One-sided Radial Fundamental Matrix Estimation

José Henrique Brito[12]
josehbrito@gmail.com

Christopher Zach[4]
chzach@microsoft.com

Kevin Köser[3]
kevin.koeser@inf.ethz.ch

Manuel João Ferreira[2]
mjf@dei.uminho.pt

Marc Pollefeys[3]
marc.pollefeys@inf.ethz.ch

[1] Instituto Politécnico do Cávado e do Ave
Barcelos, Portugal

[2] Universidade do Minho
Guimarães, Portugal

[3] Computer Vision and Geometry Group
ETH Zurich, Switzerland

[4] Machine Learning and Perception
Microsoft Research Cambridge, UK

Abstract

For modern consumer cameras often approximate calibration data is available, making applications such as 3D reconstruction or photo registration easier as compared to the pure uncalibrated setting. In this paper we address the setting with calibrated-uncalibrated image pairs: for one image intrinsic parameters are assumed to be known, whereas the second view has unknown distortion and calibration parameters. This situation arises e.g. when one would like to register archive imagery to recently taken photos. A commonly adopted strategy for determining epipolar geometry is based on feature matching and minimal solvers inside a RANSAC framework. However, only very few existing solutions apply to the calibrated-uncalibrated setting. We propose a simple and numerically stable two-step scheme to first estimate radial distortion parameters and subsequently the focal length using novel solvers. We demonstrate the performance on synthetic and real datasets.

1 Introduction

Since fully automatic registration of photos has been shown in panoramic scenarios [4] and later for photo collections [21], the basic techniques of first obtaining a number of tentative correspondences and then pairwise geometric verification have become very successful for a number of applications. The verification step removes mismatches by estimating parameters of a model (e.g. homography, epipolar geometry) that is assumed to be fulfilled by all valid correspondences and so inconsistent matches can be rejected. The choice of the right model is the crucial step in these applications, since an overly general model might accept wrong matches, and an overly narrow model may classify good correspondences as outliers. While the above mentioned approaches [4, 21] for automatic registration are very successful with almost ideal pinhole cameras without much distortion, they cannot cope with

Figure 1: Illustrative result of applying our method compared to the results of the standard 8-point algorithm; red are the inliers found by both methods; green are the extra inliers found by our method; blue are inliers found by the standard 8-point not found by our method.

omni-directional images and wide field-of-view images with significant lens distortion. In the literature, several camera models and techniques have been proposed to model such distortion [3, 7, 10, 13, 19, 24]. However, for automatic registration of images obtained from internet sources or archives, an off-line camera calibration phase is not feasible. Hence, lens distortion has to be considered directly in the multi-view geometry estimation stage, which is typically based on some variant of RANSAC [9]. The aim is to generate a good model hypothesis explaining the putative correspondences by stochastically searching for an outlier-free sample set.

The probability to draw a set of n correct correspondences from a set of all potential matches decreases exponentially with n, and authors have tried to construct so called minimal solvers in order to keep the computational complexity low, e.g. [6, 10, 15, 16]. Due to the algebraic nature of the underlying geometric problem, these solvers are typically based on systems of polynomial equations. Variable elimination techniques such as Gröbner bases [8] often lead to high-degree univariate polynomials, and solutions are not always stable in the presence of noise. Sometimes derived solutions are even only solvable using exact arithmetics on integers [17] but not when using finite-precision floating point numbers. We consider the case of matching an image with unknown focal length and radial distortion, to an image with known intrinsics. This setting has received little attention so far (see [5] for a solution not considering the lens distortion), but is relevant e.g. when matching archive imagery to recent photos captured by modern digital cameras or in sequential reconstruction. Our approach is a numerically stable two-step method by first estimating a one-sided radial fundamental matrix (leading to a cubic polynomial formulation) and subsequent extraction of the focal length (via a quadratic polynomial) in order to obtain a metric reconstruction and a good starting point for bundle adjustment. Our solution is not purely minimal and relaxes the original problem to achieve a much more practical method. We briefly considered a minimal method to estimate e.g. a one-sided radial essential matrix (known camera intrinsics but unknown one-parameter distortion model for one image), but such an approach leads to a 26 degree polynomial (which is consistent with the up to 52 solutions for the two-sided radial essential matrix reported in [17]). Our proposed method only requires solving a cubic polynomial and therefore performs similarly, in terms of speed and numerical robustness, to the standard 7-point algorithm for fundamental matrix estimation.

