UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

CORRECTION OF RADIALLY ASYMMETRIC LENS DISTORTION WITH A CLOSED FORM SOLUTION AND INVERSE FUNCTION

by

Jason Peter de Villiers

(99191319)

Study Leader: Dr Ronelle Geldenhuys

Submitted in partial fulfilment for the degree

Master of Engineering (Electronic Engineering)

in the

Faculty of Engineering, the Built Environment & Information Technology

UNIVERSITY OF PRETORIA

October 2007

SUMMARY

CORRECTION OF RADIALLY ASYMMETRIC LENS DISTORTION WITH CLOSED FORM SOLUTION AND INVERSE FUNCTION
by
Jason Peter de Villiers
Study Leader: Dr Ronelle Geldenhuys
Department of Electrical, Electronic & Computer Engineering, Univ. of Pretoria
Master of Engineering (Electronic Engineering)

The current paradigm in the lens distortion characterization industry is to use simple radial distortion models with only one or two radial terms. Tangential terms and the optimal distortion centre are also seldom determined. Inherent in the models currently used is the assumption that lens distortion is radially symmetrical. The reason for the use of these models is partly due to the perceived instability of more complex lens distortion models.

This dissertation shows, in the first of its three hypotheses, that higher order models are indeed beneficial, when their parameters are determined using modern numerical optimization techniques. They are both stable and provide superior characterization. Although it is true that the first two radial terms dominate the distortion characterization, this work proves superior characterization is possible for those applications that may require it.

The third hypothesis challenges the assumption of the radial symmetry of lens distortion. Building on the foundation provided by the first hypothesis, a sample of lens distortion models of similar and greater complexity to those found in literature are modified to have a radial gain, allowing the distortion corrections to vary both with polar angle and distance from the distortion centre. Four angular gains are evaluated, and two provide better characterization. The elliptical gain was the only method to both consistently improve the characterization and not ‘skew’ the corrected images. This gain was shown to improve characterization by as much as 50% for simple (single radial term) models and by 7% for even the most complex models.

To create an undistorted image from a distorted image captured through a lens which has had its distortion characterized, one needs to find the corresponding distorted pixel for each undistorted pixel in the corrected image. This is either done iteratively or using a simplified model typically based on the Taylor expansion of a simple (one or two radial coefficients) distortion model. The first method is accurate yet slow and the second, the opposite. The second hypothesis of this research successfully combines the advantages of

both methods without any of their disadvantages. It was shown that, using the superior characterization of high order radial models (when fitted with modern numerical optimization methods) together with the 'side-effect' undistorted image points created in the lens distortion characterization, it is possible to fit a 'reverse' model from the undistorted to distorted domains. This reverse characterization is of similar complexity to the simplified models yet provides characterization equivalent to the iterative techniques. Compared to using simplified models the reverse mapping yields an improvement of more than tenfold - from the many tenths of pixels to a few hundredths.

Keywords:
Lens distortion, distortion correction, numerical optimization, radial distortion, asymmetric radial distortion, inverse distortion

UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

OPSOMMING

CORRECTION OF RADIALLY ASYMMETRIC LENS DISTORTION WITH CLOSED FORM SOLUTION AND INVERSE FUNCTION
deur
Jason Peter de Villiers
Studieleier: Dr Ronelle Geldenhuys
Departement Elektriese-, Elektroniese & Rekenaar Ingenieurswese, Univ van Pretoria
Meester in Ingenieurswese (Elektronise Ingenieurswese)

Die huidige standaard vir die karakterisering van lensvervorming, is om eenvoudige radiale vervormingsmodelle met slegs een of twee radiale terme te gebruik. Raaklynige terme en die optimale vervormingsmiddelpunt word selde bepaal. Hierdie modelle is gebaseer op die implisiete aanname dat lensvervorming radiaal simmetries is. Die motivering vir die gebruik van hierdie modelle, kan deels toegeskryf word aan die persepsie dat meer komplekse modelle onstabiel is.

In hierdie verhandeling word aangetoon, in die eerste van drie hipoteses, dat hoër orde modelle wel voordelig is, wanneer passing met behulp van moderne numeriese optimeringsmetodes geskied. Die modelle is stabiel en verskaf aansienlik beter karakterisering van lensvervorming. Die eerste twee radiale terme domineer wel die karakterisering, maar dit word aangetoon dat beter resultate bereik kan word in gevalle waar dit benodig word/geregverdig is.

In die derde hipotese word die aanname van radiaal simmetriese lensvervorming bevraagteken. Deur die eerste hipotese verder te neem, word modelle met vergelykbare en groter kompleksiteit as dié in die literatuur aangepas/uitgebrei om radiale aanwins in te sluit. Die modelle laat beide die poolhoek en afstand van die vervormingsmiddelpunt van die korreksies toe om te varieer. Vier metodes is ge-evalueer. Twee van hierdie metodes lei tot beter karakterisering. Die koniese/elliptiese aanwins metode is die enigste wat deurgaans die karakterisering verbeter en nie ’n skewe gekorrigeerde beeld lewer nie. Die aanwins verbeter die karakterisering met tot 50% vir eenvoudige (enkel radiale term) modelle en tot 7% vir die mees komplekse modelle.

Deur die gepaardgaande vervormde beeldelement vir elke onvervormde beeldelement te bepaal, kan ‘n onvervormde beeld van ‘n vervormde een gevorm word, waar die betrokke lens se vervorming gekarakteriseer is. Dit kan iteratief gedoen word, of deur die gebruik

van ‘n vereenvoudigde model wat gebasseer is op ‘n Taylor uitbreiding van ‘n eenvoudige (een of twee radiale koëffisiënte) vervormingsmodel. Die eerste metode is akkuraat, maar stadig, en die tweede metode presies die teenoorgestelde. Die tweede hipotese in die verhandeling kombineer die voordele van beide metodes sonder enige van die nadele. Dit word aangetoon dat dit moontlik is om ‘n ‘inverse’ model te pas van die onvervormde na die vervormde areas, deur beter karakterisering van hoër orde radiale modelle (wanneer gepas met moderne numeriese optimeringsmetodes) tesame met die ‘newe-effek’ van onvervormde beeldpunte wat geskep is tydens karakterisering lensvervormings. Hierdie inverse karakterisering is vergelykbaar in kompleksiteit met die vereenvoudigde modelle, maar lewer ekwivalente resultate aan iteratiewe tegnieke. In vergelyking met vereenvoudigde modelle, lewer die inverse passing ‘n tienvoudige verbetering – van tiendes van ‘n beeldelement na slegs ‘n paar honderstes.

Sleutelwoorde:
Lens vervorming, vervormingskorreksie, numeriese optimering, radiale vervorming, asimmetriese radiale vervorming, inverse vervorming

Table of Contents

List of Tables

List of Figures

LIST OF ABBREVIATIONS AND ACRONYMS

2D	Two Dimensional
3D	Three Dimensional
CCD	Charge Coupled Device
FOV	Field Of View
FFT	Fast Fourier Transform
FR	Fletcher-Reeves Algorithm
LMA	Levenberg-Marquardt Algorithm
RMS	Root Mean Square
SD	Steepest Descent
TBD	To Be Determined

CHAPTER 1.

RESEARCH OVERVIEW

This section describes the research that was done, it places the research in context and discusses the contribution made. Finally, the methodology of the research is described and an outline of the rest of this dissertation provided.

1.1. Introduction

Photogrammetry is a measurement technique whereby three dimensional (3D) measurements are made from one or more two dimensional (2D) images taken of the object to be measured. With the advent of more and more inexpensive camera equipment (which make use of cheaper lenses with higher distortion), particularly digital photography equipment, cameras are being used for increasingly varied and diverse applications such as machine vision [1–9], image stitching [10, 11], product defect detection [12], motion capture [2], better video compression [13], 3D measurement/reconstruction [14-18], and the digital archiving of documents and archaeological artefacts [19, 20]. Previously, expensive lenses which optically reduced distortion, or expensive equipment to measure and characterize lens distortion were used [7, 14, 16, 18]. Much work has been done to cheaply and easily calibrate lenses using everyday images or images of simple calibration patterns, examples of this can be found in [2, 4, 8, 9, 12, 14, 15, 17, 19-23].

Figure 1.1 illustrates the relationship between a point in the image plane, and the distortion and image centres, and forms the basis for the overview of the different parameters involved in camera calibration.

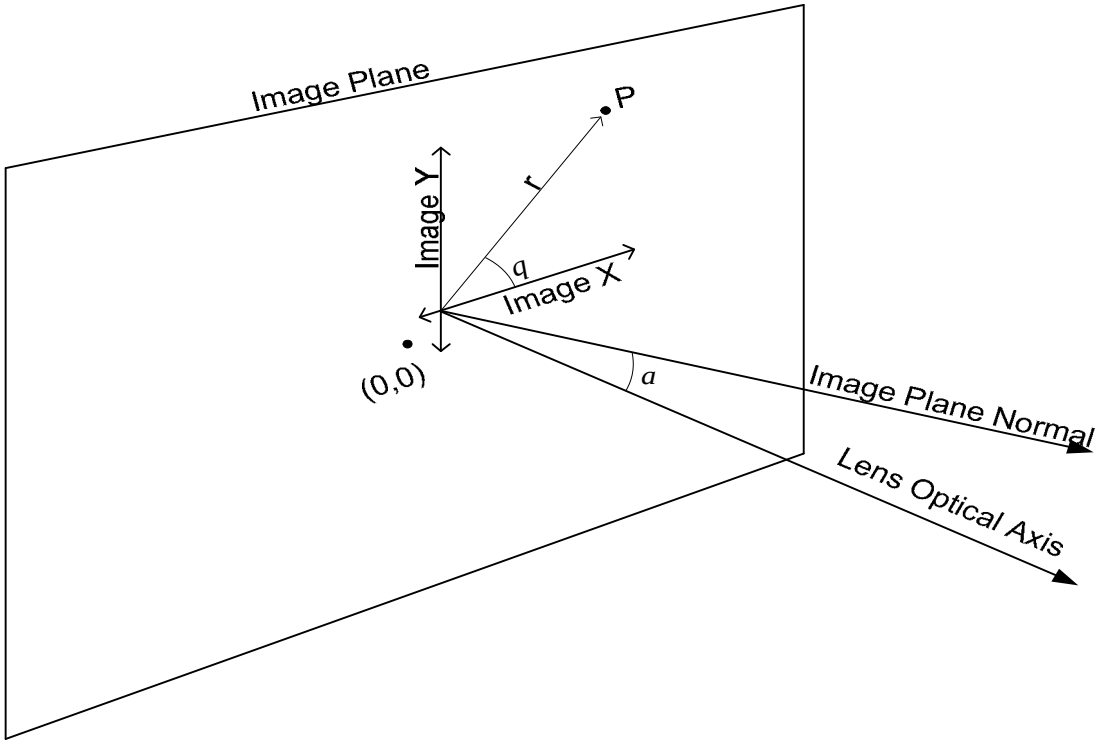

Figure 1.1 Image Plane Diagram depicting the point P in the image plane and its polar coordinates from the principle point in the image. The difference between principle point and image centre, and the image plane normal and lens optical axis are also shown.

With reference to Figure 1.1, *P* is the projection of a point in space onto the image plane. *(0, 0)* is the centre of the image plane, note that the intersection of the lens optical axis does not (in general) intersect the plane at this point, and is also offset from the normal of the image plane by an angle α. θ is the angle of the line (in the image plane) joining the optical axis intersection point and the projected point *P*, *r* is the length of this line segment.

The relationship between a point in free space and its projection on the image plane is dependant on the following parameters [1, 3, 5-9, 16, 24, 25]:

- Camera Extrinsic Parameters. These parameters transform a point in world coordinate to camera coordinate, and include the position and attitude of the camera in the world. Afterwards the point is still 3D but is now referenced to the camera.
- Camera Intrinsic Parameters. These parameters transform a point in camera coordinates to a pixel position, and include focal length, X scale Factor, Y scale factor, X/Y orthogonality and the optical axis intersection point.
- Lens Distortion. All lenses induce a non-linear mapping between a point and its projection on the image plane, the amount of nonlinearity is dependant on various

parameters including the field of view (FOV) and the exact optical geometry of complex lenses [25, 26]. This difference between the actual image position of a point and the point predicted by the extrinsic and intrinsic camera parameters is the distortion.

It is thus necessary to correct for lens distortion if accurate measurements are to be made using the linear intrinsic and extrinsic camera parameters [1-27]. This is the focus of this dissertation, as the determination of extrinsic and intrinsic parameters is a mature science. This dissertation does not study the effect of aberrations which are points in the image where the effective focal length differs from the nominal focal length (refer to §3.2).

1.2. Problem statement

Current lens distortion techniques do not fully explain all distortion evident in some lenses, refer to §4.3 for an example of this. Additionally, many proposed distortion correction models choose, for reasons of numerical stability of parameter determination, to ignore known distortion contributions. Examples of this are ignoring higher order radial distortion terms [e.g. 7, 17, 21-23]; ignoring the offset between distortion centre and image centre [e.g. 7, 23]; and ignoring tangential distortion [e.g. 7 -9]. Table 4.1 quantifies the distortion contributions considered by the models found in literature. Many lens distortion models require numerical iteration even after their distortion parameters have been determined, in order to calculate inverse distortion.

The aim of this research is to attempt to alleviate the above phenomena.

1.3. Scope of research

This research will focus on three primary goals:

1. To determine what numerical methods are used to determine the parameters for the various distortion models, and investigate how this can be made more robust, specifically with the aim of being stable with the inclusion of additional model parameters.
2. To determine if a suitably accurate estimation of the inverse distortion (distortion correction) can be repeatably and robustly determined.
3. To investigate if the implicit assumption inherent to non-cartesian models that distortion is radially symmetric is valid, and whether extending models to allow for radial asymmetry will improve distortion correction.

1.4. Research context

Photogrammetry and lens distortion are phenomena that have been studied and modelled for many decades and are mature sciences. The determination of camera intrinsic and extrinsic parameters has had much work done [1, 3, 5-9, 16, 24, 25], and is also mature. Lens distortion has had less work done on it, and with the exception of a few landmark papers such as those by Conrady [26] and Brown [14, 27] only received attention with the advent of computing power and digital imaging. However, much work has been done since, notably by Tsai [7], Stein [17] and Zhang [8, 9]. This research aims to build on this knowledge and advance it a step further, specifically to facilitate real-time precision distortion correction and inversion.

1.5. Research Approach

The approach taken in this research is one of study and hypothesis formulation followed by experimental verification and analysis as described below.

Research Questions

During the process of conducting the literature study, the following questions were used to determine the hypotheses and to verify them:

- What assumptions, both implicit and explicit, are being made?
- What parallels to other fields can be drawn?
- What advances in technology and knowledge have occurred?
- How can this be tested/verified?

Literature Study

An extensive literature study was undertaken, in which seminal papers, ideas and persons were identified. An understanding of lens distortion, and the different models thereof was sought. These models were then classified in terms of distortion type, complexity, relevance and effectiveness taking into account both the time they were derived and their intended context of use.

Hypothesis Formulation

Using the questions listed above and the understanding gained from the literature study, hypotheses were determined.

Experiment set-up and data capturing

For each hypothesis, the data required to verify it was determined. Given this required data an experiment to obtain it was derived, as well as a measure to determine the effectiveness of the hypotheses.

Hypothesis verification/validation

Using the data provided by the experiments each hypothesis was studied and tested to determine its effect.

Analysis of results

Using the effectiveness measure and the results of each hypothesis' implementation, the effect on lens distortion and thus the ultimate viability of each hypothesis was determined.

1.6. Notation

This section defines the mathematical notation used throughout this dissertation.

$\bar{x}$ or $\boldsymbol{x}$	A vector
$\bar{x}_k$ or $\boldsymbol{x}_k$	The value of vector $\boldsymbol{x}$ at iteration k.
x or X	A scalar
$\|\bar{a}\|_x$	The Xth norm of the vector $\boldsymbol{a}$.

1.7. Organization of the dissertation

This rest of this dissertation is organised in the following manner:

Chapter 2: Literature study summarizes the results of the literature study that was undertaken. The current state of knowledge of lens distortion is presented in a manner applicable to the research undertaken. The need for distortion modelling is presented followed by an overview of the history of photogrammetry. Thereafter, distortion models found in literature are used to characterize lens distortion first by method of modelling (polar vs. Cartesian) and then by time of use (before, during or after image capture). The need for determining the inverse mapping from the undistorted to distorted domains is

explained, and finally a thorough discussion of numerical optimisation as pertinent to lens distortion characterization is provided with emphasis on four particular algorithms.

Chapter 3: Configuration contains the configuration of the experiments used to gather the data which was used to verify the hypotheses. The equipment used to capture the experimental data, the camera and lens, and the optical reference jigs are described. Line straightness measures based on gradients and fitting errors are derived.

Chapter 4: Hypothesis 1: Parameter estimation is robust when determined with modern optimization methods shows that higher order radial models can be used to better characterize distortion provided sound optimization principles are followed..

Chapter 5: Hypothesis 2: Correction of barrel distortion can be modelled as pincushion distortion shows that inverse distortion need not be a trade off between speed of execution and accuracy, a method with both characteristics is devised..

Chapter 6: Hypothesis 3: Distortion is not necessarily radially symmetrical quantatively shows that distortion is not perfectly symmetrical and that allowing asymmetry can enhance even the superior characterization of Chapter 1.

Chapter 7: Conclusions the results of the research are placed in context and scope for future work is provided. Each hypothesis is discussed, including its impact on current lens distortion paradigms, how they may be applied in practice and when they are applicable.

CHAPTER 2.

LITERATURE STUDY

2.1. The need for lens distortion modelling

Every photogrammetric and machine vision application has different requirements. If the requirement is to make an image seem undistorted to cursory inspection, then correction to the nearest pixel will suffice as it did for Fernandes et al [12]. Similarly if one is required to stitch images together to form a larger image, then elaborate distortion modelling is unnecessary as shown by Hsu and Sawhney [11]. However there are applications where cameras are used to make highly precise measurements, and where as much of the distortion as possible is desired to be modelled and accounted for.

In 1992 Shih et al [34] did a study on the effect of different noise sources on camera calibration and also included the effects of ignoring lens distortion. Clearly lenses that exhibit substantial distortion, such as the fish-eye lens, can never have their distortion ignored, and Shih et al derived an error bound which can be used to determine if the distortion is acceptable for a particular application and lens.

Jeong et al [24], purposefully chose to neglect the effect of lens distortion in determining extrinsic and intrinsic camera parameters. Their focus was the effect of using different line widths and they found that the different calibration parameters varied by between 1% and 3%. The determined orientation of the camera was found to vary by as much as 2°. This is despite using an undisclosed lens that did not appear to exhibit significant distortion in the sample images in their paper, thus emphasising the need for lens distortion modelling.

2.2. Overview of the history of photogrammetry

Photogrammetry is as old as photography itself - the ambition being to make accurate measurements from 2D images. Perspective transformations however could not fully account for the positioning of real-world image points in the resultant image. This difference is due to lens distortion. Clark and Fryer [35] provide an excellent history of the evolution of lens distortion modelling and correction, this is summarized below with the addition of a few points pertinent to this dissertation.

Initially the applications for photogrammetry were primarily cartographic in nature, resulting largely from military surveillance needs. This progressed from terrestrial cameras to aerial cameras with the advent of reliable aircraft in World War I. The era from 1950 through 1979 saw much development in the fields of lens distortion modelling, advancing it from the presiding qualitative approach to a more quantitative geometric approach based on an understanding of the operation of lens systems.

Prior to 1950, lens distortion was treated by projecting each captured image through lens with which it was captured. Thereafter the primary issues of concern were the calibrated focal length (or principal distance) and the principle point. The latter is the point around which the distortion is most symmetrical radially. Distortion radial asymmetries and tangential components were deemed unimportant as long as they were within tolerance (typically 15-30μm in the image). These distortions were attributed to "thin prism distortion," which equated the distortion to that that a hypothetical thin prism placed in front of the lens would induce.

In 1965 at the Annual Convention of the American Photogrammetric Society Brown shifted the paradigm, he presented previously classified work which showed that the tangential distortion was attributable to the decentring of the lenses in the lens stack. He further proved that Conrady had been correct in his largely ignored ray tracing paper [26] published in 1919. Brown later published these findings in his scathing 1966 landmark paper: "Decentering Distortion of Lenses" [27].

In 1971 Brown published another seminal paper: "Close Range Camera Calibration" [14]. In this paper Brown both provided a robust, computer-implementable algorithm to determine the lens distortion parameters for a certain object distance and how to interpolate between distortion parameters for two different object distances to produce the parameters for a third distance. Brown's method to characterize a lens is known as the "plumb line method" which essentially formally states the truism: 'straight lines are straight.' Thus any curvature apparent in an image of straight lines is due to lens distortion, as the perspective transformation is linear. Brown's lens distortion formula is given as:

$$x_u = x_d + (x_d - x_c)(K_1 r^2 + K_2 r^4 + \ldots) + (P_1(r^2 + 2(x_d - x_c)^2) + 2P_2(x_d - x_c)(y_d - y_c))(1 + P_3 r^2 \ldots)$$
$$y_u = y_d + (y_d - y_c)(K_1 r^2 + K_2 r^4 + \ldots) + (2P_1(x_d - x_c)(y_d - y_c) + P_2(r^2 + 2(y_d - y_c)^2))(1 + P_3 r^2 \ldots)$$

where:

(x_u, y_u) = undistorted image point,

(x_d, y_d) = distorted image point,

(x_c, y_c) = centre of distortion,

K_n = Nth radial distortion coefficient,

P_n = Nth tangential distortion coefficient,

$r = \sqrt{(x_d - x_c)^2 + (y_d - y_c)^2}$, and

"…" indicates an infinite series.

Eq 2-1

Eq 2-1 has been used unchanged since. With the advent of low cost computing and Charge Coupled Devices (CCD) making digital imagery commercially and increasingly cheaply available, Brown's equations have been applied to such realms as computer vision and process automation. Clark and Fryer note that this has been done successfully despite digital cameras and their lenses differing in size, focal length, image size, mass, quality and focus (fixed versus variable/zoom) from their analogue counterparts that were purpose built for photogrammetry.

Eq 2-1 is an infinite series, however in practice a finite number of radial and tangential parameters are used. Typically one or two radial terms and often zero tangential terms, in this dissertation the term "higher order" when referring to Eq 2-1 means that more three or more radial terms are modelled and probably two or more tangential terms.

Since the 1980's the primary focus on lens distortion research has been in the digital domain [35]. The focus was on calibration using inexpensive equipment, resulting in computationally efficient distortion corrections.

In 1987 Tsai [7] published a paper describing a two step method for camera calibration, where the principle point is fixed at the centre of the image, and a single radial parameter and no tangential parameters are modelled. Tsai claimed that higher order versions of Eq 2-1 result in numerical instability, without specifying the exact method used to attempt to determine the parameters.

In 1999 Zhang [8] published a method whereby a simple (laser printable) coplanar pattern can be used to characterize a lens' distortion from three unknown views of the pattern. Zhang presented his method for two radial and no tangential distortion parameters claming that the remaining terms were comparably insignificant.

2.3. Types of lens distortion models

Two types of lens distortion models are evident in literature, with the primary difference being whether distortion is modelled with polar or Cartesian coordinates.

2.3.1. Radial lens distortion models

The vast majority of lens distortion models are of the radial type, this is due to most lenses having spherical elements. Radial distortion models assume that points move radially in or out towards or away from a distortion centre roughly in the centre of the image. The amount of distortion is solely dependant on the distance of the point from the distortion centre (i.e. the distortion is radially symmetrical). Some models allow for tangential distortion, although the direction and magnitude of the distortion is also solely dependant on the distance from the distortion centre. This is summarized in Figure 2.1.

