Fast Solvers for Minimal Radial Distortion Relative Pose Problems

Magnus Oskarsson
Centre for Mathematical Sciences
Lund University, Sweden
magnus.oskarsson@math.lth.se

Abstract

In this paper we present a unified formulation for a large class of relative pose problems with radial distortion and varying calibration. For minimal cases, we show that one can eliminate the number of parameters down to one to three. The relative pose can then be expressed using varying calibration constraints on the fundamental matrix, with entries that are polynomial in the parameters. We can then apply standard techniques based on the action matrix and Sturm sequences to construct our solvers. This enables efficient solvers for a large class of relative pose problems with radial distortion, using a common framework. We evaluate a number of these solvers for robust two-view inlier and epipolar geometry estimation, used as minimal solvers in RANSAC.

1. Overview and motivation

We will in this paper present a framework for handling a large class of relative pose problems. Specifically, we will look at problems where at least one of the cameras has unknown radial distortion. Our own motivation comes from the need for bootstrapping solvers in robust feature matching. In most standard Structure-from-motion (Sfm) pipelines – based on unordered images – a very expensive and time-consuming part is the feature matching across views. This is due to the combinatorial explosion of possible matches over all frames [1, 32, 33]. In order to alleviate this, one can devise various heuristics. One such, is the use of a number of keyframes, where the matching is only done from all images to these keyframes. In Figure 1 part of such a scenario is depicted, where we want to match the two bottom frames to the top frame. For the keyframe scenario to work, we need to be able to match a large number of images to each keyframe. For this reason, the keyframes should ideally cover a large field-of-view. One way of accomplishing this is to use a lens with large radial distortion. However, this also complicates the camera model, and in turn gives more complicated relative pose problems. Depending on the knowledge of partial or full calibration of the cameras, we will encounter different relative pose formulations. Our goal in this paper is to try to formulate a large number of such problems using the same basic language. This will enable us to solve the problems in a more systematic way than has been previously done. We will throughout the paper use the division model as described by Fitzgibbon [9], where he used it to formulate a linear method for constant radial distortion estimation. There are of course more involved models for radial distortion, but our main goal here is to use the solvers for robust inlier estimation, and hence we prioritize simplicity over accuracy.

Figure 1: A structure from motion scenario; Top shows a keyframe with heavy radial distortion, and bottom shows two example frames that we want to match to the keyframe. The lines show found inliers matches using our proposed solvers.

The constant radial distortion relative problem was also solved in [26], based on a hidden variable formulation, but no calibration constraints on the fundamental matrix were enforced in these formulations. The method was later extended to fisheye cameras [28]. Minimal solutions for the same problem were then developed in [19, 21]. A non-minimal solution to the same problem, but with varying radial distortion, was proposed in [3] . The solution was based on using an extended 4×4 fundamental matrix, and the linear solution then needed at least 15 point correspondences. The corresponding minimal problem was solved in [20] but used exact arithmetic to solve it. A Gröbner based method was presented in [4]. A number of other minimal problems have been formulated and solved over the years [14–16] and also improved on using more recent methods for automatic solving of polynomial equations [18, 22, 25]. An overview of the current state-of-the-art solvers is given in Table 1. Solvers for radial distortion models have also been developed for other problems than relative pose, such as e.g. stitching [31] and absolute pose [17, 24, 30].

The main contribution of this paper is a common framework that can be used for a large class of radial distortion relative pose problems. Using the division model, we show in Section 3 how one can systematically reduce the number of parameters of the given problem. In Section 4 a systematic overview of solvers to all possible minimal cases is given. We show how the reduced number of parameters can be used to construct Sturm sequence solvers based on the action matrix method. These solvers can all be constructed using the same basic principles, and gives us the means to easily produce solvers to a number of unsolved cases, as well as producing competitive solvers in terms of execution speed. The Matlab-Mex implementations are publicly available[1].

2. Problem description

We will use upper case letters to denote matrices, and lower case bold to denote vectors. The general radial distortion relative pose problem can be formulated in the following way,

Problem 1 (Radial Distortion Relative Pose)**.** *Given a number of two-dimensional image point correspondences* $(\mathbf{x}_i, \mathbf{x}'_i)$ *for* $i = 1, \ldots, n$*, find the fundamental matrix* F *and the radial distortion parameters* λ *and* λ'*, such that*

$$g(\mathbf{x}'_i, \lambda')^T F g(\mathbf{x}_i, \lambda) = 0, \quad i = 1, \ldots n, \tag{1}$$

$$F \in \mathcal{P}, \tag{2}$$

where $\mathcal{P}$ *defines the calibration-manifold of possible* F*, and* g *defines the homogeneous radial coefficients using the division model [9],*

$$g(\mathbf{x}, \lambda) = \begin{bmatrix} \mathbf{x} \\ 1 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \lambda \|\mathbf{x}\|^2 \end{bmatrix}. \tag{3}$$

We are especially interested in *minimal cases* where Problem 1 has a finite number (larger than zero) of solutions. For such cases we need the number of point constraints to equal the number of degrees of freedom of $\mathcal{P}$ plus the number of radial distortion parameters, i.e. $n = \dim(\mathcal{P}) + \#\lambda$, where we let $\#\lambda$ denote the number of radial distortion parameters.