2 Previous Work

Our work is related to the general multi-view geometry and self-calibration literature, for which Hartley and Zisserman [14] is the standard reference. However, we also consider radial distortion, which is usually only addressed by an off-line calibration step. Based on the classical radial distortion model [8], several authors investigated in the combination of lens distortion and epipolar geometry, e.g. [1, 10, 13, 19, 24]. Among these, Fitzgibbon [10] has proposed a formulation that is quite easy to use in multi-view relations, upon which Barreto *et al.* [1] build to propose an extended fundamental matrix. Fitzgibbon [10] also proposes a (non-minimal) solution for estimating the fundamental matrix plus a common radial distortion, a setting for which Kukelova *et al.* [16] subsequently derived the minimal solution. While we are largely inspired by Fitzgibbon's formulation, this solution does not apply to the calibrated-uncalibrated setting (with different radial distortions). Later, Byrod *et al.* [6] and Kukelova *et al.* [18] derive a solution to estimate the fundamental matrix plus two different, unknown distortions using 9 point matches. In principle, in our setting, one could neglect the known distortion for one image and apply the minimal solver for the radial fundamental matrix directly, but this comes at the expense of obtaining up to 24 solutions, whereas our proposed method returns only up to three possible values. To the best of our knownledge, the "one-sided" calibrated-uncalibrated setting has only been addressed in [5] (without considering radial lens distortion). If lens distortion is known (or neglected), it is possible to estimate a fundamental matrix and then to extract the focal length, as proposed by Urbanek and Sturm [25]. We also propose a novel way to extract the focal length from the obtained fundamental matrix, enabling the integration of the full method into large-scale systems such as [22] or [11], which is our ultimate goal, namely to register historic and present image collections. The other solution proposed in [5] for the semicalibrated case uses the Groebner basis method to solve a system of polynomial equations, also formed with the image matches.

3 The Proposed Two-Step Approach

In this section we derive our method to estimate the epipolar geometry between images taken by a fully calibrated camera and an uncalibrated one exhibiting radial lens distortion with a known distortion center. Since a direct formulation using the trace constraint on essential matrices turns out to be impractical (due to the high degree of the resulting polynomial), we propose a substantially simpler two step approach. The first step estimates the fundamental matrix and the parameter of the one unknown lens distortion, and the second step extracts the focal length useful to initialize a metric reconstruction.

3.1 One-sided Radial-Fundamental Matrix

In the following, we discuss the one-sided radial fundamental matrix describing the epipolar geometry between an undistorted and a radially distorted image, and derive a corresponding estimation method from given feature matches. We start with the epipolar relation,

$$q^T F p_u = 0 \qquad (1)$$

where $q = (q_x, q_y, 1)^T$ is an image point from a distortion-corrected image, and p_u is the undistorted version of an observed image point $p_d = (x_d, y_d, 1)^T$ from an image with un-

known radial distortion. By using the distortion model proposed in [10], we have

$$p_u \propto \begin{pmatrix} x_d \\ y_d \\ 1 + \lambda r_d^2, \end{pmatrix} \qquad (2)$$

where $r_d^2 = (x_d - u)^2 + (y_d - v)^2$ for a known distortion center $(u,v)^T$, which we assume to coincide with the image center, and λ is an unknown distortion parameter. The epipolar constraint can be written as

$$q^T F p_u = q^T F \begin{pmatrix} x_d \\ y_d \\ 1 + \lambda r_d^2 \end{pmatrix} = q^T \left(F \begin{pmatrix} x_d \\ y_d \\ 1 \end{pmatrix} + \lambda F \begin{pmatrix} 0 \\ 0 \\ r_d^2 \end{pmatrix} \right) = q^T \underbrace{\left[F \mid \lambda F_3 \right]}_{=: \hat{F}} \begin{pmatrix} x_d \\ y_d \\ 1 \\ r_d^2 \end{pmatrix}, \qquad (3)$$

