Correcting Radial Distortion by Circle Fitting

Rickard Strand and Eric Hayman
Computational Vision and Active Perception Laboratory (CVAP)
School of Computer Science and Communication
Royal Institute of Technology (KTH), SE-100 44 Stockholm, Sweden
hayman@nada.kth.se

Abstract

A powerful and popular approach for estimating radial lens distortion parameters is to use the fact that lines which are straight in the scene should be imaged as straight lines under the pinhole camera model. This paper revisits this problem using the *division model* to parameterise the lens distortion. This turns out to have significant advantages over the more conventional parameterisation, especially for a single parameter model. In particular, we demonstrate that the locus of distorted points from a straight line is a circular arc. This allows distortion estimation to be reformulated as circle-fitting for which many algorithms are available. We compare a number of suboptimal methods offering closed-form solutions with an optimal, iterative technique which minimises a cost function on the *actual* image plane as opposed to existing techniques which suffer from a bias due to the fact that they optimise a geometric cost function on the *undistorted* image plane.

1 Introduction

Virtually all algorithms based on camera geometry assume that the camera can be approximated by the pinhole model. In reality, however, almost all lenses suffer from small or large amounts of distortion. Failing to compensate for the distortion can lead to severe errors for instance when making metric measurements from images [4], performing self-calibration [16] or computing homographies or matching tensors from multiple views.

Typically lens distortion is split into two components, radial and tangential, of which the radial component is most significant. Radial distortion is a nonlinear transformation of the image along directions from the image centre to the pixel in question, giving either barrel or pin-cushion distortion, as illustrated in fig. 1. We neglect the tangential component and focus on single-parameter radial distortion models since studies [2, 18] have noted that this is sufficient with standard lenses for many tasks in machine vision.

Broadly speaking, methods for calculating distortion parameters may be split into three groups (see e.g. [15]). The first, the plumbline method, assumes the presence of straight lines in the scene [2, 5, 1, 9, 14]. Distortion parameters are sought which lead to the lines being imaged as straight in the undistorted image. Intuitively, this source of information is very strong, provided straight lines can be reliably detected. The second class of approach assumes the availability of correspondences between points in the scene with known 3D coordinates and their projections [18, 19]. With a large number of points this should also be accurate, yet the need for a calibration object or scene limits the flexibility of this approach. The final method seeks distortion parameters that produce multiple-view matching constraints with low residual [8, 12, 16]. Although flexible in the sense that no knowledge of the scene is required, this scheme is not applicable with a single static camera mounted for instance on a wall.

BMVC 2005 doi:10.5244/C.19.9

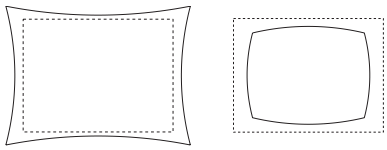

Figure 1: The transformation (solid line) of a rectangle (dotted line) under pin-cushion (left) and barrelling (right) radial distortion.

Our application is in sports analysis where multiple static cameras cover e.g. a football pitch. Correcting distortion is necessary before computing distances on the ground plane, and for creating mosaics from two or more cameras mounted on the same tripod. This scenario is well served by the plumbline method. Markings on the pitch provide excellent input since they can be detected reliably, and are known in advance to be straight.

In [8] Fitzgibbon demonstrated that a less conventional parameterisation of radial distortion, the *division model* leads to linear techniques for the simultaneous estimation of lens distortion and multiple view matching relations such as fundamental matrices or homographies. Inspired by his work, we have applied the same distortion model to the plumbline method, with considerable benefit. Various closed-form methods are derived both when considering a single or multiple lines in the image. Moreover, with the division model, we show that the locus of distorted points forms a circular arc. Thus the problem of distortion calibration is equivalent to circle-fitting, for which an optimal iterative solution is easy to implement. We show that both our closed-form and iterative methods perform better than an "industry-standard" technique [5] since that algorithm is biased because it optimises a distance in the *undistorted* as opposed to the *original* image plane.

The remainder of the paper is organised as follows. Sec. 2 reviews relevant literature. Sec. 3 contains the proof that a distorted line is a circular arc. Algorithms and experiments on synthetic data for single and multiple lines are presented in secs. 4 and 5 respectively. Fully automatic distortion correction on sports imagery is illustrated in sec. 6 before sec. 7 draws conclusions and discusses future work.

More details about this work may be found in [13].

