Robust techniques for the estimation of structure from motion in the uncalibrated case

Michael J. Brooks,[1] Wojciech Chojnacki,[1] Anton van den Hengel,[1]
Luis Baumela[2]

[1] Department of Computer Science, University of Adelaide
Adelaide, SA 5005, Australia
{mjb,wojtek,hengel}@cs.adelaide.edu.au

[2] Departamento de Inteligencia Artificial, Universidad Politécnica de Madrid
Campus de Montegancedo s/n, 28660 Boadilla del Monte, Madrid, Spain
lbaumela@fi.upm.es

Abstract. Robust techniques are developed for determining structure from motion in the uncalibrated case. The structure recovery is based on previous work [7] in which it was shown that a camera undergoing unknown motion and having an unknown, and possibly varying, focal length can be self-calibrated via closed-form expressions in the entries of two matrices derivable from an instantaneous optical flow field. Critical to the recovery process is the obtaining of accurate numerical estimates, up to a scalar factor, of these matrices in the presence of noisy optical flow data. We present techniques for the determination of these matrices via least-squares methods, and also a way of enforcing a dependency constraint that is imposed on these matrices. A method for eliminating outlying flow vectors is also given. Results of experiments with real-image sequences are presented that suggest that the approach holds promise.

1 Introduction

In this paper we present robust techniques for estimating structure from motion in the uncalibrated case. The structure recovery is based on a method for self-calibrating a single moving camera from instantaneous optical flow, as developed in [7]. Here, self-calibration amounts to automatically determining the unknown instantaneous ego-motion and intrinsic parameters of the camera, and is analogous to self-calibrating a stereo vision set-up from corresponding points [9, 17]. The method of self-calibration rests on an equation that we term the *differential epipolar equation for uncalibrated optical flow*, which relates optical flow to the ego-motion and intrinsic parameters of the camera. This equation incorporates two matrices which encode information about the ego-motion and internal geometry of the camera. Any sufficiently large subset of an optical flow field determines the composite ratio of some of the entries of these matrices, and, under certain assumptions, the moving camera can be self-calibrated by means of closed-form expressions evolved from this ratio.

Once self-calibration is completed, scene reconstruction, up to a scalar factor, can be carried out by employing a procedure, also described in [7], based on the results of self-calibration and the optical flow data. In this way, the structure recovery is essentially reduced to the determination of the critical composite ratio.

In this paper, we present various least-squares techniques and an outlier rejection scheme that facilitate robust estimation of the critical composite ratio. The self-calibration and reconstruction methods are tested on sparse optical flow fields derived from real-world image sequences of a calibration grid and an office.

The present work is most closely related to that of Viéville and Faugeras [25], which was the first to tackle the problem of interpreting optical flow arising from an uncalibrated camera. The approach of Ohta and Kanatani [18] is also related to this work, but is concerned with the calibrated case and hence is not immediately applicable to our more general situation. For examples of other work dealing with the ego-motion of a calibrated camera see [10, 11]. A contrasting approach is given by Beardsley et al. [5], involving the computation of projective and affine (rather than Euclidean) structure from motion. Other related papers include [1–4, 15, 19, 23, 26, 27].