The fully calibrated fundamental matrix—the essential matrix—is traditionally denoted E. For this reason we will, when we speak in general terms, denote the fundamental matrix F (with entries f_{ij}), but when we address the fully calibrated case denote it E (with entries e_{ij}). For a thorough investigation of different calibration cases, and how they are related please see e.g. [12]. We are mainly interested in four different calibration cases involving different knowledge of the focal lengths of the two cameras, namely

- $\mathcal{P}_E$ Fully calibrated essential matrix.
- $\mathcal{P}_{fE}$ One-sided unknown focal length.
- $\mathcal{P}_{fEf}$ Two-sided unknown constant focal length.
- $\mathcal{P}_{f'Ef}$ Two unknown focal lengths.

The 3×3 fundamental matrix is determined up to an arbitrary non-zero scale factor. All the different calibration cases can be described using a set of additional polynomial constraints in the entries of F. For the calibrated case, these are the well-known trace constraints [7, 27, 34]

$$\mathcal{P}_E = \{E \in \mathbb{P}^{3\times3} | \det E = 0,\ 2EE^TE - \mathrm{tr}(EE^T)E = 0\}. \tag{4}$$

The case with two unknown focal lengths is equivalent to the general projective case, i.e.

$$\mathcal{P}_{f'Ef} = \mathcal{P}_F = \{F \in \mathbb{P}^{3\times3} | \det F = 0\}. \tag{5}$$

In [18] the constraints for $\mathcal{P}_{fE}$ and $\mathcal{P}_{fEf}$ were given. Note that due to the symmetry in (1) the constraints for $\mathcal{P}_{fE}$ are essentially the same as for $\mathcal{P}_{Ef}$. The manifold for $\mathcal{P}_{fE}$ is defined by three quartics in addition to the determinant constraint. The manifold for $\mathcal{P}_{fEf}$ is defined by a quintic polynomial and the determinant constraint. For the remainder of the paper, we assume that any known calibration parameter is set to its canonical value, i.e. $f = 1, \lambda = 0$.

Depending on the knowledge of the radial distortion parameters, we end up with a large number of different possible calibration cases. We denote the different cases by the

[1] The code for all presented solvers is publicly available at *https://github.com/hamburgerlady/fast-radial-solvers*

unknown calibration parameters, e.g. if we have unknown radial distortion in the second camera, and a constant focal length we denote this by $\lambda f E f$. This example and other combinations might seem exotic, but many combinations may occur in practice. However, we believe that there is a merit in enumerating all possible combinations for completeness. All possible combinations are summarized in Table 1.

We will now describe some important properties of our approach on a given example problem. It follows several previous approaches for solving radial distortion problems [4, 14].

Example 1 (One-sided radial distortion calibrated)**.** *Let's assume that we have two images taken by cameras with known calibration, but where camera one has an unknown radial distortion, i.e. the case $E\lambda$. This means that in our case a point correspondence $(\mathbf{x}_i, \mathbf{x}'_i)$ should fulfill the epipolar constraint*

$$g(\mathbf{x}'_i, 0)^T E g(\mathbf{x}_i, \lambda) = 0. \tag{6}$$

Since each point correspondence gives one constraint, and E has five degrees of freedom and we have an additional parameter λ we would minimally need six point correspondences to solve for E and λ. The constraints in (6) are non-linear if we consider all variables, but we can still write these equations as an under-determined linear system in the variable monomials,

$$C_{6\times 12}(\mathbf{x}, \mathbf{x}') \begin{bmatrix} e_{11} \\ \vdots \\ e_{32} \\ e_{13} \\ e_{23} \\ e_{33} \\ \lambda e_{13} \\ \lambda e_{23} \\ \lambda e_{33} \end{bmatrix} = \mathbf{0}, \tag{7}$$

where we have purposefully ordered the monomials in a special way. Since we have six equations, we can linearly express the first six of the monomials in terms of the last six, by means of Gauss-elimination operations on C. This gives us a parametrization of E in only the last six monomials of (7). If we in addition set $e_{33} = 1$ to fixate the scale we end up with only three unknown parameters $(e_{13}, e_{23}, \lambda)$. We further know that a valid essential matrix should fulfill a number of non-linear polynomial constraints, namely

$$\mathcal{P}_E = \{\det E = 0,\, 2EE^TE - \operatorname{tr}(EE^T)E = 0\}. \tag{8}$$

Inserting our newly parametrized E into $\mathcal{P}_E$ gives ten polynomial constraints with total degrees of five and six in the three unknowns. Previous approaches [15, 25] use action-matrix methods to solve this set of equations. Typically, the most time consuming part of action-matrix based solvers is solving the elimination template, and hence the size of the template often reflects the complexity of the solver. The state-of-the-art solver in terms of speed was reported in [25] with a template of size 14×40.

In this paper we propose to solve these types of relative pose problems by reformulating the elimination stage in order to reduce the number of parameters in a systematic and optimal way. The previous example is arguably one of the simpler of the cases, and it is maybe clear how to do the elimination in the best way. However, for many of the more complicated cases there are different possible ways of ordering the monomials that lead to very different parametrizations.

For this example there are in general 26 solutions. Many of these solutions are often complex, but we're usually not interested in those. For most of the relative pose problems that we investigate in this paper, we have a rather large number of possible solutions. For this reason, we throughout base our solvers on Sturm sequences, and only search for real solutions.