where we introduced the 3×4-matrix $\hat{F}$. F_3 denotes the 3rd column of F. The task is now to determine the matrix $\hat{F}$ from a number of correspondences. $\hat{F}$ can be estimated linearly from 11 correspondences as the respective null vector of the stacked epipolar constraints. This solution ignores intrinsic constraints of $\hat{F}$, which are $\det(\hat{F}[1:3,1:3]) = \det(F) = 0$ (the fundamental matrix is rank deficient) and $\hat{F}_4 \propto \hat{F}_3$ (the last column is a multiple of the third one). Consequently, only 8 correspondences are necessary for a minimal solution (11 unknown minus 3 constraints). Given 8 (non-degenerate) correspondences, $\hat{F}$ has to reside in a 4-dimensional nullspace and can be written as

$$\hat{F} = x\hat{X} + y\hat{Y} + z\hat{Z} + w\hat{W}, \qquad (4)$$

where $\hat{X}, \ldots, \hat{W}$ corresponds to the basis of the nullspace. The freedom of scale can be fixed by setting $w = 1$, leaving 3 unknowns to be determined, x, y and z. Expanding the determinant constraint and the $\hat{F}_4 \propto \hat{F}_3$ constraints leads to a polynomial system of equations in x, y and z. Using Macaulay 2 one can verify that up to 8 real solutions for $\hat{F}$ are obtained. By using a ninth correspondence, $\hat{F}$ can be estimated in a much simpler way: we drop the determinant constraint (as also done in the 8-point algorithm [12] to determine the standard fundamental matrix). By using 9 correspondences, the nullspace $\hat{F}$ is an element of is only three-dimensional, i.e.

$$\hat{F} = x\hat{X} + y\hat{Y} + z\hat{Z}. \qquad (5)$$

We can again fix z to 1 due to the scale ambiguity of $\hat{F}$. The constraints $\hat{F}_4 \propto \hat{F}_3$, i.e. $\lambda \hat{F}_4 = \hat{F}_3$, now read as

$$x\hat{X}_{i4} + y\hat{Y}_{i4} + \hat{Z}_{i4} = \lambda \left(x\hat{X}_{i3} + y\hat{Y}_{i3} + \hat{Z}_{i3} \right) \qquad (6)$$

for $i = 1, 2, 3$. This polynomial system of equations is simple enough to perform symbolic elimination using the resultant method. Elimination of x and y (e.g. by using Maxima's `eliminate` function) yields a quartic polynomial in λ, which is unfortunately not minimal in its degree. Hence, we slightly alter the elimination method: first, we can eliminate λ by taking ratios (which corresponds to elimination of λ using the resultant method), leading to 3 polynomial equations in x and y only,

$$p_{ij}(x,y) \stackrel{\text{def}}{=} \left(x\hat{X}_{i4} + y\hat{Y}_{i4} + \hat{Z}_{i4} \right) \left(x\hat{X}_{j3} + y\hat{Y}_{j3} + \hat{Z}_{j3} \right)$$
$$- \left(x\hat{X}_{j4} + y\hat{Y}_{j4} + \hat{Z}_{j4} \right) \left(x\hat{X}_{i3} + y\hat{Y}_{i3} + \hat{Z}_{i3} \right) \stackrel{!}{=} 0 \qquad (7)$$

for $(i, j) \in \{(1,2),(1,3),(2,3)\}$. Application of the resultant method again to eliminate e.g. y returns a 4th order polynomial in x, which is not optimal. Hence we compute two resultants (e.g. combining p_{12} with p_{13}, and p_{12} with p_{23}, respectively) leading to two degree 4 polynomials in x,

$$q_1(x) \overset{\text{def}}{=} a_1 x^4 + b_1 x^3 + c_1 x^2 + d_1 x + e_1 \overset{!}{=} 0$$
$$q_2(x) \overset{\text{def}}{=} a_2 x^4 + b_2 x^3 + c_2 x^2 + d_2 x + e_2 \overset{!}{=} 0.$$