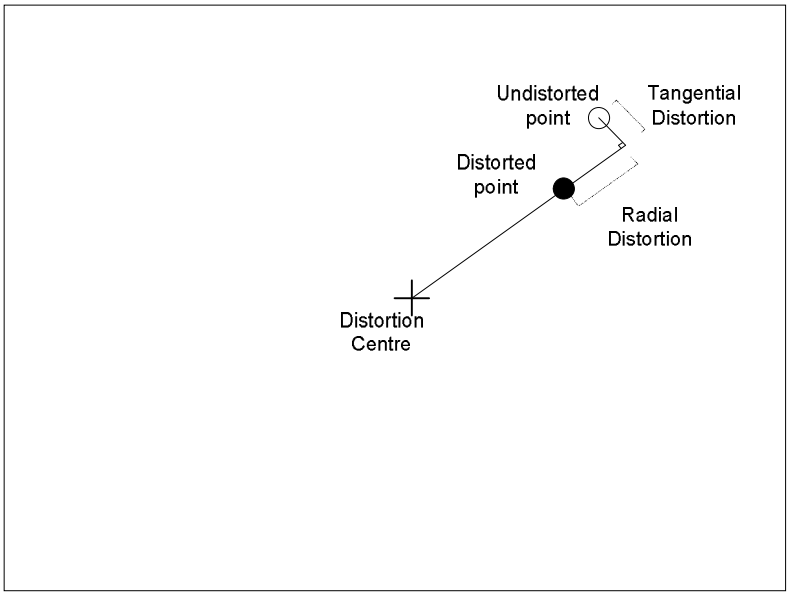

Figure 2.1 Radial and tangential distortion depicted in the image plane. The black circle is the distorted (measurable) image point and the white circle its corresponding undistorted image point, the outline is the image plane.

Depending on whether the radial component of the distortion pulls points closer to the centre (typical of wide angle lenses [10]) or pushes the points out (typical of telephoto

lenses [10]) the distortion is referred to as barrel distortion or pin cushion distortion respectively. Figure 2.2 depicts this, and it also shows how a barrel distorted image is smaller than its corresponding undistorted image, and how pincushion distorted images are larger.

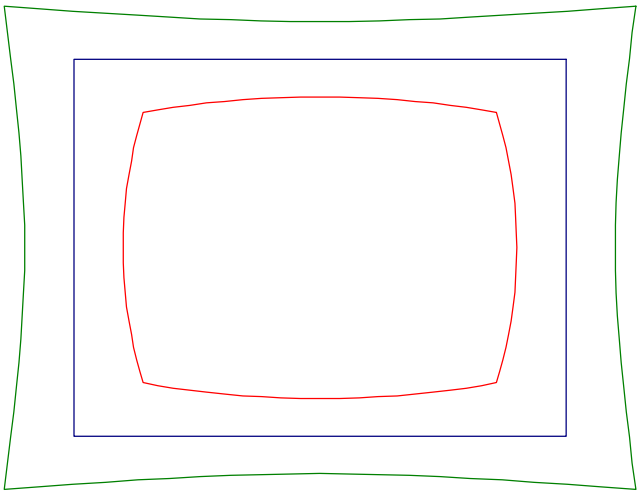

Figure 2.2 Barrel and pincushion distortion. The blue rectangle indicates the boundaries of the undistorted image; the red shape depicts the blue rectangle subject to barrel distortion; and the green shape the blue rectangle subject to pincushion distortion.

The pre-eminent radial distortion model is that of Brown [14, 27], however other models do exist, and tend to based on rigorous study of the physical properties of lenses. Note that Brown's model is only strictly radial if no tangential terms are modelled, however this degenerate case is commonly used. The next sections discuss each of these types in further detail.

2.3.1.1. Curve fitting methods

These models are those characterized by the Brown lens distortion model. All of these methods optimize the parameters of their model until the distortion is at a minimum.

One of the most common distortion measures is the straightness of straight lines. Pioneered by Brown in 1971 [14] with his 'plumb-line method' this distortion measure has been used extensively since. Essentially, points on a line (that should be straight) are found, either manually as Brown did in 1971 or automatically by edge detection algorithms on digital images. Given these points, one possible distortion measure is the sum of the perpendicular

(or vertical) distances of the points from the best fit straight line through them. Brown used a polar representation of the straight lines and the perpendicular distance from the lines, as did [4]. [21 -23] used the sum of the differences between the gradients calculated from adjacent pairs or points along the line and thus did not need to actually determine the best fit line. Cucchiara et al [2] used the Hough Transform and took the maximum of the Hough Space as an inverse indication of the amount of distortion. Clearly these methods will be adversely affected if real-life curves are mistakenly identified as lines which should be straight in the image.

A second measure, which requires a known physical set up, is where 3D reference points (whose positions in space are precisely known) are projected onto the image plane, via the camera's extrinsic and intrinsic parameters (resulting in the undistorted point) and then compared to the actual image points distorted by the radial distortion model. The sum of the distances (in pixels) between the projected and the actual images of the points (after having distortion applied to them) is the distortion measure. These methods (e.g. [1, 3, 16]) are dependant on the prior knowledge of the extrinsic and intrinsic parameters, or require their simultaneous determination. Additionally, these methods are sensitive to the error/noise with which the reference points' positions are known and the accuracy with which the (sub)pixel position of the reference points can be found in the image.

Finally, if multiple images of a scene with repeatably (algorithmically) identifiable markers are taken by a camera whose intrinsic parameters remain constant, then it is possible to use epipolar geometry to determine the relative camera positions (i.e. extrinsic parameters) and project the markers from one image to the other. The difference between the projection of a marker onto an image and the actual position is due to distortion. The distortion measure is then simply the sum of distances between the projected markers and actual markers. This technique was pioneered by [17] and has been used extensively since [8-9, 33].

2.3.1.2. Physical modelling methods

These models are based on physical and optical phenomena, directly observable in a lens system. Distortion models of this nature are less common than the curve fitting variety.

Conrady's model [26] was the first such model to correctly model radial and tangential distortion, with the latter being the direct result of the non-perfect relative centering of the lenses in the optical assembly. Conrady's model was based on rigorous analytical ray tracing of light through decentered perfect spherical lenses.

A further example of such models is that of Perš and Kovačič [5]. This model is based on the comparison of the fringes of a wide FOV camera's image to that obtained by an idealised pinhole camera which is non-perpendicular to the image plane. A unique feature of this model is that it has no parameters specific to distortion, the radial distortion is a function only of the focal length. The results presented by Perš and Kovačič show a significant improvement in distortion, although the Root Mean Square (RMS) pixel distance from the best fit straight lines through the corrected points are still in the order of 1 pixel in the image centre and 3 pixels at the image edge, for the 768 pixels by 576 pixels image. The focal lengths were a particularly short 6.5mm and 8.5mm, however these are still longer than the camera used in this dissertation (refer to §3.1).

2.3.2. Cartesian lens distortion models

The second broad variety of lens distortion models is the Cartesian class. These models typically do not assume that the distortion is radially symmetrical.

Examples of these methods include those of: Sagawa, et al [28], Nijmeijer, et al [32], and Tsatsakis et al [18]. In this class of model a direct relationship between a pixel in the (distorted) image and its corresponding projection vector is sought. This is done by determining a look up table for a subset of the pixels (typically a grid of every (e.g.) 50 pixels across and down) and then interpolating between these pixels for the rest of the pixels in the image.

Arguably some of the most ingenious work in lens distortion revolves around the exact manner in which the corresponding vectors for the chosen subset of pixels are determined. Sagawa et al placed the camera to be calibrated a known distance in front of a high resolution plasma panel and displayed (alternately) vertical and horizontal bars and used their intersections. By varying the widths and locations of the bars a very dense subset grid was obtained thus minimizing the need for, and error induced by interpolation.

Tsatsakis et al placed the (analogue video) camera-to-be-calibrated in front a line grid whose intersections are known. This grid was mounted on a linear transducer and could move back and forth relative to the camera. So (assuming the camera's roll is aligned to that of the grid) for any given pixel there is a distance such that a horizontal grid row will pass perfectly through it. The same is true for a second distance with a second, different, grid line (provided the grid density and size and transducer range of movement suffice). Thus using the known vertical distance between the grid lines and the two displacements of

the grid relative to the camera, the elevation of the pixel's unit vector can be calculated. Similarly the vertical grid components can be used to determine the azimuth of the pixel's unit vector.

Claus and Fitzgibbon [15] took a different approach. They chose to fit a separate adjustment model for the *X* and *Y* coordinates. The model is a rational function which consists of the full 2-variable 2nd order polynomial ($c_0x^2 + c_1xy + c_2y^2 + c_3x + c_4y + c_5$) of *X* and *Y*, over another such polynomial. The *X* and *Y* adjustments have the same denominator polynomial in common, thus giving a total of 18 parameters. The parameters are determined analytically using epipolar geometry and reprojection as discussed in §2.3.1.1.

An interesting model in this class is Candocia's scale preserving model [10]. Candocia was interested in keeping the size of the distorted image constant after distortion correction, in order to aide combining the images into a larger image. As a starting point he took Brown's model and chose to make the *X* adjustment a function only of the square of the distorted *Y* position, and vice versa, thus changing the model from a radial one to a Cartesian one. Candocia supplied subjective proof of the distortion correction and its superiority for its intended purpose. Candocia did not divulge how the single distortion correction parameter was determined, nor did he provide a quantative measure of the distortion correction - such as the RMS pixel distance of edge lines to the best fit straight line through them.

2.4. Methods of lens distortion parameter determination

There are three points in time when information to characterize lens distortion can be collected: before the picture is taken, while the picture is taken, or after the picture is taken. These are discussed in the following sections.

2.4.1. Prior characterisation – camera calibration

This class can be broken down in to two types: those that require expensive, highly calibrated laboratory setups, and those that require only simple setups.

Multiple examples of the first type are available in literature [1, 3, 6, 7, 13, 16, 25, 28]. All of these methods have 3D reference points whose positions are precisely known (typically) to within 1mm relative to each other. Depending on the method, single or multiple images are captured, and these may or may not be from known camera positions. These methods then use the extrinsic and intrinsic parameters to project each 3D reference point onto the

image plane, and adjust the distortion model's parameters until the points best coincide with the actual image of the points. This obviously requires knowledge of the extrinsic and intrinsic parameters; this must either be determined before hand, or determined simultaneously with distortion parameters.

There are two predominant types of laboratory target setups: planar checkerboards and planar arrays of circles. The 3D reference points are then the checker intersections or circle centres respectively. Both checkerboards and circle arrays can be manufactured to extremely tight tolerances and have their intersections/centres measured very accurately. If non-planar patterns are desired these are made up of multiple planar patterns, such as Heikkila and Silven's corner cube with circles [16]

On the image side, checker intersections are determined by finding (typically using least squares) the best fit second order, or higher, surface as a function of intensity versus X and Y in a window roughly centred around the intersection of interest. The saddle point of the surface (i.e. the point which is both the minimum intensity on the diagonal from white checker to white checker and the maximum intensity form black checker to black checker) is then the intersection pixel position. Lucchese and Mitra [36] show how this may be done and obtain a sub-pixel accurate result. The two dominant methods of circle centre determination are to use the centroid (intensity centre of gravity) or to find the best fit ellipse and use its centre, Redert et al [37] provide an algorithm to do the latter and find the centre to within a less than a pixel.

The quality of calibration obtained using laboratory setups as described above is typically superior to other methods due to the precision of the measurements used, due to the precision of the knowledge of the physical setup. However this precision needs to be balanced against the cost of the calibration and equipment and the acceptable level of distortion for the application.

The second type of prior distortion determination involves the use of less elaborate setups, typically the positions of the references are not known precisely, only their geometry. The majority of these methods cannot determine the camera intrinsic and extrinsic parameters. The quintessential example of these methods is Brown's plumb line method [14] which uses only lengths of cord pulled taught by weights suspended in oil, with pins and the like used to get lines which are non-horizontal. Fernandes et al [12] used a grid which had concentric circles and radii radiating outwards from the circle's centre. Zhang [8, 9] uses a simple printed array of non-connected squares.

All of these methods are based on the "straight lines are straight" curve fitting methodology described in §2.3.1.1. Points along a line are determined either manually from an enlarged print of the image as Brown did in 1971 or via an edge detection algorithm. Fernandes et al used the concentric circles and radii to determine the distortion centre by moving the grid relative to the camera until none of the radiating lines appeared to be distorted, thereafter the grid lines were used to determine the distortion. Zhang used an edge detection algorithm to identify the (collinear) corners of the squares on his calibration pattern. Once the points on the lines have been determined, the distortion parameters (starting with an initial estimate) can be numerically refined by the methods described in §2.6.

2.4.2. Live characterisation – photo calibration

Of the three methods this seems to be the least common, only Lin and Fuh [20] of the articles surveyed used this method. Live determination involves placing sufficient reference data around the object to be photographed so that lens distortion characterisation can be performed using this data. This is inherently unsuitable for many applications (e.g. aerial photography).

Lin and Fuh were concerned with the digital archival of ancient scrolls. They placed the scrolls on a grid of horizontal and vertical lines, such that parts of the grid are visible all around the scroll, and the camera is normal to the grid/scroll plane. A Fast Fourier Transform (FFT) is performed horizontally and vertically on the image parts that have grid lines visible to identify the grid lines. Each identified grid line is then matched to its corresponding grid line in reality. Using the known grid spacing, the distortion from the ideal pin hole camera model for each grid line intersection is determined in X and Y (i.e. this is a Cartesian model). The distortion for the rest of the image is then interpolated using the nearest known (intersection) points. Lin and Fuh did not provide any quantative measure of the distortion correction.

2.4.3. Post characterisation – photo correction

This method is used when an image (or footage) has been captured with an unknown camera and one wants to remove the lens distortion. There are two types of post characterization methods: "straight lines are straight" and multi-image reprojection.

The first method involves the manual or automatic identification of lines in the (distorted) image which are thought to be straight lines in reality. Thereafter, points along these lines

are identified and one the “straight lines are straight” curve fitting distortion measures (refer to §2.3.1.1) is used to minimize the distortion. Clearly this method requires straight lines to be present in the image, so a close up of flower could not be corrected. Depending on the contrast, the identification of points by edge detection methods may suffer significant noise. Finally, the incorrect identification of real-life curves as lines-which-should-be-straight will corrupt the distortion characterization.

Examples of this method are the trio of related papers by Ahmed, Farag and El-Melegy. In 2001, Ahmed and Farag [21] provided an analytic solution based on a distortion measure which is the sum of the difference of the gradient over each line to be straightened (as discussed in §2.3.1.1). This solution is then numerically refined. Melegy and Farag [23] refined this in 2003 by adding the ability to move to higher order distortion models and to eliminate noisy data using the least median of squares technique. In 2005, Ahmed and Farag [22] showed that modelling tangential distortion alleviated the effect of not accurately determining the distortion centre, they further claimed that high order models over fit the data for a particular image, and thus cannot be used for other images with the same camera. This could have been rectified by performing the calibration using multiple images from the same camera.

Cucchiara et al [2] used the Hough transform of the identified points and then used the maximum of the Hough Space as their distortion measure. The Hough space that they used was the standard rho-theta space, although they weighted the bins with the intensity of the point found by their line detection algorithm. They used a course-to-fine search technique to characterise their distortion, a method which worked because they elected to use only a single distortion parameter: first order radial distortion. More parameters would have increased the search space such that this method would not be viable, and a gradient based method would be required. Having only a single radial distortion parameter also allowed for an exact analytical inverse function.

The second type of post characterization is using re-projection from multiple images. In this method feature points, such as the corners of windows, doors, etc, are identified. One or more further images are then required that have same feature points. Additionally the physical relationship between the feature points needs to be constant – so one cannot use this method with images of moving objects. Finally, the intrinsic parameters need to be constant for all the images taken. The repeatability and accuracy with which the feature points can be found is of paramount importance if the distortion is to be accurately

modelled. Epipolar lines are used to determine the relative positions of the cameras, and so the feature points can be projected from one image onto another. The difference between the detected feature point and the projected feature point is the distortion measure.

An example of this type of distortion correction is Claus and Fitzgibbon's rational function model [15], where they fit 18 parameters to perform their Cartesian correction.

2.5. The need for the inverse of distortion

§2.3 and §2.4 above give an overview of how lens distortion is characterized. All of these methods allow one to take a point (pixel) in the distorted image, and determine where that point would be if the image were not distorted.

However, in practise, one generates an undistorted image in the opposite manner. One starts with a blank slate for the undistorted image, and for every pixel determines which distorted pixel to use. In general it is unlikely that the calculated source distorted pixel has integer coordinates, and so some interpolation between the four adjacent pixels is necessary. Typically bilinear interpolation is used, as depicted and explained in Figure 2.3 and Eq 2-2 respectively. For non-grayscale images the interpolation is performed for red, green and blue individually

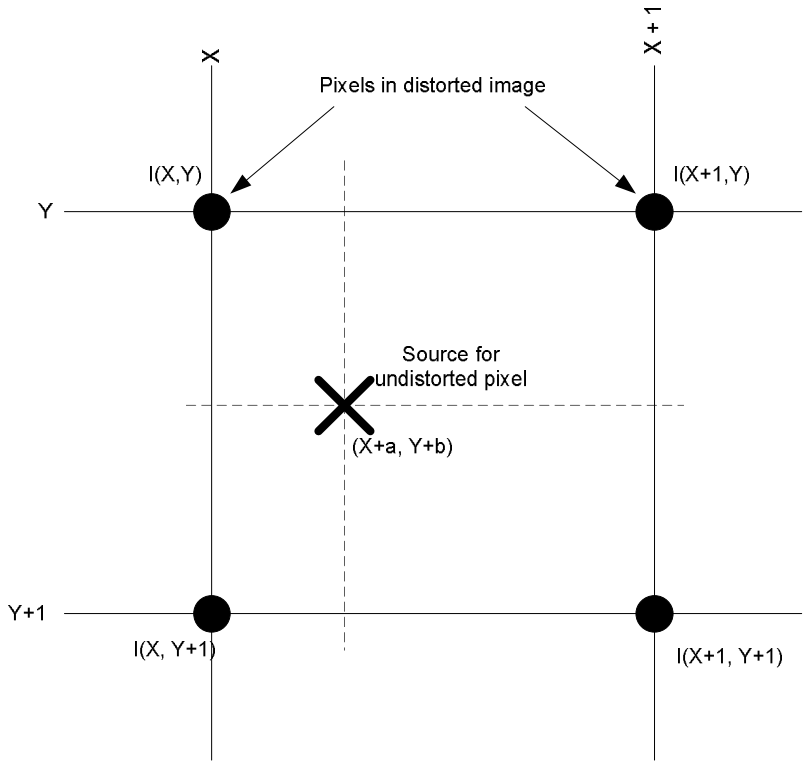

Figure 2.3 Bilinear interpolation data sources. The spot marked X, falls between the integer coordinates at which the pixel values are known, Eq 2-2 shows how to determine the value at X.

$$I(X+a, Y+b) = (1-b)((1-a)I(X,Y)+aI(X+1,Y))+$$
$$b((1-a)I(X,Y+1)+aI(X+1, Y+1))$$

Eq 2-2

where:

$I(m, n)$ = is the intensity of the pixel at (m, n),

a = sub-pixel distance from closest pixel to the left and where $a \in [0,1)$, and

b = sub-pixel distance from closest pixel above and where $b \in [0,1)$.

There is a second reason why the inverse of distortion is needed. It is often desirable to know where in the image an object in free space would appear. This point can be projected onto the image plane via the extrinsic and intrinsic parameters, thus obtaining the undistorted position, and the inverse of the distortion model must be used to obtain the distorted (i.e. image) position. An example of an application that requires this is a tracking system where a moving reference is in the field of view of a camera, one would then like to be able to predict where the reference will be in the next image so that either the area of the image searched for the reference is reduced and/or comparison of the predicted and detected reference points can be done as a believability check.

The ability to find the distorted pixel which corresponds to an undistorted picture is known variously as undistortion, distortion correction and inverse distortion. Inspection of Brown's distortion model in Eq 2-1 shows that except for trivial first order radial implementations with no tangential modelling the equation is nonlinear in terms of r^2. Thus although finding a distorted point's corresponding undistorted point is simple, the opposite is not true. Typically one guesses where the undistorted point is (possibly by calculating the correction of a distorted point at the undistorted point's position, and subtracting instead of adding that correction) and then adjusts that point until the distance between its corresponding undistorted point and that original undistorted point is acceptably close to zero. This involves an iterative numerical refinement of the guessed distorted position. This is neither a quick nor a deterministic procedure, making it undesirable for real-time applications.

Although Candocia's scale preserving model [10] (which (refer to §2.3.2) is a simplified version of Brown's model) has an analytic inverse distortion, it requires solving two 5^{th} order polynomials (one each for X and Y) for each pixel in the undistorted domain. Since a fourth order polynomial is the highest order whose roots can be solved analytically, this too requires an iterative solution to find at least one of the roots.

2.6. Nonlinear multi-dimensional optimisation techniques

Many algorithms have been developed to find the minimum of a given function. The choice of algorithm depends on the exact characteristics of the function such as: whether it is quadratic, the number of parameters, whether it is differentiable, whether the solution must be constrained to a certain space, the accuracy of the initial guess and even whether the function is known in closed form or whether it is a "black-box" (perhaps the result of an intricate and time consuming simulation or experiment). This last also determines whether the derivatives are known or must be numerically approximated.

A representative sample of four techniques suitable to lens distortion is presented here, starting with the earliest and most basic technique and progressing through to the more advanced techniques.

2.6.1. Steepest Descent (SD)

SD is the simplest multi-dimensional non-linear optimisation technique. It is a first order optimisation technique, which means that only the first order derivative information of the

function to be minimised is used. Optimisation theory states that the gradient (∇f - vector of first order partial derivatives) provides the direction of steepest ascent, and therefore the chosen search direction is $-\nabla f$ normalised to a unit vector.

With the search direction ($\boldsymbol{d}_k$) known, the problem is reduced to a single parameter search from the original N parameters, i.e. the problem becomes solve $\alpha>0$ such that $f(\boldsymbol{x}_k+\alpha\boldsymbol{d}_k)$ is a minimum. Various techniques to do this line search are available in any textbook on numerical techniques/optimisation such as that of Snyman [38] and Burden and Faires [39]. The method chosen for this dissertation is Powell's quadratic interpolation. In this method the function is evaluated at the current position, and at the maximum step size taken in the search direction. Halfway between these positions the function is re-evaluated, and if this is not lower than the function at the current position, the distance between that point and the current point is halved until the function is lower, or the point becomes closer than the tolerance to the current point. Assuming that this last does not happen, a quadratic is fitted between the three points, as a function of distance from the current point. The turning point of the function is taken as the next point, although it is clipped to the maximum step size.

The algorithm for SD is given below:

$$\bar{d}_k = \frac{-\nabla f}{\|\nabla f\|_2},$$

Find α such that

$$\begin{aligned} &\left[f\left(\bar{x}_k + a\bar{d}_k\right)\right]_{\min} \\ &\bar{x}_{k+1} = \bar{x}_k + a\bar{d}_k \end{aligned}, \qquad \textbf{Eq 2-3}$$

where:

f = function of N variables to be minimised,

∇f = partial derivative vector of f,

$\boldsymbol{d}_k$ = the search direction of steepest descent,

α = the step in the direction of $\boldsymbol{d}_k$ (which minimises f.), and

k = the iteration number.

It is worth noting that SD will always converge to a local minimum (if it exists), but since the minimum is found for each $\boldsymbol{d}_k$, it follows that each $\boldsymbol{d}_{k+1}$ is perpendicular to $\boldsymbol{d}_k$, i.e. they are orthogonal. Therefore SD 'zig-zags' to the final solution and requires a theoretical infinite amount of iterations to find the minimum for quadratic problems. Thus, to stop the iterative technique, a maximum number of iterations is specified, and the process will

terminate if the norm of the gradient or the step-size (α) drop below specified thresholds or if the maximum number of iterations is exceeded.

The implementation of SD used for this dissertation uses the line search algorithm suggested by [39] which is an adaptation of Powell's quadratic interpolation.

