

A Convex Optimization Approach to Robust Fundamental Matrix Estimation *

Y. Cheng J. A. Lopez O. Camps M. Sznaier

Electrical and Computer Engineering
Northeastern University, Boston, MA 02115

{cheng.yong,lopez.jo}@husky.neu.edu, {camps,msznaier}@coe.neu.edu

Abstract

This paper considers the problem of estimating the fundamental matrix from corrupted point correspondences. A general nonconvex framework is proposed that explicitly takes into account the rank-2 constraint on the fundamental matrix and the presence of noise and outliers. The main result of the paper shows that this non-convex problem can be solved by solving a sequence of convex semi-definite programs, obtained by exploiting a combination of polynomial optimization tools and rank minimization techniques. Further, the algorithm can be easily extended to handle the case where only some of the correspondences are labeled, and, to exploit co-occurrence information, if available. Consistent experiments show that the proposed method works well, even in scenarios characterized by a very high percentage of outliers.

1. Introduction

The fundamental matrix provides a scene-independent algebraic representation of the epipolar geometry of a stereo camera system, that depends only on the cameras' intrinsic parameters and relative pose [6]. Given the key role that this matrix plays in a large number of problems in computer vision applications, such as stereo camera calibration, 3D reconstruction, etc., the problem of estimating it from experimental data (typically consisting of point correspondences from two images) has been the subject of a large research effort. In general, existing techniques reduce the problem to a constrained optimization, where the main difficulties

stem from the need to impose that the resulting matrix must be rank deficient, and from the presence of outliers, which if not properly handled can substantially skew the estimate.

Existing methods fall roughly into one of two classes, depending on how they tackle the non-convex rank-2 constraint. The first type of methods uses a two-step approach, for instance, methods based on the well-known eight-point algorithm and its various extensions [12, 6, 3], that starts by finding a (sub)optimal estimate without taking into account the rank constraint and then refines the result by reducing its smallest singular value to 0. However, in [13] it has been shown that ignoring the rank constraint can degrade the accuracy in terms of the covariance matrix of the first-order variation of the solution. On the other hand, the second type of methods considers the rank constraint explicitly. In [16] a Levenberg-Marquard (LM) approach is proposed to optimize the singular value decomposition (SVD) of the fundamental matrix. In [20] and [1] the rank constraint is imposed by setting its determinant to 0, leading to a 3rd-order polynomial constraint. Alternatively, in [2] and [21] the estimation problem is reduced to one or several constrained polynomial optimization problems by imposing the constraint that the null space of the solution must contain a non-zero vector. The resulting optimization problems are solved by resorting to various optimization techniques, such as branch-and-bound approaches [20, 21] or moments based convex relaxations [2, 1].

Although the above methods perform well under the assumption of small, suitably distributed (for instance Gaussian) measurement noise, their performance substantially degrade in the presence of even a few outliers (i.e., point mismatches). Thus, a large number of robust methods have been developed to explicitly take into account the presence of outliers. These include random-

*This work was supported in part by NSF grants IIS-1318145 and ECCS-1404163; AFOSR grant FA9550-12-1-0271, and the Alert DHS Center of Excellence under Award Number 2008-ST-061-ED0001.

ized methods such as RANdom Sampling Consensus estimator (RANSAC) [5] and its variants, which attempt to find outlier-free data by repeatedly randomly selecting the minimal number of correspondences needed to generate a solution, and selecting the best one, according to some optimality criteria. For instance, RANSAC selects the solution with largest support on the complete dataset. MSAC [17], a redecending M-estimator [8], penalizes both the squared fitting error of inliers and number of outliers. Least Median of Squares (LMS) [11] selects the estimate which gives the least median fitting error. Finally, Maximum Likelihood Estimation SAmple Consensus (MLEsAC) [18] attempts to find the maximum likelihood estimate of the true position of the points. Random sampling based methods are attractive due to their simplicity and the existence of theoretical bounds on the number of samples required to guarantee a given probability of success. However, they suffer from several weaknesses. Firstly, for a given probability of success, the number of needed iterations grows very fast with the number of outliers. Secondly, since the bounds explicitly depend on the number of outliers, this quantity must be known or estimated accurately, since stopping the algorithm prematurely can lead to arbitrarily bad solutions. Finally, these methods cannot directly impose the rank deficiency constraint. Rather, this is done a posteriori, by projecting the solution onto the manifold of rank-2 matrices. However, as indicated before, this step can lead to substantial performance degradation.

1.1. Main Idea and Paper Contributions

Motivated by the challenges noted above, in this paper we propose a novel single-step framework for robustly estimating the fundamental matrix from point correspondences corrupted by noise and outliers. The main idea is to introduce binary variables that indicate whether a given correspondence is an inlier (and hence should be used in the estimation), or an outlier (and hence should be discarded), and to explicitly impose the rank-2 constraint by searching for the epipoles. This formulation leads to a polynomial optimization over a semi-algebraic set that can be solved by appealing to recent results on sparse polynomial optimization. Specifically, the advantages of the proposed approach vis-à-vis existing techniques include the abilities to:

- Explicitly impose rank-deficiency and handle noise and a very large percentage of outliers, without the need for additional assumptions such as bounds on the number of outliers.
- Certify that a given convex relaxation has indeed

found an optimal estimate of the fundamental matrix (the ground truth for noiseless data and the one that maximizes the number of inliers in the case of noisy measurements).

