Simple, Fast and Accurate Estimation of the Fundamental Matrix Using the Extended Eight-Point Schemes

Mohammed E. Fathy
m.e.fathy(at)gmail.com
Ashraf S. Hussein
ashrafh(at)acm.org
Mohammed F. Tolba
fahmytolba(at)gmail.com

Faculty of Computer and Information Sciences,
Ain Shams University,
Cairo, Egypt.

Abstract

The eight-point scheme is the simplest and fastest scheme for estimating the fundamental matrix (FM) from a number of noisy correspondences. As it ignores the fact that the FM must be singular, the resulting FM estimate is often inaccurate. Existing schemes that take the singularity constraint into consideration are several times slower and significantly more difficult to implement and understand. This paper describes extended versions of the eight-point (8P) and the weighted eight-point (W8P) schemes that effectively take the singularity constraint into consideration without sacrificing the efficiency and the simplicity of both schemes. The proposed schemes are respectively called the extended eight-point scheme (E8P) and the extended weighted eight-point scheme (EW8P). The E8P scheme was experimentally found to give exactly the same results as Hartley's algebraic distance minimization scheme while being almost as fast as the simplest scheme (i.e., the 8P scheme). At the expense of extra calculations per iteration, the EW8P scheme permits the use of geometric cost functions and, more importantly, robust weighting functions. It was experimentally found to give near-optimal results while being 8-16 times faster than the more complicated schemes such as Levenberg-Marquardt schemes. The FM estimates obtained by the E8P and the EW8P schemes perfectly satisfy the singularity constraint, eliminating the need to enforce the rank-2 constraint in a post-processing step.

1 Introduction

Numerous schemes have been proposed for estimating the FM from a set of noisy correspondences. Prominent examples are the (normalized) eight-point scheme (8P) [7], the iteratively re-weighted eight-point scheme (W8P) [27, 28], the Levenberg-Marquardt (LM) iterative minimization of the Sampson distance (LMS) [27, 28], Hartley's LM-based algebraic distance minimization scheme (ADMS) [8], the renormalization scheme (RNS) [11], the fundamental numerical scheme (FNS) [2], the heteroscedastic-errors-in-variable scheme (HEIV) [20], the constrained FNS (CFNS) [4], and the extended FNS (EFNS) [26]. Most of these schemes suffer from one or more of the following issues:

BMVC 2010 doi:10.5244/C.24.7

1. **Implementation difficulty**: Apart from the 8P and W8P schemes, most FM estimation schemes are relatively difficult to implement and understand for non-experts.
2. **Improper incorporation of the singularity constraint**: Some schemes do not take the singularity constraint into account while estimating the FM. Among these are the 8P and W8P schemes, FNS, HEIV, and RNS. An additional post-processing step is performed to correct the obtained estimate $\mathbf{F}_3$ so that it satisfies the singularity constraint. This is performed either by the *singular value decomposition (SVD)* correction [9] or by iteratively correcting the estimates obtained by FNS or HEIV [11, 20, 26]. Whereas SVD-correction gives inaccurate results [21, 26, 28], iterative-correction takes longer to run and makes the FM estimation process more complicated.
3. **Minimization of non-geometric cost functions**: Some schemes minimize the algebraic cost function to gain more speed at the expense of lower accuracy [8]. From accuracy point of view, however, it was experimentally and mathematically found that the switch from a geometric to an algebraic cost function has a smaller effect on accuracy than that incurred by the inappropriate incorporation of the singularity constraint [21, 26].
4. **Slow convergence**: All FM estimation schemes run several times slower than the 8P scheme. For instance, Sugaya and Kanatani [26] found in an FM estimation experiment involving $n = 100$ correspondences that FNS (followed by SVD correction), FNS (followed by iterative correction), CFNS, LMS, and EFNS take 13, 15, 25, 22, and 37 times, respectively, longer to run than the 8P scheme.
5. **Lack of robustness**: Apart from LM schemes, all the FM estimation schemes that take the singularity constraint into account (such as EFNS and CFNS) produce poor results with data sets contaminated with outliers and/or badly-located data. This is because they are non-robust least-squares techniques [27] that consider all data, including outliers and badly-located ones, equally important.
6. **Divergence possibility**: Some iterative schemes may exhaust all iterations and terminate at a very bad result, even worse than the initial estimate. Such schemes are not very reliable to be used in practice. We show in our experimental evaluation (Sect. 5) that (the stable version of) FNS [4] and EFNS [26] may diverge. In addition, we provide the data set that makes them diverge as part of the supplemental material [6].

This paper presents a simple, efficient, and effective approach for taking the rank-2 constraint into consideration during the estimation of FMs using the 8P and the W8P schemes. The two resulting schemes are called the *extended eight-point scheme (E8P)* and the *extended weighted eight-point scheme (EW8P)*, respectively. The E8P scheme minimizes an algebraic cost function whereas the EW8P scheme minimizes a weighted algebraic cost function. Given an initial guess $\mathbf{F}_k$ of the optimal FM estimate, the proposed approach replaces the nonlinear determinant function by a Taylor linear approximation at $\mathbf{F}_k$. A better estimate $\mathbf{F}_{k+1}$ is then obtained by solving the resulting linear system of equations, and the process is repeated until convergence. This is in contrast to the 8P and the W8P schemes which totally ignore the singularity constraint. The resulting iterative procedure is similar in concept to Newton's root finder and indeed has a similar rapid convergence.

The two proposed schemes have several advantages:

1. They are as simple to understand and implement as the traditional 8P schemes.
2. The E8P scheme was experimentally found to give exactly the same results as ADMS while being almost as fast as the simplest scheme (i.e., the 8P scheme).

