Two-View Focal Length Estimation for All Camera Motions Using Priors

Manmohan Krishna Chandraker

Abstract

Direct solutions to systems of polynomial equalities have been established in literature pertaining to two-view geometry to extract the focal lengths of a camera pair from the fundamental matrix, when other intrinsic calibration parameters are known. But these formulae break down near degeneracies characterized as critical motions, which occur quite commonly in practical scenarios. The contribution of this report is twofold. First, a novel method of imposing priors on the calibration matrix is presented that allows us to formulate the problem of evaluating the focal length as a constrained minimization and arrive at a correct estimate even for degenerate configurations. Second, we explore solution techniques that achieve a global minimum for polynomial objective functions with feasible region constrained to be a semialgebraic set defined by polynomial equalities. Preliminary experimental results are presented on simulated noiseless data to demonstrate the validity of our theory, both in achieving a global minimum and doing so even for critical motions.

1. Introduction

Self-calibration is the process of determination of the camera exterior orientation and intrinsic parameters using constraints on the entries of the camera calibration matrix (such as zero skew and rigidity constraints). It is known that the seven degrees of freedom of the fundamental matrix in the two-view case can yield estimates of the relative orientation of the cameras and the two focal lengths. The basic constraints for the two-view case are the *Kruppa equations* which can be considered to be an algebraic representation of the correspondence of epipolar lines tangent to the absolute conic on the plane at infinity. The Kruppa equations lend two independent constraints and all two-view focal length estimation procedures can be derived from them.

But there is a class of degenerate configurations for which the focal length cannot be estimated in an auto-calibration set-up where all other internal parameters are known. These configurations do not depend on specific values of the camera parameters, but are generic singularities which only depend on the relative orientation of the two cameras with respect to the scene points. So, these critical configurations are more aptly termed critical motions. Standard relations between the focal length and the entries of the fundamental matrix, such as [3, 11, 1, 9] break down at these degeneracies and give arbitrary estimates in their vicinity.

These critical motions are quite common in practice, especially in the two-view case. For instance, one camera motion that results in a degeneracy in the solution to two-view self-calibration is when the two camera axes and the baseline are coplanar. Several industrial application stereo heads are in fact designed such that their sensors "look at" a common scene point. Also common is the degeneracy in stereo placement in many robotic vision systems - the sensors are designed to look straight ahead, thus, their optical axes are parallel and both are coplanar with the line joining their centers. We would also expect near-degenerate configurations in applications that involve a photographic sampling of the scene, such as in panoramic mosaic construction.

We propose to impose priors on the camera calibration matrix to regularize the focal length estimation problem and pose it as a optimization problem. More specifically, we derive a robust formulation in terms of a polynomial objective function to be minimized subject to polynomial constraints. The minimization can be easily solved in a Levenberg-Marquardt iterative subroutine using Lagrange multipliers.

The drawback of such iterative techniques is that they easily get stuck in local minima. But solution techniques exist that can achieve a true global minimum, or approach a global minimum through convex relaxations, for a polynomial objective function constrained within a semialgebraic set by polynomial inequalities. We demonstrate the applicability of recent work in algebraic geometry that reduces the polynomial minimization problem to a positive semidefinite program by using convex linear matrix inequality (LMI) relaxations.

We show the correctness of our theory by simulation experiments on synthetic, noiseless two-view data. It is verified that imposing priors enables us to extract the correct focal length even for degenerate configurations. But iterative estimation procedures can get trapped in local minima for bad initializations, this menace is demonstrably alleviated by using a global optimization routine for polynomials.

2. Background

2.1. Previous Work

An early algorithm for focal length estimation from the fundamental matrix is presented in [3]. Other approaches, such as the SVD method in [11] and the procedure in [1], use the Kruppa equations more explicitly. [4] provides a comprehensive overview of concepts from multi-view geometry and the Kruppa equations in particular.

A linear algorithm for focal length estimation from the dual image of the absolute conic is developed in [9]. It also provides a complete characterization of critical motions for the two-view case. The case of n-view projective reconstructions is approached through subgroup constraints in [7] and completely characterized in [6] .

Global minimization of polynomials over not necessarily convex, compact sets is treated in [8]. An exciting new field of study for extrema analysis of polynomial functions is the concept of Gröbner bases [2] that we plan to use in our future implementations.

2.2. Geometry of Two Views

To clarify the notation, we will use $\mathbf{K}$ for the upper triangular 3×3 intrinsic calibration matrix, $\mathbf{R}$ to denote the rotation and $\mathbf{t}$ for the translation of one camera relative to the other. f stands for focal length, (u, v) denotes the principal point. $\omega^* = \mathbf{K}\mathbf{K}^\top$ is the dual image of the absolute conic. The reader is referred to the exposition in [4] for a review of these concepts.