- Exploit co-occurrence priors to improve the estimate.
- Handle partially known correspondences.
- Explicitly exploit the underlying sparse structure of the problem to reduce the computational burden.

In addition, we provide theoretical results showing that, if in the optimization above a certain matrix containing only variables related to the fundamental matrix has rank 1, then the first order relaxation of the problem achieves global optimality. Combining these ideas with rank-minimization techniques leads to a computationally efficient algorithm with complexity comparable to robust regression based techniques, while still retaining the advantages noted above. These results are illustrated with several examples where the proposed algorithm is shown to consistently outperform existing techniques.

2. Preliminaries

2.1. Notation

$\mathbf{X}, \mathbf{x}, x$	matrix, vector, scalar
$\mathbf{X}^T(\mathbf{x}^T)$	the transpose of $\mathbf{X}(\mathbf{x})$
\mathbb{R}, \mathbb{N}	set of real number and non-negative integers
\mathbf{I}	Identity matrix
$\mathbf{M} \succeq \mathbf{N}$	the matrix $\mathbf{M} - \mathbf{N}$ is positive semidefinite
$\sigma_i(\mathbf{A})$	the i -th largest singular value of matrix \mathbf{A}
$ \mathcal{I} $	cardinality of the set \mathcal{I}

2.2. Problem Statement

Given a pair of images of the same scene from two uncalibrated perspective views, the fundamental matrix $\mathbf{F} \in \mathbb{R}^{3 \times 3}$ is defined as the rank-2 matrix which satisfies the epipolar constraint

$$\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0, \forall \mathbf{x}', \mathbf{x}, \quad (1)$$

where the homogenous coordinates $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^3$ are the corresponding projections of a 3D point in the two images. \mathbf{F} has seven degrees of freedom due to the ambiguity caused by the scaling and singularity. Our goal is to develop a computationally tractable algorithm that, starting from noisy point correspondences corrupted by outliers, simultaneously estimates a rank-2 fundamental matrix that maximizes the number of inliers, and, at the

same time, explicitly identifies outliers, defined as those points whose distance from the surface defined by (1) is beyond a given bound. Specifically, we address the following problem:

Problem 1. *Given a set of noisy point correspondences, $\{\mathbf{x}_i, \mathbf{x}'_i\}$, $i = 1, \dots, n$ drawn from two images of the same scene, and a-priori bound on the fitting error $|\mathbf{x}'_i{}^T \mathbf{F} \mathbf{x}_i| \leq \epsilon$, find a fundamental matrix \mathbf{F} such that:*

- 1.- $\|\mathbf{F}\|_F = 1$;
- 2.- \mathbf{F} is rank-2;
- 3.- The number of inliers is maximized, that is $\max |\mathcal{I}(\mathbf{F})|$ where $\mathcal{I}(\mathbf{F}) \doteq \{(\mathbf{x}'_i, \mathbf{x}) : |\mathbf{x}'_i{}^T \mathbf{F} \mathbf{x}_i| \leq \epsilon\}$.

As we show in the sequel, the problem above can be recast as a constrained polynomial optimization and solved using the methods briefly described below.

2.3. Moments Based Polynomial Optimization

2.3.1 The Problem of Moments

Consider a constrained Polynomial Optimization Problem (POP) of the form

$$p_{\mathcal{K}}^* \doteq \min_{\mathbf{x} \in \mathcal{K}} p(\mathbf{x}) \quad (\text{P1})$$

where $p(\mathbf{x})$ is a multivariate polynomial expressed as a sum of monomials, that is $p(\mathbf{x}) \doteq \sum_{\alpha} p_{\alpha} \mathbf{x}^{\alpha}$, where $\alpha \doteq (\alpha_1, \dots, \alpha_n)$ with $\alpha_i \in \mathbb{N}$ is a multi-index, $\mathbf{x} \doteq (x_1, \dots, x_n)$, $x_i \in \mathbb{R}$ and, following standard notation, we have defined $\mathbf{x}^{\alpha} = \prod_{i=1}^n x_i^{\alpha_i}$. Here \mathcal{K} denotes a set defined by d inequalities of the form $g_k(\mathbf{x}) \geq 0$, $k = 1, \dots, d$, where $g_k(\cdot)$ is a multivariate polynomial also expressed as a sum of monomials, e.g. $\mathcal{K} \doteq \{\mathbf{x} \in \mathbb{R}^n : g_k(\mathbf{x}) \doteq \sum_{\beta} g_{k,\beta} \mathbf{x}^{\beta} \geq 0, k = 1, \dots, d\}$.

Problem (P1) is generically non-convex (except in a few special cases) and hence computationally challenging. Thus, we consider a related problem:

$$\tilde{p}_{\mathcal{K}}^* \doteq \min_{\mu \in \mathcal{P}(\mathcal{K})} \int p(\mathbf{x}) \mu(dx) = \min_{\mu \in \mathcal{P}(\mathcal{K})} \mathcal{E}_{\mu} [p(\mathbf{x})] \quad (\text{P2})$$

where $\mathcal{P}(\mathcal{K})$ is the space of probability measures on \mathcal{K} with $\int_{\mathcal{K}} d\mu = 1$ and $\mathcal{E}_{\mu}[\cdot]$ denotes the expectation. Although (P2) is an infinite dimensional problem, it is, in contrast to (P1), convex. As shown in [9], Problems (P1) and (P2) are equivalent, in the sense that

- $\tilde{p}_{\mathcal{K}}^* = p_{\mathcal{K}}^*$.
- For every optimal solution μ^* of (P2), $p(x) = p_{\mathcal{K}}^*$, μ^* -almost everywhere.