2.6.2. Levenberg-Marquardt Algorithm (LMA)

This is the de-facto standard optimisation technique used for lens distortion parameter determination. LMA was first formulated by Levenberg in 1943 [40], and later independently formulated by Marquardt in 1963 [41]. The LMA is a damped least-squares second order non-linear numerical optimisation technique. Its iterative formula, to solve M equations, (f_i, $i \in [1, M]$) of N variables is

$$\begin{aligned}(J_k^T(\bar{x}_k)J_k(\bar{x}_k) + l_k I)\bar{q} &= -J_k^T(\bar{x}_k)\bar{F}(\bar{x}_k)\\ \bar{x}_{k+1} &= \bar{x}_k + \bar{q}\end{aligned},$$

where:

$J_k(\boldsymbol{x}_k)$ = M by N Jacobian matrix of partial derivatives at $\boldsymbol{x}_k$,

$\boldsymbol{F}_k(\boldsymbol{x}_k)$ = M by 1 vector of values of the M equations at $\boldsymbol{x}_k$, **Eq 2-4**

λ_k = is the damping factor at iteration k and $\lambda_k > 0$,

$\boldsymbol{x}_k$ = the current estimate of the optimal point $\boldsymbol{x}^*$,

k = the iteration number, and

$\boldsymbol{q}$ = step to improve estimate of $\boldsymbol{x}^*$.

By comparison, the classical second order technique is Newton's method for non-linear systems, as per Burden and Faires [39]. The technique's basic formula is:

$$\begin{aligned}H_k(\bar{x}_k)\bar{q} &= -\nabla f(\bar{x}_k)\\ \bar{x}_{k+1} &= \bar{x}_k + \bar{q}\end{aligned},$$

where:

$H_k(\boldsymbol{x}_k)$ = N by N Hessian matrix of second partial derivatives at $\boldsymbol{x}_k$,

$\nabla f(\bar{x}_k)$ = M by 1 vector of partial derivatives of $f(\boldsymbol{x})$at $\boldsymbol{x}_k$, **Eq 2-5**

$\boldsymbol{x}_k$ = the current estimate of the optimal point $\boldsymbol{x}^*$,

k = the iteration number,

$\boldsymbol{q}$ = step to improve estimate of $\boldsymbol{x}^*$, and

$f(\boldsymbol{x})$ = function of N variables to be minimised.

Eq 2-5 is an extension of Newton's method to find the roots of a single variable function, both to higher (N) dimensions, and to find not the zero of the function, but the zero of its derivative – i.e. the minimum (or maximum).

Comparison between Eq 2-4 and Eq 2-5 reveals that the Newton's Hessian matrix has been replaced by the Jacobian, and both sides have been multiplied by the transpose of the Jacobian, this is to find the least squares zero to the M equations. Note that M has to be greater than or equal to N for $J_k^T(\mathbf{x}_k)J_k(\mathbf{x}_k)$ to be non-singular and thus invertible. In addition a multiple of the identity matrix is added to the least squares Hessian estimate. For $\lambda_k = 0$ Eq 2-4 is simply the Newtonian method applied in the least squares sense. However, the larger λ_k, the closer the resultant step direction is to the steepest descent search direction and the smaller the magnitude of that step.

Note if $\lambda_k = 0$, and $f_i(\bar{x}) = \frac{\partial}{\partial x_i} g(\bar{x})$ then M = N and Eq 2-4 and Eq 2-5 are equivalent.

So LMA is an interpolation between the Newtonian method for systems, which has quadratic convergence but is not guaranteed to converge [38], and steepest descent - which has guaranteed convergence but an extremely slow convergence rate due to its orthogonal search directions [38]. The optimality of this interpolation is dependant on λ_k and how it is changed from iteration to iteration.

2.6.3. Fletcher-Reeves (FR) Conjugate Gradient Method

It was noted in the description of SD (§2.6.1) that the direction of steepest descent was used as the direction for the line search to minimize the function *f* of N variables. It was also further stated that the strict adherence to the steepest descent direction, led to each search direction being orthogonal to the previous direction and that this slowed convergence to the final solution.

To prevent this, yet still retain the robustness of first order optimisation techniques, Fletcher and Reeves [42] suggested using vectors which were conjugate with respect to the Hessian matrix. This would allow for the successive search directions to use knowledge gained from previous search directions, resulting in successive search directions not being orthogonal.

To illustrate this, Figure 2.4 portrays the SD search directions for a function with an elliptical minimum, and Figure 2.5 depicts the FR search directions for the same function and starting point.

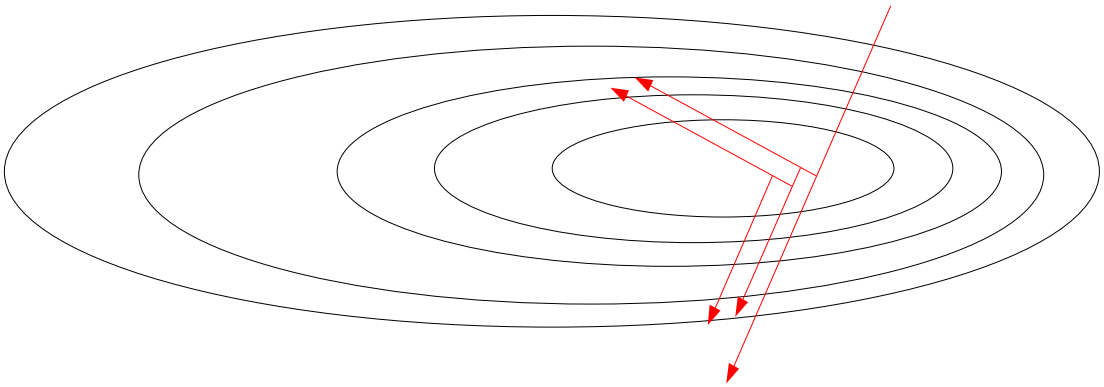

Figure 2.4 Successive SD search directions, depicting the orthogonality of the SD search directions for this simple contour diagram of a function of two variables.

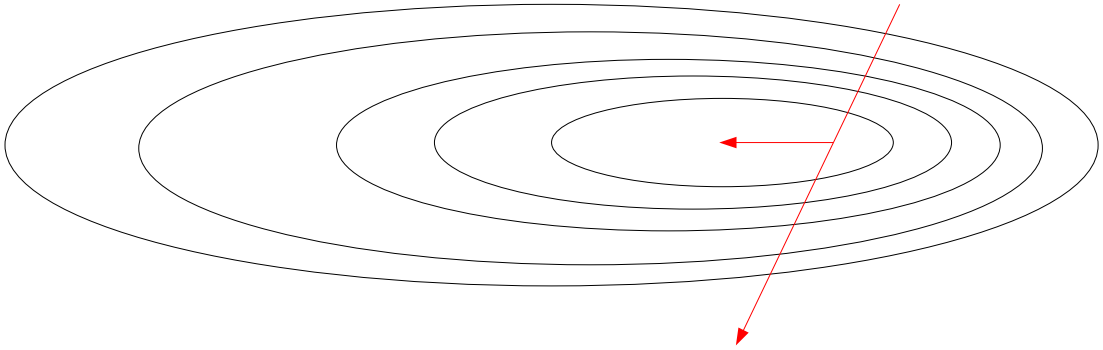

Figure 2.5 Successive FR search directions, depicting the non-orthogonal FR search directions. For quadratic error surfaces, such as this, FR finds the minimum in two iterations.

It is apparent how even the slight eccentricity in Figure 2.4and Figure 2.5 results in zig-zagging for SD, whereas the FR directions provided by

$$\bar{d}_k = \frac{-\nabla f_k}{\|\nabla f_k\|_2} + b_k \bar{d}_{k-1}$$

$$b_k = \begin{cases} \|\nabla f_k\|_2 \,/\, \|\nabla f_{k-1}\|_2 & k \bmod (N+1) \neq 0 \\ 0 & k \bmod (N+1) = 0 \end{cases}$$

Find α such that

$$\begin{matrix} \left[f\left(\bar{x}_k + a\bar{d}_k\right) \right]_{\min} \\ \bar{x}_{k+1} = \bar{x}_k + a\bar{d}_k \end{matrix}, \qquad \textbf{Eq 2-6}$$

where:

f = function of N variables to be minimised,

∇f = partial derivative vector of f,

$\boldsymbol{d}_k$ = the search direction of steepest descent,

α = the step in direction of $\boldsymbol{d}_k$ which minimises f, and

k = the iteration number.

find the minimum in two steps for this simple quadratic case. Fletcher and Reeves [42] show that for a quadratic function of N variables that Eq 3-4 converges in N steps. For non-quadratic functions the FR technique is comparable to second order techniques in terms of convergence, but does not suffer from their sensitivity. The periodic 'forgetting' of the accumulated knowledge evident in Eq 3-4 is to improve the convergence for enharmonic functions.

2.6.4. Leapfrog Dynamic Method for Unconstrained Minimization

The Leapfrog Algorithm is the most recently developed of the algorithms considered. As it uses only gradient information to find the minimum, it is a first order technique.

Figure A. 1 in Appendix A portrays the implementation of the algorithm. It simulates a particle which appears at rest at the provided starting point and is subjected to accelerations induced by the gradient of the function. It worth noting that the acceleration induces a velocity and momentum which allows the particle to go over "humps" and thus finds not the nearest local minimum, but the nearest "low local" minimum.

For more information the interested reader is referred to the original paper [43] and a guide to optimization [38] by the author of the Leapfrog Algorithm, Prof .J.A. Snyman.

2.6.5. Summary of the optimization techniques' attributes

Steepest descent does not require the function to be quadratic or to have a smooth gradient. It is guaranteed to find the closest downhill local optimum if it exists, however its convergence is slow, particularly when the function is sensitive to a subset of its variables.

Fletcher-Reeves assumes that the function is quadratic in nature in the vicinity of the minimum and uses this information to select better search directions to increase convergence to the minimum. If the function is not quadratic the search may veer off, eventually terminating in a different local minimum.

Levenberg-Marquardt also assumes that the function is quadratic, however it is more sensitive to this than is Fletcher-Reeves, as it takes a Newtonian step based on a quadratic best fit of the current area of the function and can thus step completely out the current local minimum's valley particularly if the rate of change of the gradient is steep. Its fall back to steepest descent will only be activated if the step results in a higher function evaluation, which may not happen if another shallow minimum's valley is entered from the top of a deeper (better) minimum's valley. Highly correlated parameters also cause premature convergence.

Leapfrog does not make any quadratic assumptions, and is robust in the presence of noise. Its momentum characteristic allows it to flow over ripples in the function to find a deeper minimum further down the valley, however this makes it prone to over fitting data. Leapfrog's convergence is not as fast as Fletcher-Reeves.

CHAPTER 3.

CONFIGURATION OF EXPERIMENTS

This chapter details the common methods used to capture data to test and verify the hypotheses. It also defines the measures that were used quantify the distortion of an image, which then provide the error which can be minimized with an optimization technique.

3.1. Short focal length, wide FOV camera

Carl Zeiss Optronics (http://www.zeiss.com/optronics) has developed an extremely compact camera with a custom optical lens assembly as part of its optical helmet tracker system. The specifications of this camera pertinent to this dissertation are:

- Nominal Focal Length: 5mm
- Field of View: 78° horizontally
- Image resolution 667 pixels by 502 pixels

In total 32 of these cameras were evaluated. The first 8 have a different lens design to the final 24, although both designs have the nominal characteristics stated above. The pixels of the imager used are square, and the horizontal and vertical axis are orthogonal.

3.2. Calibrated reference point jig

This jig provides circular optical reference points that are densely packed over a large physical area and can be individually activated. For proprietary reasons the details of this jig cannot be divulged except as noted below.

The jig consists of 504 circular optical references arranged into six concentric squares, with the number of references remaining constant per area instead of per concentric square. This provides more data points at the edge of the field of view where the distortion is greatest.

The optical references are not point sources but rather circular references which to alleviates the affect of any aberrations, and justifies their exclusion from this study.

A sample of 32 of the cameras discussed in §3.1 each captured the jig's optical references from four different positions, so that the entire FOV of the camera was sampled.

Figure 3.1 illustrates what the captured data for all four positions looks like, each position's captured data is in a different colour.

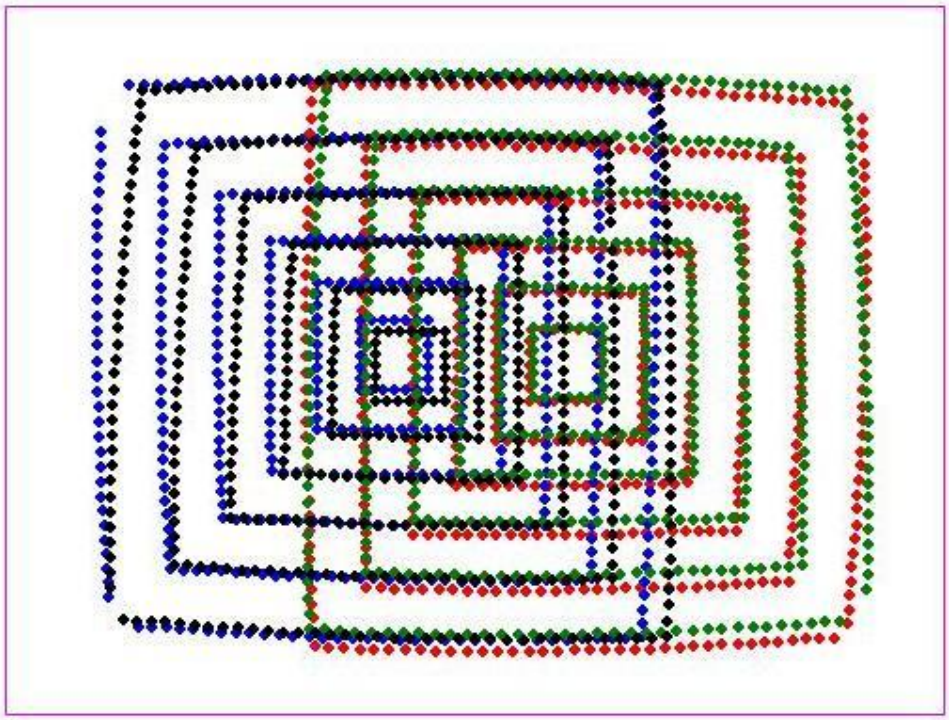

Figure 3.1 Collated image of the calibrated reference point jig as seen from four different positions covering the entire FOV showing the severe distortion. The purple outline represents the camera's FOV.

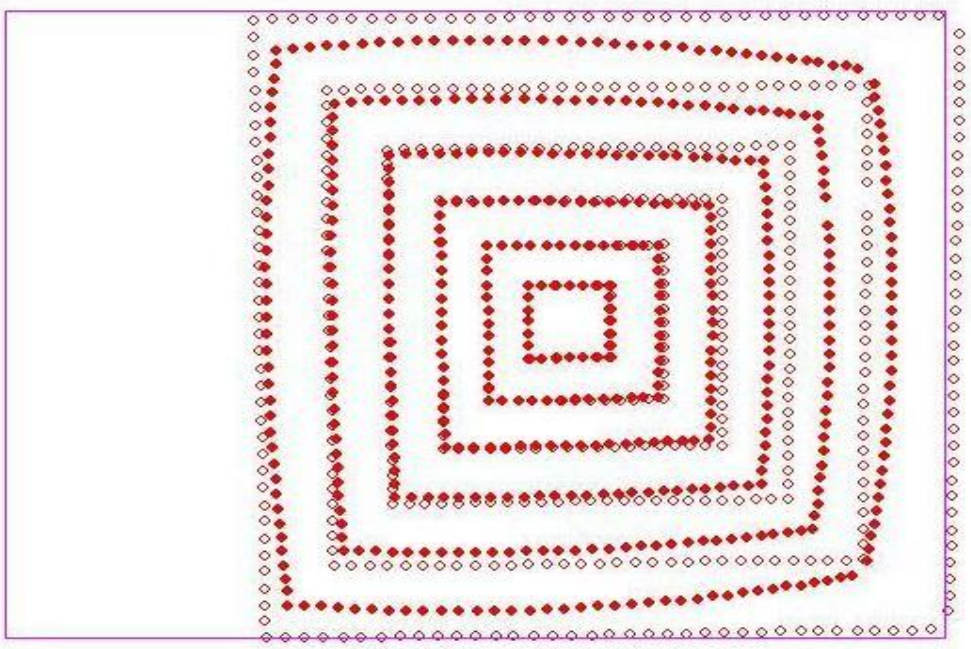

Figure 3.2 A depiction of the distorted and undistorted points of the calibrated reference jig as seen from camera position one. The solid points are the distorted points, and the outlines are undistorted points. The purple outline represents the camera's FOV.

Note the extremely high bowing of what should have been the straight lines of the concentric squares. For reference, the undistorted points (as calculated using a 3rd order radial, 3rd order tangential, variable distortion centre radial model) for the first camera position are plotted with the distorted points in Figure 3.2. Where an N*th* order implies N terms, note that a first order tangential model is therefore not possible. The displacement in the top right is equivalent to 73 pixels, or 8.7% of the diagonal FOV.

3.3. Definition of residual distortion

In order to determine how distorted an image is, one needs a quantative, repeatable measure. This allows one to measure the residual error after performing a lens distortion correction, and refine the parameters of the correction until the distortion is at a minimum.

Both measures below are of the "*straight lines are straight*" variety. This implies that an image of a straight line taken with a distortion-free camera/lens would yield a straight line in the image. Thus any curvature of straight lines is due to lens distortion.

3.3.1. RMS miss distance of data points from a straight line

Brown [14] proposed the plumb-line method in 1971 and it has been used almost universally since. Essentially, given points along a supposed-to-be-straight line, Brown

fitted the least-squares straight line through the points, the error (for that line) was the sum of the perpendicular distances from the line.

For this dissertation a slight adaptation to the above was done whereby the perpendicular distance between the point and the line is still used, but the line is expressed in the more familiar form of *y* as a function of *x*. The following equation,

$$
\begin{aligned}
&vert = c_1 horiz + c_0 \\
&\bar{d} =< 1, c_1 > \\
&\bar{d}_n = \bar{d} / \left\| \bar{d} \right\|_2 \\
&\bar{p}_l =< 0, c_0 > \\
&\bar{p}_p =< x, y > \\
&dist = \left\| (\bar{p}_p - \bar{p}_l) - ((\bar{p}_p - \bar{p}_l) \bullet \bar{d}_n) \right\|_2
\end{aligned}
\quad , \qquad \textbf{Eq 3-1}
$$

where:

c_1 = gradient of the straight line,

c_0 = vertical offset of straight line,

$\boldsymbol{d}$ = 2D direction vector of straight line,

$\boldsymbol{d_n}$ = unit vector in direction of straight line,

$\boldsymbol{p_l}$ = a point on the line, and

$\boldsymbol{p_p}$ = the point off the line to which to calculate the distance.

shows how this can be easily done by converting the equation for a straight line into its geometric equivalent of a vector pointing along the line and a point on the line. Thereafter the vector from the line-point to the point has the projection of itself along the line subtracted from it - resulting in the perpendicular vector from the line to the point. The magnitude of this vector is then clearly the desired perpendicular distance.

Table 3.1 provides the statistics of the distortion measures of the raw data used in this hypothesis for the 32 cameras and reference jig described in §3.1 and §3.2 respectively. Note that the standard deviation of the RMS values serves only to verify that there are two populations of sensors, as both the standard deviation for the first 8 sensors and for the last 24 sensors are significantly smaller than the standard deviations for all 32 sensors.

Table 3.1 RMS perpendicular distance distortion measures for camera population set prior to any distortion correction.

	RMS Distance (pixels)		
	1st 8	Last 24	All
Min	1.662	4.214	1.662
Max	1.952	4.379	4.379
Average	1.768	4.302	3.669
Std Dev	0.105	0.035	1.116

3.3.2. Sum of gradient differences

A second distortion measure was derived so as to provide a different error 'surface' to be minimized - to further test the robustness of the optimisation techniques. Multiple different options were considered; including the minimising the leading coefficients of the best-fit quadratics through the lines, and the sum of gradient differences between the left and right pairs of points for each adjacent trio of points on a line as depicted in Figure 3.3 and explained by Eq 3-2.

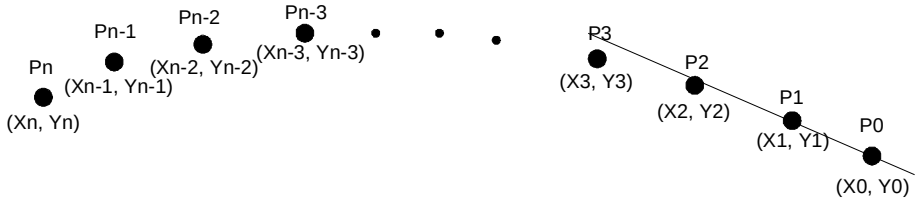

Figure 3.3 Exaggerated diagram showing points on a line and the corresponding coordinates.

$$Distortion = \sum_{lines}\sum_{i=2}^{n} \frac{y_i - y_{i-1}}{x_i - x_{i-1}} - \frac{y_{i-1} - y_{i-2}}{x_{i-1} - x_{i-2}} \qquad \textbf{Eq 3-2}$$

Eq 3-2 is similar to the distortion measure used by [21 – 23], however the sum of quadratic coefficients distortion measure proved to have a very flat error surface, and the adjacent gradient difference sum proved to be very insensitive to small perturbations in the distortion model parameters and also to have many false minimums. This is evident by considering the points on the line to be distorted to lie on a shallow arc (as is typically the case with lens distortion), the change in gradient between any three adjacent point is very small (e.g. P0, P1 and P2 in Figure 3.3) whereas the change in over the entire line is vast, as evidenced by the gradients between points Pn and Pn-1 and points P0 and P1 in Figure 3.3. Additionally Eq 3-2 is very sensitive to near-vertical lines due to the asymptotic

tendency of their gradients to infinity, thus resulting in large distortion measures for lines that are in fact reasonably straight.

To counter this, the gradient between each adjacent pair of points is compared not to the gradient for the immediately prior adjacent pair's gradient, but the gradient calculated from the first two points as evidenced in the following equation. The sensitivity to vertical lines is countered by calculating the gradient of the lines expressed as x as a function of y for lines whose gradient is greater than unity, as evidenced by Eq 3-3.

$$Distortion = \begin{cases} \sum_{lines}\sum_{i=2}^{n}\left(\frac{y_i - y_{i-1}}{x_i - x_{i-1}} - \frac{y_1 - y_o}{x_1 - x_0}\right), & \frac{y_1 - y_o}{x_1 - x_0} \le 1 \\ \sum_{lines}\sum_{i=2}^{n}\left(\frac{x_i - x_{i-1}}{y_i - y_{i-1}} - \frac{x_1 - x_o}{y_1 - y_0}\right), & \frac{y_1 - y_o}{x_1 - x_0} > 1 \end{cases} \quad \textbf{Eq 3-3}$$

Table 3.2 provides the statistics of the distortion measures of the raw data used in this hypothesis for the 32 cameras and reference jig described in §3.1 and §3.2 respectively.

Table 3.2 Uncompensated sum of gradient differences distortion measures for camera population set

	Gradient Difference		
	1st 8	**Last 24**	**All**
Min	54.053	129.536	54.053
Max	68.134	151.089	151.089
Average	59.139	137.562	117.956
Std Dev	4.615	4.816	34.819

CHAPTER 4.

HYPOTHESIS 1: PARAMETER ESTIMATION IS ROBUST WHEN DETERMINED WITH MODERN OPTIMIZATION METHODS

4.1. Hypothesis Formulation

An accepted truism evident in the articles researched is that using higher order versions of Brown's distortion model [14, 27] both yields diminishing returns, and makes the numerical determination of the distortion parameters unstable. There is also a strong contingent stating that tangential distortion is not relevant and that it is unnecessary to determine the exact centre of distortion (the image centre is close enough).

Much of this stems from the Tsai's landmark paper [7] in which he states: "However, my experience shows that for industrial machine vision application[s], only radial distortion needs to be considered, and only one term is needed. Any more elaborate modeling not only would not help but also would cause numerical instability." Quantative proof of this is not provided, nor is the numerical technique used specified, only a suggestion that the distortion parameters can be found "using [a] standard optimization scheme such as steepest descent." This hypothesis will determine if Tsai's assertion is true by testing the stability of high order models when their parameters are determined with modern optimization techniques.

