

The Space of Multibody Fundamental Matrices: Rank, Geometry and Projection*

Xiaodong Fan and René Vidal

Center for Imaging Science, Johns Hopkins University
301 Clark Hall, 3400 N. Charles St., Baltimore, MD, 21218, USA

Abstract

We study the rank and geometry of the multibody fundamental matrix, a geometric entity characterizing the two-view geometry of dynamic scenes consisting of multiple rigid-body motions. We derive an upper bound on the rank of the multibody fundamental matrix that depends on the number of independent translations. We also derive an algebraic characterization of the SVD of a multibody fundamental matrix in the case of two or odd number of rigid-body motions with a common rotation. This characterization allows us to project an arbitrary matrix onto the space of multibody fundamental matrices using linear algebraic techniques.

1. Introduction

Given two perspective views of a scene containing multiple rigidly moving objects, we consider the problem of estimating the motion of each object relative to the camera, without knowing which measurements belong to which object.

When the scene is *static*, i.e., when either the camera or a single object move rigidly, it is well-known [7] that if $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{P}^2$ are two perspective images of a point in 3-D space, then they must satisfy the *epipolar constraint*

$$\mathbf{x}_2^\top F \mathbf{x}_1 = 0, \quad (1)$$

where $F \in \mathbb{R}^{3 \times 3}$ is a rank-2 matrix called the *fundamental matrix*. The epipolar constraint can be used to estimate F and the camera motion from a set of point correspondences using linear techniques such as the eight-point algorithm. In the case of a calibrated camera, it is also known that F factors as $F = [T]_\times R$, where $[T]_\times \in so(3)$ is a skew-symmetric matrix associated with the camera translation $T \in \mathbb{R}^3$ and $R \in SO(3)$ is the camera rotation. The space $so(3) \times SO(3)$ is known as the *essential manifold* and can be characterized as the space of matrices with singular values $\{\|T\|, \|T\|, 0\}$. Such a characterization is crucial when estimating F from noisy correspondences, because it allows us to project a noisy linear estimate of F onto a geometrically correct *essential matrix*.

*This work has been funded by Johns Hopkins Whiting School of Engineering startup funds and by grants NSF CAREER ISS-0447739, ONR N000140510836, NSF DMS-042723 and NSF DMS-0219016.

The work of [14] proposes a generalization of the eight-point algorithm to the more general and challenging case of *dynamic scenes* in which both the camera and an unknown number of objects with unknown 3-D structure move independently. The paper shows that applying a polynomial embedding to the image points leads to the so-called *multibody epipolar constraint* and its associated *multibody fundamental matrix* \mathcal{F} . The method computes the number of motions from a rank constraint on the image measurements, estimates the multibody fundamental matrix using least squares, and the individual fundamental matrices using multivariate polynomial factorization or differentiation.

Unfortunately, the method is not yet reliable in the presence of noise, because of the following reasons:

1. The polynomial embedding is not invariant with respect to rotations or translations of the image data, which makes it difficult to characterize the space of multibody fundamental matrices. Such a characterization is crucial for improving the performance of linear algorithms in the presence of noisy data.
2. The multibody fundamental matrix \mathcal{F} is computed *linearly*, without taking into account nonlinear constraints dictated by its rank and geometry. Therefore, the estimate of \mathcal{F} may not be geometrically correct in the presence of noise, meaning that it may not perfectly factor into the multiple fundamental matrices associated with each one of the rigid-body motions.

In this paper, we show how to overcome these difficulties by exploiting the rank and geometry of the multibody fundamental matrix. More specifically,

1. **Rank:** we show that the rank of \mathcal{F} depends on the number of independent translational motions and on the number of times they are repeated. Our results complete the analysis in [14], which deals with the particular case of one repeated translational motion.
2. **Geometry:** we show that in the case of n rigid-body motions with common rotation, \mathcal{F} factors as the product of a symmetric (n even) or skew-symmetric (n odd) matrix times a rotation matrix. When the number

of motions is two or odd, this leads to a characterization of the SVD of \mathcal{F} . This characterization is possible thanks to a slightly new definition of the polynomial embedding that makes the singular values of the multibody fundamental matrix invariant with respect to rotations of the image data.

3. **Projection:** we show that the characterization of the SVD of \mathcal{F} can be used to project an arbitrary matrix estimated from noisy correspondences onto the space of multibody fundamental matrices using linear algebraic techniques.

To the best of our knowledge, there is no prior work studying the geometry and projection onto the space of multibody fundamental matrices. In fact, finding a linear algebraic characterization of this space is an extremely challenging problem. Therefore, although the case of two or odd number of motions with common rotations may appear to be restrictive, we believe this case is an important step toward solving the general case.

Previous work. Most prior work on dynamic scene reconstruction proceeds by first segmenting image measurements into various motion models, and then estimating a single motion model for each group of measurements, or else in an iterative manner with the aid of the EM algorithm. The number of models can also be estimated in a probabilistic framework using model selection techniques such as [10, 6]. However, the convergence of iterative/probabilistic methods to the global optimum depends strongly on correct initialization [10, 9]. This has motivated the recent development of geometric approaches to dynamic scene reconstruction which do not require initialization. Algebraic approaches include methods for multiple moving objects seen by an orthographic camera [1, 5, 17, 11], self-calibration from multiple motions [2], multiple points moving in planes [8], segmentation of two [16] and multiple [14, 15] rigid-body motions from two or three [4] perspective views.

2. Multibody Epipolar Geometry

Given a set of point correspondences $\{(\mathbf{x}_1^j, \mathbf{x}_2^j)\}_{j=1}^N$ generated from n independently and rigidly moving objects, our goal is to estimate their associated fundamental matrices $\{F_i\}_{i=1}^n$ and the object to which each image pair belongs.