It follows (see [9, 10] for more details) that problem (P1) can be reduced to a sequence of Linear Matrix Inequalities (LMI) optimization problems in the moments of the unknown probability measure μ of the form

$$p_N^* = \min_{\mathbf{m}} \sum_{\alpha} p_{\alpha} m_{\alpha} \quad \text{s.t.} \quad \begin{aligned} \mathbf{M}_N(\mathbf{m}) &\succeq 0, \\ \mathbf{L}_{N-\lceil \frac{\delta_k}{2} \rceil}(g_k \mathbf{m}) &\succeq 0, \forall k=1 \end{aligned} \quad (2)$$

where $\mathbf{m} \doteq \{m_{\alpha^{(i)}} : m_{\alpha^{(i)}} \doteq \int_{\mathcal{K}} \mathbf{x}^{\alpha^{(i)}} \mu(dx)\}$ is the moments sequence of order up to $2N$, indexed by a multi-index sequence $\alpha^{(i)} \in \mathbb{N}^n$ arranged in a graded reverse lexicographic order, and \mathbf{M}_N and $\mathbf{L}_{N-\lceil \frac{\delta_k}{2} \rceil}$ are the truncated moment and localizing matrices given by

$$\begin{aligned} \forall_{i=1, j=1}^{S_N} : \mathbf{M}_N(\mathbf{m})(i, j) &= m_{\alpha^{(i)} + \alpha^{(j)}} \\ \forall_{i=1, j=1}^{S_{N-\lceil \frac{\delta_k}{2} \rceil}} : \mathbf{L}_{N-\lceil \frac{\delta_k}{2} \rceil}(g_k \mathbf{m})(i, j) &= \sum_{\beta} g_{k,\beta} m_{\beta + \alpha^{(i)} + \alpha^{(j)}}, \end{aligned} \quad (3)$$

where δ_k is the degree of polynomial g_k , and $S_N = \binom{N+n}{n}$ (e.g. the number of moments in \mathbb{R}^n up to order N).

Further, it can be shown that as N increases, p_N^* in (2) monotonically increases to $p_{\mathcal{K}}^*$ from below. The necessary and sufficient conditions to guarantee the equivalence between (2) and (P1) are either

- positive semi-definiteness of the infinite dimensional matrices \mathbf{M}_N and $\mathbf{L}_{N-\lceil \frac{\delta_k}{2} \rceil}$ as N increases to infinity; or
- for a finite N , the flat extension [4] property holds, that is, for $\mathbf{M}_N \succeq 0$, $\mathbf{M}_{N+1} \succeq 0$, and $\text{rank}(\mathbf{M}_N) = \text{rank}(\mathbf{M}_{N+1})$.

2.3.2 Exploiting an Underlying Sparse Structure

In many cases of practical interest, including the problem in this paper, the objective function and the polynomials defining the constraint set \mathcal{K} exhibit a sparse structure that can be exploited to substantially reduce the computational burden entailed in solving (2). To this effect, start by introducing the concept of *running intersection*:

Definition 1. *Consider problem (P1) and let $\mathcal{I}_k \subset \{1, \dots, n\}$ be the set of indices of variables such that each $g_j(\mathbf{x})$ contains variables only from some \mathcal{I}_k . Assume that the objective function $p(\mathbf{x})$ can be partitioned as $p(\mathbf{x}) = p_1(\mathbf{x}) + \dots + p_l(\mathbf{x})$ where each p_k contains*

only variables from \mathcal{I}_k . Problem (P1) is said to satisfy the running intersection property if there exists a re-ordering $\mathcal{I}_{k'}$ of \mathcal{I}_k such that for every $k' = 1, \dots, l-1$:

$$\mathcal{I}_{k'+1} \cap \bigcup_{j=1}^{k'} \mathcal{I}_j \subseteq \mathcal{I}_s \text{ for some } s \leq k' \quad (4)$$

As shown in [10], when the running intersection property holds, one can construct a convergent hierarchy of semidefinite programs of smaller size:

$$p_N^* = \min_{\mathbf{m}} \sum_{j=1}^l \sum_{\alpha^{(j)}} p_{j,\alpha^{(j)}} m_{\alpha^{(j)}} \quad (5)$$

s.t. $\mathbf{M}_{k,N}(\mathbf{m}_{\mathcal{I}_k}) \succeq 0, \forall k=1, \dots, l,$
 $\mathbf{L}_{j,N-\lceil \frac{\delta_j}{2} \rceil}(g_j \mathbf{m}_{\mathcal{I}_k}) \succeq 0, \forall j=1, \dots, l,$

where $p_{j,\alpha^{(j)}}$ is the coefficient of the $\alpha^{(j)}$ -th monomial in the polynomial p_j , $\mathbf{M}_{k,N}(\mathbf{m}_{\mathcal{I}_k})$ denotes the moment matrix and $\mathbf{L}_{j,N-\lceil \frac{\delta_j}{2} \rceil}$ the localizing matrix for the subset of variables in \mathcal{I}_k . Notably, for a given N , this approach requires considering moments and localizing matrices containing $O(\kappa^{2N})$ variables, where κ is the maximum cardinality of \mathcal{I}_k , rather than $O(n^{2N})$. Since in the problems considered in this paper $\kappa \ll n$ this results in substantial computational complexity reduction.

