



Non-Linear Estimation of the Fundamental Matrix With Minimal Parameters

Adrien Bartoli, Peter Sturm

► To cite this version:

Adrien Bartoli, Peter Sturm. Non-Linear Estimation of the Fundamental Matrix With Minimal Parameters. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2004, 26, pp.426-432. hal-00092887

HAL Id: hal-00092887

<https://hal.science/hal-00092887>

Submitted on 12 Sep 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Non-Linear Estimation of the Fundamental Matrix With Minimal Parameters

Adrien Bartoli and Peter Sturm

Affiliation: INRIA

Address:

INRIA Rhône-Alpes

655, avenue de l'Europe

38334 Saint Ismier cedex, France

first.last@inria.fr

September 12, 2006

The purpose of this paper is to give a very simple method for non-linearly estimating the fundamental matrix using the minimum number of 7 parameters. Instead of minimally parameterizing it, we rather update what we call its orthonormal representation, which is based on its singular value decomposition. We show how this method can be used for efficient bundle adjustment of point features seen in two views. Experiments on simulated and real data show that this implementation performs better than others in terms of computational cost, i.e. convergence is faster, although methods based on minimal parameters are more likely to fall into local minima than methods based on redundant parameters.

Index Terms: Structure-from-motion, Bundle Adjustment, Minimal Parameterization, Fundamental Matrix.

1 Introduction

The fundamental matrix has received a great interest in the computer vision community, see e.g. [5, 6, 11, 12, 20, 23, 24]. It encapsulates the epipolar geometry or the projective motion between two uncalibrated perspective cameras and can be used for 3D reconstruction, motion segmentation, self-calibration and so forth. Accurately estimating the fundamental matrix is therefore a major research issue. Most of the time, point correspondences between the two images are used. A linear solution is obtained using the 8 point algorithm [5, 11], optionally embedded in a robust estimation scheme [20, 23]. This estimate is then non-linearly refined by minimizing a physically meaningful criterion that may involve reconstructed 3D point coordinates as well (in particular for bundle adjustment). However, non-linearly estimating the fundamental matrix suffers from the lack of a simple technique to represent it efficiently. This paper, which is an extension of [2], provides such a technique in §3, based on the orthonormal representation of the fundamental matrix that we introduce. We show in §4 how this method can be used to refine the fundamental matrix by bundle adjustment of point features. We demonstrate experimentally in §§5.1 and 5.2 that the resulting algorithm performs better than existing ones in terms of computational cost.

2 Notations and Relation to Previous Work

The fundamental matrix denoted as F is an homogeneous (i.e. defined up to scale) (3×3) rank-2 matrix. It has therefore 9 entries but only 7 degrees of freedom.

There have been many attempts to minimally parameterize it, i.e. to represent it with 7 parameters. Most of the previous works deal with directly parameterizing the epipolar geometry. The fundamental matrix F is decomposed into the epipoles e and e' and the epipolar transformation, which is a 1D projective transformation relating the epipolar pencils, represented by an homogeneous (2×2) -matrix g [4, 12, 23].

Representing these entities with minimal parameters requires eliminating their arbitrary scale factors. This can be done by fixing e.g. the 2-norm of homogeneous entities, but then the parameterization would not be minimal. Another solution is to freeze one entry of each homogeneous entity (in practice, the largest entry), which yields 3 possibilities for each epipole and 4 for the epipolar transformation, so $3 \cdot 3 \cdot 4 = 36$ possible parameterizations.

In [12], the authors propose to restrict the two-view configurations considered to the cases where both epipoles are finite and can therefore be expressed in affine coordinates. Consequently, this parameterization can be used only when both epipoles do not lie at infinity. Due to the homogeneity of the epipolar transformation, 4 distinct parameterizations are still necessary for g . A total of 4 parameterizations are then needed to represent this restricted set of fundamental matrices.

The method has been extended in [23] to the general case, i.e. when the epipoles can be either finite or infinite. In this case, it is shown that all 36 distinct parameterizations are necessary. This leads to a cumbersome and error-prone implementation of the optimization process.

Note that there are 9 different possibilities to form the fundamental matrix — or any other 2D entity such as the extended epipolar transformation [4] or the canonic plane homography H^* [13] — from e , e' and g [23].

In [4, 24], the method has been revised so as to reduce the number of parameterizations using image transformations. In [4], the image transformations used are metric and the number of distinct parameterizations is restricted to 3 plus one bilinear constraint on the entries of g , while in [24], the transformations used are projective, which allows one to reduce the number of parameterizations to 1. The main drawback is that in the transformed image space, the original noise model on the image features is not preserved. A means to

preserve it, up to first order approximation, has been proposed in [24] for the gradient-weighted criterion, which is not the one used for bundle adjustment.

Another solution is the point-based parameterization of [19]. The idea is to represent the fundamental matrix by a set of 7 point correspondences. Minimal optimization can then be conducted by varying one coordinate for each point correspondence. The fundamental matrix is obtained at each minimization step by computing the standard 7 point solution, which means that the null-space of a (7×9) matrix has to be computed and a cubic equation has to be solved. There may be up to 3 solutions. The one giving the lowest residual error is kept. The disadvantage of this parameterization is that it is costly to obtain the fundamental matrix given its parameters (i.e. the 7 point correspondences). Also, analytic differentiation is not possible.