The coefficients[1] are large, but still manageable expressions of the nullspace matrices $\hat{X}$, $\hat{Y}$ and $\hat{Z}$. The leading monomial x^4 can now be eliminated by one step of Gaussian elimination leading to a final cubic polynomial,

$$r(x) \overset{\text{def}}{=} a_2 q_1(x) - a_1 q_2(x) \overset{!}{=} 0. \tag{8}$$

This can be solved in closed form leading to one or three real solutions. For each possible value of x, a corresponding y can be extracted by a similar procedure. Two of the p_{ij} polynomials (which are quadratic) yield a linear equation in y after one Gaussian elimination step. The extended fundamental matrix is given by $\hat{F} = x\hat{X} + y\hat{Y} + \hat{Z}$, and λ can be obtained as the ratio $\lambda = \hat{F}_{14}/\hat{F}_{13} = \hat{F}_{24}/\hat{F}_{23} = \hat{F}_{34}/\hat{F}_{33}$. By construction all those ratios are equal.

Since we dropped the rank constraint, the estimated fundamental matrix (i.e. the 3×3 submatrix $\hat{F}[1:3, 1:3]$) will generally be of full rank. We enforce rank-2 using the SVD as in the 8-point algorithm.

Normalization: Similarly to the 8-point algorithm [12] we normalize the image measurements. For the calibrated image we use the inverse of the camera intrinsics for the normalization, and for the uncalibrated one we use an initial estimate of the focal length, $f_{\text{guess}} = \frac{W/2}{\tan(fov_{\text{guess}}/2)}$ where W is the image width and $fov_{\text{guess}} = 50°$ is an a-priori estimate of the focal length. Further, the image points are centered to have the principal point in the origin.

3.2 Extracting the Focal Length from the Fundamental Matrix in a Partially Calibrated Setup

The method detailed in the previous section allows to compensate for the lens distortion in an uncalibrated image (at least to a large extent). Nevertheless, for metric reconstructions and bundle adjustment the knowledge of camera intrinsics, i.e. mostly the focal length, is desirable. This section describes how the focal length can be extracted in partially calibrated settings and proposes a different approach than [23, 25].

Let F be a fundamental matrix, and K and K' camera intrinsics such that $E = (K')^T F K$ is an essential matrix. K is assumed to be known and K' is of the shape $\text{diag}(f, f, 1)$ for an unknown focal length f, hence we can incorporate K into F, yielding $E = \text{diag}(f, f, 1)F$. Plugging this expression into the trace constraint for essential matrices, $2EE^T E - \text{tr}(EE^T)E = 0$ (see e.g. [20]), leads to a corresponding matrix constraint in terms of f,

$$G(f) \overset{\text{def}}{=} 2\,\text{diag}(f, f, 1)FF^T \text{diag}(f^2, f^2, 1)F$$
$$- \text{tr}\left(\text{diag}(f, f, 1)FF^T \text{diag}(f, f, 1)\right) \text{diag}(f, f, 1)F \overset{!}{=} 0. \tag{9}$$