3. Parameter elimination

We will now show how the process in the previous example can be formulated in a more general way. The key idea is to separate the unknowns of the problem into radial distortion parameters and fundamental matrix parameters. We start by noting that we can reformulate Problem 1 in the following way

Problem 2 (Bilinear Formulation)**.** *Given a number of two-dimensional image point correspondences $(\mathbf{x}_i, \mathbf{x}'_i)$ for $i = 1, \ldots, n$, find F, λ and λ' such that*

$$\Lambda^T X_i \mathbf{f}, \quad i = 1, \ldots n, \tag{9}$$

$$F \in \mathcal{P}, \tag{10}$$

where $\mathcal{P}$ defines the calibration-manifold of possible F, and $\mathbf{f}$ contains the column stacked entries of F,

$$\mathbf{f} = \begin{bmatrix} f_{11} & f_{21} & \cdots & f_{23} & f_{33} \end{bmatrix}^T. \tag{11}$$

The radial distortion parameters are contained in

$$\Lambda = \begin{bmatrix} 1 & \lambda & \lambda' & \lambda\lambda' \end{bmatrix}^T, \tag{12}$$

and X_i only depends on the point correspondence $(\mathbf{x}_i, \mathbf{x}'_i)$.

For each point correspondence $(\mathbf{x}_i, \mathbf{x}'_i)$, $\Lambda^T X_i$ is a 1×9 vector. If we stack these vectors, for all correspondences, we get an $n \times 9$ matrix $A(\lambda, \lambda')$, so that

$$A(\lambda, \lambda')\mathbf{f} = 0. \tag{13}$$

Table 1: Overview of all minimal radial distortion relative pose problems and solvers. First column shows the number of point correspondences for the minimal case. The next three columns denote the names of the problems, with symmetric cases in the middle. Then comes the number of solutions for the minimal problem, the previous state-of-the-art solver template size and the execution time, respectively. The four last columns show our contribution, in terms of the number of parameters that the problem can be formulated in, using (14), the new number of solutions, the new template size and the execution time.

# Pts	Name of problem			# Sol.	SOTA template	t (ms)	# Var.	# Sol.	New template	t (ms)
6	λE		$E\lambda$	26	$\mathbf{14 \times 40}$ [25]	0.20	3	26	$\mathbf{14 \times 40}$	**0.050**
6		$\lambda E \lambda$		52	$\mathbf{53 \times 115}$ [25]	**1.18**	3	∞	-	-
7	$\lambda f E$		$E f \lambda$	19	24×43 [25]	0.11	2	27	$\mathbf{8 \times 35}$	**0.033**
7	$\lambda E f$		$f E \lambda$	23	-		2	32	$\mathbf{13 \times 45}$	**0.045**
7	$\lambda f E f$		$f E f \lambda$	37	-		2	52	$\mathbf{28 \times 80}$	**0.14**
7	$\lambda f E \lambda$		$\lambda E f \lambda$	42	-		2	66	$\mathbf{19 \times 85}$	**0.215**
7		$\lambda f E f \lambda$		68	581×658 [22]	19.2	2	113	$\mathbf{51 \times 164}$	**1.01**
7		$\lambda' E \lambda$		76	-		3	76	$\mathbf{851 \times 927}$	**47.3**
8	$\lambda f' E f$		$f' E f \lambda$	8	7×16 [25]	0.024	1	8	$\mathbf{1 \times 9}$	**0.0033**
8	λF		$F \lambda$	8	7×16 [25]	0.024	1	8	$\mathbf{1 \times 9}$	**0.0033**
8		$\lambda f' E f \lambda$		16	32×48 [22]	0.080	1	16	$\mathbf{1 \times 17}$	**0.0055**
8		$\lambda F \lambda$		16	32×48 [22]	0.080	1	16	$\mathbf{1 \times 17}$	**0.0055**
8	$\lambda' f E \lambda$		$\lambda' E f \lambda$	56	-		2	128	$\mathbf{26 \times 154}$	**3.20**
8		$\lambda' f E f \lambda$		104	-		2	224	$\mathbf{102 \times 326}$	**11.3**
9		$\lambda' f' E f \lambda$		24	84×117 [25]	0.30	2	48	$\mathbf{51 \times 99}$	**0.21**
9		$\lambda' F \lambda$		24	84×117 [25]	0.30	2	48	$\mathbf{51 \times 99}$	**0.21**

By performing row-operations (i.e. Gauss elimination) on (13) we can eliminate a number of unknowns in $\mathbf{f}$. In this way we can use as many as possible of the point constraints to eliminate unknowns. This gives a new system, equivalent to (13),

$$B(\lambda, \lambda')\mathbf{f} = 0, \tag{14}$$

where the entries of $B(\lambda, \lambda')$ are rational polynomials in (λ, λ'). Since F is only determined up to scale, we can multiply F with the greatest common divisor of $B(\lambda, \lambda')$ (assuming this is non-zero) ending up with a polynomial expression in (λ, λ'). The polynomial degree in (λ, λ') will depend both on the addressed problem and how we order the entries. In order to have as low degree as possible, we want to order the columns in A in increasing powers of (λ, λ') from left to right. After elimination we have increased the degree in the radial distortion parameters, going from $A(\lambda, \lambda')$ to $B(\lambda, \lambda')$. However, (14) will still be *linear in F*. For a minimal case we have $n = \dim(\mathcal{P}) + \#\lambda$ constraints in (13). This means that we can linearly eliminate up to $\dim(\mathcal{P}) + \#\lambda$ parameters from the initial eight parameters of F, using (14). In addition to the remaining parameters of F we have a number of radial distortion parameters $\#\lambda$. We have now used all the linear constraints, and the only constraints left are the calibration constraints defined by $\mathcal{P}$. This gives us the following