Given two views $\mathcal{I}$ and $\mathcal{I}'$ of a scene and a sufficient set of homogeneous point correspondences $\mathbf{x} \leftrightarrow \mathbf{x}'$ between them, the fundamental matrix, $\mathbf{F}$, between the two views is a homogeneous 3×3, rank 2 matrix that satisfies

$$\mathbf{x}'^\top \mathbf{F}\mathbf{x} = 0 \tag{1}$$

For any point $\mathbf{x} \in \mathcal{I}$, the corresponding epipolar line $\mathbf{l}' \in \mathcal{I}'$ is given by $\mathbf{l}' = \mathbf{F}\mathbf{x}$. The right null vector of $\mathbf{F}$ is the epipole in the first image, $\mathbf{e}$ and the left null-vector is the epipole in the second image, $\mathbf{e}'$.

The geometry of two views is compactly encapsulated in the Kruppa equations

$$[\mathbf{e}']_\times \, \omega^{*\prime} \, [\mathbf{e}']_\times = \mathbf{F}\omega^* \mathbf{F}^\top \tag{2}$$

which are a set of nine equations, although due to homogeneity and symmetry, only two of them are independent.

2.3. Focal Length Calibration

The fundamental matrix has seven degrees of freedom (dof) - nine elements in the 3×3 homogeneous transformation with eight independent ratios and one dof less for rank deficiency. Out of these, five degrees of freedom encode the relative exterior orientation $(\mathbf{R}, \mathbf{t})$ of the two cameras. The remaining two dof can be used to estimate the focal lengths of the two cameras, assuming knowledge of the other intrinsic parameters.

Let the intrinsic calibration matrix of the first camera be of the form,

$$\mathbf{K} = \begin{bmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

that is, the skew is zero, aspect ratio unity and the position of the principal point is at the origin (which, most of the time, are fair assumptions in practice). Similarly, the calibration matrix for the second view is assumed $\text{diag}(f', f', 1)$. If the coordinate system of the two images are rotated such that both the epipoles lie on the positive u-axis, that is, $\mathbf{e} = (e_1, 0, e_3)^\top$ and $\mathbf{e}' = ({e'}_1, 0, {e'}_3)^\top$, then the fundamental matrix has a special form:

$$F \sim \begin{bmatrix} {e'}_3 & & \\ & 1 & \\ & & -{e'}_1 \end{bmatrix} \begin{bmatrix} a & b & a \\ c & d & c \\ a & b & a \end{bmatrix} \begin{bmatrix} e_3 & & \\ & 1 & \\ & & -e_1 \end{bmatrix} \tag{3}$$

for some $a, b, c, d \in \mathbb{R}$.

With this representation, direct formulae may be computed for the two focal lengths [3]:

$$f^2 = \frac{-ace_1^2}{ace_3^2 + bd}\,, \quad {f'}^2 = \frac{-abe'^2_1}{abe'^2_3 + cd} \tag{4}$$

An alternative approach is presented in [11], where the focal lengths are assumed to be the same for the two views and a Singular Value Decomposition of the fundamental matrix is used to solve the Kruppa Equations.

Let $\mathbf{F} = \mathbf{U}\mathbf{D}\mathbf{V}^\top$ be the SVD of $\mathbf{F}$, where $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3]$ and $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3]$ are orthonormal matrices and $\mathbf{D} = \text{diag}(\sigma_1, \sigma_2, 0)$ is the diagonal matrix of singular vectors, with the last singular value zero as the matrix $\mathbf{F}$ is rank-deficient. Note that $\mathbf{v}_3$ is the same as the right epipole $\mathbf{e}$, $\mathbf{u}_3$ is the left epipole $\mathbf{e}'$. Substituting these along with the SVD representation of $\mathbf{F}$ into (2), we get three equations

$$\frac{\mathbf{u}_2^\top \omega^{*\prime} \mathbf{u}_2}{\sigma_1^2 \mathbf{v}_1^\top \omega^* \mathbf{v}_1} = \frac{-\mathbf{u}_1^\top \omega^{*\prime} \mathbf{u}_2}{\sigma_1 \sigma_2 \mathbf{v}_1^\top \omega^* \mathbf{v}_2} = \frac{\mathbf{u}_1^\top \omega^{*\prime} \mathbf{u}_1}{\sigma_2^2 \mathbf{v}_2^\top \omega^* \mathbf{v}_2} \tag{5}$$

only two of which are linearly independent. A solution for f can be obtained for the case where $f = f'$ and the other intrinsic parameters are known (that is, $\omega^* = \omega^{*\prime} = \text{diag}(f^2, f^2, 1)$).

For the more general case of unequal focal lengths and principal point not at the origin, the focal length can still be (a bit surprisingly!) computed directly [1]:

$$f = \sqrt{\frac{-\mathbf{p}'^\top [\mathbf{e}']_\times \tilde{\mathbf{I}}\mathbf{F}\mathbf{p}\mathbf{p}^\top \mathbf{F}^\top \mathbf{p}'}{\mathbf{p}'^\top [\mathbf{e}']_\times \tilde{\mathbf{I}}\mathbf{F}\tilde{\mathbf{I}}\mathbf{F}^\top \mathbf{p}'}} \tag{6}$$

where $\mathbf{p} = (u, v)^\top$ and $\mathbf{p}' = (u', v')^\top$ are the principal points in the two images and $\tilde{\mathbf{I}} = \mathrm{diag}(1, 1, 0)$.

