

Structure-From-Motion Using Lines: Representation, Triangulation and Bundle Adjustment

Adrien Bartoli

LASMEA
24, avenue des Landais
63177 Aubière cedex
France

Adrien.Bartoli@gmail.com

Peter Sturm

INRIA Rhône-Alpes
655, avenue de l'Europe
38334 St Ismier cedex
France

Peter.Sturm@inria.fr

Abstract

We address the problem of camera motion and 3D structure reconstruction from line correspondences across multiple views, from initialization to final bundle adjustment. One of the main difficulties when dealing with line features is their algebraic representation.

First, we consider the triangulation problem. Based on Plücker coordinates to represent the 3D lines, we propose a Maximum Likelihood algorithm, relying on linearizing the Plücker constraint and on a Plücker correction procedure, computing the closest Plücker coordinates to a given 6-vector.

Second, we consider the bundle adjustment problem, which is essentially a nonlinear optimization process on camera motion and 3D line parameters. Previous overparameterizations of 3D lines induce gauge freedoms and / or internal consistency constraints. We propose the orthonormal representation, which allows handy nonlinear optimization of 3D lines using the minimum 4 parameters with an unconstrained optimization engine.

We compare our algorithms to existing ones on simulated and real data. Results show that our triangulation algorithm outperforms standard linear and bias-corrected quasi-linear algorithms, and that bundle adjustment using

our orthonormal representation yields results similar to the standard Maximum Likelihood trifocal tensor algorithm, while being usable for any number of views.

1 Introduction

The goal of this paper is to give methods for reconstruction of line features from image correspondences over multiple views, from initialization to final bundle adjustment. Reconstruction of line features is an important topic since it is used in areas such as scene modeling, augmented reality and visual servoing. Bundle adjustment is the computation of an optimal visual reconstruction of camera motion and 3D scene structure, where optimal means Maximum Likelihood in terms of reprojected image error. We make no assumption about the calibration of the cameras. We assume that line correspondences over at least three views are available¹.

While the multiple-view geometry of lines is well-understood, see *e.g.* [5, 11], there is still a need for practical structure and motion algorithms. The factorization algorithms [15, 18, 25] yield reliable results but requires all lines to be visible in all views. We focus on the common three-stage approach, see *e.g.* [11, §17.5], consisting in (i) computing camera motion using inter-image matching tensors, (ii) triangulating the features and (iii) running bundle adjustment.

There exist reliable algorithms for step (i). In particular, it can be solved by computing trifocal tensors for triplets of consecutive images, using *e.g.* the automatic computation algorithm described in [11, §15.6], and registering the triplets in a manner similar to [6]. Other integrated motion estimation systems are [20], based on Kalman filtering techniques and [26], registering each view in turn.

In steps (ii) and (iii), one of the main difficulties when dealing with line features arises: the algebraic representation. Indeed, there is no minimal, complete and globally non singular parameterization of the 4-dimensional set of 3D lines, see *e.g.* [11, §2.2]. Hence, they are often overparameterized, *e.g.* as the join of two points or as the meet of two planes (8 parameters), or by the 6 coefficients of their Plücker coordinates, which must satisfy the bilinear Plücker constraint. Another overparameterization is two images of the line (6 parameters). The most appropriate representation depends upon the problem considered. For example, the algorithm in [11, §15.2] shows that the ‘two image lines’ representation is

¹Line correspondences over two views do not constrain the camera motion.

well-adapted to the computation of the trifocal tensor, while the sequential algorithm of [20] is based on Plücker coordinates.

Concerning step (ii), many of the previous works assume calibrated cameras, e.g. [14, 21, 23, 27] and use specific Euclidean representations. The linear three view algorithm of [27] and the algorithm of [23] utilize a ‘closest point+direction’ representation, while [21] uses the projections of the line on the $x = 0$ and the $y = 0$ planes, which has obvious singularities. These algorithms yield sub-optimal results in that none of them maximizes the individual likelihood of the reconstructed lines.

Bundle adjustment, step (iii), is a nonlinear procedure involving camera and 3D line parameters, attempting to maximize the likelihood of the reconstruction, corresponding to minimizing the reprojection error when the noise on measured features has an identical and independent (i.i.d.) normal distribution. Previously-mentioned overparameterizations are not well-adapted to standard nonlinear optimization engines. The ‘two point’ and the ‘two plane’ overparameterizations have 4 degrees of internal gauge freedoms² which may induce numerical instabilities. The ‘two image lines’ parameterization has 2 degrees of internal gauge freedoms and implies that one may have to choose different images for different lines since all lines may not be visible in all images. Also, one must check that the chosen images are not too close to each other. Finally, direct optimization of Plücker coordinates makes sense only if a constrained optimization technique is used to enforce the bilinear Plücker constraint. An appropriate representation would not involve internal constraint or gauge freedom.

To summarize, there is a need for an efficient optimal triangulation algorithm, and a representation of 3D lines well-adapted to nonlinear optimization. We address both of these problems through the following contributions.

In §3, we give an overview of various 3D line representations and their characteristics.

In §4, we propose triangulation methods based on using Plücker coordinates to represent the lines. A simple and optimal algorithm is obtained based on linearizing the bilinear Plücker constraint within an iteratively reweighted least squares approach.

In §5, we propose a nonlinear representation of 3D lines that we call the *orthonormal representation*. This representation allows efficient nonlinear optimization since only the minimum 4 parameters are computed at each step which allows

²For the former one, the position of the points along the line, and the free scale factor of the homogeneous representation of these points.

the use of a standard unconstrained optimization engine. With this representation, there is no internal gauge freedom or consistency constraint, and analytic differentiation of the error function is possible.

Finally, §6 validates our algorithms and compares them to existing ones. The next section gives some preliminaries and notations and states the problem.

2 Preliminaries and Notation

Notation. We make no formal distinction between coordinate vectors and physical entities. Everything is represented in homogeneous coordinates. Equality up to scale is denoted by \sim , transposition and transposed inverse by T and $^{-T}$. Vectors are typeset using bold fonts (\mathbf{L} , \mathbf{l}), matrices using sans-serif fonts (S , A , R) and scalars in italics. Bars represent inhomogeneous leading parts of vectors or matrices, *e.g.* $\mathbf{M}^T \sim (\bar{\mathbf{M}}^T \mid m)$. The \mathcal{L}_2 -norm of vector \mathbf{v} is denoted $\|\mathbf{v}\|$. The identity matrix is denoted I . $SO(2)$ and $SO(3)$ denote the 2D and 3D rotation groups.

The 2D orthogonal (Euclidean) distance between point \mathbf{q} and line \mathbf{l} weighted by q_3 is:

$$d_{\perp}^2(\mathbf{q}, \mathbf{l}) = (\mathbf{q}^T \mathbf{l})^2 / (l_1^2 + l_2^2). \quad (1)$$

Matrix factorization. We make use of the Singular Value Decomposition of matrices, dubbed SVD. The SVD of matrix A is $A_{m \times n} = U_{m \times n} \Sigma_{n \times n} V_{n \times n}^T$, where U and V are orthonormal, and Σ is diagonal, containing the singular values of A in decreasing order. The QR factorization of matrix A is $A_{m \times n} = Q_{m \times m} R_{m \times n}$, with Q orthonormal and R upper triangular. More details on these matrix factorizations can be read in *e.g.* [7].

Maximum likelihood estimation. As noted in [11, §15.7.2], no matter how many points are used to represent an image line \mathbf{l} , the quadratic error function on it can be expressed in the form $d_{\perp}^2(\mathbf{x}, \mathbf{l}) + d_{\perp}^2(\mathbf{y}, \mathbf{l})$ for two weighted points \mathbf{x} , \mathbf{y} on \mathbf{l} . We will use this representation for simplicity. If we have 3D lines $\mathcal{S} = \{\mathbf{L}^1, \dots, \mathbf{L}^m\}$ and cameras $\mathcal{M} = \{\mathbf{P}^1, \dots, \mathbf{P}^n\}$, the negative log likelihood function $\mathcal{E}(\mathcal{S}, \mathcal{M})$ for the reconstruction, corresponding to the reprojection error,

can be written in terms of individual reprojection errors $\mathcal{E}(\mathbf{L}^j, \mathcal{M})$ for each line j :

$$\mathcal{E}(\mathcal{S}, \mathcal{M}) = \sum_{j=1}^m \mathcal{E}(\mathbf{L}^j, \mathcal{M}) \quad (2)$$

$$\mathcal{E}(\mathbf{L}^j, \mathcal{M}) = \sum_{i=1}^n (d_{\perp}^2(\mathbf{x}^{ij}, \mathbf{l}^{ij}) + d_{\perp}^2(\mathbf{y}^{ij}, \mathbf{l}^{ij})). \quad (3)$$

3 Representing 3D Lines

We describe several representations for 3D lines in projective space and their characteristics. Some of these representations are ‘partial’ in the sense that they can only represent a subset of all 3D lines. For example, some work on metric reconstruction, particularly in photogrammetry, assume that the reconstructed lines do not lie at infinity. The goal of this study is to choose a representation for the triangulation and bundle adjustment problems. Concerning the triangulation, the most important criterion is that the reprojected lines is a linear function of the 3D line. Bundle adjustment is a nonlinear procedure allowing more flexibility in the choice of the parameterization. The quality of the parameterization is assessed based on criteria such as the number of internal gauge freedoms or internal constraints. A summary of the reviewed representations is finally provided. The first representation that we describe is the Plücker coordinates. We link all the other representations to Plücker coordinates.