3. The EW8P scheme not only improves the accuracy of the W8P scheme, but also runs about twice as fast. Experimental results indicate that it gives near-optimal results (its reprojection error is within 0.1% of the best result most of the time) while being 8-16 times faster than the more complicated schemes such as Levenberg-Marquardt schemes.
4. Similar to the LM scheme and the W8P scheme, the EW8P scheme offers a very flexible structure that permits the use of geometric cost functions and/or robust weighting functions.
5. The FM estimates obtained by the E8P and the EW8P schemes have unit Frobenius norm[1] and perfectly satisfy the singularity constraint, eliminating the need for a post-processing step to enforce the rank-2 constraint.

2 Background

Let I and I' be two perspective images of the same scene taken from different viewpoints. The epipolar constraint specifies that there exists a 3×3 rank-2 matrix $\mathbf{F}$ such that if $\boldsymbol{p}$ and $\boldsymbol{p}'$ are the image coordinates of any 3D point $\boldsymbol{P}$ on respectively I and I', then

$$\tilde{\boldsymbol{p}}^T \mathbf{F} \tilde{\boldsymbol{p}}' = 0. \tag{1}$$

where $\tilde{\boldsymbol{p}} = \begin{pmatrix} \boldsymbol{p}^T & 1 \end{pmatrix}^T$. The matrix $\mathbf{F}$ is called the Fundamental Matrix (FM) and the pair of points $\boldsymbol{M} = (\boldsymbol{p}, \boldsymbol{p}')$ is referred to as a correspondence. We refer to the quantity $d_A(\boldsymbol{M}) = d_A(\boldsymbol{p}, \boldsymbol{p}') = \tilde{\boldsymbol{p}}^T \mathbf{F} \tilde{\boldsymbol{p}}'$ as the algebraic distance. Simple algebraic manipulations show that $d_A(\boldsymbol{M})$ can be written in the following dot-product form:

$$d_A(\boldsymbol{M}) = \boldsymbol{m}^T \boldsymbol{f}. \tag{2}$$

where $\boldsymbol{m} = \begin{pmatrix} xx' & xy' & x & yx' & yy' & y & x' & y' & 1 \end{pmatrix}^T$ is called the measurement vector and $\boldsymbol{f} = \begin{pmatrix} f_1 & f_2 & \dots & f_9 \end{pmatrix}^T$ is the 9-vector containing the 9 entries of $\mathbf{F}$ in row-major order.

In theory, 7 correspondences between I and I' are generally enough to estimate $\mathbf{F}$. In practice, we are given a number $n > 7$ of noisy correspondences $\{\boldsymbol{M}_i = (\boldsymbol{p}_i, \boldsymbol{p}'_i)\}_{i=1}^n$. In this case, there is no a 3×3 rank-2 matrix that exactly satisfy the epipolar constraint (1) for all these noisy correspondences. So, we select the FM $\mathbf{F}$ that best fits these correspondences by solving an optimization problem of the following form:

$$\operatorname{argmin}_{\boldsymbol{f}} \quad C(\boldsymbol{f}) = \sum_{i=1}^{n} r_i^2, \quad \text{subject to} \quad D(\boldsymbol{f}) = 0. \tag{3}$$

where r_i measures the residual of the correspondence $\boldsymbol{M}_i$ as a function of the FM $\mathbf{F}(\boldsymbol{f})$, and $D(\boldsymbol{f})$ is the determinant of $\mathbf{F}(\boldsymbol{f})$:

$$D(\boldsymbol{f}) = \det \mathbf{F}(\boldsymbol{f}) = f_1 f_5 f_9 - f_1 f_6 f_8 - f_2 f_4 f_9 + f_2 f_6 f_7 + f_3 f_4 f_8 - f_3 f_5 f_7. \tag{4}$$

Different FM estimation schemes may differ in the particular definition they adopt for the residual function r_i, the way in which they enforce the rank-2 constraint, and/or the numerical scheme they use to solve the optimization problem in Eq. (3). Ideally, the residual function r_i should measure the Euclidean distance d_{Ei} in 4D space between the noisy correspondence $\boldsymbol{M}_i$ and the hyper-surface $\mathbb{S}$ implicitly defined by the epipolar constraint (1). The Euclidean distance d_E is also referred to as the reprojection error [9]. The rank-2 matrix

[1] After applying the denormalization transformation, the Frobenius norm of the obtained FM estimate may not remain one.

$\mathbf{F}_{ML}$ that minimizes $\sum_i d_{Ei}^2$ is the maximum-likelihood (ML) estimate [9] and the schemes that carry out this optimization, such as bundle adjustment (BA) schemes, are regarded as the gold standard (GS). Since they are relatively expensive, GS techniques mainly serve as a benchmark against which other approximate schemes are evaluated.

A more efficient, yet accurate alternative to d_E is Sampson distance d_S. It has the following closed-form expression:

$$d_{Si}^2 = w_{Si}^2 d_{Ai}^2 = \quad d_{Ai}^2 / |\nabla d_{Ai}|^2 \quad = d_{Ai}^2/(l_1^2 + l_2^2 + l_1'^2 + l_2'^2) \tag{5a}$$

where w_{Si} is the Sampson weighting [27], l_j is the jth component of the epipolar line $\boldsymbol{l} = \mathbf{F}\tilde{\boldsymbol{p}}_i'$, and l_j' is the jth component of the epipolar line $\boldsymbol{l}' = \mathbf{F}^T \tilde{\boldsymbol{p}}_i$. Sampson distance d_S^2 has the advantage of being a first-order approximation of d_E^2 [9].