---

[1]Matlab code is available at http://www.cvg.ethz.ch/research/distortion-in-multiple-view-geometry/
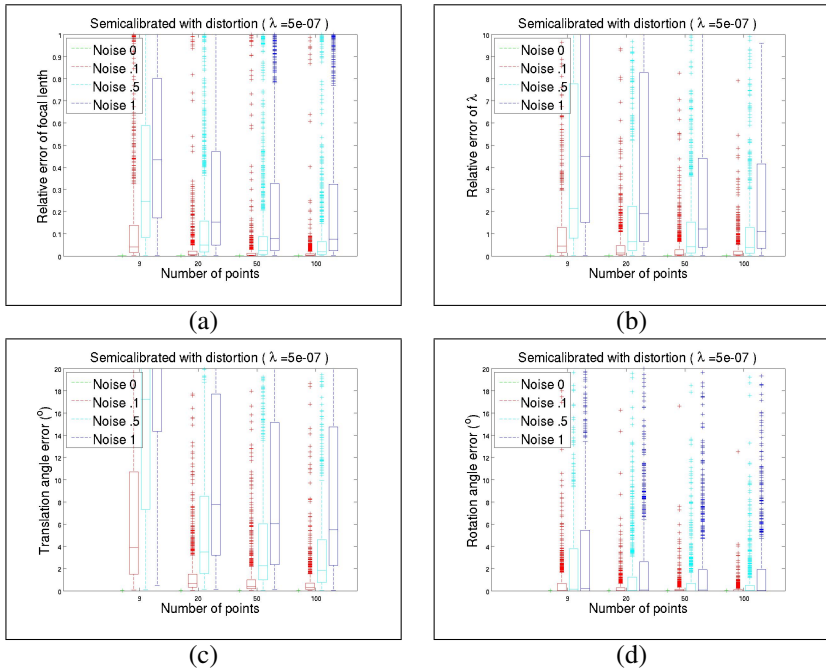
Figure 2: Error measurements for random camera poses with distortion parameter λ of $5 \cdot 10^{-7}$ and different noise levels: (a) relative error of focal length; (b) relative error of lambda; (c) angular error for the translation vector between the cameras; (d) error for the rotation angle between the cameras

We determine f by minimizing the algebraic error, $\|G(f)\|_F^2$. First order optimality conditions, $d\|G(f)\|_F^2/df = 0$, yields a polynomial in f^5, f^3 and f. Since we can exclude the degenerate solution $f = 0$, a double quadratic polynomial in f^4 and f^2 can be obtained, which is trivial to solve after substituting $w = f^2$. Since f has to be strictly positive, up to two possible values for f need to be checked for optimality.

4 Experiments

In this section we evaluate the numerical stability and behaviour of the proposed algorithms. To test the algorithms, we first analyse their behaviour with synthetic data and then with real images. In the experiments with synthetic data we use random camera configurations, different noise levels, different distortion parameters. Due to the lack of earlier methods for our setting, the results are compared to those obtained with the method from Kukelova *et al*. [13], although this latter method estimates distortion for both cameras whereas, in our setting, one of the images in each image pair has known intrinsics and no distortion. We then test the algorithms with real world images, for which we have ground truth information for focal length and distortion. Results are also compared with those obtained with the standard method of computing the fundamental matrix with the 8-point algorithm (disregarding distortion).

Figure 3: Error comparison with the method from Kukelova *et al.* [18] for random camera poses with distortion parameter λ of $5 \cdot 10^{-7}$ and different noise levels: (a) relative error of focal length; (b) relative error of lambda; (c) angular error for the translation vector between the cameras; (d) error for the rotation angle between the cameras