2 Differential epipolar equation and a cubic constraint

Consider a camera with an associated coordinate frame such that the origin of the frame coincides with the camera's optical centre, two basis vectors span the focal plane, and the other basis vector passes through the optical axis. Suppose that the camera undergoes smooth motion. Denote by $\boldsymbol{v}$ and $\boldsymbol{\omega}$ the camera's instantaneous *translational velocity* and instantaneous *angular velocity* with respect to the camera frame. Let P be a static point in space, let $\boldsymbol{x} = [x_1, x_2, x_3]^T$ be the coordinates of the vector connecting C with P in the vector basis of the camera frame, and let $\boldsymbol{p} = [p_1, p_2, p_3]^T$ be the coordinates of the perspective projection of P, through C, onto the image plane $\{\boldsymbol{x} \in \mathbb{R}^3 \mid x_3 = -f\}$, relative to the camera frame; here f is the focal length. To account for the characteristics of the camera, we adopt a separate coordinate frame in the image plane. Let $[m_1, m_2]^T$ be the coordinates of the image of P in the vector basis of this frame. If we append to $[m_1, m_2]^T$ an extra entry equal to 1 to yield the vector $\boldsymbol{m} = [m_1, m_2, 1]^T$, then the relation between $\boldsymbol{p}$ and $\boldsymbol{m}$ can be written as $\boldsymbol{p} = \boldsymbol{A}\boldsymbol{m}$, where $\boldsymbol{A}$ is a 3×3, possibly time-dependent, invertible matrix called the *intrinsic-parameter matrix*. As the camera moves, the position of P relative to the image frame will change accordingly and will be recorded in the function $t \mapsto \boldsymbol{m}(t)$. The associated function $t \mapsto [\boldsymbol{m}^T(t), \dot{\boldsymbol{m}}(t)^T]^T$ satisfies the *differential epipolar equation for uncalibrated optical flow*

$$\boldsymbol{m}^T \boldsymbol{W} \dot{\boldsymbol{m}} + \boldsymbol{m}^T \boldsymbol{C} \boldsymbol{m} = 0, \tag{1}$$

where $\boldsymbol{C}$ and $\boldsymbol{W}$ are 3×3 matrices explicitly expressible in terms of $\boldsymbol{v}$, $\boldsymbol{\omega}$, $\boldsymbol{A}$ and $\dot{\boldsymbol{A}}$. The function $t \mapsto [\boldsymbol{m}^T(t), \dot{\boldsymbol{m}}(t)^T]^T$ is a single trajectory of an *optical flow field*. An optical flow field at a given instant t is a set of vectors of the form $[\boldsymbol{m}(t)^T, \dot{\boldsymbol{m}}(t)^T]^T$ that describe the instantaneous position and velocity of the images of various elements of the scene. More precisely, such a set constitutes a *true image motion field* which, as is usual, we assume to correspond to an *observed image velocity field* or optical flow field. Eq. (1) relates optical flow, ego-motion and camera parameters. A similar constraint, termed the *first-order expansion of the fundamental motion equation*, was derived using quite different means by Viéville and Faugeras [25]. In contrast with the

above, however, it takes the form of an approximation rather than a strict equality. A simpler form of (1), applying to the calibrated case, was first derived by Maybank [16] and was subsequently exploited by many authors (cf. [2, 3, 12]).

The matrix $\boldsymbol{C}$ is symmetric, and hence it is uniquely determined by the entries $c_{11}, c_{12}, c_{13}, c_{22}, c_{23}, c_{33}$. The matrix $\boldsymbol{W}$ is antisymmetric, and so it is uniquely determined by the entries w_{12}, w_{13}, w_{23}. Define the *joint projective form* $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ of $\boldsymbol{C}$ and $\boldsymbol{W}$ to be the composite ratio $(c_{11} : c_{12} : c_{13} : c_{22} : c_{23} : c_{33} : w_{12} : w_{13} : w_{23})$. Eq. (1) is essentially a constraint on $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$, the equation remaining valid whenever $\boldsymbol{C}$ and $\boldsymbol{W}$ are multiplied by a common scalar factor. It turns out that

$$\boldsymbol{w}^T \boldsymbol{C} \boldsymbol{w} = 0, \tag{2}$$

where $\boldsymbol{w} = [-w_{23}, w_{13}, -w_{12}]^T$. Again (2) can be interpreted as a constraint on $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$—more precisely, as a cubic constraint. It implies that $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ lies on a seven-dimensional manifold that does not depend on the optical flow entering (1), a fact already noted in [25].

As indicated in the Introduction, the differential epipolar equation (1) forms the basis for a method of self-calibration. We use this equation along with (2) to determine $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ from noisy optical flow. Knowing $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ allows in turn recovery of some of the parameters describing the ego-motion and internal geometry of the camera. For details, the reader is referred to [7].