To this end, let $(\mathbf{x}_1, \mathbf{x}_2)$ be an arbitrary image pair associated with *any* of the n moving objects. Then, there exists a fundamental matrix $F_i \in \mathbb{R}^{3 \times 3}$ such that the *epipolar constraint* $\mathbf{x}_2^\top F_i \mathbf{x}_1 = 0$ is satisfied. Therefore, regardless of the object associated with the image pair, the following *multibody epipolar constraint* [14] must be satisfied by the fundamental matrices $\{F_i\}_{i=1}^n$ and the image pair $(\mathbf{x}_1, \mathbf{x}_2)$

$$\text{MEC}(\mathbf{x}_1, \mathbf{x}_2) \doteq \prod_{i=1}^n (\mathbf{x}_2^\top F_i \mathbf{x}_1) = 0. \quad (2)$$

The multibody epipolar constraint (MEC) is a homogeneous polynomial of degree n in each of \mathbf{x}_1 or \mathbf{x}_2 . Therefore, if we let $\mathbf{x}_1 = [x_1, y_1, z_1]^\top$, then equation (2) viewed as a function of \mathbf{x}_1 can be written as a linear combination of the following $M_n \doteq (n+1)(n+2)/2$ independent monomials $\{x_1^n, x_1^{n-1}y_1, x_1^{n-1}z_1, \dots, z_1^n\}$. After collecting all these monomials into a vector

$$\nu_n(\mathbf{x}_1) = [\dots, \gamma_{n_1, n_2, n_3} x_1^{n_1} y_1^{n_2} z_1^{n_3}, \dots]^\top \in \mathbb{R}^{M_n}, \quad (3)$$

where $\gamma_{n_1, n_2, n_3} = \sqrt{\frac{n!}{n_1! n_2! n_3!}}$ with $0 \leq n_1, n_2, n_3 \leq n$, $n_1 + n_2 + n_3 = n$, the MEC can be written as the following a bilinear expression in $\nu_n(\mathbf{x}_1)$ and $\nu_n(\mathbf{x}_2)$ (see [14]):

$$\nu_n(\mathbf{x}_2)^\top \mathcal{F} \nu_n(\mathbf{x}_1) = 0. \quad (4)$$

The matrix $\mathcal{F} \in \mathbb{R}^{M_n \times M_n}$ is called the *multibody fundamental matrix*, and is a natural generalization of the fundamental matrix $F \in \mathbb{R}^{3 \times 3}$ to the case of n moving objects. The embedding $\nu_n : \mathbb{R}^3 \rightarrow \mathbb{R}^{M_n}$ is known in algebraic geometry as the Veronese map of degree n [3].

Remark 1 (Rotation invariant) *Notice that our definition of the Veronese map is slightly different from the one in [14], as we deliberately multiply the monomial $x_1^{n_1} y_1^{n_2} z_1^{n_3}$ by the coefficient γ_{n_1, n_2, n_3} . As we will show in Theorem 2, this new definition of the Veronese map makes it rotation invariant, a property that will be shown to be crucial for characterizing the space of multibody fundamental matrices.*

Thanks to the Veronese map, we can write the epipolar constraint for all N point correspondences as

$$\mathbf{V}_n \mathbf{f} \doteq [\nu_n(\mathbf{x}_2^1) \otimes \nu_n(\mathbf{x}_1^1) \cdots \nu_n(\mathbf{x}_2^N) \otimes \nu_n(\mathbf{x}_1^N)]^\top \mathbf{f} = \mathbf{0}, \quad (5)$$

where $\mathbf{f} \in \mathbb{R}^{M_n^2}$ is the stack of the rows of \mathcal{F} and \otimes represents the Kronecker product. Given \mathcal{F} , which can be computed as the least squares solution of (5), the individual fundamental matrices $\{F_i\}_{i=1}^n$ are obtained by factorizing the bi-homogeneous polynomial

$$\nu_n(\mathbf{x}_2)^\top \mathcal{F} \nu_n(\mathbf{x}_1) = \prod_{i=1}^n (\mathbf{x}_2^\top F_i \mathbf{x}_1) = 0 \quad (6)$$

into a product of bilinear forms [14], or from the second order derivatives of the MEC [12].

Notice that the multibody fundamental matrix \mathcal{F} is determined by the fundamental matrices of the individual rigid motions $\{F_i\}_{i=1}^n$. Since these fundamental matrices are of rank two and/or belong to the essential manifold, the multibody fundamental matrix is not an arbitrary matrix in $\mathbb{R}^{M_n \times M_n}$, but must satisfy some nonlinear constraints, such as rank constraints and/or geometric constraints. Such constraints are clearly not exploited by the linear algorithm of

[14]. Therefore, the linear estimate of the multibody fundamental matrix may not be geometrically correct in the presence of noise, meaning that its associated MEC may not perfectly factor as a product of epipolar constraints.

Such problems motivate our development in the rest of this paper.

3. Rank of the Multibody Fundamental Matrix

It is well-known [7] that the rank of a fundamental matrix F is two. The vector e in its left null space is called the *epipole* and satisfies the following relationship $e^\top F = 0$.

In the case of n rigid-body motions, there exist n epipoles $\{e_i\}_{i=1}^n$ such that $e_i^\top F_i = 0$. This implies that

$$(e_1^\top F_1 \mathbf{x}) (e_2^\top F_2 \mathbf{x}) \cdots (e_n^\top F_n \mathbf{x}) = \nu_n(e_i)^\top \mathcal{F} \nu_n(\mathbf{x}) = 0,$$

for all $\mathbf{x} \in \mathbb{P}^2$. Since the vector $\nu_n(\mathbf{x})$ spans all of \mathbb{R}^{M_n} when \mathbf{x} ranges over \mathbb{P}^2 ,¹ we immediately have [14]

$$\nu_n(e_i)^\top \mathcal{F} = 0 \quad \text{for } i = 1, \dots, n. \quad (7)$$

Therefore, the multibody fundamental matrix \mathcal{F} is also rank deficient, because the n embedded epipoles $\{\nu_n(e_i)\}_{i=1}^n$ lie in its left null space. Notice, however, that the dimension of the null space of \mathcal{F} need not be n , because the embedded epipoles may not be linearly independent. For instance, if two different rigid-body motions have the same translation, but different rotation, then they have the same epipole, hence the same embedded epipole.

The purpose of this section is to characterize the null space of \mathcal{F} as a function of the number of motions n , the number of different epipoles $n_e \leq n$ (different up to a scale factor) and the number of times $\{k_i\}_{i=1}^{n_e}$, with $\sum_{i=1}^{n_e} k_i = n$, that each epipole is repeated.² More specifically, we prove the following theorem.

Theorem 1 (Null space of \mathcal{F}) *Let \mathcal{F} be the multibody fundamental matrix generated by n fundamental matrices. Let n_e be the number of different epipoles and $k_i, i = 1, \dots, n_e$, be the number of times each different epipole is repeated. The rank of the multibody fundamental matrix is bounded by*

$$\text{rank}(\mathcal{F}) \leq M_n - \sum_{i=1}^{n_e} M_{k_i-1} \leq M_n - n, \quad (8)$$

where the inequality on the right hand side is true regardless of whether the epipoles are repeated or not.