4.1 Evaluation with synthetic data

The first set of tests for the semicalibrated case was performed using synthetic data. All tests with synthetic data were performed with a set of random 3D points and 1000 generated random camera poses. The first camera was placed at the origin, with fixed parameters, pointed towards the set of 3D points. The 1000 random poses were generated for the second camera, by generating random translations, random rotations and random focal lengths, varying between 1/2 and 2x the focal length of the first camera. For each camera pose we projected the 3D points on both cameras, distorted the points on the second camera according to the division distortion model[10] using values for λ ranging from strong distortion ($5 \cdot 10^{-7}$ on 1024x1024 image) to almost no distortion ($5 \cdot 10^{-9}$ and also 0), and added noise with different values of standard deviation σ. For each of the poses we computed the error in the estimation of λ, focal length of the second camera, the angular error for the translation vector and the error for the rotation between the two cameras. Same as with the 8 point algorithm, we can also use more than 9 correspondences in our formulation to obtain the null space in eq.5. and so we vary also the number of point correspondences used. All computations of the fundamental matrix were performed using point correspondences with normalised coordinates according to section 3.1. Results are presented in Fig. 2, where we can see that as the number of used correspondences rises the error values decrease. We also found that changing the value of the ground truth lambda does not have a perceivable impact on the estimate errors for the focal length, translation error and rotation, while for the estimation of lambda, the smaller ground truth lambda gets, the harder it is to estimate it, and so the estimate error
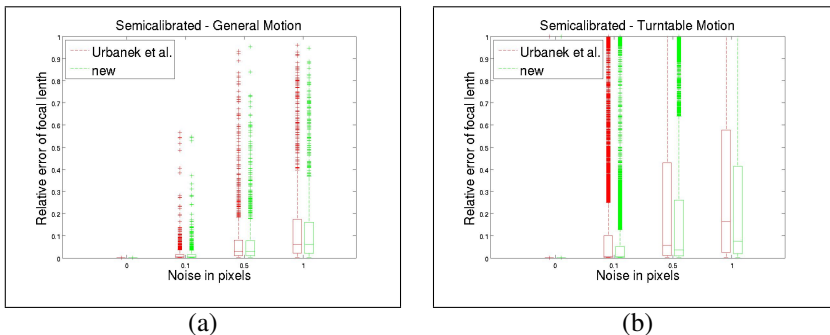
Figure 4: Error of estimates of the focal length from the fundamental matrix with our method and the method from Urbanek *et al.*: (a) random configurations; (b) turntable motion

grows. We don't include those results for different values of λ because of lack of space.

We then compared the performance of our method to that from Kukelova *et al.* [18], so for each pose we computed the errors using a minimal set of 9 points, while on one hand using our method, and on the other hand computing the fundamental matrix with the method from [18]. The focal length is extracted in both cases using the algorithm in section 3.2. Results are presented in Fig. 3 for ground truth λ equal to $5 \cdot 10^{-7}$. We can see that if we use a set of 9 points as one would use in RANSAC, our method is able to better handle increasing noise levels. It should also be noted that applying the method from Kukelova *et al.* to our setting should lead to estimates of λ_1 for the first image close to zero. However, this is not the case. In many tests, especially with noise, the estimate of λ_1 is very different from zero, which could indicate that the method is not stable for distortion values close to zero (i.e. when no distortion is present). Of course, in a practical situation, one would sample subsets of 9 correspondences in a RANSAC framework, and also perform optimization on the solution computed with the inliers. However, in these experiments we were trying to measure the sensitivity to noise of both algorithms for a random minimal set of correspondences.

To validate the second part of our method, which is the extraction of the focal length from the fundamental matrix, we compared the performance of our method to that of the method proposed by Urbanek at el. in [25]. To do this we conducted a separate test in which we generated 1000 random configurations, projected a set of 3D points in the images and added different levels of noise, but without adding distortion. Then we estimated the focal length from the fundamental matrix (computed with the standard 8-point algorithm) with both methods and compared them to the ground truth focal length. In randomly generated configurations the results do not significantly differ, but in specific settings, like turntable motion, our method produces more accurate estimates of the focal length. For this setting results are presented in Fig. 4.