3 Estimating $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$

We now turn to the problem of estimating $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ from optical flow that is perturbed by noise and outliers.

Ohta and Kanatani [18] (see also [13, Chap. 12]) present a method for estimating ego-motion and structure from noise-contaminated calibrated optical flow. It relies on a model of noise in the measurement of flow velocity adopted from the outset, and takes the form of an optimal estimation algorithm accompanied by a measure of the reliability of the estimates. Our problem is not directly amenable to Ohta and Kanatani's sophisticated analysis for two reasons. First, we deal with uncalibrated flow, and, second, we cannot assume the position of a flow vector to be unperturbed. Our goal is to develop methods capable of dealing with sparse optical flow obtained via the tracking of features, and not only with dense optical flow calculated across a grid. Nevertheless, we draw upon various ideas present in [18].

Let $\mathcal{S} = \{[\boldsymbol{m}_i{}^T, \dot{\boldsymbol{m}}_i{}^T]^T \mid i = 1, \ldots, n\}$ with $n \geq 7$ be a data set representing measurements of a portion of instantaneous optical flow. We shall consider ways in which estimates of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ can be derived from $\mathcal{S}$. Note that a pair of matrices $(\boldsymbol{C}, \boldsymbol{W})$ is uniquely determined by the vector

$$\boldsymbol{\theta}_{\boldsymbol{C}, \boldsymbol{W}} = [c_{11}, c_{12}, c_{13}, c_{22}, c_{23}, c_{33}, w_{12}, w_{13}, w_{23}]^T.$$

When $(\boldsymbol{C}, \boldsymbol{W})$ is understood, $\boldsymbol{\theta}_{\boldsymbol{C}, \boldsymbol{W}}$ will often be contracted to $\boldsymbol{\theta}$. Furthermore, observe that the ratio $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ can be identified, to within sign, with the vector $\boldsymbol{\theta}_{\boldsymbol{C}, \boldsymbol{W}}$ satisfying the normalisation condition $\|\boldsymbol{\theta}_{\boldsymbol{C}, \boldsymbol{W}}\| = 1$, where $\|\boldsymbol{\theta}_{\boldsymbol{C}, \boldsymbol{W}}\| = ({c_{11}}^2 + {c_{12}}^2 +$

$\cdots + {w_{23}}^2)^{1/2}$. It is therefore always possible to represent an estimate of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ as a normalised vector $\widehat{\boldsymbol{\theta}} = \boldsymbol{\theta}_{\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}}$ for some pair $(\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}})$. Here, adhering to a custom in statistics, we use $\widehat{x}$ to denote an estimate of x.

3.1 Seven-point estimator

If $n = 7$, then an estimate of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ can be obtained by solving a system of seven equations

$$\boldsymbol{m}_i^T \boldsymbol{W} \dot{\boldsymbol{m}}_i + \boldsymbol{m}_i^T \boldsymbol{C} \boldsymbol{m}_i = 0 \tag{3}$$

and equation (2). All these equations are homogeneous in the components of $\boldsymbol{\theta}$, and effectively provide seven constraints for the ratio $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$. Note that $\boldsymbol{m}^T \boldsymbol{W} \dot{\boldsymbol{m}} + \boldsymbol{m}^T \boldsymbol{C} \boldsymbol{m}$ depends linearly on $\boldsymbol{\theta}$, being expressible as

$$\boldsymbol{m}^T \boldsymbol{W} \dot{\boldsymbol{m}} + \boldsymbol{m}^T \boldsymbol{C} \boldsymbol{m} = \boldsymbol{\theta}^T \boldsymbol{u}_{\boldsymbol{m}, \dot{\boldsymbol{m}}}, \tag{4}$$

where

$$\boldsymbol{u}_{\boldsymbol{m}, \dot{\boldsymbol{m}}} = \begin{bmatrix} m_1^2 \\ 2m_1 m_2 \\ 2m_1 m_3 \\ m_2^2 \\ 2m_2 m_3 \\ m_3^3 \\ m_1 \dot{m}_2 - m_2 \dot{m}_1 \\ m_1 \dot{m}_3 - m_3 \dot{m}_1 \\ m_2 \dot{m}_3 - m_3 \dot{m}_2 \end{bmatrix}.$$