3 The Extended Eight-Point (E8P) Scheme

3.1 Derivation

The E8P scheme aims to minimize the same algebraic cost function of the 8P scheme but subject to the additional constraint that the estimated matrix is singular:

$$\text{argmin}_{\boldsymbol{f}} \quad C_A(\boldsymbol{f}) = \sum_i (\boldsymbol{m}_i^T \boldsymbol{f})^2, \ \text{ subject to } \quad |\boldsymbol{f}|^2 = 1, \quad D(\boldsymbol{f}) = 0. \tag{6}$$

Let the measurement matrix $\mathbf{M}$ be the $n \times 9$ matrix whose ith row is $\boldsymbol{m}_i$. It follows that $C_A(\boldsymbol{f}) = |\mathbf{M}\boldsymbol{f}|^2 = \boldsymbol{f}^T \mathbf{M}^T \mathbf{M} \boldsymbol{f} = \boldsymbol{f}^T \mathbf{A} \boldsymbol{f}$. We hereafter refer to the 9×9 positive (semi-)definite matrix $\mathbf{A}$ as the moment matrix. Let $g_1(\boldsymbol{f}) = |\boldsymbol{f}|^2 - 1$ and $g_2(\boldsymbol{f}) = D(\boldsymbol{f})$. We shall refer to g_1 and g_2 as the constraint functions. It follows that Eq. (6) can be re-written as follows:

$$\text{argmin}_{\boldsymbol{f}} \ C_A(\boldsymbol{f}) = |\mathbf{M}\boldsymbol{f}|^2, \quad \text{subject to} \quad g_1(\boldsymbol{f}) = 0, \quad g_2(\boldsymbol{f}) = 0. \tag{7}$$

Introducing Lagrange multipliers converts (7) into the following unconstrained optimization problem[2]:

$$\text{argmin}_{\boldsymbol{f},\lambda_1,\lambda_2} \ L(\boldsymbol{f},\lambda_1,\lambda_2) = |\mathbf{M}\boldsymbol{f}|^2 + 2\lambda_1 g_1(\boldsymbol{f}) + 2\lambda_2 g_2(\boldsymbol{f}). \tag{8}$$

At the optimal point $(\hat{\boldsymbol{f}},\lambda_1,\lambda_2)$, the partial derivatives ($\boldsymbol{\partial}_{\boldsymbol{f}} L$, $\boldsymbol{\partial}_{\lambda_1} L$, and $\boldsymbol{\partial}_{\lambda_2} L$) must vanish. This leads to the following system of equations:

$$\mathbf{M}^T \mathbf{M} \hat{\boldsymbol{f}} + \lambda_1 \boldsymbol{\partial}_{\boldsymbol{f}} g_1(\hat{\boldsymbol{f}}) + \lambda_2 \boldsymbol{\partial}_{\boldsymbol{f}} g_2(\hat{\boldsymbol{f}}) = \mathbf{0}, \quad g_1(\hat{\boldsymbol{f}}) = 0, \quad g_2(\hat{\boldsymbol{f}}) = 0. \tag{9}$$

Assume that we are given a guess $\boldsymbol{f}_k$ of the minimizer $\hat{\boldsymbol{f}}$. Let ℓ_i be the first-order Taylor series approximation of the constraint function g_i at $\boldsymbol{f}_k$. Then, ℓ_i is given as follows:

$$\ell_i(\boldsymbol{f}) = g_i(\boldsymbol{f}_k) + \boldsymbol{\partial}_{\boldsymbol{f}}^T g_i(\boldsymbol{f}_k)(\boldsymbol{f} - \boldsymbol{f}_k). \tag{10}$$

The linearization ℓ_i can be used to approximate the behavior of g_i at any point $\boldsymbol{f}$, but the approximation is highly accurate when $\boldsymbol{f}$ is sufficiently close to $\boldsymbol{f}_k$ and when g_i is a low-order polynomial that is free from exponentials, trigonometric functions, or any highly nonlinear functions. Consequently, if the guess $\boldsymbol{f}_k$ is provided sufficiently close to $\hat{\boldsymbol{f}}$, we can replace each constraint function g_i in Eq. (9) by its linearization ℓ_i to obtain:

$$\mathbf{M}^T \mathbf{M} \hat{\boldsymbol{f}} + \lambda_1 \boldsymbol{\partial}_{\boldsymbol{f}} \ell_1(\hat{\boldsymbol{f}}) + \lambda_2 \boldsymbol{\partial}_{\boldsymbol{f}} \ell_2(\hat{\boldsymbol{f}}) = \mathbf{0}, \quad \ell_1(\hat{\boldsymbol{f}}) = 0, \quad \ell_2(\hat{\boldsymbol{f}}) = 0. \tag{11}$$

[2] Using arbitrary non-zero constants to the left of Lagrange multipliers does not alter the optimization problem.

Substituting with the definitions of ℓ_i into Eq. (11) yields

$$\begin{aligned} \mathbf{M}^T\mathbf{M}\hat{\boldsymbol{f}} + \lambda_1 \boldsymbol{\partial}_{\boldsymbol{f}} g_1(\boldsymbol{f}_k) + \lambda_2 \boldsymbol{\partial}_{\boldsymbol{f}} g_2(\boldsymbol{f}_k) &= \mathbf{0}, & \text{(12a)} \\ g_1(\boldsymbol{f}_k) + \boldsymbol{\partial}^T_{\boldsymbol{f}} g_1(\boldsymbol{f}_k)(\hat{\boldsymbol{f}} - \boldsymbol{f}_k) &= 0, & \text{(12b)} \\ g_2(\boldsymbol{f}_k) + \boldsymbol{\partial}^T_{\boldsymbol{f}} g_2(\boldsymbol{f}_k)(\hat{\boldsymbol{f}} - \boldsymbol{f}_k) &= 0. & \text{(12c)} \end{aligned}$$

where $\boldsymbol{\partial}_{\boldsymbol{f}} g_1(\boldsymbol{f}) = 2\boldsymbol{f}$. The expression of $\boldsymbol{\partial}_{\boldsymbol{f}} g_2(\boldsymbol{f})$ was omitted for space limitations. It can be found elsewhere [12, p. 284] or easily by taking the partial derivatives of the expression in Eq. (4).

