input from video files or cameras connected to the computer via USB2 or USB3.
- Provides full control of video playback with options for pause, fast forward, rewind, step forward by a frame or step back by a frame.
- Displays disparity between the rectified images using one of two matching algorithms (BM or SGBM).
- Tunes stereo matching parameters to improve disparity quality.
- Applies a colour-map onto the disparity to allow clear visualisation of depth.

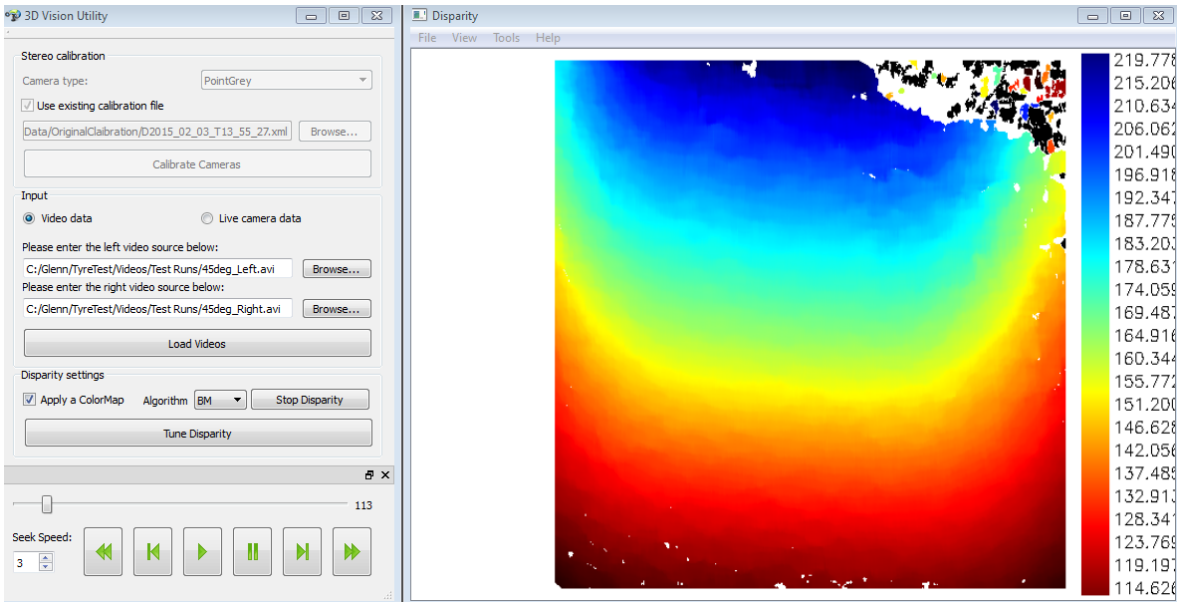

Figure 3.16: Screenshot of the calibration and tuning application with a disparity region being displayed

The application calibrates the cameras using the plane based calibration model discussed in subsection 3.3.1 and allows the cameras to be calibrated before being mounted in the tyre. After calibration the cameras are placed in the tyre, the lighting system is turned on and the matching parameters are tuned to produce the best disparity possible. Figure 3.16 shows the application after calibration has been completed playing back a test run and showing the disparity map with an appropriate colourmap.

3.3.5 Post Processing Application

The final piece of software developed is a post processing program to obtain useful results from the data captured by the stereo cameras. The program uses the calibration file generated by the calibration and tuning application to undistort and rectify the images captured in the video data. Disparity between the two images is determined using the SGBM algorithm. Using the disparity map and calibration data the 3D points corresponding to the pixels in the left image are calculated using triangulation.

The 3D points are displayed in a 3D display window using the OpenGL library (Khronos, 2015) and some code originally written by Botha (2014). This 3D window is scrollable, zoomable and rotatable much like the CAD interfaces that many of us are accustomed to.

Due to the fact that each 3D point corresponds to a pixel in the image captured by the left camera, the points can be coloured accordingly to produce a video effect as each pair of frames is processed and its 3D points plotted and coloured. Additionally, the colour of each of the points can be related to its vertical height creating a colour-map with blue at its lowest point and red at its highest point.

3.4 CONCLUSION

The software developed in this chapter implements the various computer vision techniques discussed in section 3.2 to capture images, calibrate cameras, and extract depth information from image pairs; additionally post processing software was developed to meaningfully visualise the 3D points generated from the depth information. This software forms one half of the stereo-vision tyre deformation measurement system. The other half of the measurement system is the hardware itself which is described in chapter 4.

CHAPTER 4

MEASUREMENT SYSTEM DESIGN AND IMPLEMENTATION

The development of the software required to determine 3D displacement from stereo images was described in chapter 3. This chapter describes the design and assembly of the camera stabilisation, data acquisition and power systems which allows these images to be captured. The stabilisation system has been designed to allow the cameras to be trained towards an area of interest within the tyre (most often the contact patch) and will stay in this orientation while the rim and tyre rotate about their primary axis. The positioning system is entirely mechanical using gear ratios to ensure the cameras view the desired area.

The cameras used to capture the images are a pair of PointGrey Flea3 (FL3-U3-13Y3M-C) high speed USB3 cameras (PointGrey, 2015) which were fitted with Kowa wide angle lenses (Kowa, 2015) to maximise the field of view. The cameras were connected (via USB3) to an IEI Hyper-KBN single chip computer (IEI, 2015) which captured the images using custom written software and was accessed via a WiFi connection. A lighting system was built using three 3W LEDs positioned behind the cameras to eliminated glare.

Power is provided (at twelve volts) to the system through two pairs of slip rings, one pair to get the power from the stationary hub and another to get from the rotating rim to the stationary ring inside the tyre.

4.1 STABILISING SYSTEM DESIGN

4.1.1 Stabilising Principle

The stabilisation system was designed to allow the cameras to be oriented to view any area of interest inside the tyre and locked to view that area until the setting is changed. The system is entirely mechanical relying on a planetary gear system to keep the camera orientation fixed.

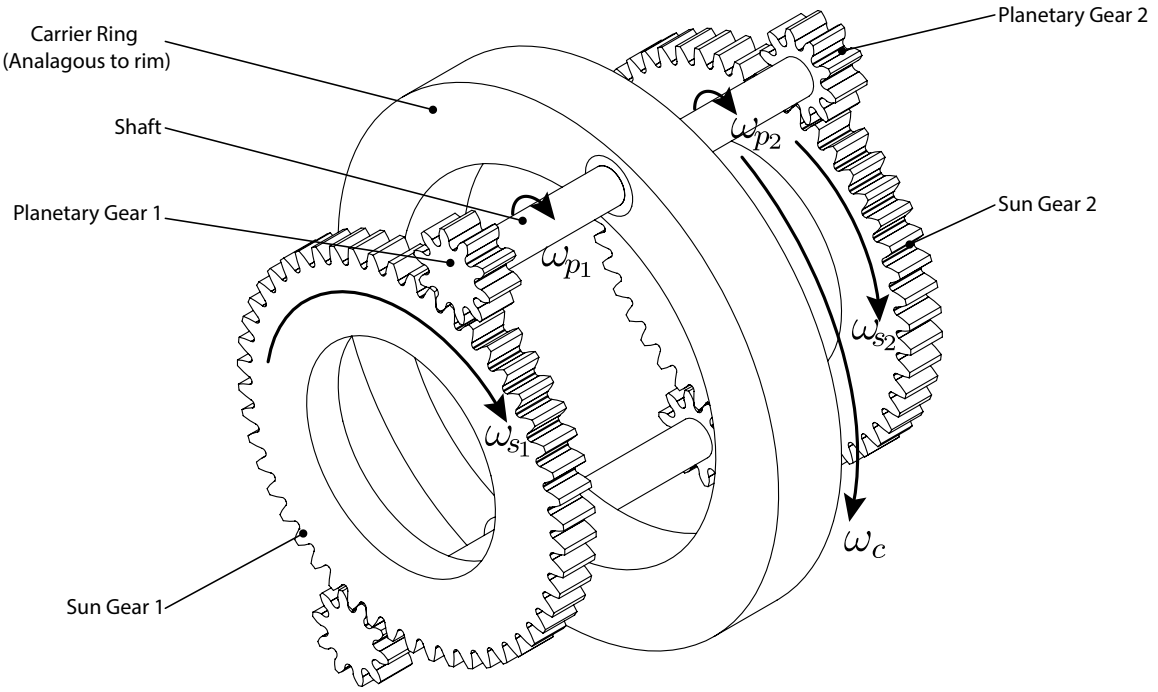

Figure 4.1: Simplified diagram of the mechanism principle

Figure 4.1 shows schematic diagram of the planetary gear drive-train operating without an annulus gear (a large internally toothed gear which would normally interface with the planetary gears). The carrier ring, which constrains the movement of the planetary gears, and the sun gears rotate about the same central axis while the planetary gear on either side of the carrier ring is coupled together by a shaft ensuring that they rotate at the same angular velocity. This planetary gear system is described by (4.1) (Freudenstein and Yang, 1972)

$$\mathrm{N}_s\omega_s + \mathrm{N}_p\omega_p - (\mathrm{N}_s + \mathrm{N}_p)\omega_c = 0 \tag{4.1}$$

where N_s and N_p are the number of teeth on the sun and planetary gear respectively; while ω_s, ω_p and ω_c are the angular velocity of the sun gear, planetary gear and carrier ring

respectively. Since the drive-train in Figure 4.1 effectively has two sun and planetary systems coupled together (4.1) can be modified into (4.2).

$$N_{s_1}\omega_{s_1} + N_{p_1}\omega_{p_1} - (N_{s_1} + N_{p_2})\omega_c = N_{s_2}\omega_{s_2} + N_{p_2}\omega_{p_2} - (N_{s_2} + N_{p_2})\omega_c \tag{4.2}$$

The drive train in this example uses identical sun and planetary gears on either side of the carrier ring therefore $N_{s_1} = N_{s_2}$ and $N_{p_1} = N_{p_2}$; additionally since the planetary gears are coupled together $\omega_{p_1} = \omega_{p_2}$. Substituting these relations into (4.2) gives (4.3)

$$N_{s_1}\omega_{s_1} + N_{p_1}\omega_{p_1} - (N_{s_1} + N_{p_2})\omega_c = N_{s_1}\omega_{s_2} + N_{p_1}\omega_{p_1} - (N_{s_1} + N_{p_1})\omega_c \tag{4.3}$$

which after simplification yields (4.4)

$$\omega_{s_1} = \omega_{s_2} \tag{4.4}$$

Equation 4.4 proves that in this drive-train the two sun gears will always have the same rotational velocity irrespective of the carrier ring velocity and therefore if one of the sun gears is held stationary the other will also be stationary.

4.1.2 Mechanism Design

Using the principle developed in subsection 4.1.1 a stabilising mechanism system has been designed to fit a 16 inch off-road vehicle rim (shown in Figure 4.2) while having the space to accommodate a separately designed six component Wheel Force Transducer (WFT) (as described in Els (2012)) and corresponding telemetry system. The WFT will be used for further research with the system to correlate wheel force to tyre deformation however in this it serves purely structural purposes. The mounted WFT and telemetry system can be seen in Figure 4.3

The stabilising mechanism consists of a pair of identical rings which have an inner diameter large enough to accommodate the rim in their centre. Each ring houses a length of 5mm pitch HTD belt which transforms the rings into internally toothed gears that perform the function of the sun gears in (4.4).

Two sets of identical planetary high torque drive (HTD) pulleys (Figure 4.4(a)), coupled together with shafts, interface with the belted inner surface of the rings and orbit the rotational axis of the rim with the rim assuming the role of the carrier ring in Figure 4.1. The shafts

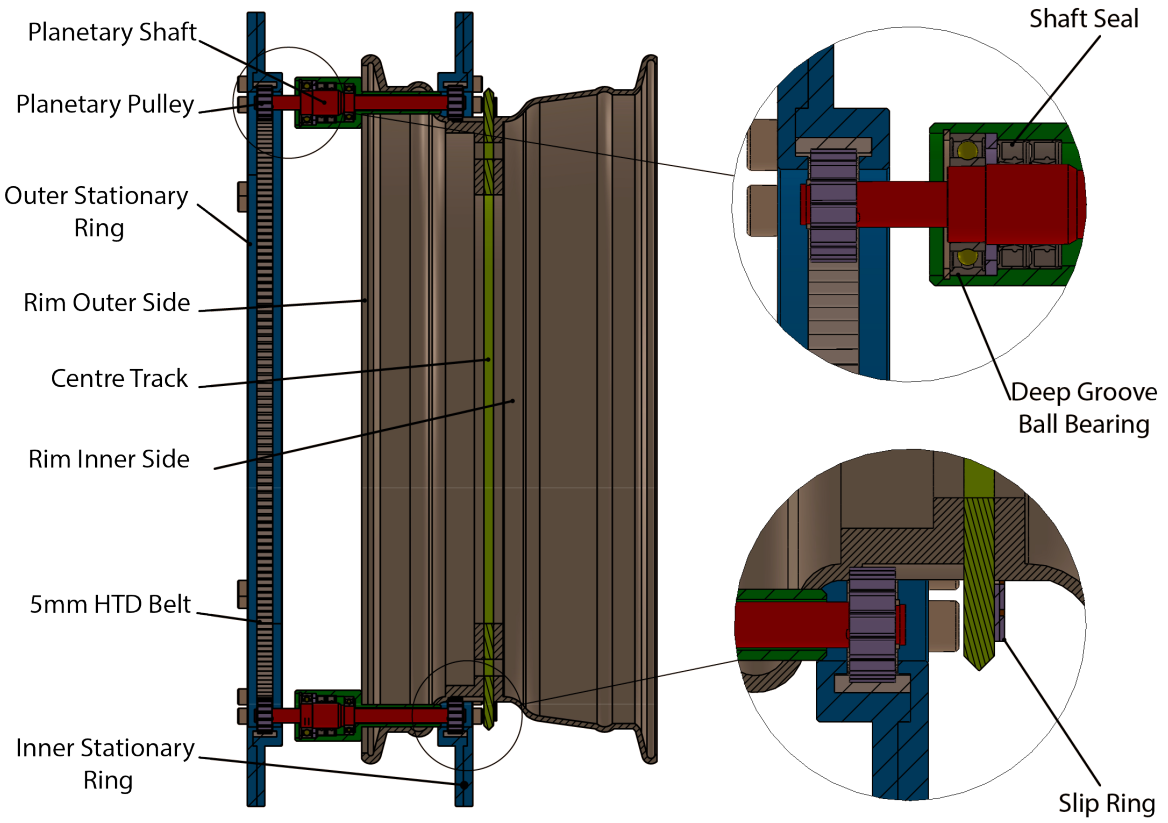

Figure 4.2: The stabilisation mechanism with detailed views

are located by bearings inside cylindrical housings which penetrate the rim and are sealed by high pressure shaft seals.

The rim splits in the centre for ease of assembly and clamped between either side is a chamfered track. The interface area between the two sides of the rim and the track is sealed using O-ring cord. White paint was applied to the rim to increase light diffusion and improve lighting conditions at the area being measured.

Three caster wheels with a V-profile (Figure 4.4(b)) are spaced at 120° intervals and interface with the track to ensure that the rim rotates concentrically relative to the inner stationary ring. The caster wheels are manufactured from plastic to reduce vibration in the system and contain high-quality deep groove ball bearings rotating on shoulder bolts.