4.2 Tests on real images

To test the method on real images we first matched a set of uncalibrated/distorted images to an image with known calibration parameters using different datasets. In Fig. 5 (a) and (b) we show some images of two of the datasets we used. From the evaluation of the previous section we conclude that our method produces more accurate results for a minimal solution from nine points and therefore we consider it to be more suitable to be used in a RANSAC framework. Our algorithm has a complexity comparable to that of the standard 7-point
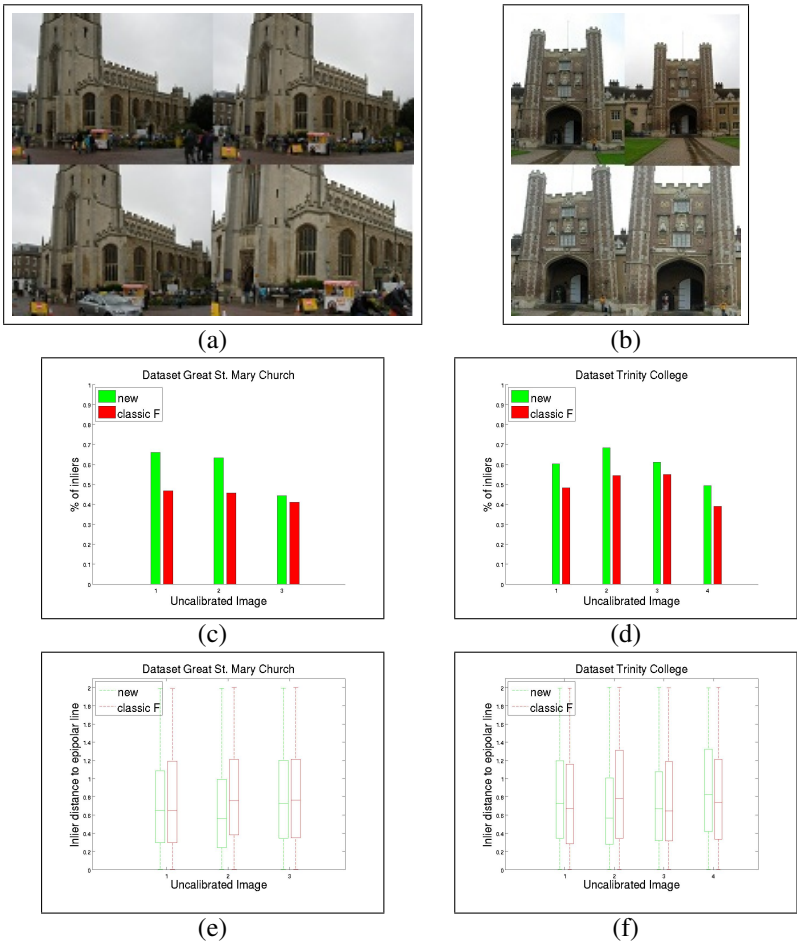
Figure 5: Some images from two of the datasets of real images used in the tests: (a) Great St. Mary Church; (b) Trinity College. (c) and (d) ratio of number of inliers vs. number of putative matches; (e) and (f) epipolar error of inliers

algorithm for classical fundamental matrix, since the main steps are finding the null space created by the nine correspondences (rather than the seven correspondences in the 7-point algorithm) and solving a quadratic polynomial. We will now compare whether the solution presented is effective in a RANSAC scheme and really helps as compared to not considering lens distortion. For this, we evaluate whether an overall practical system using our 9 point solver in a RANSAC scheme improves the number of retrieved matches and how the average error behaves as compared to a standard 8-point-method for classical fundamental.

To extract features in the images we used SURF[2], and then we computed a number of putative matches in each image pair by standard feature space matching. This produced a number of matches for each image pair, not all of which were correct correspondences. We then ran the two methods in a RANSAC framework with same parameters and sample sets. In the end we computed the ratio of inliers and the average epipolar error of the inliers. To obtain the epipolar error (used for classifying outliers) we computed the distance in pixels between a point and the epipolar line in the undistorted image. Results for two of the tested

datasets are shown in Fig. 5 (c-f) where we can see that our method tends to use a higher number of inliers and that these inliers have an equal or lower average epipolar error. In Fig. 1 we can see two typical situations were the standard 8-point method would use only the correspondences where the points are at the center of the image and therefore not heavily affected by radial distortion, whereas our method would be able to use more correspondences also at the edge of the image, where radial distortion is more severe.

5 Conclusion

We have presented a numerically stable and efficient solution to the calibrated-uncalibrated image registration problem with unknown focal length and radial distortion. The solution significantly improves the number of inliers in presence of distortion. Compared to more general solvers that consider radial distortion but that cannot make strong assumptions about the first camera, we could reduce the degree of the problem from 24 to 3, which means that we can obtain the solution in closed form without having to worry about numerical problems. At the same time complexity is reduced and we do not have to evaluate 24 different hypothesis generated from the 9 correspondences. Finally, we extract the focal length and the relative motion from the fundamental matrix, allowing for a metric reconstruction and integration of the uncalibrated, distorted image e.g. into previously calibrated, photo collections.