Define the 2×9 matrix $\mathbf{J}_k = \begin{bmatrix} \boldsymbol{\partial}_{\boldsymbol{f}} g_1(\boldsymbol{f}_k) & \boldsymbol{\partial}_{\boldsymbol{f}} g_2(\boldsymbol{f}_k) \end{bmatrix}^T$, the 2×1 vector $\boldsymbol{\lambda} = \begin{pmatrix} \lambda_1 & \lambda_2 \end{pmatrix}^T$, and the 2×1 vector $\boldsymbol{c}_k = \mathbf{J}_k \boldsymbol{f}_k - \begin{pmatrix} g_1(\boldsymbol{f}_k) & g_2(\boldsymbol{f}_k) \end{pmatrix}^T$. The above system of equations can then be compactly rewritten as follows:

$$\mathbf{M}^T\mathbf{M}\hat{\boldsymbol{f}} + \mathbf{J}_k^T \boldsymbol{\lambda} = \mathbf{0}, \quad (13)$$

$$\mathbf{J}_k \hat{\boldsymbol{f}} = \boldsymbol{c}_k. \quad (14)$$

This is a linear system of 11 equations in the vector of 11 unknowns $\boldsymbol{x} = \begin{pmatrix} \hat{\boldsymbol{f}}^T & \boldsymbol{\lambda}^T \end{pmatrix}^T$. Finally, we can rewrite Eq. (13) in the following more elegant form:

$$\mathbf{L}_k \boldsymbol{x} = \boldsymbol{b}_k \quad (15)$$

where the 11×11 symmetric matrix $\mathbf{L}_k$ and the 11-vector $\boldsymbol{b}_k$ are defined as follows:

$$\mathbf{L}_k = \begin{bmatrix} \mathbf{M}^T\mathbf{M} & \mathbf{J}_k^T \\ \mathbf{J}_k & \mathbf{0}_{2\times 2} \end{bmatrix}, \boldsymbol{b}_k = \begin{pmatrix} \mathbf{0}_{9\times 1} \\ \boldsymbol{c}_k \end{pmatrix} \quad (16)$$

The linear system of equations (15) can be solved for $\boldsymbol{x}$ using standard elimination techniques such as the LU decomposition [6]. The first nine entries of $\boldsymbol{x}$ give the desired estimate $\hat{\boldsymbol{f}}$, whereas the two other entries give the less-interesting value of the Lagrange multiplier $\boldsymbol{\lambda}$. It should be stressed that we do not make any use of $\boldsymbol{\lambda}$ in any subsequent computation.

The estimate of $\boldsymbol{f}$ obtained from Eq. (15) may not exactly be equal to the minimizer $\hat{\boldsymbol{f}}$ of (7) because the linearizations ℓ_i may not have approximated the behavior of the original constraint functions g_i at $\hat{\boldsymbol{f}}$ exactly. Nevertheless, this estimate can be used as a better guess $\boldsymbol{f}_{k+1}$ of the minimizer $\hat{\boldsymbol{f}}$, allowing us to obtain more accurate linearizations ℓ_i of the constraints g_i and solve Eq. (15) again. Equation (15) should then be written in the following recurrent form:

$$\mathbf{L}_k \boldsymbol{x}_{k+1} = \boldsymbol{b}_k \quad (17)$$

This process is repeated until we find that subsequent estimates of $\boldsymbol{f}$ are sufficiently close.

Since the E8P approach relies on linear Taylor approximations near the solution, it exhibits a similar behavior to the Newton-Raphson scheme for root-finding: Once the current estimate is sufficiently close to the solution, the convergence becomes highly rapid. This is especially true in our case because the constraints being linearized are low-order polynomials that involve no exponentials or trigonometric functions.

It is not difficult to show that, upon convergence, the obtained estimate $\boldsymbol{f}_k$ is a critical point (and hopefully the global minimizer) of the Lagrange version (8) of the optimization problem given in Eq. (7). Indeed, substituting $\hat{\boldsymbol{f}} = \boldsymbol{f}_k$ into the expanded update equation (12) reveals that $\boldsymbol{f}_k$ satisfies the equation set in (9). This equation set defines the necessary conditions for the critical points of the optimization problem of Eq. (8). The satisfaction of the last of these equations ($g_2(\boldsymbol{f}_k) = 0$) does also imply that $\boldsymbol{f}_k$ satisfies the singularity constraint.

A key advantage of this scheme is that the count of flops it makes in each iteration is independent of n. Indeed, it takes only $(2d'^3/3)$ flops to solve a linear system such as (17) every iteration [6], where $d' = 11$. Yet, a more efficient update procedure can be obtained if the 9×9 moment matrix $\mathbf{A} = \mathbf{M}^T\mathbf{M}$ is invertible. This procedure further reduces the count of flops performed in each iteration to just $(2d^2)$, where $d = 9$.

The more efficient procedure can be derived as follows. If $\mathbf{A}$ is invertible, we can multiply both sides in Eq. (13) by $\mathbf{A}^{-1}$ to obtain the following formula for $\boldsymbol{f}_{k+1}$ in terms of $\boldsymbol{\lambda}$:

$$\boldsymbol{f}_{k+1} = -\mathbf{A}^{-1}\mathbf{J}_k^T\boldsymbol{\lambda}. \tag{18}$$

Substituting Eq. (18) into Eq. (14) yields:

$$-\mathbf{J}_k\mathbf{A}^{-1}\mathbf{J}_k^T\boldsymbol{\lambda} = \boldsymbol{c}_k \tag{19}$$

Define the 9×2 matrix $\mathbf{T}_k = \mathbf{A}^{-1}\mathbf{J}_k^T$ and the 2×2 symmetric matrix $\mathbf{N}_k = \mathbf{J}_k\mathbf{T}_k = \mathbf{J}_k\mathbf{A}^{-1}\mathbf{J}_k^T$. We can solve Eq. (19) by multiplying both sides by $-\mathbf{N}_k^{-1}$ to obtain $\boldsymbol{\lambda} = -(\mathbf{N}_k^{-1}\boldsymbol{c}_k)$. Substituting into Eq. (18) yields the following more efficient update equation:

$$\boldsymbol{f}_{k+1} = \mathbf{T}_k(\mathbf{N}_k^{-1}\boldsymbol{c}_k). \tag{20}$$

An efficient and robust implementation should be able to switch between the two alternative solution procedures (17) and (20) based on the rank of $\mathbf{M}$. When it discovers that $\mathbf{M}$ has rank 9, it should calculate $\mathbf{A}^{-1}$ once in a pre-processing step and use Eq. (20) to update the solution at each iteration. If it discovers that $\mathbf{M}$ does not have rank 9, it should calculate $\mathbf{A}$ once in a pre-processing step and use Eq. (17) to advance the solution. We describe one such implementation at the end of this section.