3.1 Complete Representations

Plücker coordinates. Given two 3D points $\mathbf{M}^T \sim (\bar{\mathbf{M}}^T \mid m)$ and $\mathbf{N}^T \sim (\bar{\mathbf{N}}^T \mid n)$, one can represent the line joining them by a homogeneous ‘Plücker’ 6-vector $\mathbf{L}^T \sim (\mathbf{a}^T \mid \mathbf{b}^T)$, see *e.g.* [11, §2.2]:

$$\begin{cases} \mathbf{a} &= \bar{\mathbf{M}} \times \bar{\mathbf{N}} \\ \mathbf{b} &= m\bar{\mathbf{N}} - n\bar{\mathbf{M}}. \end{cases} \quad (4)$$

Other conventions for Plücker 6-vectors are also possible. Each comes with a bilinear constraint that the 6-vector must satisfy in order to represent valid line coordinates. For our definition, the constraint is:

$$\mathcal{C}(\mathbf{L}) = 0 \quad \text{where} \quad \mathcal{C}(\mathbf{L}) = \mathbf{a}^T \mathbf{b}. \quad (5)$$

Similarly, one can construct the Plücker coordinates of a line defined as the meet of two planes. The Plücker coordinates of a line defined as the meet of two planes $\mathbf{P}^T \sim (\bar{\mathbf{P}}^T \mid p)$ and $\mathbf{Q}^T \sim (\bar{\mathbf{Q}}^T \mid q)$ are given by:

$$\begin{cases} \mathbf{a} &= p\bar{\mathbf{Q}} - q\bar{\mathbf{P}} \\ \mathbf{b} &= \bar{\mathbf{P}} \times \bar{\mathbf{Q}}. \end{cases} \quad (6)$$

As an example, triangulation from two views has the following closed-form solution. Let \mathbf{P}^1 and \mathbf{P}^2 be the two projection matrices and \mathbf{l}^1 and \mathbf{l}^2 the two imaged lines. The Plücker coordinates of the corresponding 3D line are given as the meet of the two viewing planes $\pi^i \sim \mathbf{P}^{iT} \mathbf{l}^i$.

Given a standard (3×4) perspective projection matrix $\mathbf{P} \sim (\bar{\mathbf{P}} \mid \mathbf{p})$, a (3×6) matrix projecting Plücker line coordinates [2, 5] is given by:

$$\tilde{\mathbf{P}} \sim (\det(\bar{\mathbf{P}})\bar{\mathbf{P}}^{-T} \mid [\mathbf{p}]_{\times}\bar{\mathbf{P}}). \quad (7)$$

It can be easily derived by expanding the expression of the 2D line joining the projections of two points:

$$\begin{aligned} \mathbf{l} &\sim \mathbf{m} \wedge \mathbf{n} \\ &\sim (\mathbf{P}\mathbf{M}) \wedge (\mathbf{P}\mathbf{N}) \\ &\sim (\bar{\mathbf{P}}\bar{\mathbf{M}} + m\mathbf{p}) \wedge (\bar{\mathbf{P}}\bar{\mathbf{N}} + n\mathbf{p}) \\ &\sim (\bar{\mathbf{P}}\bar{\mathbf{M}}) \wedge (\bar{\mathbf{P}}\bar{\mathbf{N}}) + m\mathbf{p} \wedge (\bar{\mathbf{P}}\bar{\mathbf{N}}) - n\mathbf{p} \wedge (\bar{\mathbf{P}}\bar{\mathbf{M}}) \\ &\sim \det(\bar{\mathbf{P}})\bar{\mathbf{P}}^{-T}(\bar{\mathbf{M}} \wedge \bar{\mathbf{N}}) + [\mathbf{p}]_{\wedge}\bar{\mathbf{P}}(m\bar{\mathbf{N}} - n\bar{\mathbf{M}}) \\ &\sim \tilde{\mathbf{P}}\mathbf{L}. \end{aligned}$$

Seo *et al.* [20] use the Plücker coordinates representation for sequential Structure-From-Motion with a Kalman filtering technique. Pottmann *et al.* [17] use these coordinates for 3D shape reconstruction and understanding from 3D data.

Pair of points or pair of planes. These are two dual representations, described in details in [11, §2.2.2]. In the first case, the line is defined as the join of two points \mathbf{M} and \mathbf{N} , and in the second case, it is defined as the intersection of two planes \mathbf{P} and \mathbf{Q} . These representations have similar characteristics. They have 8 parameters, hence 4 degrees of gauge freedom, the position of the points along the line (respectively the position of the planes in the pencil of planes around the line) and the scale factors in the homogeneous coordinates of the points or the planes.

For metric reconstruction, if one drops the lines at infinity, the two point representation has 6 parameters. There is a direct link with Plücker coordinates using equations (4) and (6). The projected line l is a bilinear function of the entries of the point or the plane coordinates. For example, for the two point representation, $l \sim (PM) \times (PN)$. Hartley [10] proposes a triangulation algorithm based on these representations. Habib *et al.* [9] use the two point representation for bundle adjustment. They consider that the line is not at infinity. The ambiguity on the position of the points along the line is fixed by constraining them to reprojected near the end-points observed in one of the images.

3.2 Partial Representations

Closest point and direction. A 3D line is represented by its closest point to the origin, with coordinates $Q^T \sim (\bar{Q}^T 1)$, and its direction, with coordinates $Q_\infty^T \sim (\bar{Q}_\infty^T 0)$, giving a total of 6 parameters. This representation does not include lines at infinity and hence can not be used in projective space. The link with the Plücker line coordinates L is given by:

$$L \sim \begin{pmatrix} \bar{Q} \times \bar{Q}_\infty \\ \bar{Q}_\infty \end{pmatrix}.$$

Reprojecting the line with the camera matrix $P \sim (\bar{P} \mathbf{p})$ is a bilinear function of the line parameters: $l \sim (\bar{P}\bar{Q} + \mathbf{p}) \times (\bar{P}\bar{Q}_\infty)$. The line reconstruction algorithms proposed by Weng *et al.* [27] for three views and by Taylor and Kriegman [23] for multiple views use this representation. In the field of photogrammetry, Tomaselli and Lugnani [24] use this representation for bundle adjustment. Mulawa and Mikhail [16] use the additional constraint $\|\bar{Q}_\infty\| = 1$.

Two projections. A 3D line can be represented by two projections [10, 21]. This is related to the fact that reconstructing a line from two views has in general a unique solution.

Spetsakis and Aloimonos [21] use the intersection of two planes, one parallel to the plane $x = 0$, and the other one parallel to the plane $y = 0$. These two planes are formulated using 4 parameters a , b , c and d by $x = az + b$ and $y = cz + d$

respectively. The pencil of points \mathbf{Q} on the 3D line is parameterized by the z coordinate:

$$\mathbf{Q} \sim \begin{pmatrix} az + b \\ cz + d \\ z \\ 1 \end{pmatrix}.$$

This representation has obvious singularities: lines which are parallel to the plane $z = 0$ can not be represented. Indeed, the points lying on such lines have a constant z coordinate, and since the points are parameterized by this coordinate, one always gets the same point if the z coordinate is constant. One can link this representation to the Plücker coordinates \mathbf{L} of the line by considering any two points lying on the line, *e.g.* for $z = 0$ and $z = 1$, and equation (4), giving:

$$\mathbf{L} \sim \begin{pmatrix} d \\ -b \\ bc - da \\ a \\ c \\ 1 \end{pmatrix}.$$

Ayache and Faugeras [4] use this representation for mobile robot navigation. In the field of photogrammetry, Habib [8] extends this representation by using different pairs of planes depending on the 3D line, in order to avoid the singularities.

Hartley [10] uses two images of the line. This representation has the following singularities: all 3D lines lying in an epipolar plane induced by the two cameras have the same images in both views. The 3D lines that can not be uniquely represented thus form a Linear Line Complex, see *e.g.* [22]. Note that these singularities can be encountered in practice. The Plücker coordinates corresponding to this representation can be calculated by intersecting the two viewing planes induced by the two image lines using equation (6). Hartley shows that the reprojection of the line in other views is a bilinear function of the parameters.

The Denavit-Hartenberg parameters. The Denavit-Hartenberg representation [3] has become the standard way of representing robots and modeling their motions. The idea is to relate each joint to the next by using the minimal 4 parameters, namely two distances and two angles. A general 3D Euclidean transformation, between two Euclidean coordinate frames, has 6 degrees of freedom. For using the

Denavit-Hartenberg representation, the x -axis of one coordinate frame has to be aligned with the line orthogonal to the z axes of both coordinate frames, which cancels out 2 degrees of freedom, 1 in rotation and 1 in translation. This suggests to represent a 3D line by the z -axis of a coordinate frame, and to parameterize it by the 4 Denavit-Hartenberg parameters with respect to a reference coordinate frame, *e.g.* the world coordinate frame. The Plücker coordinates corresponding to these parameters can be obtained by *e.g.* applying the coordinate transformation given by the 4 parameters to the z -axis $\mathbf{L}_z^T \sim (0 \ 0 \ 0 \ 0 \ 0 \ 1)$ of the reference frame using a 3D line rigid displacement matrix [2]. The projection equation is nonlinear in the Denavit-Hartenberg parameters since it involves products and trigonometric operators.

One problem with this parameterization is that two distances are used as parameters, which prevents from representing the lines at infinity. There is also an indeterminacy in the choice of one of the coordinate frame when the line is parallel to the z -axis of the reference coordinate frame.

Roberts [19] proposes to model 3D lines using two distances and two angles. His representation has drawbacks similar to those described above.

Note that there are other representations for modeling robots. For example, Hayati *et al.* [12] introduce an extra rotation parameter to the Denavit-Hartenberg representation to model the error due to near parallel axes. This representation is thus not minimal.

3.3 Summary

Table 1 summarizes the characteristics of the aforementioned representations, and of the orthonormal representation that we propose in §5. We observe that the only representation for which the reprojected lines is a linear function of the 3D line parameters is the Plücker coordinates. It is also seen that besides our orthonormal representation, no other complete representation allows a minimal update with 4 parameters, which is due to gauge freedoms and / or internal consistency constraints. Minimal update is an important criterion for using a representation within bundle adjustment.

4 Triangulation

This section discusses computation of structure given camera motion. We propose direct linear and iterative nonlinear methods to recover Plücker line coordinates.

representation	comp.	# gauge	# cstr	reprojection	min. up.
closest point and direction	no	1	1	bilinear	no
two image lines	no	2	0	bilinear	no
Denavit-Hartenberg	no	0	0	non-linear	yes
two points or two planes	yes	4	0	bilinear	no
Plücker coordinates	yes	1	1	linear	no
orthonormal representation	yes	0	0	non-linear	yes

Table 1: Summary of different representations for 3D lines with their characteristics. The ability of the representation to cover all lines in \mathbb{P}^3 is on the column ‘comp.’ (completeness). The number of gauge freedoms (column ‘# gauges’) and internal constraints (column ‘# cstr’) are strongly linked. The ‘reprojection’ column is about the equation for reprojecting the 3D line with a perspective camera. The column ‘min. up.’ indicates if the representation can be updated with 4 parameters.