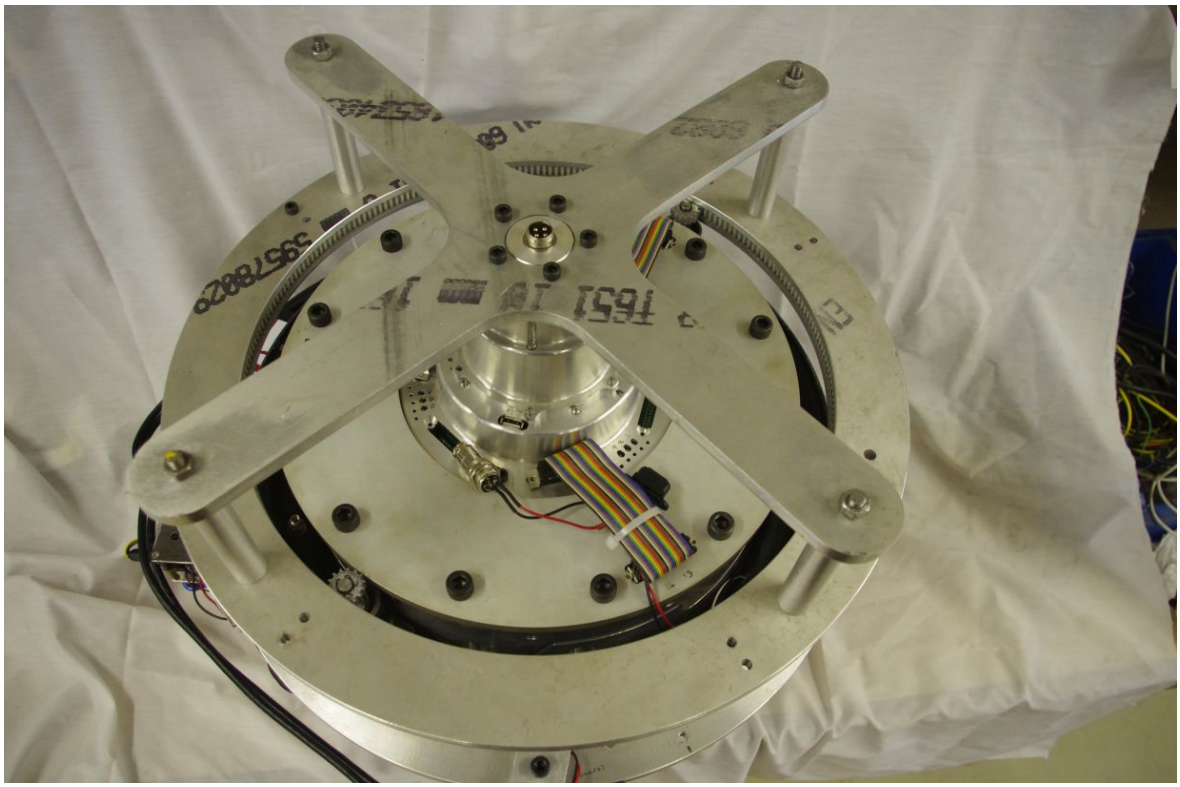

Figure 4.3: Wheel force transducer telemetry system with power input and mounting system for the outer stationary ring.

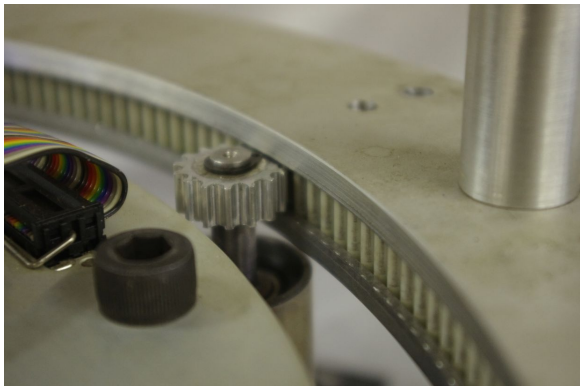

(a) Cameras mounted on their bracket with LEDs positioned behind the lenses to prevent glare

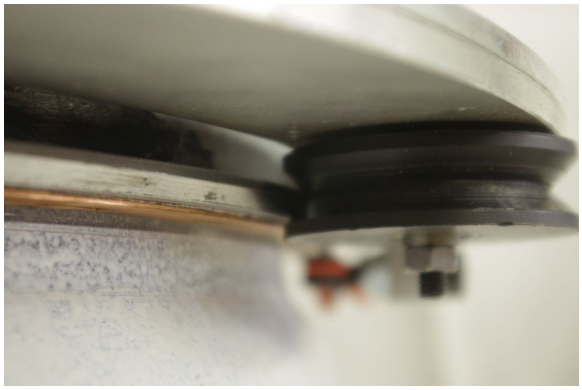

(b) Caster wheel interfacing with the central track

Figure 4.4: Component Detail Views

The inside of the tyre is sprayed in a speckle pattern to increase the optical features of the surface and thereby improve the quality of the stereo matching. Lines were drawn on the outer sidewall at sixty degree increments and numbered from $0-5$. Corresponding lines were drawn on the inner sidewall (where they would pass within view of the cameras) and numbered identically. These lines were drawn to so that the camera's orientation inside the tyre could be determined by comparing the location of the outside lines to lines in the

camera's field of view.

4.2 POWER SYSTEM DESIGN

The data acquisition system (DAQ) inside the tyre been designed to be powered by 13.8V provided by the vehicle alternator. The power is introduced into the rotating rim via slip rings in the wheel force transducer telemetry system (Figure 4.3) and is transferred through the rim wall and onto the slip rings. A set of brushes (mounted on the stationary inner ring) slides along the slip rings to provide power to the rest of the system. The slip rings and brushes are shown in Figure 4.2 and Figure 4.5

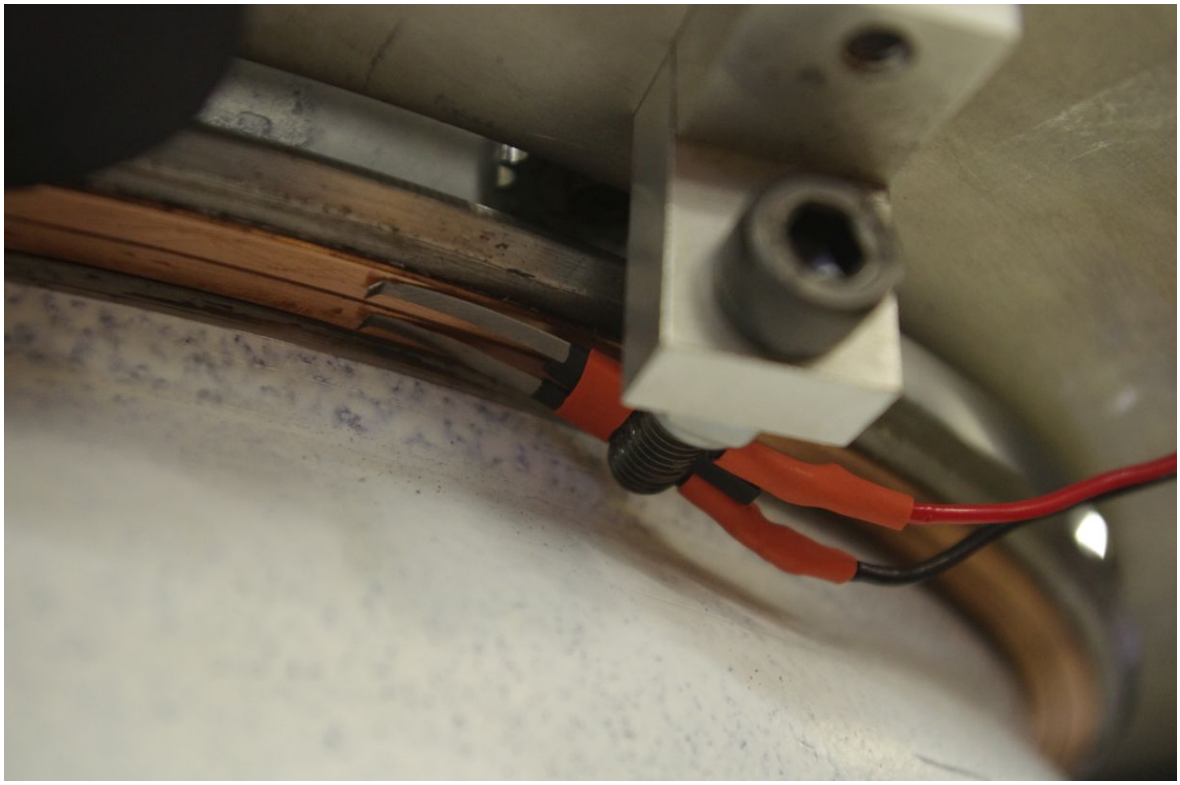

Figure 4.5: Brushes mounted on the stationary ring.

Input voltage fluctuates if the brushes lose contact with the slip rings due to vibration or slight surface defects therefore a 12V battery is integrated into the system (shown in Figure 4.6(a)) to smooth out any fluctuations. Having a battery in the circuit solves the fluctuation problem however it means that the system cannot be shut down or rebooted by cycling the power since the battery will keep the system running until it is completely discharged. To prevent this a delay circuit (shown in Figure 4.6(b)) is included which prevents the DAQ from shutting down by short losses of power, while still allowing the system to be powered down completely

by simply disconnecting the input voltage source for a longer than about 3 seconds.

Figure 4.7 shows the layout of the delay circuit. The input voltage (from the brushes) of 13.8V charges the battery through a diode. This diode prevents the battery discharging into the vehicle when the engine is off. The input voltage additionally charges a 2200µF capacitor and provides a gate voltage to the N-type MOSFET. When the input voltage drops due to a slip ring problem the capacitor discharges slowly through the 1MΩ resistor and the gate voltage across the MOSFET drops proportionally to the discharge rate of the capacitor. Even though the input voltage in off the gate voltage across the MOSFET is still high for a while and the source-drain channel stays open, keeping the DAQ circuit connected to ground and therefore powered on for a few seconds supplied by the battery. When the gate voltage drops too low the source-drain channel closes and the DAQ powers down.

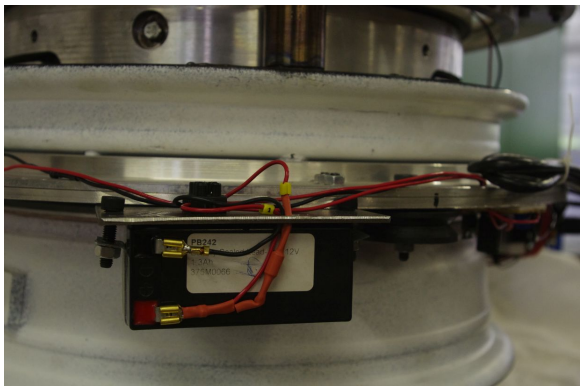

(a) The battery used to prevent power interruptions.

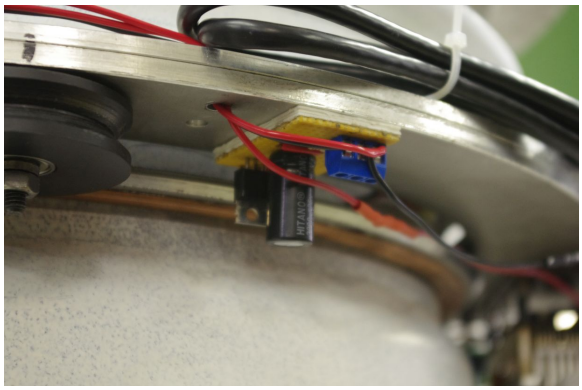

(b) Delay switch to administrate the battery system in case of power loss.

Figure 4.6: Electrical system detail views

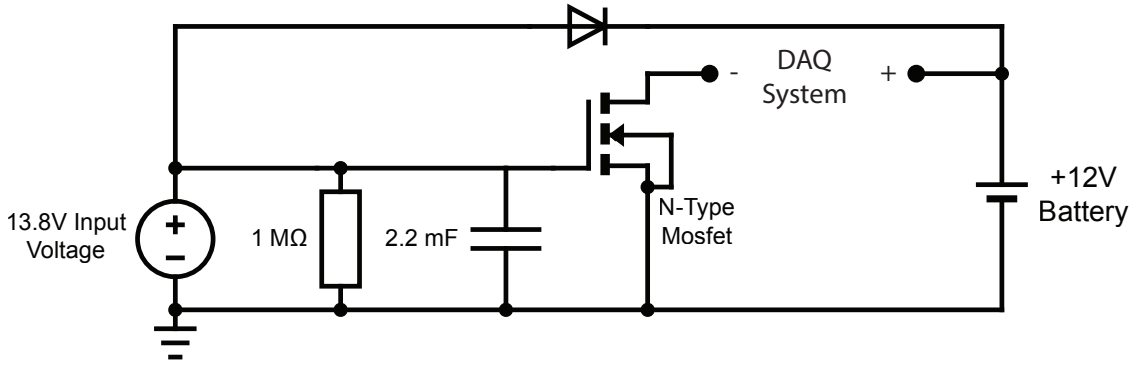

Figure 4.7: Switching circuit to prevent power interruptions

4.3 DATA ACQUISITION SYSTEM DESIGN

The DAQ consists of a IEI Hyper-KBN single chip computer (specifications given in Table 4.1), 2 PointGrey Flea3 (FL3-U3-13Y3M-C) high speed USB3 cameras and a 9W LED lighting system.

Table 4.1: Hyper-KBN hardware specifications

Specification	**Specification Value**
CPU	AMD GX-415GA (1.5GHz, quad-core)
RAM	2GB DDR3 1333MHz
I/O Interface	1 × RS-232 1 × SATA 6Gb/s 2 × USB 3.0 (on rear I/O) 2 × USB 2.0 (by pin header)
Power Supply	12V input only 1 × DC power jack (5.5mm)
Power Usage	17.4W when 2GB 1333MHz DDR3 is installed

The Hyper-KBN was chosen specifically for this task because of its compact size, powerful processor and inclusion of two USB3 ports. This computer is capable of running any x86 operating system however for this project Ubuntu Linux 14.04 was installed due to its ease of installation and customisation. Ubuntu is also actively supported by PointGrey (the camera manufacturer) and therefore drivers are readily available.

The computer is mounted to the inner stationary ring (Figure 4.8(a)) and set up to act as a WiFi access point using one of the USB2 ports and an off the shelf WiFi dongle. Using WiFi and Linux secure shell (ssh) the Hyper-KBN can be connected to and controlled using normal Linux terminal commands.

The PointGrey cameras (whose specs are given in Table 4.2) were mounted alongside one another on a specially designed bracket as shown in Figure 4.8(b). The cameras were chosen

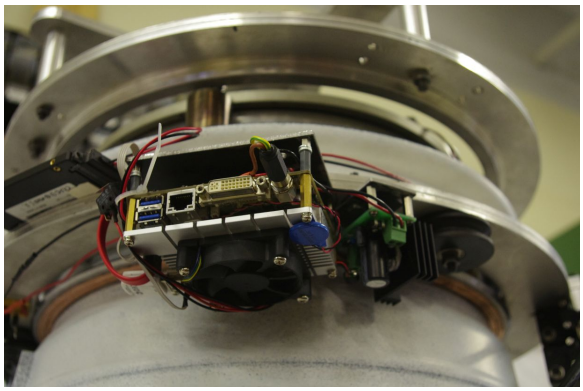

(a) The Hyper-KBN mounted to the stationary ring with the green LED voltage regulator in the background

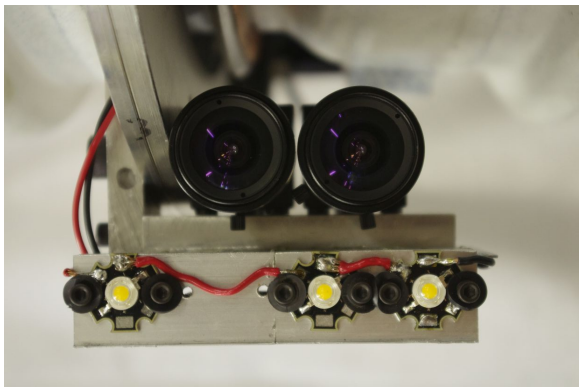

(b) Cameras mounted on their bracket with LEDs positioned behind the lenses to prevent glare

Figure 4.8: DAQ hardware components

because they are small, have a high frame rate and global shutter. Power for the computer and cameras was provided from the same 12 V source and data was exchanged via USB3 cables.