3 Non-Linear Optimization With 7 Parameters

In contrast to the existing work, we do not try to represent the entire set of fundamental matrices using 7 parameters. We rather locally update it with 7 parameters. Before going further, we illustrate this idea by considering the case of the non-linear estimation of 3D rotations, which is simpler and, as will be seen later, has similarities with the case of the fundamental matrix.

3.1 The Case of 3D Rotations

There exist many representations of 3D rotations, see e.g. [18], including Euler angles, the Gibbs vector, Cayley-Klein parameters, Pauli spin matrices, axis-and-angle systems, $SO(3)$ matrices¹ and unit quaternions. None of these representations is able to uniquely represent all 3D rotations with the minimum 3 parameters. For that reason, the following scheme is often used for their non-linear estimation, see e.g. [1, 7, 21]. The rotation is represented by an $SO(3)$ matrix R and is locally updated using 3 parameters by any well-behaved (locally non-singular) representation, such as 3 Euler angles $\boldsymbol{\theta}^T = (\theta_1 \ \theta_2 \ \theta_3)$ as:

$$R \leftarrow R R(\boldsymbol{\theta}), \quad (1)$$

where $R(\boldsymbol{\theta}) = R_x(\theta_1) R_y(\theta_2) R_z(\theta_3)$ is the $SO(3)$ matrix representation of the 3D rotation corresponding to $\boldsymbol{\theta}$ with $R_x(\theta_1) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_1 & -\sin \theta_1 \\ 0 & \sin \theta_1 & \cos \theta_1 \end{pmatrix}$, $R_y(\theta_2) = \begin{pmatrix} \cos \theta_2 & 0 & \sin \theta_2 \\ 0 & 1 & 0 \\ -\sin \theta_2 & 0 & \cos \theta_2 \end{pmatrix}$ and

¹ $SO(3)$ is the Lie group of (3×3) matrices R satisfying $R^T R = I$ and $\det R = 1$.

$R_z(\theta_3) = \begin{pmatrix} \cos \theta_3 & -\sin \theta_3 & 0 \\ \sin \theta_3 & \cos \theta_3 & 0 \\ 0 & 0 & 1 \end{pmatrix}$. At the end of each iteration, R is updated and θ is reset to zero. Hence, at each iteration, the estimated Euler angles are small (initialized as zero), which makes this representation non-singular.

3.2 Minimal Update

Following the example of 3D rotations, we propose the *orthonormal representation* of the fundamental matrix where more parameters than degrees of freedom are needed, but that can be easily updated using the minimum 7 parameters.

Given an estimate of the fundamental matrix F obtained using e.g. the 8 point algorithm, consider its singular value decomposition $F \sim U \Sigma V^T$, where U and V are $O(3)$ matrices² and Σ a diagonal one containing the singular values of F . Since F has rank 2, $\Sigma \sim \text{diag}(\sigma_1, \sigma_2, 0)$, where $\sigma_1 \geq \sigma_2 > 0$ [22]. We can scale Σ such that $F \sim U \text{diag}(1, \sigma, 0) V^T$ where $\sigma = \sigma_2/\sigma_1$ ($\sigma_1 \neq 0$ since $F \neq 0$) and $1 \geq \sigma > 0$.

This decomposition shows that any fundamental matrix can be represented by (U, V, σ) , i.e. two $O(3)$ matrices and a scalar, which form what we call its orthonormal representation. Note that in the case $\sigma = 1$, i.e. when the fundamental matrix is an essential matrix [8], the orthonormal representation is not unique (see below).

The orthonormal representation is consistent in that it yields $3 + 3 + 1 = 7$ degrees of freedom. The fundamental matrix can be recovered as:

$$F \sim \mathbf{u}_1 \mathbf{v}_1^T + \sigma \mathbf{u}_2 \mathbf{v}_2^T, \quad (2)$$

where \mathbf{u}_i and \mathbf{v}_i are the columns of U and V respectively.

This representation suggests the following update scheme. Each $O(3)$ matrix can be updated using an $SO(3)$ matrix, using equation (1) as in the case of 3D rotations, while σ can be included as such into the optimization:

$$U \leftarrow U R(\mathbf{x}) \quad V \leftarrow V R(\mathbf{y}) \quad \sigma \leftarrow \sigma + \delta_\sigma. \quad (3)$$

Here, \mathbf{x} and \mathbf{y} are 3-vectors of Euler angles. Intuitively, the orthonormal representation should be intrinsically well-conditioned since U and V are $O(3)$ matrices.

Completeness. A first question that immediately follows about the above-proposed method is whether all two-view configurations are covered. Clearly, any fundamental matrix can

² $O(3)$ is the Lie group of (3×3) matrices R satisfying $R^T R = I$.

be decomposed into two $O(3)$ matrices and a scalar. The question arises from the fact that U and V are $O(3)$ matrices, which may have positive or negative determinants, and are updated using $SO(3)$ matrices, $R(x)$ and $R(y)$ respectively, which have positive determinants. Actually, this is not a problem since the signs of U and V can be freely switched, which accordingly switches the signs of their determinants, while leaving the corresponding F invariant: $F \sim (\pm U)\Sigma(\pm V)^T$.

Ensuring bounds on σ . A second remark is about the bounds on σ : $0 < \sigma \leq 1$. There are several possibilities to ensure them while leaving the corresponding F invariant. However, we have found during our experiments that in practice, this does not affect the behaviour of the underlying optimization process.