¹This is simply because the M_n monomials in $\nu_n(\mathbf{x})$ are linearly independent.

²The particular case in which one epipole is repeated k times, and the other $n - k$ epipoles are different can be found in [14].

The formal proof of the theorem is organized as follows. In Section 3.1, we show that if an epipole e_i is repeated k_i times, then all the derivatives of ν_n of order less than k_i evaluated at e_i lie in the left null space of \mathcal{F} . In Section 3.2, we show that only M_{k_i-1} of these derivatives are linearly independent, thus each different epipole contributes with an M_{k_i-1} -dimensional subspace to $\text{null}(\mathcal{F})$. In Section 3.3 we show that these n_e subspaces are independent, meaning that they intersect only at $\mathbf{0}$. Therefore, the dimensionality of the null space of \mathcal{F} is at least $\sum_{i=1}^{n_e} M_{k_i-1} \geq n$.

3.1. Partial derivatives at repeated epipoles

In this subsection, we show that when an epipole e_i is repeated k_i times, not only $\nu_n(e_i)$ is in the null space of \mathcal{F} , as shown by equation (7), but also the derivatives of $\nu_n(\mathbf{x})$ of order less than k_i at e_i . Before proving this, we need the following technical lemma, which allows us to express the derivatives of the n th order MEC as a linear combination of MECs of lower order.

Lemma 1 *Let $\mathcal{F}^{(n)}$ be the multibody fundamental matrix generated by F_1, \dots, F_n . Let $\mathcal{F}_j^{(n-l)}$ be a multibody fundamental matrix generated by a choice of $n - l$ out of the n fundamental matrices for $j = 1, \dots, \binom{n}{l}$. Then $\forall (l_1, l_2, l_3)$, such that $l_1 + l_2 + l_3 = l, \forall \mathbf{x} = [x, y, z]^\top, \forall \mathbf{y} \in \mathbb{P}^2$, we have*

$$\frac{\partial^l (\nu_n(\mathbf{x})^\top \mathcal{F}^{(n)} \nu_n(\mathbf{y}))}{\partial x^{l_1} \partial y^{l_2} \partial z^{l_3}} = \sum_{j=1}^{\binom{n}{l}} \alpha_j \nu_{n-l}(\mathbf{x})^\top \mathcal{F}_j^{(n-l)} \nu_{n-l}(\mathbf{y}),$$

where the coefficient $\alpha_j \in \mathbb{R}$ depends on $\mathcal{F}^{(n)}$ and \mathbf{y} , but is independent of \mathbf{x} .

We are now ready to show that the derivatives of ν_n at a repeated epipole lie in the left null space of \mathcal{F} .

Lemma 2 *If $e_i \in \mathbb{P}^2$ is an epipole that is repeated k_i times, and $\mathbf{x} = [x, y, z]^\top$, then $\forall (l_1, l_2, l_3)$, such that $l_1 + l_2 + l_3 = l \leq k_i - 1$, we have*

$$\left. \frac{\partial^l \nu_n(\mathbf{x})^\top}{\partial x^{l_1} \partial y^{l_2} \partial z^{l_3}} \right|_{e_i} \mathcal{F} = \mathbf{0}. \quad (9)$$

Proof. Since e_i is repeated k_i times, there are k_i fundamental matrices whose left null space is e_i . Then any choice of $n - l$ fundamental matrices with $l \leq k_i - 1$ will contain at least one fundamental matrix whose left null space is e_i . From (7) we have that e_i is an epipole for each one of the multibody fundamental matrices $\mathcal{F}_j^{(n-l)}$ with $l \leq k_i - 1$, i.e., $\nu_{n-l}(e_i)^\top \mathcal{F}_j^{(n-l)} = \mathbf{0}$. This, together with Lemma 1, implies that for all $\mathbf{y} \in \mathbb{P}^2$ and for all (l_1, l_2, l_3) such that $l_1 + l_2 + l_3 = l \leq k_i - 1$

$$\left. \frac{\partial^l \nu_n(\mathbf{x})^\top}{\partial x^{l_1} \partial y^{l_2} \partial z^{l_3}} \right|_{e_i} \mathcal{F} \nu_n(\mathbf{y}) = \mathbf{0}.$$

Since this is true for all $\mathbf{y} \in \mathbb{P}^2$, the claim follows. ■

3.2. Dimension of the subspaces spanned by the partial derivatives

In this subsection, we show that an epipole repeated k_i times contributes to the null space of \mathcal{F} with a subspace of dimension at least M_{k_i-1} . The result is a consequence of the following facts: 1) the subspace spanned by the partial derivatives of order l is included in any of the subspaces spanned by higher order partial derivatives; and 2) the dimension of the subspace spanned by the derivatives of order l is M_l .

As for the first fact, notice that each entry of $\nu_n(\mathbf{x})$ is of the form $\gamma_{n_1, n_2, n_3} x^{n_1} y^{n_2} z^{n_3}$ with $n_1 + n_2 + n_3 = n$. After some simple algebraic calculations, we can show that

$$(n-l) \frac{\partial^l \nu_n(\mathbf{x})}{\partial x^{l_1} \partial y^{l_2} \partial z^{l_3}} \quad (10)$$

$$= \left[\frac{\partial^{l+1} \nu_n(\mathbf{x})}{\partial x^{l_1+1} \partial y^{l_2} \partial z^{l_3}}, \frac{\partial^{l+1} \nu_n(\mathbf{x})}{\partial x^{l_1} \partial y^{l_2+1} \partial z^{l_3}}, \frac{\partial^{l+1} \nu_n(\mathbf{x})}{\partial x^{l_1} \partial y^{l_2} \partial z^{l_3+1}} \right] \mathbf{x}.$$

Therefore, if we let $A_l(\mathbf{x})$ be the span of the l th order partial derivatives of $\nu_n(\mathbf{x})$, then (10) implies that $A_l(\mathbf{x}) \subseteq A_{l+1}(\mathbf{x})$ for all $0 \leq l < n$. By simple induction we have that if \mathbf{e}_i is an epipole that is repeated k_i times, then

$$A_0(\mathbf{e}_i) \subseteq A_1(\mathbf{e}_i) \subseteq \dots \subseteq A_{k_i-1}(\mathbf{e}_i). \quad (11)$$

As a consequence of (11), studying the dimension of the subspace spanned by all the partial derivatives at a repeated epipole *up to* a certain order, boils down to finding the dimension of the subspace spanned by the partial derivatives of *exactly* that order. The following lemma shows that all the derivatives of a fixed order are linearly independent, hence the dimension of $A_l(\mathbf{x})$ is M_l .