Table 4.2: PointGrey FL3-U3-13Y3M-C hardware specifications

Specification	**Specification Value**
Max Resolution	1280 × 1024
Max Frame Rate	150 FPS
Megapixels	1.3
Sensor Type	CMOS
Sensor Format	1/2"
Shutter Type	Global
Lens Type	Kowa LM4NCL
Lens Angle of View	103.6° vertical, 76.6° horizontal

As can be seen in Figure 4.8(b) the lighting system consisted of three 3 W LEDs connected in series and mounted on a heat sink below the cameras. These LEDs are powered by a DC-DC (the green circuit in Figure 4.8(a)) linear voltage regulator to lower the input voltage to the required level of 10.8V and provide directional illumination of the area being captured while providing no glare due to their placement below and slightly behind the lenses.

4.4 COMPLETE STABILISATION SYSTEM

The complete stabilisation system, showing the location of all the components discussed in earlier sections, is given in Figure 4.9. The computer and battery are positioned at the top of the inner stationary ring while the cameras are placed at the bottom of the ring to capture the contact patch. Located slightly below and behind the stereo cameras, the LED lighting system illuminates the contact patch. It should be noted that the cameras are mounted off-centre on the stationary ring and are set at an angle rather than vertically; this was done in order to increase the field of view of the cameras and consequently the size of the measured area. However, this did introduce rotation in the processed data as discussed in section 4.5.

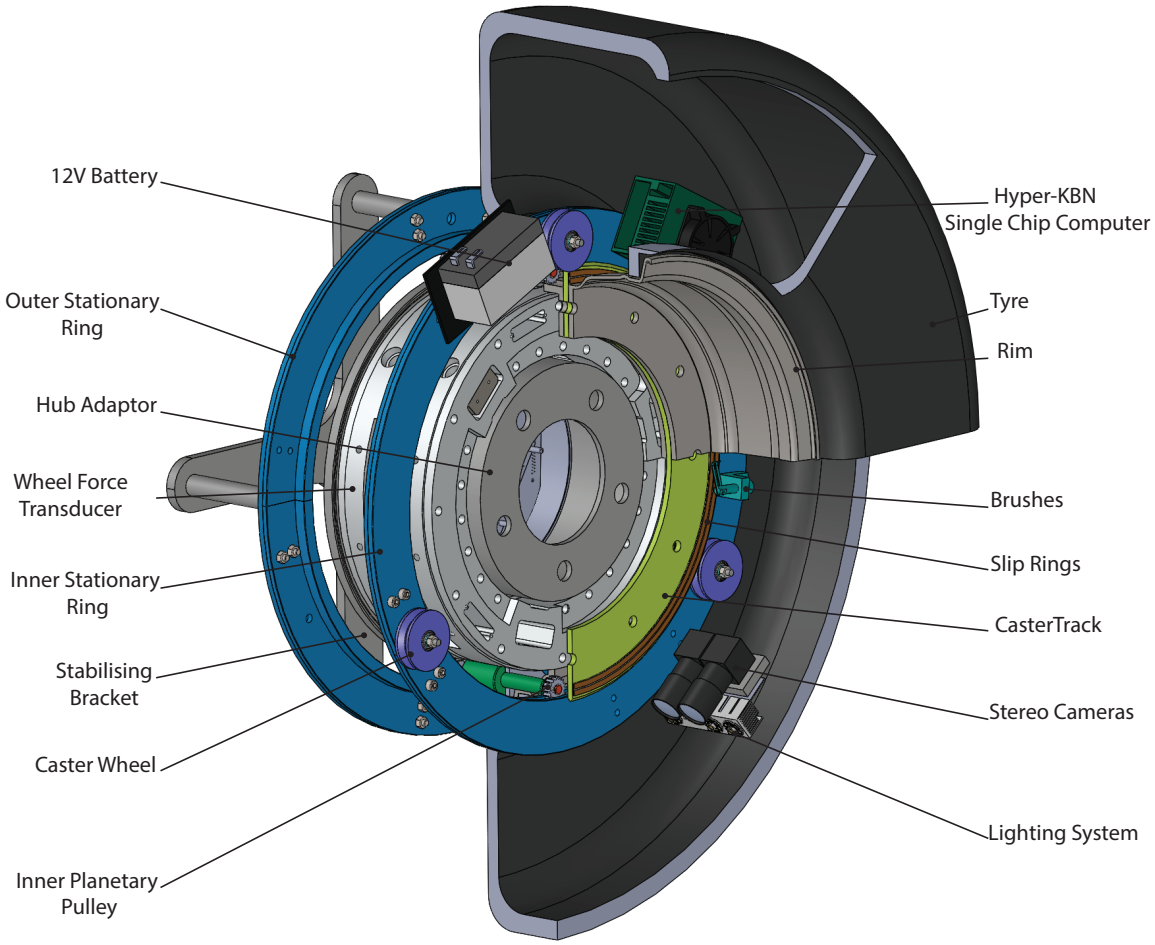

Figure 4.9: The complete stabilisation system showing important components

Twelve bolt holes are located around the rim perimeter to hold the WFT to the rim and caster track. A hub adaptor (visible in the centre of the rim) is used to connect the WFT to the wheel hub. The slip ring with its brushes (mounted on the inner stationary ring) are visible

above the cameras and the planetary pulley interfaces with the inner ring's belt to the left of the cameras. Three caster wheels are located at 120° intervals from one another and run along the chamfered caster track. The stabilising bracket locates the outer stationary ring concentrically with the rim and rotates on a bearing in the WFT telemetry system.

4.5 ROTATION OF OUTPUT DATA

The 3D points are triangulated from the data contained in the disparity maps and therefore all 3D points are relative to the left optical centre O_l (as described in subsection 3.2.3). This is problematic because the cameras are positioned at an angle relative to the flat, central region of the contact patch as shown in Figure 4.9. Figure 4.10(a) shows how the 3D points

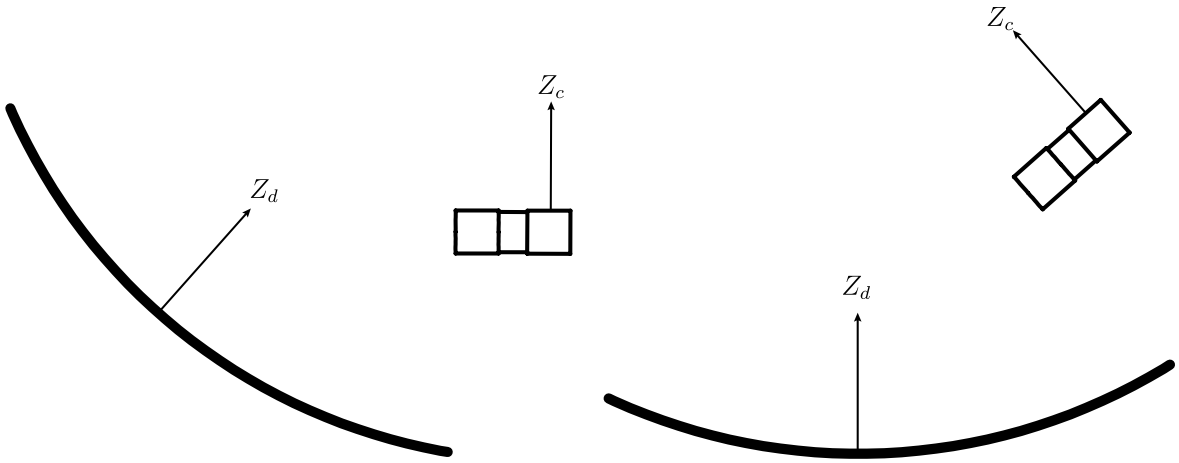

(a) The data is captured relative to the camera coordinate system therefore the z-axis of the data will be Z_c

(b) The data is rotated so that an axis placed perpendicular to the centre of the contact patch (Z_d) is the new z-axis

Figure 4.10: Rotation of 3D data

reflect this angular orientation by having the data oriented according to Z_c (the z axis of the camera) rather than Z_d (the z axis relative to the road plane). This problem can be overcome by using the component rotation matrices of (4.5) which describe rotation about the x, y and z axes.

$$R_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{bmatrix} R_y(\theta) = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix} R_z(\theta) = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (4.5)$$

When these matrices are multiplied together in (4.6) they form the rotation matrix R_{data} which is used to rotated the data about the x axis by 157° until the Z_d (perpendicular to the

ground plane) is parallel to the z axis of the world coordinate system. The transformation from one coordinate system to another is shown in Figure 4.10(b).

$$R_{\text{data}} = R_z(\theta_z)\, R_y(\theta_y)\, R_x(\theta_x) \tag{4.6}$$

4.6 CONCLUSION

The hardware developed in this chapter solves the mechanical problems inherent in measuring deformation inside a tyre. Namely, constraining rotation to constantly ensure a constant view of the deformed area, powering DAQ & lighting systems, and remotely accessing the data captured on the DAQ. These problems were solved by designing and manufacturing a DAQ with wireless capabilities, a pair of slip rings to supply power to the DAQ & lights and a stabilisation system based on planetary gearing principles. This hardware (along with the software developed in chapter 3) completes the stereo-vision tyre deformation measurement system which is now ready for testing.

CHAPTER 5

TESTING AND RESULTS

Testing was performed on the system described in section 3.2 and chapter 4 to determine both it's accuracy and performance in measuring full field contact patch deformations. The testing methodology, obstacles used as well as the results are shown in the sections below.

The accuracy of using stereo cameras for measurement was tested by measuring the geometry of a known surface and comparing the measured coordinates to the known ones as discussed in section 5.1. Full field deformation measurements were captured in section 5.3 as the tyre negotiated a number of different obstacles. The size as well as the layout of the obstacles was changed for each test (subsection 5.2.1). A test to measure the relative accuracy and sensitivity of the system was performed in subsection 5.3.7 by driving over obstacles which decreased in size to discover what level of deformation was measurable.

The results of all tests are discussed and compared to one another to gain a full understanding of the strengths and weaknesses of the stereo camera measurement system.

5.1 ACCURACY OF THE CAMERA SYSTEM

The software that was developed in section 3.3 needed to be tested to ensure that measurements produced by the system were accurate. Accuracy was tested by measuring an inclined surface with known geometry and comparing the measurements to an idealised "true" surface.

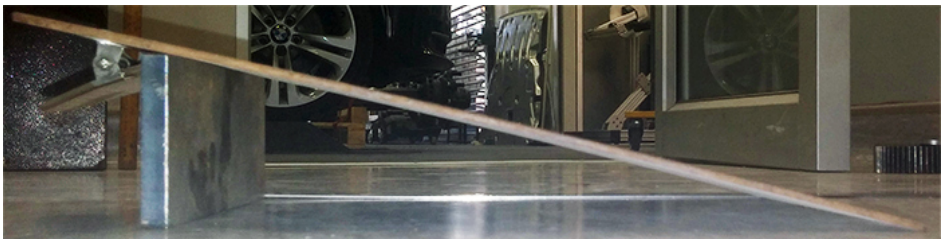

(a) Side view of the measured surface

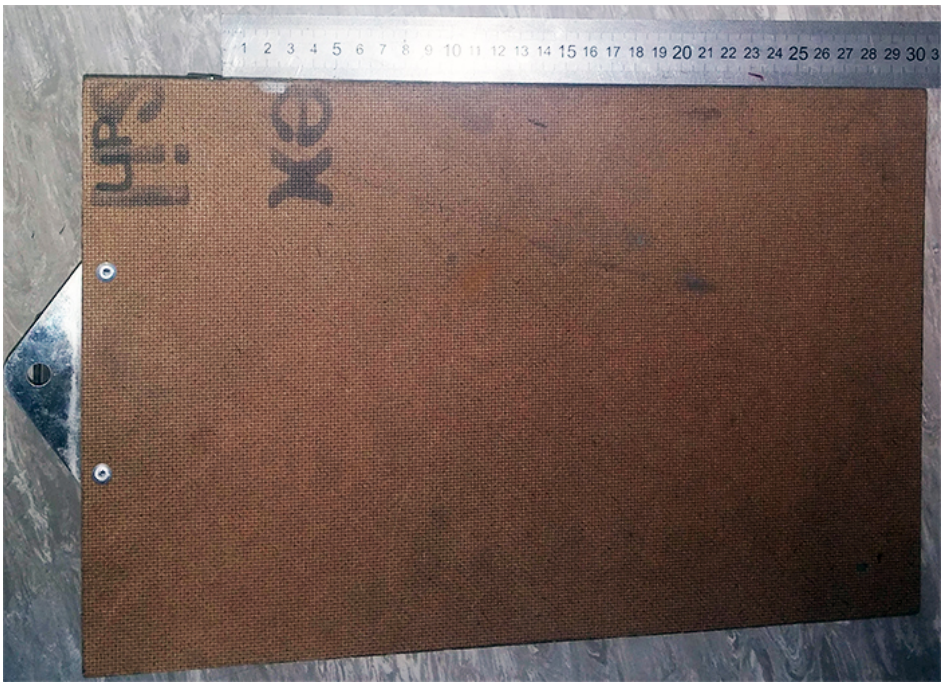

(b) Top view of the measured surface

Figure 5.1: Measured Surface orientation

An inclined surface (shown in Figure 5.1(a) and Figure 5.1(b)) was captured at a resolution of 640×480 using the same PointGrey Flea3 cameras and Kowa lenses that were used in the rest of the system. The inclined surface was chosen as a reference because it has a gradual increasing vertical height value. This is a more accurate analogue for the data captured in this study than a surface with discontinuous height changes (height discontinuities introduce errors when one camera's view is blocked by an edge).

The video of the inclined surface was processed using the SGBM algorithm described in subsection 3.2.5 and the measured data was interpolated onto a regular grid with 1 mm resolution. This interpolation was performed so that the true geometry and the measured geometry could be directly compared at each point. The measured and true surfaces are plotted in Figure 5.2 (with equal axes) and clearly have very similar gradients and dimensions. The measured surface is missing some points where the data is inaccurate; therefore, these points were not included in the error calculation.

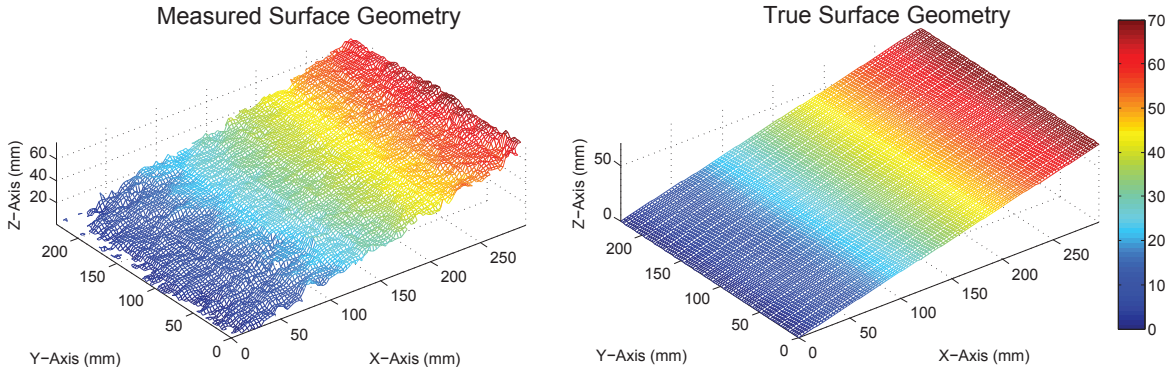

Figure 5.2: Measured geometry (left) and true geometry (right) of the surface

Figure 5.3 shows the mean, measured vertical height and the mean, true vertical height along the y-axis. The absolute mean error between the two surfaces is plotted at the bottom of the figure and shows a constant trend of approximately 2mm across the entire surface.