Therefore the space of solutions to system (3) is spanned by two normalised linearly independent solutions $\widehat{\boldsymbol{\theta}}_1$ and $\widehat{\boldsymbol{\theta}}_2$. Accordingly, an un-normalised solution $\widehat{\boldsymbol{\eta}}$ to the full system of equations can be represented as $\alpha \widehat{\boldsymbol{\theta}}_1 + (1 - \alpha) \widehat{\boldsymbol{\theta}}_2$ for some scalar parameter α. Substituting $\alpha \widehat{\boldsymbol{\theta}}_1 + (1 - \alpha) \widehat{\boldsymbol{\theta}}_2$ into (2) leads to a cubic constraint on α. This equation has either one or three real solutions, which in turn give rise to one or three normalised estimates $\widehat{\boldsymbol{\theta}} = \widehat{\boldsymbol{\eta}} / \|\widehat{\boldsymbol{\eta}}\|$.

3.2 Least squares estimator based on algebraic distances

If $n \geq 8$, then the linear homogeneous equations forming system (3) provide $n - 1 \geq 7$ constraints for $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$. We can use these equations to estimate $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ without immediate recourse to the cubic constraint (2). The redundancy in system (3) suggests a least squares solution. In order to develop such a solution, a cost function has to be specified. Given a single piece of data $[\boldsymbol{m}^T, \dot{\boldsymbol{m}}^T]^T$ and a parameter $\boldsymbol{\theta}$ to be fit, we take the *algebraic distance* $|\boldsymbol{m}^T \boldsymbol{W} \dot{\boldsymbol{m}} + \boldsymbol{m}^T \boldsymbol{C} \boldsymbol{m}|$ between $[\boldsymbol{m}^T, \dot{\boldsymbol{m}}^T]^T$ and the hypersurface

$$\mathcal{M}_{\boldsymbol{\theta}} = \{ [\boldsymbol{n}^T, \dot{\boldsymbol{n}}^T]^T \mid \boldsymbol{n}^T \boldsymbol{W} \dot{\boldsymbol{n}} + \boldsymbol{n}^T \boldsymbol{C} \boldsymbol{n} = 0 \}$$

for the *residual* or the measure of the agreement between $[\boldsymbol{m}^T, \dot{\boldsymbol{m}}^T]^T$ and $\boldsymbol{\theta}$. The *residual square error*

$$J_1(\boldsymbol{\theta}; \mathcal{S}) = \sum_{i=1}^{n} |\boldsymbol{m}_i^T \boldsymbol{W} \dot{\boldsymbol{m}}_i + \boldsymbol{m}_i^T \boldsymbol{C} \boldsymbol{m}_i|^2$$

then serves as a natural cost function. This choice can be viewed as a generalisation of the criterion used by Bookstein [6] for conic fitting. The J_1-based *best-fit estimate* of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ is a unique (up to a sign) normalised vector $\widehat{\boldsymbol{\theta}}$ such that

$$J_1(\widehat{\boldsymbol{\theta}}; \mathcal{S}) = \min_{||\boldsymbol{\theta}|| = 1} J_1(\boldsymbol{\theta}; \mathcal{S}). \tag{5}$$