3. Main Results

In this section we present the main theoretical results of the paper: (i) a reformulation of Problem 1 into a constrained polynomial optimization problem, (ii) a convergent sequence of convex relaxation, and (iii) a sufficient condition, given in terms of the rank of a small matrix, for the first element of this sequence to attain global optimality, leading to a computationally efficient algorithm.

3.1. A Constrained POP Reformulation

By introducing indicator variables $s_i \in \{0, 1\}$, Problem 1 can be reformulated as

$$p^* = \min_{\mathbf{q}, \mathbf{F}, s_i} \sum_{i=1}^n (1 - s_i) \text{ subject to:} \quad (6a)$$

$$\|\mathbf{F}\|_F^2 = 1 \quad (6b)$$

$$\mathbf{F}\mathbf{q} = 0, \mathbf{q}^T \mathbf{q} = 1 \quad (6c)$$

$$s_i^2 = s_i, \forall i=1, \dots, n \quad (6d)$$

$$s_i |\mathbf{x}_i'^T \mathbf{F} \mathbf{x}_i| \leq s_i \epsilon \quad (6e)$$

In this formulation (6c) guarantees that \mathbf{F} is rank deficient, while (6d) forces s_i to be binary. This last equation, combined with (6e) enforces that $s_i = 0$ for outliers (that is, points where $|\mathbf{x}_i'^T \mathbf{F} \mathbf{x}_i| > \epsilon$). Thus, the cost function (6a) is meant to minimize the number of outliers.

Remark 1. Note that in this formulation, rank deficiency of \mathbf{F} is enforced by imposing the existence of a nonzero vector \mathbf{q} in its null space, leading to a second order polynomial constraint. For comparison, the approach in [1], imposes $\det(\mathbf{F}) = 0$, which in this case results in a third order polynomial constraint. Thus, the lowest order relaxation in that formulation requires considering moments of order up to four, while in our formulation the lowest order relaxation involves only moments of order up to 2. Since the size of the optimization problem is combinatorial in the order of the relaxation, our approach, combined with theoretical results guaranteeing that this relaxation is exact, leads to a substantial reduction in computational complexity.

Remark 2. It can be easily shown that (6) exhibits the running intersection property. This observation will be used in the next section to efficiently construct a hierarchy of convex relaxations.

Remark 3. Under the assumption that all the correspondences are inliers, then (6) reduces to a feasibility problem about \mathbf{F} and \mathbf{q} with $s_i = 1, \forall i = 1, \dots, n$.

3.2. Moments Based Relaxations

Clearly, (6) is nonconvex due to the bilinear terms and the binary variables in (6b)-(6e). However, the moments based polynomial optimization techniques introduced in Section 2.3, can be used to obtain a sequence of convex relaxations of the form

$$p_N^* = \min_{\mathbf{M}, \mathbf{L}} n - \sum_{i=1}^n m(s_i) \quad (7)$$

s.t. $\mathbf{M}_{i,N} \succeq 0, \forall i=1, \dots, n,$
 $\mathbf{L}_{i,N-1} \succeq 0, \forall i=1, \dots, n,$

where $\mathbf{M}_{i,N}$ denotes the N -th order moment matrix consisting of moments of $\{\mathbf{F}, \mathbf{q}, s_i\}$ (variables associated with $(\mathbf{x}_i, \mathbf{x}_i')$) of order up to $2N$, and where $\mathbf{L}_{i,N-1}$ denotes the localizing matrices corresponding to the constraints (6b)-(6e) that define the feasibility set in (6).

The results from Section 2.3.2 guarantee that this sequence of relaxations converges monotonically (from below) to the optimum. Nevertheless, from a practical stand point its applicability is limited to relatively small problems, due to the fact that the number of variables increases combinatorially with N (even when exploiting the underlying sparse structure), and that, in principle the value of N required to achieve a flat extension, and hence certify optimality, can be large. In addition, once an optimal N has been found, extracting the solution \mathbf{F} from the corresponding moments matrix \mathbf{M} is far from trivial, unless $\text{rank}(\mathbf{M})$ is low (see [7] for details). To circumvent these difficulties, in the sequel

we will exploit the fact that moment matrices associated with atomic measures having a single atom have rank 1, since in this case the moments simply correspond to the powers of the variables, evaluated at the location of the atom, and in turn the variables are equal to their first order moments. In principle, one could try to exploit this observation by adding a low rank constraint on each of the moment matrices $\mathbf{M}_{i,N}$ to (7). However, this constraint would be computationally hard to enforce due to the large number of matrices involved. Surprisingly, as shown by the result below, enforcing a low rank constraint *only* on the moment matrix associated with \mathbf{F} and \mathbf{q} is enough to guarantee that the relaxation corresponding to $N = 1$ attains global optimality. Further, in this case the elements of \mathbf{F} can be read directly from its associated moment matrix, without the need for additional computations.