Essential matrices. As pointed out previously, in the case of $\sigma = 1$, where the fundamental matrix considered is an essential matrix, the proposed orthonormal representation is not unique: if U and V represent F , then also $U R_z(\alpha)$ and $V R_z(\alpha)$ for any α . This induces that the Jacobian matrix (6) has rank 6, as shown in §4.3. We propose two ways to deal with this singularity.

First, one can use a non-linear optimization technique that handles singular parameterizations, e.g. damped Newton-type techniques. Using Levenberg-Marquardt, we found in our experiments that the singularity does not induce numerical instabilities.

Second, one can avoid singular configurations by properly normalizing the image points. Indeed, an essential matrix arises usually from a semi-calibrated configuration where the origin of the coordinate frame in the image lies close to the principal point and where the image coordinates have been scaled by approximately the inverse focal length. In practice, the principal point position is unknown, but it is likely to be close to the image center. Hence, singular configurations can be avoided by translating the origin of the coordinate frame off the image center.

4 Bundle Adjustment

In this section, we show how the orthonormal representation can be used for bundle adjustment of point features $\mathbf{q}_i \leftrightarrow \mathbf{q}'_i$, $i \in 1 \dots m$ seen in two views, through the minimization of the reprojection error. Similar results can be derived for other criteria, such as the minimization of the distances between points and epipolar lines or the gradient-weighted

criterion [12, 23]. However, in order to obtain the maximum likelihood estimate of the fundamental matrix, one has also to estimate corrected point positions $\hat{\mathbf{q}}_i \leftrightarrow \hat{\mathbf{q}}'_i$, i.e. which satisfy exactly the epipolar geometry and therefore, correspond to 3D points \mathbf{Q}_i .

Bundle adjustment consists in minimizing a cost function described in §4.1 over structure and motion parameters. In projective space, there are 15 inherent degrees of gauge freedom, due to the coordinate-frame ambiguity. In [9], a general framework consisting in incorporating gauge constraints up to first order in numerical estimation is introduced. The method of [15] falls in that category. Another technique is to let the gauge free to drift, sometimes partially, while it is ensured that it does not move too far at each iteration. These methods are compared to ours in §5.

When the motion is represented by the fundamental matrix, the gauge is completely eliminated. We call any pair of camera matrices \mathbf{P} and \mathbf{P}' a *realization*. In §4.2, we give analytical formulae to compute a realization from the orthonormal representation of \mathbf{F} (as opposed to [12, 19, 23, 24]).

The algorithm is summarized in table 1.

4.1 Cost Function

Bundle adjustment consists in solving the following optimization problem, see e.g. [15, 21, 23]: $\min_{\mathbf{a}, \mathbf{b}} \sum_j r_j^2$, where \mathbf{a} and \mathbf{b} are respectively motion and structure parameters (or parameters used to update them), \mathbf{r} is the $4m$ -vector of residual errors defined by:

$$\mathbf{r}_{(4m \times 1)}^T = (\dots \quad \mathbf{q}_{i1} - \hat{\mathbf{q}}_{i1} \quad \mathbf{q}_{i2} - \hat{\mathbf{q}}_{i2} \quad \mathbf{q}'_{i1} - \hat{\mathbf{q}}'_{i1} \quad \mathbf{q}'_{i2} - \hat{\mathbf{q}}'_{i2} \quad \dots),$$

where $\hat{\mathbf{q}}_i \sim \mathbf{P}\mathbf{Q}_i$ and $\hat{\mathbf{q}}'_i \sim \mathbf{P}'\mathbf{Q}_i$ are predicted image points.

4.2 Computing a Realization

Due to the projective frame ambiguity, there exists a 15-parameter family of realizations for a given fundamental matrix. A common choice is the canonic projection matrices given by [13]:

$$\mathbf{P} \sim (\mathbf{I}_{(3 \times 3)} \quad \mathbf{0}_{(3 \times 1)}) \text{ and } \mathbf{P}' \sim (\mathbf{H}^* \quad \gamma \mathbf{e}'), \quad (4)$$

where \mathbf{e}' is the second epipole, given by the left null-vector of \mathbf{F} , $\mathbf{F}^T \mathbf{e}' \sim \mathbf{0}_{(3 \times 1)}$ and $\mathbf{H}^* \sim [\mathbf{e}']_{\times} \mathbf{F}$ is the canonic plane homography [13]. The arbitrary scalar γ fixes the relative scale between \mathbf{H}^* and \mathbf{e}' . Without loss of generality, we assume that $\gamma = \|\mathbf{e}'\| = 1$. Any

Two-view projective bundle adjustment expressed within the framework of [7, p.574] (algorithm A4.1). The initial guess of the fundamental matrix is F_0 .

Add the following steps:

- (i') Initialize the orthonormal representation (U, V, σ) by a scaled singular value decomposition of F_0 : $F_0 \sim U \text{diag}(1, \sigma, 0) V^T$.
- (ii') Turn the full $(r \times 12)$ camera Jacobian matrix $A = \bar{A}$ into the minimal $(r \times 7)$ Jacobian matrix of the orthonormal representation: $A \leftarrow A A^\perp$, where A^\perp is given by equations (6,7).

Change the parameter update step as:

- (viii) Update the orthonormal representation as:

$$U \leftarrow U R(\mathbf{x}) \quad V \leftarrow V R(\mathbf{y}) \quad \sigma \leftarrow \sigma + \delta_\sigma,$$

where $\delta_a^T = (\mathbf{x}^T \mathbf{y}^T \delta_\sigma)$ are the 7 motion update parameters, update the structure parameters by adding the incremental vector δ_b and compute the new residual vector.