With errors ascribed to the $\boldsymbol{u}_{\boldsymbol{m}_i, \dot{\boldsymbol{m}}_i}$, the J_1-based minimisation can be viewed as a method of *orthogonal regression*, also termed *total least squares regression* or *principal component regression*. The estimate $\widehat{\boldsymbol{\theta}}$ can be computed explicitly. In fact, in view of (4),

$$|\boldsymbol{m}^T \boldsymbol{W} \dot{\boldsymbol{m}} + \boldsymbol{m}^T \boldsymbol{C} \boldsymbol{m}|^2 = \boldsymbol{\theta}^T \boldsymbol{M}_{\boldsymbol{m}, \dot{\boldsymbol{m}}} \boldsymbol{\theta}, \tag{6}$$

where

$$\boldsymbol{M}_{\boldsymbol{m}, \dot{\boldsymbol{m}}} = \boldsymbol{u}_{\boldsymbol{m}, \dot{\boldsymbol{m}}} \boldsymbol{u}_{\boldsymbol{m}, \dot{\boldsymbol{m}}}^T$$

and further

$$J_1(\boldsymbol{\theta}; \mathcal{S}) = \boldsymbol{\theta}^T \boldsymbol{D} \boldsymbol{\theta},$$

where

$$\boldsymbol{D} = \sum_{i=1}^{n} \boldsymbol{M}_{\boldsymbol{m}_i, \dot{\boldsymbol{m}}_i}.$$

The gradient of J_1 with respect to $\boldsymbol{\theta}$, traditionally written as a row vector,

$$\nabla_{\boldsymbol{\theta}} J_1(\boldsymbol{\theta}; \mathcal{S}) = \left[\frac{\partial J_1}{\partial c_{11}}(\boldsymbol{\theta}; \mathcal{S}), \frac{\partial J_1}{\partial c_{12}}(\boldsymbol{\theta}; \mathcal{S}), \ldots, \frac{\partial J_1}{\partial w_{23}}(\boldsymbol{\theta}; \mathcal{S}) \right]$$

is given by

$$[\nabla_{\boldsymbol{\theta}} J_1(\boldsymbol{\theta}; \mathcal{S})]^T = 2 \boldsymbol{D} \boldsymbol{\theta}. \tag{7}$$

The minimisation condition (5) implies that

$$[\nabla_{\boldsymbol{\theta}} J_1(\widehat{\boldsymbol{\theta}}; \mathcal{S})]^T = 2 \lambda \widehat{\boldsymbol{\theta}}$$

for some (unknown) scalar *Lagrange multiplier* λ. This together with (7) shows that $\widehat{\boldsymbol{\theta}}$ is an eigenvector of $\boldsymbol{D}$ with eigenvalue λ. If ξ is a normalised eigenvector of $\boldsymbol{D}$ with eigenvalue μ, then

$$\mu = \mu \xi^T \xi = J_1(\xi; \mathcal{S}) \geq J_1(\widehat{\boldsymbol{\theta}}; \mathcal{S}) = \widehat{\boldsymbol{\theta}}^T \lambda \widehat{\boldsymbol{\theta}} = \lambda,$$

showing that λ is the smallest eigenvalue of $\boldsymbol{D}$. Therefore $\widehat{\boldsymbol{\theta}}$ is recognised as the normalised eigenvector (or the principal component) corresponding to the smallest eigenvalue of $\boldsymbol{D}$. Both $\widehat{\boldsymbol{\theta}}$ and λ can be efficiently calculated by employing the method of singular value decomposition.

3.3 Least squares estimator based on Euclidean distances

The algebraic distance underlying the J_1-based estimation method has no geometric significance. In contrast, the expression

$$\delta(\boldsymbol{\theta}; \boldsymbol{m}, \dot{\boldsymbol{m}}) = \frac{|\boldsymbol{m}^T \boldsymbol{W} \dot{\boldsymbol{m}} + \boldsymbol{m}^T \boldsymbol{C} \boldsymbol{m}|}{\sqrt{\|2\boldsymbol{C}\boldsymbol{m} + \boldsymbol{W}\dot{\boldsymbol{m}}\|^2 + \|\boldsymbol{W}\boldsymbol{m}\|^2}} \tag{8}$$