Lemma 3 For $\mathbf{x} \in \mathbb{P}^2$ and $l < n$, all the l th order partial derivatives $\left\{ \frac{\partial^l \nu_n(\mathbf{x})}{\partial x^{l_1} \partial y^{l_2} \partial z^{l_3}} \right\}_{l_1+l_2+l_3=l}$ are linearly independent. Hence, the dimension of $A_l(\mathbf{x})$ is M_l .

Proof. Our goal is to show that

$$\sum_{l_1+l_2+l_3=l} \alpha_{l_1, l_2, l_3} \frac{\partial^l \nu_n(\mathbf{x})}{\partial x^{l_1} \partial y^{l_2} \partial z^{l_3}} = \mathbf{0} \quad (12)$$

if and only if $\alpha_{l_1, l_2, l_3} = 0$ for all (l_1, l_2, l_3) such that $l_1 + l_2 + l_3 = l$. Since $\mathbf{x} = [x, y, z]^T \neq \mathbf{0}$, without loss of generality let us assume that $x \neq 0$. Notice each entry of $\frac{\partial^l \nu_n(\mathbf{x})}{\partial x^{l_1} \partial y^{l_2} \partial z^{l_3}}$ is of the form $x^{n_1-l_1} y^{n_2-l_2} z^{n_3-l_3}$. Hence the first entry of (12) has the form $\alpha_{l, 0, 0} x^{n-l} = 0$, and therefore $\alpha_{l, 0, 0} = 0$. By sequentially applying the same reasoning to entries of $\nu_n(\mathbf{x})$ of the form $x^{n_1} y^{n_2} z^{n_3}$, where $(n_1, n_2, n_3) = (n+l_1-l, l_2, l_3)$ for $l_1 = l-2, l-3, \dots, 0$, we obtain $\alpha_{l_1, l_2, l_3} x^{n-l} = 0$, and so $\alpha_{l_1, l_2, l_3} = 0$ as claimed. ■

3.3. Independence of subspaces corresponding to different epipoles

In this subsection, we show that the subspaces associated with different epipoles are independent, in the sense that they intersect only at $\mathbf{0}$. Therefore, the dimension of the left null space of \mathcal{F} , which contains the union of the subspaces associated with each one of the n_e different epipoles, is lower bounded by the sum of the dimensions of these subspaces. The main result is summarized in Lemma 4.

Lemma 4 Given two different epipoles \mathbf{e}_1 and \mathbf{e}_2 that are repeated k_1 and k_2 times, respectively, the span of the partial derivatives at \mathbf{e}_1 and \mathbf{e}_2 intersect only at $\mathbf{0}$, i.e.,

$$A_{k_1-1}(\mathbf{e}_1) \cap A_{k_2-1}(\mathbf{e}_2) = \{\mathbf{0}\}. \quad (13)$$

This completes the proof of the rank constraint on the multibody fundamental matrix $\text{rank}(\mathcal{F}) \leq M_n - \sum_{i=1}^{n_e} M_{k_i-1}$. Furthermore, because $\sum_{i=1}^{n_e} M_{k_i-1} \geq n$ when $\sum_{i=1}^{n_e} k_i = n$, we immediately know that $\text{rank}(\mathcal{F}) \leq M_n - n$ regardless of whether the epipoles are repeated or not.

4. Geometry of the Space of Multibody Fundamental Matrices

Recall from Section 2 that given enough point correspondences, one can compute the corresponding multibody fundamental matrix \mathcal{F} by solving the linear system $\mathbf{V}_n \mathbf{f} = \mathbf{0}$ in (5). With perfect data, the linearly estimated \mathcal{F} will automatically satisfy the rank constraints studied in the previous section. However, with noisy data the so-computed \mathcal{F} may not be geometrically correct, because the rank constraints are not taken into account.

In this section, we propose to enforce these constraints by first estimating the null space of \mathbf{V}_n *ignoring the internal algebraic structure* of \mathcal{F} , and then *projecting* the matrix thus obtained onto the space of multibody fundamental matrices. Our analysis applies to both the uncalibrated case, in which we only need to enforce rank constraints on \mathcal{F} , as well as the calibrated case, in which the singular values of \mathcal{F} must satisfy additional constraints due to the geometry of $so(3) \times SO(3)$. First, we introduce some invariant properties of the Veronese map and multibody fundamental matrix. Then, we use these properties to characterize the singular values of \mathcal{F} under some constrained scenarios. Later, we propose a linear algebraic technique to project the linearly estimated \mathcal{F} onto the multibody fundamental space, by exploiting both the rank and singular value constraints.

4.1. Invariance properties of the Veronese map

Before studying the geometry of the multibody fundamental matrix, let us first explore some invariance properties of the Veronese map, which will be important for the theoretical development in the following subsections.

Theorem 2 *The Veronese map as defined in (3) has the following properties for all $\mathbf{x}, \mathbf{y} \in \mathbb{P}^2$:*

- **Inner product invariance:** $\nu_n(\mathbf{y})^\top \nu_n(\mathbf{x}) = (\mathbf{y}^\top \mathbf{x})^n$.
- **Linear invariance:** For all $A \in \mathbb{R}^{3 \times 3}$ there exists an $A \in \mathbb{R}^{M_n \times M_n}$ such that for all \mathbf{x} , $\nu_n(A\mathbf{x}) = A\nu_n(\mathbf{x})$.
- **Rotation invariance:** For all $R \in SO(3)$ there exists $\mathcal{R} \in SO(M_n)$ such that for all \mathbf{x} , $\nu_n(R\mathbf{x}) = \mathcal{R}\nu_n(\mathbf{x})$.

Note that the rotation invariance property of the Veronese map implies that if the image measurements $\{(\mathbf{x}_1^j, \mathbf{x}_2^j)\}_{j=1}^N$ are related by a multibody fundamental matrix \mathcal{F} , then the rotated image measurements $\{(R_1\mathbf{x}_1^j, R_2\mathbf{x}_2^j)\}_{j=1}^N$, where $R_1, R_2 \in SO(3)$, are related by a multibody fundamental matrix $\mathcal{F}' = \mathcal{R}_2^\top \mathcal{F} \mathcal{R}_1$, where $\mathcal{R}_1, \mathcal{R}_2 \in SO(M_n)$. This is because

$$\nu_n(R_2\mathbf{x}_2)^\top \mathcal{F}' \nu_n(R_1\mathbf{x}_1) = \nu_n(\mathbf{x}_2)^\top \mathcal{R}_2^\top \mathcal{F} \mathcal{R}_1 \nu_n(\mathbf{x}_1).$$

Therefore, \mathcal{F} and \mathcal{F}' share the same singular values. This property is crucial for characterizing the singular values of \mathcal{F} , as we show in the next subsection.