Many researchers have either accepted Tsai's statement as truth or come to a similar conclusion independently. Table 4.1 summarizes the order of radial and tangential distortion modeled and whether the distortion centre is numerically determined or assumed to be either the image centre or principle point. The table also lists whether the justification for the order of the model is due to the numerical instability of higher order models or their perceived lack of benefit. Finally, a mark in the last column indicates that the only justification provided for a low order model is a reference to Tsai.

Table 4.1 Distortion model complexity and justification summary

Reference	Num Radial	Num Tangential	Determine Distortion Centre	Instability	Not Beneficial	Justified only by Tsai ref
Ahmed & Farag [21]	1	1	X		X	X
Ahmed & Farag [22]	1	1	X		X	X
Bacakoglu & Kamel [1]	1	2			X	
Cucchiara, et al [2]	1	0			X	X
Fernandes, et al [12]	1	0			X	
Graf & Hanning [3]	2	0		X	X	X
Hsu, and Hawney [11]	1	0				
Heikkila & Silven [16]	2	2	X			
Karras, et al. [19]	2	0			X	
Mallon & Whelan [31]	2	0			X	
McLean [4]	1	0	X		X	
El-Melegy & Farag [23]	Variable	2		X	X	
Meng & Zhuang [25]	1	0	X		X	
Stein [17]	2	0	X			
Tsai[7]	1	0		X	X	X
Yu, Chung [29]	1	0				
Yu [30]	1	0				
Zhang [9]	2	0		X	X	
Zheng, et al. [13]	2	0			X	X

In 2005 Graf and Hanning [3] showed that different optimisation techniques provided different solutions given the same starting point for lens distortion characterization. In particular, LMA performed poorly.

Many researchers chose LMA to perform their numerical optimisation, a superficially sound decision. The algorithm converges quickly, and is supposed to combine the guaranteed convergence of steepest descent with the rapid convergence of the Newton's method. There are two subtle reasons why one should not use LMA for lens distortion characterization. Newton's method, in using the matrix of 2^{nd} order partial derivatives, implicitly assumes that the error surface is quadratic, and can rapidly diverge if this is not true. Secondly, the algorithm performs poorly when the parameters to be minimized are highly correlated (a fact stated in Marquardt's original derivation [41]), in particular this causes the damping factor to increase drastically, which in turn decreases the step size until it drops below the convergence threshold, yielding a false claim to have achieved convergence.

Few of the papers researched (with the notable eption Brown's [14, 27]) presented details of their numerical optimisation, often stating that (as did Tsai [7]) that any optimisation technique could be used or perhaps stating that LMA or Newtonian iteration was used. Current optimization theory [38] states that each parameter needs to be scaled so that the error's gradient vector's magnitude is equally sensitive to changes in any parameter. Additionally, many advances in numerical techniques have occurred since the Newtonian, Steepest Descent and LMA algorithms were developed, the last of these three (LMA) was developed in 1963.

It is suggested that improving the optimization methods – both via modern optimisation algorithms and proper scaling – will allow for robust and repeatable lens distortion characterization, even of high order models.

4.2. Verification

To test the theory that the determination of radial lens distortion's parameters is stable and robust when modern numerical optimisation techniques are used, four different numerical optimisation techniques were implemented.

The first technique is the simplest and original technique: steepest descent. Refer to §2.6.1 or handbooks on numerical techniques [39] or optimisation [38] for details of the algorithm.

The second algorithm implemented is the Levenberg-Marquardt algorithm, a popular choice for numerical determination of lens distortion parameters. Refer to §2.6.2 or the original papers by Levenberg [40] and Marquardt [41] for details on the algorithm. In the implementation of the algorithm, the algorithm is aborted if the damping factor exceeds 10^{300}, which is approaching the limits of 64bit floating-point arithmetic computation.

In keeping with the chronological order of the algorithms, the third is the Fletcher-Reeves conjugate gradient algorithm. Refer to §2.6.3 or the original paper by Fletcher and Reeves [42]. Most handbooks on numerical optimisation (e.g. [38]) will give a brief description, although interestingly may not mention the periodic "forgetting" of the gradient information that vastly enhanced FR's performance for lens distortion parameter determination. The line search technique used for FR is identical to that described for SD.

The final, and most recently developed, algorithm implemented is the Leap Frog algorithm, so named after the integration technique it uses. Despite being a gradient technique, this algorithm performs no explicit line searches. For more details refer to §2.6.4 or Snyman's paper [43].

All the algorithms require gradient information, and LMA requires second order gradient information too. In order to keep the numerical optimisation algorithms generic, they receive only a pointer to the error function to be minimised (one of the line straightness measures of §3.3.1 and §3.3.2) and thus all gradient information must be estimated numerically. The chosen method to do this is the so called 'central gradient.' This entails evaluating the function at the point of interest plus a small positive delta in the variable of interest, and then evaluating the function at the point of interest minus the same small delta. The delta used to numerically determine the gradient was 10^{-8}. The gradient is then the difference between the two evaluations divided by twice the delta. Second order derivatives entail four function evaluations, and a bilinear interpolation of their results. As the number of function calls required will be monitored, and in the interest of fair-play, the symmetric nature of the second order partial derivative matrix is exploited to minimize the number of function calls made by LMA.

In terms of convergence criteria, the four optimization techniques were all given the following criteria, where the step size is in scaled units as per Table 4.3 and the gradient is that of the error surface being modelled:

- Maximum number of iterations: 1000
- Minimum step size: 10^{-12}.
- Target gradient norm: 10^{-6}.

In total, seven different radial distortion models were tested, covering the entire range of those found in the literature study (§Chapter 2) as well as the addition of a few models of even higher order. Table 4.2 summarises the models implemented and indicates to which reference articles they are similar. The model naming convention also incorporates its basic parameters: a number after the R specifies the number of radial terms if greater than one, the number of tangential terms (if any) are indicated by a P and the number of terms, and DC indicates that the optimal distortion centre is found.

Table 4.2 Summary of the radial distortion models implemented, indicating their similarity to models found in literature.

Num	Name	Num Radial	Num Tangential	Distortion Centre	Similar to
1	R	1	0	N	[2], [7], [11], [12], [20], [29], [30]
2	RDC	1	0	Y	[4]
3	R2DC	2	0	Y	[17], [31]
4	R3P2	3	2	N	None
5	R3P2DC	3	2	Y	None
6	R3P3DC	3	3	Y	[14]
7	R5	5	0	N	None

In order to test the robustness of the algorithms, the starting point was chosen to represent no distortion. This is very far from the truth as indicated by Figure 3.2. This starting point

tests the ability and speed of the algorithms to converge from a distant starting point, and also allows for the possibility of false local minimums between the starting point and the global optimum. The distortion centre starting point is not the expected (333.5; 251) due to the blanking of a few rows and columns of the imager. Table 4.3 lists the starting points for each distortion parameter as defined by Brown's model [14, 27] given in Eq 2-1.

As mentioned in §4.1, scaling of variables is important to successful numerical optimisation. Essentially the aim of scaling is to make the error function (in the vicinity of the minimum) equally sensitive to small delta's in any variable, i.e. the gradient vector's elements should be of similar magnitude. This will ensure that the direction of each step is due to the shape of the function and not due to a misleading (un)sensitivity to one or more parameters. The correct method would be to obtain a particularly good starting point for a camera that is representative of the whole sample, and use the gradients at that point as the scale factors. However, as the purpose of this hypothesis is to test the robustness of the parameter determination, such thorough efforts were not undertaken. Instead, the scaling factors were manually adjusted until the all the models converged for a sample camera and the gradients appeared to be roughly of the same order of magnitude. The same scaling factors were used for all the algorithms. Table 4.3 lists the scale factors used.

Table 4.3 Parameter Starting points and scale factors

Parameter	Starting Point	Scale Factor
K_1	0.0	10^{-6}
K_2	0.0	10^{-12}
K_3	0.0	10^{-18}
K_4	0.0	10^{-24}
K_5	0.0	10^{-30}
P_1	0.0	10^{-10}
P_2	0.0	10^{-10}
P_3	0.0	10^{-2}
x_c	336.0	10^{0}
y_c	247.0	10^{0}

Finally, a discussion on the usage of the input data is pertinent. An explanation of the reference jig is provided in §3.2 which explains that for each camera all the visible reference points on the jig were captured for four different positions such that data over the entire FOV was captured. Five optimisations were performed for each camera/distortion model/error function combination. With reference to Figure 3.1 the data for the five optimisations are:

1. Red and blue data points,
2. Red and black data points,
3. Green and blue data points,
4. Green and black data points,
5. Red, green, blue and black data points.

The above combinations were chosen so that each optimisation input combination contains data in both the left and right half of the FOV, so as not to converge to a local minimum applicable for only one half of the image.

4.3. Analysis of results

As discussed in §4.2, 5 optimisations per camera/distortion model/error function combination were performed for a total of 1280 optimisations. Each optimisation yielded the optimisation success/error value, optimised parameters (between one and eight depending on the distortion model), optimised error value, and number of iterations required. Clearly the raw results are far too extensive to be explicitly provided in this dissertation (although they are available upon request) and so only the analysis of the raw results is provided herein.

For each camera/distortion model/error model combination, the standard deviation of each parameter over the five optimisations (refer to §4.2) was calculated. This standard deviation is then averaged over the 32 cameras and all the parameters for the distortion model to yield a single relative value - which indicates how sensitive the optimisation to fit parameters for a particular distortion model is to variances in the input data when using that particular optimisation function and error function. The smaller this sensitivity measure is, the less sensitive the model/numerical technique combination is to input perturbations.

The average optimised error value for each distortion model/optimisation value was also calculated. This information is given in Table 4.4 for the distortion measure provided in §3.3.1 and in Table 4.6 for the distortion measure provided in §3.3.2. The a priori error for Table 4.4 is 3.669 pixels RMS and 117.956 for Table 4.6. This same data is given pictorially in Figure 4.1 and Figure 4.2 for Table 4.4 and in Figure 4.3 and Figure 4.4 for Table 4.6. The smaller the average error is, the better the distortion characterization is. Finally, a summary of the time taken (to the nearest second) to perform all the optimisations for each distortion model/optimisation method is provided in Table 4.5 and Table 4.7 for the distortion measures of §3.3.1 and §3.3.2 respectively.

Table 4.4 Summary of optimisation efficiency and repeatability per distortion model and optimisation method for the RMS perpendicular distance distortion measure.

Distortion Model	Steepest Descent		Levenberg Marquardt		Fletcher Reeves		Leapfrog	
	Ave Error (Pixels)	Sensitivity	Ave Error (Pixels)	Sensitivity	Ave Error (Pixels)	Sensitivity	Ave Error (Pixels)	Sensitivity
R	0.2810	0.010	1.8593	0.79	0.2810	0.010	0.2810	0.010
RDC	0.2774	0.80	2.2858	1.4	0.2809	0.49	0.2659	1.4
R2DC	0.0764	0.10	2.3670	1.2	0.0802	0.084	0.0719	0.27
R3P2	0.0770	0.041	2.2159	0.75	0.0758	0.066	0.0742	0.097
R3P2DC	0.0753	0.065	2.5708	0.71	0.0789	0.052	0.0710	0.21
R3P3DC	0.0739	0.069	2.7033	0.63	0.0750	0.056	0.0703	0.26
R5	0.0771	0.042	1.2676	0.91	0.0917	0.097	0.0742	0.12

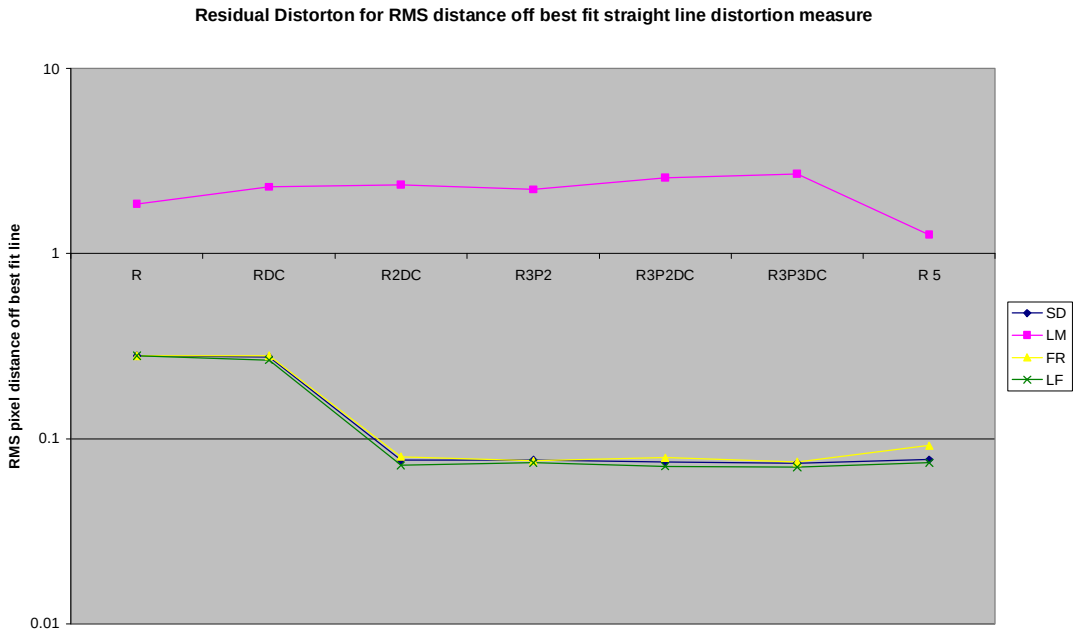

Figure 4.1 Average error for each distortion model as determined by each optimisation method tested for the RMS perpendicular distance distortion measure.

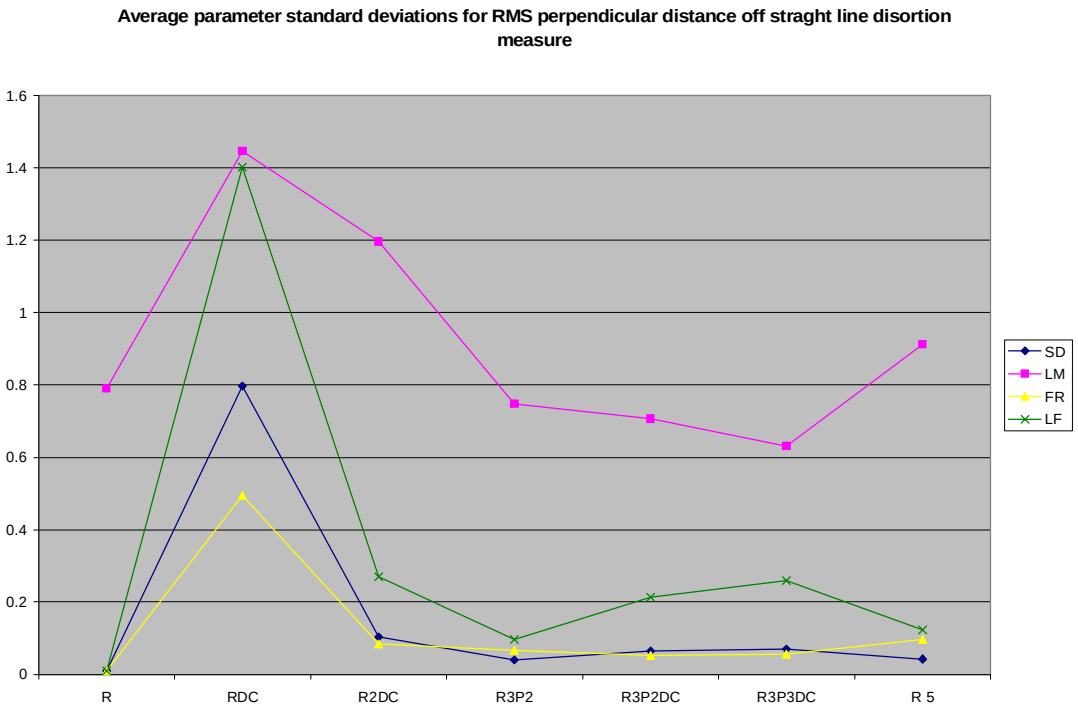

Figure 4.2 Average sensitivity measured for each optimisation method tested for the RMS perpendicular distortion measure.

Table 4.5 Time taken to perform optimisations, per distortion model and optimisation method for RMS perpendicular distance distortion measure.

Distortion Model	Steepest Descent	Levenberg Marquardt	Fletcher Reeves	LeapFrog	Total per Model
R	4	3	5	4	16
RDC	59	7	13	125	204
R2DC	304	6	127	154	591
R3P2	355	8	176	214	753
R3P2DC	444	17	224	339	1024
R3P3DC	490	19	244	402	1155
R5	347	11	163	207	728
Total per Optimisation	2003	71	952	1445	4471

Table 4.6 Summary of optimisation efficiency and repeatability per distortion model and optimisation method for the sum of gradient differences distortion measure.

Distortion Model	Steepest Descent		Levenberg Marquardt		Fletcher Reeves		Leapfrog	
	Ave Error (pixels)	Sensitivity	Ave Error (pixels)	Sensitivity	Ave Error (pixels)	Sensitivity	Ave Error (pixels)	Sensitivity
R	19.747	0.018	57.725	0.073	19.747	0.018	19.747	0.018
RDC	19.739	0.053	57.642	0.072	19.745	0.072	19.627	0.68
R2DC	12.644	0.24	57.670	0.20	12.767	0.15	12.457	0.63
R3P2	12.790	0.090	57.725	0.058	12.801	0.098	12.624	0.48
R3P2DC	12.605	0.17	57.653	0.15	12.734	0.12	12.247	0.60
R3P3DC	12.410	0.19	57.670	0.13	12.553	0.15	12.047	0.49
R5	12.789	0.093	57.724	0.054	12.800	0.11	12.590	0.66

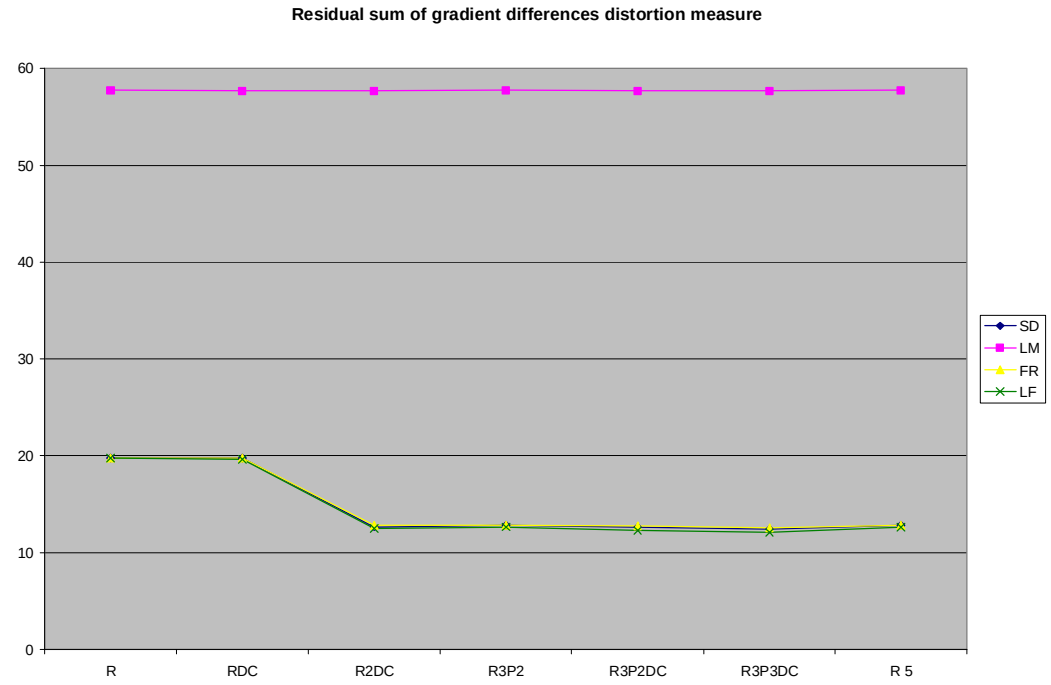

Figure 4.3 Average error for each distortion model as determined by each optimisation method tested for the sum of gradient differences distortion measure.

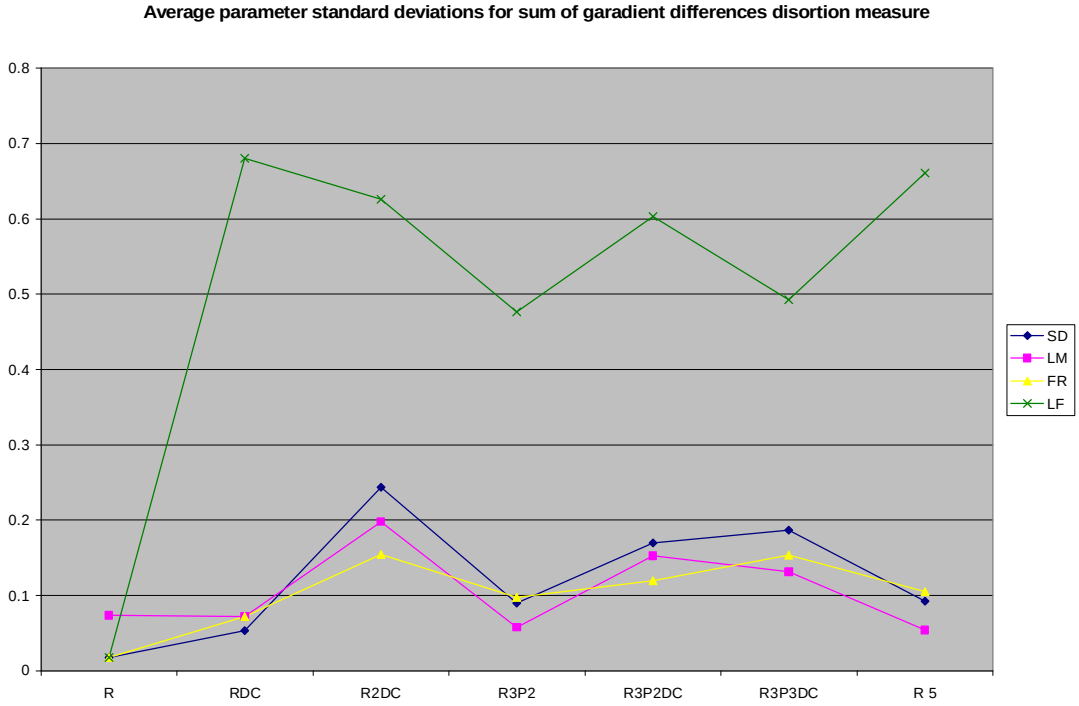

Figure 4.4 Average sensitivity measured for each optimisation method tested for the sum of gradient differences distortion measure.

Table 4.7 Time taken to perform optimisations, per distortion model and optimisation method for the sum of gradient differences distortion measure.

Distortion Model	Steepest Descent	Levenberg Marquardt	Fletcher Reeves	LeapFrog	Total per Model
R	3	1	4	9	17
RDC	35	3	5	167	210
R2DC	248	5	38	229	520
R3P2	130	7	43	321	501
R3P2DC	400	11	58	449	918
R3P3DC	421	13	76	527	1037
R5	127	7	41	310	485
Total per Optimisation	1364	47	265	2012	3688

The average error columns of Table 4.4 and Table 4.6 provide a measure for how well the model characterizes the distortion evident in the cameras, lower values are better. Similarly the parameter standard deviation columns provide a measure of how much the parameters of the models varied when calculated using different combinations of the input data. Lower standard deviation values indicate the model/method combination is less sensitive to changes in input data (including noise).

Evaluation of Table 4.4 and Table 4.6 (and their graphs if so desired) reveal the following:

- LMA was by the far the worst performing algorithm (a point confirmed by [3]). This is due to its sensitivity to fitting models with highly correlated parameters. As the distortion model complexity increased, it produced poorer distortion characterization for the RMS perpendicular distance measure, and similar distortion characterizations for the sum of gradient differences distortion. LMA is excluded from consideration in the remarks below. When LMA converged, it met the convergence criteria of either step size, or gradient. The magnitude of the damping factor was not recorded.
- The first radial distortion coefficient dominates the distortion characterization, as is evidenced by contrasting the first and second rows of Table 4.4 and Table 4.6 with Table 3.1and Table 3.2.