2 Previous work

The most commonly used model for radial distortion is the polynomial model (e.g. [2, 5])

$$p = x(1 + \lambda_1 r^2 + \lambda_2 r^4 + \ldots) \quad , \quad q = y(1 + \lambda_1 r^2 + \lambda_2 r^4 + \ldots) \tag{1}$$

where (p,q) is an undistorted point which satisfies the pinhole camera model, and (x,y) is the corresponding distorted point in the actual image. $\{\lambda_1,\ \lambda_2 ...\}$ are the coefficients of radial distortion and $r^2 = x^2 + y^2$ is the radial distance from the principal point. The assumption here is that the principal point is known and at (0,0), that skew is zero, and that the aspect ratio is unity. The latter two assumptions are reasonable for most digital cameras, and for the purpose of estimating radial distortion, it is safe to set the principal point to the centre of the CCD array (trying to estimate it can be ill-conditioned).

In [8] Fitzgibbon used the division model to parameterise the distortion,

$$p = \frac{x}{1 + \lambda r^2} \quad , \quad q = \frac{y}{1 + \lambda r^2} \ . \tag{2}$$

For barrelling distortion, $\lambda < 0$. If λ is small, the division and single-parameter polynomial models are more or less equivalent, which is easily shown by performing a binomial

expansion of eqn (2). In fact, fitting the division model to real image data proved if anything slightly superior to the conventional polynomial model in experiments in [8, 10].

For estimating distortion coefficients from images of straight lines, the most popular method is to conduct a non-linear optimisation (e.g. Devernay and Faugeras [5]). At each iteration the current estimate of distortion coefficients is used to remove the distortion, and the cost function is then the sum of squared distances from each i of N_j undistorted points $\mathbf{p}_{i,j}$ to each j of M straight lines $\mathbf{l}_j$ fitted by orthogonal regression, i.e.

$$\mathscr{F}_{DF} = \sum_{j=1}^{M} \sum_{i=1}^{N_j} d\,(\,\mathbf{l}_j\,,\,\mathbf{p}_{i,j}\,)^2 \quad . \tag{3}$$

Although optimal with respect to the *undistorted* points, it is not optimal with respect to the points in the *original image*, there is a bias towards coefficients which "shrink" the undistorted image to give a lower cost function. [14] addressed this issue with an iterative scheme based on minimising the distance between measured points and the distorted line. Yet to make the problem tractable their cost function was only approximate, and the method was still rather expensive. Our circle-fitting approach is similar in spirit, but both faster and more accurate due to our *exact* knowledge of the locus of distorted points.

Ahmed and Farag [1] proposed closed-form solutions for estimating radial distortion from straight lines. One method exploits the notion that the slope of the best-fit line to the distorted points is usually close to the slope of the unknown undistorted line. [6] extends the approach by considering model selection and outlier elimination.

3 Distorted straight lines are circular arcs

In this section we prove that under the division model of eqn (2), the distorted images of straight lines form circular arcs. Following [8], eqn (2) may be written, up to a scale factor s, in homogeneous coordinates as

$$s\,\mathbf{p} = s\begin{pmatrix} p \\ q \\ 1 \end{pmatrix} = \begin{pmatrix} x \\ y \\ 1+\lambda(x^2+y^2) \end{pmatrix} = \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} + \lambda \begin{pmatrix} 0 \\ 0 \\ r^2 \end{pmatrix} = \mathbf{x} + \lambda \mathbf{z} \quad . \tag{4}$$

The undistorted point, $\mathbf{p}$, must lie on a straight line denoted $\mathbf{l} = (l_1\ l_2\ l_3)^\top$ in homogeneous coordinates, therefore $\mathbf{l}^\top \mathbf{p} = 0$. Inserting the expression for $\mathbf{p}$ from eqn (4) yields

$$l_1 x + l_2 y + l_3(1+\lambda(x^2+y^2)) = 0 \quad . \tag{5}$$

Assuming (i) that the line does not pass through the origin, $l_3 \neq 0$, and (ii) the presence of radial distortion, $\lambda \neq 0$, we may divide by λl_3, and by completing the square it is easy to show that this represents the equation of a circle $(x-x_0)^2 + (y-y_0)^2 = R^2$ where

$$x_0 = -\frac{l_1}{2\lambda l_3} \quad , \quad y_0 = -\frac{l_2}{2\lambda l_3} \quad , \quad R^2 = \frac{l_1^2+l_2^2}{4\lambda^2 l_3^2} - \frac{1}{\lambda} = x_0^2 + y_0^2 - \frac{1}{\lambda} \quad . \tag{6}$$

The circle degenerates to a line if either $\lambda = 0$ or $l_3 = 0$.