The E8P scheme requires an initial estimate $\boldsymbol{f}_0$ of the minimizer $\hat{\boldsymbol{f}}$. As with the traditional 8P scheme, we can use the right singular vector $\boldsymbol{v}_9$ corresponding to the smallest singular value d_9 of the measurement matrix $\mathbf{M}$. Among all unit vectors, $\boldsymbol{v}_9$ is the minimizer of $|\mathbf{M}\boldsymbol{f}|^2$ when the singularity constraint is ignored. With this particular choice, it takes the E8P scheme only 3-4 iterations to converge in most of the cases. We have also considered different choices of the initial estimate $\boldsymbol{f}_0$. When we used the estimate obtained by running the 7-point scheme [9] on a random subset of 7 correspondences, it took our technique about 5-6 iterations to converge. In all our experiments, the E8P scheme never diverged to a bad solution. This indicates that our technique converges rapidly and reliably to the correct solution whenever provided with a *reasonable* initial estimate $\boldsymbol{f}_0$.

3.2 Implementation

We present an efficient implementation of the E8P scheme in a MATLAB-like pseudocode. The implementation switches to the more efficient iterative formulation (20) if it discovers that the rank of $\mathbf{M}$ is 9. This will almost always be the case in practice because we are provided with $n >> 8$ noisy correspondences. It is assumed that we are given as input a set of correspondences $\mathbb{C} = \{\boldsymbol{M}_i\}_{i=1}^{i=n}$, a convergence constant δ, a maximum number of iterations *max*, and an optional initial estimate $\boldsymbol{f}_0$. It is also assumed that the coordinates of the input correspondences have been normalized as described in [7].

$\mathbf{M} \leftarrow$ BuildMeasurementMatrix($\mathbb{C}$);
$(\boldsymbol{d}, \mathbf{V}) \leftarrow$ SVD($\mathbf{M}$);[3]
if no $\boldsymbol{f}_0$ was provided, **set** $\boldsymbol{f}_0 \leftarrow \mathbf{V}(:,9)$;

[3] We do not require the left orthogonal factor $\mathbf{U}$ of the SVD of $\mathbf{M}$ and thus it is not calculated. This makes SVD run faster [6].

if $n > 8$ AND $\boldsymbol{d}(9)/\boldsymbol{d}(1) \geq 1 \times 10^{-10}$ **then**
  Precalculate $\mathbf{A}^{-1} \leftarrow \mathbf{V} * \text{diag}(\boldsymbol{d}(1)^{-2}, \boldsymbol{d}(2)^{-2}, \ldots, \boldsymbol{d}(9)^{-2}) * \mathbf{V}^T$;
  Use update equation (20) to advance the solution;
  Stop when *max* iterations are exhausted or when $\left|\boldsymbol{f}_{k+1} - \boldsymbol{f}_k\right|^2 \leq \delta^2$;
else
  Precalculate $\mathbf{A} = \mathbf{M}^T\mathbf{M}$;
  Use update equation (17) to advance the solution;
  Stop when *max* iterations are exhausted or when $\left|\boldsymbol{f}_{k+1} - \boldsymbol{f}_k\right|^2 \leq \delta^2$;
return $\boldsymbol{f}_{conv}$;

4 The Extended Weighted Eight-Point (EW8P) Scheme

The (traditional) weighted eight-point scheme (W8P) attempts to minimize a cost function of the following form:

$$C(\boldsymbol{f}) = \sum_{i=1}^{n} (w_i d_{Ai})^2 = \sum_{i=1}^{n} (w_i \boldsymbol{m}_i^T \boldsymbol{f})^2, \quad \text{subject to } |\boldsymbol{f}|^2 = 1. \tag{21}$$

where each weight w_i depends on both the correspondence $\boldsymbol{M}_i$ and the FM $\boldsymbol{f}$. The scheme carries out the optimization in the following iterative fashion. At iteration k, the FM estimate $\boldsymbol{f}_k$ is used to evaluate the weight w_i for each correspondence $\boldsymbol{M}_i$. The factor w_i is then used to weight the corresponding row $\boldsymbol{m}_i^T$ in the measurement matrix $\mathbf{M}$. The optimization problem can then be written as follows:

$$\text{argmin}_{\boldsymbol{f}} \; L(\boldsymbol{f}, \lambda) = |\mathbf{M}_k \boldsymbol{f}|^2 - \lambda(|\boldsymbol{f}|^2 - 1). \tag{22}$$

where the $n \times 9$ matrix $\mathbf{M}_k = \mathbf{W}_k\mathbf{M}$ and $\mathbf{W}_k$ is the $n \times n$ diagonal matrix $\text{diag}(w_1, w_2, \ldots, w_n)$. The W8P scheme can be made to minimize the Sampson geometric cost function by setting $w_i = w_{Si} = 1/|\nabla d_{Ai}|$ [27, 28]. The scheme that results when this particular form of w_i is used is called the *Sampson scheme* [3]. It can also be made to minimize a robust geometric cost function as in [27] by setting $w_i = w_{Si}w_{Hi}$ where w_{Hi} is the robust Huber weighting that helps lower the effect of badly-located correspondences on the estimated FM.