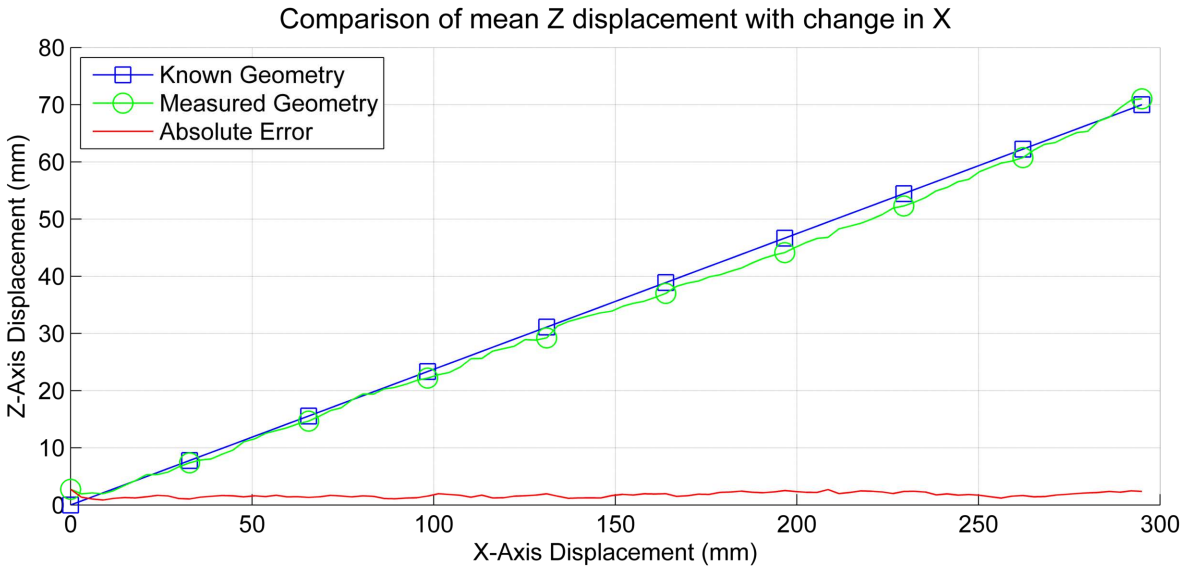

Figure 5.3: Cross section showing mean surface values at every millimetre and absolute error between these values

The error in the measured data could be reduced by increasing the resolution of the cap-

tured images. Higher resolution would allow more feature points to be compared, therefore increasing accuracy but would also increase processing time for marginal gains. As tyre displacements are expected to be large, this error was deemed acceptable for the current purpose.

5.2 TEST SETUP

The measurement system was bolted to the rear right hub of a Land Rover Defender 110 off-road vehicle. A cross shaped stabilising bracket (which holds the outer stationary ring concentric with the inner ring) was bolted to a torque arm (shown in Figure 5.4) to prevent the ring from rotating with the wheel. The setup was checked to ensure that the stabilising rings could rotate unimpeded, the tyre was pressurised to 2 bar and power was supplied to the system. Once the DAQ had successfully booted and WiFi was available, system calibration was performed as discussed in subsection 5.2.2.

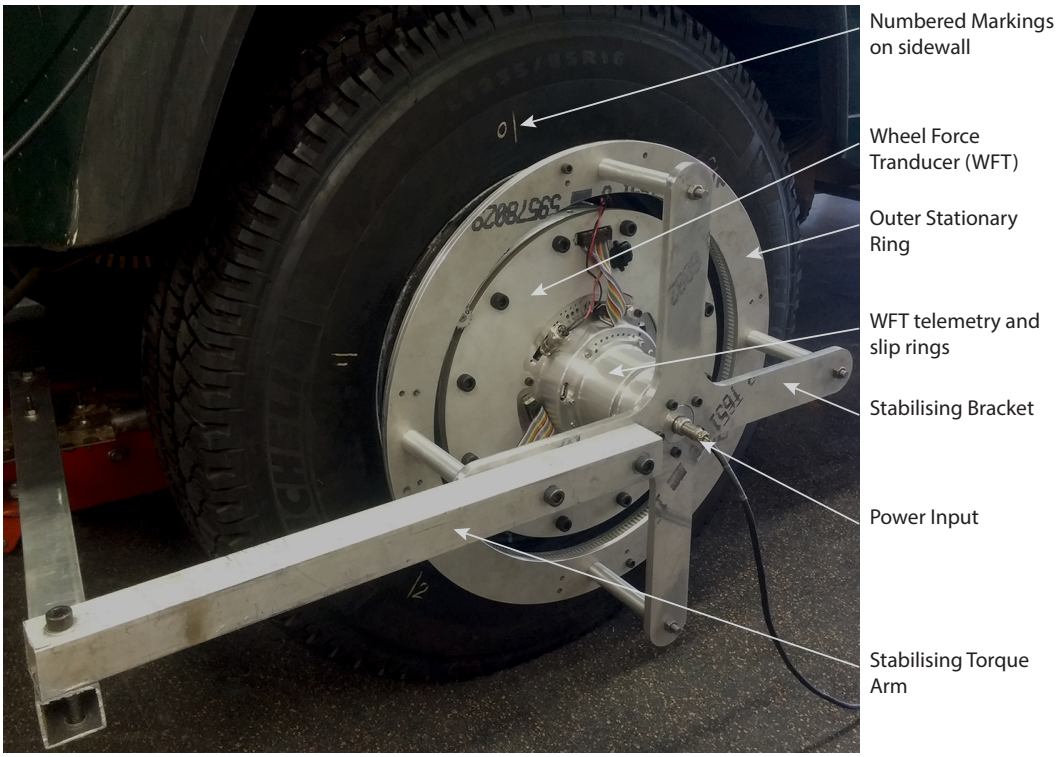

Figure 5.4: The mechanism mounted on the rear right wheel hub showing the position of all important components

5.2.1 Obstacles

The obstacles used in the dynamic tests were cleats with various cross-sections, chosen to approximate the terrain which an off-road vehicle may have to negotiate and to provide different amounts of tyre deformation to test the accuracy of the measurement system. The cleats chosen were a 38 mm square section, a 50 mm square section, and a 40 mm angle iron lying with the hypotenuse of the section coincident with the floor. These cross-sections of the cleats are shown (with dimensions) in Figure 5.5

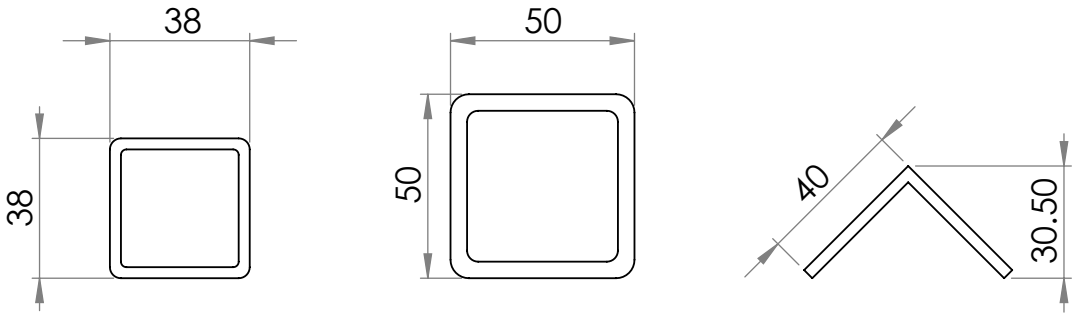

Figure 5.5: Dimensioned cross-sections of the cleats used in the dynamics tests

A second set of obstacles consisting of nuts which decrease in size (shown in Figure 5.6) was used for the final test run to determine the smallest obstacle which would create a measurable deformation to the tyre. The nuts were M24, M20, M14 and M10 with a nyloc torque ring.

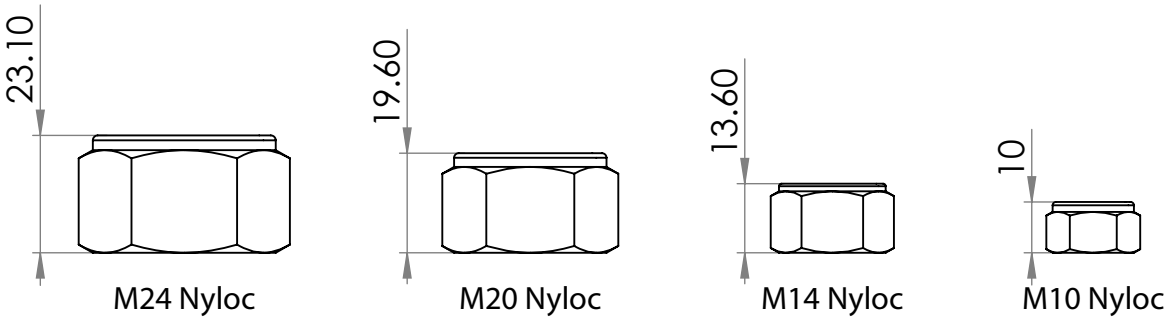

Figure 5.6: Dimensioned drawing of the nuts used to test sensitivity

5.2.2 System Calibration

A number of runs were performed to align the cameras so that they viewed the contact patch. This was done by rotating the tyre until the line numbered zero on the sidewall was perpendicular with the ground beneath the tyre and then rotating the cameras around the

tyre. The videos were examined and when the zero line on the inside of the tyre was in the middle of the camera's field of view the torque arm was bolted in place.

5.2.3 Test Procedure

Dynamic tests were performed under lab conditions, on a flat surface. The same starting position was used for every run and the vehicle was driven over the obstacles at low speed (less than 10 km/h). The wheel containing the measurement system navigated the obstacles while all three other wheels remained on the ground leading the vehicle to pitch as weight was transferred.

5.3 RESULTS

The results of each of the dynamic tests are given in the sections below using identical post processing settings. For each of the tests, the first figure shows the full deformed profile as the tyre rolls over the obstacles and the second figure is a cross-section of the data captured at a plane cutting the data laterally at $x = 150$ mm or a longitudinal plane cutting the data at $y = 0$ mm depending on the test performed. All data has been rotated to account for camera positioning, with a constant angle being assumed. However, in a few cases it was clear during post-processing that the angle of the flat contact patch area with respect to the xy-plane was not zero indicating that the original rotation angle was incorrect.

It was observed in videos of the tests that the torque arm holding the system stationary was not rigid enough and tended to deflect by a few degrees when the vehicle acceleration changed. This rotation means that although all the cross-sections are taken at the same longitudinal position the data might be over or under rotated relative to the ground plane. This means that the different tests cannot be meaningfully compared as the cross-sections will not lie on the same vertical plane and can easily add or subtract a few millimetres from the true deformation.

It should also be noted that the vehicle pitched as it negotiated obstacles and since the torque arm was attached to the vehicle the angle of rotation relative to the ground changed. This pitch was not measured in these tests and added an unknown that could not be corrected for using the rotation technique described in section 4.5. Datasets with rotational problems

have been discussed in the results presented below.

5.3.1 Baseline Undeformed Profile

A run was performed with the instrumented wheel lifted off the ground so that an undeformed profile could be measured. This profile was used to tune the post-processing program so that the undeformed data's minimum z value was at $z = 0\,\text{mm}$ (just touching the xy-plane). The amount of z shift required to move the lowest point to zero was stored and the value was applied to all future runs as a baseline to see the relative deformation despite the rotational problem discussed earlier.

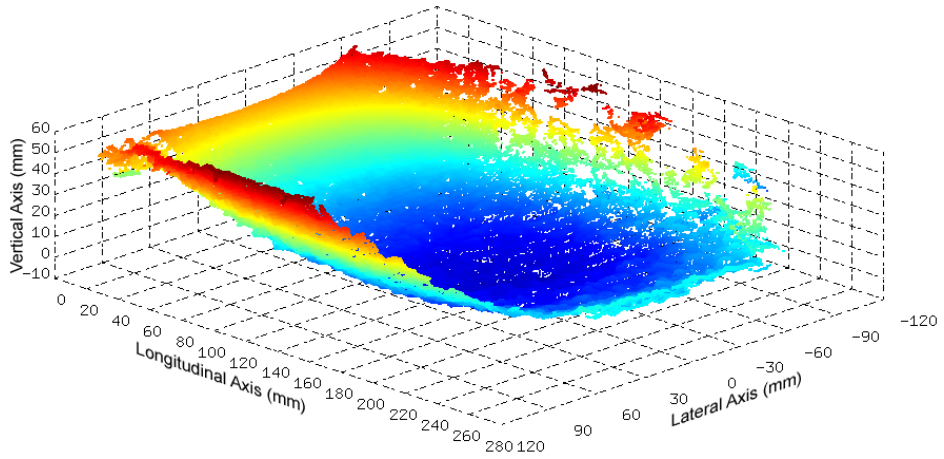

Figure 5.7: Undeformed profile of the inflated tyre.

The effectiveness of the stereo cameras as a measurement system is clearly evident in the size of the undeformed profile (Figure 5.7) with a measured area of $200\,\text{mm} \times 280\,\text{mm}$ and corresponding to an approximately 40° section of tyre arc. This area is large enough to capture the entire contact patch if the cameras are oriented correctly. The measured area is smooth with relatively few discontinuities as the cameras have a clear, well lit view of the undeformed section.

The cross-section of the undeformed data is given in Figure 5.8 and shows a smooth arc with a flat central section. As expected the lowest point in the data is at $z = 0\,\text{mm}$.

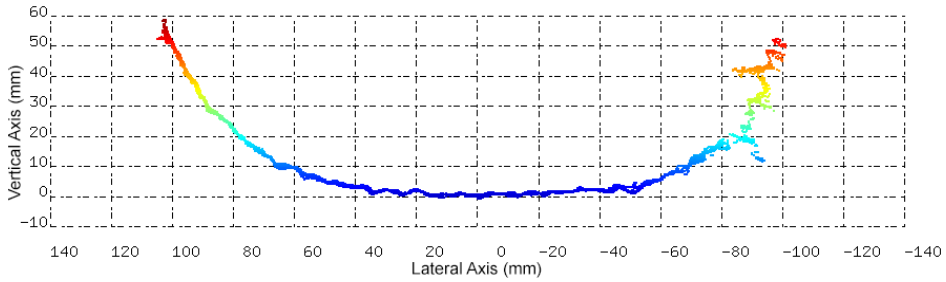

Figure 5.8: Undeformed cross-section of the inflated tyre.

5.3.2 Deformation on Flat Ground

The profile of the tyre as it rolled along a flat piece of ground, with constant velocity, was measured to illustrate the effect of vertical loading. The full profile (Figure 5.9) shows the tyre has been compressed and an almost complete oval shaped flat contact patch has formed and is easily measurable. This test was not affected by rotation in the torque arm as there were no obstacles to jerk the system and as such a reliable comparison can be made between these results and the undeformed results.