Alternatively the proof and degeneracies may easily be seen by rewriting eqn (5) as the equation of a conic in homogeneous coordinates,

$$\mathbf{x}^\top \mathtt{C} \mathbf{x} = 0 \quad , \quad \mathtt{C} = \begin{pmatrix} c_1 & 0 & c_2 \\ 0 & c_1 & c_3 \\ c_2 & c_3 & c_4 \end{pmatrix} \quad , \tag{7}$$

where $\mathbf{x} = (x_i\ y_i\ 1)^\top$ and in this case $\mathtt{C}$ is not a general conic but a circle.

Interestingly, the images of lines are also circles for paraboloid catadioptric cameras [11] where rays reflected off a parabolic mirror are projected orthographically. Indeed, one may show that such an imaging system gives rise to a projection equation of the same form as with a system comprised of a perspective camera followed by barrelling distortion parameterised by the division model. It is intriguing that two imaging systems which appear so different in their geometry can give rise to similar projection equations.

4 Estimating lens distortion from a single line

In this section we develop a number of techniques for estimating the radial distortion coefficient from a single line. Methods for multiple lines will be discussed later in sec. 5.

4.1 Optimal Circle-Fitting approach (OCF)

Using the result from sec. 3, and assuming isotropic, Gaussian noise in the actual image plane, the optimal value of λ and the line parameters $\mathbf{l}$ may be found by fitting a circle to the data in the original image, minimising the sum of squared distances, d_i, from each i of N points to the circle

$$\mathscr{F}_1 \quad = \quad \sum_{i=1}^{N} d_i^2 = \sum_{i=1}^{N} \left(\sqrt{(x_i - x_0)^2 + (y_i - y_0)^2} - R \right)^2 \quad . \tag{8}$$

Although a problem with three parameters $\{R, x_0, y_0\}$ it is well known that the problem reduces to a two-parameter minimisation by writing the radius R as a function of the circle centre $\{x_0, y_0\}$ by setting the partial derivative of $\mathscr{F}_1$ with respect to R to zero, and inserting the result into the cost function in eqn (8). After minimising this function, the radial distortion parameter λ may be found from the rightmost equation in eqns (6). In the remainder of this paper this method will be referred to as Optimal Circle-Fitting (OCF).

Although one would expect this optimisation to be stable, a good initial estimate of the circle parameters is useful to ensure convergence to the correct minimum, and to reduce the number of iterations. We will therefore derive sub-optimal closed-form solutions. Additionally, such solutions will be much faster than carrying out the non-linear minimisation, and if sufficiently accurate might avoid the need for iterative methods altogether.

4.2 Linear Solution # 1 (LS1)

A first linear approach, which we will denote LS1, is inspired by the algorithm of Fitzgibbon [8] for simultaneously computing the lens distortion and fundamental matrix. For each point i on the line, eqn. (5) gives

$$(x_i \ \ y_i \ \ 1)\mathbf{l} + \lambda (0 \ \ 0 \ \ r_i^2)\mathbf{l} = 0 \ . \tag{9}$$

Stacking equations from N points together gives

$$(\mathtt{D}_1 + \lambda \mathtt{D}_2)\mathbf{l} = 0 \tag{10}$$

where D_1 and D_2 are $N \times 3$ matrices. Premultiplying eqn (10) by $\mathtt{D}_1$ does not alter the solution, and yields a generalised eigenvalue problem

$$(\mathtt{D}_1{}^\top \mathtt{D}_1 + \lambda \mathtt{D}_1{}^\top \mathtt{D}_2)\mathbf{l} \quad = \quad 0 \ . \tag{11}$$

Given the small size of the matrices, and the simple form of $\mathtt{D}_2$, it is readily shown by setting the matrix determinant to zero that the solution for λ is

$$\lambda = \frac{-d_3(d_2d_5 - d_4d_3) - d_5(d_3d_2 - d_5d_1) - d_6(d_1d_4 - d_2^2)}{e_1(d_2d_5 - d_4d_3) + e_2(d_3d_2 - d_5d_1) + e_3(d_1d_4 - d_2^2)} \tag{12}$$

where the coefficients on the right hand side are given by matrices ${\mathtt{D}_1}^\top\mathtt{D}_1$ and ${\mathtt{D}_1}^\top\mathtt{D}_2$ as

$${\mathtt{D}_1}^\top\mathtt{D}_1 = \begin{pmatrix} d_1 & d_2 & d_3 \\ d_2 & d_4 & d_5 \\ d_3 & d_5 & d_6 \end{pmatrix} \quad , \quad {\mathtt{D}_1}^\top\mathtt{D}_2 = \begin{pmatrix} 0 & 0 & e_1 \\ 0 & 0 & e_2 \\ 0 & 0 & e_3 \end{pmatrix} \quad . \tag{13}$$

It turns out that the method derived above is equivalent to Coope's linear method for circle-fitting [3][1]. Although, he derived it through a change of coordinates, his method is most easily explained via the use of homogeneous coordinates which readily turn circle-fitting into a linear problem via the conic equation in eqn (7).