Theorem 1. *A minimal Problem 1 can be written using the constraints of $\mathcal{P}(F)$, where the entries of F are polynomials in*

$$k = \#\lambda + \max(0, 8 - \dim(\mathcal{P}) - \#\lambda)$$

parameters, given by (14).

These results give us tools for finding low parametric formulations in a systematic way for radial relative pose problems. The number of parameters k varies between one and three, depending on the specific problem. The different numbers of parameters are given in Table 1.

We will now show how to use the proposed approach on a given example.

Example 2 (One sided radial distortion uncalibrated $F\lambda$)**.** *A solution to this case was given in [15] and later improved upon using the technique of [25]. Now, according to Theorem 1, we should be able to directly express this problem using $k = \#\lambda + \max(0, 8 - \dim(\mathcal{P}) - \#\lambda) = 1 + \max(0, 8 - 7 - 1) = 1$ parameter. This means that we should be able to write the problem as a univariate polynomial in the radial distortion parameter λ. In this minimal case we need eight point correspondences. We have only one radial distortion parameter, so $\Lambda = [1\ \lambda]^T$, and $A(\lambda)$ is an $8 \times 9-$matrix. Due to the special structure of X in Problem 2, we can re-order the columns of the matrix A so that it can be written as*

$$A = A_1 + \lambda \begin{bmatrix} \mathbf{0}_{8\times 6} & A_2 \end{bmatrix}. \tag{15}$$

After row operations we get

$$B = \begin{bmatrix} 1 & 0 & \dots & 0 & b_1 \\ 0 & & & & \\ \vdots & & \ddots & & \\ 0 & & & 1 & b_8 \end{bmatrix}, \tag{16}$$

where each

$$b_i = \frac{p_i(\lambda)}{q(\lambda)}, \tag{17}$$

and where the different p_i are third and second degree polynomials in λ. The denominator q is of degree two. Note that the denominator is the same for all rows. Using B we can express eight of the entries of F in terms of the last entry of F. We fix the scale of F by setting $f_{33} = 1$. This means that after multiplication with q, the entries of F can all be written as third or second degree polynomials in λ. We can now find λ by inserting our F into $\mathcal{P}_F$, i.e. $\det(F) = 0$*, which directly gives an eight-degree polynomial in λ. The fundamental matrix F can then be linearly extracted using the solution of λ. If we are considering the unknown calibrated case $f'Ef\lambda$, then f and f' can be extracted from F using the technique described in Appendix A.*

A similar approach can be used for all possible relative pose problems with unknown radial distortion. We will in the next section show how we apply our technique to the different cases. There is one issue with the approach, namely the denominator q in (17). When we multiply with q we may introduce spurious solutions, corresponding to $q = 0$.. The severity of the introduced extra solutions varies with the problem – for some only a small number of extra solutions are introduced but for others an infinite number is added. The total number of possible solutions (including both spurious and complex) is given in Table 1. We will address this issue in the next section.

4. A comprehensive list of radial distortion relative pose problems

In the previous section we saw that we can write the equations of all considered relative pose problems in a small number of variables (ranging from one to three). One way of solving these problems would be to use these set of equations in an automatic generator such as [16, 22]. The typical way of extracting the solutions is to construct the action matrix and then solve the corresponding Eigenvalue problem. This will also find all complex solutions. For problems with many solutions, and where many of the solutions are complex a non-negligible time is spent on finding (uninteresting) complex solutions. This fact is exacerbated in our approach where we have additional spurious solutions (indeed if the number of spurious solutions is infinite we cannot directly use the action matrix method at all).

We have investigated two approaches for constructing our solvers based on the proposed new parametrizations.

4.1. Resultant based solvers

The low number of variables makes it possible to use elimination techniques to reduce the problems to univariate polynomials. Here, the extra solutions can be handled efficiently at runtime.

For the simple case of two variables, there is a direct formula for eliminating one variable, based on resultant theory [5]. The theory of resultants can be generalized to more variables, but it directly becomes much more complicated. Note that for our problems we need to eliminate at most two variables for all cases. The final univariate polynomial is given by some determinant of a matrix with entries that are polynomials in the hidden variable. Although we can readily describe this determinant, it can be difficult to actually compute. In many cases it is much faster to compute the determinant at runtime based on the input data, and not in closed form. One such approach is described in [11]. Here the authors used it to formulate hidden variable solutions to five point relative calibrated pose and six point relative pose with an unknown focal length. The determinant is calculated at runtime (based on methods described in [2, 6]), and then the univariate polynomial is solved using Sturm sequences.