Theorem 1. *Consider the first order moment relaxation (7) of Problem (6) given by*

$$\begin{aligned} \tilde{p}_1^* = \min \quad & n - \sum_{i=1}^n m(s_i) \\ \text{s.t.} \quad & \mathbf{M}_{i,1} \succeq 0, \forall_{i=1}^n \text{ and } \mathbf{L}_{i,0} \succeq 0, \forall_{i=1}^d \end{aligned} \quad (8)$$

Let $\mathbf{M}_1^c \doteq \begin{bmatrix} 1 & m(\mathbf{f})^T & m(\mathbf{q})^T \\ m(\mathbf{f}) & m(\mathbf{f}\mathbf{f}^T) & m(\mathbf{f}\mathbf{q}^T) \\ m(\mathbf{q}) & m(\mathbf{q}\mathbf{f}^T) & m(\mathbf{q}\mathbf{q}^T) \end{bmatrix}$. Then, if $\text{rank}(\mathbf{M}_1^c) = 1$, the relaxation (8) is exact. Further, in this case the corresponding optimal first order moments sequence $\{m(\mathbf{F})^*, m(\mathbf{q})^*, m(s_i)^*\}$ is also an optimal solution to the original non-convex problem (6).

Proof. (Sketch) For each pair $(\mathbf{x}_i, \mathbf{x}'_i)$, $\forall i = 1, \dots, n$, the first order moment matrix $\mathbf{M}_{i,1}$ corresponding to its associated variables \mathbf{f} , \mathbf{q} and s_i is of the form

$$\begin{aligned} \mathbf{M}_{i,1} &= m \left(\begin{bmatrix} 1 \\ \mathbf{f} \\ \mathbf{q} \\ s_i \end{bmatrix} \begin{bmatrix} 1 & \mathbf{f}^T & \mathbf{q}^T & s_i \end{bmatrix} \right) \\ &= \begin{bmatrix} 1 & m(\mathbf{f})^T & m(\mathbf{q})^T & m(s_i) \\ m(\mathbf{f}) & m(\mathbf{f}\mathbf{f}^T) & m(\mathbf{f}\mathbf{q}^T) & m(\mathbf{f}s_i) \\ m(\mathbf{q}) & m(\mathbf{q}\mathbf{f}^T) & m(\mathbf{q}\mathbf{q}^T) & m(\mathbf{q}s_i) \\ m(s_i) & m(\mathbf{f}^T s_i) & m(\mathbf{q}^T s_i) & m(s_i) \end{bmatrix} \end{aligned} \quad (9)$$

where \mathbf{f} denotes the vectorized \mathbf{F} , and where the last entry of $\mathbf{M}_{i,1}$ results from the constraint $s_i = s_i^2$. Since $\text{rank}(\mathbf{M}_1^c) = 1$, then the first order moment variables $m(\mathbf{f})$ and $m(\mathbf{q})$ satisfy (6b) and (6c). Next we show that they also satisfy (6e).

Let $\mathbf{v} = [v(j)]_{j=1}^2 = [\mathbf{f}^T \quad \mathbf{q}^T]^T$. Due to the positive semi-definiteness of $\mathbf{M}_{i,1}$, its principal minors should also be positive semi-definite. Hence:

$$\mathbf{M}_m = \begin{bmatrix} 1 & m(v(j)) & m(s_i) \\ m(v(j)) & m(v(j)^2) & m(v(j)s_i) \\ m(s_i) & m(v(j)s_i) & m(s_i) \end{bmatrix} \succeq 0 \quad (10)$$

If \mathbf{M}_1^c is rank 1, then $m(v(j)^2) = m^2(v(j))$, furthermore, from

$$\det(\mathbf{M}_m) = -[m(v(j)s_i) - m(v(j))m(s_i)]^2 \geq 0, \quad (11)$$

we have

$$m(v(j)s_i) = m(v(j))m(s_i). \quad (12)$$

At the first order relaxation, the localizing matrix associated with (6e) reduces to:

$$\begin{cases} \mathbf{x}'_i{}^T m(s_i \mathbf{F}) \mathbf{x}_i \leq m(s_i) \epsilon; \\ \mathbf{x}'_i{}^T m(s_i \mathbf{F}) \mathbf{x}_i \geq -m(s_i) \epsilon. \end{cases} \quad (13)$$

Substituting (12) into (13) leads to

$$\begin{cases} \mathbf{x}'_i{}^T m(s_i) m(\mathbf{F}) \mathbf{x}_i \leq m(s_i) \epsilon; \\ \mathbf{x}'_i{}^T m(s_i) m(\mathbf{F}) \mathbf{x}_i \geq -m(s_i) \epsilon. \end{cases} \quad (14)$$

Thus, for any $m(s_i) > 0$, (13) reduces to

$$\begin{cases} \mathbf{x}'_i{}^T m(\mathbf{F}) \mathbf{x}_i \leq \epsilon \\ \mathbf{x}'_i{}^T m(\mathbf{F}) \mathbf{x}_i \geq -\epsilon \end{cases}, \quad (15)$$

implying that the pair $(\mathbf{x}'_i, \mathbf{x}_i)$ is an inlier associated with the fundamental matrix $m(\mathbf{F})$. Next, note that, since (8) seeks to minimize $-\sum_{i=1}^n m(s_i)$, then any non-zero $m(s_i)^*$ term in the optimal solution will automatically increase to 1. Therefore, (8) is exactly equivalent to (6) (by replacing \mathbf{F} , \mathbf{q} , s_i with $m(\mathbf{F})$, $m(\mathbf{q})$, $m(s_i)$). Finally, note that the combination of the binary variables $m(s_i)^*$ and $\text{rank}(\mathbf{M}_1^c) = 1$ guarantees that $\mathbf{M}_{i,1}$ is also rank 1, from where it follows that the first order moment sequence $\{m(\mathbf{F})^*, m(\mathbf{q})^*, m(s_i)^*\}$ is indeed an optimal solution of the original nonconvex problem (6). \square