4.3 Linear Solutions # 2 and # 3 (LS2 and LS3)

We now present two further linear methods based on circle-fitting.

Eqn (7) may be rewritten as a set of linear equations $\mathtt{A}\mathbf{c} = 0$ in the unknowns $\mathbf{c} = (c_1\ c_2\ c_3\ c_4)^\top$ where row i of the $N \times 4$ matrix $\mathtt{A}$ is $\begin{bmatrix} r_i^2 & 2x_i & 2y_i & 1 \end{bmatrix}$. Recalling that $\mathtt{C}$ and hence $\mathbf{c}$ is only defined up to a scale factor, we may apply the constraint that $\|\mathbf{c}\| = 1$, in which case the solution that minimises algebraic error is found from the least singular vector of $\mathtt{A}$. This procedure is commonplace in the literature on circle and conic fitting. We denote this by Linear Solution # 2 (LS2).

Applying instead the constraint $c_4 = 1$ gives the set of equations $\mathtt{B}\mathbf{c}' = -\mathbf{1}$ where $\mathtt{B}$ contains the first three columns from $\mathtt{A}$ and $\mathbf{c}' = (c_1\ c_2\ c_3)^\top$. The solution to this over-determined set of equations may be found using the pseudo-inverse. It is safe to apply this constraint since $c_4 = 0$ implies that $l_3 = 0$ and the straight line and circle pass through the origin, yet in practise we will not use any lines close to the origin since they provide very little information (they remain straight even with arbitrarily large amounts of radial distortion). We call this method Linear Solution # 3 (LS3).

4.4 Closed-Form, Approximate Geometric solution (CFAG)

With noise, $\varepsilon = \mathbf{l}^\top\mathbf{p}$ represents the signed distance ε between point $\mathbf{p}$ and the line $\mathbf{l}$, provided $l_1^2 + l_2^2 = 1$ and the third coordinate of $\mathbf{p}$ is unity. With a one-parameter division model, it turns out that it is possible to generate a closed-form solution that minimises an *approximation* to the geometric cost function of eqn (3) commonly used in the literature,

$$\mathscr{F}_2 = \sum_{i=1}^{N} \left(\mathbf{l}^\top(\mathbf{x}_i + \lambda\mathbf{z}_i)\right)^2 \approx \sum_{i=1}^{N} \varepsilon_i^2 \tag{14}$$

subject to the constraint that $l_1^2 + l_2^2 = 1$. As before $\mathbf{l} = (l_1\ \ l_2\ \ l_3)^\top$ describes the line in homogeneous coordinates, $\mathbf{x}_i$ is each point, also in homogeneous coordinates, $(x_i\ \ y_i\ \ 1)^\top$, and $\mathbf{z}_i = (0\ \ 0\ \ r_i^2)^\top$. It is only an approximation to eqn (3) because the third coordinate of the points $(\mathbf{x}_i + \lambda\mathbf{z}_i)$ is $1 + \lambda r_i^2$ which is close to, but not identical to, unity. However, it turns out that this induces a bias towards expanding the undistorted image plane which counteracts the bias of eqn (3) which shrinks the undistorted image plane.

[1] In section 2 of his paper, his matrix $\mathtt{B}$ is our $\mathtt{D}_1$, his vector $\mathbf{d}$ is the third column of our $\mathtt{D}_2$, and his vector of unknowns $\mathbf{y}$ is simply $\mathbf{l}/(\lambda l_3)$. Thereafter it easily follows that his set of equations is identical to our eqn (10).