References

[1] João P. Barreto and Kostas Daniilidis. Fundamental matrix for cameras with radial distortion. In *International Conference on Computer Vision (ICCV)*, pages 625–632, 2005.

[2] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. *Computer Vision and Image Understanding (CVIU)*, 110:346–359, June 2008.

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

[4] M. Brown and D. Lowe. Recognising panoramas. In *International Conference on Computer Vision (ICCV)*, volume 2, pages 1218–1225, Nice, October 2003.

[5] Martin Bujnak, Zuzana Kukelova, and Tomas Pajdla. 3d reconstruction from image collections with a single known focal length. In *International Conference on Computer Vision (ICCV)*, pages 1803–1810, 2009.

[6] Martin Byröd, Zuzana Kukelova, Klas Josephson, Tomas Pajdla, and Kalle Aström. Fast and robust numerical solutions to minimal problems for cameras with radial distortion. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1 –8, june 2008.

[7] D. Claus and A. W. Fitzgibbon. A rational function lens distortion model for general cameras. In *Computer Vision and Pattern Recognition (CVPR)*, pages 213–219, 2005.

[8] D. Cox, J. Little, and D. O'Shea. *Ideals, varieties, and algorithms*. Springer, 1997.

[9] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, June 1981.

[10] A.W. Fitzgibbon. Simultaneous linear estimation of multiple view geometry and lens distortion. In *Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 125–132, 2001.

[11] J Frahm, P. Fite-Georgel, D. Gallup, T. Johnson, R. Raguram, C. Wu, Y. Jen, E. Dunn, B. Clipp, S. Lazebnik, and M. Pollefeys. Building rome on a cloudless day. In *European Conference on Computer Vision (ECCV)*, pages 368–381, 2010.

[12] R. Hartley. In defense of the eight-point algorithm. *Pattern Analysis and Machine Intelligence (PAMI)*, 19(6):580–593, 1997.

[13] R. Hartley and Sing Bing Kang. Parameter-free radial distortion correction with center of distortion estimation. *Pattern Analysis and Machine Intelligence (PAMI)*, 29(8): 1309 –1321, August 2007.

[14] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.

[15] Z. Kukelova and T. Pajdla. A minimal solution to the autocalibration of radial distortion. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1 –7, june 2007.

[16] Z. Kukelova and T. Pajdla. Two minimal problems for cameras with radial distortion. In *International Conference on Computer Vision (ICCV)*, pages 1 –8, 2007.

[17] Zuzana Kukelova and Tomas Pajdla. Two minimal problems for cameras with radial distortion. In *OMNIVIS*, 2007.

[18] Zuzana Kukelova, Martin Byröd, Klas Josephson, Tomas Pajdla, and Kalle Aström. Fast and robust numerical solutions to minimal problems for cameras with radial distortion. *Computer Vision and Image Understanding (CVIU)*, 114(2):234 – 244, 2010.

[19] B. Micusik and T. Pajdla. Estimation of omnidirectional camera model from epipolar geometry. In *Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 485–490, june 2003.

[20] J. Philip. A non-iterative algorithm for determining all essential matrices corresponding to five point pairs. *The Photogrammetric Record*, 15(88):589–599, 1996.

[21] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Photo tourism: Exploring photo collections in 3d. In *SIGGRAPH*, pages 835–846, 2006.

[22] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Modeling the world from Internet photo collections. *International Journal of Computer Vision (IJCV)*, 80(2):189–210, November 2008.

[23] Peter Sturm. On focal length calibration from two views. In *Computer Vision and Pattern Recognition (CVPR)*, pages 145–150, 2001.

[24] S. Thirthala and M. Pollefeys. Multi-view geometry of 1d radial cameras and its application to omnidirectional camera calibration. In *International Conference on Computer Vision (ICCV)*, volume 2, pages 1539–1546, 2005.

[25] Magdalena Urbanek, Radu Horaud, and Peter Sturm. Combining off- and online calibration of a digital camera. In *International Conference on 3D Digital Imaging and Modeling (3DIM)*, pages 99–106, 2001.