is geometrically meaningful being an approximation of the *Euclidean distance* between $[\boldsymbol{m}^T, \dot{\boldsymbol{m}}^T]^T$ and $\mathcal{M}_{\boldsymbol{\theta}}$. More generally, if $\mathcal{M}$ is a hypersurface in $\mathbb{R}^k$ defined by $\mathcal{M} = \{x \in \mathbb{R}^k \mid f(x) = 0\}$ and $z \in \mathbb{R}^k$ is a point close to $\mathcal{M}$, then the Euclidean distance between z and $\mathcal{M}$ is to a first order approximation equal to $|f(z)|/\|\nabla f(z)\|$—a fact exploited in vision-related statistical formulations first by Sampson [22] and later by a number of authors (cf. [13, 14, 24, 28]). Adopting $\delta(\boldsymbol{\theta}; \boldsymbol{m}, \dot{\boldsymbol{m}})$ as a statistically more adequate residual, we introduce a cost function

$$J_2(\boldsymbol{\theta}; \mathcal{S}) = \sum_{i=1}^{n} |\delta(\boldsymbol{\theta}; \boldsymbol{m}_i, \dot{\boldsymbol{m}}_i)|^2.$$

A similar cost function can be devised starting from the principles underpinning Kanatani's approach to statistical optimisation (cf. [13, Chap. 12] and [18]). With errors ascribed to the $[\boldsymbol{m}_i{}^T, \dot{\boldsymbol{m}}_i{}^T]^T$ rather than to the $\boldsymbol{u}_{\boldsymbol{m}_i, \dot{\boldsymbol{m}}_i}$, the J_2-based minimisation can be viewed as a variant of the method of orthogonal regression. Due to the complicated way in which $\boldsymbol{C}$ and $\boldsymbol{W}$ enter J_2, it is not clear whether the J_2-based estimate of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ can be given an explicit form. To evolve a numerical estimate, various standard iterative techniques can be employed like the Newton-Raphson method, the steepest descent method, the Levenberg-Marquardt method, the downhill simplex method, and various direction set methods of Powell (cf. [20]). Alternatively, one might attempt to develop algorithms especially adopted to cope with the task of minimisation of J_2. Two such algorithms are proposed next.

3.4 Iteratively reweighted least squares estimator

Let $\widehat{\boldsymbol{\theta}}$ be a normalised estimate corresponding to a pair $(\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}})$, and let

$$J_3(\boldsymbol{\theta}; \widehat{\boldsymbol{\theta}}, \mathcal{S}) = \sum_{i=1}^{n} |\widetilde{\delta}(\boldsymbol{\theta}; \widehat{\boldsymbol{\theta}}, \boldsymbol{m}_i, \dot{\boldsymbol{m}}_i)|^2,$$

where

$$\widetilde{\delta}(\boldsymbol{\theta}; \widehat{\boldsymbol{\theta}}, \boldsymbol{m}, \dot{\boldsymbol{m}}) = \frac{|\boldsymbol{m}^T \boldsymbol{W} \dot{\boldsymbol{m}} + \boldsymbol{m}^T \boldsymbol{C} \boldsymbol{m}|}{\sqrt{\|2\widehat{\boldsymbol{C}}\boldsymbol{m} + \widehat{\boldsymbol{W}}\dot{\boldsymbol{m}}\|^2 + \|\widehat{\boldsymbol{W}}\boldsymbol{m}\|^2}}.$$

The denominator in the expression for $\widetilde{\delta}$ does not depend on $\boldsymbol{\theta}$, and so minimisation of $J_3(\boldsymbol{\theta}; \mathcal{S})$ subject to the constraint $\|\boldsymbol{\theta}\|^2 = 1$ leads to an estimator falling into the category of weighted least squares techniques. Let

$$\boldsymbol{D}_{\widehat{\boldsymbol{\theta}}} = \sum_{i=1}^{n} \lambda_i(\widehat{\boldsymbol{\theta}}) \boldsymbol{M}_{\boldsymbol{m}_i, \dot{\boldsymbol{m}}_i} \tag{9}$$

where

$$\lambda_i(\widehat{\boldsymbol{\theta}}) = \left(\|2\widehat{\boldsymbol{C}}\boldsymbol{m}_i + \widehat{\boldsymbol{W}}\dot{\boldsymbol{m}}_i\|^2 + \|\widehat{\boldsymbol{W}}\boldsymbol{m}_i\|^2 \right)^{-1}.$$