The derivation, given in full in [13], is not entirely trivial and is omitted here due to lack of space. Instead we merely state the result which is

$$\lambda = -\frac{e_1\, l_1 + e_2\, l_2 + e_3 l_3}{f l_3} \quad , \quad l_1 = \pm \mathrm{sign}(\alpha_1) \frac{\sqrt{1 \mp \frac{\alpha_2}{\sqrt{4\alpha_1^2+\alpha_2^2}}}}{\sqrt{2}} \tag{15}$$

$$l_2 = \sqrt{1-l_1^2} \quad , \quad l_3 = \frac{-(e_1e_2 - d_2 f)(l_1^2 - l_2^2) - (e_2^2 - e_1^2 + (d_1 - d_4)f) l_1 l_2}{(e_2 e_3 - d_5 f) l_1 - (e_1 e_3 - d_3 f)\, l_2} \tag{16}$$

$$\alpha_1 = -e_3^2 d_2 + e_2 e_3 d_3 + e_1 e_3 d_5 - e_1 e_2 d_6 - d_3 d_5 f + d_2 d_6 f \;, \tag{17}$$

$$\alpha_2 = e_3^2(d_1 - d_4) - 2e_3(e_1 d_3 - e_2 d_5) + (e_1^2 - e_2^2) d_6 + [d_3^2 - d_5^2 - (d_1 + d_4) d_6] f \;, \tag{18}$$

where d_k and e_k were defined in eqns (13) and f is the (3,3) element of the matrix ${D_2}^\top D_2$. There are two solutions, representing the maximum and minimum of the cost function.

4.5 Experiments

We now present experiments on synthetic data comparing the five algorithms presented in this section and the iterative method of Devernay and Faugeras [5] discussed in sec. 2.

A first experiment (fig. 2a-c) tests robustness to increasing levels (standard deviation σ) of isotropic Gaussian noise applied to the distorted points. The true distortion λ_{true} is constant at -10^{-7} which represents a fairly small amount of barrelling distortion, typical of lenses we often use. We generated lines of random orientation, length 300 – 768 pixels, and distance from the origin of 40 – 384 pixels. Data points along each line were separated by 1 pixel, giving many points from which to compute λ. Points outside a 960×960 image plane were removed. Results from 2000 trials are shown as (a) RMS relative error in recovery of λ: $(\lambda - \lambda_{\text{true}})/\lambda_{\text{true}})$, (b) RMS residual distance from image points to circles found from each method, and (c) RMS error in the predicted location of undistorted image points, relative to those obtained with λ_{true} and without noise. For evaluating (c) a regular grid of points across the entire image was used. In a second set of experiments, (fig. 2d), noise was constant at 1 pixel, and now the amount of distortion λ varied at logarithmic intervals from extreme barrel to extreme pin-cushion distortion.

Almost all algorithms give good results, confirmed especially by the fact that the residuals in fig. 2b are indistinguishable from the theoretical lower bound which here is well approximated by $\sigma/\sqrt{2}$. (The factor of $\sqrt{2}$ arises because we added noise such that the expected distance between original and noisy points was σ in a random direction, yet for these experiments it is only the direction perpendicular to the line which is relevant.) The performance of Linear Solutions #2 and #3 (LS2, LS3) is remarkable, and little seems to be gained by applying the the optimal solution (OCF). We confirmed that there is a bias in the geometric solution of Devernay and Faugeras (DF). Our Closed-Form Approximate Geometric solution (CFAG) is better since its two biases almost cancel each other out. Yet both DF and CFAG prove inferior to Linear Solutions #2 and #3.

Notice, however, the very poor performance of Linear method # 1 (LS1), equivalent to Coope's method of circle-fitting, which disappears off the graphs for anything other than very small noise levels. This may be explained by splitting the cost function of that method into the product of two components, as in [3],

$$\mathscr{F}_3 = \sum_i^N |R - \|\mathbf{x}_i - \mathbf{x}_0\|_2 \,| \times |R + \|\mathbf{x}_i - \mathbf{x}_0\|_2 \,| \;\;, \tag{19}$$

where $\mathbf{x}_i = (x_i, y_i)^\top, \mathbf{x}_0 = (x_0, y_0)^\top$ are inhomogeneous coordinates and $\|\cdot\|_2$ indicates the Euclidean distance. The first component is the orthogonal distance to the circle, which

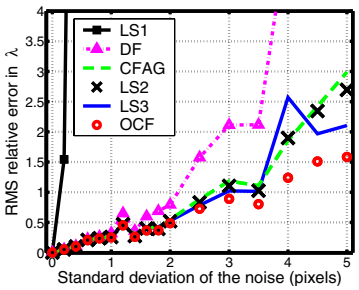

(a) Varying noise: RMS error in $\frac{\lambda-\lambda_{\text{true}}}{\lambda_{\text{true}}}$