4.2. Action matrix based solvers

Instead of using resultant theory to do elimination we can also base our solvers on action matrix methods. To this end we use an automatic generator [22] to construct a solver that estimates the action matrix from our new parametrized problem. In order to make the estimation fast (i.e. using small elimination templates) we use the basis selection methodology described in [25]. Then, instead of directly solving the Eigenvalue problem corresponding to the action matrix, we use numerical methods to estimate the characteristic polynomials from the action matrix. We use the method of Danilevsky [8] based on similar matrices, but there are also other possible approaches such as Krylov's method [13]. The real roots of the characteristic polynomial can then be found using Sturm Sequences.

4.3. Overview of solvers

We will now list the possible minimal relative pose problems that may arise, for different combinations of radial distortion and calibration on the two cameras. The different calibrations that we will consider are fully calibrated cameras, constant unknown focal length, one-sided unknown focal length and two unknown focal lengths. The different radial distortions that we will consider are: constant unknown distortion, one-sided unknown distortion and two unknown distortions. In total this will give 24 different

Table 2: Parametrizations for the different relative pose cases

Problems				Parameters
λE,	$E\lambda$			λ, e_{31}, e_{32}
$\lambda E\lambda$				λ, e_{31}, e_{32}
λfE,	$Ef\lambda$			λ, f_{32}
λEf,	$fE\lambda$			λ, f_{32}
λfEf,	$fEf\lambda$			λ, f_{32}
$\lambda fE\lambda$,	$\lambda Ef\lambda$			λ, f_{32}
$\lambda fEf\lambda$				λ, f_{32}
$\lambda' E\lambda$				λ, λ', e_{32}
$\lambda f'Ef$,	$f'Ef\lambda$,	λF,	$F\lambda$	λ
$\lambda f'Ef\lambda$,	$\lambda F\lambda$			λ
$\lambda' fE\lambda$,	$\lambda' Ef\lambda$			λ, λ'
$\lambda' fEf\lambda$				λ, λ'
$\lambda' f'Ef\lambda$,	$\lambda' F\lambda$			λ, λ'

cases, but due to the symmetry in cameras we only need to consider 16 problems. An overview of the different problems can be seen in Table 1. Since $\mathcal{P}_{f'Ef} = \mathcal{P}_F$ we further can reduce the number of unique problems to 13. We have previously described our approach applied to $E\lambda$ and λE (Example 1), $f'Ef\lambda$, $\lambda f'Ef$, $F\lambda$ and λF (Example 2) and $\lambda E\lambda$ (Example 3). An overview of the different parametrization for all cases is given in Table 2. From these parametrizations we construct the final solvers. We found that the approach described in Section 4.2 for all cases gave much faster solvers than the approach in Section 4.1. So, our proposed solvers, (as described in Table 1) are all based on the method in Section 4.2. For two cases we end up with an infinite number of solutions for the given parametrizations. In these two cases the spurious solutions lead to that we can not directly use the standard methods. We can handle this by applying the saturation technique proposed in [23]. This will in general lead to larger elimination templates, and for the previously solved case $\lambda E\lambda$ our parametrization does not lead to a solver that is faster than the current state-of-the-art.

5. Evaluation of solvers

We will in this section give results on a number of radial distortion solvers. We will show some numerical properties in the next subsection. Then, in the subsection after that, we will show how the solvers can be used for robust inlier estimation, based on real images. The investigated solvers were implemented in Matlab, with some coefficient manipulation, action matrix generation and the Sturm sequence solver implemented as Mex C++ code.The average execution time of the full solvers are based on a 2,5 GHz Intel Core i7 Macbook Pro, and given in Table 1.

5.1. Numerical accuracy

In order to test the numerical accuracy of our solvers we generate random problem instances, run our solvers, and evaluate the results. We generate random minimal point sets, and run the solvers.

We evaluate how well we can estimate the ground truth (random) radial distortion parameter. In Figure 2 histograms over the logarithm of the absolute relative error in λ is shown. One can see that we get a large spread in error. One reason for this is that for many instances the separation between focal length and radial distortion is ill-posed, and inherent to the problem. If our main goal is to use the solvers for inlier estimation this is not any major problem, but if we actually want to extract calibration statistics one would in many cases need further non-linear refinement or preprocessing of input data to the solvers. This is an area for further research.

5.2. RANSAC test on real images

In order to perform a somewhat controlled proof of concept experiment using real images, we did the following. We used ten images of the San Marco square in Venice, taken by a camera with known calibration, and negligible radial distortion. In addition to this we have one image (taken at a different time) with large radial distortion. We extracted features and descriptors using the learned affine regions of [29]. We then matched features using a similarity threshold of 20% and a ratio for ambiguous matches of 0.8. Note that finding good initial feature matches in the presence of radial distortion is in itself a difficult problem, and the end results will of course depend on the quality of the initial matches. The number of initial matches is given in Table 3. In order to have a baseline we used the known calibration ($f = 2396$) of the first images, and the exif tag of the radial distortion image as an estimate of the second focal length ($f' = 2032$). We also manually estimated a division model radial distortion coefficient for the second image ($\lambda = -0.37$). Using this information, we calibrated both image feature correspondences, and ran an inlier estimation using RANSAC (1 000 iterations, epipolar inlier bound 0.002 in normalized coordinates, used throughout the experiments). The resulting number of inliers is shown in Table 3. In many cases we do not have access to full calibration information, or it is unreliable. In Table 3 we show the result of running our solvers with varying degrees of known calibration. Numbers in bold indicate estimated values. These are estimated purely based on the results from the minimal solvers, and no additional refinement or bundle adjustment was performed. One can see that we get more or less the same number of inliers for the varying cases, indicating the stability of the solvers. Examples of inlier set is shown in Figure 3, when we have used the λEf-solver.