- Finding the optimum distortion centre improves the distortion characterization. This is evident by comparing the R and RDC results and R3P2 and R3P2DC results. This is also the most sensitive parameter to different input combinations, an observation which is both evident in the average parameter standard deviations for models which consider the distortion centre, and intuitive since the distortion centre will affect all the other distortion parameters.
- Adding a second radial distortion coefficient (R2DC compared with RDC) drastically improves the distortion characterisation by approximately 72% and 35% for Table 4.4 and Table 4.6 respectively (unless, as noted above, the LMA is used). For the RMS perpendicular distance measure, it also improves the parameter standard deviations.
- Increasing the number of radial distortion coefficients is less effective than compensating for the distortion centre (by comparison of the R5 and R2DC models' average errors).
- Modelling tangential distortion does improve the distortion characterization (i.e. the average error is less), this is evident both in the better characterisation of R3P2 vs R5, and R3P3DC vs R3P2DC. This improvement is in the order of 1-6% and 1.5-3% for the RMS distance and gradient difference error functions respectively.

In conjunction with Table 4.5 and Table 4.7 the following conclusions can also be made.

- The distortion characterization does not become noticeably less robust with model complexity, as evident by the marginal variances in parameter standard deviations (with the exception of the RDC model), although the time to fit the parameters increases with model complexity. This is attributable to the modern optimization methods applied in this study.
- Steepest descent and Fletcher-Reeves have similar results, except that Fletcher Reeves converges two- to five-fold quicker depending on the distortion measure. In fact steepest descent often failed to converge within the maximum (1000) iterations, as expected due to the slow convergence from its orthogonal search directions.
- The Leapfrog algorithm in all cases obtained better distortion characterizations than the other methods (except in the simplest R model where it produced identical results to steepest descent and Fletcher-Reeves) albeit at the expense of processing

time similar to steepest descent. However its parameter standard deviations are noticeably worse than steepest descent's and (especially) Fletcher-Reeves', this is because of Leapfrog's ability to find the nearest low-local minimum, essentially finding a better optimum for which ever data was used. Leapfrog is thus trading an increased susceptibility to input perturbations for a superior characterization of the input data it is supllied.

- It is evident that the improvement in characterization with increased model complexity over and above R2DC is not significant for the lenses tested. These models are thus only desirable for high accuracy applications where it is useful to model and remove every systematic error possible.
- The accuracy to which each parameter is known was not explicitly determined, however it can be stated that for each additional parameter tested in the seven distortion models the characterization improved. Thus it can be said that distortion model was sensitive to the additional parameter and it is that it was determined to within an order of the smallest allowable step size. A first order indication of the sensitivity of the parameters can be obtained from their scaling values in Table 4.3.
- Despite the fact the characterization data that was captured consists of points on predominantly vertical or horizontal lines (this is a characteristic of the reference jig (§3.2)), the characterization will straighten lines in the image that are at any angle. This is because no use of the verticality or horizontalness of the lines is made, the best fit straight lines through the points are of generic form. Figure 4.5 depicts both the distorted points and undistorted points for captured data of the reference jig taken with a camera with approximately 45°. The distortion characterization was performed using the normal set of four captured reference jig centroids, and the resultant parameters used on the centroids of the rolled image. The R3P3DC distortion model was used.

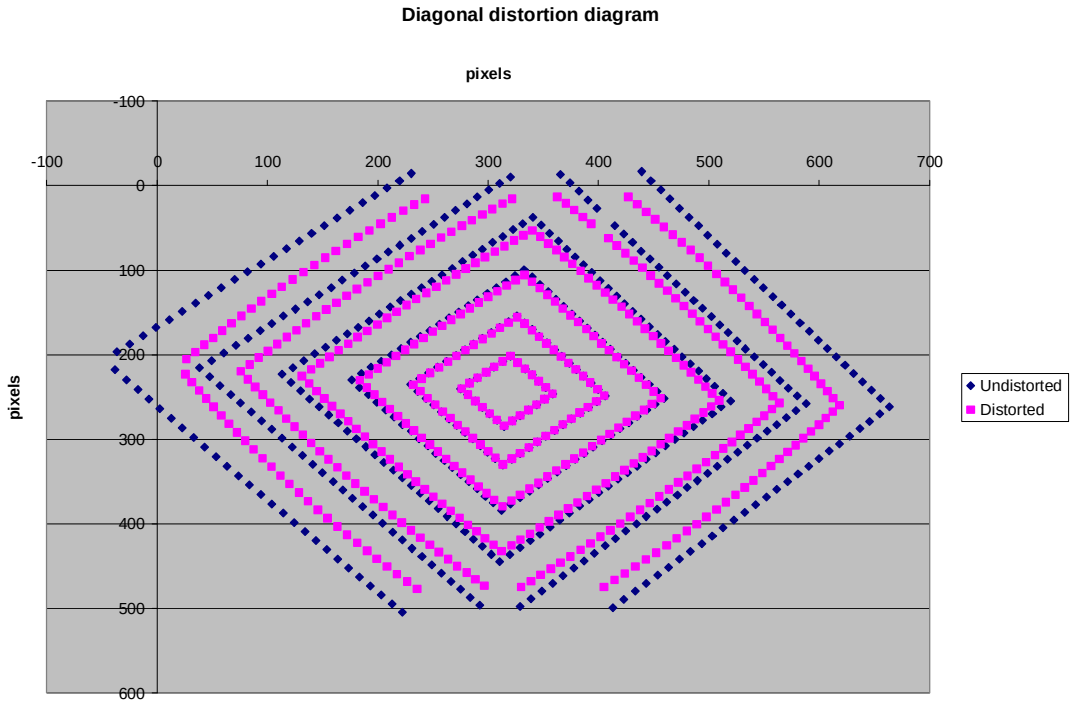

Figure 4.5 Corrected image of reference jig taken with a rolled camera

CHAPTER 5.

HYPOTHESIS 2: CORRECTION OF BARREL DISTORTION CAN BE MODELLED AS PINCUSHION DISTORTION

5.1. Hypothesis formulation

The very nature of lens distortion correction yields a model to map distorted image points to their corresponding undistorted points in the image plane. §2.5 explains the need for inverse distortion, which maps undistorted image points back to their distorted positions.

To undistort an image (once its distortion parameters have been modelled) one must determine, for every pixel in the undistorted image, what the corresponding distorted pixel is. There are many methods of doing this:

1. The most obvious is also the most computationally intensive. This is to guess the distorted point, and iteratively refine it by using the fitted distortion function to see how closely each iteration's distorted position's corresponding undistorted point coincides with the desired undistorted point. Although this produces accurate results (with sufficient iterations) it is a processing and time demanding task and thus may not be suitable for real time applications.
2. A second method is to pre-calculate and store a look up table containing each undistorted pixel's corresponding distorted pixel. This table would be populated as described in the previous paragraph. However it requires a vast amount of memory as each distorted pixel will be a non-integer value. This may not always be possible in an embedded real-time application.
3. The final method found in literature is to create a custom inverse distortion model, and analytically determine its parameters from those of the distortion model. This is typified by [31], where despite only using a second order radial distortion model with no tangential parameters, the resulting undistortion model has 6 parameters! Although these methods require neither extensive memory nor extensive processing, they are not as accurate. [31] obtained a mean accuracy of 0.42 and 0.32 pixels for a 6mm and 8mm lens respectively.

The Brown lens model can correct both barrel and pincushion distortion. Barrel distortion pulls image points closer to the distortion centre, and pincushion distortion pushes image points further out. It is thus suggested that using the Brown lens model the undistortion of a barrel distorted image can be modelled as pincushion distortion, optimally combining the accuracy of the numerically refined method with the processing and memory efficiency of the inverse distortion method. This will yield a model requiring only multiplications and summations, with eight (the maximum tested) or fewer predetermined parameters, and goal accuracy in the order of a tenth of a pixel.

5.2. Verification

Undistortion requires mapping undistorted image points to their corresponding distorted points. This requires knowledge of the undistorted image points' coordinates - which are not generally known. An expensive 3D setup may be made whereby the position of reference points are known precisely relative to the focal point of the lens being calibrated. If the extrinsic and intrinsic camera parameters are known then the points can be projected into the image plane yielding the undistorted image points. Comparison with distorted image points acquired by taking an image through the lens may then take place.

However this requires an expensive outlay, where the calibration of lens distortion is moving to simpler equipment such as checker boards. The proposed solution entails the reuse of the lens distortion data. All lens distortion characterizations create a set of supposedly undistorted image points from the distorted points. These points are then evaluated to obtain a measure of the residual distortion (eg §3.3.1 and §3.3.2), which is used to refine the distortion models parameters. The end result is a set of points which are deemed sufficiently undistorted for the application under consideration, together with their corresponding distorted points.

A simple measure of the efficiency of the undistortion of an image is the sum of the square of the distances between the distorted positions calculated from the undistorted points and the true distorted points. This is represented by Eq 5-1:

$$Distortion = \sum_i \left\| \langle x_{di}, y_{di} \rangle - f(x_{ui}, y_{ui}) \right\|_2 \Big],$$

where:

(x_{di}, y_{di}) = i[th] distorted point,

(x_{ui}, y_{ui}) = i[th] undistorted point, and

$f()$ = inverse distortion function returning the distorted coordinate as a two element vector.

Eq 5-1

The distorted and undistorted points used for the verification of this hypothesis will be those resulting from the verification of hypothesis 1. For each of the 32 sample cameras, the 7 distortion models of Table 4.3 were fitted using 4 different numerical techniques, and using 2 different distortion measures. For each of the camera/model/technique/measure combinations the resultant undistorted and distorted points were output. If one only uses the data from the technique that yielded the best distortion characterization, this leaves 448 camera/model/measure sample sets of approximately 2000 corresponding undistorted and distorted image point pairs. Further electing to only use the data calculated using the prevailing distortion measure defined in §3.3.1 reduces this to 224 data sets. This data will be used to test this hypothesis.

For each camera/model/measure combination, the same distortion model that was used to characterize the distortion will be used to characterize the undistortion. Thereafter the distortion model which best characterizes the distortion across all 32 cameras, will be selected to re-characterize the inverse distortion of all 224 camera/model combinations.

These results will allow us to determine how well each distortion model is able to characterize its own inverse distortion, and will also show the accuracy of which the best inverse characterization of a set of distorted and undistorted is capable.

5.3. Analysis of results

In the analysis LF was used to determine the optimal parameters for the undistortion model. Iteration continued for a maximum of 10000 iterations or until convergence was achieved either via obtaining a gradient with magnitude less than 10^{-6} or a step size of less than 10^{-6}. The number of iterations may seem excessive however it was required as the maximum step size was constrained to 10^{-2}.

Figure 5.1 depicts the accuracy to which an undistorted pixel's corresponding distorted pixel was determined, when the same distortion model that was used to characterize the

distortion was used to characterize the undistortion. The two different types of lenses are clearly evident, with the first 8 cameras performing better and worse for simpler and more complex distortion models respectively. It is perhaps surprising that R3P3DC (refer to Table 4.2 for distortion notation definition), which characterized distortion the best (refer to §4.3), yielded worse results than the other tangential models for the last 24 sensors. The reason for this is because R3P3DC often failed to converge within the maximum number of iterations (it did however not diverge it typically was either oscillating or converging slowly when computation ceased), as shown by Table 5.1. The reason for this lack of convergence is due primarily to a casually chosen starting point and to the decreased magnitudes of the elements of the gradient vector required to achieve a specified magnitude as the number of elements increases. Specifically, a gradient vector with one element requires this element to be 10^{-6} in order to have a magnitude of 10^{-6}, a two element vector requires the elements be 1.41 times smaller or 0.707 x 10^{-6} on average to achieve this magnitude, the eight element R3P3DC error gradient requires each element be 2.90 times smaller, or about 0.345 x 10^{-6} on average.

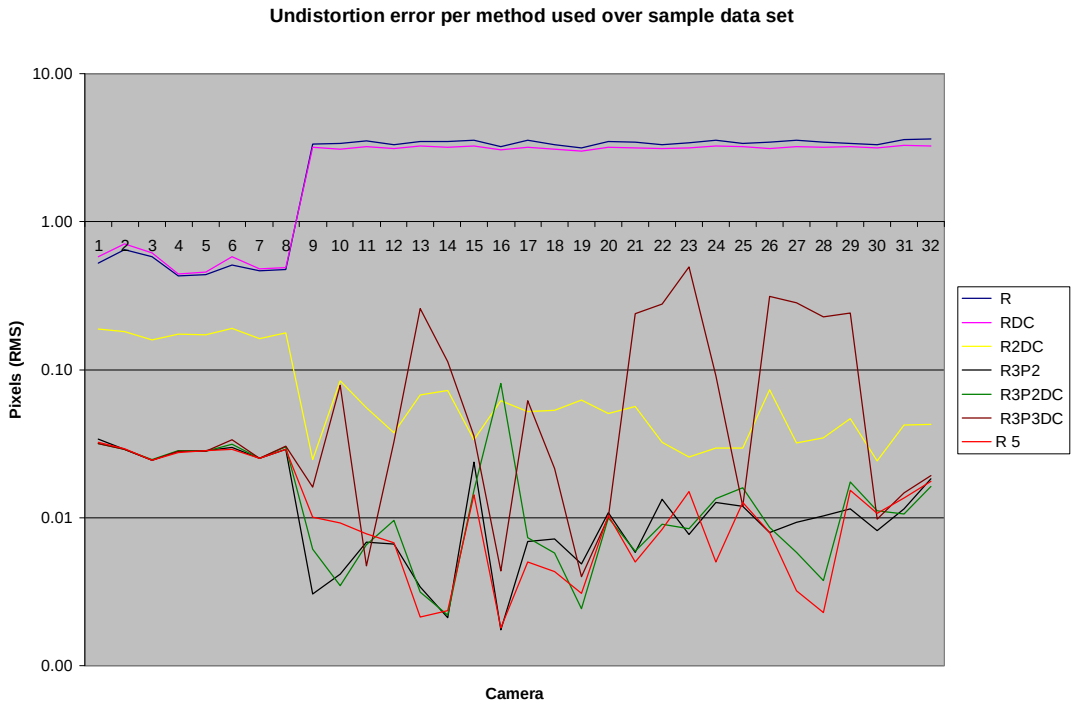

Figure 5.1 Plot of RMS undistortion error, when the model used to undistort the image is the same as the model used to characterize the distortion

Table 5.1 Number of times convergence was not achieved within 10000 iterations over the 32 camera sample set

R	RDC	R2DC	R3P2	R3P2DC	R3P3DC	R 5
0	0	0	0	2	11	8

Despite this, it is clear that accuracies far exceeding those of 0.42 and 0.32 pixels reported in [31] have been obtained, when models with at least two radial coefficients were used. This is shown in Figure 5.2.

R3P2, which best modelled its own undistortion, was then used to model the undistortion for all the data as generated by the 7 distortion models of Table 4.3. Figure 5.3 depicts the results. In comparison with Figure 5.1 and Figure 5.2 it is evident that the R3P2 better models the undistortion for the simpler distortion models than they do themselves. It is on par with itself and the similarly complex R3P2DC model, however it is worse for models of higher complexity, markedly so for R3P3DC. The improvements obtained when using R3P2 are summarized in Figure 5.4.

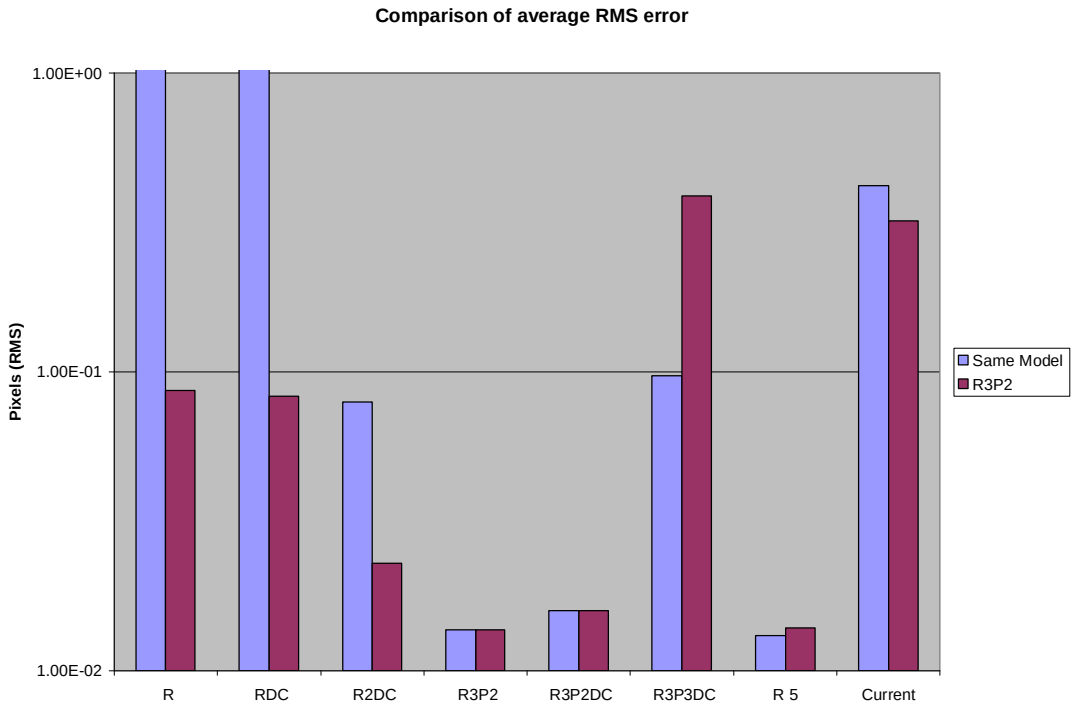

Figure 5.2 Comparison of average RMS values when using the same method for undistortion as was used to characterize the distortion, to using R3P2. The values 0.42 and 0.32 are also plotted for reference purposes.

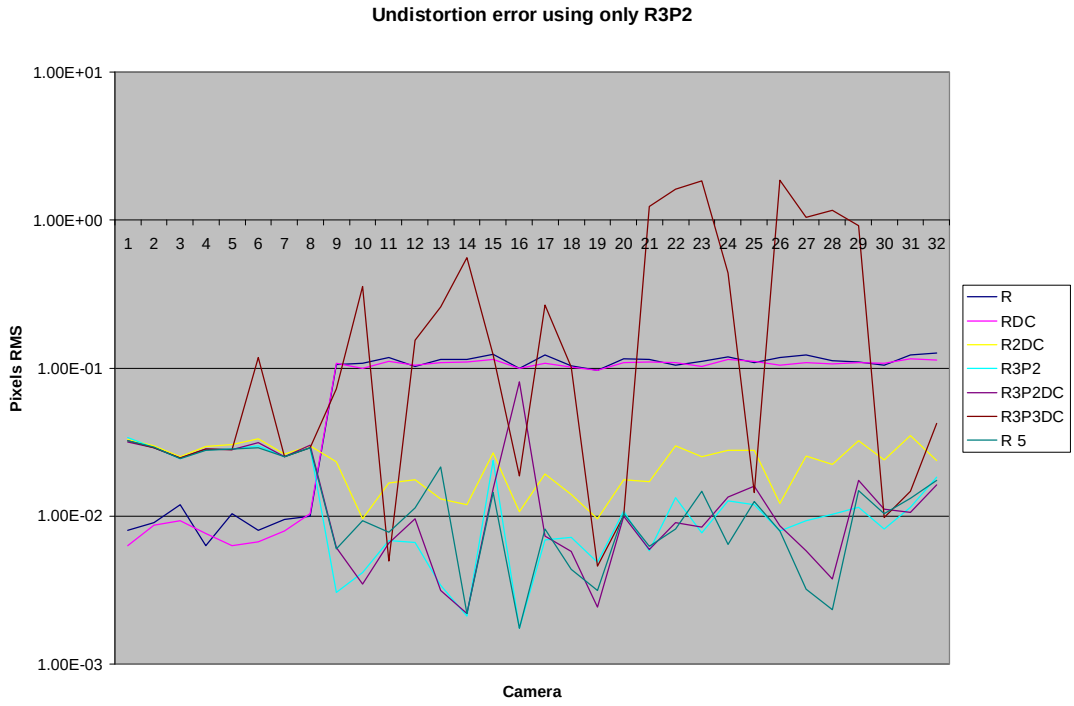

Figure 5.3 Plot of RMS undistortion error, when the model used to undistort the image is R3P2 instead of the model used to characterize the distortion

This is partly due to the fact that convergence was not always achieved, particularly due when R3P3DC was modelled, as shown in Table 5.2. An additional factor is that distortion characterization inherent in the R3P3DC model's data for the 32 cameras is of a more complex nature - as evident by R3P3DC's superior performance over R3P2 in distortion characterization (refer to §4.3) – than R3P2 can sufficiently model.

Table 5.2 Number of times convergence was not achieved within 10000 iterations over the 32 camera sample set

R	RDC	R2DC	R3P2	R3P2DC	R3P3DC	R 5
8	8	0	0	2	20	0

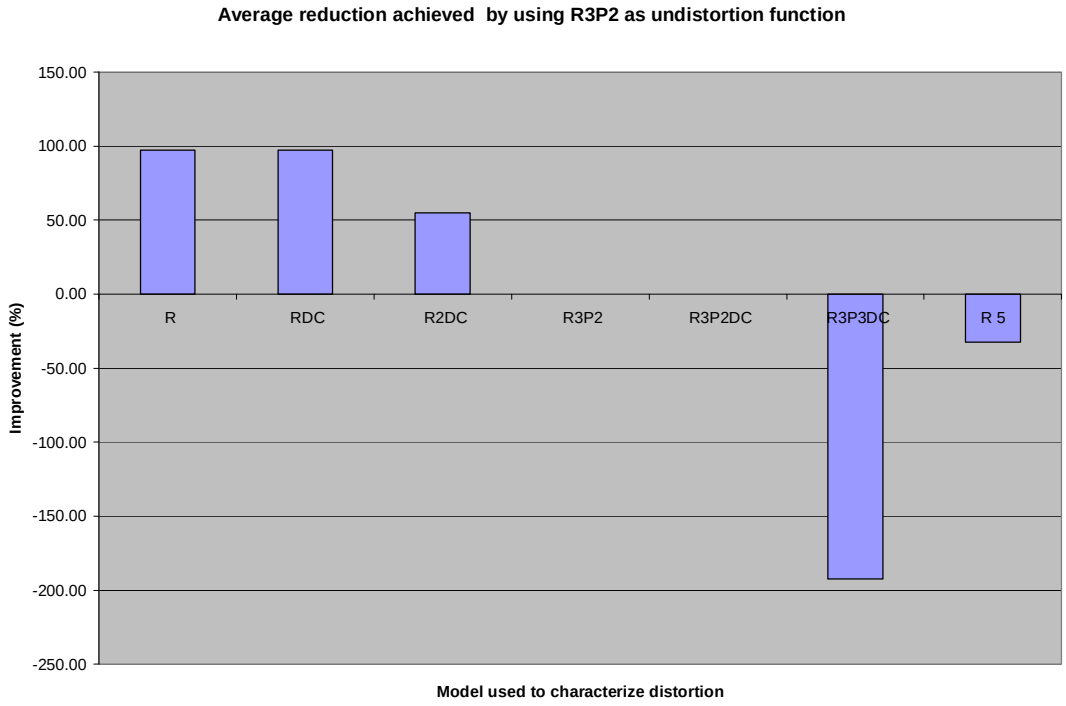

Figure 5.4 Average improvement in undistortion when using R3P2 instead of the model used for characterization, as a percentage of the original undistortion accuracy achieved using the same model undistortion and characterization.

It has been shown that by applying sound optimization techniques, and using higher order distortion models (which were proven to be stable in §4.3) that inverse distortion can be effectively modelled. In particular the error involved in undistortion has been reduced by 96% from that reported in literature [31] to 1.3 hundredths of a pixel for R3P2. This is despite the considerable distortion or up to 73 pixels inherent in the cameras used for this study (refer to §3.1).