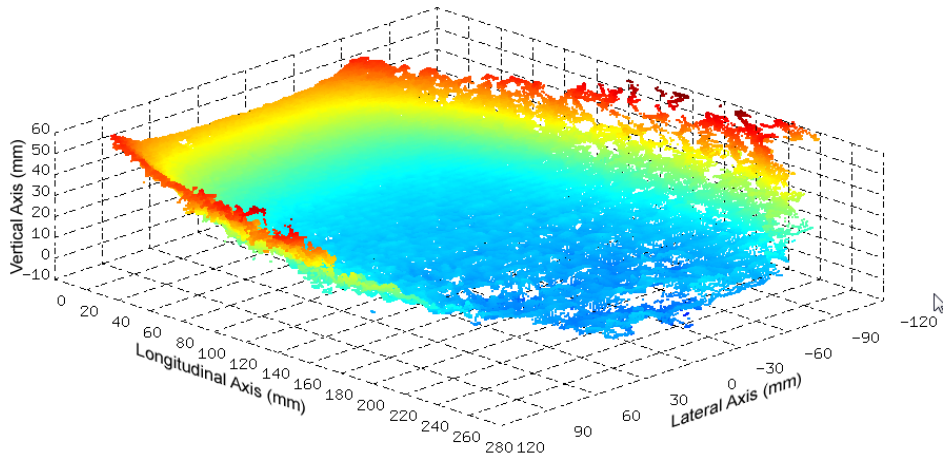

Figure 5.9: Deformation when driving on a flat piece of ground.

The cross-section of the flat ground deformation (Figure 5.10) shows that the lowest point on the tyre has deformed upwards by approximately 18 mm when compared to the undeformed profile and the central area seen in the undeformed cross-section is enlarged.

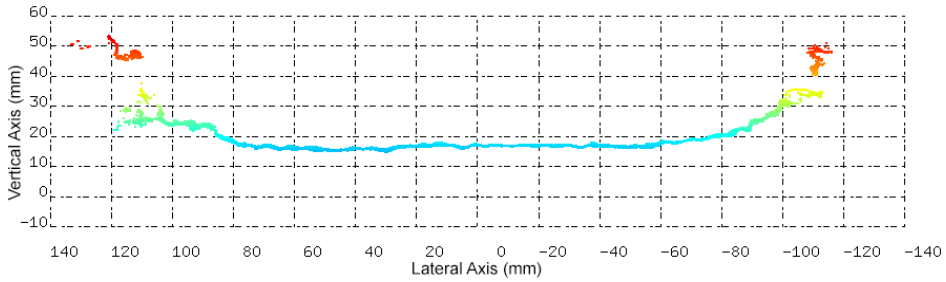

Figure 5.10: Cross section of the tyre when driving on a flat piece of ground.

5.3.3 Deformation Over Lateral Cleats

The lateral cleat test was performed by positioning the three cleats discussed in subsection 5.2.1 perpendicular to the vehicle's direction of travel and driving forward at low speed. The cross-sections shown below were all captured at the plane defined by the vertical and longitudinal axes. Profiles were captured at the point of maximum deformation over each of the obstacles.

Evidence that rotational problems were present in this dataset were visible from the profile of the data driving on the flat ground prior to mounting the first cleat. The data showed that the contact patch was not parallel to the ground plane and as such direct comparison to the undeformed data is not possible, however trends were still present and will be discussed below.

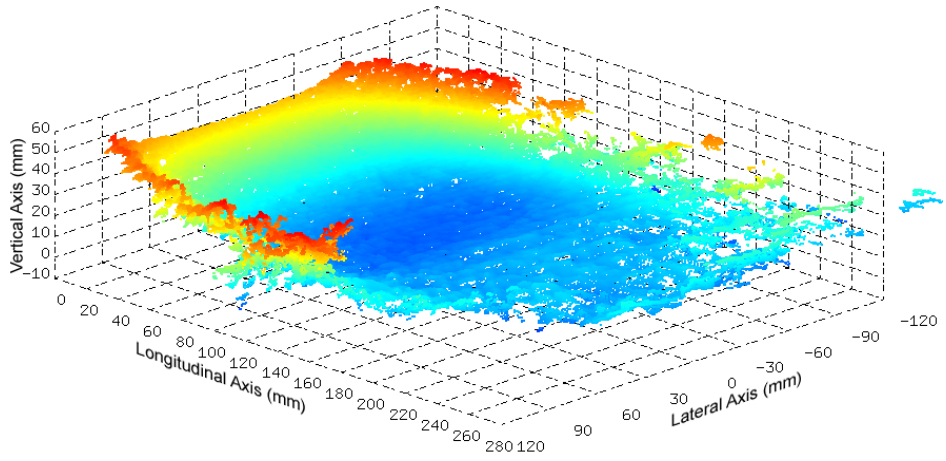

Figure 5.11: Deformation when driving over the 38 mm square section.

The profile as the tyre negotiated the 38 mm square section is shown in Figure 5.11. This profile shows hardly any deflection however when the longitudinal cross-section in Figure 5.12 is examined it is clear a vertical deflection of approximately 10 mm with reference to the flat

ground results is present. The cross-section's highest point shows a slightly flatter region for a length that is approximately 40 mm. This measurement corresponds with the width of the top face of the cleat and illustrates that the deformation is being measured accurately.

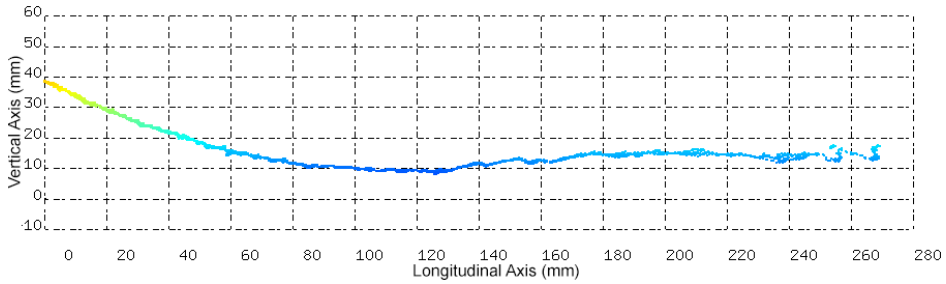

Figure 5.12: Longitudinal cross section of the tyre when driving over the 38 mm square section.

Due to the rotational issues in the data the vertical deflection over this cleat is most likely incorrect however the size of the object being negotiated is still well defined.

Figure 5.13 shows the tyre as it deforms over the 50 mm square section and shows a much larger deformation with respect to the profile in Figure 5.11. The cross-section (Figure 5.14) shows a vertical displacement of 20 mm which is significantly greater than the deflection seen with the 38 mm square section and a raised central area of about 50 mm which agrees once again with the real size of the cleat.

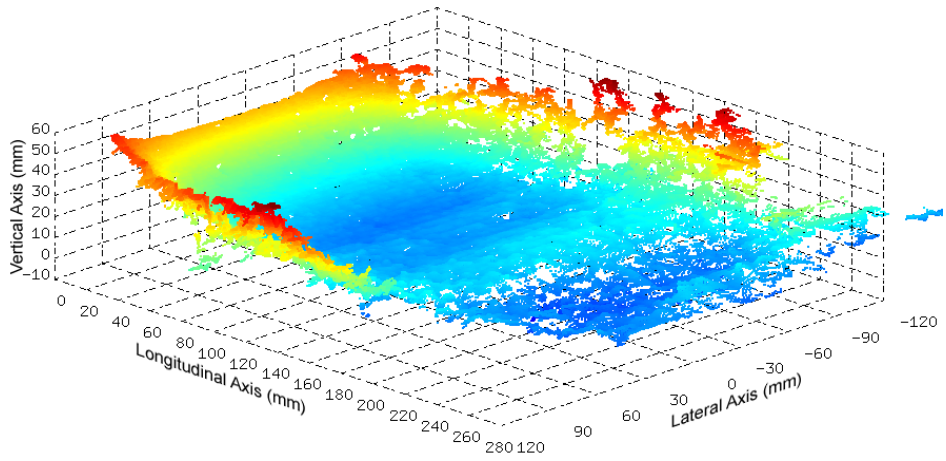

Figure 5.13: Deformation when driving over the 50 mm square section.

The profile of the tyre as it rolls over the 40 mm angle iron (Figure 5.15) has a pronounced, narrow raised region in the centre clearly showing the cleat passing beneath the tyre. The cross-section (Figure 5.16) indicates that the lowest point in this cross-section is significantly lower than the lowest point when traversing the 38 mm square section leading to the conclusion that the system rotated slightly relative to the earlier results.

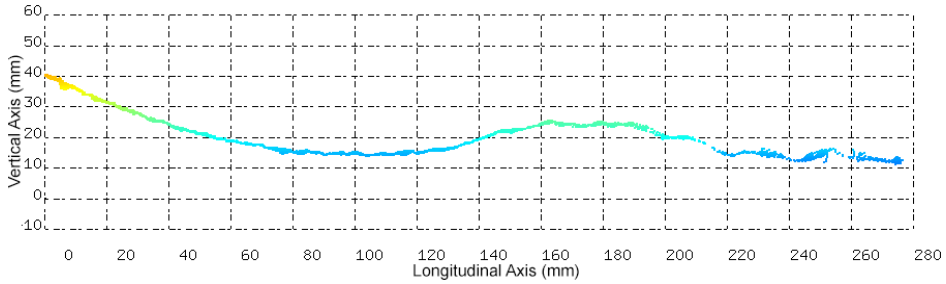

Figure 5.14: Longitudinal cross section of the tyre when driving over the 50 mm square section.

Even with this rotational problem the deformation in the cross-section is still interesting as it shows a much narrower region of high vertical displacement with reference to the 38 mm square section. The narrower area of vertical deformation corresponds with the triangular shape of the cleat.

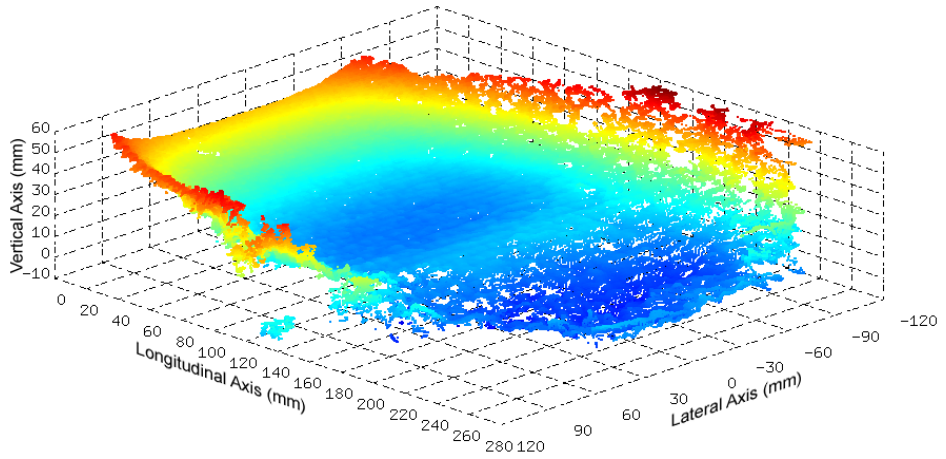

Figure 5.15: Deformation when driving over the 40 mm angle iron.

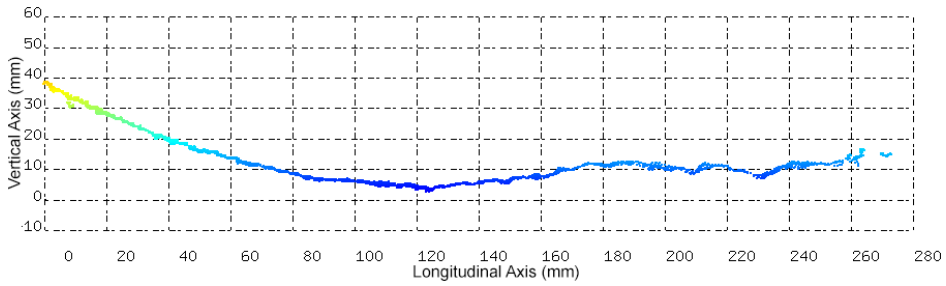

Figure 5.16: Longitudinal cross section of the tyre when driving over the 40 mm angle iron.

All the tests performed managed to capture the deformation over the obstacles with reasonable accuracy despite the rotation issues which plagued these runs. The cross-sections showed deformations that were slightly larger than the objects causing the deformation which is expected due to the thickness and rigidity inherent in the tyre carcass.

5.3.4 Deformation Over Longitudinal Cleat

The longitudinal test was performed by positioning the 38 mm square section parallel to the vehicle's direction of travel and in the path of the tyre. The results in this test were not severely affected by the rotational problems discussed earlier and can therefore be compared to the undeformed profile.

The profile of the tyre as tyre mounts the cleat (Figure 5.17) shows how the deformation starts at the front of the contact patch and radiates outward from the point of highest deformation to the point of lowest deformation which now has less vertical deformation than it would in the regular contact patch profile.

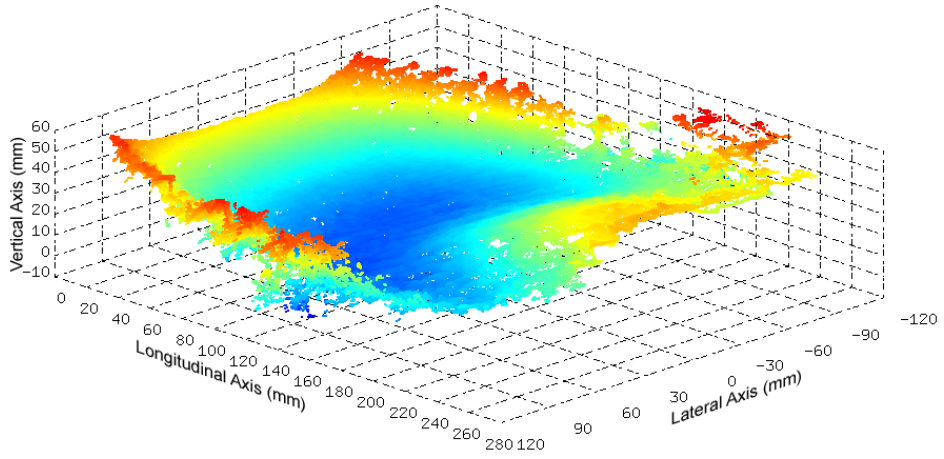

Figure 5.17: Deformation when driving onto the 38 mm square section.

The cross-section as the tyre mounts the cleat (Figure 5.18) shows the cleat entering the contact area slightly off centre and how the tyre begins to envelop the cleat on either side.

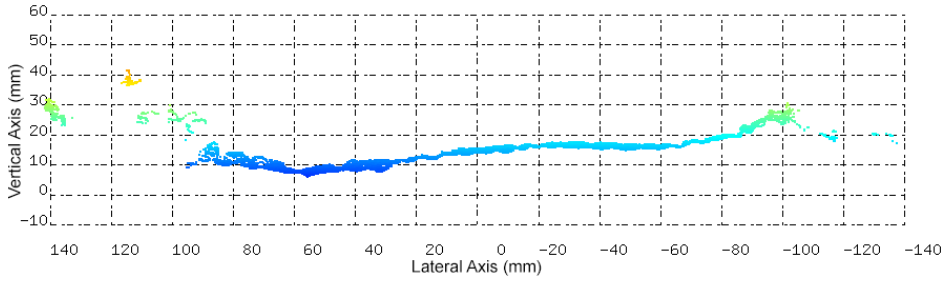

Figure 5.18: Lateral cross-section of the tyre when driving onto the cleatA.