4.2. SVD of the multibody essential matrix

In the case of one motion, if we further assume that the camera calibration parameters are known, then F is usually called the *essential matrix* and can be expressed as [7]:

$$F = [T]_\times R, \quad \text{with } [T]_\times \in so(3), \quad R \in SO(3), \quad (14)$$

where $[T]_\times \in se(3)$ is a skew-symmetric matrix generating the cross product by $T \in \mathbb{R}^3$. This property allows us to characterize the singular values of F as $\{\|T\|, \|T\|, 0\}$.

In this subsection, we aim to generalize this result to n rigid-body motions by characterizing the singular values of the multibody essential matrix \mathcal{F} . To the best of our knowledge, there is no prior work addressing this problem, which we believe to be very challenging. Therefore, we restrict our attention to the case of rigid-body motions with common rotation. This case shows up, e.g., when a rigidly moving camera observes multiple translating objects. We show that in this case \mathcal{F} can be written as the product of a symmetric (n even) or skew-symmetric (n odd) matrix with a rotation matrix as stated by the following theorem.

Theorem 3 *Let $\{(R, T_i) \in SE(3)\}_{i=1}^n$ be n independent rigid-body motions sharing a common rotation matrix R . Their multibody essential matrix \mathcal{F} can be expressed as*

$$\mathcal{F} = \mathcal{T} \mathcal{R}, \quad (15)$$

where $\mathcal{T} \in \mathbb{R}^{M_n \times M_n}$ is a multibody fundamental matrix corresponding to purely translational motions $\{T_i\}_{i=1}^n$ which is either symmetric (n even) or skew-symmetric (n odd), and $\mathcal{R} \in SO(M_n)$ is a rotation matrix in \mathbb{R}^{M_n} .

Proof. Notice that $F_i = [T_i]_\times R$, hence for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{P}^2$, and $\mathbf{x}'_1 = R\mathbf{x}_1$, the multibody epipolar constraint can be written as

$$\begin{aligned} \nu_n^\top(\mathbf{x}_2) \mathcal{F} \nu_n(\mathbf{x}_1) &= \prod_{i=1}^n (\mathbf{x}_2^\top [T_i]_\times R \mathbf{x}_1) = \prod_{i=1}^n (\mathbf{x}_2^\top [T_i]_\times \mathbf{x}'_1) = \\ \nu_n^\top(\mathbf{x}_2) \mathcal{T} \nu_n(\mathbf{x}'_1) &= \nu_n^\top(\mathbf{x}_2) \mathcal{T} \nu_n(R\mathbf{x}_1) = \nu_n^\top(\mathbf{x}_2) \mathcal{T} \mathcal{R} \nu_n(\mathbf{x}_1), \end{aligned}$$

where the last step follows from the rotation invariance property of the Veronese map. Therefore, $\mathcal{F} = \mathcal{T} \mathcal{R}$, where $\mathcal{T} \in \mathbb{R}^{M_n \times M_n}$ is a multibody fundamental matrix corresponding to purely translational motions $\{T_i\}_{i=1}^n$, as claimed. Furthermore, note that \mathcal{T} is the symmetric tensor product of n essential matrices associated with n purely translational motions. Since such essential matrices are skew-symmetric, we have that $\mathcal{T}^\top = (-1)^n \mathcal{T}$, hence \mathcal{T} is symmetric when n is even and skew-symmetric otherwise. ■

Thanks to Theorem 3, we can characterize the SVD of a multibody essential matrix with a common rotation for an odd number of motions, as stated by the following theorem.

Theorem 4 (Singular values of a multibody essential matrix with odd number of motions) *Let n be an odd number of independent rigid-body motions $\{R, T_i\}_{i=1}^n$ with a common rotation $R \in SO(3)$. The corresponding multibody essential matrix has a SVD $\mathcal{F} = U \Sigma V^\top$, with*

$$\Sigma = \text{diag}\{\sigma_1, \sigma_1, \sigma_2, \sigma_2, \dots, \sigma_m, \sigma_m, 0, \dots, 0\}, \quad (16)$$

where $\sigma_1 \geq \dots \geq \sigma_m \geq 0$ and $m = \lfloor \frac{M_n - n}{2} \rfloor$, where $\lfloor x \rfloor$ is the largest integer that is less than or equal to $x \in \mathbb{R}$.

Proof. Let \mathcal{T} be the multibody fundamental matrix corresponding to n translations $\{T_i\}_{i=1}^n$. Based on Theorem 3, \mathcal{T} shares the same singular values with \mathcal{F} . \mathcal{T} is a skew-symmetric matrix when n is odd, and it is well-known that all non-zero singular values of a skew-symmetric matrix must appear in pairs. Furthermore, by Theorem 1, the rank of \mathcal{T} is upper bounded by $M_n - n$, which means that \mathcal{T} has at least n or $n+1$ zero singular values, depending on whether $M_n - n$ is a multiple of two or not, respectively. ■

Unfortunately, the above singular value characterization does not generalize to an even number of motions, because in this case the multibody essential matrix for purely translational motions is symmetric. However, when the number of independent motions is two, we are still able to completely specify the singular values of the multibody essential matrix, though using another method. More precisely, we have the following result.

Theorem 5 (Singular values of two-body essential matrix)

Let \mathcal{F} be the multibody essential matrix corresponding to two independent motions (R, T_1) and (R, T_2) with a common rotation R . Its singular values $\sigma_1 \geq \dots \geq \sigma_6$ are

$$\begin{cases} \sigma_1 = \sigma_2 = \frac{\sqrt{2\|T_1\|^2\|T_2\|^2 + 2(T_1 \cdot T_2)^2}}{2} \\ \sigma_3 = \frac{\|T_1\|\|T_2\|}{2} - \frac{T_1 \cdot T_2}{2}, \quad \sigma_4 = \frac{\|T_1\|\|T_2\|}{2} + \frac{T_1 \cdot T_2}{2} \\ \sigma_5 = \sigma_6 = 0 \end{cases} \quad (17)$$

Furthermore, $\sigma_1^2 = \sigma_2^2 = \sigma_3^2 + \sigma_4^2$.