Table 1: Implementing our minimal estimator within the bundle adjustment Levenberg-Marquardt-based framework given in [7, p.574] (algorithm A4.1). Note that r is the number of residuals and that the second projection matrix has to be extracted from the orthonormal representation using equation (5) (for e.g. computing the residual vector).

other realization can then be obtained by post-multiplying P and P' by a non-singular 3D homography.

Computing the canonic projection matrices (4) can be achieved directly from the orthonormal representation of F . The second epipole is the last column of U : $\mathbf{e}' \sim \mathbf{u}_3$ ($\|\mathbf{u}_3\| = 1$), so the canonic plane homography can be formulated as:

$$H^* \sim [\mathbf{e}']_\times F \sim [\mathbf{u}_3]_\times (\mathbf{u}_1 \mathbf{v}_1^T + \sigma \mathbf{u}_2 \mathbf{v}_2^T).$$

Since U is an $O(3)$ matrix, $[\mathbf{u}_3]_\times \mathbf{u}_1 = \pm \mathbf{u}_2$ and $[\mathbf{u}_3]_\times \mathbf{u}_2 = \mp \mathbf{u}_1$ which yields $H^* \sim \mathbf{u}_2 \mathbf{v}_1^T - \sigma \mathbf{u}_1 \mathbf{v}_2^T$, and thus the particularly simple and direct form of the second projection

matrix:

$$\mathbf{P}' \sim \left(\mathbf{u}_2 \mathbf{v}_1^\top - \sigma \mathbf{u}_1 \mathbf{v}_2^\top \mid \mathbf{u}_3 \right). \quad (5)$$

4.3 Analytical Differentiation

Many non-linear optimization methods necessitate computing the Jacobian matrix $\mathbf{J} = (\mathbf{A} \mid \mathbf{B})$ of the residual vector \mathbf{r} with respect to motion and structure parameters \mathbf{a} and \mathbf{b} . While this can be achieved numerically using e.g. finite differences [16], it may be better to use an analytical form for both computational efficiency and numerical accuracy. We focus on the computation of $\mathbf{A} = \frac{\partial \mathbf{r}}{\partial \mathbf{a}}$ since $\mathbf{B} = \frac{\partial \mathbf{r}}{\partial \mathbf{b}}$ only depends upon structure parameterization. Let $\mathbf{p}' = \text{vect}(\mathbf{P}')$, where $\text{vect}(\cdot)$ is the row-wise vectorization). We decompose $\mathbf{A}_{(4m \times 7)} = \frac{\partial \mathbf{r}}{\partial \mathbf{p}'} \frac{\partial \mathbf{p}'}{\partial \mathbf{a}} = \bar{\mathbf{A}}_{(4m \times 12)} \mathbf{A}_{(12 \times 7)}^\perp$. Only the 12 entries of \mathbf{P}' are considered since \mathbf{P} is fixed in the canonic reconstruction basis (4). The matrix $\bar{\mathbf{A}} = \frac{\partial \mathbf{r}}{\partial \mathbf{p}'}$ depends on the chosen realization of the fundamental matrix, i.e. on the coordinate frame employed. We have chosen the canonic projection matrices (4). This Jacobian matrix is employed directly for the overparameterization proposed in [6]. Deriving its analytical form is straightforward. We therefore concentrate on deriving a closed-form expression for \mathbf{A}^\perp . One of the advantages of the update rule (3) is that there exists a simple closed-form expression for \mathbf{A}^\perp . Non-linear least squares with analytical differentiation can be applied based on \mathbf{A}^\perp .

Let us consider the orthonormal representation $(\mathbf{U}, \mathbf{V}, \sigma)$. The motion update parameters are minimal and defined by $\mathbf{a}^\top = (x_1 \ x_2 \ x_3 \ y_1 \ y_2 \ y_3 \ \sigma)$, where $\mathbf{x}^\top = (x_1 \ x_2 \ x_3)$ and $\mathbf{y}^\top = (y_1 \ y_2 \ y_3)$ are used to update \mathbf{U} and \mathbf{V} respectively. Since \mathbf{U} and \mathbf{V} are updated with respect to the current estimate, \mathbf{A}^\perp is evaluated at $(\mathbf{U}, \mathbf{V}, \sigma)$, i.e. at $\mathbf{a}^\top = \mathbf{a}_0^\top = (\mathbf{0}_{(6 \times 1)}^\top \ \sigma)$. Equation (5) is used to derive a closed-form expression of the second canonic projection matrix after updating. By expanding, differentiating and evaluating this expression at \mathbf{a}_0 , we obtain:

$$\mathbf{A}^\perp = \frac{\partial \mathbf{p}'}{\partial \mathbf{a}} = \left(\left(\frac{\partial \mathbf{p}'}{\partial x_1} \right) \cdots \left(\frac{\partial \mathbf{p}'}{\partial y_3} \right) \left(\frac{\partial \mathbf{p}'}{\partial \sigma} \right) \right), \quad (6)$$

where:

$$\begin{aligned}
\partial \mathbf{p}' / \partial x_1 &= \text{vect}(\mathbf{u}_3 \mathbf{v}_1^\top \mid -\mathbf{u}_2) \\
\partial \mathbf{p}' / \partial x_2 &= \text{vect}(\sigma \mathbf{u}_3 \mathbf{v}_2^\top \mid \mathbf{u}_1) \\
\partial \mathbf{p}' / \partial x_3 &= \text{vect}(-\mathbf{u}_1 \mathbf{v}_1^\top - \sigma \mathbf{u}_2 \mathbf{v}_2^\top \mid \mathbf{0}_{3 \times 1}) \\
\partial \mathbf{p}' / \partial y_1 &= \text{vect}(-\sigma \mathbf{u}_1 \mathbf{v}_3^\top \mid \mathbf{0}_{3 \times 1}) \\
\partial \mathbf{p}' / \partial y_2 &= \text{vect}(-\mathbf{u}_2 \mathbf{v}_3^\top \mid \mathbf{0}_{3 \times 1}) \\
\partial \mathbf{p}' / \partial y_3 &= \text{vect}(\mathbf{u}_2 \mathbf{v}_2^\top + \sigma \mathbf{u}_1 \mathbf{v}_1^\top \mid \mathbf{0}_{3 \times 1}) \\
\partial \mathbf{p}' / \partial \sigma &= \text{vect}(-\mathbf{u}_1 \mathbf{v}_2^\top \mid \mathbf{0}_{3 \times 1}).
\end{aligned} \tag{7}$$

In the general case, $\text{rank}(\mathbf{A}^\perp) = 7$ but when $\sigma = 1$, $\text{rank}(\mathbf{A}^\perp) = 6$ since $\partial \mathbf{p}' / \partial x_3 + \partial \mathbf{p}' / \partial y_3 = 0$.

If the minimal method of e.g. [23] were used, 36 different Jacobian matrices, one for each parameterization, would have to be derived.

4.4 Particular Configurations

The epipolar geometry can be decomposed as a pair of epipoles and the 3-degrees of freedom epipolar transformation [12, 23]. If one or two of these components are a priori known, it may be convenient to leave them invariant during optimization of the fundamental matrix. Such features are easily added to our estimation method, as follows.

Leaving an epipole invariant. Consider e.g. the second epipole, encapsulated in the orthonormal representation as the third column of \mathbf{U} . The update $\mathbf{U} \leftarrow \mathbf{U} \mathbf{R}(\mathbf{x})$ does not affect \mathbf{u}_3 if $x_1 = x_2 = 0$. Therefore, freezing the left or the right epipole can be done by removing x_1, x_2 or y_1, y_2 respectively from the estimation and updating as $\mathbf{U} \leftarrow \mathbf{U} \mathbf{R}_z(x_3)$ or $\mathbf{V} \leftarrow \mathbf{V} \mathbf{R}_z(y_3)$ respectively.

Leaving the epipolar transformation invariant. The epipoles are encapsulated by the x_1, x_2 and the y_1, y_2 update parameters. Hence, the 3 degrees of freedom of the epipolar transformation are contained in the remaining update parameters: x_3, y_3 and σ . Removing them from the optimization freezes the underlying epipolar transformation.

5 Experimental Results

We compare an algorithm based on the orthonormal representation to other algorithms. We use simulated and real data, in §§5.1 and 5.2 respectively. Below, we give details about the compared methods, the measured quantities, the computation of an initial sub-optimal solution for structure and motion and the non-linear optimization scheme we use.

Compared methods. We compare the following motion parameterizations:

- **FREE** directly optimizes the 24 entries of the camera matrices. The gauge is left free to drift. The $24 - 7 = 17$ extra parameters are the homogeneous factors of each camera matrix and the 15-dimensional projective basis.
- **NORMALIZED** [15] is similar to **FREE**, but the gauge is fixed since a normalized coordinate frame is used. This is done by renormalizing the reconstruction before each step of the non-linear minimization and by including first-order gauge constraints into the minimization. The reconstruction basis, as well as the homogeneous scale of the camera matrices are constrained.
- **PARFREE** [6] partially fixes the gauge by optimizing only the entries of the second camera matrix, while keeping $P \sim (I \ 0)$. The $12 - 7 = 5$ extra parameters are the homogeneous scale of the second camera matrix, the global scene scale and the position of the plane at infinity.
- **MAPS** [3, 23] is a minimal parameterization based on multiple maps.
- **ORTHO** uses the orthonormal representation proposed in this paper.

Measured quantities. We measure two quantities characteristic of a bundle adjustment process, computational cost, i.e. CPU time to convergence, and the error at convergence.

Structure parameterization. We use the structure parameterization proposed in [7] which consists in scaling the reconstructed points such that their third element is unity. The 3 remaining free elements are then optimized. Note that this parameterization can be used only when a canonical basis enforcing $P \sim (I \ 0)$ is used. Therefore, methods **FREE** and **NORMALIZED** have their own structure parameterization: they optimize the 4 elements of each point.

Initialization. We compute an initial solution for the motion using the normalized 8 point algorithm [5]. Image point coordinates are standardized such that they lie in $[-1 \dots 1]$. Each point is reconstructed by minimizing its reprojection error.