Similar to the 8P scheme, the W8P can be extended to directly incorporate the rank-2 constraint into the optimization. We refer to the resulting scheme as the *extended weighted eight-point (EW8P)* scheme. It differs from the E8P scheme in that the measurement matrix $\mathbf{M}$ has to be re-weighted in each iteration k by the weight matrix $\mathbf{W}_k$, yielding $\mathbf{M}_k = \mathbf{W}_k\mathbf{M}$. Consequently, the update equation (17) changes into:

$$\begin{bmatrix} \mathbf{A}_k & \mathbf{J}_k^T \\ \mathbf{J}_k & \mathbf{0}_{2\times 2} \end{bmatrix} \begin{pmatrix} \boldsymbol{f}_{k+1} \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{0}_{9\times 1} \\ \boldsymbol{c}_k \end{pmatrix}, \text{ or shortly into} \tag{23a}$$

$$\mathbf{H}_k \boldsymbol{x}_{k+1} = \boldsymbol{b}_k. \tag{23b}$$

The implementations of the EW8P and the 8P schemes differ in two ways. First, the EW8P scheme calculates in each iteration a set of different weights and subsequently a new moment matrix $\mathbf{A}_k = \mathbf{M}_k^T\mathbf{M}_k$. Second, the EW8P scheme does not use the solution procedure given in Eq. (20) because it needs to compute a different inverse $\mathbf{A}_k^{-1}$ every iteration which wastes any possible gain in efficiency. Apart from these two differences, the EW8P and the 8P schemes have quite similar implementations.

Unlike the E8P scheme, the count of flops made in each iteration of the EW8P scheme is no more independent of n. This is due to the need to calculate a set of different weights and subsequently calculate a new moment matrix $\mathbf{A}_k = \mathbf{M}_k^T\mathbf{M}_k$ every iteration. This scheme may not, consequently, be as efficient as the E8P scheme. Practical experiments, however, show that the EW8P scheme runs many times faster than the other more complicated schemes (such as FNS, EFNS, LM). It is even faster than the less-accurate W8P scheme and Hartley's ADMS (as long as n is not very large).

5 Results and Discussion

The accuracy and efficiency of the proposed schemes were examined against several other existing schemes using 6 pairs of real images with varying numbers of correspondences. The schemes that were included in the evaluation are (a) the 8P scheme with SVD correction, (b) the E8P scheme, (c) the ADMS [8], (d) the W8P scheme with Sampson weighting and SVD correction, (e) the EW8P scheme with Sampson weighting, (f) the stable version of FNS with SVD correction (FNSS) [2, 4], (g) the stable version of FNS with iterative rank correction (FNSI) [26], (h) the EFNS [26], and (i) the LM minimization of $\sum_i d_{Si}^2$ using the SVD 7-parameterization of Bartoli and Sturm (LMBS7) [1]. We have set the maximum iteration count of all schemes to 200. The normalization step proposed by Hartley [7] was applied on the input set of matches to improve the conditioning of the calculations with all schemes except for the EFNS where the points were uniformly scaled down with a factor $f_0 = 1/600$. The experiment and all schemes were implemented in C++ using double-precision arithmetic. The experiment was carried out on a Dell Vostro notebook equipped with an Intel Core 2 Duo T7500 (2.2 GHz) processor and 2.0 GB RAM.

For each pair of images, correspondences were established using the *scale invariant feature transform (SIFT)* [18]. Outlier matches were then removed using RANSAC along with the seven-point scheme [9]. The resulting set of inliers was then used as input to all the FM estimation schemes under evaluation. For each scheme, we recorded the root-mean-square reprojection error E of the obtained FM estimate $\mathbf{F}$, the computing time T (in milliseconds), the least singular value σ_3 of the version of $\mathbf{F}$ having unit Frobenius norm, and the number of iterations i taken to converge. The quantity σ_3 is the singularity distance or the distance between the FM $\mathbf{F}$ and the hyper-surface in 9D of singular matrices. It is used to measure the degree by which $\mathbf{F}$ deviates from satisfying the singularity constraint. The reprojection error of a given correspondence $\boldsymbol{M}_i$ with respect to $\mathbf{F}$ was measured by first correcting $\boldsymbol{M}_i$ into $\hat{\boldsymbol{M}}_i$ with Kanatani's iterative correction [13] and measuring the distance between $\boldsymbol{M}_i$ and $\hat{\boldsymbol{M}}_i$.

The 6 image pairs used in our experiment are shown in Fig. 1. The measurements taken on these pairs are shown in Table 1 and are discussed below.

Across all data sets, the E8P scheme runs almost as fast as the 8P scheme while achieving a lower reprojection error. It gives almost the same results as ADMS while being much simpler and 2-22 times faster. Indeed, the two schemes are just different ways of minimizing the same algebraic cost function subject to the same set of constraints. Although these schemes are not as accurate as LMBS7, the accuracy gap shrinks with the increase in the number of correspondences (as in the Castle and Moebius sequences).

The EW8P scheme has an impressive performance. It achieves almost the same accuracy (quantity E in Table 1) as the more complicated schemes while being, on average, 11 times faster than LMBS7, 120 times faster than EFNS, and 12 times faster than FNSI[4]. Not only does it improve the accuracy of the W8P scheme but its running time as well. It also runs faster than ADMS except when the number of correspondences n gets very large (as in the Moebius sequence).

[4]This excludes the timings of EFNS and FNS in the Mosque B sequence where both diverge.

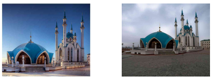

(a) Mosque A (36)[16, 17]

(b) Sphinx (72)[23, 24]

(c) Mosque B (172)[14, 15]

(d) Corridor (412)[9]

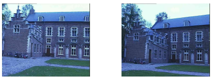

(e) Castle (524)

(g) Moebius (2411)[10]

Figure 1: The image pairs used in our experiment sorted ascendingly based on the number of matches (shown between brackets). The Castle sequence contains images 0 and 5 of the Lueven castle sequence [19, 22].