These algorithms are general in the sense that they can be used with calibrated, partially calibrated or uncalibrated cameras.

First, we describe a somehow trivial linear algorithm where a biased error function (compared to the reprojection error) is minimized. This algorithm is subject to the same kind of drawback as the eight-point algorithm for computing the fundamental matrix: due to possible noise in the data, the resulting 6-vectors do not generally satisfy the bilinear Plücker constraint (5), similarly to the matrix computed by the eight-point algorithm not being rank deficient [11, §10.2]. We propose what we call a *Plücker correction* procedure, which allows to compute the closest Plücker coordinates to a 6-vector.

Second, we propose an algorithm where the reprojection error of the line is minimized. The cornerstone of this algorithm is the linearization of the Plücker constraint.

Since the reconstruction of each line is independent from the others, we drop the j index in this section.

4.1 Linear Algorithm

We describe a linear algorithm, ‘LIN’. In the reprojection error (3), each term is based on the square of the 2D point-to-line orthogonal distance d_{\perp} , defined by equation (1). The denominator of this distance is the cause of the nonlinearity.

Ignoring this denominator leads to an algebraic distance denoted d_a , biased compared to the orthogonal distance. It is linear in the predicted line \mathbf{l} and defined by $d_a^2(\mathbf{q}, \mathbf{l}) = d_{\perp}^2(\mathbf{q}, \mathbf{l}) w^2 = (\mathbf{q}^T \mathbf{l})^2$, where the scalar factor w encapsulates the bias as $w^2 = l_1^2 + l_2^2$:

$$(w^i)^2 = \left((\tilde{\mathbf{P}}^i \mathbf{L})_1 \right)^2 + \left((\tilde{\mathbf{P}}^i \mathbf{L})_2 \right)^2. \quad (8)$$

We define the biased linear least squares error function:

$$\mathcal{B}(\mathbf{L}, \mathcal{M}) = \sum_{i=1}^n \left((\mathbf{x}^i{}^T \tilde{\mathbf{P}}^i \mathbf{L})^2 + (\mathbf{y}^i{}^T \tilde{\mathbf{P}}^i \mathbf{L})^2 \right) \quad (9)$$

$$= \|\mathbf{A}_{(2n \times 6)} \mathbf{L}\|^2 \quad \text{with} \quad \mathbf{A} = \begin{pmatrix} \dots \\ \mathbf{x}^i{}^T \tilde{\mathbf{P}}^i \\ \mathbf{y}^i{}^T \tilde{\mathbf{P}}^i \\ \dots \end{pmatrix}. \quad (10)$$

Since \mathbf{L} is an homogeneous vector, we add the constraint $\|\mathbf{L}\|^2 = 1$. The \mathbf{L} that minimizes $\mathcal{B}(\mathbf{L}, \mathcal{M})$ is then given by the singular vector of \mathbf{A} associated to its smallest singular value, that we compute using SVD. Due to noise, the recovered 6-vector does not in general satisfy the Plücker constraint (5).

4.2 Plücker Correction

The Plücker correction procedure is analogous to the standard rank correction of the fundamental matrix based on SVD: the eight-point algorithm linearly computes a full-rank matrix \mathbf{F} , whose smallest singular value is nullified to obtain the rank-two matrix $\hat{\mathbf{F}}$, see *e.g.* [11]. Matrix $\hat{\mathbf{F}}$ is the closest rank-two matrix to \mathbf{F} , in the sense of the Frobenius norm. It is used to initialize nonlinear algorithms.

The Plücker correction procedure computes the closest Plücker coordinates to a given 6-vector, where closest is to be understood in the sense of the \mathcal{L}_2 -norm, equivalent to the matrix Frobenius norm. It is also equivalent to the Euclidean distance between two points in \mathbb{R}^6 . This correction is necessary to initialize the nonlinear algorithms from the solution provided by linear methods ignoring the Plücker constraint. Pottmann *et al.* [17] use the Euclidean distance between Plücker coordinate vectors to compare 3D lines. They underline the facts that this distance is practical for minimization purposes and is in accordance with visualization in the region of interest, *i.e.* near the origin.

More formally, let $\mathbf{L}^T \sim (\mathbf{a}^T \mid \mathbf{b}^T)$ be a 6-vector that does not necessarily satisfy the Plücker constraint (5), *i.e.* $\mathbf{a}^T \mathbf{b}$ might be non-zero. We seek

$\widehat{\mathbf{L}}^T \sim (\mathbf{u}^T \mid \mathbf{v}^T)$, defined by $\min_{\widehat{\mathbf{L}}, \mathbf{u}^T \mathbf{v} = 0} \|\widehat{\mathbf{L}} - \mathbf{L}\|^2$. This is a linear least squares optimization problem under a nonlinear constraint. Although it has a clear and concise formulation, it is *not* trivial.

Obviously, one can modify one entry of the Plücker coordinates in accordance with the Plücker constraint, *e.g.* set $a_1 = -(a_2 b_2 + a_3 b_3)/b_1$. This simple solution has the disadvantage that the entry must be chosen depending on the actual values of the coordinates since the correction rule uses a division. Also, all entries are clearly not treated uniformly.

By comparison, our solution orthogonally projects the 6-vector on the Klein quadric and treats all its entries the same way. Kanatani [13] proposes a general iterative scheme for projecting points on nonlinear manifolds, such as projecting points in \mathbb{R}^6 on the Klein quadric. Our algorithm performs this projection in a non iterative manner, which thus guarantees that the optimal projected point on the Klein quadric, *i.e.* the optimal 3D line, is found. Its derivation is quite tricky but it can be readily implemented with few lines of code from its summary shown in table 2.

- Compute the Singular Value Decomposition $(\mathbf{a} \ \mathbf{b}) = \bar{\mathbf{U}} \bar{\Sigma} \bar{\mathbf{V}}^T$.
- Let $\bar{\mathbf{Z}} = \bar{\Sigma} \bar{\mathbf{V}}^T$, form matrix $\mathbf{T} = \begin{pmatrix} z_{21} & z_{22} \\ z_{12} & -z_{11} \end{pmatrix}$.
- Compute the singular vector $\widehat{\mathbf{v}}$ associated to the smallest singular value of matrix \mathbf{T} .
- Define $\bar{\mathbf{V}} = \begin{pmatrix} \hat{v}_1 & -\hat{v}_2 \\ \hat{v}_2 & \hat{v}_1 \end{pmatrix}$ and set $(\mathbf{u} \ \mathbf{v}) \sim \bar{\mathbf{U}} \widehat{\mathbf{V}} \text{diag} \left(\widehat{\mathbf{V}}^T \bar{\Sigma} \bar{\mathbf{V}}^T \right)$.

Table 2: The *Plücker correction* algorithm. Given a 6-vector $\mathbf{L}^T \sim (\mathbf{a}^T \mid \mathbf{b}^T)$, this algorithm computes the closest Plücker coordinates $\widehat{\mathbf{L}}^T \sim (\mathbf{u}^T \mid \mathbf{v}^T)$, *i.e.* $\mathbf{u}^T \mathbf{v} = 0$, in the sense of the \mathcal{L}_2 -norm, *i.e.* $\|\widehat{\mathbf{L}} - \mathbf{L}\|^2$ is minimized.

A geometric interpretation. We interpret the 3-vectors \mathbf{a} , \mathbf{b} , \mathbf{u} and \mathbf{v} as coordinates of 3D points. These points are not directly linked to the underlying 3D line. This interpretation is just intended to visualize the problem. The Plücker constraint $\mathbf{u}^T \mathbf{v} = 0$ corresponds to the fact that the lines induced by the origin with \mathbf{u} and \mathbf{v} are perpendicular. The correction criterion is the sum of squared Euclidean distances between \mathbf{a} and \mathbf{u} and between \mathbf{b} and \mathbf{v} . Hence, the problem may be formulated as finding two points \mathbf{u} and \mathbf{v} , as close as possible to \mathbf{a} and \mathbf{b} respectively and such that the lines induced by the origin with \mathbf{u} and \mathbf{v} are perpendicular. We

begin by rotating the coordinate frame such that \mathbf{a} and \mathbf{b} are transferred on the $z = 0$ plane. This is the *reduction of the problem*. We solve the reduced problem, by finding two points on the $z = 0$ plane, minimizing the correction criterion and satisfying the Plücker constraint. Finally, we *express the solution* back to the original space.

Reducing the problem. Let us define the (3×2) matrices $\bar{\mathbf{C}} \sim (\mathbf{a} \ \mathbf{b})$ and $\hat{\mathbf{C}} \sim (\mathbf{u} \ \mathbf{v})$. The Plücker constraint is fulfilled if and only if the columns of matrix $\hat{\mathbf{C}}$ are orthogonal. We rewrite the correction criterion as :

$$\mathcal{O} = \|\mathbf{L} - \hat{\mathbf{L}}\|^2 = \|\bar{\mathbf{C}} - \hat{\mathbf{C}}\|^2.$$

Using the following SVD $\bar{\mathbf{C}}_{(3 \times 2)} = \bar{\mathbf{U}}_{(3 \times 2)} \bar{\mathbf{\Sigma}}_{(2 \times 2)} \bar{\mathbf{V}}_{(2 \times 2)}^T$:

$$\mathcal{O} = \|\bar{\mathbf{U}} \bar{\mathbf{\Sigma}} \bar{\mathbf{V}}^T - \hat{\mathbf{C}}\|^2 = \|\bar{\mathbf{\Sigma}} \bar{\mathbf{V}}^T - \bar{\mathbf{U}}^T \hat{\mathbf{C}}\|^2,$$

since $\bar{\mathbf{U}}$ has orthonormal columns. We define $\bar{\mathbf{Z}} = \bar{\mathbf{\Sigma}} \bar{\mathbf{V}}^T$ and $\hat{\mathbf{Z}} = \bar{\mathbf{U}}^T \hat{\mathbf{C}}$. Matrix $\bar{\mathbf{V}}$ is orthonormal and $\bar{\mathbf{\Sigma}}$ is diagonal, hence the rows of $\bar{\mathbf{Z}}$ are orthogonal (*i.e.* $\bar{\mathbf{Z}} \bar{\mathbf{Z}}^T$ is diagonal, but not $\bar{\mathbf{Z}}^T \bar{\mathbf{Z}}$). Note that $\hat{\mathbf{Z}} = \bar{\mathbf{U}}^T \hat{\mathbf{C}}$ implies $\hat{\mathbf{C}} = \bar{\mathbf{U}} \hat{\mathbf{Z}}$, even if $\bar{\mathbf{U}} \bar{\mathbf{U}}^T$ is not the identity³. The problem is reduced to finding a column-orthogonal⁴ matrix $\hat{\mathbf{Z}}$, as close as possible to the row-orthogonal matrix $\bar{\mathbf{Z}}$.