With regards to the processing time required, only a single model call is necessary, this is equivalent to one third of an iteration in the guess-and-refine scenario (where a minimum of three model calls would be necessary to determine the error gradient). Compared to analytical undistortion models, it has one fewer parameter than [31] (for R3P2) and is more suitable to hardware implementation as it does not required a floating point division. Undistortion in [10] requires solving a fifth order polynomial for every undistorted pixel, this cannot be done analytically and is therefore much more computationally intensive. Finally in [2] it is shown that even for a single radial parameter undistortion would require two cube roots, two square roots and 6 divisions if the analytical solution was to be used.

CHAPTER 6.

HYPOTHESIS 3: DISTORTION IS NOT NECESSARILY RADIALLY SYMMETRICAL

6.1. Hypothesis formulation

All the radial lens distortion models found in literature are based on Brown's lens Distortion model [14, 27] as provided by Eq 2-1. In order to use Eq 2-1, each image point is converted from pixel (Cartesian) coordinates to polar coordinates with respect to the distortion centre. However, Eq 2-1 assumes radial symmetry for the K and thus only the polar radius is used and the polar angle is discarded. To this radially symmetrical distortion is added the asymmetrical tangential distortion components.

This assumption of radial symmetry is entirely valid if the image plane is exactly perpendicular to the optical axis, and the lenses have perfectly spherical surfaces and are aligned flawlessly. The tangential distortion component of Eq 2-1 allows for some decentring of the lenses, i.e. their optical axes are parallel yet not necessarily collinear.

In practice, it is not possible to achieve the conditions stated above, especially if one is unable to afford top of the range lenses. It is thus proposed that a "radial gain" be incorporated as per Eq 6-1 where the radial components have an asymmetry factor which is independent of the tangential distortion components:

$$x_u = x_d + f(q)(x_d - x_c)(K_1 r^2 + K_2 r^4 + \ldots) + (P_1(r^2 + 2(x_d - x_c)^2) + 2P_2(x_d - x_c)(y_d - y_c))(1 + P_3 r^2 \ldots)$$

$$y_u = y_d + f(q)(y_d - y_c)(K_1 r^2 + K_2 r^4 + \ldots) + (2P_1(x_d - x_c)(y_d - y_c) + P_2(r^2 + 2(y_d - y_c)^2))(1 + P_3 r^2 \ldots)$$

where:

(x_u, y_u) = undistorted image point,

(x_d, y_d) = distorted image point,

(x_c, y_c) = centre of distortion, **Eq 6-1**

K_n = Nth radial distortion coefficient,

P_n = Nth tangential distortion coefficient,

$r = \sqrt{(x_d - x_c)^2 + (y_d - y_c)^2}$,

$\theta = \tan^{-1}\left(\frac{y_d - y_c}{x_d - x_c}\right)$,

$f(\theta)$ = radial gain, and

"…" indicates an infinite series.

This radial gain function has to be continuous over the range [0, 2π) otherwise there will be distinct discontinuities evident in undistorted image lines. This also implies that $f(0) = f(2\pi)$. The following sections define the candidate functions with this property that were evaluated in this study.

6.1.1. Non-orthogonal conic slice

Consider a cone whose tip is at the focal point and whose axis coincides with the lens' axis. The intersection of the axis with the image plane is the principle point. The surface of the cone is the set of light rays which are all a certain angle off the optical axis. If the image plane is perfectly normal to the optical/cone axis, then these rays will produce a circle on the image plane. However, if this is not the case (which in the strictest sense it never can be), then these rays will create an ellipse. This implies that image coordinates of rays with the same angle between them and the optical axis, are different distances from the intersection point of the optical axis with the image plane depending on their polar angle in the image plane (θ in Figure 1.1). Figure 6.1 depicts an exaggerated example of this.

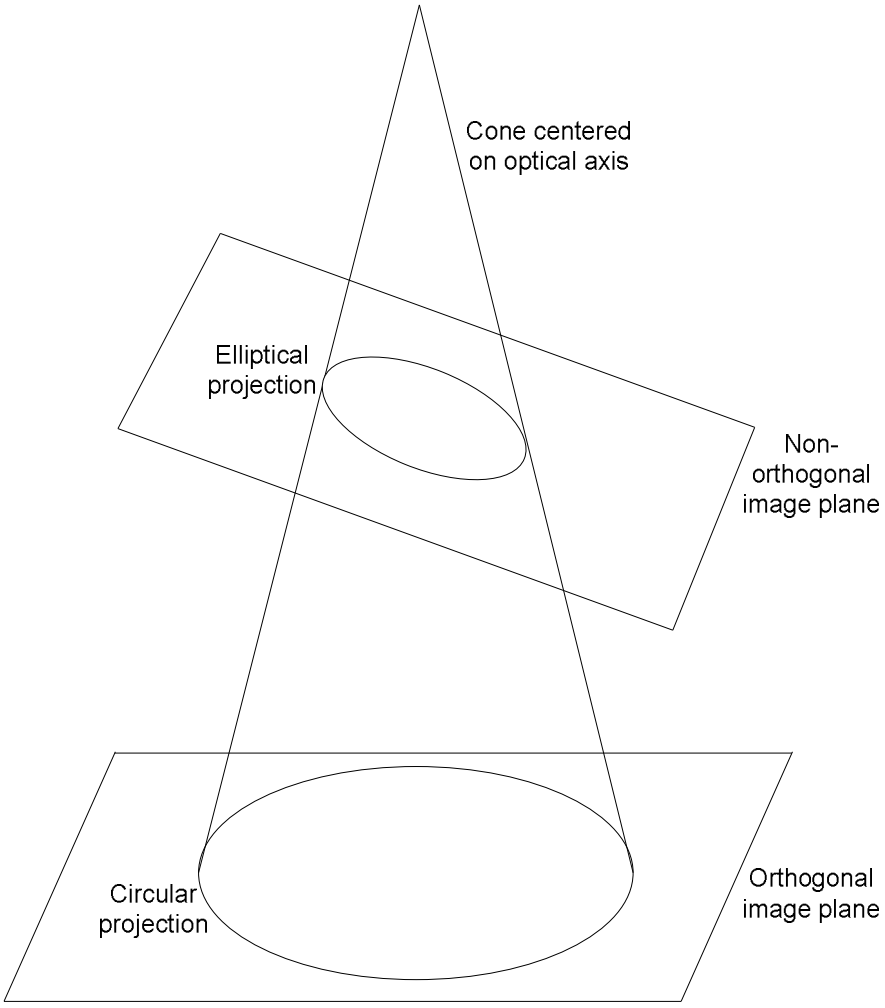

Figure 6.1 Depiction of the elliptical projection resulting from image plane non-orthogonality

With reference to Figure 6.2, Eq 6-2 defines the radius (as a function of angle from the positive X axis) of an ellipse that is centred at the origin and whose major axis is parallel to the X axis.

$$r = \sqrt{a^2 \cos^2(q) + b^2 \sin^2(q)},$$

where:

a = major axis,

b = minor axis,

r = distance from origin to ellipse, and

θ = polar angle.

Eq 6-2

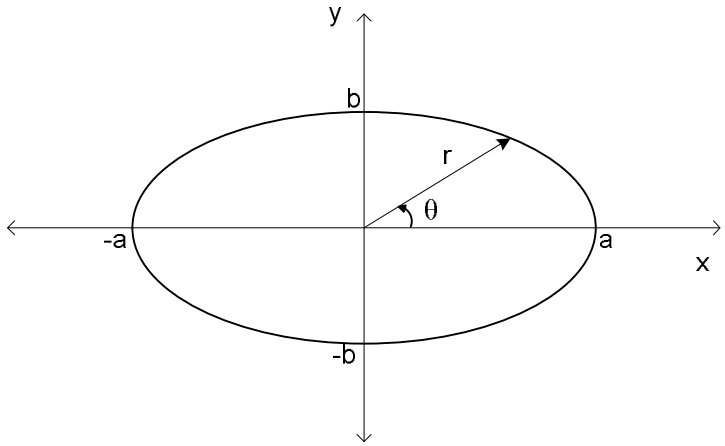

Figure 6.2 Ellipse with major axis parallel to the X axis

However, in general, the elliptical conical slice is unlikely to have its major axis parallel to the image plane X axis. Indeed, this would only be the case if the image plane was tilted about a vector parallel to the Y axis. Bearing this in mind and assuming (with out loss of generality) that the major axis has a length of 1, the final radial gain function can be expressed as

$$f(q) = \sqrt{\cos^2(q-a) + b^2 \sin^2(q-a)},$$

where:

b = minor axis and is less than 1,

r = distance from origin to ellipse

θ = polar angle, and

α = angle between the x axis and the major axis of the ellipse.

Eq 6-3

The parameters of this radial gain function to be optimised are b - the minor axis, and α – the angle between the X axis and the major axis of the ellipse.

6.1.2. Zero order clamped polynomial

The simplest, best understood and most tractable functions are polynomials, and so the simplest polynomial adhering to the requirements for an angular gain function is tested. Consider the simple polynomial $x^2 - 2px$, it has roots at $x = 0$ and $x = 2\pi$ and is smooth for all x. It thus meets the requirements for a radial gain function but has a fixed form. Eq 6-4 is a more generic form of this equation allowing the function to be more optimally fitted. Figure 6.3 depicts a possible radial gain for illustrative purposes.

$$f(q) = a((q-a)^2 - 2p(q-a)) + b,$$

where:

a is the gain,

b is the offset, and

α is the angular offset.

Eq 6-4

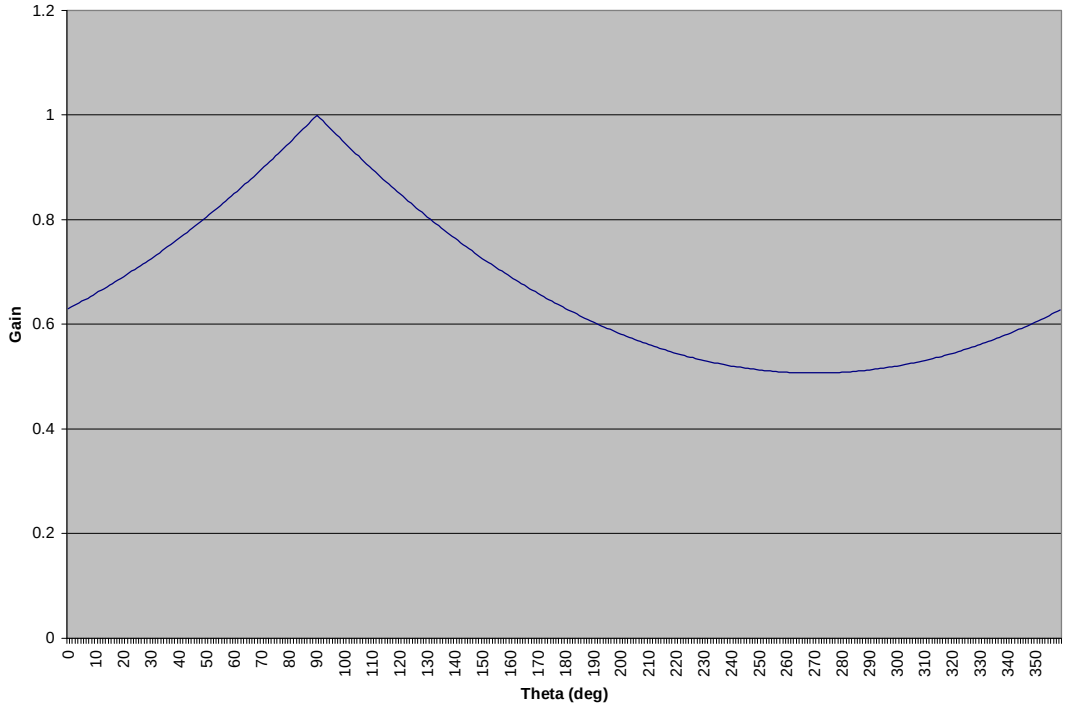

Figure 6.3 Example zero order clamped polynomial radial gain, where *a* = 0.05, *b* = 1.0, and α = π/2.

Figure 6.3 portrays the basic characteristics of this model: α pans the graph left or right, *a* controls its magnitude and *b* shifts the graph up and down.

6.1.3. First order clamped polynomial

As an improvement on the zero order clamped polynomial, an extra root was inserted in the domain of interest and allowed to be adjustable, as indicated by

$$f(q) = a(q-a)((q-a)-b)((q-a)-2p)+c,$$

where:

a is the gain,

c is the offset,

b is the second root, and

α is the angular offset.

Eq 6-5

This will allow one to determine if higher order polynomials are a better fit for angular asymmetry. Figure 6.4 depicts the general shape of the first order clamped polynomial radial graph, whose parameters function similar to the corresponding parameters of §6.1.2, except for b which controls where the graph crosses the vertical offset. Note that the first order clamped polynomial does not have the first derivative discontinuity evident at the angular offset point, unlike the zero order version.

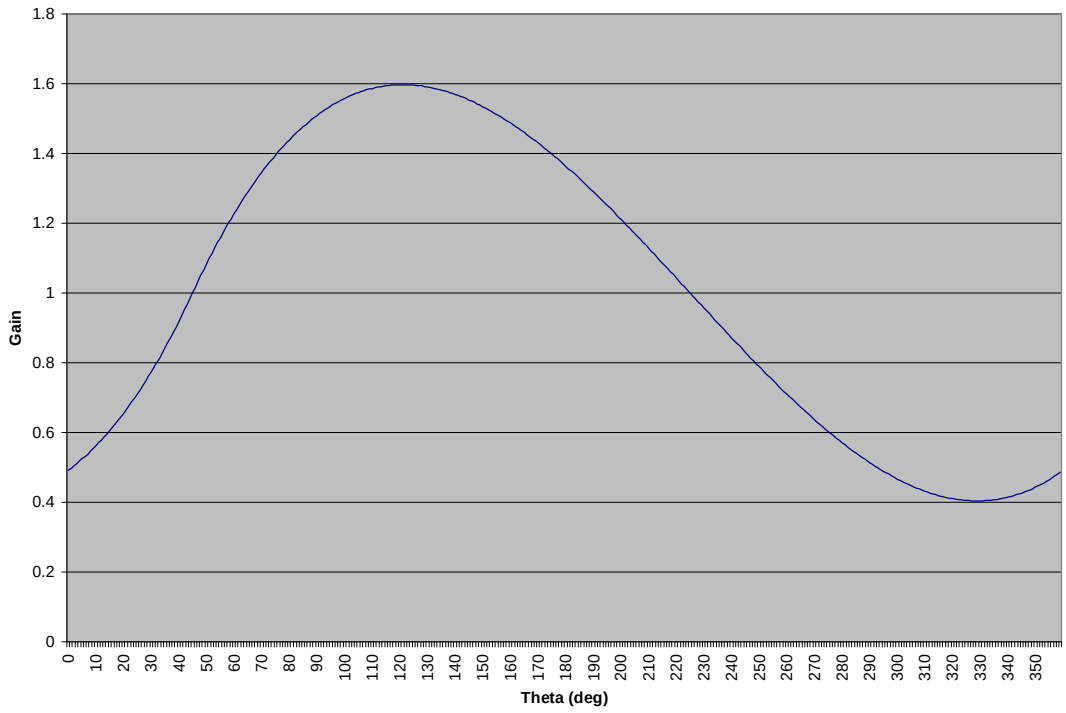

Figure 6.4 Example first order clamped polynomial radial gain, where a = 0.05, b = π, c = 1.0 and α = π /4

6.1.4. Floating Sine

The final angular gain considered in this study is the classical sine wave, as it meets all the criteria stipulated for the angular gain, and will determine if fitting the harmonics of the

sinusoid (i.e. a Fourier series) is plausible. Holding the frequency of the wave constant, the general formula for a sine wave becomes.

$$f(q) = a\sin(q - a) + b,$$

Eq 6-6

where:

a is the gain,

b is the vertical offset, and

α is the angular offset.

Figure 6.5 depicts the familiar sine wave shape, the three parameters of the floating sine angular gain function are precisely the same as those of the zero order clamped polynomial angular gain function of §6.1.2.

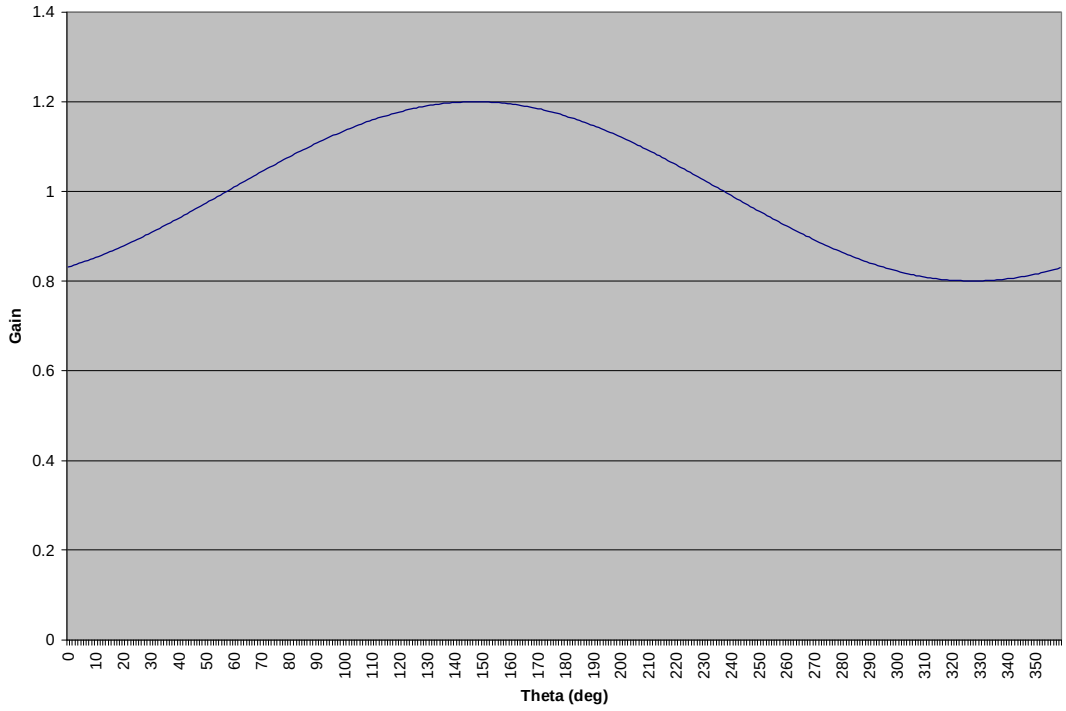

Figure 6.5 Example floating sine angular gain function, where a = 0.2, b = 1.0 and α = 1 rad.

6.2. Verification

The data used to test this hypothesis are the captured centroids of the optical reference jig (§3.2) for the 32 cameras described in §3.1. Approximately 2000 reference points for each camera were captured and are known to form straight lines. Unlike the numerical robustness study (§Chapter 1), where different combinations of the 2000 reference points were used to test sensitivity to input perturbations, all the points will always be used for each camera.

In order to test the hypothesis, each camera had five different optimisations performed as described below.

1. Unity angular gain, corresponding to the standard brown lens model of Eq 2-1.
2. Elliptical angular gain, corresponding to the angular gain function of §6.1.1
3. Zero order clamped polynomial angular gain, as described in §6.1.2.
4. First order clamped polynomial angular gain, as described in §6.1.3.
5. Floating sine angular gain as per §6.1.4.

This will allow one to compare the effects of allowing different types of angular gain on the resultant optimally distortion free image, where the distortion measure used is the RMS perpendicular distance measure as described in §3.3.1.

Each of the optimizations will be performed using the Leapfrog numerical optimization algorithm where the maximum number of iterations is restricted to 10000, the maximum step size is 0.1, and the convergence criteria are a step size of 10^{-8} or a gradient of less than $\sqrt{n}10^{-6}$, where n is the number parameters being optimized. This last serves to ensure that each element of the gradient vector is, on average, less than 10^{-6} regardless of the number of the parameters, and thus does not require smaller gradient vector elements for convergence as the number of parameters being optimized increases. Finally the gradient calculation method used will be the central gradient as expressed by Eq 6-7 where the gradient step size was chosen to be 10^{-8} (i.e. the same order of magnitude as the minimum step size).

$$\nabla f(\bar{x}) = \begin{bmatrix} f_1(\bar{x}) \\ f_2(\bar{x}) \\ \vdots \\ f_n(\bar{x}) \end{bmatrix}$$

$$f_i(\bar{x}) = \frac{f(x_1, x_2, .., x_i + d, x_{i+1}, .., x_n) - f(x_1, x_2, .., x_i - d, x_{i+1}, .., x_n)}{2d}$$

Eq 6-7

where:

f is the function whose gradient vector is being determined,

f_i is the i-th element of the gradient vector,

$\boldsymbol{x}$ is the position at which the gradient is being estimated,

x_j is the j-th value of the vector $\boldsymbol{x}$,

δ is the gradient calculation step size, and

n is the number of parameters on which f is dependant.

5 optimisations will be done for each of the 32 cameras for each of the 7 distortion functions specified in Table 4.2. This will allow a comparison between the general applicability of each of the proposed angular gain functions, and determine if non-uniform angular gains can improve lens distortion characterization.

It was seen during preliminary testing that angular gains could result in a skewness of the resulting undistorted domain. The skewness transforms a rectangular collection of corrected image points, into a trapezoidal or rhombic collection of points as illustrated by Figure 6.6 and Figure 6.7. In order to quantify this effect, each undistorted line, as generated by the variable angular gain distortion function, will be compared to the corresponding undistorted line as calculated when using the same distortion function with out any angular gain. The maximum angle between any of the 96 pairs of such lines (4 sets of images of 6 concentric squares) will then be taken as the skewness measure, i.e.:

$$\bar{p}_i = \frac{(\bar{p}_{i,end} - \bar{p}_{i,begin})}{\left\|\bar{p}_{i,end} - \bar{p}_{i,begin}\right\|}, \bar{q}_i = \frac{(\bar{q}_{i,end} - \bar{q}_{i,begin})}{\left\|\bar{q}_{i,end} - \bar{q}_{i,begin}\right\|},$$
$$skewness = \max_i \left|\cos^{-1}(\bar{p}_i \bullet \bar{q}_i)\right| \quad \textbf{Eq 6-8}$$

where:

$i \in [1, 96]$ is the line number being considered,

$\boldsymbol{p}_{i,\,end}$ is the last existing undistorted point in the variable radial gain line,

$\boldsymbol{p}_{i,begin}$ is the first existing undistorted point in the variable radial gain line,

$\boldsymbol{q}_{i,\,end}$ is the last existing undistorted point in the constant radial gain line, and

$\boldsymbol{q}_{i,begin}$ is the first existing undistorted point in the constant radial gain line.

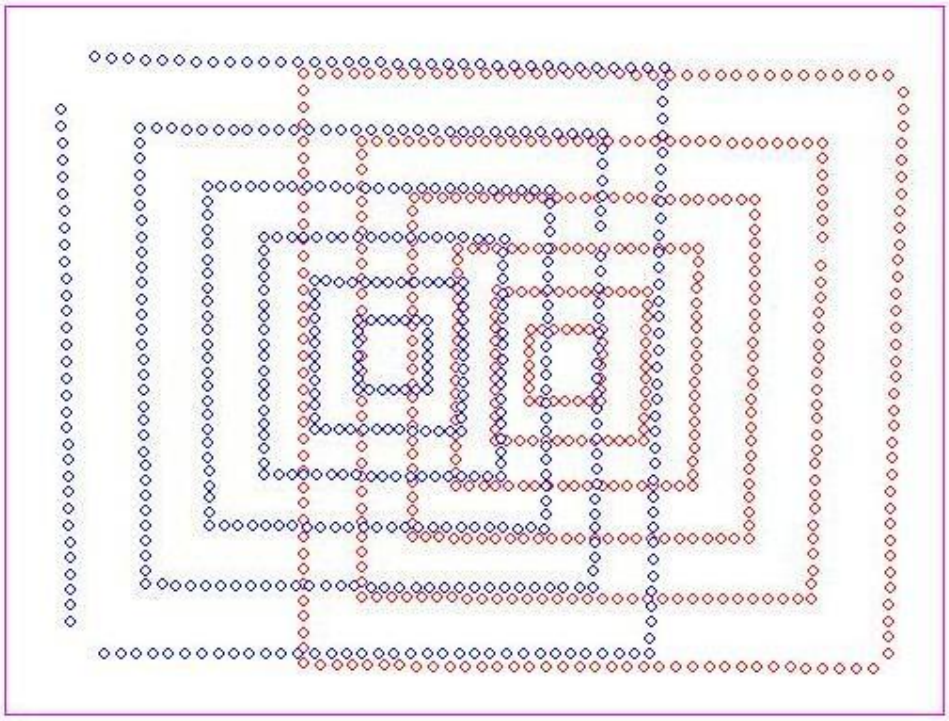

Figure 6.6 Depiction of undistorted points from two views of the reference jig when no angular gain is used. The outline represents the camera's FOV.