2.4. Critical Motions

It is well-known that self-calibration fails for certain special classes of camera motion, called critical motions. These critical motions depend on the constraints imposed on the autocalibration algorithm, such as zero skew, but are independent of the specific values of the internal parameters. In particular, we are interested in critical motions arising in stereo self-calibration scenarios where the intrinsic parameters of the cameras are known modulo an unknown focal length.

The following classic result from two-view geometry provides a complete characterization of critical motions in two-view autocalibration:

Theorem 1. *All degenerate configurations for the problem of focal length recovery from the fundamental matrix with other internal parameters known can be reduced to one of the following two configurations:*

1. *The optical axes of the two cameras and the baseline are coplanar.*
2. *One optical axis, the baseline and the vector perpendicular to the baseline and the other optical axis are coplanar.*

Conversely, any two-view imaging system with the sensors in either of the above two configurations cannot uniquely determine the focal length from image data alone.

The interested reader is referred to a specific proof in [9] and a more general treatment in terms of subgroup constraints on the camera calibration matrix in [7].

3. Priors in Focal Length Estimation

As discussed in the previous section, traditional Kruppa constraint based methods fail to estimate the correct focal length around the degeneracy. But often in practice, one has some idea of an approximate value of the focal length. This can be used to impose a prior, f_{prior} on the focal length and the problem can be formulated as a minimization problem. From here onwards, we would consider the case where $f = f'$. Not only does this assumption keep the dimensionality of the estimation problem lower, but also it is actually the more common case in practice. For example, it is common to have both the cameras in a stereo pair calibrated to common focal length, or to keep a common focal length between successive snapshots when a panoramic mosaic is being constructed. But we emphasize that the treatment is analogous for the case of unequal focal lengths, the only change would be that the polynomials involved would have an extra variable.

3.1. A Naïve Approach

As a first attempt, consider the following minimization problem

$$\min_f w_1^2(f - f_1)^2 + w_2^2(f - f_2)^2 + w_3^2(f - f_{\mathrm{prior}})^2 \quad (7)$$

where f_1 anf f_2 are the values returned by the formulae such as (4) and w_i^2 are some positive weights, for $i = 1, 2, 3$. (For camera configurations far from a degeneracy, f_1 and f_2 would be the same.)

Consider a case where the camera is (nearly) in a critical configuration. The erroneous values of focal lengths f_1 and f_2 computed by methods discussed in Section 2.3 would lead to arbitrary answers for the focal length estimate, f^*, in (7) as the minimum is just at the weighted average of f_1, f_2 and f_{prior}. Consequently, the minimization in (7) is not robust for configurations near critical.

Moreover, the goodness of the estimate is critically dependent on how close the guessed f_{prior} is to the true focal length, f^*. Thus, the optimal estimate is greatly influenced by the initialization, which is not desirable in a practical situation.

3.2. Problem Formulation

The observation we exploit is that although the prior on the focal length can be devious at times, we know that the principal point is exactly at the origin. Thus, we can introduce new variables for the (x, y)-coordinates of the principal point, with a strongly weighted prior at $(0, 0)$ to impart a greater degree of robustness to the optimization problem.

The calibration matrix under consideration now has the form

$$\mathbf{K} := \mathbf{K}(f, u, v) = \begin{bmatrix} f & 0 & u \\ 0 & f & v \\ 0 & 0 & 1 \end{bmatrix} \quad (8)$$

and the corresponding dual image of the absolute conic could be parametrized as

$$\omega^* := \omega^*(f, u, v) = \mathbf{K}\mathbf{K}^\top = \begin{bmatrix} f^2 + u^2 & uv & u \\ uv & f^2 + v^2 & v \\ u & v & 1 \end{bmatrix} \quad (9)$$

The estimation problem can now be formulated as an unconstrained minimization

$$\mathbb{U} \mapsto \min_{f,u,v} \sum_{i=1}^{2} w_i^2(f - f_i)^2 + w_f^2(f - f_{\mathrm{prior}})^2 + w_{uv}^2 \|(u, v) - (0, 0)\|^2 \quad (10)$$

Solution to $\mathbb{U}$ involves evaluating the initial values of the focal lengths using one of the approaches in Section 2.3, for instance, the formula (6). But note that these formulae are derived for theoretical scenarios, they do not make allowances for ill-conditioning one might encounter in practical situations. Therefore, even in a noiseless case, near a degeneracy, these formulae break down and predict meaningless values.