Solving the reduced problem. We parameterize the column-orthogonal matrix $\hat{\mathbf{Z}}$ as $\hat{\mathbf{Z}} = \hat{\mathbf{V}} \hat{\mathbf{\Sigma}}$, where $\hat{\mathbf{V}}$ is orthonormal and $\hat{\mathbf{\Sigma}}$ is diagonal. Hence :

$$\mathcal{O} = \|\bar{\mathbf{\Sigma}} \bar{\mathbf{V}}^T - \hat{\mathbf{V}} \hat{\mathbf{\Sigma}}\|^2 = \|\hat{\mathbf{V}}^T \bar{\mathbf{\Sigma}} \bar{\mathbf{V}}^T - \hat{\mathbf{\Sigma}}\|^2.$$

The diagonal matrix $\hat{\mathbf{\Sigma}}$ which minimizes this expression is given by the diagonal entries of $\hat{\mathbf{V}}^T \bar{\mathbf{\Sigma}} \bar{\mathbf{V}}^T$, and does not depend on the solution for $\hat{\mathbf{V}}$. The orthonormal

³Indeed, denote \mathbf{u}_i the columns of matrix $\bar{\mathbf{U}}$ and form $\mathbf{U} = (\mathbf{u}_1 \ \mathbf{u}_2 \ \mathbf{u}_1 \times \mathbf{u}_2)$. We have $\mathbf{U}^T \bar{\mathbf{U}} = (\mathbf{I}_{(2 \times 2)} \ \mathbf{0}_{(2 \times 1)})^T$. Let us multiply the correction criterion by \mathbf{U}^T : $\mathcal{O} = \|(\bar{\mathbf{V}} \bar{\mathbf{\Sigma}} \ \mathbf{0}_{(2 \times 1)})^T - \mathbf{U}^T \hat{\mathbf{C}}\|^2$. Denote $\mathbf{Y}_{(3 \times 2)} = \mathbf{U}^T \hat{\mathbf{C}}$. The optimal solution has the form $\mathbf{Y}^T = (\hat{\mathbf{Z}}^T \ \mathbf{0}_{(2 \times 1)})$, since, according to the geometric interpretation, the corrected points \mathbf{u} and \mathbf{v} must lie on the plane defined by points \mathbf{a} , \mathbf{b} and the origin, the plane $z = 0$. Therefore, we obtain $\hat{\mathbf{C}} = \mathbf{U} \mathbf{Y} = \bar{\mathbf{U}} \bar{\mathbf{Y}}$.

⁴The fact that matrix $\hat{\mathbf{Z}} = \bar{\mathbf{U}}^T \hat{\mathbf{C}}$ is column-orthogonal is induced from the Plücker constraint. Indeed, this constraint implies that $\hat{\mathbf{C}}$ is column-orthogonal, hence $\hat{\mathbf{C}}^T \hat{\mathbf{C}}$ is diagonal. Matrix $\mathbf{U}^T \hat{\mathbf{C}}$, where $SO(3) \ni \mathbf{U} = (\mathbf{u}_1 \ \mathbf{u}_2 \ \mathbf{u}_1 \times \mathbf{u}_2) = (\bar{\mathbf{U}} \ \bar{\mathbf{u}})$, is also column-orthogonal. Observe that $\hat{\mathbf{C}}^T \mathbf{U} \mathbf{U}^T \hat{\mathbf{C}} = \hat{\mathbf{C}}^T \bar{\mathbf{U}} \bar{\mathbf{U}}^T \hat{\mathbf{C}} + \hat{\mathbf{C}}^T \bar{\mathbf{u}} \bar{\mathbf{u}}^T \hat{\mathbf{C}} = \hat{\mathbf{C}}^T \bar{\mathbf{U}} \bar{\mathbf{U}}^T \hat{\mathbf{C}}$ since $\bar{\mathbf{u}}^T \hat{\mathbf{C}} = \mathbf{0}^T$. Hence, matrix $\bar{\mathbf{U}}^T \hat{\mathbf{C}}$ is column-orthogonal.

matrix $\widehat{V} = (\widehat{v}_1 \ \widehat{v}_2)$ is given by minimizing the sum of squares of the off-diagonal entries of $\widehat{V}^T \bar{Z}$, with $\bar{Z} = \bar{\Sigma} \bar{V}^T = (z_1 \ z_2)$:

$$\mathcal{O} = (\widehat{v}_1^T z_2)^2 + (\widehat{v}_2^T z_1)^2.$$

Define the 2D rotation matrix with angle $\pi/2$ by $M = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ and parameterize the orthonormal matrix \widehat{V} by a unit vector \widehat{v} , as :

$$\begin{cases} \widehat{v}_1 &= \widehat{v} \\ \widehat{v}_2 &= M\widehat{v}, \end{cases}$$

The correction criterion can be rewritten as :

$$\mathcal{O} = (\widehat{v}^T z_2)^2 + (\widehat{v}^T M^T z_1)^2 = \|\mathbb{T}\widehat{v}\|^2 \quad \text{with} \quad \mathbb{T} = \begin{pmatrix} z_2^T \\ z_1^T M \end{pmatrix}.$$

The unit vector \widehat{v} minimizing this expression is given by the singular vector associated to the smallest singular value of matrix \mathbb{T} .

Expressing the solution. From vector \widehat{v} which solves the reduced problem, we form the orthonormal matrix $\widehat{V} = \begin{pmatrix} \widehat{v}_1 & -\widehat{v}_2 \\ \widehat{v}_2 & \widehat{v}_1 \end{pmatrix}$. The diagonal matrix $\widehat{\Sigma}$ is given by $\widehat{\Sigma} = \text{diag}(\widehat{V}^T \bar{\Sigma} \bar{V}^T)$.

4.3 Quasi-Linear Algorithms

We describe algorithms ‘QLIN1’ and ‘QLIN2’, that consider the reprojection error (3). They are based on an iterative bias-correction, through reweighting of the biased error function (9). Such algorithms are coined quasi-linear.

We showed previously that the orthogonal and the algebraic distances are related by a scalar factor, given by equation (8), depending on the 3D line. The reprojection error and the biased error functions are therefore related by a set of such factors, one for each image of the line. The fact that these factors depend on the unknown 3D line suggests an iterative reweighting scheme.

The first approach that comes to mind is ‘QLIN1’. The linear system considered for method LIN is formed and solved. The resulting 6-vector L_0 is corrected to be valid Plücker coordinates. This yields a biased estimate of the 3D line. Using this estimate, weight factors that contain the bias of the linear least squares error function are computed, and used to reweight the equations. The process is iterated to compute successive refined estimates L_k until convergence, where k is

the iteration counter. Convergence is determined by thresholding the difference between two consecutive errors. It is typically reached in 3 or 4 iterations.

Experimental results show that this naive approach performs very badly, see §6. This is due to the fact that the Plücker constraint is enforced afterhand and is not taken into account while solving the linear least squares system.

To remedy to this problem, we propose ‘QLIN2’, that linearizes and enforces the Plücker constraint (5), as follows. The algorithm is summarized in table 3. Rewrite the constraint as $\mathcal{C}(\mathbf{L}) = \mathbf{L}^\top \mathbf{G} \mathbf{L}$ where $\mathbf{G}_{(6 \times 6)} = \begin{pmatrix} 0 & \mathbf{I} \\ \mathbf{I} & 0 \end{pmatrix}$. By expanding this expression to first order around the estimate \mathbf{L}_k , and after some minor algebraic manipulations, we obtain the following linear constraint on \mathbf{L}_{k+1} :

$$\mathcal{C}_k(\mathbf{L}_{k+1}) = \mathbf{L}_k^\top \mathbf{G} \mathbf{L}_{k+1} = 0.$$

We follow the constrained linear least squares optimization method summarized in [11, §A3.4.3] to enforce this linearized constraint, as well as $\|\mathbf{L}_{k+1}\| = 1$. The idea is to find an orthonormal basis of all possible vectors satisfying the constraint and to solve for a 5-vector γ expressed in this basis. Such an orthonormal basis is provided by computing the nullspace of $\mathbf{L}_k^\top \mathbf{G}$ using SVD. Let $\bar{\mathbf{V}}$ be a (6×5) orthonormal matrix whose columns span the basis (i.e. $\mathbf{L}_k^\top \mathbf{G} \bar{\mathbf{V}} = \mathbf{0}$), we define $\mathbf{L}_{k+1} = \bar{\mathbf{V}} \gamma$, hence $\mathcal{C}_k(\mathbf{L}_{k+1}) = \mathbf{L}_k^\top \mathbf{G} \bar{\mathbf{V}} \gamma = 0$ and $\|\mathbf{L}_{k+1}\| = \|\gamma\|$. We solve for γ by substituting in equation (10) ($\|\mathbf{A} \mathbf{L}_{k+1}\|^2 = \|\mathbf{A} \bar{\mathbf{V}} \gamma\|^2$). The singular vector associated to the smallest singular value of matrix $\mathbf{A} \bar{\mathbf{V}}$ provides the solution vector with unit \mathcal{L}_2 -norm such that $\mathcal{B}(\mathbf{L}_{k+1}, \mathcal{M})$ is minimized.

1. *Initialization:* Form the linear least squares system \mathbf{A} from equation (10), compute \mathbf{L}_0 by minimizing $\|\mathbf{A} \mathbf{L}_0\|^2$, see §4.1, and by applying the Plücker correction procedure described in §4.2. Set $k = 0$.
2. *Constraint linearization:* Compute the Singular Value Decomposition $\mathbf{L}_k^\top \mathbf{G} \sim \mathbf{u}^\top \text{diag}(1, 0, 0, 0, 0, 0) (\mathbf{v}_{(6 \times 1)} \mid \bar{\mathbf{V}}_{(6 \times 5)})^\top$.
3. *Estimation:* Compute $\min_{\gamma, \|\gamma\|^2=1} \|\mathbf{A} \bar{\mathbf{V}} \gamma\|^2$ and set $\mathbf{L}_{k+1} = \bar{\mathbf{V}} \gamma$.
4. *Bias-correction:* Reweight the linear system \mathbf{A} by computing the weights according to equation (8).
5. *Iteration:* Iterate steps 2, 3 and 4 until convergence.