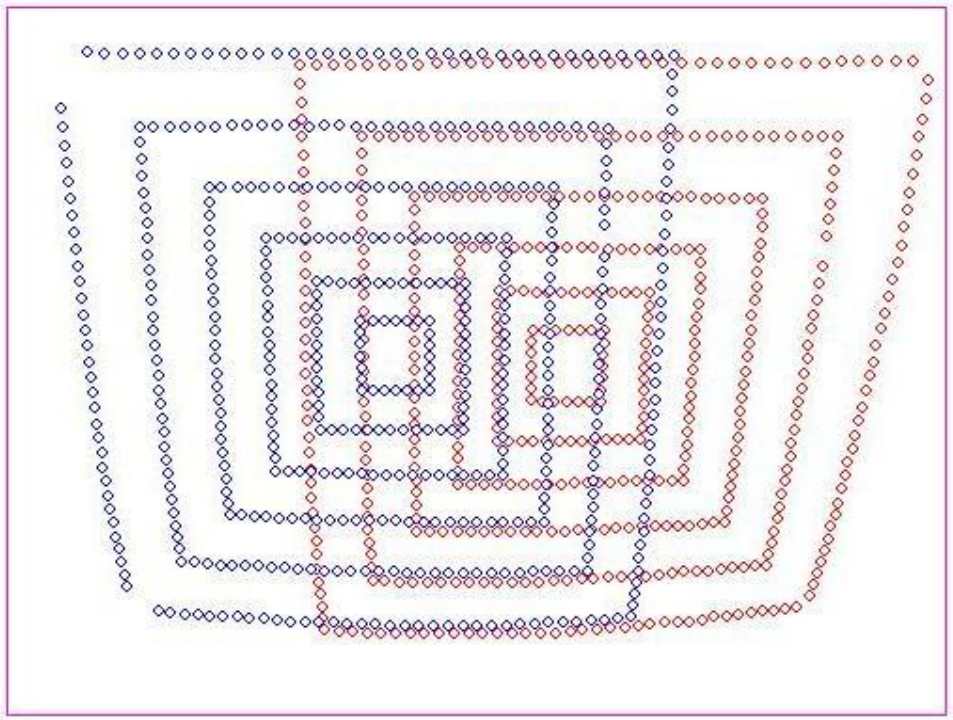

Figure 6.7 Depiction of undistorted points from two views of the reference jig, exhibiting skewness due to poor angular gain. The outline represents the camera's FOV.

6.3. Analysis of results

The raw results of all the optimizations performed are far too extensive to be presented here, and thus only the statistics for each of the distortion model/angular combinations are presented. The complete raw results are available upon request.

As mentioned in §3.1 the 32 sample cameras contain two different lens designs. As the results differ for each of the lens designs, they are presented separately.

Figure 6.8 to Figure 6.14 below present that data for the different distortion models as a series of box plots. The whiskers on the plot extend to the minimum and maximum improvements for a lens/model/gain combination. The box itself extends from the 25^{th} to 75^{th} percentiles. All improvements are expressed as a percentage compared to having unity angular gain, where a positive value indicates an improvement, and negative values deterioration. The angular gains listed on the X axis of the graphs: conic, zero order, first order and floating sine correspond to those described paragraphs §6.1.1, §6.1.2, §6.1.3, and §6.1.4 respectively.

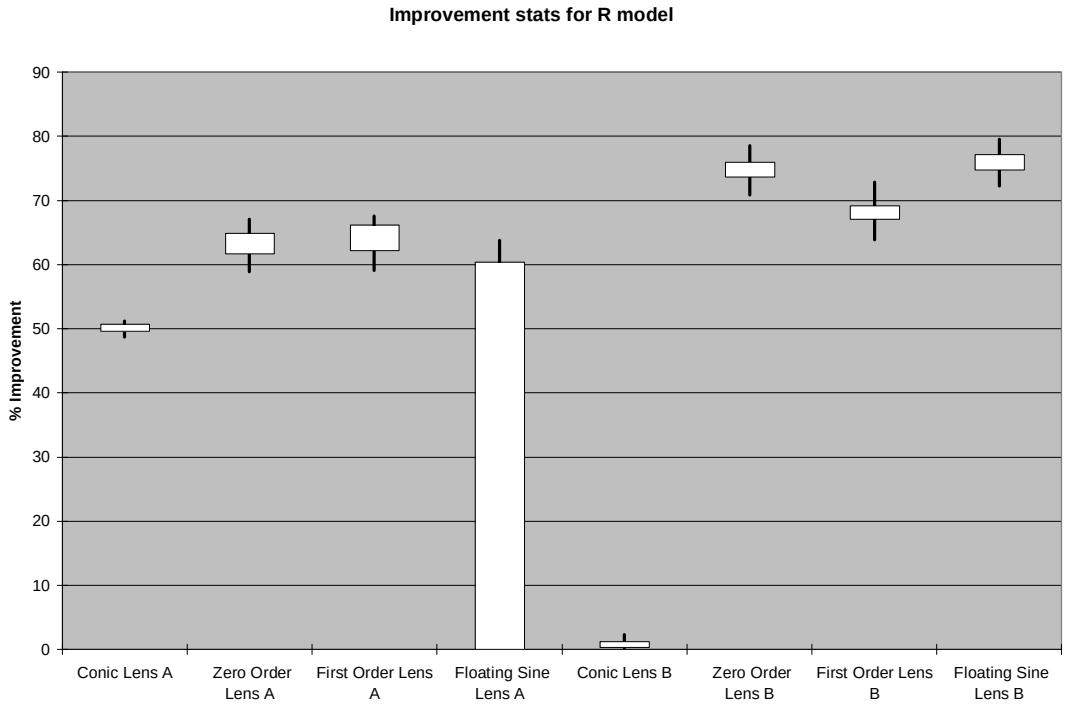

Figure 6.8 Box plot presentation of the improvements obtained by the different angular gains over unity gain for the R distortion model.

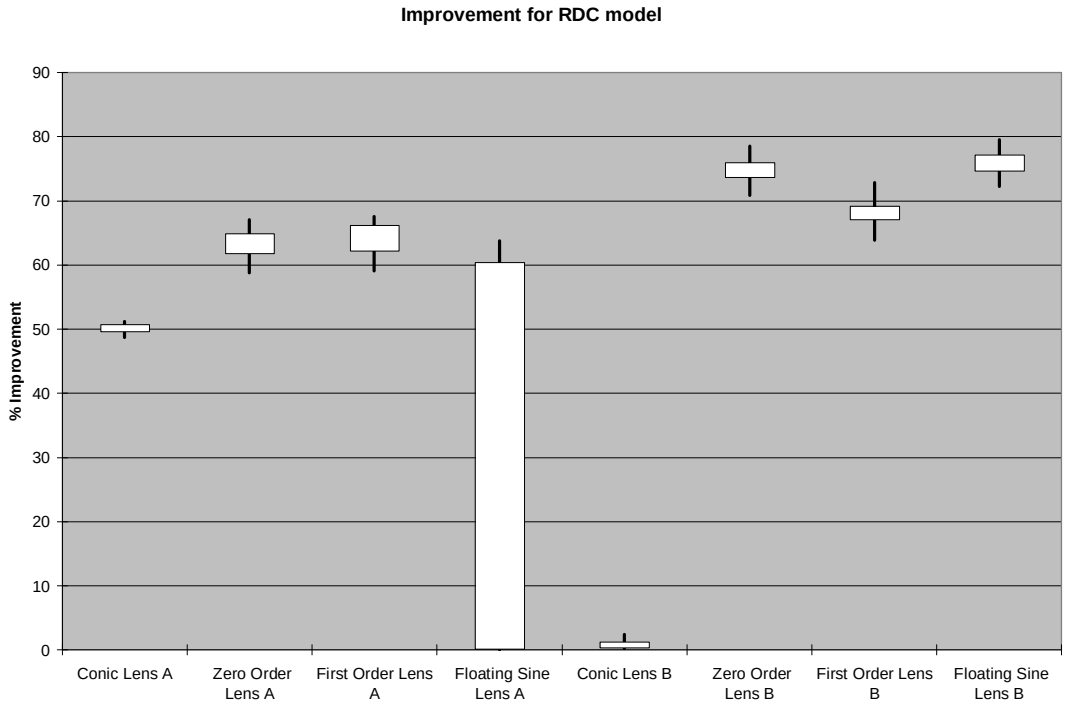

Figure 6.9 Box plot presentation of the improvements obtained by the different angular gains over unity gain for the RDC distortion model.

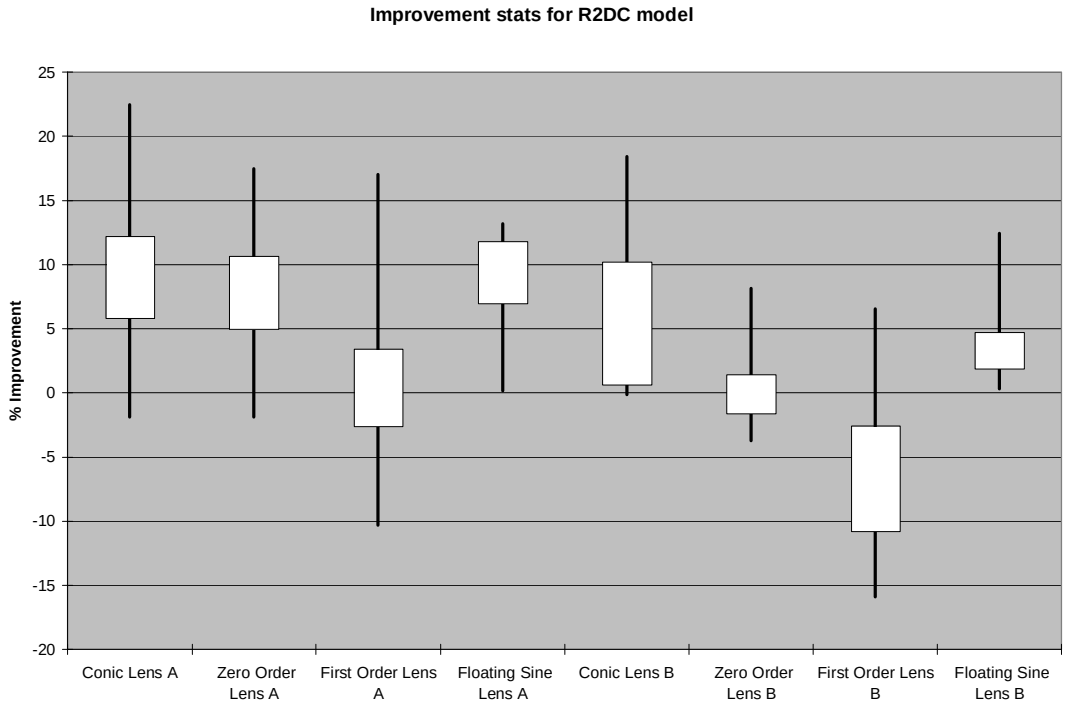

Figure 6.10 Box plot presentation of the improvements obtained by the different angular gains over unity gain for the R2DC distortion model.

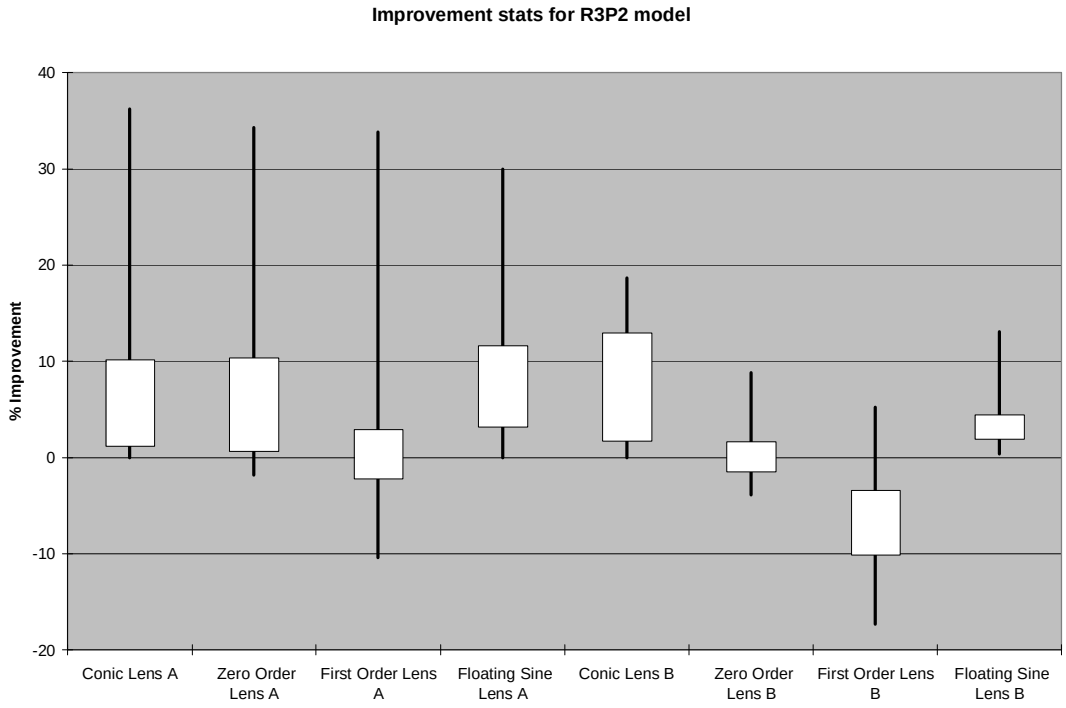

Figure 6.11 Box plot presentation of the improvements obtained by the different angular gains over unity gain for the R3P2 distortion model.

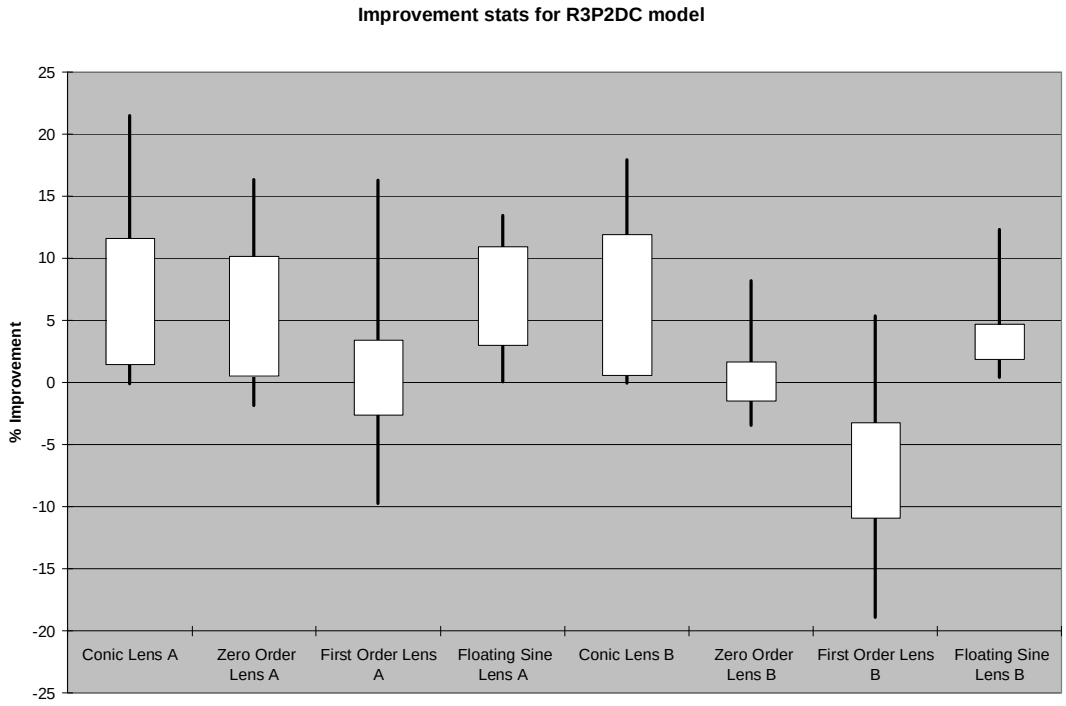

Figure 6.12 Box plot presentation of the improvements obtained by the different angular gains over unity gain for the R3P2DC distortion model.

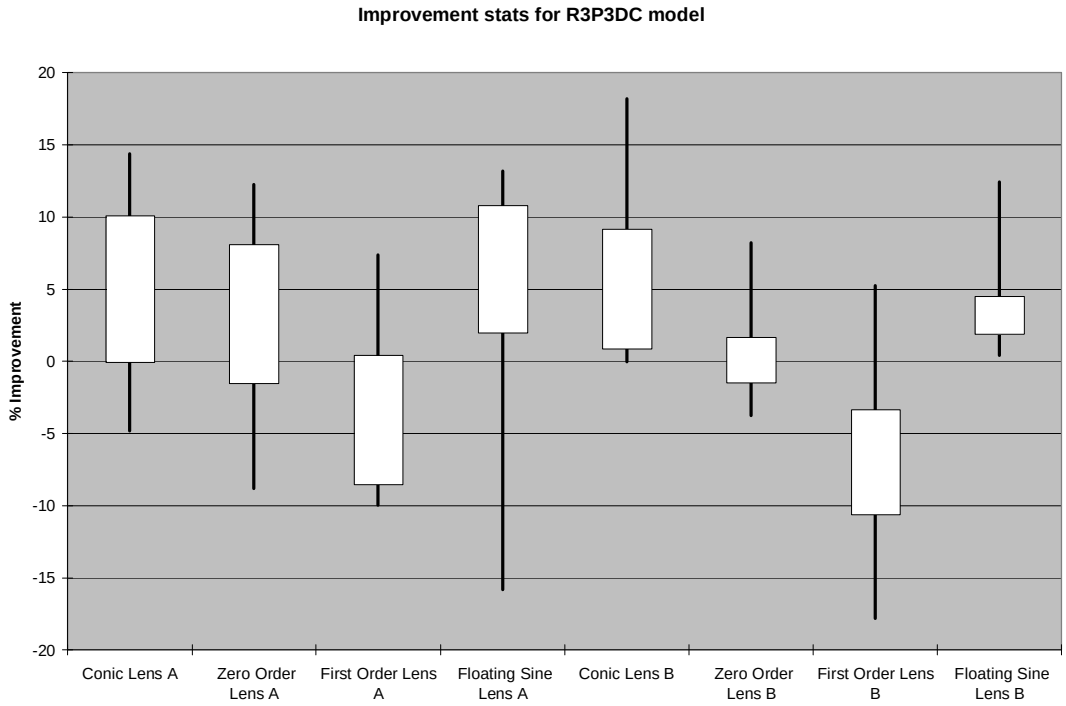

Figure 6.13 Box plot presentation of the improvements obtained by the different angular gains over unity gain for the R3P3DC distortion model.

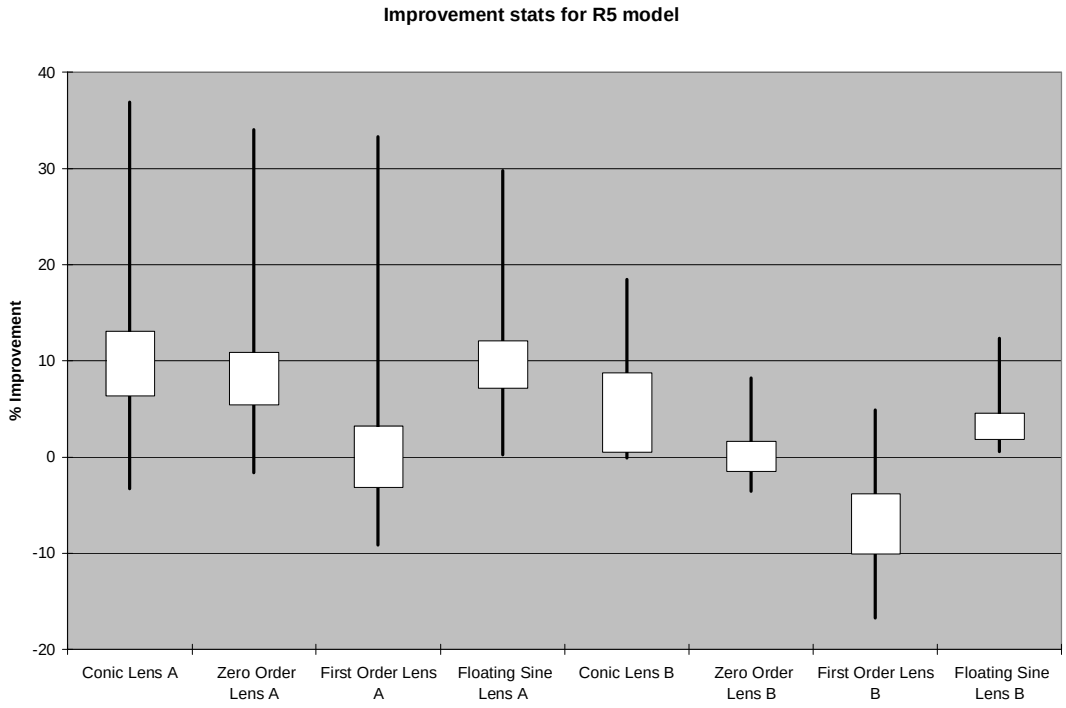

Figure 6.14 Box plot presentation of the improvements obtained by the different angular gains over unity gain for the R5 distortion model.

The average skewness, as per Eq 6-3, over all 32 cameras for each distortion model/ angular gain combination is presented in tabular form in Table 6.1 and pictorially in Figure 6.15.

Table 6.1 Average skewness in degrees of each distortion model/angular gain combination

	Conic	Zero Order	First Order	Floating Sine
R	0.226	11.839	12.398	8.585
RDC	0.233	11.778	12.395	8.593
R2DC	0.218	3.605	7.364	0.836
R3P2	0.225	3.633	7.472	0.755
R3P2DC	0.213	3.618	7.500	0.712
R3P3DC	0.241	3.633	7.646	0.743
R5	0.217	3.625	7.463	0.751

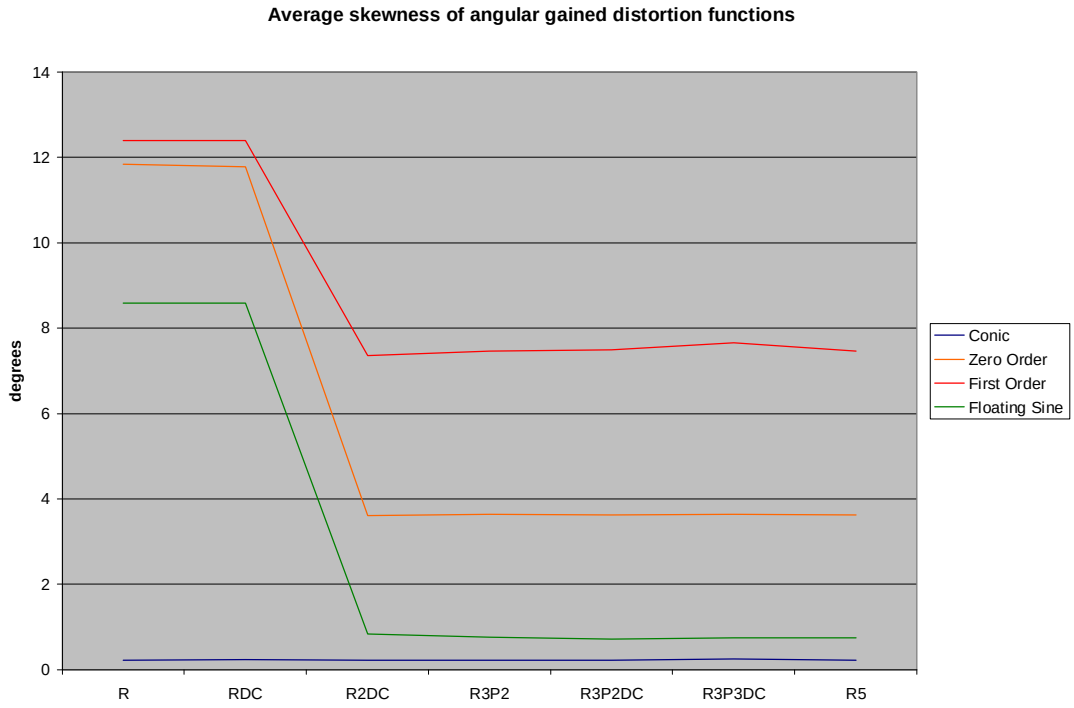

Figure 6.15 Average skewness of each distortion model/angular gain combination.