A better approach is to formulate the above as a constrained minimization problem. Any relation that evaluates the focal length in a two-view case must be based on the Kruppa equations. The Kruppa equations in (5), with $\omega^{*\prime} = \omega^*$ enforced, are rewritten as polynomial functions of f, u and v:

$$\begin{aligned} p_1(f,u,v) &:= \sigma_1(\mathbf{v}_1^\top \omega^* \mathbf{v}_1)(\mathbf{u}_1^\top \omega^* \mathbf{u}_2) \\ &\quad + \sigma_2(\mathbf{v}_1^\top \omega^* \mathbf{v}_2)(\mathbf{u}_2^\top \omega^* \mathbf{u}_2) = 0 \\ p_2(f,u,v) &:= \sigma_1(\mathbf{v}_1^\top \omega^* \mathbf{v}_2)(\mathbf{u}_1^\top \omega^* \mathbf{u}_1) \\ &\quad + \sigma_2(\mathbf{v}_2^\top \omega^* \mathbf{v}_2)(\mathbf{u}_1^\top \omega^* \mathbf{u}_2) = 0 \\ p_3(f,u,v) &:= \sigma_1^2(\mathbf{v}_1^\top \omega^* \mathbf{v}_1)(\mathbf{u}_1^\top \omega^* \mathbf{u}_1) \\ &\quad - \sigma_2^2(\mathbf{v}_2^\top \omega^* \mathbf{v}_2)(\mathbf{u}_2^\top \omega^* \mathbf{u}_2) = 0 \end{aligned} \quad (11)$$

Thus, the constrained version of $\mathbb{U}$ is our problem formulation:

$$\mathbb{C} \mapsto \begin{cases} \min\limits_{f,u,v} \quad w_{\text{prior}}^2 (f - f_{\text{prior}})^2 + w_{uv}^2 \|(u,v)\|^2, \\ \quad p_1(f,u,v) = 0 \\ \quad p_2(f,u,v) = 0 \\ \quad p_3(f,u,v) = 0 \end{cases} \quad (12)$$

The minimization problem $\mathbb{C}$ can be solved using standard iterative techniques such as Gauss-Newton or Levenberg-Marquardt by first converting to an unconstrained problem, $\mathbb{L}$, using Lagrange Multipliers. Explicitly,

$$\begin{aligned} \mathbb{L} \mapsto \min_{f,u,v,\lambda_1,\lambda_2,\lambda_3} & w_f^2 (f - f_{\text{prior}})^2 + w_{uv}^2 \|(u,v)\|^2 \\ & + \lambda_1 |p_1(f,u,v)| + \lambda_2 |p_2(f,u,v)| \\ & + \lambda_3 |p_3(f,u,v)| \end{aligned} \quad (13)$$

But as the following section discusses, better solution techniques are available at our disposal that are not limited to local optimality.

4. Global Minimization Using Convex LMI Relaxations

Note that the objective function as well as the constraints in minimization problem $\mathbb{C}$ are low degree polynomials in a small number of variables. Recent advances in algebraic geometry [2] and semidefinite programming [8] show that it is possible to find a global minimum for such objective functions, constrained by polynomial inequality constraints.

The optimization problem that we are dealing with has the form

$$\mathbb{P}_K \to p_K^* := \min_{\mathbf{x} \in \mathbb{R}^n} \{p(\mathbf{x}) \mid p_k(\mathbf{x}) \geq 0,\ k = 1, \cdots, m\} \quad (14)$$

where K stands for the compact set defined by the inequality constraints, that is,

$$K := \{\mathbf{x} \in \mathbb{R}^n | p_k(\mathbf{x}) \geq 0, k = 1, \cdots, n\}$$

Here $\mathbf{x} = (x_1, \cdots, x_n)$ is a point in n-space and $p(\mathbf{x}) = \sum\limits_\alpha p_\alpha(\mathbf{x}^\alpha)$ is a degree m real-valued polynomial expressed by a coefficient vector $p = \{p_\alpha\} \in \mathbb{R}^{s(m)}$ with respect to the $s(m)$-cardinality set of basis monomials

$$\{\mathbf{x}^\alpha := x_1^{\alpha_1} \cdots x_n^{\alpha_n} | \alpha_i \geq 0, \sum_i \alpha_i \leq m, i = 1, \cdots, n\} \quad (15)$$

Just to demonstrate the approach, we will review here the simpler unconstrained problem (16). (The approach is similar in the constrained case and the reader is pointed towards [8] for the relevant details.)

$$\mathbb{P} \to p^* := \min_{\mathbf{x} \in \mathbb{R}^n} p(x) \quad (16)$$

The idea is to consider the dual problem

$$\mathcal{P} \mapsto p^* := \min_{\mu \in \mathcal{P}(\mathbb{R}^n)} \int p(\mathbf{x})\, \mu(d\mathbf{x}) \quad (17)$$

Recall from probability theory that for some $s(2m)$-vector $\mathbf{y} := y_\alpha$, the $(i+j)$-th order moment is defined by $y_{i,j} = \int \mathbf{x}^i \mathbf{y}^j \mu(d(\mathbf{x}, \mathbf{y}))$ and the dimension $s(m)$ moment matrix $M_m(\mathbf{y})$ is defined as the block matrix $\{M_{i,j}(\mathbf{y})\}_{0 \leq i,j \leq 2m}$ where

$$M_{i,j}(\mathbf{y}) = \begin{bmatrix} y_{i+j,0} & y_{i+j-1,0} & \cdots & y_{i,j} \\ y_{i+j-1,1} & y_{i+j-2,2} & \cdots & y_{i-1,j+1} \\ \vdots & \cdots & & \vdots \\ y_{j,i} & y_{i+j-1,1} & \cdots & y_{0,i+j} \end{bmatrix}$$