Table 3: The quasi-linear algorithm ‘QLIN2’ for optimal triangulation.

5 Bundle Adjustment

Bundle adjustment is the nonlinear minimization of the reprojection error (2), over camera and 3D line parameters. We focus on the parameterization of 3D lines. Parameterizing the camera motion has been addressed in *e.g.* [1, 11, §A4.6].

5.1 Problem Statement

As said in the introduction, there are various possibilities to overparameterize the 4-dimensional set of 3D lines. In the context of nonlinear optimization, choosing an overparameterized representation may induce the following problems. First, the computational cost of each iteration is increased by superfluous parameters. Second, artificial freedoms in the parameter set (internal gauge freedoms) are induced and may give rise to numerical instabilities. Third, some internal consistency constraints, such as the Plücker constraint, may have to be enforced.

These reasons motivate the need for a representation of 3D lines allowing nonlinear optimization with the minimum 4 parameters. In that case, there is no free scale induced by homogeneity or internal consistency constraints, and an unconstrained optimization engine can be used.

5.2 The Orthonormal Representation

The orthonormal representation has been introduced in [1] for the nonlinear optimization of the fundamental matrix with the minimum 7 parameters. It consists in finding a representation involving elements of $SO(n)$ and scalars (hence the term ‘orthonormal representation’). In particular, no other algebraic constraints should be necessary, such as the rank-two constraint of fundamental matrices or the bilinear Plücker constraint. Using orthonormal matrices implies that the representation is well-conditioned. Based on such a representation, local update using the minimum number of parameters is possible.

Commonly used nonlinear optimization engine, *e.g.* Newton type such as Levenberg-Marquardt, often require the Jacobian matrix of the error function with respect to the update parameters. In the orthonormal representation framework, we split it as the product of the Jacobian matrix of the error function considered with respect to the ‘standard’ entity representation, *e.g.* the fundamental matrix or Plücker coordinates, and the *orthonormal Jacobian matrix*, *i.e.* for the ‘standard’ representation with respect to the update parameters.

Example: representing \mathbb{P}^1 . We derive the orthonormal representation of the 1-dimensional projective space \mathbb{P}^1 . This is used in §5.3 to derive the orthonormal representation of 3D lines. Let $\boldsymbol{\sigma} \in \mathbb{P}^1$. Such a 2-vector is defined up to scale and has therefore only 1 degree of freedom. We represent it by an $SO(2)$ matrix W defined by:

$$W = \frac{1}{\|\boldsymbol{\sigma}\|} \begin{pmatrix} \sigma_1 & -\sigma_2 \\ \sigma_2 & \sigma_1 \end{pmatrix}. \quad (11)$$

The first column of this matrix is $\boldsymbol{\sigma}$ itself, normalized to unit-norm. Let θ be the update parameter. A local update step is $W \leftarrow WR(\theta)$ where $R(\theta)$ is the 2D rotation matrix of angle θ . The Jacobian matrix $\frac{\partial \boldsymbol{\sigma}}{\partial \theta}$ evaluated at $\theta_0 = 0$ (the update is with respect to a base rotation) is given by:

$$\left. \frac{\partial \boldsymbol{\sigma}}{\partial \theta} \right|_{\theta_0} = \left. \frac{\partial \mathbf{w}_1}{\partial \theta} \right|_{\theta_0} = \begin{pmatrix} -\sigma_2 \\ \sigma_1 \end{pmatrix} = \mathbf{w}_2, \quad (12)$$

where \mathbf{w}_i is the i -th column of W .

Updating $SO(3)$. A matrix $U \in SO(3)$ can be locally updated using 3 parameters by any well-behaved (locally non singular) representation, such as 3 Euler angles $\boldsymbol{\theta}^\top = (\theta_1 \mid \theta_2 \mid \theta_3)$ as:

$$U \leftarrow UR(\boldsymbol{\theta}) \quad \text{with} \quad R(\boldsymbol{\theta}) = R_x(\theta_1)R_y(\theta_2)R_z(\theta_3), \quad (13)$$

where $R_x(\theta_1)$, $R_y(\theta_2)$ and $R_z(\theta_3)$ are $SO(3)$ matrices representing 3D rotations around the x -, y - and z -axes with angle θ_1 , θ_2 and θ_3 respectively. The Jacobian matrix is derived as follows. As in the $SO(2)$ case, the update is with respect to a base rotation. The orthonormal Jacobian matrix is therefore evaluated at $\boldsymbol{\theta}_0 = \mathbf{0}_{(3 \times 1)}$:

$$\left. \frac{\partial U}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}_0} = \left(\left. \frac{\partial U}{\partial \theta_1} \right|_{\boldsymbol{\theta}_0} \mid \left. \frac{\partial U}{\partial \theta_2} \right|_{\boldsymbol{\theta}_0} \mid \left. \frac{\partial U}{\partial \theta_3} \right|_{\boldsymbol{\theta}_0} \right).$$

After minor algebraic manipulations, we obtain:

$$\begin{aligned} \left. \frac{\partial U}{\partial \theta_1} \right|_{\boldsymbol{\theta}_0} &= \left. \frac{\partial (UR_x(\theta_1)R_y(\theta_2)R_z(\theta_3))}{\partial \theta_1} \right|_{\boldsymbol{\theta}_0} \\ &= (\mathbf{0}_3 \mid \mathbf{u}_3 \mid -\mathbf{u}_2), \end{aligned} \quad (14)$$

where \mathbf{u}_i is the i -th column of \mathbf{U} . Similarly:

$$\left. \frac{\partial \mathbf{U}}{\partial \theta_2} \right|_{\theta_0} = (-\mathbf{u}_3 \mid \mathbf{0}_3 \mid \mathbf{u}_1) \quad (15)$$

$$\left. \frac{\partial \mathbf{U}}{\partial \theta_3} \right|_{\theta_0} = (\mathbf{u}_2 \mid -\mathbf{u}_1 \mid \mathbf{0}_3). \quad (16)$$

These expressions are vectorized to form the orthonormal Jacobian matrix.

5.3 The Case of 3D Lines

The case of 3D lines is strongly linked with the cases of $SO(2)$ and $SO(3)$, as shown by the following result:

Any (projective) 3D line \mathbf{L} can be represented by:

$$(\mathbf{U}, \mathbf{W}) \in SO(3) \times SO(2),$$

where $SO(2)$ and $SO(3)$ are the Lie groups of respectively (2×2) and (3×3) rotation matrices. (\mathbf{U}, \mathbf{W}) is the orthonormal representation of the 3D line \mathbf{L} .

The proof of this result is obtained by showing that any 3D line has an orthonormal representation $(\mathbf{U}, \mathbf{W}) \in SO(3) \times SO(2)$, while any $(\mathbf{U}, \mathbf{W}) \in SO(3) \times SO(2)$ corresponds to a unique 3D line. The next paragraph illustrates this by means of Plücker coordinates.

Note that this result is consistent with the fact that a 3D line has 4 degrees of freedom, since an element of $SO(2)$ has one degree of freedom and an element of $SO(3)$ has 3 degrees of freedom.

Using this representation of 3D lines, we show that there exists a locally non singular minimal parameterization. Therefore, 3D lines can be locally updated with the minimum 4 parameters. The update scheme is inspired from those given above for 2D and 3D rotation matrices, and can be plugged into most of the existing nonlinear optimization algorithms. These results are summarized in table 4.

Relating Plücker coordinates and the orthonormal representation. The orthonormal representation of a 3D line can be computed from its Plücker coordi-

nates $\mathbf{L}^\top \sim (\mathbf{a}^\top \mid \mathbf{b}^\top)$, as follows. Let $\bar{\mathbf{C}}_{(3 \times 2)} \sim (\mathbf{a} \mid \mathbf{b})$ be factored as :

$$\bar{\mathbf{C}} \sim \underbrace{\begin{pmatrix} \mathbf{a} & \mathbf{b} & \mathbf{a} \times \mathbf{b} \\ \|\mathbf{a}\| & \|\mathbf{b}\| & \|\mathbf{a} \times \mathbf{b}\| \end{pmatrix}}_{SO(3)} \underbrace{\begin{pmatrix} \|\mathbf{a}\| & \\ & \|\mathbf{b}\| \end{pmatrix}}_{(\|\mathbf{a}\| \|\mathbf{b}\|)^\top \in \mathbb{P}^1}.$$

In practice, we use QR decomposition, $\bar{\mathbf{C}}_{(3 \times 2)} = \mathbf{U}_{(3 \times 3)} \Sigma_{(3 \times 2)}$. The special form of matrix Σ , *i.e.* the zero at the (1, 2) entry is due to the Plücker constraint. While $\mathbf{U} \in SO(3)$, the two non-zero entries of Σ defined up to scale can be represented by an $SO(2)$ matrix \mathbf{W} , as shown in §5.2.

Going back from the orthonormal representation to Plücker coordinates is trivial. The Plücker coordinates of the line are obtained from its orthonormal representation (\mathbf{U}, \mathbf{W}) as:

$$\mathbf{L}^\top \sim (w_{11} \mathbf{u}_1^\top \mid w_{21} \mathbf{u}_2^\top), \quad (17)$$

where \mathbf{u}_i is the i -th column of \mathbf{U} .

A 4-parameter update. Consider $(\mathbf{U}, \mathbf{W}) \in SO(3) \times SO(2)$, the orthonormal representation of a 3D line. Since $\mathbf{U} \in SO(3)$, as reviewed in §5.2, it can not be minimally parameterized but can be locally updated using equation (13), as $\mathbf{U} \leftarrow \mathbf{U}R(\boldsymbol{\theta})$ where $\boldsymbol{\theta} \in \mathbb{R}^3$. Matrix $\mathbf{W} \in SO(2)$ can be updated as $\mathbf{W} \leftarrow \mathbf{W}R(\theta)$, where $\theta \in \mathbb{R}$. We define the update parameters by the 4-vector $\mathbf{p}^\top \sim (\boldsymbol{\theta}^\top \mid \theta)$.