Non-linear optimization. We use the Levenberg-Marquardt technique with analytic differentiation. This is a damped Gauss-Newton method. Let J be the Jacobian matrix and $H = J^T J$ the Gauss-Newton approximation of the Hessian matrix. The damping consists in augmenting the normal equations $H\delta = -J^T r$ to be solved at each iteration: $H \leftarrow H + W(\lambda)$. The parameter $\lambda \in \mathbb{R}$ is tuned heuristically, as described in [7, 21]. We try two approaches for the step control strategy, i.e. the choice of matrix $W(\lambda)$. First, in [21], the authors recommend $W(\lambda) = \lambda I$. This is the original idea of the Levenberg-Marquardt algorithm [10, 14]. This will be referred to as LM. Second, in [7], the authors recommend $W(\lambda) = (1 + \lambda) \text{diag}(H)$, i.e. multiply the diagonal entries of H by $1 + \lambda$. This strategy is recommended in [16] and is due to [17]. This will be referred to as SEBER.

Note that gauge freedoms cause $H = J^T J$ to be rank-deficient, but that the damped matrix is guaranteed to have full-rank. Hence, Levenberg-Marquardt iterations change both the actual estimated geometry as well as the gauge.

We take advantage of the sparse structure of H and J to efficiently solve the augmented normal equations, as described in [7, 21]. More precisely, the sparseness of the structure parameters is exploited, and the complexity of the computation is $\mathcal{O}(mp^3)$, where m is the number of points and p is the number of motion parameters. Hence, we can expect the computational cost for an iteration to be similar for all parameterizations when the number of points is very large, and to be very different when the number of points is low.

We stop the estimation when the difference between two consecutive residual errors is lower than a threshold ξ , chosen typically in the range $10^{-8} \leq \xi \leq 10^{-4}$.

5.1 Simulated Data

Experimental setup. We simulate points lying in a cube with 1 meter side length, observed by two cameras looking at the center of the cube. The standard configuration is the following: the focal length of the cameras is 1000 (expressed in number of pixels). They are situated 10 meters away from the center of the cube and the baseline between them is 1 meter. The number of simulated points is 50. We add a centered Gaussian noise on true point positions with a 2 pixel variance. The normal equations are augmented using

method LM. Each parameter of the above-described setup is independently varied to compare the parameterizations in different situations. The results are averaged over 100 trials. Computing the median gives similar results.

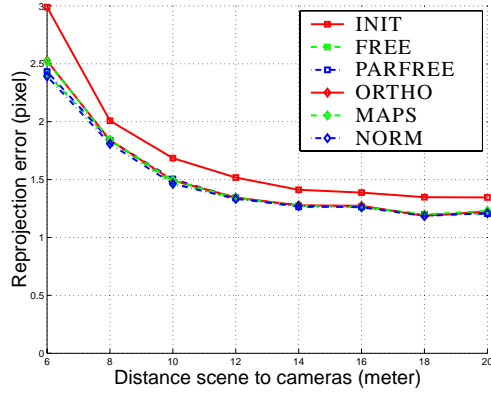
Results. Figure 1 shows the results. We observe that all methods have roughly the same accuracy, i.e. they give the same reprojection errors, up to small discrepancies. Further comments on these discrepancies are given in the next paragraph.

On the other hand, there are quite large discrepancies between the computational costs of each method. The methods that have the highest computational costs are NORM and FREE, followed by PARFREE. The minimal methods MAPS and ORTHO have the lowest computational cost, roughly the same. These discrepancies are explained by the fact that redundant methods have more unknowns to estimate than minimal ones. Solving the normal equations is therefore more expensive (see below). These observations are valid for other experiments (not shown here) where the focal length of the cameras is varied from 500 to 2000 pixels and where the baseline is varied from 1 to 3 meters. We also conduct the same experiments while augmenting the normal equation using SEBER. The same observations as above are valid. The results for all methods, compared to the LM augmentation, are worse in terms of both computational cost and reprojection error, while the discrepancies between the different methods for the reprojection error are reduced.

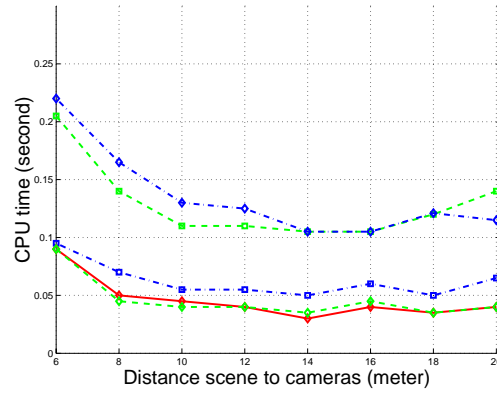
We observe that in our C implementation, the computational cost of each iteration is dominated by the resolution of the normal equations, whose size is directly linked to the number of parameters. We measure the computational cost of an iteration for the different parameterizations. As said above, the complexity is linear in the number of points and cubic in the number of motion parameters. For different numbers of points, we obtain :

	number of points			
iteration cost (second)	10	100	1,000	10,000
ORTHO	0.0045	0.0251	0.2152	2.0658
MAPS	0.0046	0.0251	0.2151	2.0658
PARFREE	0.0056	0.0307	0.2591	2.0753
FREE	0.0120	0.0589	0.5729	2.3231
NORM	0.0130	0.0664	0.6791	2.4148

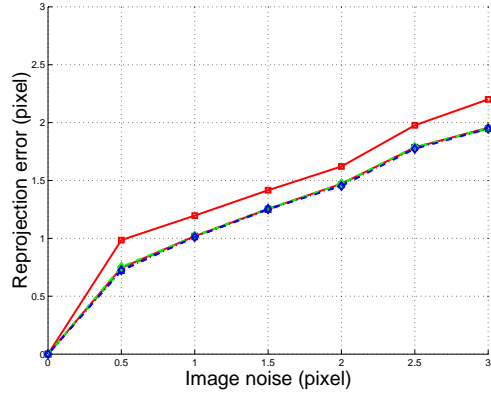
These results show that the differences in computational cost are largely dominated by the number of motion parameters. The discrepancies become smaller when the number of points increases beyond 10,000, which is very large in the case of structure from motion for two views.