It is evident from Figure 6.8 and Figure 6.9 that one is able to dramatically improve the residual distortion of these simple models. This is particularly true for Lens A, however Lens B could not be significantly improved by the conical angular gain, which was the only gain function not to induce significant skewness. The remarkable similarity between Figure 6.8 and Figure 6.9 is attributable to the close relationship between the two distortion methods which only differ by virtue of the second method allowing a variable distortion centre.

With regards to the five distortion models that have a second or higher order radial component (i.e. Figure 6.10 through Figure 6.14), only the conical and floating sine gains consistently improved the distortion characterization. This is true for both lens designs. The conical gain provided both a wider spread and higher average improvement compared to the floating sine method. Both of these methods also provided minimal skewnesses of less than one degree. Considering that this skewness is essentially a comparison with the non-perfect distortion correction obtained when using unity angular gain; both of the skewnesses of 0.22° and 0.75° obtained on average by the conic and floating sine methods respectively are acceptable.

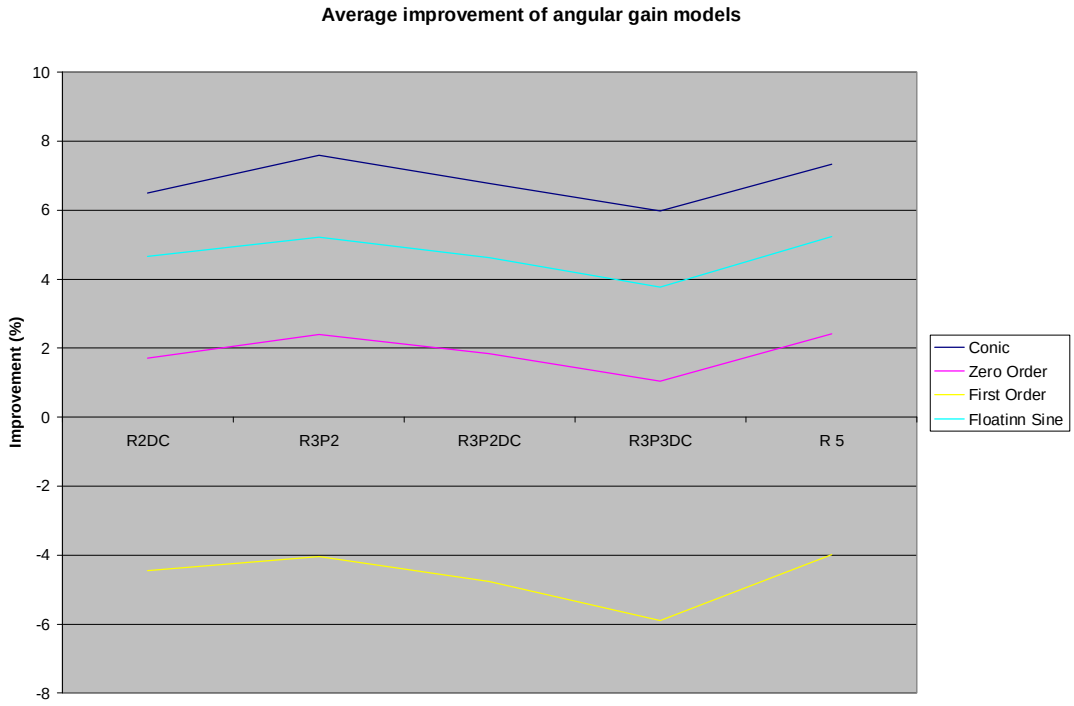

Figure 6.16 Average percentage improvement of each angular method compared to unity angular gain.

With reference to Figure 6.16, it can be stated that incorporating the angular component of the polar coordinate of an image point in the lens distortion function can improve the distortion characterization. It is evident from Figure 6.16 that the improvement seems to lessen with model complexity, which is expected as higher order models can better characterize the distortion. In particular the conical gain, which was derived from a physical phenomenon, seems to perform particularly well.

It has been shown that angular gain can improve lens distortion characterization. As the purpose of this hypothesis was not to determine the optimal angular model but merely to ascertain if such a model would be beneficial, one can safely say that it has been verified.

CHAPTER 7.

CONCLUSIONS

7.1. Comparison of results with previous research

The first hypothesis challenged the prevailing belief that high order radial distortion models were not suitable to modern digital photogrammetric applications. This was due both to a perceived instability in the fitting of these models; and a belief that such models actually worsened the characterization of lens distortion. It was shown via the application of modern numerical optimization methods that higher order models can indeed be successfully used for digital photogrammetry. The fitting of the models was shown to be as stable as for lower order models, whilst they were shown to provide better characterization of lens distortion.

Building on a stable foundation provided by the first hypothesis, the second hypothesis queried the established inverse distortion methods used. Previously one had to choose between an iterative (i.e. slow) yet accurate technique or a quick but less accurate model based on simplifications of the inverse of low order radial models. It was shown that a high order radial model could successfully be fitted to model the inverse distortion by reusing left-over distortion characterization data. Such an inverse characterization is of the same order of processor loading as the simplified inverse low order models, yet attains accuracies comparable to the iterative technique of 0.013 pixels RMS.

The third hypothesis queried the core assumption of radial symmetry used in all radial lens distortion models. It was shown that multiplying the corrections of current radial models with a simple scale factor that is a function of the polar angle of the image point in question, improves lens distortion characterization. This improvement varies from 50% with the simplest distortion models to 6% for the most complex models which prior to this work (in the first hypothesis) had not even been deemed possible to use.

7.2. Applicability of results

The improved characterization accuracies yielded by this work were not vast. The superior characterization offered by the highest order model tested (a third order radial, third order tangential model with decentering) is only 3% (for the two lenses considered) when compared to the second order radial model commonly used in industry. The improvement

provided by allowing variable angular gain is then a further 6%. This improvement is only likely to be useful for high accuracy photogrammetric applications. Applications such as robotic vision, image stitching and distortion correction for visually aesthetic reasons are unlikely to find these improvements useful.

The improvements to inverse distortion correction however are more dramatic. With only a single call to a model (which is of less complexity than current models), accuracy that is comparable to iterative procedures can attained. Whereas a *minimum* of 3 function calls are required to find the partial derivatives with respect to X and Y for each iteration (and this is a poor forward-difference approximation of the gradient – a central difference gradient with five models calls would be better) only one call needs to be made. This allows accurate inverse distortion modelling, enabling images to be undistorted accurately in real time. This is relevant to fields such as robotic vision, automatic defect detection, image registration and, of course, high accuracy photogrammetric measurements.

7.3. Contribution made by this work

In addition to definitive proof of the applicability of high order radial distortion models, two numerical optimization techniques were highlighted that are suitable to the task of fitting these models to the captured data. The Fletcher-Reeves method converges rapidly to a solution and is not affected by input perturbations. It is thus recommended when the data is noisy or limited in number or FOV coverage. The Leapfrog algorithm consistently finds a better fit for the models (therefore providing a better characterization) but is more sensitive to input perturbations, it is thus recommended when there is confidence in the quality of the data and it abundantly covers the entire FOV.

Inverse distortion modelling has been advanced significantly. Using the undistorted positions created as a side effect in the distortion characterization, it was shown that the Leapfrog algorithm can fit a high order model such that with a single model call the distorted point corresponding to an undistorted point can be found to within a few hundredths of a pixel. This is an increase in accuracy of 25 fold compared to the current inverse distortion models and uses a third of the processing required for a single iteration of the iterative techniques. Low cost, low power, real-time distortion correction is now feasible.

The accuracy of lens distortion characterization, already improved by using high order models, has been further enhanced by allowing distortion to be radially asymmetrical. This

has allowed images to be straightened so that points along straight lines are within 7 hundredths of a pixel (RMS) of the line. This is an improvement of 5580% compared to the distorted image. The ability provided by the conic slice radial gain to algorithmically correct for non-orthogonality between the imager and optical assembly further paves the way for the use of less expensive equipment for photogrammetric applications.

As two different lens types were used and the results varied for each type, this research highlights the need to determine what order of model is best suited to the lens type being used in the application at hand. It was shown that the different lenses had different properties concerning their initial distortion characteristics, responded differently to different distortion (and inverse distortion) characterization models, and differed in their improvements induced by variable radial gains.

It has thus been shown that the characterization of any lens might be significantly improved by selection of an appropriate distortion model (possibly with a variable radial gain) and optimization technique.

7.4. Future work to be done

Better models for radial gain need to be developed. This work has successfully shown that such models will be beneficial to high accuracy photogrammetric applications.

The conic slice radial gain model, which was the best performing radial gain model, would also benefit from further work. In particular the relationship between the eccentricity of the ellipse, its two foci, and its angle in the focal plane may possibly be exploited to determine the normal of the optical axis relative to the image plane.

The comparable performance of the floating sine angular gain to the conic slice radial gain model suggests that future work using a Fourier series may be beneficial.

Although the polynomial functions performed poorly, it may be beneficial to attempt using a set of orthogonal polynomials such as Legendre or Chebyshev polynomials.

CHAPTER 8.

REFERENCES

[1] H. Bacakoglu, and M. Kamel, “An optimized two-step calibration method,” in *Proceedings of the 1997 IEEE International Conference on Robotics and Automation, 20-25 April, 1997*, vol. 2, pp. 1347-1352.

[2] R. Cucchiara, C. Grana, A. Pratzi, and R. Vezzani, “A Hough transform-based method for radial lens distortion correction,” in *Proceedings of the 12th International Conference on Image Analysis and Processing*, 17*–19 Sept, 2003,* vol. 1, pp. 182-187.

[3] S. Graf, and T. Hanning, “Analytically solving radial distortion parameters,” in *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 20*–25 June, 2005,* vol. 2, pp. 1104-1109.

[4] G. McLean, “Image warping for calibration and removal of lens distortion,” in *Proceedings of the IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, 19-21 May, 1993,* vol. 1, pp. 170-173.

[5] J. Perš, and S. Kovačič, "Nonparametric, Model-Based Radial Lens Distortion Correction Using Tilted Camera Assumption," in *Proceedings of the Computer Vision Winter Workshop, Feb, 2002*, vol. 1, pp. 286-295.

[6] K. Sirisantismrid, T. Matsuura, and K. Tirasesth, “A simple technique to determine calibration parameters for coplanar camera calibration,” in *Proceedings of the 2004 IEEE Region 10 Conference, 21-24 Nov, 2004,* vol. A, pp. 677-680.

[7] R. Y. Tsai, “A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf tv cameras and lenses,” *IEEE Journal of Robotics and Automation*, vol. 3, no. 4, pp. 323–344, 1987.

[8] Z. Zhang, “Flexible camera calibration by viewing a plane from unknown orientations,” in *Proceedings of the Seventh IEEE International Conference on Computer Vision, 20-27 Sept, 1999,* vol. 1, pp. 666-673.

[9] Z. Zhang, “A flexible new technique for camera calibration,” *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 22, no. 11, pp. 1330-1334, 2000.

[10] F.M. Candocia, “A scale-preserving lens distortion model and its application to image registration,” in *Proceding of the 2006 Florida Conference in Recent Advances in Robotics (FCRAR 2006), 23-26 May, 2006,* vol. 1, pp. 1-6.

[11] S.C. Hsu, and H.S. Sawhney, "Influence of global constraints and lens distortion on pose and appearance recovery from a purely rotating camera," in *Proceedings of the 4th IEEE Workshop on Applications of Computer Vision*, 19–*21 Oct, 1998,* vol. 1, pp. 154–159.

[12] J.C.A. Fernandes, M.J.O. Ferreira, J.A.B.C. Neves, and C.A.C. Couto, "Fast correction of lens distortion for image applications," in *Proceedings of the IEEE International Symposium on Industrial Electronics*, 7-*11 July, 1997*, vol. 2, pp. 708-712.

[13] W. Zheng, Y. Shishikui, Y. Kanatsugu, Y.Tanaka, and I Yuyama, "A high-precision camera operation parameter measurement system and its application to image motion inferring," *IEEE Transactions on Broadcasting,* vol. 47, no. 1, pp. 46-55.

[14] D.C. Brown, "Close Range Camera Calibration," *Photogrammetric Engineering,* vol. 37, no. 8, pp. 855-866, 1971.

[15] D. Claus, and A.W. Fitzgibbonn, "A rational function lens distortion model for general cameras," in *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 20–*25 June, 2005,* vol. 1, pp. 213-219.

[16] O. Silven, and J. Heikkila, "Calibration procedure for short focal length off-the-shelf CCD cameras," in *Proceedings of the 13th International Conference on Pattern Recognition*, 25–*29 Aug, 1996,* vol. 1, pp. 166-170.

[17] G.P. Stein, "Lens distortion calibration using point correspondences," in P*roceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97), 17-19 June, 1997*. vol. 1, pp. 602-608.

[18] I. Tsatsakis, E. Kayafas, V. Loumos, and G. Cambourakis, "Using low cost video cameras in automation: a close range photogrammetry approach," in *Proceedings of the IEEE International Symposium on Industrial Electronics, 10-14 July, 1995,* vol. 2, pp. 523-528.

[19] G.E. Karras, G. Mountrakis, P. Patias, and E. Petsa, "Modelling distortion of super-wide-angle lenses for architectural and archeological applications," *International Archives of Photogrammetry and Remote Sensing*, vol. 32, no. 5, pp. 570-573, 1998.

[20] Y.C. Lin, and C.S. Fuh, "Distortion for digital cameras," in *Proceedings of the International Symposium on Computer Graphics, Image Processing and Vision, 20-23 Oct, 1998,* vol. 1, pp. 396-401

[21] M. Ahmed, and A. Farag, "Nonmetric calibration of camera lens distortion," in *Proceedings of the 2001 International Conference on Image Processing, 7–10 Oct, 2001,* vol 2, pp. 157-160.

[22] M. Ahmed, and A. Farag, "Nonmetric calibration of camera lens distortion: differential methods and robust estimation," *IEEE Transactions on Image Processing*, 2005 vol. 8, no. 14, pp. 1215-1230.

[23] M. T. El-Melegy, and A.A. Farag, "Statistically Robust Approach to Lens Distortion Calibration with Model Selection," in *Proceedings of the International Conf. on Computer Vision and Pattern Recognition (CVPR-03), 16-22 June, 2003*, vol. 1, pp. 150-156.

[24] J.I. Jeong, S.Y. Moon, S.G. Choi, and D Rho, "A study on the flexible camera calibration method using a grid type frame with different line widths," in *Proceedings of the 41st SICE Annual Conference*, 5–*7 Aug, 2002*, vol. 2, pp. 1319-1324.

[25] Y. Meng, and H. Zhuang, "What you see is what you get (self-calibrating camera lens distortion)," *IEEE Robotics and Automation Magazine*, vol. 11, no. 4, pp. 123-127, 2004.

[26] A.E. Conrady , "Decentered lens systems," *Monthly Notices of the Royal Astronomical Society*, Vol. 79, pp. 384-390, 1919

[27] D.C. Brown, "Decentering distortion of lenses," *Photogrammetric Engineering*, vol. 32, no 7, pp. 444-462, 1966.

[28] R. Sagawa, M. Takatsuji, T. Echigo, and Y. Yagi, "Calibration of lens distortion by structured light scanning," in *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2-6 Aug, 2005*, vol. 1, pp 832-837.

[29] W. Yu, and Y. Chung, "An embedded camera lens distortion correction method for mobile computing applications," in *Proceedings of the 2003 IEEE International Conference on Consumer Electronics, 17-19 June, 2003*, vol. 1, pp. 400-401.

[30] W. Yu, "An embedded camera lens distortion correction method for mobile computing applications," *IEEE Transactions on Consumer Electronics*, vol. 49, no 4, pp. 894-901, 2003.

[31] J. Mallon, P.F. Whelan, "Precise Radial Un-distortion of Images," in *Proceedings of the 17th International Conference on Pattern Recognition (ICPR2004), 23-26 August, 2004*, Vol 1, pp. 18-21.

[32] A.G.J. Nijmeijer, et al, "Correction of lens-distortion for real-time image processing systems," in *Proceedings of the Workshop on VLSI Signal Processing, 20-22 Oct, 1993*, vol. 1, pp. 316-324.

[33] K. Hwang, and M.G. Kang, "Correction of lens distortion using point correspondences," in *Proceedings of the IEEE Region 10 Conference, 15–17 Sept, 1999*, vol. 1, pp. 690–693.

[34] S.W. Shih, Y.P. Hung, and W.S. Lin, "Accuracy assessment on camera calibration method not considering lens distortion," in *Proceedings of the International Conference. on Computer Vision and Pattern Recognition (CVPR '92), 15-18 June, 1992*, vol. 1, pp. 755-757.

[35] T.A. Clarke, and J.G. Fryer, "The Development of Camera Calibration Methods and Models," *Photogrammetric Record*, vol. 16, no, 91, pp. 51-66, 1998.

[36] L. Lucchese, S.K. Mira, "Using saddle points for subpixel feature detection in camera calibration targets," in *Proceedings of the Asia-Pacific Conference on Circuits and Systems*, 2002, vol. 2. pp. 191-195

[37] A. Redert, E. Hendriks, and J. Biemond, "Accurate and robust marker localization algorithm for camera calibration," in *Proceedings of the First International Symposium on 3D Data Processing Visualization and Transmission, 19-21 June, 2002*, vol. 1, pp. 522-525.

[38] J.A. Snyman, *Practical Mathematical Optimization*, 1st Ed. New York, USA: Springer, 2005.

[39] R.L. Burden, and J.D. Faires, *Numerical Analysis*, 7th Ed. Pacific Grove, USA: 2001.

[40] K. Levenberg, "A method for the solution of certain non-linear problems in least squares," *Quarterly Applied Mathematics*, No. 2, pp. 164-168, 1944.

[41] D.W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *J. Soc. Indust. Appl. Math.*, vol. 11. no. 2, pp. 431-441, 1963.

[42] R. Flectcher and C.M. Reeves, "Function minimization by conjugate gradients," *Computer Journal*, vol. 7, pp. 140-154, 1964.

[43] J.A. Snyman, "An improved version of the original leap-frog dynamic method for unconstrained minimization: LFOP1(b)," *Applied Mathematics and Modelling*, vol. 7, no. 7, pp. 216-218, 1983.

Appendix A. Leapfrog Algorithm

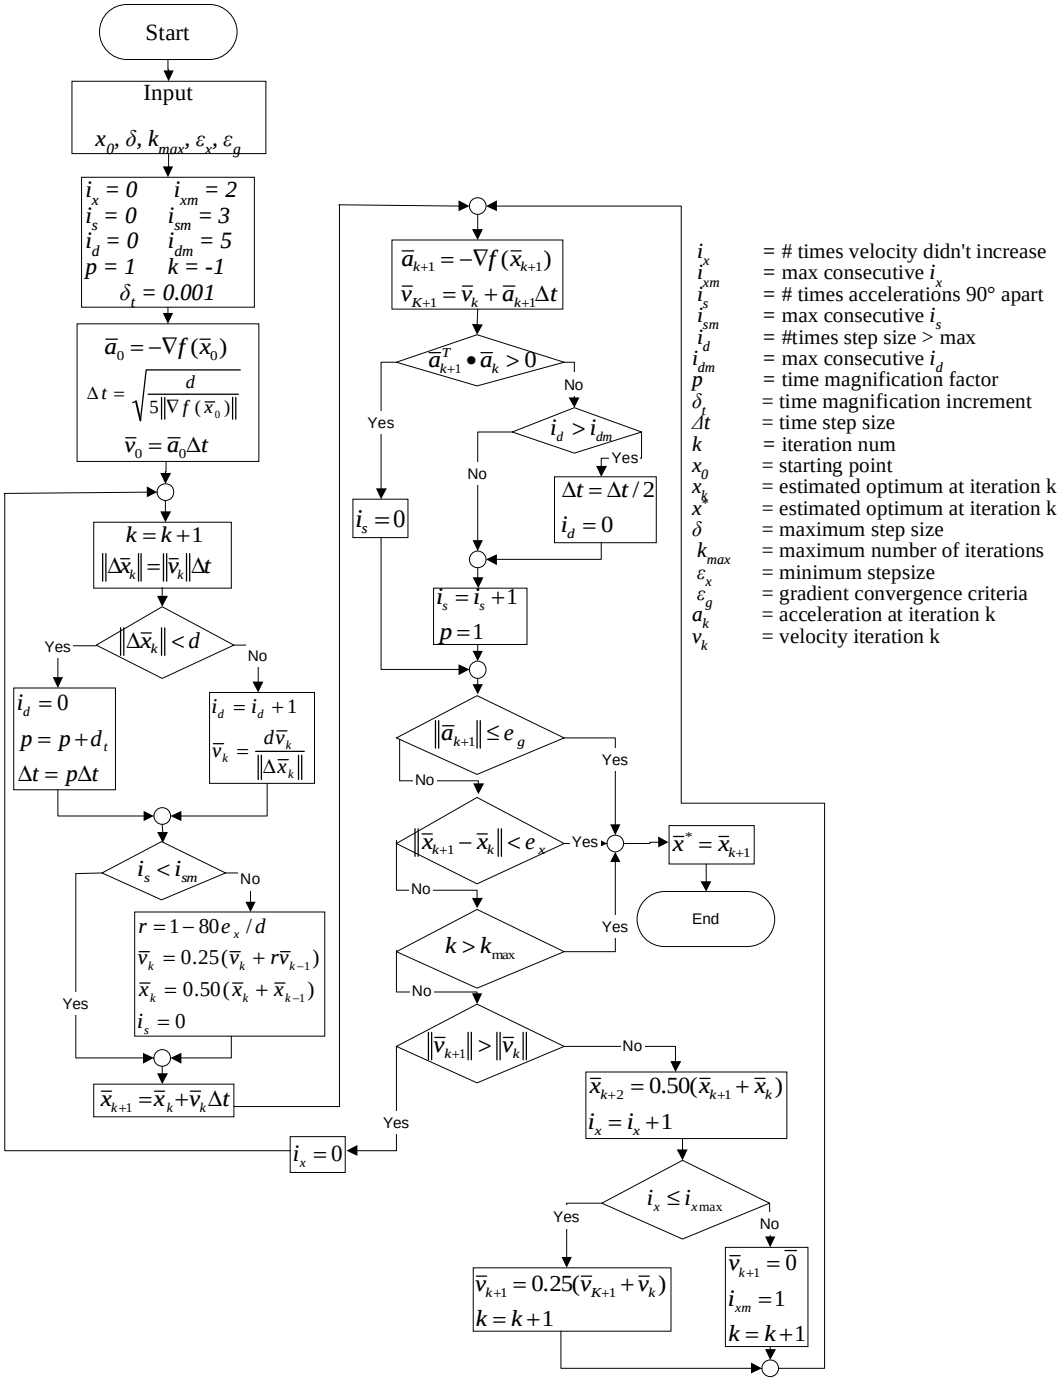

Figure A. 1 Leapfrog Algorithm flow diagram, adapted from [38] showing how LF simulates the motion of a charged particle subject to an N-dimensional field determined by the error gradient.