Now consider the positive semidefinite program:

$$\mathbb{Q} \mapsto \begin{cases} \inf\limits_{\mathbf{y}} \sum\limits_\alpha p_\alpha y_\alpha, \\ M_m(\mathbf{y}) \succeq 0 \end{cases} \quad (18)$$

The following theorem, proved in [8], establishes that the minimization problem in $\mathcal{P}$ (and hence, $\mathbb{P}$) is equivalent to the convex LMI problem $\mathbb{Q}$ when the polynomial $p(\mathbf{x})$ is expressible as a sum of squares of polynomials up to a constant.

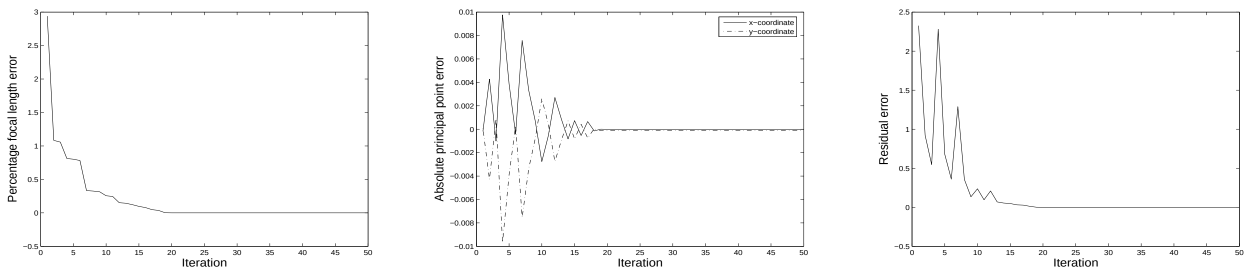

(a) Focal length estimation

(b) Principal point estimation

(c) Principal point estimation

Figure 1: *Plots of (a) focal length error, (b) principal point error and (c) residual error for the iterative Levenberg-Marquardt algorithm with the cameras in a general, non-degenerate configuration. The estimates converge to a zero error minimum for both the parameters and the residual. The prior used here for the focal length has a* 12.5% *negative offset relative to the true focal length.*

Theorem 2. *Let* $p(x) : \mathbb{R}^n \to \mathbb{R}$ *be a degree* $2m$ *polynomial with a global minimum,* $\min \mathbb{P}$*. If the non-negative polynomial* $p(x) - p^*$ *is a sum of squares of polynomials and* x^* *is the global minimizer of* $\mathbb{P}$*, then*

$$y^* = (x_1^*, \cdots, x_n^*, (x_1^*)^2, x_1^* x_2^*, \cdots, (x_n^*)^{2m}) \qquad (19)$$

is a minimizer of $\mathbb{Q}$ *and* $\min \mathbb{P} = \min \mathbb{Q}$*.*

Conversely, if the dual to $\mathbb{Q}$ *has a feasible solution, then* $\min \mathbb{P} = \min \mathbb{Q}$ *only if* $p(x) - p^*$ *is a sum of squares of polynomials.*

It can further be shown that for the general case where $p(x) - p^*$ is not a sum of squares, the global optimum can be achieved through solving a series of convex LMI relaxations of increasing order, but of a form similar to (18). It relies on the following result from polynomial theory shows that it is always possible for a non-negative polynomial to be "approached" by polynomials that are sums of squares.

Theorem 3. *Let* $\mathcal{A}$ *be the space of real-valued polynomials* $p(x) : \mathbb{R}^n \to \mathbb{R}$ *equipped with the norm* $\|p(x)\|_{\mathcal{A}} = \|p\|$ *where* p *is the finite dimensional vector of coefficients of the polynomial* $p(x)$ *in the basis (15). Then the cone of polynomials which are sums of squares is dense (for the norm* $\|\cdot\|_{\mathcal{A}}$*) in the set of polynomials that are non-negative on* $[-1, 1]^n$*.*

Thus, even when we have a polynomial for which $p(\mathbf{x}) - p^*$ is not expressible as a sum of squares and the relaxations do not converge, it is possible to perturbate $p(\mathbf{x})$ by adding higher order terms that do not have a significant impact on the value of the polynomial in the domain under consideration, yet make the resulting polynomial expressible as a sum of squares. So, LMI relaxations of a correspondingly higher order would yield a global extremum for the semidefinite program corresponding to the perturbed polynomial.