3.3. A Reweighted Heuristic Based Algorithm

Theorem 1 suggests that a computationally attractive algorithm can be obtained by simply adding the constraint $\text{rank}(\mathbf{M}_c) = 1$ to (8). Unfortunately, the resulting problem is no longer convex. Nevertheless, a tractable convex relaxation can be obtained by using a (weighted) nuclear norm as a surrogate for rank^1 ,

¹This relaxation amounts to a local linearization of the matrix's log-det function.

and iteratively solving a sequence of regularized convex problems[14], of the form:

$$\begin{aligned} \tilde{p}_1^{(k)*} = \min \quad & n - \sum_{i=1}^n m(s_i) + \lambda \text{trace}(\mathbf{W}^{(k)} \mathbf{M}_1^c) \\ \text{s.t.} \quad & \mathbf{M}_{i,1} \succeq 0, \forall_{i=1}^n \text{ and } \mathbf{L}_{i,0} \succeq 0, \forall_{i=1}^d \end{aligned} \quad (16)$$

favoring low rank solutions through the second term in the cost, leading to Algorithm 1 outlined below. Note that each iteration has computational complexity roughly equivalent to that of regularized robust regression and robust low rank factorization based methods. Further, consistent numerical experience shows that typically only a few iterations are needed for convergence.

Algorithm 1 Robust Fundamental Matrix Estimation

Initialize: $k = 0, \mathbf{W}^{(1)} = \mathbf{I}$;

repeat

$k = k + 1$;

Solve (16) and obtain $\mathbf{M}_1^{c*(k)}$;

Update $\mathbf{W}^{(k+1)} = [\mathbf{M}_1^{c*(k)} + \sigma_2(\mathbf{M}_1^{c*(k)})\mathbf{I}]^{-1}$

until $\sigma_2(\mathbf{M}_1^{c*(k)}) < 10^{-6}\sigma_1(\mathbf{M}_1^{c*(k)})$.

4. Extensions: Co-occurrences and Partially Known Correspondences

A salient feature of the proposed framework is that it can be easily extended to handle co-occurrences and partially known correspondences. The former simply entails using the same variable s_i for correspondences known to have the same label (either all are outliers or all are inliers, although the actual label is unknown). Partially known correspondences arise in cases where labeling is expensive, and thus only a few labels are available. Consider first a scenario without outliers. In this case, the problem can be formalized as: given a set of n labeled point correspondences $(\mathbf{y}_i, \mathbf{y}'_i), i = 1, \dots, n$, m unlabeled points $\{\mathbf{x}_j\}_{j=1}^m$ from image 1 and m unlabeled points $\{\mathbf{x}'_j\}_{j=1}^m$ from image 2 of the same scene, find the fundamental matrix \mathbf{F} and a set of m pairs $(\mathbf{x}'_i, \mathbf{x}_j)$ such that $|\mathbf{x}'_i \mathbf{F} \mathbf{x}_j| \leq \epsilon$. By introducing suitable variables $l_{i,j}$ the problem can be reduced to the following polynomial feasibility form:

$$\begin{aligned} l_{i,j} |\mathbf{x}'_i \mathbf{F} \mathbf{x}_j| &\leq l_{i,j} \epsilon, \forall_{i=1}^m \forall_{j=1}^m \\ |\mathbf{y}'_i \mathbf{F} \mathbf{y}_i| &\leq \epsilon, \forall_{i=1}^n \\ \sum_{i=1}^m l_{i,j} &= 1, \forall_{j=1}^m \\ \sum_{j=1}^m l_{i,j} &= 1, \forall_{i=1}^m \end{aligned} \quad (17)$$

where the last two equations in (17) enforce the fact that each unlabeled point must be assigned to exactly one

correspondence. Finally, the case where the set of labeled data contains outliers can be handled as before, by introducing variables s_i and minimizing $n - \sum s_i$.




Figure 1. Sample image pairs from the VGG dataset.

5. Experiments

Next, we describe a set of experiments used to verify our theoretical results. These experiments consist of the estimation of fundamental matrices from points correspondences from six pairs of stereo images: *House*, *Merton I*, *Merton II*, *Merton III*, *Library*, and *Wadham*, given by VGG, University of Oxford, shown in Figure 1.

5.1. Experimental Protocol

In each experiment, before computing the fundamental matrix, the data was normalized as in [6] and [3].

Ground truth data. Given a pair of images, we first calculated a fundamental matrix from the correspondences $(\mathbf{x}_{i,t}, \mathbf{x}'_{i,t})$ provided with the dataset by minimizing the algebraic error. We considered the resulting fundamental matrix \mathbf{F}_t , as the ground truth, with the threshold ϵ given by $\epsilon \doteq \max_i |\mathbf{x}'_{i,t} \mathbf{F}_t \mathbf{x}_{i,t}|$. Then, the VLFeat toolbox [19] was used to obtain a set of SIFT features from the two images, and correspondences $(\mathbf{x}_i, \mathbf{x}'_i)$ were defined by those pairs of points whose ℓ_2 norm is ≤ 0.5 . Finally, inliers and outliers were selected as those correspondences satisfying $\max_i |\mathbf{x}'_i \mathbf{F}_t \mathbf{x}_i| \leq \epsilon$, and those violating this bound, respectively.