We denote \mathbf{J} the (6×4) Jacobian matrix of the Plücker coordinates, with respect to the orthonormal representation. Matrix \mathbf{J} must be evaluated at $\mathbf{p}_0 = \mathbf{0}_{(4 \times 1)}$:

$$\mathbf{J}|_{\mathbf{p}_0} = \left(\left. \frac{\partial \mathbf{L}}{\partial \theta_1} \right|_{\mathbf{p}_0} \mid \left. \frac{\partial \mathbf{L}}{\partial \theta_2} \right|_{\mathbf{p}_0} \mid \left. \frac{\partial \mathbf{L}}{\partial \theta_3} \right|_{\mathbf{p}_0} \mid \left. \frac{\partial \mathbf{L}}{\partial \theta} \right|_{\mathbf{p}_0} \right).$$

By using the orthonormal representation to Plücker coordinates equation (17) and the Jacobian matrices for $SO(2)$ and $SO(3)$, as defined by equations (12,14,15,16), we obtain, after minor algebraic manipulations:

$$\mathbf{J}_{(6 \times 4)} = \begin{pmatrix} \mathbf{0}_{(3 \times 1)} & -\sigma_1 \mathbf{u}_3 & \sigma_1 \mathbf{u}_2 & -\sigma_2 \mathbf{u}_1 \\ \sigma_2 \mathbf{u}_3 & \mathbf{0}_{(3 \times 1)} & -\sigma_2 \mathbf{u}_1 & \sigma_1 \mathbf{u}_2 \end{pmatrix}. \quad (18)$$

Initialization. The initial guess is given by the Plücker coordinates $\mathbf{L}_0^\top \sim (\mathbf{a}_0^\top \mid \mathbf{b}_0^\top)$.

- Compute the orthonormal representation $(\mathbf{U}, \mathbf{W}) \in SO(3) \times SO(2)$ of \mathbf{L}_0 by QR decomposition $(\mathbf{a}_0 \mid \mathbf{b}_0) = \mathbf{U} \begin{pmatrix} \sigma_1 & \sigma_2 \\ \sigma_2 & -\sigma_1 \end{pmatrix}$ and set $\mathbf{W} = \begin{pmatrix} \sigma_1 & -\sigma_2 \\ \sigma_2 & \sigma_1 \end{pmatrix}$.
- The 4 optimization parameters are $\mathbf{p}^\top = (\boldsymbol{\theta}^\top \mid \theta)$ where the 3-vector $\boldsymbol{\theta}$ and the scalar θ are used to update \mathbf{U} and \mathbf{W} respectively.

Update. (i.e. one optimization step)

- Current line is $\mathbf{L}^\top \sim (w_{11}\mathbf{u}_1^\top \mid w_{21}\mathbf{u}_2^\top)$ and $\partial\mathbf{L}/\partial\mathbf{p}$ is given by equation (18).
- Compute \mathbf{p} by minimizing some criterion.
- Update \mathbf{U} and \mathbf{W} : $\mathbf{U} \leftarrow \mathbf{U}\mathbf{R}(\boldsymbol{\theta})$ and $\mathbf{W} \leftarrow \mathbf{W}\mathbf{R}(\theta)$.

Table 4: Elements for 3D line optimization using the minimal 4 parameters through the orthonormal representation.

Geometric interpretation. Each of the 4 above-defined update parameters \mathbf{p} has a geometric interpretation. Matrix \mathbf{W} encapsulates the ratio $\|\mathbf{a}\|/\|\mathbf{b}\|$, hence the distance d from the origin \mathbf{O} to \mathbf{L} . Thus, parameter θ acts on d . Matrix \mathbf{U} is related to a 3D coordinate frame attached to \mathbf{L} . Parameter θ_1 rotates \mathbf{L} around a circle with radius d , centered on \mathbf{O} , and lying on the plane defined by \mathbf{O} and \mathbf{L} . Parameter θ_2 rotates \mathbf{L} around a circle with radius d , centered on \mathbf{O} , and lying in a plane containing \mathbf{O} , the closest point \mathbf{Q} of \mathbf{L} to \mathbf{O} , and perpendicular to \mathbf{L} . Parameter θ_3 rotates \mathbf{L} around the axis defined by \mathbf{O} and \mathbf{Q} . For the last three cases, the angles of rotation are the parameters themselves. This interpretation allows to easily incorporate a priori knowledge while estimating a line. For example, to leave the direction of the line invariant, one may use the 2 update parameters θ_2 and θ , while to leave the distance of the line to the origin invariant, one may use the 3 update parameters $\boldsymbol{\theta}$. This allows to solve constrained line estimation cases, as summarized in the table below, indicating which update parameters to optimize in which case:

scenario	θ_1	θ_2	θ_3	θ
fixed direction		×		×
fixed normal to the plane formed with the origin	×			×
fixed distance to the origin	×	×	×	

6 Experimental Results

6.1 Simulated Data

Our simulated experimental setup consists of a set of cameras looking inwards at 3D lines randomly chosen in a sphere with a 1 meter radius. Cameras are spread widely around the sphere, at a distance of roughly 10 meters away from the centre of the sphere. We fix the focal length of the cameras to 1000 (in number of pixels). Note that this information is not used in the rest of the experiments. The end-points of all lines are projected in all views, where their positions are corrupted by an additive Gaussian noise. We vary the parameters of this setup to assess and compare the quality of the different estimators on various scene configurations.

We compare the 4 methods given in this paper: LIN, QLIN1, QLIN2 and MLE (bundle adjustment based on our orthonormal representation of 3D lines), as well as the method given in [11, §15.4.1], denoted by ‘MLE_HARTLEY’. This method consists in nonlinearly computing the trifocal tensor as well as reconstructed lines by minimizing the reprojection error (2) and parameterizing the 3D lines by two of their three images. We also compare QLIN2 to a direct Levenberg-Marquardt-based minimization of the reprojection error, dubbed NLIN: the two methods gave undistinguishable results in all our experiments. Note that most existing methods, *e.g.* [14, 21, 23, 27] can be applied only when camera calibration is available.

We measure the quality of an estimate using the *estimation error*, as described in [11, §4], which also provides the theoretical lower bound. The estimation error is equivalent to the value of the negative log likelihood (2) (*i.e.* the reprojection error).

The results are shown on graphs on figures 1 and 2. We observe that the different methods are always in the same order. Three distinct behaviours can be seen. Methods LIN and QLIN1 give similar results since they are subject to the same bias induced by ignoring the Plücker constraint until the final correction. Methods QLIN2 and NLIN are undistinguishable. They give better results than the biased methods. Finally, methods MLE and MLE_HARTLEY are hardly ever distinguishable. Their results are the best since they adjust the camera positions

along with the 3D line parameters.

In more details, we vary the added noise level from 0 to 2 pixels, while considering 20 lines and 3 views on figure 1 (a). One observes that, beyond 1 pixel noise, methods LIN and QLIN1 behave very badly. This is mainly due to the bias introduced by the Plücker correction procedure. Methods QLIN2, MLE and MLE_HARTLEY degrade gracefully as the noise level increases. Method QLIN2 gives reasonable results. Methods MLE and MLE_HARTLEY give undistinguishable results, very close to the theoretical lower bound.

We vary the number of lines from 15 to 60, while considering a 1 pixel noise and 3 views on figure 1 (b). Similar conclusions as for the previous experiment can be drawn, except for the fact, that when more than 30 lines are considered, methods LIN and QLIN1 give reasonable results. Also, methods MLE and MLE_HARTLEY give results undistinguishable from the theoretical lower bound when more than 45 lines are considered.

Figure 2 (a) shows the results when the number of images is varied from 3 to 12. The algorithms that do not optimize the cameras, namely LIN, QLIN1, QLIN2 and NLIN, have an error which increases with the number of images, whereas the bundle adjustment algorithms, namely MLE and MLE_HARTLEY, have an error which decreases. This is due to the fact that when the number of images increases, the initial camera estimation degrades, which is characteristic of the camera initialization algorithm.

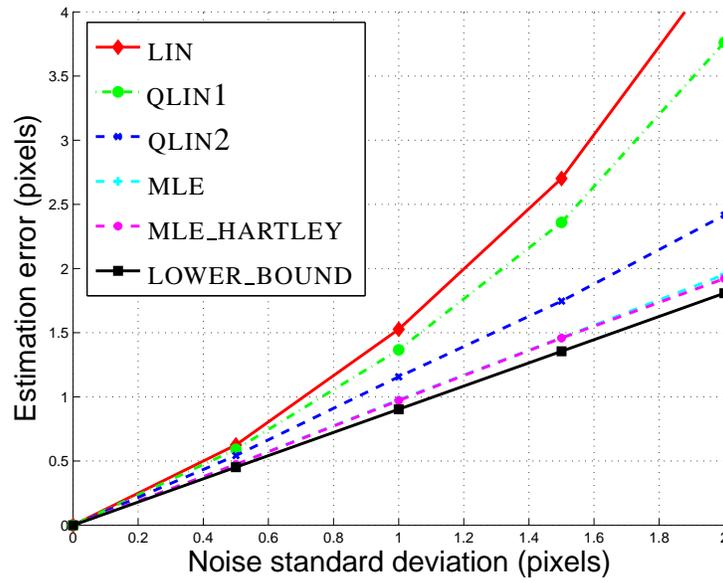
When the distance between the lines and the cameras increases, figure 2 (b) shows that the error decreases for all methods. This is explained by the fact that the cloud of 3D lines gets smaller and smaller in the images, which decrease the estimation error, but does not mean that the estimate is better.

We observed that the quasi-linear methods always converge within 5 iterations.