5. Experiments

5.1. Iterative Focal Length Estimation

The objective function in $\mathbb{L}$ is optimized over the focal length, principal point and the Lagrange Multiplier variables using an implemetation of Levenberg-Marquardt iterative estimation procedure. A slightly modified analytical equivalent of the objective function in (13) is used:

$$\begin{aligned}\mathbb{L} \mapsto \min_{f,u,v,\lambda_1,\lambda_2,\lambda_3} & w_f^2 (f - f_{\text{prior}})^2 + w_{uv}^2 (u^2 + v^2) \\ & + \lambda_1 p_1^2(f,u,v) + \lambda_2 p_2^2(f,u,v) \\ & + \lambda_3 p_3^2(f,u,v) \qquad (20)\end{aligned}$$

For a configuration where the relative camera positions are not in a critical configuration, the minimization is quite good for an initialization not very far from the true minimum. As Figure 1 shows, even for a prior that differs from the true focal length by 12.5%, a zero-error minimum is achieved. The residual after several iterations was very low, of the order of 10^{-7}. This is expected as the cameras were in a general configuration for this experiment and the minimization routine implicitly weighted the Kruppa equations, which must be exactly satisfied, more than the prior terms.

A more interesting result is when the cameras are in a critical configuration. The critical configuration was generated by ensuring a rotation that makes the optical axes intersect on the z-axes. It is a simple exercise to show that the rotation matrix in such a case must have its third row equal to $(0, 1/\sqrt{2}, 1/\sqrt{2})$. We can complete the rotation matrix by selecting appropriate rotations or using Householder reflections. The camera centers were then translated by a vector of the order of the focal length to ensure that the principal axes intersect at some scene point and not at the origin or at infinity. The residual after several iterations of the Levenberg-Marquardt routine stabilized at around 0.02 and the corresponding error in focal length estimation was 0.2%. See Figure 2.

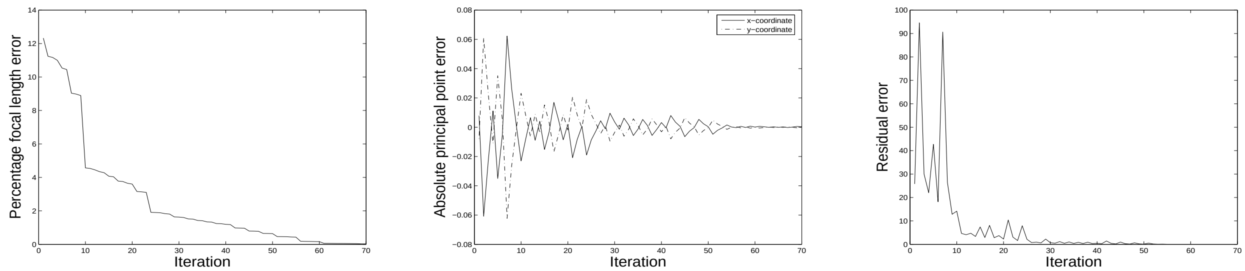

(a) Focal length estimation (b) Principal point estimation (c) Principal point estimation

Figure 2: *Plots of (a) focal length error, (b) principal point error and (c) residual error for the iterative Levenberg-Marquardt algorithm with the cameras in a critical configuration. The estimates converge to a zero-error minimum for both the parameters and the residual. The prior used here for the focal length has a* 12.5% *negative offset relative to the true focal length.*

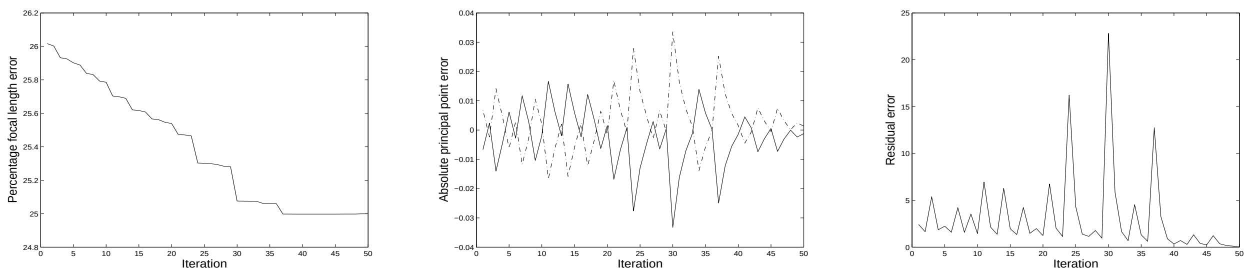

(a) Focal length estimation (b) Principal point estimation (c) Principal point estimation

Figure 3: *Plots of (a) focal length error, (b) principal point error and (c) residual error for the iterative Levenberg-Marquardt algorithm with the same camera configuration as in Fig 1 but a different initialization. The estimates converge to a local minimum for both the parameters and the residual. The prior used here for the focal length has a* 12.5% *negative offset relative to the true focal length.*

But as we have pointed out earlier, a drawback of iterative estimation procedures is that they guarantee only local optimality. Thus, for a slightly worse initialization, but with the same prior on the focal length, the algorithm converges to a local minimum (Figure 3) where the error in focal length estimate is nearly 25%.