Proof. From Theorem 3, it is sufficient to characterize the singular values of the multibody essential matrix \mathcal{T} corresponding to two translational motions T_1 and T_2 . To this end, let R_0 be a rotation matrix that maps T_1 to $T'_1 = R_0 T_1 = \|T_1\| [1, 0, 0]^\top$, and let $T'_2 = R_0 T_2$. Then the multibody fundamental matrix associated with the two translational motions T'_1 and $T'_2 = [u, s, v]^\top$

$$T' = \|T_1\| \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{v}{2} & -\frac{\sqrt{2}u}{2} \\ 0 & 0 & 0 & -\frac{\sqrt{2}v}{2} & \frac{u}{2} & 0 \\ 0 & 0 & -\frac{\sqrt{2}v}{2} & 0 & 0 & 0 \\ 0 & \frac{v}{2} & \frac{u}{2} & 0 & -s & 0 \\ 0 & -\frac{\sqrt{2}u}{2} & 0 & s & 0 & 0 \end{bmatrix}$$

is much simpler, yet shares the same SVD with \mathcal{T} and \mathcal{F} due to the rotation invariance property. The proof of the theorem follows by direct calculation of the singular values of T' . ■

4.3. Projection onto multibody essential space

Given a characterization of the space of multibody essential matrices, our remaining task is to enforce these constraints in the estimation of \mathcal{F} . We achieve this by projecting the linearly estimated multibody fundamental matrix onto the multibody essential manifold. The projection consists of two main steps. First (Theorem 6), we show that the closest matrix B (in Frobenius norm) to an arbitrary matrix A can be obtained by minimizing the sum-of-squares distance among their corresponding singular values. Second (Theorems 7 and 8) we show how to find the optimal singular values for each one of the characterizations of the multibody essential matrix.

Theorem 6 *Let the singular values of $A, B \in \mathbb{R}^{m \times m}$ be $\sigma_1(A) \geq \dots \geq \sigma_m(A)$ and $\sigma_1(B) \geq \dots \geq \sigma_m(B)$. Then*

$$\min_{U^\top U=I, V^\top V=I} \|A - UBV^\top\|_f^2 = \sum_{i=1}^m (\sigma_i(A) - \sigma_i(B))^2,$$

where $\|A\|_f = \sqrt{\text{trace}(A^\top A)}$ is the Frobenius norm. Furthermore, if $A = U_0 \text{diag}\{\sigma_1(A), \dots, \sigma_m(A)\} V_0^\top$, then the minimizing matrices U^* and V^* are such that $U^* B V^{*\top} = U_0 \text{diag}\{\sigma_1(B), \dots, \sigma_m(B)\} V_0^\top$.

Based on this projection theorem, one can find a matrix B with constrained singular values that is closest in Frobenius norm to an arbitrary matrix A by replacing the singular values of A by those with the desired structure.

In the case of uncalibrated cameras, as shown in Theorem 1, the rank of the multibody fundamental matrix \mathcal{F} is upper bounded by $M_n - n$ when there are n independent motions, hence \mathcal{F} has at least n zero singular values. Such a rank constraint can be enforced by simply setting the n smallest singular values of the estimated \mathcal{F} to 0.

In the case of calibrated cameras, the next two theorems show how to enforce both rank and singular value constraints for the motions with common rotation. The proof of the theorems is very straightforward: We solve an optimization problem in the spirit of Theorem 6 with constraints as specified in Theorems 4 and 5 respectively, using the method of Lagrange multipliers.

Theorem 7 (Projection onto the n -body essential space for an odd number of motions with common rotation)

Let $\hat{\mathcal{F}} \in \mathbb{R}^{M_n \times M_n}$ be the estimate of an n -body essential matrix with common rotation and n odd. Let the SVD of $\hat{\mathcal{F}}$ be $\hat{\mathcal{F}} = U_0 \text{diag}\{\hat{\sigma}_1, \dots, \hat{\sigma}_{M_n}\} V_0^\top$, $\hat{\sigma}_1 \geq \dots \geq \hat{\sigma}_{M_n}$. The n -body essential matrix \mathcal{F} which minimizes the error $\|\hat{\mathcal{F}} - \mathcal{F}\|_f$ is given by $\mathcal{F} = U_0 \text{diag}\{\sigma_1, \dots, \sigma_{M_n}\} V_0^\top$, where $\sigma_{2i-1} = \sigma_{2i} = \frac{\hat{\sigma}_{2i-1} + \hat{\sigma}_{2i}}{2}$ for $1 \leq i \leq \lfloor \frac{M_n - n}{2} \rfloor$ and 0 otherwise.

Theorem 8 (Projection onto two-body essential space)

Let $\hat{\mathcal{F}} \in \mathbb{R}^{6 \times 6}$ be the estimate of a two-body essential matrix with common rotation. Let the SVD of $\hat{\mathcal{F}}$ be $\hat{\mathcal{F}} = U_0 \text{diag}\{\hat{\sigma}_1, \dots, \hat{\sigma}_6\} V_0^\top$, where $\hat{\sigma}_1 \geq \dots \geq \hat{\sigma}_6$. The two-body essential matrix \mathcal{F} which minimizes the error $\|\hat{\mathcal{F}} - \mathcal{F}\|_f$ is given by $\mathcal{F} = U_0 \text{diag}\{\sigma_1, \dots, \sigma_6\} V_0^\top$, where $\sigma_1 = \sigma_2 = \beta \sqrt{\hat{\sigma}_3^2 + \hat{\sigma}_4^2}$, $\sigma_3 = \beta \hat{\sigma}_3$, $\sigma_4 = \beta \hat{\sigma}_4$, $\sigma_5 = \sigma_6 = 0$ and $\beta = \frac{1}{3} \left(\frac{\hat{\sigma}_1 + \hat{\sigma}_2}{\sqrt{\hat{\sigma}_3^2 + \hat{\sigma}_4^2}} + 1 \right)$.

5. Experiments

In our experiments, we compare the following algorithms:

1. *GPCA*: The Generalized Principal Component Analysis (GPCA) method [12, 13] is specifically designed for purely translational motions. In this case, the epipolar constraint reduces to the linear equation

$$\mathbf{x}_2^\top [T_i]_\times \mathbf{x}_1 = T_i^\top (\mathbf{x}_2 \times \mathbf{x}_1) = T_i^\top \boldsymbol{\ell} = 0, \quad (18)$$

where $\boldsymbol{\ell} = (\mathbf{x}_2 \times \mathbf{x}_1) \in \mathbb{R}^3$. Consequently, the segmentation of 3-D translational motions is equivalent to clustering data $\{\boldsymbol{\ell}\}$ lying on a collection of hyperplanes in \mathbb{R}^3 whose normal vectors are $\{T_i\}_{i=1}^n$.