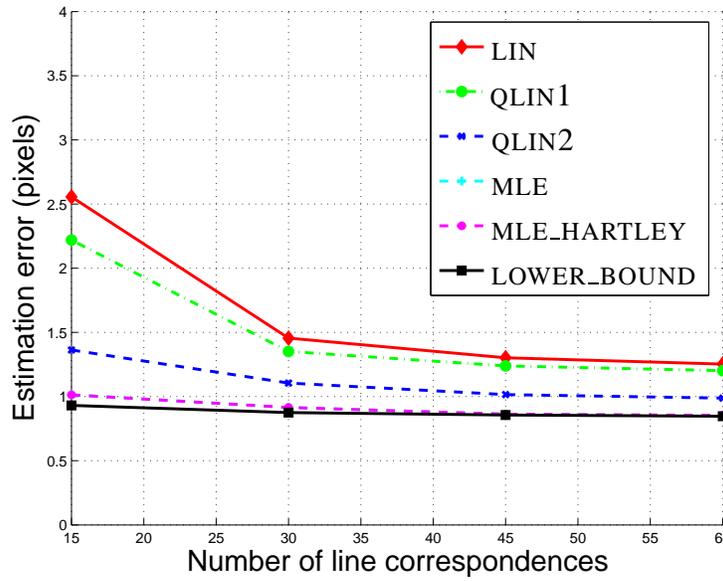
6.2 Real Data

We tested our algorithms on several image sequences. For two of them, we show results. We compared methods LIN, QLIN1, QLIN2 and MLE, since MLE_HARTLEY is for 3 views only.

We observed that QLIN1 generally needs more iterations to converge than QLIN2. This is due to the Plücker correction step that significantly modifies the estimate in QLIN1, while in QLIN2, since the constraint is linearized and enforced in the estimation, the correction applied to the estimate is less important.

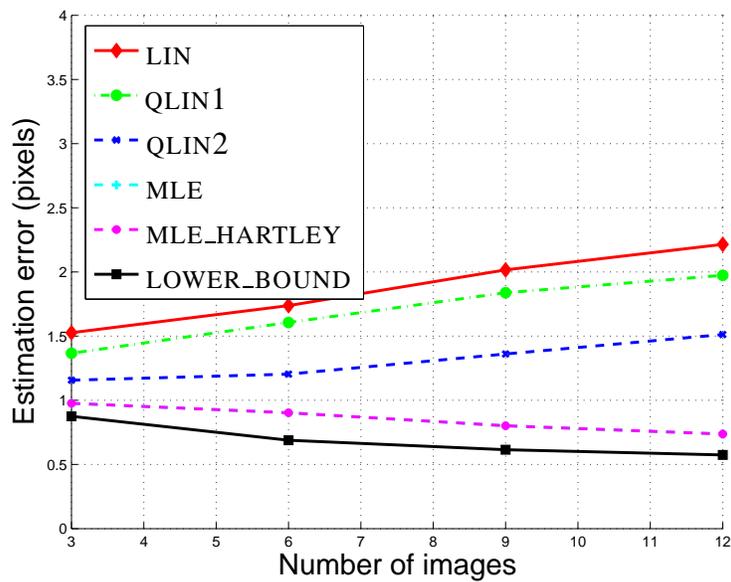


(a)

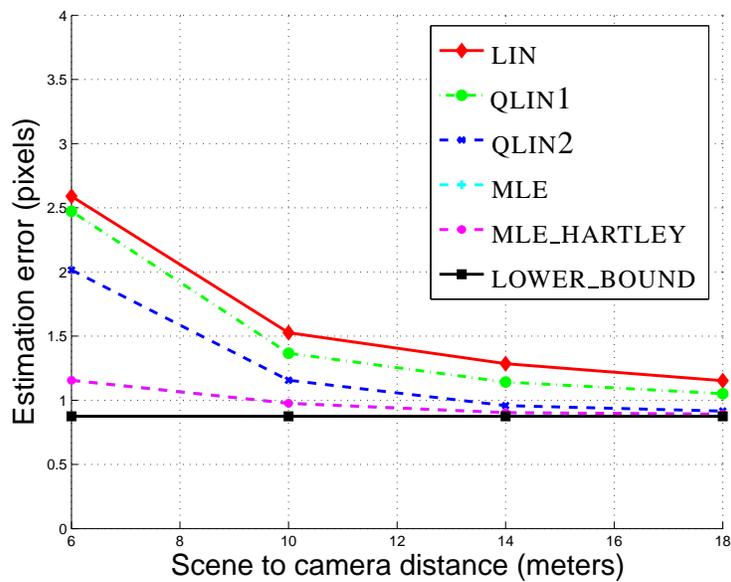


(b)

Figure 1: Estimation error for different methods when varying the variance of added noise on image end-points (a) and the number of lines considered (b).



(a)



(b)

Figure 2: Estimation error for different methods when varying the number of images (a) and the scene to camera distance (b).

The ‘books’ sequence. Figure 3 shows images from this 5-frame sequence. We provided 45 line correspondences by hand. Note that some of them are visible in two views only. We used these line correspondences to compute the trifocal tensor corresponding to each subsequence formed by triplets of consecutive images, using the linear method described in *e.g.* [11, §15.2]. We used method QLIN2 to reconstruct the lines associated with each triplet. We registered these subsequences by using the method given in [2]. At this point, we had a suboptimal guess of metric structure and motion. We further refined it using our triangulation algorithms, to reconstruct each line by taking into account all of its images. The corresponding estimation errors are, respectively for LIN, QLIN1 and QLIN2, 2.30, 2.27 and 1.43 pixels. Note the significant improvement of QLIN2 compared to the biased methods LIN and QLIN1. Methods QLIN1 and QLIN2 respectively took 4 and 3 iterations to converge.

We used the result of QLIN2 to initialize our Maximum Likelihood estimator for structure and motion based on the proposed orthonormal representation together with a metric parameterization of the camera motion, which ends up with a 0.9 pixel estimation error.

For each estimation, we reconstructed the end-points corresponding to the first view (shown on the left of figure 3). The Maximum Likelihood end-points are given by orthogonally projecting their images onto the image of the corresponding line.

These results are visible on figure 4. Note the significant improvement of method MLE over methods LIN, QLIN1 and QLIN2. The lines predicted by MLE and the original lines are undistinguishable. Figure 5 shows the cameras and lines reconstructed by MLE. There is visually no difference with the reconstruction provided by algorithm QLIN2, but that reconstructions provided by LIN and QLIN1 appear distorted.

The ‘laptop’ sequence. Figure 6 shows sample images for the 8-frame ‘laptop’ sequence, overlaid with the 40 manually-entered line correspondences. We performed 3D reconstruction by applying the same algorithms as for the ‘books’ sequence. We obtained the following estimation errors for the triangulation algorithms, namely LIN: 1.34 pixels, QLIN1: 1.29 pixels and QLIN2: 1.04 pixels. Methods QLIN1 and QLIN2 took respectively 7 and 5 iterations to converge. For the bundle adjustment algorithms, we obtained an estimation error of 0.82 pixels. Figure 7 shows snapshots of the reconstructed 3D models.

These results show that accurate reconstructed models can be obtained on real

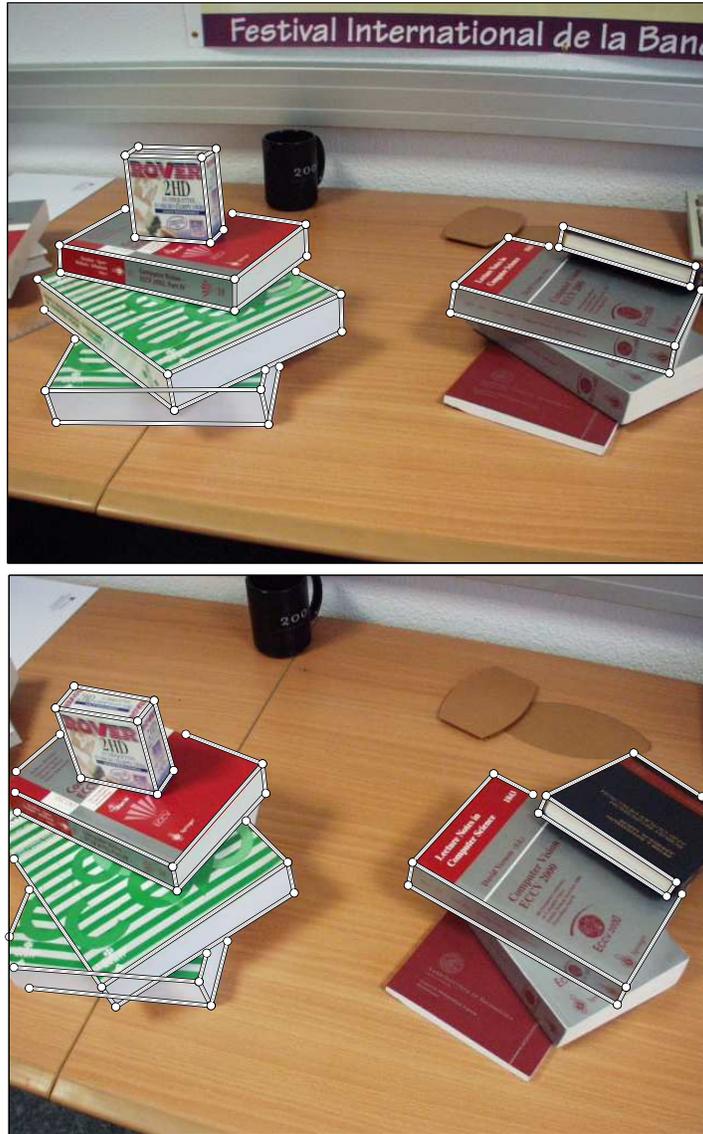
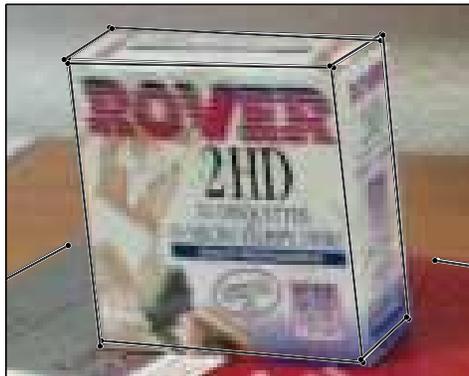


Figure 3: Sample images out of the 5-frame ‘books’ sequence overlaid with manually-provided lines. Note that the optical distortion is not corrected.



LIN & QLIN1



QLIN2



MLE

Figure 4: Zoom on some original (white) and reprojected lines (black) for the 'books' sequence for different methods.