Experimental set-up. For each pair of images, we ran seven sets of experiments, with 100 correspondences and outliers ranging from 10% to 70%. For each set of experiments we ran 50 times by randomly picking N_{out} outliers and N_{in} inliers from $(\mathbf{x}_i, \mathbf{x}'_i)$ and compared the results with several state-of-the-art methods.

Evaluation criteria. We compared four performance measurements as follows:

- $Precision = \frac{\text{Ground truth inliers} \cap \text{Identified inliers}}{\text{Identified inliers}} \times 100$;
- $Recall = \frac{\text{Ground truth inliers} \cap \text{Identified inliers}}{\text{Ground truth inliers}} \times 100$;
- $Hmeans = \sqrt{Precision \times Recall} \times 100$;
- $Similarity = |\text{trace}(\mathbf{F}_t \mathbf{F}^T)|$. The closer $Similarity$ is to 1, the smaller the distance between the identified fundamental matrix \mathbf{F} and the ground truth \mathbf{F}_t .

Computational platform. All algorithms were implemented in MATLAB and run on a 3.4GHz iMac with 32G memory. The SDP solver used was SeDuMi [15].

5.2. Results

In these experiments we compared the performance of the proposed algorithm with that of RANSAC and several of its variants, i.e., MSAC, MLESAC and LMEDS. For all these methods, the number of iterations was set to 500, and in each iteration the fundamental matrix was calculated using the eight-point algorithm. The results are summarized in Tables 1-4, and illustrated in Figure 2. As shown there, the proposed algorithm was consistently more robust than the SAC algorithms, in the sense that it identified a larger number of inliers and yielded a fundamental matrix closer to the ground truth, in terms of a larger inner product. Note that while in Tables 1-3 the proposed method leads to larger variance, the corresponding mean value is higher, indicating that the other methods have consistently lower performance. Indeed, a quick computation assuming gaussian distributions shows that our method yields a higher objective value with probability > 0.7 even in the most unfavorable case. In terms of computational efficiency, since our algorithm requires solving semi-definite optimization problems, each iteration is more time consuming than those of SACs. On the other hand, consistent numerical experience shows that only a few iterations are needed for convergence (typically no more than 14). Since each iteration requires about 10 seconds when using 100 points, the overall computational cost remains competitive vis-à-vis randomized methods, specially in cases involving large number of outliers.

Table 1. Mean and Standard Deviation of Precision (%)

N_{out}	RANSAC	MSAC	LMEDS	MLESAC	Proposed
10	97.7586	97.6864	97.8189	97.9042	98.2640
	1.7967	1.7241	1.9095	1.7129	1.9726
20	94.7582	94.6092	94.5228	94.6768	96.6838
	3.0774	2.3474	3.0246	3.0948	2.8109
30	90.3102	90.6371	90.1515	90.6361	93.5078
	5.1555	3.8062	4.8416	5.0145	4.6797
40	83.8319	83.6164	83.9802	84.0462	90.1450
	6.8161	5.6156	7.3236	6.6029	6.5219
50	74.9536	74.7678	74.7990	75.8945	83.8643
	10.7474	6.6748	10.0023	9.5612	8.7383
60	58.3182	59.4671	60.3656	60.9045	75.8507
	12.7762	9.1357	12.4958	10.9836	14.0279
70	43.5151	42.9147	44.8255	44.0817	62.6045
	13.0902	12.2906	12.9395	13.3457	16.8592

Table 2. Mean and Standard Deviation of Recall (%)

N_{out}	RANSAC	MSAC	LMEDS	MLESAC	Proposed
10	75.6074	71.9333	69.1296	72.9741	89.2111
	4.8220	4.5412	5.1152	6.2715	7.8245
20	74.1917	70.8500	70.6417	71.6167	85.5804
	5.5170	4.7135	5.3528	6.5717	8.9735
30	72.1571	70.0000	70.6714	69.9000	80.1794
	6.3000	5.1970	6.1241	7.3354	10.3759
40	69.4556	67.1944	68.7500	67.6278	77.3000
	7.7685	7.7287	8.4810	7.8525	10.9068
50	64.9600	62.4600	62.3733	63.4733	71.9000
	10.5049	9.8250	10.8597	9.6938	11.1134
60	57.1500	56.9583	55.7667	57.3917	66.7817
	13.4839	10.6995	13.8929	10.8910	15.0884
70	50.2222	47.5000	45.7000	48.8556	62.0917
	14.3284	16.3911	14.0125	14.7437	18.6231

Table 3. Mean and Standard Deviation of Hmean (%)

N_{out}	RANSAC	MSAC	LMEDS	MLESAC	Proposed
10	85.9271	83.7512	82.1771	84.4428	93.5321
	2.9264	2.8941	3.2714	3.7745	4.7464
20	83.7898	81.7846	81.6613	82.2578	90.8371
	3.6833	2.7879	3.7210	4.2533	5.7553
30	80.6565	79.5672	79.7593	79.4978	86.4264
	4.9393	3.8809	4.8613	5.4053	7.2137
40	76.2084	74.8478	75.8708	75.2759	83.2783
	6.3890	6.1565	7.0401	6.1798	7.8551
50	69.6849	68.2157	68.1352	69.2968	77.4227
	10.0422	7.4362	9.4998	8.8621	8.5625
60	57.6164	58.0626	57.7907	59.0019	70.8664
	12.6155	9.4320	12.2474	10.2399	13.3851
70	46.6358	45.0231	45.0193	46.2706	62.0409
	13.3631	14.0348	12.6929	13.5550	16.7669