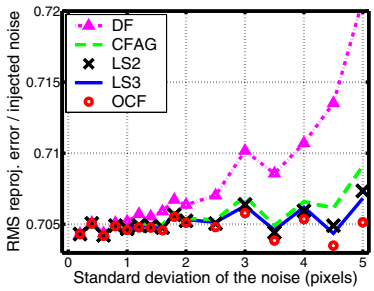

(b) Varying noise: RMS reprojection error
LS1 was off the scale of this graph

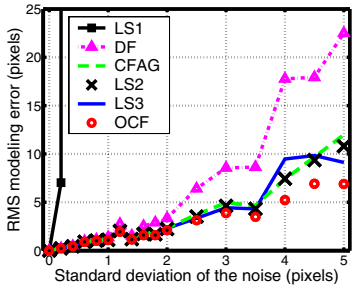

(c) Varying noise: RMS prediction error

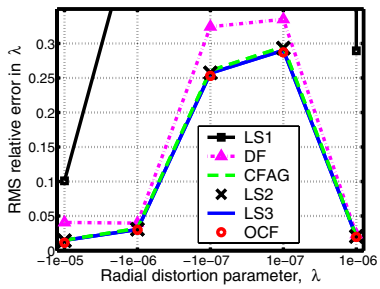

(d) Varying λ: RMS error in $\frac{\lambda-\lambda_{\text{true}}}{\lambda_{\text{true}}}$

Figure 2: Experiments on synthetic data, estimating radial distortion from a single line. (a-c) represent experiments with constant $\lambda = -10^{-7}$ and varying noise, (d) with varying λ and constant noise of 1 pixel. Algorithm key: DF (Devernay Faugeras [5]), LS1–3 (Linear Solution # 1–3, this paper), CFAG (Closed-Form Approximate Geometric, this paper), OCF (Optimal circle-fitting, this paper)

is clearly desirable to minimise. The second term is the distance from the point $\mathbf{x}_i$ to the furthest point on the circle. The problem is that our application is based on very small arcs, such that a considerable decrease in the radius of the circle yields a reduction in the second term capable of overpowering the increase in orthogonal distance in the first term.

5 Estimating lens distortion from multiple lines

5.1 Optimal technique based on circle-fitting (OCF)

The optimal circle-fitting approach readily extends to the case where multiple straight lines are detected. Now the cost function is a sum over all points on all of M lines,

$$\mathscr{F}_4 = \sum_{j=1}^{j=M} \sum_{i=1}^{N_j} d_{i,j}^2 \quad , \tag{20}$$

where $d_{i,j}^2$ is the orthogonal distance from point i out of N_j on the jth circle to that circle. We have $2M+1$ degrees of freedom, 2 for each line and 1 for λ. The non-linear minimisation may for instance be parameterised in terms of $\{x_0, y_0\}$ for each circle and λ. It is easy to show that the Jacobian is sparse since the circle centre for line j' does not depend on points from a circle $j \neq j'$. This may be used to provide a speed-up, in the first instance

by computing the Jacobian explicitly rather than relying on numerical differentiation, and second by using sparse techniques (as in bundle-adjustment [17]).

5.2 Closed-Form Multiple Line solution (CFML)

With multiple lines it is harder to produce closed-form solutions without over-parameterising the problem (which essentially corresponds to finding a separate value of λ for each line). We will, however, derive one closed-form solution. Although suboptimal, it turns out that this method has desirable properties, and in practise yields excellent performance.

While [1] argued for assuming the *orientation* of the line is known, we instead assume that the *distance from the origin*, l_3 is known. We minimise the cost function

$$\mathscr{F}_5 = \sum_{j=1}^{j=M} \sum_{i=1}^{N_j} \left(\mathbf{l}_j{}^\top \left(\mathbf{x}_{i,j} + \lambda \mathbf{z}_{i,j} \right) \right)^2 \tag{21}$$

subject to the constraint that $l_{3,j}$ is known for each line j. The derivation of an expression for λ is relatively straightforward, and may be found in [13]:

$$\lambda = -\frac{\sum_{j=1}^{j=M} l_{3,j}^2 \left(\frac{e_{1,j}(d_{2,j}d_{5,j} - d_{3,j}d_{4,j}) + e_{2,j}(d_{2,j}d_{3,j} - d_{1,j}d_{5,j})}{d_{4,j}d_{1,j} - d_{2,j}^2} + e_{3,j} \right)}{\sum_{j=1}^{j=M} l_{3,j}^2 \left(\frac{e_{1,j}(d_{2,j}e_{2,j} - d_{4,j}e_{1,j}) + e_{2,j}(d_{2,j}e_{1,j} - d_{1,j}e_{2,j})}{d_{4,j}d_{1,j} - d_{2,j}^2} + f_j \right)} \tag{22}$$

The subscript j denotes the line, the coefficients $d_{k,j}$ and $e_{k,j}$ depend only on the measurements and were defined in eqn (13), though we now have one matrix $D_{1,j}$ for each line, and the same applies to $D_{2,j}$. f_j is the (3,3) element of the matrix ${D_{2,j}}^\top D_{2,j}$.