The profile created when the tyre has completely mounted the cleat (Figure 5.19) shows a high ridge slightly off centre and two troughs where the sidewalls force the tyre down on

either side of the cleat. The trough on the left hand side of the cleat is more pronounced because the cleat was not in the centre of the tyre in this run.

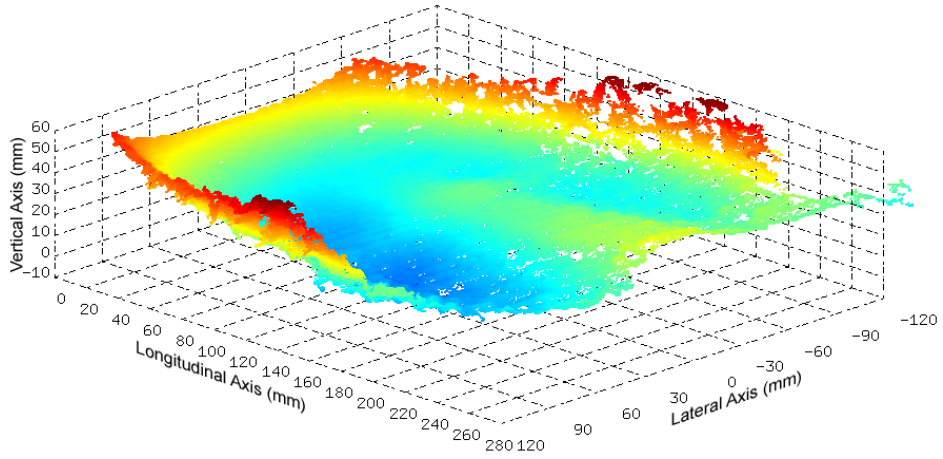

Figure 5.19: Deformation when driving in the middle section of the 38 mm square section.

The cross-section of the tyre (Figure 5.20) shows that the maximum deflection, relative to the undeformed profile, is approximately 28 mm which is less than the height of the cleat (38 mm). This discrepancy indicates that both sides of the tyre had left the ground and all tyre forces were being transferred through the area on top of the cleat. The width of the area of maximum deflection in the centre of the cross-section is approximately 40 mm which shows strong agreement with the width of the cleat's top face.

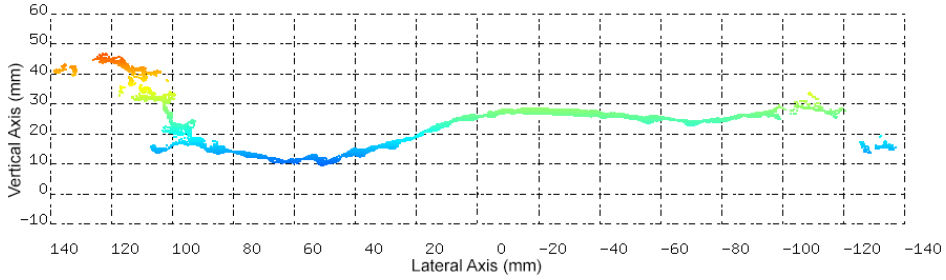

Figure 5.20: Cross section of the tyre when driving in the middle of the 38 mm square section.

The profile of the tyre as it dismounts the cleat (Figure 5.21) is similar to the profile when the tyre mounts the tyre except that the lowest point is in front of the cleat. The lowest point is lower than the expected contact patch deflection once again, due to the tyre regaining its undeformed shape before touching the ground and being compressed once more.

The cross-section as the tyre dismounts the cleat is shown in Figure 5.22 and is quite similar to the cross-section in Figure 5.18 except that it is slightly more complete with less measurement errors due to occlusion.

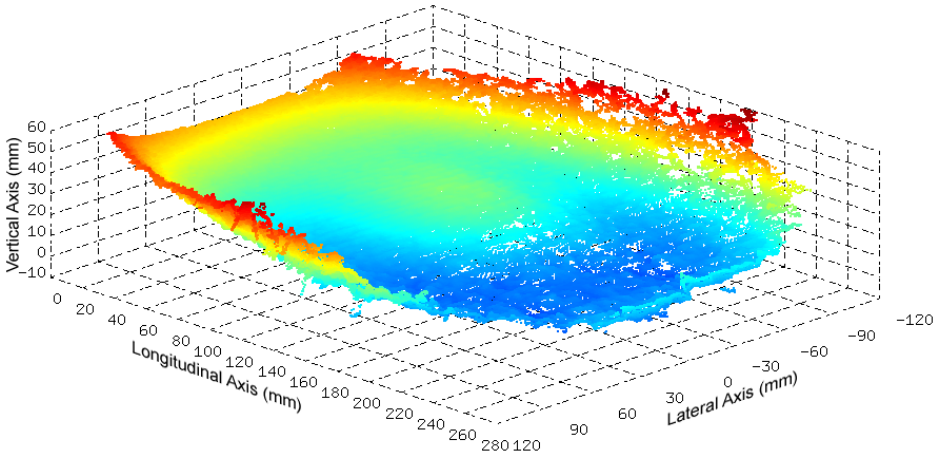

Figure 5.21: Deformation when driving off the 38 mm square section.

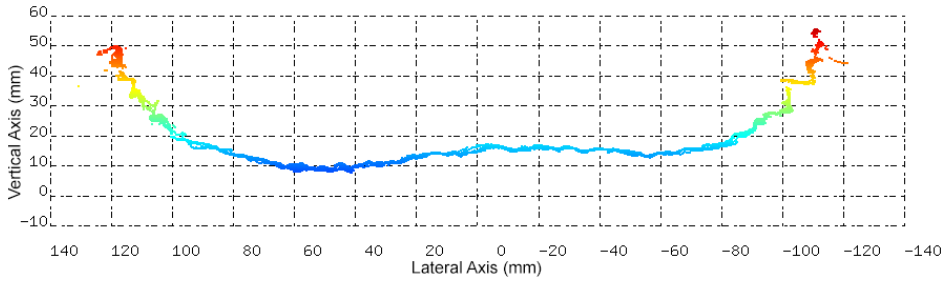

Figure 5.22: Lateral cross-section of the tyre when driving off the 38 mm square section.

In these tests the deformations captured with the tyre traversing the longitudinal cleat indicate that the system is working well and when there are no rotational errors all deflections are captured accurately.

5.3.5 Deformation Over 45° Cleats

The 45° cleat test was performed by positioning the cleats at 45° angles to the direction of travel of the vehicle. The results collected in this test showed some small rotation when they were processed so the cross-section will not lie at the same plane and therefore cannot be compared to one another or to the undeformed profile.

The three profiles produced when the tyre mounted the 38 mm square section, 50 mm square section, and 40 mm angle iron are shown in Figure 5.23, Figure 5.25, and Figure 5.27 respectively. The profiles all show a deformed area oriented at 45° across the centre of the contact patch with the 50 mm square section showing the highest deformation and the 38 mm square section, and 40 mm angle iron having similar deformations which agree with the geometry of the obstacles.

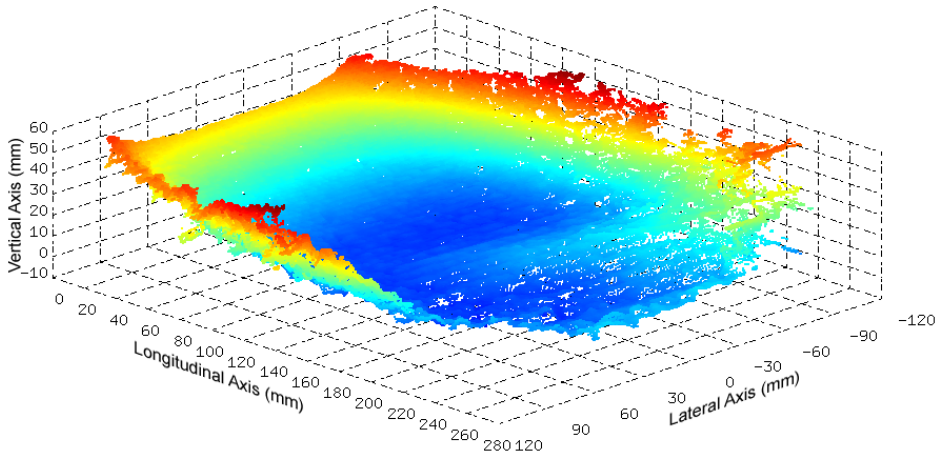

Figure 5.23: Deformation when driving over the 38 mm square section.

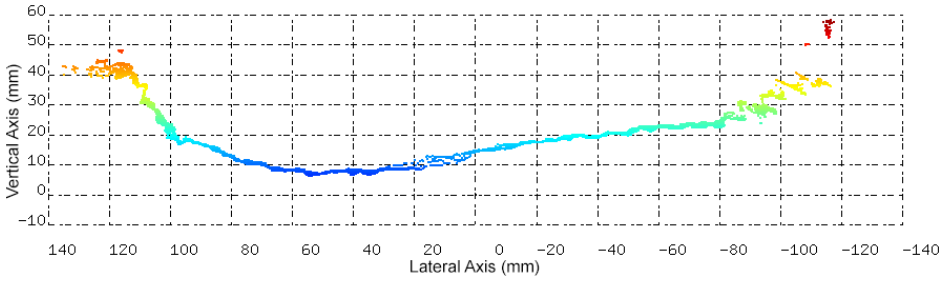

Figure 5.24: Cross section of the tyre when driving over the 38 mm square section.

Cross sections generated from the profiles are shown in Figure 5.24, Figure 5.26, and Figure 5.28 and reflect a similar trend to the one seen in the profiles except that the deformation over the 38 mm square section appears to be much larger than the deformation over 40 mm angle iron when the cross-sections are expected to agree. This is due to unquantifiable rota-

tion in the system causing the cross-sections to not lie on the same plane across datasets and which has been a major issue in this study.

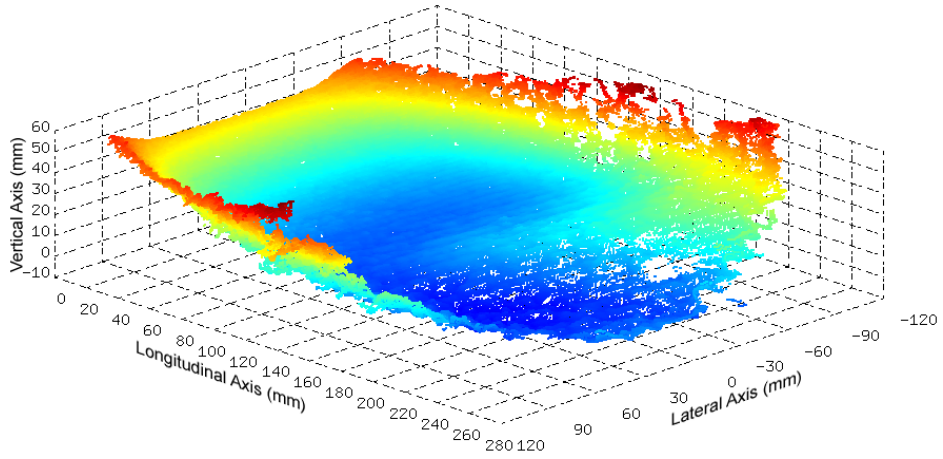

Figure 5.25: Deformation when driving over the 50 mm square section.

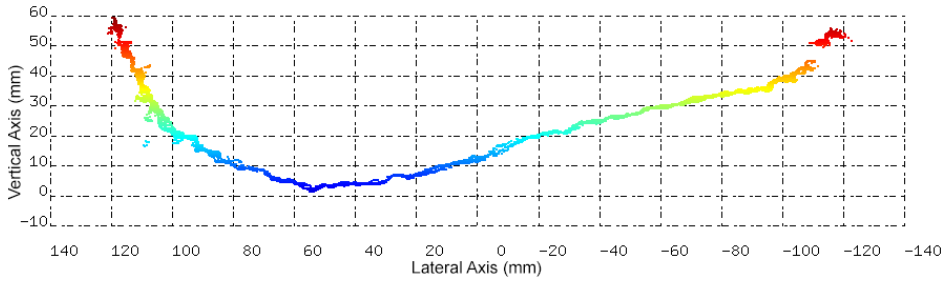

Figure 5.26: Cross section of the tyre when over the 50 mm square section.

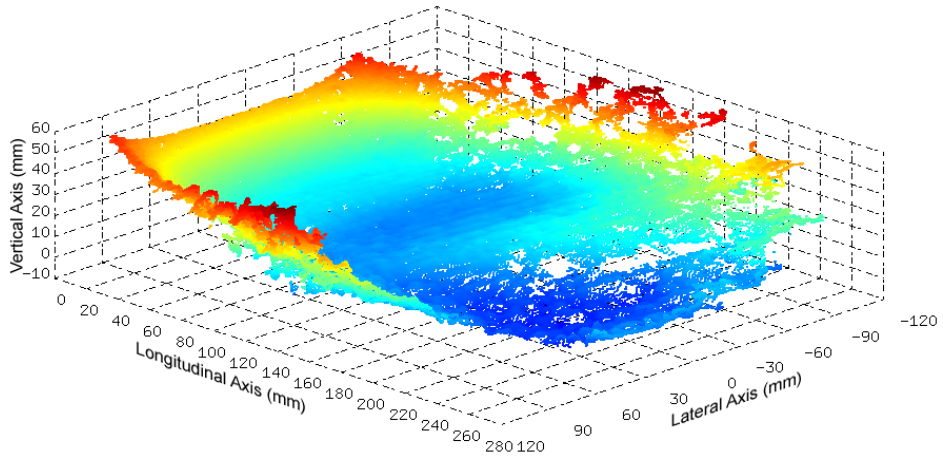

Figure 5.27: Deformation when driving over the 40 mm angle iron.

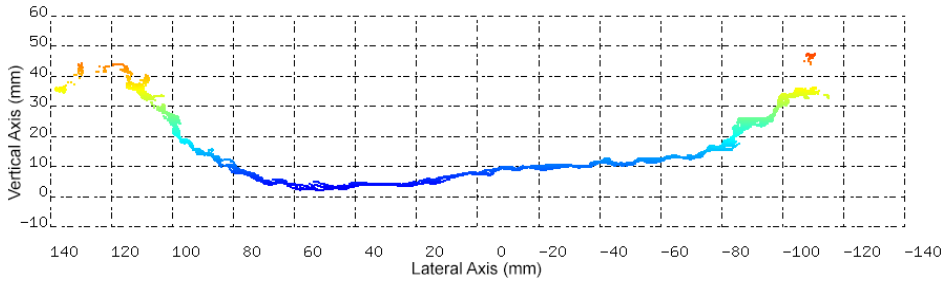

Figure 5.28: Cross section of the tyre when driving over the 40 mm angle iron.

5.3.6 Deformation Over Offset Lateral Cleats

The deformation when the tyre negotiated cleats which were only partially under the contact patch was measured by placing end of the cleats offset laterally from contact patch centre and driving forward at low speed. The results collected in this test did not seem to be affected by rotational issues when they were processed therefore the cross-section will lie on the same plane and can be compared to one another or to the undeformed profile.