5.2. Global Minimization using GloptiPoly

Before analyzing the results of the global minimization, we provide a brief description of the routines used for the implementation. *GloptiPoly* [5] is a MATLAB toolbox that solves convex linear matrix inequality relaxations to minimize a multivariable polynomial function subject to polynomial inequality or integer constraints. The theory discussed in presented in [8] and discussed in Section 4 is used to generate a series of lower bounds that monotonically converge to the global optimum.

The MATLAB semidefinite program solver *SeDuMi* [10] is used as a subroutine in GloptiPoly to solve LMI relaxations over symmetric cones. GloptiPoly can be interfaced with the MATLAB *Symbolic Toolbox* to provide a convenient syntax similar to Maple.

A drawback of GloptiPoly is that it is designed only for medium-scale polynomials. It has a design limit of up to 19 variables, but is difficult to parametrize for higher order polynomials in far fewer variables. But the minimization problem we are concerned with consists of up to degree 4 polynomials in 3 variables, so GloptiPoly works perfectly well. GloptiPoly is designed for inequality polymial constraints, but handles equality constraints $p(\mathbf{x}) = 0$ by splitting into two inequalities $p(\mathbf{x}) \leq 0$ and $p(\mathbf{x}) \geq 0$.

The dependence of focal length and principal point variation with respect to the relative weights of their prior terms is studied in Figure 4. Not surprisingly, the estimates converge more closely to the correct ground truth values as when the principal point is heavily weighted, as the principal point is known *a priori* to be at the origin, while the prior on the focal length is only approximate.

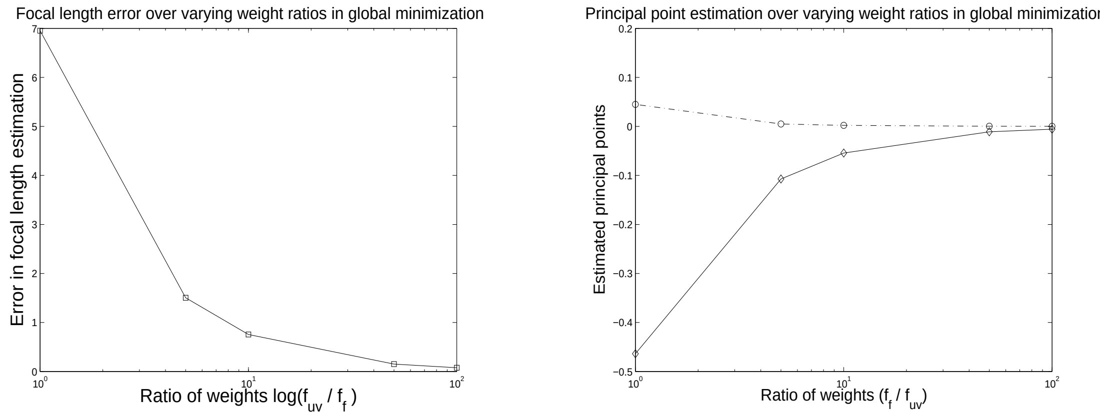

(a) Focal length estimation

(b) Principal point estimation

Figure 4: *Plots of (a) focal length and (b) principal point as the ratio of the weights assigned to the two terms in the objective function is varied. Notice that the estimates converge towards true values (zero error) as the weighting for the term corresponding to the principal points is increased relative to the weighting for the term corresponding to the focal length. This is expected as the principal point location is known exactly, while the prior used here for the focal length has a* 33% *negative offset relative to the true focal length.*

Focal length estimation of a non-degenerate camera pair is analyzed in Figure 5. A large erroneous prior is applied to the principal point as the prior on the focal length is varied. There is a smooth degradation of the performance of our formulation as the prior offset is increased in either direction of zero, with a perfect estimation for a prior close to ground truth.

The most interesting result is, of course, for a critical camera motion, constructed as described in the previous subsection. As the prior offset as varied, we observe a well-behaved degradation of the focal length error estimate. Compare this with the case of direct formulae for focal length computation, which output arbitrary values for configurations close to degeneracy. An accurate global minimum is obtained for priors close to the true focal length.

6. Discussions

We have discussed in this report a novel approach to two-view focal length estimation by imposing priors that performs well even for critical camera motions. The problem is formulated as a polynomial minimization subject to polynomial constraints, so mathematical tools for achieving a global minimum can be exploited.

Further work in progress includes a noise and robustness analysis. An obvious generalization currently being investigated is to the three-view case, it leads to new challenges in formulating the problem as a polynomial minimization. Finally, we plan to use of Gröbner bases for exact global minimization, rather than the relaxation approach of semidefinite programming.

References

[1] S. Bougnoux. From projective to euclidean space under any practical situation, a criticism of self-calibration. In *ICCV*, pages 790–798, 1998.

[2] D. Cox, J. Little, and D. O'Shea. *Using Algebraic Geometry*. Springer-Verlag, 1998.

[3] R. Hartley. Extraction of focal lengths from the fundamental matrix. Technical report, G.E.CRD, Schenectady, NY, 1993.