Table 4. Mean and Standard Deviation of Similarity ($\times 100$)

N_{out}	RANSAC	MSAC	LMEDS	MLESAC	Proposed
10	99.7772	99.6793	99.4775	99.9650	99.9828
	0.7316	0.0126	2.2484	0.0908	0.0716
20	99.6139	99.1964	99.2437	99.9289	99.9860
	1.7647	0.0129	3.3507	0.3425	0.0697
30	99.4306	99.2144	99.6560	99.8947	99.9920
	2.6753	0.0207	1.4154	0.3682	0.0176
40	99.6427	99.0665	99.5319	99.8482	99.9918
	1.0935	0.0189	2.2583	0.3942	0.0187
50	98.5472	96.9745	99.4611	99.6875	99.9876
	5.1293	0.0331	3.4161	1.3350	0.0327
60	97.6169	94.9443	99.0108	98.8553	99.9878
	8.7403	0.0570	5.0077	5.3757	0.0289
70	96.6981	94.5070	98.7259	98.5653	99.9931
	9.0199	0.2751	5.6968	5.0236	0.0078




Figure 2. Performance comparison (Hmean and Similarity) for the image pairs *House*, *Library* and *Merton II*. The rest of the experiments using VGG data show similar behavior.

6. Conclusions

This paper considers the problem of robustly estimating the fundamental matrix from point correspondences corrupted by noise and outliers. Its main result is a computationally tractable algorithm that explicitly enforces the rank deficiency constraint, without assumptions on the percentage or distribution of the outliers. In addition, the algorithm can be easily extended to handle co-occurrence information or scenarios where only some

of the correspondences are known. These results are illustrated with experiments, showing that the proposed method performs well, even in scenarios characterized by a very high percentage of outliers, consistently outperforming existing techniques. Research is currently under way seeking implementations based on first order methods (as opposed to the interior point methods used by conventional SDP solvers), in order to further reduce the computational burden.

References

- [1] F. Bugarin, A. Bartoli, D. Henrion, J.-B. Lasserre, J.-J. Orteu, and T. Sentenac. Rank-constrained fundamental matrix estimation by polynomial global optimization versus the eight-point algorithm. *arXiv preprint arXiv:1403.4806*, 2014. 1, 4
- [2] G. Chesi, A. Garulli, A. Vicino, and R. Cipolla. Estimating the fundamental matrix via constrained least-squares: A convex approach. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(3):397–401, 2002. 1
- [3] W. Chojnacki, M. J. Brooks, A. V. D. Hengel, and D. Gawley. Revisiting hartley’s normalized eight-point algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25:1172–1177, 2003. 1, 6
- [4] R. E. Curto and L. A. Fialkow. Truncated k-moment problems in several variables. *arXiv preprint math/0507067*, 2005. 3
- [5] M. A. Fischler and R. 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, 1981. 2
- [6] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 1, 6
- [7] D. Henrion, J.-B. Lasserre, and J. Löfberg. Gloptipoly 3: moments, optimization and semidefinite programming. *Optimization Methods & Software*, 24(4-5):761–779, 2009. 4
- [8] P. J. Huber. Projection pursuit. *The annals of Statistics*, pages 435–475, 1985. 2
- [9] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001. 3
- [10] J. B. Lasserre. Convergent sdp-relaxations in polynomial optimization with sparsity. *SIAM Journal on Optimization*, 17(3):822–843, 2006. 3, 4
- [11] A. M. Leroy and P. J. Rousseeuw. Robust regression and outlier detection. *Wiley Series in Probability and Mathematical Statistics, New York: Wiley*, 1987, 1, 1987. 2
- [12] H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Readings in Computer Vision: Issues, Problems, Principles, and Paradigms, MA Fischler and O. Firschein, eds*, pages 61–62, 1987. 1
- [13] T. Migita and T. Shakunaga. Evaluation of epipole estimation methods with/without rank-2 constraint across algebraic/geometric error functions. In *Computer Vision and Pattern Recognition, 2007. CVPR’07. IEEE Conference on*, pages 1–7. IEEE, 2007. 1
- [14] K. Mohan and M. Fazel. Iterative reweighted algorithms for matrix rank minimization. *J. of Machine Learning Research*, 13:3441–3473, 2012. 6
- [15] J. F. Sturm. Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. *Optimization methods and software*, 11(1-4):625–653, 1999. 7
- [16] Y. Sugaya and K. Kanatani. High accuracy computation of rank-constrained fundamental matrix. In *BMVC*, pages 1–10, 2007. 1
- [17] P. H. Torr and D. W. Murray. The development and comparison of robust methods for estimating the fundamental matrix. *International journal of computer vision*, 24(3):271–300, 1997. 2
- [18] P. H. Torr and A. Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78(1):138–156, 2000. 2
- [19] A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms. <http://www.vlfeat.org/>, 2008. 6
- [20] Y. Zheng, S. Sugimoto, and M. Okutomi. A branch and contract algorithm for globally optimal fundamental matrix estimation. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2953–2960. IEEE, 2011. 1
- [21] Y. Zheng, S. Sugimoto, and M. Okutomi. A practical rank-constrained eight-point algorithm for fundamental matrix estimation. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 1546–1553. IEEE, 2013. 1