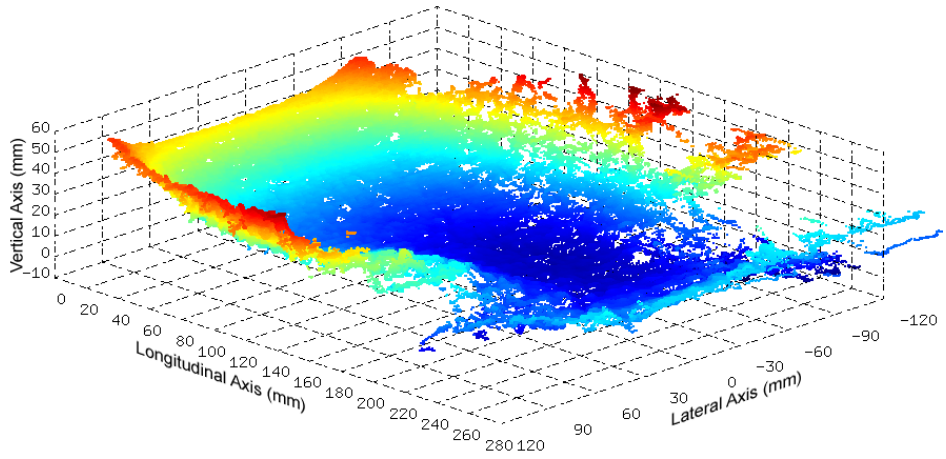

Figure 5.29: Deformation when driving over the 38 mm square section.

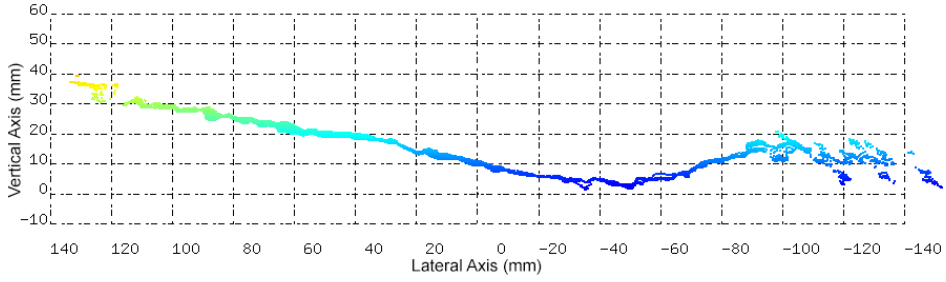

Figure 5.30: Cross section of the tyre when driving over the 38 mm square section.

The profiles of the tyre driving over the 38 mm square section, the 50 mm square section, and the 40 mm angle iron are shown in Figure 5.29, Figure 5.31, and Figure 5.33 respectively.

Deformation is visible on the right edge of the contact patch with the 50 mm square section having the largest vertical deflection and the 38 mm square section, and the 40 mm angle iron having very similar maximum deflection values.

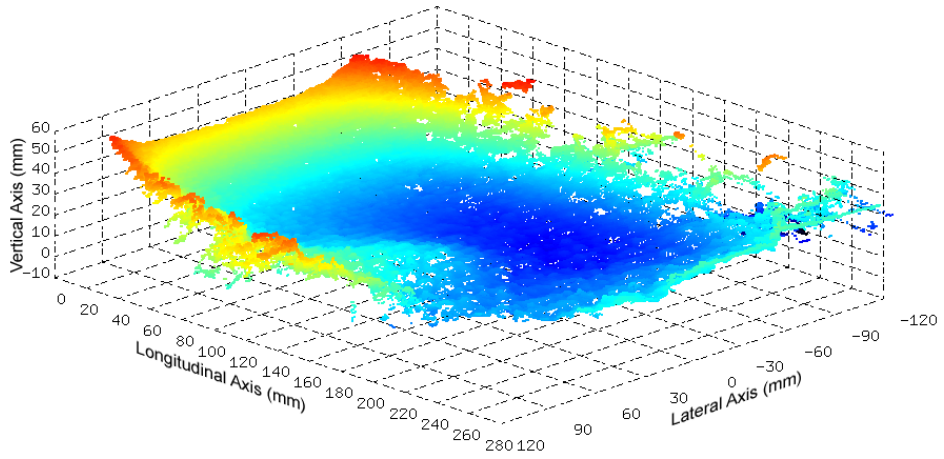

Figure 5.31: Deformation when driving over the 50 mm square section.

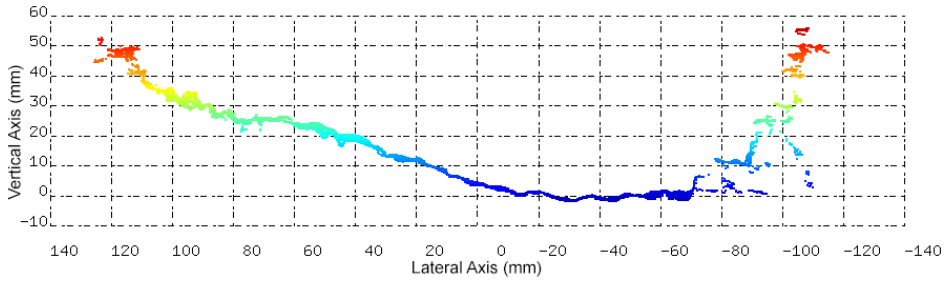

Figure 5.32: Cross section of the tyre when over the 50 mm square section.

The cross-sections shown in Figure 5.30, Figure 5.32, and Figure 5.34 have the same trend as the profiles with the 38 mm square section. The minimum value in all cross-sections is close to the minimum value of the undeformed profile because the vertical load on the tyre is supported in the deflected region and forcing the wheel upwards. This results in the tyre decompressing in the region not in contact with the cleat.

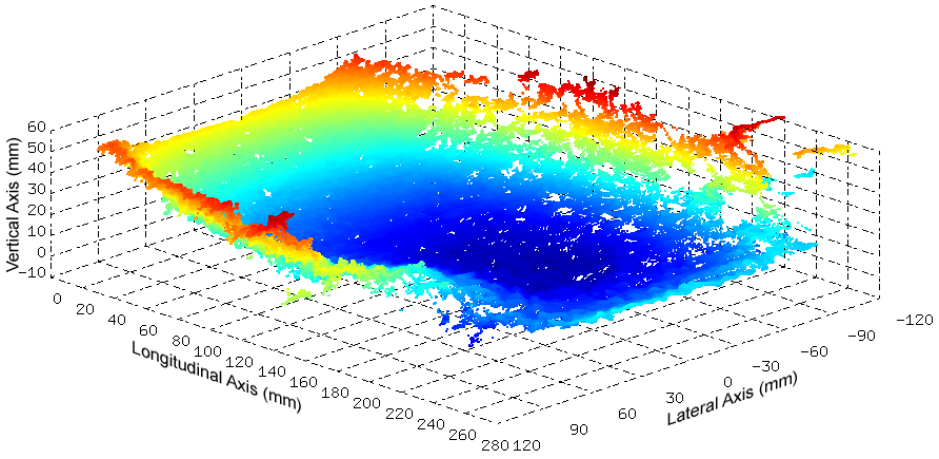

Figure 5.33: Deformation when driving over the 40 mm angle iron.

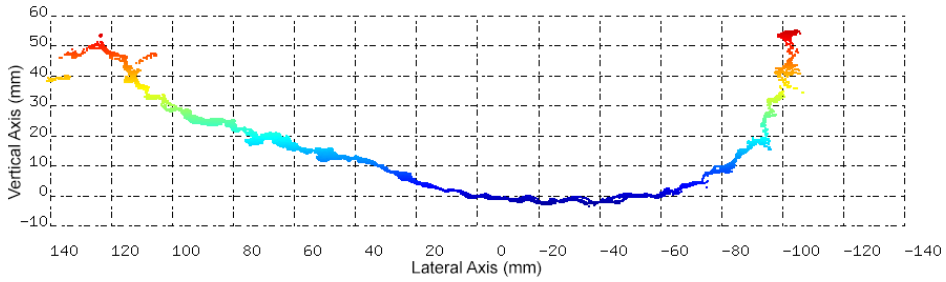

Figure 5.34: Cross section of the tyre when driving over the 40 mm angle iron.

5.3.7 Deformation Over Decreasing Size Nuts

Deformation of the tyre was captured as it passed over (at low speed) a series of nuts arranged in the direction of travel of the vehicle. The nuts decreased in size (shown in Figure 5.6) to test the sensitivity of the system to small obstacles. Results collected in this test were not affected by rotational issues because the measurements were taken during a single run over all the obstacles. The cross-sections will therefore lie on the same plane and can be compared to one another.

Profiles as the tyre passed over the M24, M20, M14, and M10 nuts are shown in Figure 5.35, Figure 5.37, Figure 5.39 and Figure 5.41 respectively. The nuts were not centred in the tyre's path therefore the maximum vertical deformation is not at the centre of the profile but rather to the right in the figures. The magnitude of the vertical deflection is observed to decrease as the size of the nuts decreases (as is expected) and the effect of the M10 nut is still visible in the data.

The cross-sections of the nuts are shown in Figure 5.36, Figure 5.38, Figure 5.40, and Fig-

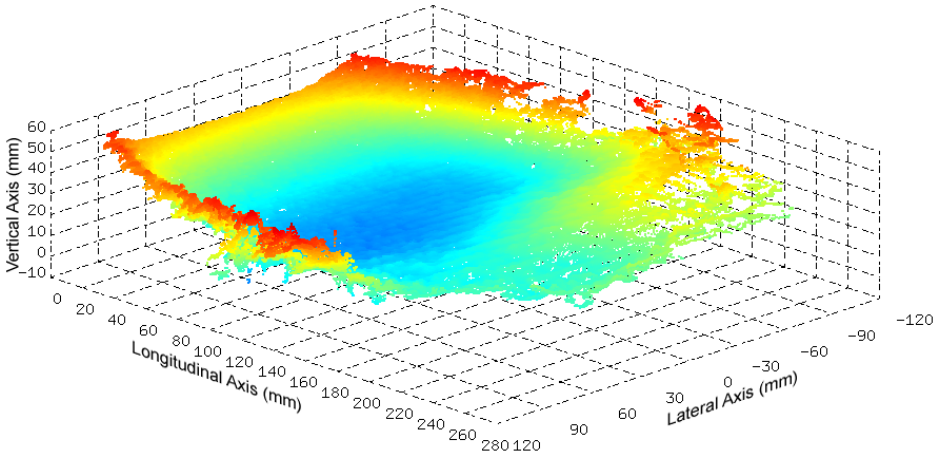

Figure 5.35: Deformation when driving over the M24 Nut.

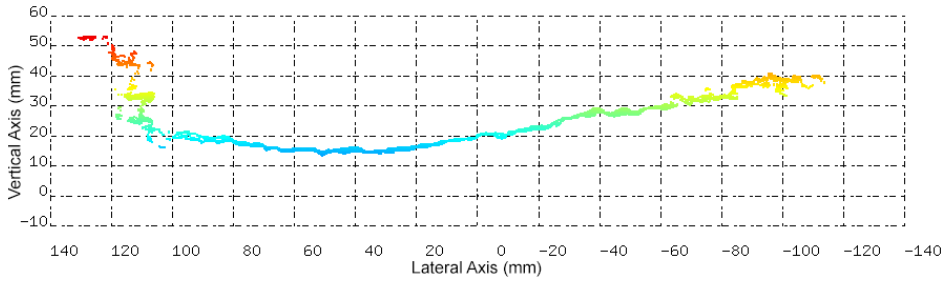

Figure 5.36: Cross section of the tyre when driving over the M24 Nut.

ure 5.42 and illustrate the vertical deformation observed in the profiles in more quantifiable manner. Maximum deformation was available for each cross section and since there were no rotational problems the measured deformation relative to the undeformed zero is compared with the obstacle height. The difference between the measured deformation and the obstacle height is calculated in Table 5.1 for each of the nuts.

Table 5.1: Maximum deformation in cross-sections relative to obstacle height

	M24	**M20**	**M14**	**M10**
Measured Deformation (mm)	30.0	26.0	19.0	16.0
Obstacle Height (mm)	23.1	19.6	13.6	10.0
Difference (mm)	5.9	6.4	5.4	6.0

The difference across all obstacles in this test changed by a maximum of 1 mm and was offset from the true height of the obstacles by approximately 6 mm. One would expect that the difference between the measured deformation and obstacle height would be closer to zero however it is believed that this discrepancy is due to the assumed centre of the contact patch

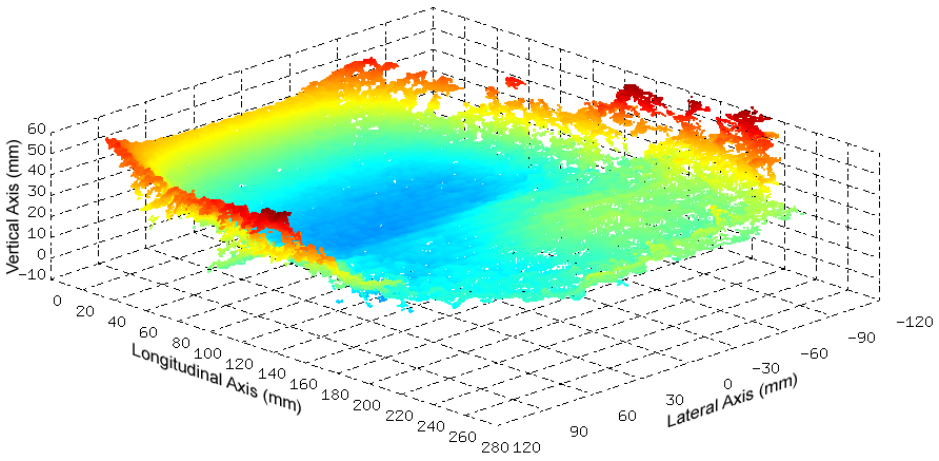

Figure 5.37: Deformation when driving over the M20 Nut.

(150mm in this case) being offset from the true centre, rather than a measurement error. It should be noted that the maximum deformation was read directly from cross-section graphs and is therefore slightly subjective. Some differences in measured deformation may be present in the data if the obstacle is small enough that it is partly in contact with a tread element and partly in one of the tread gaps. In this case the deformation would not reflect the height of the obstacle accurately.

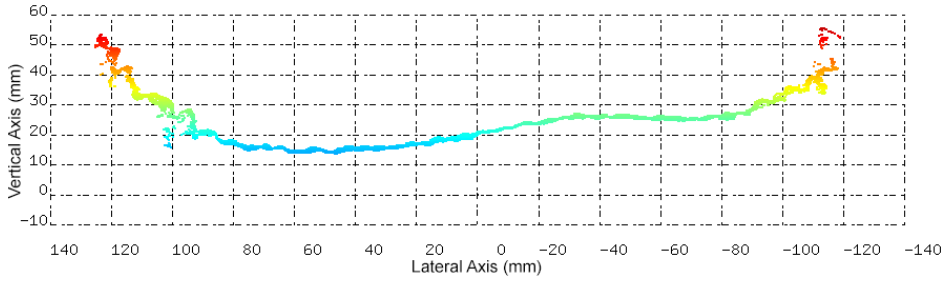

Figure 5.38: Cross section of the tyre when driving over the M20 Nut.

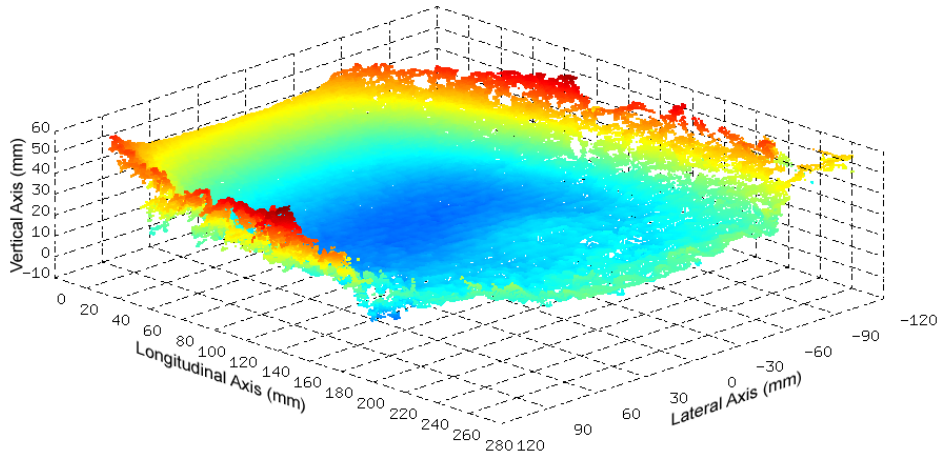

Figure 5.39: Deformation when driving over the M14 Nut.