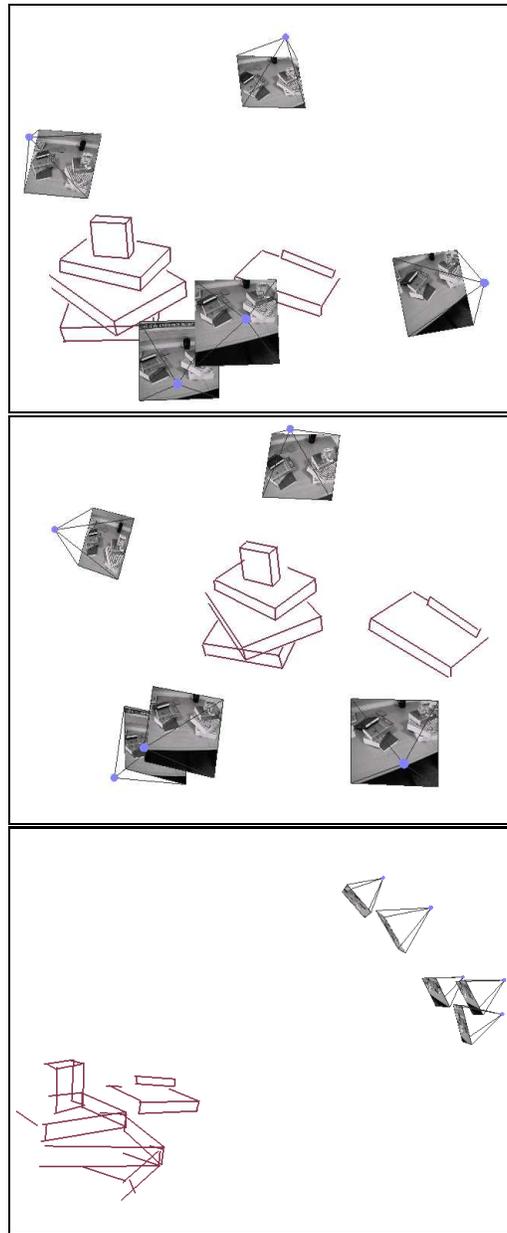


Figure 5: Snapshots of the cameras and lines reconstructed by method MLE for the ‘books’ sequence. The images shown in figure 3 correspond to the top- and bottom-most cameras.

images taken by amateur digital cameras. They also show the importance of running a final bundle adjustment after initial triangulation.

7 Conclusion

We addressed the problem of structure and motion recovery from line correspondences across multiple views.

First, we proposed an optimal triangulation algorithm. Given camera motion, the Plücker coordinates of the 3D lines are estimated by minimizing the reprojection error. The algorithm relies on an iteratively reweighted least squares scheme. We linearized the bilinear Plücker constraint to incorporate it up to first order in the estimation process. A Plücker correction procedure is proposed to find the nearest Plücker coordinates to a given 6-vector.

Second, we proposed the orthonormal representation of 3D lines, which allows nonlinear optimization with the minimal 4 parameters within an unconstrained optimization engine, contrarily to previously proposed overparameterizations. This representation is well-conditioned and allows analytic differentiation.

Experimental results on simulated and real data show that the standard linear method and its naive bias-corrected extension perform very badly in many cases and should only be used to initialize a nonlinear estimator. Our bias-corrected algorithm including the Plücker constraint performs as well as direct Levenberg-Marquardt-based triangulation. It is therefore a good solution to initialize subsequent bundle adjustment. Based on our orthonormal representation, bundle adjustment gives results close to the theoretical lower bound and undistinguishable from the three-view maximum likelihood estimator of [11, §15.4.1], while being usable with any number of views.

References

- [1] 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.
- [2] A. Bartoli and P. Sturm. The 3D line motion matrix and alignment of line reconstructions. *International Journal of Computer Vision*, 57(3), May/June 2004.



Figure 6: Sample images out of the 8-frame ‘laptop’ sequence overlaid with manually-provided lines. Note that the optical distortion is not corrected.

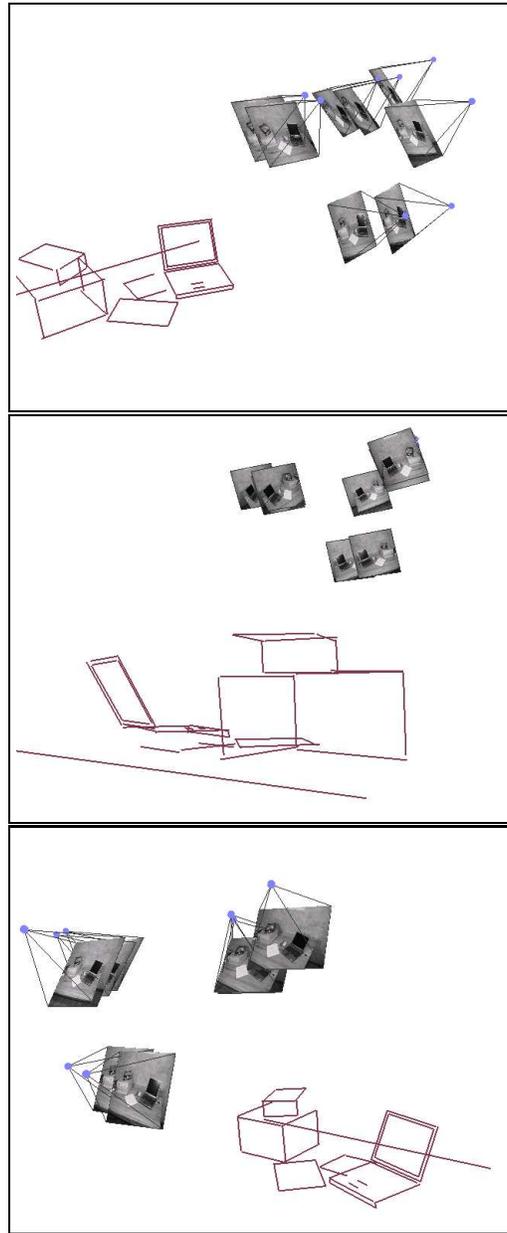


Figure 7: Snapshots of the cameras and lines reconstructed by method MLE for the 'laptop' sequence.

- [3] J. Denavit and R. S. Hartenberg. A kinematic notation for lower pair mechanisms based on matrices. *ASME Journal of Applied Mechanics*, 22:215–221, 1955.
- [4] N. Ayache et Faugeras. Maintaining representations of the environment of a mobile robot. *IEEE Transactions on Robotics and Automation*, 5(6):804–819, 1989.
- [5] O. Faugeras and B. Mourrain. On the geometry and algebra of the point and line correspondences between n images. In *Proceedings of the 5th International Conference on Computer Vision, Cambridge, Massachusetts, USA*, pages 951–956, June 1995.
- [6] A.W. Fitzgibbon and A. Zisserman. Automatic camera recovery for closed or open image sequences. In *European Conference on Computer Vision*, pages 311–326, June 1998.
- [7] G.H. Golub and C.F. van Loan. *Matrix Computation*. The Johns Hopkins University Press, Baltimore, 1989.
- [8] A. Habib. Motion parameter estimation by tracking stationary three-dimensional straight lines in image sequences. *International Archives of Photogrammetry and Remote Sensing*, 53, 1998.
- [9] A. Habib, M. Morgan, and Y.-R. Lee. Bundle adjustment with self-calibration using straight lines. *Photogrammetric Record*, October 2002.
- [10] R.I. Hartley. Lines and points in three views and the trifocal tensor. *International Journal of Computer Vision*, 22(2):125–140, 1997.
- [11] R.I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, June 2000.
- [12] S. A. Hayati and M. Mirmirani. Improving the absolute positioning accuracy of robot manipulators. *Journal of Robotic Systems*, 2(4):397–441, 1985.
- [13] K. Kanatani. *Statistical Optimisation for Geometric Computation: Theory and Practice*. Elsevier Science, 1996.
- [14] Y. Liu and T.S. Huang. A linear algorithm for motion estimation using straight line correspondences. *Computer Vision, Graphics and Image Processing*, 44(1):35–57, October 1988.

- [15] D. Martinec and T. Pajdla. Line reconstruction from many perspective images by factorization. In *Proceedings of the Conference on Computer Vision and Pattern Recognition, Madison, Wisconsin, USA*, volume I, pages 497–502. IEEE Computer Society Press, June 2003.
- [16] D. C. Mulawa and E. M. Mikhail. Photogrammetric treatment of linear features. *International Archives of Photogrammetry and Remote Sensing*, 27:383–393, 1988.
- [17] H. Pottmann, M. Hofer, B. Odehnl, and J. Wallner. Line geometry for 3D shape understanding and reconstruction. In *Proceedings of the European Conference on Computer Vision*, 2004.
- [18] L. Quan and T. Kanade. Affine structure from line correspondences with uncalibrated affine cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(8):834–845, August 1997.
- [19] K. Roberts. A new representation for a line. In *Proceedings of the Conference on Computer Vision and Pattern Recognition, San Diego, California, USA*, pages 635–640, 1988.
- [20] Y. Seo and K. S. Hong. Sequential reconstruction of lines in projective space. In *Proceedings of the 13th International Conference on Pattern Recognition, Vienna, Austria*, pages 503–507, August 1996.
- [21] M. Spetsakis and J. Aloimonos. Structure from motion using line correspondences. *International Journal of Computer Vision*, 4:171–183, 1990.
- [22] G. P. Stein and A. Shashua. On degeneracy of linear reconstruction from three views: Linear line complex and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(3):244–251, 1999.
- [23] C.J. Taylor and D.J. Kriegman. Structure and motion from line segments in multiple images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(11):1021–1032, November 1995.
- [24] A. Tommaselli and J. Lugnani. An alternative mathematical model to collinearity equations using straight features. *International Archives of Photogrammetry and Remote Sensing*, 27:765–774, 1998.

- [25] B. Triggs. Factorization methods for projective structure and motion. In *Proceedings of the Conference on Computer Vision and Pattern Recognition, San Francisco, California, USA*, pages 845–851, 1996.
- [26] T. Viéville, Q.T. Luong, and O.D. Faugeras. Motion of points and lines in the uncalibrated case. *International Journal of Computer Vision*, 17(1), 1995.
- [27] J. Weng, T.S. Huang, and N. Ahuja. Motion and structure from line correspondences: Closed-form solution, uniqueness, and optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(3):318–336, 1992.