2. *Multibody epipolar constraint without projection (MEC-noprojection)*: This method is described in [14]. Based on the multibody epipolar constraint (MEC) in (4), the method estimates the multibody fundamental matrix \mathcal{F} linearly by solving (5). Then, it computes the epipolar line in the second view ℓ_2^j associated with each image pair $(\mathbf{x}_1^j, \mathbf{x}_2^j)$ from the derivatives of the MEC, and obtains the epipole in the second view \mathbf{e}_i for each independent motion by clustering all the epipolar lines using the GPCA algorithm [13]. The segmentation of the image data is obtained by assigning image pair $(\mathbf{x}_1^j, \mathbf{x}_2^j)$ to the i th motion whose associated

epipole e_i is closest to the epipolar line ℓ_2^j . Finally, the single-body fundamental matrix F_i and the rigid-body motions $\{(R_i, T_i)\}_{i=1}^n$ are computed using the standard 8-point algorithm.

3. *Multibody epipolar constraint with projection (MEC-projection)*: This algorithm is essentially the same as MEC-noprojection, except for two main differences. First, the linearly estimated multibody fundamental matrix \mathcal{F} is projected onto the multibody essential manifold by enforcing the rank and singular value constraint, as described in Section 4.3. Second, the segmentation of the image features is obtained by first computing the epipole e^j for each image pair $(\mathbf{x}_1^j, \mathbf{x}_2^j)$ by a polynomial differentiation method as described in [14], and then assigning the image pair $(\mathbf{x}_1^j, \mathbf{x}_2^j)$ to the i th motion whose associated epipole e_i is closest to the epipole e^j of $(\mathbf{x}_1^j, \mathbf{x}_2^j)$.

Synthesized data: First, we evaluate the performance of our motion segmentation algorithm as a function of the amount of noise in the image measurements on the synthesized data. More specifically, we randomly pick $n = 2$ collections of $N = 100$ feature points each and apply a different rigid-body motion $(R, T_i) \in SE(3)$, with $R \in SO(3)$ the rotation and $T_i \in \mathbb{R}^3$ the translation ($i = 1, \dots, n$). Zero-mean Gaussian noise with standard deviation (std) from 0 to 2.0 pixels is added to the images \mathbf{x}_1 and \mathbf{x}_2 independently. The image size is 500×500 pixels, and we run 500 trials for each noise level. For each trial, the error between the true motions $\{(R_i, T_i)\}_{i=1}^n$ and the estimates $\{(\hat{R}_i, \hat{T}_i)\}_{i=1}^n$ is computed as³

$$\begin{aligned} \text{Translation error} &= \frac{1}{n} \sum_{i=1}^n \text{acos} \left(\frac{T_i^\top \hat{T}_i}{\|T_i\| \|\hat{T}_i\|} \right) \text{ (degrees)}. \\ \text{Rotation error} &= \frac{1}{n} \sum_{i=1}^n \text{acos} \left(\frac{\text{trace}(R_i \hat{R}_i^\top) - 1}{2} \right) \text{ (degrees)}. \end{aligned}$$

The segmentation error is estimated by computing the percentage of incorrectly classified feature points.

The left column of Figures 1(a)-(c) plots the mean error in translation, rotation and segmentation as a function of noise (standard deviation in pixels), in the case of two purely translational motions with translation directions $T_1 = [1, 0, 1]^\top$ and $T_2 = [1, 0, -1]^\top$. We also generated data undergoing two rigid-body motions with translations T_1 and T_2 and a common rotation R . The rotation axis is chosen at random, and the rotation angle is gradually increased from 0 to 30 degrees. The noise level is fixed at 2.0 pixel standard deviation. Similarly to the pure translational case, we plot the mean error in translation, rotation and segmentation as a function of the rotation angle for the GPCA and MEC-based algorithms in the right column of Figure 1.

³We do not compute the rotation error for the GPCA method because it assumes that each rigid-body motion is pure translational.

Notice that in the case of purely translational motions, the GPCA-based algorithm outperforms the MEC-based algorithms. This is because GPCA is specifically designed for purely translational motions. However, notice that as the amount of rotation increases, the performance of GPCA deteriorates very quickly, while the MEC-based algorithms have almost constant error. This is expected, as the GPCA algorithm is only applicable to translational motions.

Comparing the result of MEC-projection to that of MEC-noprojection, we can clearly observe an improvement in the recovered rigid-body motion in the sense of reducing the rotation and translation errors. However, on average, the error in the segmentation of the point correspondences increases with the projection. This is because our method minimizes the sum of the squares of the MECs, which does not depend on the segmentation. Therefore, the projection step is meant to improve the estimation of motion parameters, which does not necessarily guarantee that feature classification will improve.

Real images: We also evaluate the performance of our motion segmentation algorithm on 22 pairs of real images taken from the three-car sequence shown in Figure 2. Originally, there are three independent motions between every two frames, among which the first two are pure translational. In order to use the MEC-projection algorithm, we artificially create a situation where there is a common rotation between every pair of frames as follows: Let $\{(\mathbf{x}_1^j, \mathbf{x}_2^j)\}$ be a set of point correspondences between two frames, and \mathcal{J}_1 , \mathcal{J}_2 and \mathcal{J}_3 be three sets of indices of feature points on the three independently moving objects respectively. We first compute the rigid-body motion (R, T_3) of the third object from its own image correspondences $\{(\mathbf{x}_1^j, \mathbf{x}_2^j) : j \in \mathcal{J}_3\}$ using the standard eight-point algorithm, i.e.,

$$K \mathbf{x}_2^j \sim R K \mathbf{x}_1^j + T_3, \quad j \in \mathcal{J}_3,$$

where K is the camera calibration matrix. Because the first and second motions are pure translational, we know that

$$K \mathbf{x}_2^j \sim K \mathbf{x}_1^j + T_i, \quad j \in \mathcal{J}_i \quad \text{and} \quad i \in \{1, 2\}.$$

Therefore, if we let $\tilde{\mathbf{x}}_1^j = K^{-1} R^\top K \mathbf{x}_1^j$ ($j \in \mathcal{J}_1 \cup \mathcal{J}_2$), clearly the rotated point correspondences $(\tilde{\mathbf{x}}_1^j, \mathbf{x}_2^j)$ ($j \in \mathcal{J}_i$) undergo a new rigid-body motion (R, T_i) , where $i \in \{1, 2\}$.