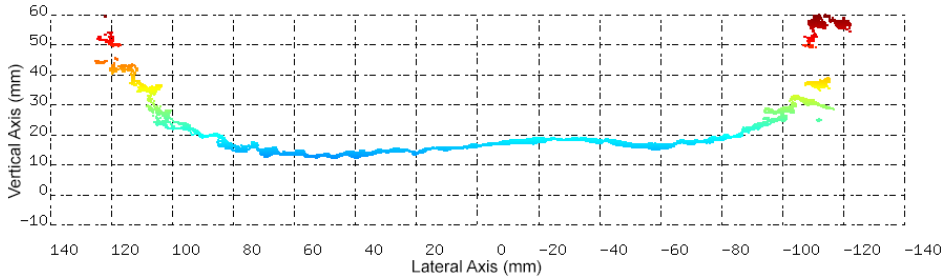

Figure 5.40: Cross section of the tyre when driving over the M14 Nut.

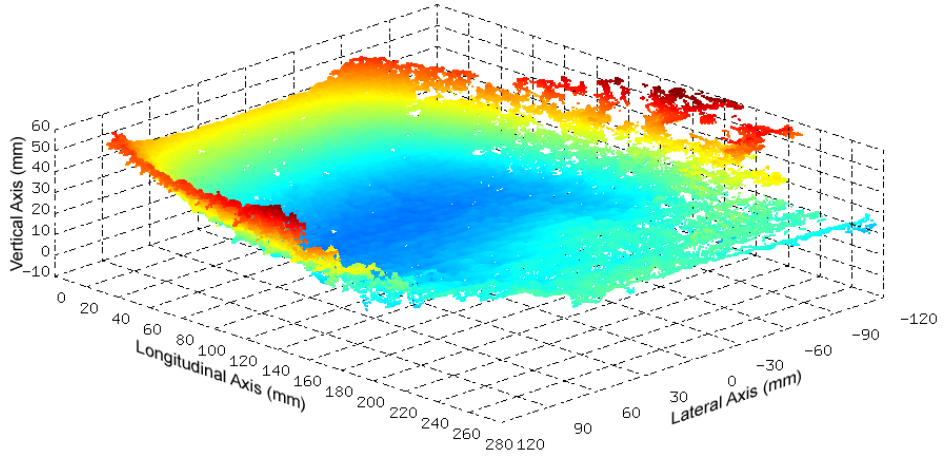

Figure 5.41: Deformation when driving over the M10 Nut.

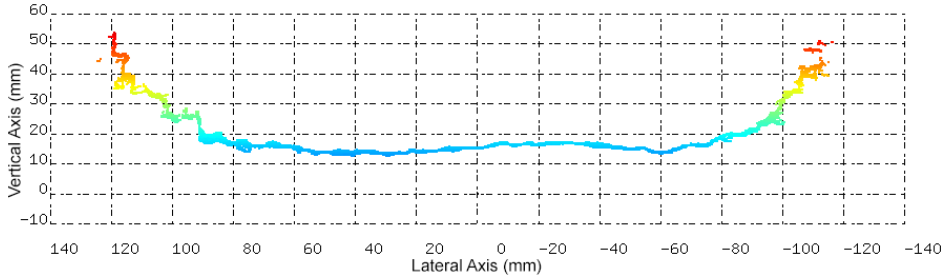

Figure 5.42: Cross section of the tyre when driving over the M10 Nut.

5.4 DISCUSSION OF RESULTS

The tests showed promising results and the system succeeded in allowing accurate full field measurement of the inner surface of the tyre while the wheel was rotating. Using stereo cameras in this manner provides possibilities not possible before, such as allowing the visualisation of the full 3D displacement field in an easily interpreted manner or measuring the entire contact patch at all times during a dynamic manoeuvre.

The problem encountered with unquantified rotation of the stabilising torque arm decreases the value of the current results, however it is easily remedied by improving the rigidity of the arms or measuring the angle using an encoder. Despite this mechanical issue the deformations captured by the system were sensible and showed expected features across all datasets.

The system accomplished its goal of measuring the contact patch while the tyre rotated and showed the potential of having a non rotating measurement system inside a rotating wheel. The system has only been put through basic tests in this study and there exist many possibilities for future investigation such as:

- Extensive static laboratory tests to find the ideal point to place the cross-sectional plane and to validate the contact patch size and deformation.
- Investigating the possibility of synchronising the data captured by a wheel force transducer and the data from the cameras.
- Placing multiple discrete objects are under the tyre at the same moment to determine both the deformation and the filtering effect the tyre has on discontinuities on the road surface.
- Synchronising measurements from inside the tyre with measurements of deformable terrain (such as sand) in front of and behind the tyre to investigate tyre terrain interaction on soft soil.
- Attempting to recover the strain in the tyre carcass (and potentially to the tread surface) by using consecutive images captured by the system.
- Using images to estimate the slip and friction experienced by the tyre in the contact region.

All of these investigations would yield interesting and relevant results for the vehicle dynamics and terra-mechanics communities and would not be possible without the camera system.

CHAPTER 6

CONCLUSION AND RECOMMENDATIONS

6.1 SUMMARY

It was determined from a review of the literature that the contact patch is a very important region when modelling vehicle dynamics. The size of the contact patch is determined by the vertical load on the vehicle and the inflation pressure of the tyre. The literature also indicated that while the contact patch has been measured statically, only single point deformation has been measured dynamically and then only once per revolution. Since the dynamic and static contact patches are not identical there exists a need for dynamic contact patch measurement over varied terrain. New data about the dynamic contact patch could prove invaluable in both the development and validation of tyre models.

This study proposed and tested a novel method for constantly measuring the contact patch and general tyre deformation through all rotations by placing stereo cameras inside the tyre. The system consists of a stabilisation mechanism based on planetary gear's ratios, a pair of stereo cameras, a lighting rig, and a DAQ which recorded video data for post processing. Software required calibrate the system and process the videos was developed based on the OpenCV library to support the testing of the measurement system.

Stereo videos were captured of the tyre deforming over various obstacles while the DAQ was recording and the data was post processed using the software mentioned above. The results were compared against one another and trends were identified and discussed.

6.2 SYSTEM PERFORMANCE AND LIMITATIONS

The test results showed great promise managing to capture the full field deformation in most cases with high accuracy and consistency however a few problems were experienced and should be corrected in future iterations of the system. The biggest problem was with unmeasured rotation of the camera rig relative to the ground plane, the rotations resulted from the torque arm not being rigid enough to prevent all rotation in the system. These rotations often made it difficult to accurately compare results over different obstacles as the zero reference was not constant and cross-sections could not be guaranteed to lie on the same plane.

6.3 RECOMMENDATIONS

The system that has been developed opens up many avenues of investigation in vehicle dynamics by providing a relatively inexpensive means of measuring physical phenomenon inside the tyre. One possible use for the system, other than conventional deformation measurement, would be to extract strain information across the captured region using image correlation techniques to track the relative motion of discrete points between frames.

While the current system proved capable of measuring deformations a few teething problems were experienced, therefore the following features should be added or upgraded in future iterations:

- The stabilising torque arm should be redesigned to ensure that it is rigid to increase the reliability of the measured data.
- An improved method of finding the contact patch should be developed so that calibration is more straightforward and the angle of the cameras can be known with greater accuracy.
- Changing the position of the cameras so that they are further apart. This should increase accuracy of the matching algorithms.
- More advanced stereo matching algorithms can be used to increase the quality of the depth maps.
- Investigation should be done on whether it is possible to have the cameras capture

images only at discrete rotation intervals rather than at a constant number of frames per second. This would allow the system to capture at a fixed rotational displacement independent of the rotational velocity of the wheel.

- The software for post processing the results should be improved to allow direct finer control over video playback and more output options.
- Investigating the possibility of measuring strain and wheel slip angle with the same image data gathered by the cameras.

6.4 CONCLUSION

The aim of this study was to design, implement and test a system which would solve some inherent problems of dynamic contact patch measurement and provide a proof of concept for using stereo cameras inside a tyre. It is successful in this regard as the implemented design was able to capture a much larger area of the contact patch than any previous research while the tyre rolled over various obstacles with high consistency and sensitivity. The problems experienced with the system were minor and easily remedied in a future design iteration. Data produced using this measurement technique sheds new light on tyre phenomenon and future research could discover more uses for the data over-and-above simple contact patch measurement. This system is extremely versatile and holds great potential for measuring not just tyre deformation but many other difficult to quantify tyre properties in the future.

BIBLIOGRAPHY

Andor Technology Ltd. 2015. *Rolling and Global Shutter*. [Online]. Available at: http://www.andor.com/learning-academy/ [Accessed]: 15/01/2015.

Botha, T. R., P. S. Els and C. Becker. 2012. 'Measuring Tyre Contact Patch Deformation Using Stereography'. In: *Proceedings of the12th European Regional Conference of the ISTVS*. Pretoria.

Botha, T. R. 2014. 'Digital Image Correlation : Applications in Vehicle Dynamics'. PhD thesis. University of Pretoria.

Bradski, G. and A. Kaehler. 2008. *Learning OpenCV: Computer Vision with the OpenCV Library*. 1st ed. O'Reilly Media, p. 555.

Brown, D. C. 1966. 'Decentering distortion of lenses'. In: *Photometric Engineering* 32.3, pp. 444–462.

Cosin. 2015. *F-Tire Product Brief*. [Online]. Available at: http://www.cosin.eu/ [Accessed]: 15/03/2015.

Cyganek, B. and J. P. Siebert. 2011. *An introduction to 3D computer vision techniques and algorithms*. John Wiley & Sons.

Digia. 2015. *Qt Project*. [Online]. Available at: http://qt-project.org/ [Accessed]: 07/03/2014.

Durini, D. 2014. *High Performance Silicon Imaging : Fundamentals and Applications CMOS and CCD Sensors*. Ed. by D. Durini. 1st. Cambridge: Woodhead Publishing is.

Els, P. S. 2012. *Wheel Force Transducer Research and Development*. Tech. rep. University of Pretoria. [Online]. Available at: https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/ADA547059.xhtml.

Freudenstein, F. and A. T. Yang. 1972. 'Kinematics and statics of a coupled epicyclic spur-gear train'. In: *Mechanism and Machine Theory* 7.2, pp. 263–275.

Green, R. W. 2011. 'A Non-contact Method for Sensing Tire Contact Patch Deformation Using a Monocular Vision System and Speckled Image Tracking'. PhD thesis. Auburn University.

Hartley, R. and A. Zisserman. 2003. *Multiple view geometry in computer vision*. Cambridge university press.

Hiraoka, N., R. Matsuzaki and A. Todoroki. 2009. 'Concurrent Monitoring of In-plane Strain and Out-of-plane Displacement of Tire Using Digital Image Correlation Method'. In: *Journal of Solid Mechanics and Materials Engineering* 3.11, pp. 1148–1159.

Hirschmuller, H. 2005. 'Accurate and efficient stereo processing by semi-global matching and mutual information'. In: *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on* 2, 807–814 vol. 2.

IEI. 2015. *IEI World*. [Online]. Available at: `http://www.ieiworld.com/` [Accessed]: 24/05/2015.

Khronos. 2015. *OpenGL Main Page*. [Online]. Available at: `https://www.opengl.org/`.

Kowa. 2015. *Kowa LM4NCL*. [Online]. Available at: `http://www.kowa.eu/lenses/en/products.php` [Accessed]: 24/05/2015.

Magori, V., V. R. Magori and N. Seitz. 1998. 'On-line determination of tyre deformation, a novel sensor principle'. In: *Ultrasonics Symposium, 1998. Proceedings., 1998 IEEE*. Vol. 1. IEEE, pp. 485–488.

Matsuzaki, R. and A. Todoroki. 2007. 'Wireless flexible capacitive sensor based on ultra-flexible epoxy resin for strain measurement of automobile tires'. In: *Sensors and Actuators A: Physical* 140.1, pp. 32–42.

Michael, M., J. Salmen, J. Stallkamp and M. Schlipsing. 2013. 'Real-time stereo vision: Optimizing Semi-Global Matching'. In: *2013 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1197–1202.

Nagappan, S. 2007. *Ambient-Light Sensors Mimic the Eye*. [Online]. Available at: `http://www.photonics.com/Article.aspx?AID=28345` [Accessed]: 09/01/2015.

OpenCV. 2014. *Open Computer Vision Library*. [Online]. Available at: `http://opencv.org/` [Accessed]: 01/09/2014.

Pohl, A., R. Steindl and L. Reindl. 1999. 'The "intelligent tire" utilizing passive SAW sensors measurement of tire friction'. In: *Instrumentation and Measurement, IEEE Transactions on* 48.6, pp. 1041–1046.

PointGrey. 2015. *Point Grey Flea3 1.3 MP Mono USB3 Camera*. [Online]. Available at: `http://www.ptgrey.com/flea3-13-mp-mono-usb3-vision-vita-1300-camera` [Accessed]: 24/05/2015.

Potenger, M. G. 2006. 'Contact Patch (Footprint) Phenomena'. In: *The Pneumatic Tire*. Ed. by J. Walter and A. Gent. February. Washington D.C: National Highway Traffic Safety Administration, pp. 231–285.

Radke, R. J. 2012. *Computer Vision for Visual Effects*. Cambridge University Press.

Sandu, C., E. Pinto, S. Naranjo, P. Jayakumar and B. Ross. 2012. 'Soft Soil Tire Model Development and Experimental Testing'. In: *Proceedings of the12th European Regional Conference of the ISTVS*. DTIC Document.

Scharstein, D. 2002. 'A Taxonomy and Evaluation of Dense Two-Frame Stereo'. In: *International Journal of Computer Vision* 47.1, pp. 7–42.

Scharstein, D. and R. Szeliski. 2014. *Middlebury Stereo Vision Page*. [Online]. Available at: `vision.middlebury.edu/stereo/` [Accessed]: 21/09/2014.

Scharstein, D. and R. Szeliski. 2003. 'High-accuracy stereo depth maps using structured light'. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 1, pp. 195–202.

Sergio, M., N. Manaresi, M. Tartagni, R. Canegallo and R. Guerrieri. 2006. 'On a road tire deformation measurement system using a capacitive–resistive sensor'. In: *Smart materials and structures* 15.6, p. 1700.

Sharma, A. K. and K. P. Pandey. 1996. 'A review on contact area measurement of pneumatic tyre on rigid and deformable surfaces'. In: *Journal of Terramechanics* 33.5, pp. 253–264.

Tuononen, A. 2009. 'On-board estimation of dynamic tyre forces from optically measured tyre carcass deflections'. In: *International journal of heavy vehicle systems* 16.3, pp. 362–378.

Wulfsohn, D. and S. K. Upadhyaya. 1992. 'Determination of dynamic three-dimensional soil-tyre contact profile'. In: *Journal of terramechanics* 29.4, pp. 433–464.

Zhang, Z. 2000. 'A Flexible New Technique for Camera Calibration'. In: *IEEE Transactions On Pattern Aanlysis And Machine Intelligence* 22.11, pp. 1330–1334.