Eqn (22) may be interpreted as a set of equations $l_{3,j}^2 \;\; g_j \;\; \lambda = l_{3,j}^2 \;\; h_j$ where each equation j represents a solution for λ derived from one line, and g_j and h_j are known and depend only on image measurements, and not on the unknowns $\mathbf{l}_j$ or λ. Thus the *only* effect of the $l_{3,j}$ is to weight the equations from each line. Lines close to the origin (small $l_{3,j}$) have little influence on the solution, as do short lines (small N_j). Note also that with a single line, $M = 1$, the solution is independent of l_3. Thus, we would be free to set $l_3 = 1$, which is essentially what Linear Solution # 3 for a single line does (sec. 4.3).

We obtain estimates for $l_{3,j}$ via orthogonal regression on the line parameters in the distorted image plane. Experiments will show that this crude method is remarkably accurate, even when large amounts of distortion clearly give inaccurate estimates of $l_{3,j}$. A second approach (not explored experimentally in this paper) is to perform orthogonal regression on an undistorted image where λ has been estimated from one line only using one of the closed-form solutions discussed in sec. 4.

This also suggests an iterative scheme whereby the latest value of λ is used to undistort the image, after which the $l_{3,j}$ are found by orthogonal regression. Then λ may be recomputed using these values and so forth. In practise we did not find this fruitful, presumably because results are already so accurate, so we do not report those results further.

5.3 Experiments

In figs. 3a and b we repeat the experiments of sec. 4.5, but now using 20 lines to compute the solutions. Comparing with figs. 2a and d, using multiple lines improves results. Our two new algorithms, optimal circle-fitting (OCF) and Closed-Form Multi-Line (CFML) are compared to the Devernay-Faugeras scheme (DF) and the linear solution of Ahmed

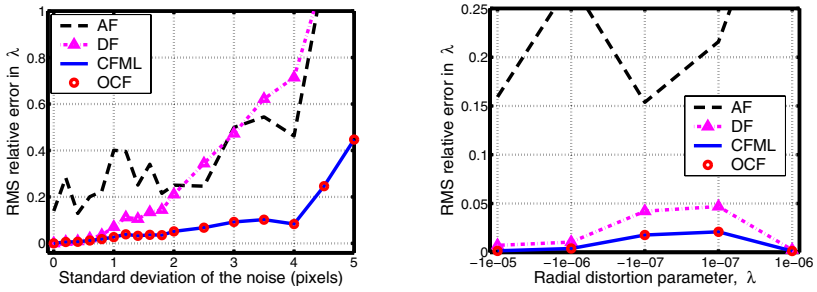

(a) Varying noise. $\lambda = \text{const.} = -10^{-7}$

(b) Varying λ. Noise = const. = 1 pixel

Figure 3: Experiments estimating radial distortion from 20 lines. Algorithm key: AF (Ahmed and Farag [1]), DF (Devernay and Faugeras [5]), CFML (Closed-Form Multi-Line, this paper), OCF (Optimal Circle-Fitting, this paper)

(a) Original image (brightened) and inlying points

(b) Corrected image

(c) Plots of the line before and after distortion correction

Figure 4: Fully-automatic correction of radial distortion on sports imagery.

and Farag (AF, [1] sec.4.2). Our novel techniques are clearly superior, and the closed-form solution is excellent, despite the crude estimates of $l_{3,j}$. AF performed poorly, but we did not implement the refinement stages they suggest, and their scheme can handle any number of distortion coefficients. Yet we doubt that their ad-hoc cost functions would be as effective as ours, and their linear solution is actually rather computationally expensive.