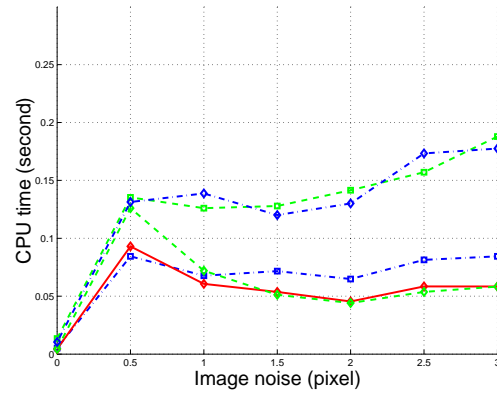
(a)




(b)




(c)



(d)



(e)



(f)

Figure 1: Reprojection error (left column) and CPU time to convergence (right column) measured against different simulation parameters: distance scene to cameras (first row), image noise (second row) and number of points (third row). Concerning the reprojection error, the curves are almost always undistinguishable, apart from the initialization. For the CPU time, methods are divided into three groups: (from top to bottom) FREE and NORM, PARFREE, then MAPS and ORTHO.

Convergence. As said above, there are small discrepancies in the reprojection errors achieved by the different methods, see in particular figure 1 (a). These small discrepancies are due to the fact that each parameterization may lead to a different local minimum of the cost function. To better characterize this phenomenon, we measure the rate of successful estimations for the different methods against the distance from the scene to the cameras. An estimation is successful if it is not improved by any of the other compared method. More precisely, let M and M' designate two methods and $\mathcal{E}_M(M')$ be the error achieved by method M initialized by the result of method M' . We define the success of an estimation made with method M as:

$$\text{Success}(M) \equiv [\forall M' \neq M, |\mathcal{E}_M(\text{INIT}) - \mathcal{E}_{M'}(M)| \leq \xi],$$

where ξ is the threshold used to stop the iterations. We obtain the following results :

		distance scene to cameras (meter)								
		success rate in %	6	8	10	12	14	16	18	20
minimal	{	ORTHO	83	97	97	100	100	98	100	99
		MAPS	88	95	96	100	100	99	98	99
redundent	{	PARFREE	100	100	94	100	100	100	100	100
		FREE & NORM	100	100	100	100	100	100	100	100

In the light of these results, we can say that methods using minimal parameters fall into local minima more often than methods based on redundant parameters. An explanation is that the minimal parameterizations are non-linear, while the overparameterizations are linear, in the entries of the projection matrices. Hence, the local quadratic approximation of the cost function used in Levenberg-Marquardt is more accurate for overparameterizations.

Essential matrix. As pointed out in §3.2, the orthonormal representation has a one-dimensional ambiguity when an essential matrix is considered. We want to check if, in the essential or near-essential cases, the orthonormal representation could induce numerical instabilities in the optimization process. For that purpose, we repeat the previous experiments, with the following two changes.

First, we map the fundamental matrix given by the 8 point algorithm to the closest essential matrix [8], and use this as an initial solution for the non-linear optimization. Hence, the target epipolar geometry, is a fundamental matrix, but the initial solution is an essential one.

Second, instead of using the coordinates of the points in the images, we use the coordinates of the points on the retina. Hence, the underlying true epipolar geometry is repre-

sented by an essential matrix. We run the experiments based on varying the geometry of the problem for both SEBER and LM.

We obtained results very similar to the previous experiments. This means that the orthonormal representation can be used for both fundamental and essential matrices, without inducing numerical instabilities, when an appropriate non-linear optimizer is employed.

5.2 Real Data

We use different pairs of the images shown in table 2, in order to cover all possibilities for the epipoles to be close to the images or at infinity, with 60 point correspondences. The results are shown on table 2. For each combination of images and each algorithm, we estimate the computational cost and the reprojection error. The last row of the table shows mean values for each algorithm over the set of image pairs. Note that for any image pair, the reprojection error is the same for all algorithms. Methods ORTHO, PARFREE and MAPS give the lowest computational costs, roughly twice as low as those of methods FREE and NORM. We obtain similar results using SEBER.

6 Conclusions

We studied the problem of estimating the fundamental matrix over a minimal set of 7 parameters. We proposed the orthonormal representation which enables to easily update an estimate of the fundamental matrix using 7 parameters. The canonic projection matrices can be directly extracted from the orthonormal representation. The method can be plugged into most of the (possibly sparse) non-linear optimizers such as Levenberg-Marquardt. We give a closed-form expression for the Jacobian matrix of the residuals with respect to the motion parameters for bundle adjustment purposes, necessary for Newton-type optimization techniques.

We conducted experiments on simulated and real data. Our conclusions are that the methods based on minimal parameter sets have lower computational cost but may be more frequently trapped in local minima.

Acknowledgements. The authors would like to thank Bill Triggs for discussions and one of the anonymous reviewers for very useful comments.