In particular, we collect the point correspondences $\{(\tilde{\mathbf{x}}_1^j, \mathbf{x}_2^j) : j \in \mathcal{J}_1\} \cup \{(\mathbf{x}_1^j, \mathbf{x}_2^j) : j \in \mathcal{J}_3\}$, which undergo a two-body motion with a common rotation following the above reasoning, and apply our motion segmentation algorithm. In Figure 2, the results as the segmentation error of MEC-projection and MEC-nonprojection are compared at 22 pairs of frames. We can see that projection improves the segmentation results significantly in most cases. This is encouraging as it suggests that a better estimation of motion parameters due to the projection indeed improves the feature classification in some situations.

6. Conclusions

We have presented a new approach for the analysis of dynamic scenes containing multiple rigidly moving objects. Our approach is based on a characterization of the space of multibody fundamental matrices in terms of its rank and geometry, which can be used to obtain a more robust estimate of the multibody fundamental matrix in the presence of noise via a suitable projection. Our characterization is restricted to the case of two or odd number of rigid-body motions with a common rotation. It remains open how to characterize the space of multibody fundamental matrices with different rotations or with an even number of motions.

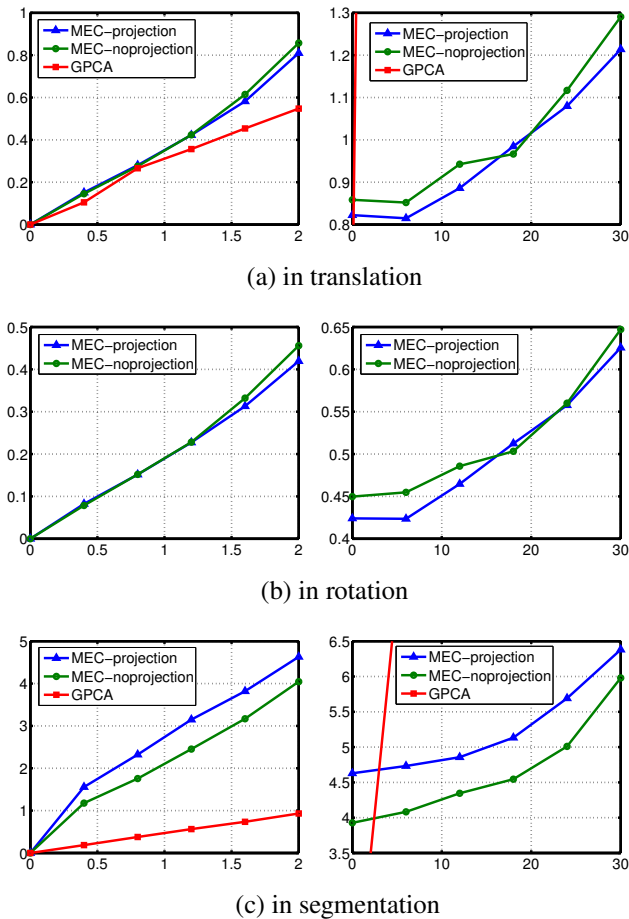


Figure 1: Error in the estimation of the translation, rotation (degrees) and segmentation (percentage). In the left column, there are two purely translational motions, with the x -axis indicating the standard deviation of the Gaussian additive noise in the image points (in pixels); in the right column, there are two independent motions with a common rotation, with the x -axis indicating the rotation angle (in degrees). The noise level is 2.0 pixel standard deviation.




Figure 2: Percentage of misclassification on a real sequence with 2 independent motions with common rotation for different frame pairs. The x -axis indicates frame pair index.

References

- [1] J. Costeira and T. Kanade. Multi-body factorization methods for motion analysis. In *IEEE International Conference on Computer Vision*, pages 1071–1076, 1995.
- [2] A. Fitzgibbon and A. Zisserman. Multibody structure and motion: 3D reconstruction of independently moving objects. In *European Conference on Computer Vision*, pages 891–906, 2000.
- [3] J. Harris. *Algebraic Geometry: A First Course*. Springer-Verlag, 1992.
- [4] R. Hartley and R. Vidal. The multibody trifocal tensor: Motion segmentation from 3 perspective views. In *IEEE Conference on Computer Vision and Pattern Recognition*, vol. I, pages 769–775, 2004.
- [5] K. Kanatani. Motion segmentation by subspace separation and model selection. In *IEEE International Conference on Computer Vision*, volume 2, pages 586–591, 2001.
- [6] K. Kanatani. Evaluation and selection of models for motion segmentation. In *Asian Conference on Computer Vision*, pages 7–12, 2002.
- [7] Y. Ma, S. Soatto, J. Kosecka, and S. Sastry. *An Invitation to 3D Vision: From Images to Geometric Models*. Springer-Verlag, 2003.
- [8] P. Sturm. Structure and motion for dynamic scenes - the case of points moving in planes. In *European Conference on Computer Vision*, pages 867–882, 2002.
- [9] P. Torr, R. Szeliski, and P. Anandan. An integrated Bayesian approach to layer extraction from image sequences. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(3):297–303, 2001.
- [10] P. H. S. Torr. Geometric motion segmentation and model selection. *Phil. Trans. Royal Society of London*, 356(1740):1321–1340, 1998.
- [11] R. Vidal and R. Hartley. Motion segmentation with missing data by PowerFactorization and Generalized PCA. In *IEEE Conf. on Computer Vision and Pattern Recognition*, vol. II, pages 310–316, 2004.
- [12] R. Vidal and Y. Ma. A unified algebraic approach to 2-D and 3-D motion segmentation. In *European Conference on Computer Vision*, pages 1–15, 2004.
- [13] R. Vidal, Y. Ma, and J. Piazzzi. A new GPCA algorithm for clustering subspaces by fitting, differentiating and dividing polynomials. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume I, pages 510–517, 2004.
- [14] R. Vidal, Y. Ma, S. Soatto, and S. Sastry. Two-view multibody structure from motion. *International Journal of Computer Vision*, 2005.
- [15] R. Vidal and S. Sastry. Optimal segmentation of dynamic scenes from two perspective views. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 281–286, 2003.
- [16] L. Wolf and A. Shashua. Two-body segmentation from two perspective views. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 263–270, 2001.
- [17] Y. Wu, Z. Zhang, T.S. Huang, and J.Y. Lin. Multibody grouping via orthogonal subspace decomposition. In *IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2, pages 252–257, 2001.