6 Application: sports images

We now demonstrate fully automatic correction of radial distortion in a practical application. Markings from a football pitch are automatically detected via a second-order Gaussian derivative operator. A potential problem is that each line may be split into separate fragments due to occlusions. Moreover, the detection procedure finds not only pitch markings, but also other lines. We apply RANSAC [7] which can both eliminate outliers and link multiple fragments of the same line. The minimal solution for RANSAC may be provided by any of the closed-form solutions, and the final value for λ is found by applying Linear Solution # 3 to the final set of inliers. Currently we only search for a single line. An example image is shown in fig. 4a along with the detected sideline of the pitch. Due to occlusions (people) the sideline is split into four segments, but the robust fitting scheme automatically identifies the segments as arising from the same straight line. The corrected image is given in fig. 4b while a graph showing the original and corrected points is provided in fig. 4c. Clearly, the distortion has been removed very well indeed.

7 Discussion

This paper revisited the plumbline method for estimating radial distortion by using the division model to parameterise the distortion. A number of techniques were developed, including closed-form solutions, some of which gave excellent performance, and an iterative scheme which is optimal under Gaussian noise in the actual image plane. This latter algorithm is based on the insight that the locus of a distorted straight line becomes the arc of a circle. Such an optimal method does not appear available when using the standard parameterisation for distortion, nor are closed-form solutions as easy to derive. We showed the superiority of our approaches over conventional methods on synthetic data, and a real-world application for sports imagery was also presented.

Although this paper assumed known aspect ratio and zero skew, it is possible to relax these constraints and even estimate these parameters from the observed distortion. This is an interesting topic for future research. With non-unity aspect ratio, the circle becomes an ellipse with principal axes in the x and y directions, and it is not difficult to derive linear or optimal iterative techniques.With non-zero skew the orientation of the ellipse changes. The principal point can easily be added to the unknowns in the optimal nonlinear algorithm. It remains an open question whether the recovery of these additional parameters will prove reliable or worthwhile in the presence of noise.

Another extension is to go beyond the single-parameter model by including two or more terms in the denominator in eqn (2). We have confirmed the existence of linear solutions in this case, but the locus of distorted points is no longer as simple, so optimal methods may not easily be obtained.

References

[1] M.T. Ahmed and A.A. Farag. Differential methods for nonmetric calibration of camera lens distortion. In *Proceedings CVPR*, pages II:477–482, 2001.

[2] D.C. Brown. Close-range camera calibration. *Photogrammetric Engineering*, 37(8):855–866, 1971.

[3] I.D. Coope. Circle fitting by linear and nonlinear least squares. *J. Opt. Th. and Applic.*, 76:381–388, 1993.

[4] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. *IJCV*, 40(2):123–148, Nov 2000.

[5] F. Devernay and O.D. Faugeras. Straight lines have to be straight. *MVA*, 13(1):14–24, 2001.

[6] M.T. El-Melegy and A.A. Farag. Nonmetric lens distortion calibration: Closed-form solutions, robust estimation and model selection. In *Proc. ICCV*, pages 554–559, 2003.

[7] M. A. Fischler and R. C. Bolles Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. Assoc. Comp. Mach.* **24**(6), 381–395, 1981.

[8] A. W. Fitzgibbon. Simultaneous linear estimation of multiple view geometry and lens distortion. In *Proc. CVPR*, 2001.

[9] S.B. Kang. Radial distortion snakes. In *IAPR Workshop on Machine Vision Applications (MVA2000)*, pages 603–606, Tokyo, Japan, November 2000.

[10] L. Ma, Y. Chen, and K.L. Moore. Rational radial distortion models with analytical undistortion formulae, 2003. http://arxiv.org/abs/cs.CV/0307047.

[11] S.A. Nene and S.K. Nayar. Stereo with mirrors. In *Proc. ICCV*, pages 1087–1094, 1998.

[12] G.P. Stein. Lens distortion calibration using point correspondences. In *Proc CVPR*, pages 602–608, 1997.

[13] R. Strand. Masters thesis, 2005. Available via `www.nada.kth.se/~hayman/BMVC2005`.

[14] R. Swaminathan and S.K. Nayar. Nonmetric calibration of wide-angle lenses and polycameras. In *Proc. CVPR*, pages II: 413–419, 1999.

[15] B. J. Tordoff. *Active Control of Zoom for Computer Vision*. PhD thesis, University of Oxford, 2002.

[16] B. J. Tordoff and D. W. Murray. Violating rotating camera geometry: The effect of radial distortion on self-calibration. In *Proc. ICPR*, pages I: 423–427, 2000.

[17] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon. Bundle adjustment — A modern synthesis. In *Proc. Vision Algorithms: Theory and Practice*, Sep 1999.

[18] R.Y. Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *RA*, 3(4):323–344, 1987.

[19] Z. Zhang. A flexible new technique for camera calibration. In *Proc. ICCV*, pages 666–673, 1999.