From the minimal solvers we directly get estimates of

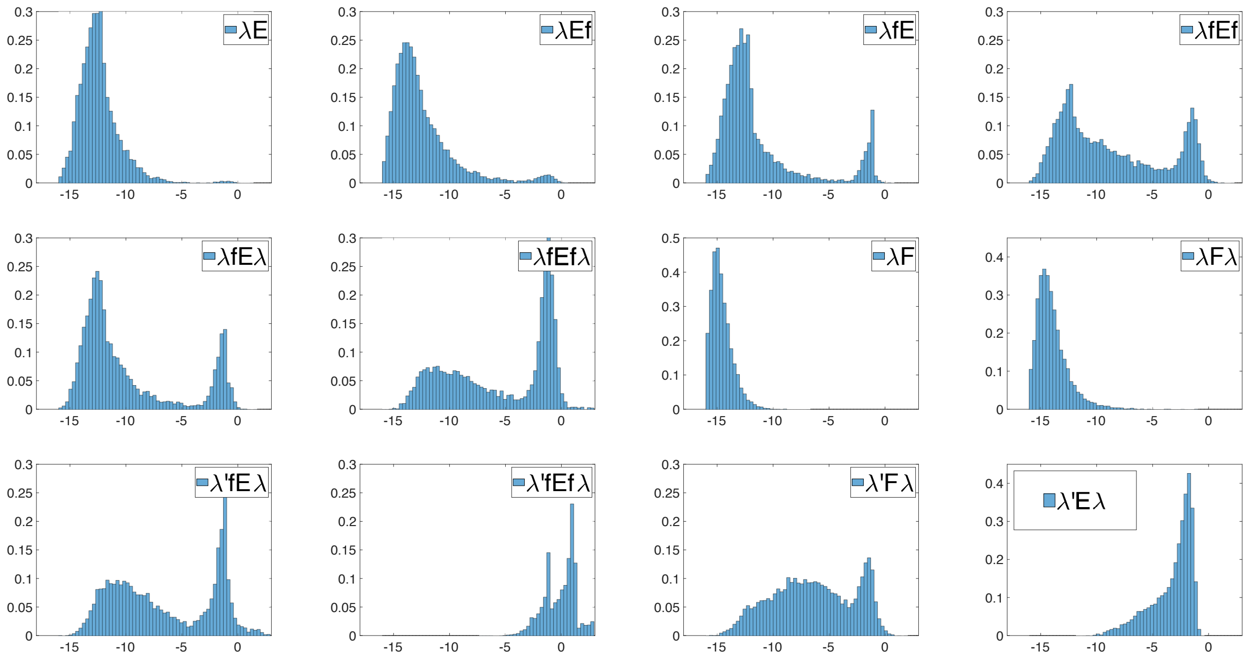

Figure 2: Histograms over residuals (x-axis shows $\log_{10}$ of absolute residuals between estimated and ground truth λ) for 10 000 synthetic random problem instances.

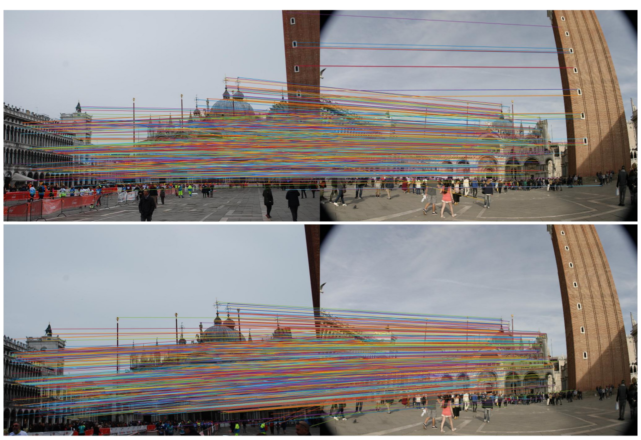

Figure 3: Example results on robust feature matching using the proposed minimal solvers in RANSAC. Figure shows inliers for the $\lambda E f$-case, for two example corresponding images.

the radial distortion parameter, and one can see that these estimates are quite stable and accurate for these cases. We can also extract the various unknown focal lengths from the estimated fundamental matrices, using the technique described in Appendix A. These are somewhat more unstable, and particularly the case of extracting two different focal lengths from a fundamental matrix is known to be an ill-posed problem.

6. Conclusion

We have in this paper presented a unifying framework for relative pose problems, where at least one camera suffers from radial distortion. Using this framework we can systematically describe the different problems that one can formulate. Some of these problems may seem exotic but most of them are actually realistic scenarios. Our formulation also gives us tools to eliminate variables, and formulate the problems as systems of polynomial equations, in at most three unknown parameters. This enables the ability to use efficient Sturm sequence solvers. We have enumerated all minimal relative pose problems based on partial unknown focal lengths and radial distortion. We have also shown how solvers to such problems can be used as components for robust inlier estimation.