\mathcal{A}



\mathcal{B}



\mathcal{C}



\mathcal{D}

epipoles		views	INIT		FREE		PARFREE		ORTHO		MAPS		NORM	
\mathbf{e}	\mathbf{e}'		\mathcal{E}	\mathcal{T}	\mathcal{E}	\mathcal{T}	\mathcal{E}	\mathcal{T}	\mathcal{E}	\mathcal{T}	\mathcal{E}	\mathcal{T}	\mathcal{E}	\mathcal{T}
$-\infty$	∞	\mathcal{A}, \mathcal{B}	0.49	-	0.47	0.99	0.47	0.54	0.47	0.64	0.47	0.65	0.47	1.10
		\mathcal{A}, \mathcal{C}	0.68	-	0.65	0.67	0.65	0.39	0.65	0.29	0.65	0.34	0.65	0.70
$-\infty$	$-\infty$	\mathcal{A}, \mathcal{D}	0.84	-	0.67	0.70	0.67	0.38	0.67	0.33	0.67	0.33	0.67	0.74
∞	∞	\mathcal{B}, \mathcal{C}	0.57	-	0.53	0.27	0.53	0.14	0.53	0.15	0.53	0.14	0.53	0.23
∞	$-\infty$	\mathcal{B}, \mathcal{D}	0.79	-	0.55	0.25	0.55	0.23	0.55	0.18	0.55	0.19	0.55	0.25
		\mathcal{C}, \mathcal{B}	0.57	-	0.53	0.30	0.53	0.10	0.53	0.12	0.53	0.21	0.53	0.20
average \mathcal{E} and \mathcal{T}			0.66	-	0.57	0.53	0.57	0.30	0.57	0.28	0.57	0.31	0.57	0.54

Table 2: Reprojection error at convergence, \mathcal{E} , and CPU time to convergence, \mathcal{T} , obtained when combining pairs of images to obtain epipoles close to the images or towards infinity.

References

- [1] K.B. Atkinson, editor. *Close Range Photogrammetry and Machine Vision*. Whittles Publishing, 1996.
- [2] A. Bartoli. On the non-linear optimization of projective motion using minimal parameters. In *Proceedings of the 7th European Conference on Computer Vision, Copenhagen, Denmark*, volume 2, pages 340–354, May 2002.
- [3] A. Bartoli and P. Sturm. Three new algorithms for projective bundle adjustment with minimum parameters. Research Report 4236, INRIA, Grenoble, France, August 2001.
- [4] A. Bartoli, P. Sturm, and R. Horaud. Projective structure and motion from two views of a piecewise planar scene. In *Proceedings of the 8th International Conference on Computer Vision, Vancouver, Canada*, volume 1, pages 593–598, July 2001.
- [5] R. Hartley. In defence of the 8-point algorithm. In *Proceedings of the 5th International Conference on Computer Vision, Cambridge, Massachusetts, USA*, pages 1064–1070, June 1995.
- [6] R.I. Hartley. Projective reconstruction and invariants from multiple images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(10):1036–1041, October 1994.
- [7] R.I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, June 2000.
- [8] T.S. Huang and O.D. Faugeras. Some properties of the E matrix in two-view motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(12):1310–1312, December 1989.
- [9] K. Kanatani and D. D. Morris. Gauges and gauge transformations for uncertainty description of geometric structure with indeterminacy. *IEEE Transactions on Information Theory*, 47(5), July 2001.
- [10] K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, pages 164–168, 1944.

- [11] H.C. Longuet-Higgins. A computer program for reconstructing a scene from two projections. *Nature*, 293:133–135, September 1981.
- [12] Q.T. Luong and O. Faugeras. The fundamental matrix: Theory, algorithms and stability analysis. *International Journal of Computer Vision*, 17(1):43–76, 1996.
- [13] Q.T. Luong and T. Vieville. Canonic representations for the geometries of multiple projective views. *Computer Vision and Image Understanding*, 64(2):193–229, 1996.
- [14] D.W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, June 1963.
- [15] P. F. McLauchlan. Gauge invariance in projective 3D reconstruction. In *Proceedings of the Multi-View Workshop, Fort Collins, Colorado, USA*, 1999.
- [16] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C - The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1992.
- [17] G. A. F. Seber and C. J. Wild. *Non-linear regression*. John Wiley & Sons, New York, 1989.
- [18] J. Stuelpnagel. On the parametrization of the three-dimensional rotation group. *SIAM Review*, 6(4):422–430, October 1964.
- [19] P. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78(1), 2000.
- [20] P.H.S. Torr, A. Zisserman, and S.J. Maybank. Robust detection of degenerate configurations while estimating the fundamental matrix. *Computer Vision and Image Understanding*, 71(3):312–333, September 1998.
- [21] B. Triggs, P.F. McLauchlan, R.I. Hartley, and A. Fitzgibbon. Bundle adjustment — a modern synthesis. In B. Triggs, A. Zisserman, and R. Szeliski, editors, *Proceedings of the International Workshop on Vision Algorithms: Theory and Practice, Corfu, Greece*, volume 1883 of *Lecture Notes in Computer Science*, pages 298–372. Springer-Verlag, 2000.

- [22] R.Y. Tsai and T.S. Huang. Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(1):13–27, January 1984.
- [23] Z. Zhang. Determining the epipolar geometry and its uncertainty: A review. *International Journal of Computer Vision*, 27(2):161–195, March 1998.
- [24] Z. Zhang and C. Loop. Estimating the fundamental matrix by transforming image points in projective space. *Computer Vision and Image Understanding*, 82(2):174–180, May 2001.