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

[5] D. Henrion and J. Lasserre. GloptiPoly: Global optimization over polynomials with Matlab and SeDuMi. Technical report, LAAS-CNRS, April 2002.

[6] F. Kahl, R. Hartley, and K. Åström. Critical configurations for n-view projective reconstruction. In *CVPR*, pages 158–163, 2001.

[7] F. Kahl, B. Triggs, and K. Åström. Critical motions for auto-calibration when some intrinsic parameters can vary. *Journal of Mathematical Imaging and Vision*, 13(2):131–146, October 2000.

[8] J. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.

[9] G. Newsam, D. Huynh, M. Brooks, and H.-P. Pan. Recovering unknown focal lengths in self-calibration: An essentially linear algorithm and degenerate configurations. In *Int. Arch. Photogrammetry & Remote Sensing*, pages 575–580, 1996.

[10] J. Sturm. Using sedumi 1. 02, a Matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11(12):625–653, 1999.

[11] P. Sturm. On focal length calibration from two views. In *CVPR*, pages 145–150, 2001.

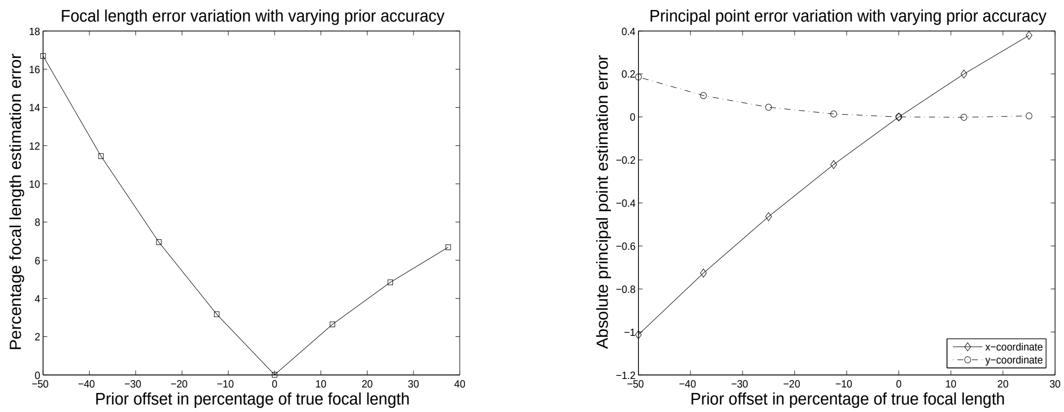

(a) Focal length estimation

(b) Principal point estimation

Figure 5: *Plots of (a) focal length error and (b) principal point error as the focal length prior is varied for a non-degenerate configuration. The initialization for the principal point is also erroneous, being (1,1) when the true value is known to coincide with the origin. Note that the estimates are nearly perfect (zero error) for a prior close to the true focal length.*

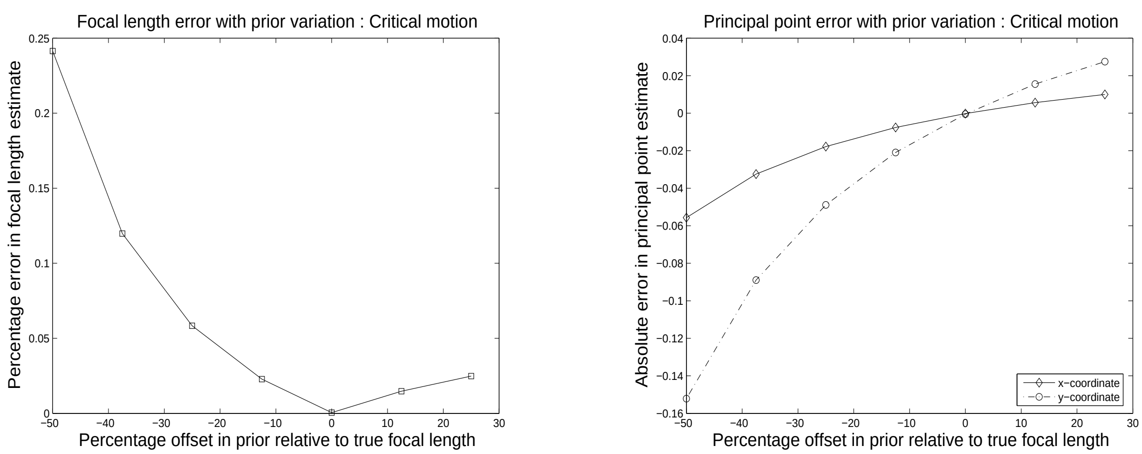

(a) Focal length estimation

(b) Principal point estimation

Figure 6: *Plots of (a) focal length error and (b) principal point error as the focal length prior is varied for a critical configuration. The initialization for the principal point is also erroneous, being (-0.1,-0.2) when the true value is known to coincide with the origin. Note that the estimates are nearly perfect (zero error) for a prior close to the true focal length, whereas they would be arbitrary (or very deviant) for the direct Kruppa constraint-based formulae.*