A. Extracting unknown focal lengths from the fundamental matrix

In this section we will describe how we can find the two unknown focal lengths from a given fundamental matrix F. This means that we would like to find a valid essential matrix E and two focal lengths f and f' such that

$$\begin{bmatrix} f' & 0 & 0 \\ 0 & f' & 0 \\ 0 & 0 & 1 \end{bmatrix} F \begin{bmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{bmatrix} = E. \tag{18}$$

Table 3: RANSAC experiment on real data. The Table shows inlier estimates for various setups, using the same initial matches in each case, but with varying knowledge of the calibration. Bold entries indicate estimated entities based purely on the tested minimal solvers.

	Images										Median
#match	1265	1526	1592	1386	1268	1241	1388	1499	1466	1526	1427
E											
#inl.	**927**	**1144**	**1186**	**1034**	**961**	**914**	**1014**	**1080**	**1087**	**1159**	**1057**
f	2396	2396	2396	2396	2396	2396	2396	2396	2396	2396	2396
f'	2032	2032	2032	2032	2032	2032	2032	2032	2032	2032	2032
λ	-0.37	-0.37	-0.37	-0.37	-0.37	-0.37	-0.37	-0.37	-0.37	-0.37	-0.37
λE											
#inl.	**920**	**1135**	**1171**	**1021**	**941**	**910**	**988**	**1062**	**1075**	**1135**	**1042**
f	2396	2396	2396	2396	2396	2396	2396	2396	2396	2396	2396
f'	2032	2032	2032	2032	2032	2032	2032	2032	2032	2032	2032
λ	**-0.46**	**-0.39**	**-0.35**	**-0.34**	**-0.32**	**-0.34**	**-0.56**	**-0.32**	**-0.43**	**-0.35**	**-0.35**
λfE											
#inl.	**920**	**1138**	**1178**	**1015**	**965**	**926**	**1007**	**1063**	**1076**	**1147**	**1039**
f	2396	2396	2396	2396	2396	2396	2396	2396	2396	2396	2396
f'	**1695**	**2172**	**1930**	**2257**	**3566**	**6313**	**1937**	**2821**	**2337**	**1774**	**2214**
λ	**-0.41**	**-0.40**	**-0.35**	**-0.26**	**-0.38**	**-0.37**	**-0.35**	**-0.35**	**-0.33**	**-0.30**	**-0.35**
λEf											
#inl.	**852**	**1103**	**1140**	**992**	**891**	**866**	**940**	**1027**	**1051**	**1114**	**1010**
f	**2508**	**2173**	**2009**	**2259**	**2613**	**2220**	**1560**	**2588**	**2447**	**1918**	**2240**
f'	2032	2032	2032	2032	2032	2032	2032	2032	2032	2032	2032
λ	**-0.32**	**-0.33**	**-0.32**	**-0.32**	**-0.30**	**-0.31**	**-0.26**	**-0.30**	**-0.31**	**-0.35**	**-0.31**
$\lambda f'Ef$											
#inl.	**869**	**1117**	**1149**	**995**	**947**	**884**	**978**	**1029**	**1062**	**1121**	**1013**
f	**1116**	**1251**	**1376**	**1031**	**2992**	**1223**	**3344**	**1362**	**1327**	**899**	**1289**
f'	**4089**	**3541**	**3887**	**2777**	**3874**	**3900**	**30000**	**4416**	**4745**	**2047**	**3894**
λ	**-0.33**	**-0.29**	**-0.36**	**-0.32**	**-0.31**	**-0.35**	**-0.25**	**-0.28**	**-0.30**	**-0.38**	**-0.31**

There exists a method by Hartley [10] to solve this problem. For each F we get two solutions, but the solutions differ only in sign, so since we know that the focal lengths should be positive, we essentially have a unique solution to the problem. The method due to Hartley involves a number of quite involved manipulations of the epipolar geometry in order to extract the focal lengths. We will here describe a simpler method based on using the trace constraints (8) directly on the left hand side of (18). These constraints can be simplified by parametrizing in $a = f^2$ and $b = f'^2$. Using this parametrization, the constraints only contain the four monomials $(ab, a, b, 1)$. If we express $\mathcal{P}_E(a, b)$ as

$$M \begin{bmatrix} ab \\ a \\ b \\ 1 \end{bmatrix} = 0, \tag{19}$$

then we find the solution by taking the null-space of M, and normalizing with the last entry. We then extract a and b and find $f = \sqrt{a}$ and $f' = \sqrt{b}$. Note that M will always be rank deficient due to the construction, but if our F doesn't follow the assumption that it can be written as (18), the solution for a and b might be negative, and we can't extract real focal lengths. This is a general issue for the formulation, and without additional regularization constraints any algebraic solution will inherently have this problem. We can use the same method for extracting one-sided focal lengths or constant focal length by setting $b = 1$ respectively $a = b$.

Acknowledgement This work was supported by the strategic research project eSSENCE and the ECSEL Joint Undertaking (JU) under grant agreement No 876019. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Germany, Netherlands, Austria, Romania, France, Sweden, Cyprus, Greece, Lithuania, Portugal, Italy, Finland, Turkey.