	Mosque A (36)				Sphinx (72)			
	E	T	σ_3	i	E	T	σ_3	i
8P	3.9867	0.045	1.4E-019	1	0.7666	0.066	1.0E-018	1
E8P	1.8410	0.044	6.0E-021	3	0.7588	0.061	4.1E-019	3
ADMS	1.8410	0.854	3.0E-020	12	0.7588	1.363	3.9E-019	23
W8P	5.2731	0.148	3.9E-020	3	0.7309	0.213	9.5E-019	3
EW8P	1.0451	0.075	7.6E-018	4	0.7259	0.101	9.5E-019	3
FNSS	5.2479	0.770	1.8E-020	5	0.7307	1.099	6.2E-019	4
FNSI	1.0456	0.929	1.7E-019	8	0.7257	1.326	8.3E-019	7
EFNS	1.0436	5.717	1.5E-017	35	0.7253	8.635	1.7E-018	32
LMBS7	1.0449	0.960	1.3E-019	7	0.7257	1.027	1.6E-019	4

	Mosque B (172)				Corridor (412)			
	E	T	σ_3	i	E	T	σ_3	i
8P	1.1939	0.122	1.1E-018	1	0.5080	0.296	6.4E-018	1
E8P	1.1367	0.177	2.7E-019	5	0.5000	0.303	1.8E-017	6
ADMS	1.1366	1.275	3.4E-019	19	0.5000	1.112	1.7E-017	7
W8P	1.0956	1.973	5.2E-019	18	0.4456	1.650	2.3E-017	6
EW8P	1.0826	0.514	4.9E-019	11	0.4451	0.842	3.8E-018	7
FNSS	7.6038	118.458	3.1E-018	200	0.4455	13.335	2.4E-017	10
FNSI	7.4437	108.950	1.6E-017	204	0.4449	14.005	1.3E-017	14
EFNS	6.5493	113.401	4.9E-010	200	0.4449	44.923	1.3E-017	35
LMBS7	1.0064	5.143	4.3E-019	9	0.4449	8.235	1.8E-017	6

	Castle (524)				Moebius (2411)			
	E	T	σ_3	i	E	T	σ_3	i
8P	0.4186	0.400	1.9E-020	1	0.2797	1.758	3.8E-018	1
E8P	0.3589	0.406	9.9E-020	4	0.2774	1.770	7.7E-018	3
ADMS	0.3589	2.737	3.4E-020	36	0.2774	3.340	1.7E-018	22
W8P	0.4226	1.250	5.1E-020	3	0.2797	6.713	4.3E-018	4
EW8P	0.3587	0.840	9.4E-017	4	0.2773	3.985	8.3E-018	4
FNSS	0.4224	5.221	1.9E-019	3	0.2797	31.206	1.6E-018	4
FNSI	0.3587	6.528	4.0E-020	6	0.2771	36.476	3.3E-018	7
EFNS	0.3587	52.777	1.1E-017	33	0.2771	1275.936	2.1E-018	179
LMBS7	0.3587	7.126	1.8E-019	4	0.2771	63.091	5.7E-018	8

Table 1: The performance of each scheme over the different data sets. E is the RMS reprojection error, T is the computing time in milliseconds, σ_3 is the singularity distance, and i is the number of iterations taken to converge.

As mentioned above, the accuracy of the EW8P scheme is almost the same as that of LMBS7 (its reprojection error E is within 0.1% of LMBS7). The only exception to this is the Mosque B sequence in which the reprojection error of EW8P is 7.6% higher than that of LMBS7. Yet, it is better than the reprojection errors achieved by all other schemes. In particular, both FNS and EFNS diverge in that particular case. Indeed, they exhaust all of the 200 iterations and terminate at a solution with a reprojection error that is relatively high. We have tried to do the estimation using the implementation written by the original authors of EFNS [25] but it also diverged. This indicates that the use of FNS and EFNS should be made with sufficient precautions to avoid their divergence possibility. As evidence, we provide, as part of the supplemental material, the set of correspondences identified in our experiment for the Mosque B sequence [5].

Table 1 also indicates that the singularity distance σ_3 of the proposed schemes is within the same small orders of magnitude of the singularity distances of the inherently rank-2 schemes (such as ADMS and LMBS7) and the SVD- or iteratively-corrected schemes (such as 8P, W8P, FNSS, and FNSI). It can therefore be deduced that the proposed schemes converge to singular FM estimates and no further post-processing is required to enforce that constraint. This observation agrees with the theoretical argument given in Sect. 3 that the convergence of the schemes implies the satisfaction of the singularity constraint.

6 Conclusion

We have presented two accurate, yet fast and simple schemes for FM estimation. The two schemes are called the *extended eight-point scheme (E8P)* and the *extended weighted eight-point scheme (EW8P)*. As their names imply, they extend the 8P and W8P schemes to effectively take the singularity constraint into consideration without sacrificing the efficiency and the simplicity of both schemes. Experimental results confirmed that the E8P scheme runs as fast as the 8P scheme while giving exactly the same accuracy of the ADMS. It is interesting to point out that the EW8P scheme improves the accuracy of the W8P scheme while running about twice as fast. Whereas the E8P scheme minimizes the sum of squared algebraic errors, the EW8P scheme minimizes the sum of weighted algebraic errors. With proper weights, the EW8P scheme can be made to minimize a geometric criterion or more importantly weight down bad data to obtain robust FM estimates.

When simplicity and/or ultimate speed are the major concern, the E8P scheme should be preferred to the 8P scheme as it offers better accuracy while being as fast and as simple. By sacrificing a little more speed, the EW8P scheme offers near-optimal results while still running 8-16 times faster than the more complicated schemes such as LMBS7. If accuracy is the major concern, the output of the EW8P scheme should be refined on using either LMBS7 or BA. When the input data set is contaminated with outliers and badly-located data, it becomes necessary to use the EW8P scheme and optionally follow it by LMBS7 or BA. In this case, a robust weighting function should be adopted and initial estimates of the FM as well as the noise standard deviation should be obtained using a robust technique like *Least Median of Squares* [27, 28]. Finally, the use of FNS and EFNS should be made with enough precautions due to their divergence possibility.

Acknowledgements
Our C++ implementations of EFNS and LMBS7 were converted from the high-quality Octave implementations provided by Kenichi Kanatani and Yasuyuki Sugaya. We extend our thanks to Richard Hartley and Hongdong Li for providing the Corridor image sequence, and to Zhengyou Zhang for helpful discussions.

References

[1] A. Bartoli and P. Sturm. Nonlinear estimation of the fundamental matrix with minimal parameters. *IEEE Trans. Pattern Anal. Machine Intell.*, 26(3):426–432, 2004.

[2] W. Chojnacki, M.J. Brooks, A.J. Van Den Hengel, and D.J. Gawley. On the fitting of surfaces to data with covariances. *IEEE Trans. Pattern Anal. Machine Intell.*, 22(11): 1294–1303, 2000.

[3] W. Chojnacki, M.J. Brooks, and A. van den Hengel. Rationalising the renormalisation method of Kanatani. *J. of Math. Imaging and Vision*, 14(1):21–38, 2001.

[4] W. Chojnacki, M.J. Brooks, A. van den Hengel, and D. Gawley. FNS, CFNS and HEIV: A Unifying Approach. *J. of Math. Imaging and Vision*, 23(2):175–183, 2005.

[5] M. E. Fathy, A. S. Hussein, and M. F. Tolba. Supplemental material: The correspondences of the mosque b sequence, 2010. `MosqueB_Data.m`.

[6] G. Golub and C. van Loan. *Matrix Computations*. Johns Hopkins Studies in Mathematical Sciences, 1996.

[7] R. Hartley. In defense of the eight-point algorithm. *IEEE Trans. Pattern Anal. Machine Intell.*, 19(6):580–593, 1997.

[8] R. Hartley. Minimizing algebraic error in geometric estimation problems. In *Int'l Conf. on Comput. Vision (ICCV)*, pages 469–476, 1998.

[9] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, March 2004. ISBN 0521540518.

[10] H. Hirschmuller and D. Scharstein. Evaluation of cost functions for stereo matching. In *Proc. IEEE Conf. on Comput. Vision and Pattern Recognition (CVPR)*, pages 1–8, 2007.

[11] K. Kanatani and N. Ohta. Comparing optimal 3-d reconstruction for finite motion and optical flow. *Electronic Imaging*, 12(03):478–488, 2003.

[12] K. Kanatani and Y. Sugaya. High accuracy fundamental matrix computation and its performance evaluation. *IEICE - Trans. Inf. Syst.*, E90-D(2):579–585, 2007.

[13] K. Kanatani, Y. Sugaya, and H. Niitsuma. Triangulation from two views revisited: Hartley-sturm vs. optimal correction. In *Proc. British Machine Vision Conf. (BMVC)*, pages 173–182, 2008.

[14] Kul Sharif Mosque. Available at myriad URLs with an unknown source, 2010. URL `http://karim74.files.wordpress.com/2009/06/kul-sharif-mosque-in-kazan-russia.jpg`.

[15] Kul Sharif Mosque. Photo by _malfet. Available under a CC license, 2010. URL `http://www.flickr.com/photos/malfet/1000391372/`.

[16] Kul Sharif Mosque. Photo by Nu.Clear.Photo. Available under a CC license, 2010. URL `http://www.flickr.com/photos/nu_clear/1475665351/`.

[17] Kul Sharif Mosque. Reprinted with permission from mikhail_k, 2010. URL http://www.flickr.com/photos/mikhailk/3370766410/.

[18] D. Lowe. Distinctive image features from scale-invariant keypoints. *Int'l J. Comput. Vision*, 60(2):91–110, November 2004.

[19] Lueven Castle Image Sequence, 2010. URL http://www.cs.unc.edu/~marc/data/castlejpg.zip.

[20] BC Matei and P. Meer. Estimation of nonlinear errors-in-variables models for computer vision applications. *IEEE Trans. Pattern Anal. Machine Intell.*, 28(10):1537–1552, 2006.

[21] T. Migita and T. Shakunaga. Evaluation of epipole estimation methods with/without rank-2 constraint across algebraic/geometric error functions. In *Proc. IEEE Conf. on Comput. Vision and Pattern Recognition (CVPR)*, pages 1–7, 2007.

[22] M. Pollefeys, L. Van Gool, M. Vergauwen, F. Verbiest, K. Cornelis, J. Tops, and R. Koch. Visual modeling with a hand-held camera. *Int'l J. Comput. Vision*, 59(3):207–232, 2004.

[23] Sphinx. Photo by Heinz Albers. Available under a CC license, 2010. URL http://commons.wikimedia.org/wiki/File:Sphinx_from_Gizeh.jpg.

[24] Sphinx. Photo by ironmanixs. Available under a CC license, 2010. URL http://www.flickr.com/photos/ironmanixs/515950995/.

[25] Y. Sugaya, April 2010. URL http://www.iim.ics.tut.ac.jp/~sugaya/public-e.html.

[26] Y. Sugaya and K. Kanatani. Highest accuracy fundamental matrix computation. In *Proc. Asian Conf. on Comput. Vision (ACCV)*, pages 311–321, 2007.

[27] PHS Torr and DW Murray. The development and comparison of robust methods for estimating the fundamental matrix. *Int'l J. Comput. Vision*, 24(3):271–300, 1997.

[28] Z. Zhang and T. Kanade. Determining the epipolar geometry and its uncertainty: A review. *Int'l J. Comput. Vision*, 27(2):161–195, 1998.