References

[1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. 1

[2] Erwin H Bareiss. Sylvester's identity and multistep integer-preserving gaussian elimination. *Mathematics of computation*, 22(103):565–578, 1968. 5

[3] João Pedro Barreto and Kostas Daniilidis. Fundamental matrix for cameras with radial distortion. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 1, pages 625–632. IEEE, 2005. 2

[4] Martin Byrod, Zuzana Kukelova, Klas Josephson, Tomas Pajdla, and Kalle Astrom. Fast and robust numerical solutions to minimal problems for cameras with radial distortion. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 2, 3

[5] David A Cox, John Little, and Donal O'shea. *Using algebraic geometry*, volume 185. Springer Science & Business Media, 2006. 5

[6] George Cybenko. The numerical stability of the levinson-durbin algorithm for toeplitz systems of equations. *SIAM Journal on Scientific and Statistical Computing*, 1(3):303–319, 1980. 5

[7] M Demazure. *Sur deux problemes de reconstruction*. PhD thesis, INRIA, 1988. 2

[8] VN Faddeev and VN Faddeeva. *Computational Methods of Linear Algebra,(trans)*. Dover Publications Inc., New York, 1959. 5

[9] Andrew W Fitzgibbon. Simultaneous linear estimation of multiple view geometry and lens distortion. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 1, pages I–I. IEEE, 2001. 1, 2

[10] Richard Hartley. Extraction of focal lengths from the fundamental matrix. *Unpublished manuscript*, 1993. 8

[11] R. Hartley and H. Li. An efficient hidden variable approach to minimal-case camera motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(12):2303–2314, Dec 2012. 5

[12] Anders Heyden and Kalle Åkström. Minimal conditions on intrinsic parameters for euclidean reconstruction. In *Asian Conference on Computer Vision*, pages 169–176. Springer, 1998. 2

[13] Alston S Householder. *The Theory of Matrices in Numerical Analysis*. Courier Corporation, 2006. 5

[14] Fangyuan Jiang, Yubin Kuang, Jan Erik Solem, and Kalle Åström. A minimal solution to relative pose with unknown focal length and radial distortion. In *Asian Conference on Computer Vision*, pages 443–456. Springer, 2014. 2, 3

[15] Yubin Kuang, Jan E Solem, Fredrik Kahl, and Kalle Astrom. Minimal solvers for relative pose with a single unknown radial distortion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 33–40, 2014. 2, 3, 4

[16] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Automatic generator of minimal problem solvers. In *Computer Vision–ECCV 2008*, pages 302–315. Springer, 2008. 2, 5

[17] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Real-time solution to the absolute pose problem with unknown radial distortion and focal length. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2816–2823, 2013. 2

[18] Zuzana Kukelova, Joe Kileel, Bernd Sturmfels, and Tomas Pajdla. A clever elimination strategy for efficient minimal solvers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4912–4921, 2017. 2

[19] Zuzana Kukelova and Tomas Pajdla. A minimal solution to the autocalibration of radial distortion. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–7. IEEE, 2007. 2

[20] Zuzana Kukelova and Tomas Pajdla. Two minimal problems for cameras with radial distortion. In *7th Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras (OMNIVIS'07)*, 2007. 2

[21] Zuzana Kukelova and Tomas Pajdla. A minimal solution to radial distortion autocalibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2410–2422, 2011. 2

[22] Viktor Larsson, Kalle Åström, and Magnus Oskarsson. Efficient solvers for minimal problems by syzygy-based reduction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 4, 5

[23] Viktor Larsson, Kalle Astrom, and Magnus Oskarsson. Polynomial solvers for saturated ideals. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2288–2297, 2017. 6

[24] Viktor Larsson, Zuzana Kukelova, and Yinqiang Zheng. Making minimal solvers for absolute pose estimation compact and robust. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2316–2324, 2017. 2

[25] Viktor Larsson, Magnus Oskarsson, Kalle Astrom, Alge Wallis, Zuzana Kukelova, and Tomas Pajdla. Beyond grobner bases: Basis selection for minimal solvers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3945–3954, 2018. 2, 3, 4, 5

[26] Hongdong Li and Richard Hartley. A non-iterative method for correcting lens distortion from nine point correspondences. *OMNIVIS 2005*, 2:7, 2005. 2

[27] Stephen Maybank. *Theory of reconstruction from image motion*, volume 28. Springer Science & Business Media, 2012. 2

[28] Branislav Micusik and Tomas Pajdla. Estimation of omnidirectional camera model from epipolar geometry. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 1, pages I–I. IEEE, 2003. 2

[29] Dmytro Mishkin, Filip Radenovic, and Jiri Matas. Repeatability is not enough: Learning affine regions via discriminability. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 284–300, 2018. 6

[30] Gaku Nakano. A versatile approach for solving pnp, pnpf, and pnpfr problems. In *European Conference on Computer Vision*, pages 338–352. Springer, 2016. 2

[31] Oleg Naroditsky and Kostas Daniilidis. Optimizing polynomial solvers for minimal geometry problems. In *2011 International Conference on Computer Vision*, pages 975–982. IEEE, 2011. 2

[32] C. Olsson and O. Enqvist. Stable structure from motion for unordered image collections. In *Scandinavian Conf. on Image Analysis*, 2011. 1

[33] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. 1

[34] P Stefanovic. Relative orientation–a new approach. *ITC Journal*, 3:417–448, 1973. 2