Employing Lagrange multipliers, we verify at once that the J_3-based estimate can be identified with the eigenvector of $\boldsymbol{D}_{\widehat{\boldsymbol{\theta}}}$ corresponding to the smallest eigenvalue. Using this observation, we can now propose the following iteratively reweighted least squares estimator that seeks to minimise J_2:

1. Compute $\widehat{\boldsymbol{\theta}}_0$ using least-square fitting based on J_1.
2. Assuming that $\widehat{\boldsymbol{\theta}}_{k-1}$ is known, compute the matrix $\boldsymbol{D}_{\widehat{\boldsymbol{\theta}}_{k-1}}$.
3. Compute a normalised eigenvector of $\boldsymbol{D}_{\widehat{\boldsymbol{\theta}}_{k-1}}$ corresponding to the smallest eigenvalue and take this eigenvector for $\widehat{\boldsymbol{\theta}}_k$.
4. If $\widehat{\boldsymbol{\theta}}_k$ is sufficiently close to $\widehat{\boldsymbol{\theta}}_{k-1}$, then terminate the procedure; otherwise increment k and return to Step 2.

3.5 Modified iteratively reweighted least squares estimator

Typically, an estimate evolved by the algorithm in the previous subsection will be different from the sought-after estimate based on minimisation of J_2—the iteratively reweighted least squares techniques are susceptible to statistical bias (cf. [13, Chap. 9], [29]). In an attempt to eliminate this shortcoming, we put forward a modified algorithm. While in some aspects the proposed technique resembles Kanatani's technique of renormalisation [13, Chap. 9], it is different from the latter in that it is formulated in a purely deterministic, probability-free fashion, and that it impinges on the standard, rather than generalised, eigenvalue analysis.

Let

$$\boldsymbol{\rho}_1 = \begin{bmatrix} 1&0&0&0&0&0&0&0&0 \\ 0&1&0&0&0&0&0&0&0 \\ 0&0&1&0&0&0&0&0&0 \end{bmatrix}, \qquad \boldsymbol{\sigma}_1 = \begin{bmatrix} 0&0&0&0&0&0&0&0&0 \\ 0&0&0&0&0&0&1&0&0 \\ 0&0&0&0&0&0&0&1&0 \end{bmatrix},$$

$$\boldsymbol{\rho}_2 = \begin{bmatrix} 0&1&0&0&0&0&0&0&0 \\ 0&0&0&1&0&0&0&0&0 \\ 0&0&0&0&1&0&0&0&0 \end{bmatrix}, \qquad \boldsymbol{\sigma}_2 = \begin{bmatrix} 0&0&0&0&0&0&-1&0&0 \\ 0&0&0&0&0&0&0&0&0 \\ 0&0&0&0&0&0&0&0&1 \end{bmatrix},$$

$$\boldsymbol{\rho}_3 = \begin{bmatrix} 0&0&1&0&0&0&0&0&0 \\ 0&0&0&0&1&0&0&0&0 \\ 0&0&0&0&0&1&0&0&0 \end{bmatrix}, \qquad \boldsymbol{\sigma}_3 = \begin{bmatrix} 0&0&0&0&0&0&0&-1&0 \\ 0&0&0&0&0&0&0&0&-1 \\ 0&0&0&0&0&0&0&0&0 \end{bmatrix}.$$

A fundamental property of these matrices is that, for each $\alpha \in \{1, 2, 3\}$,

$$\boldsymbol{\rho}_\alpha \boldsymbol{\theta} = [c_{\alpha 1}, c_{\alpha 2}, c_{\alpha 3}]^T \quad \text{and} \quad \boldsymbol{\sigma}_\alpha \boldsymbol{\theta} = [w_{\alpha 1}, w_{\alpha 2}, w_{\alpha 3}]^T.$$

Using this property, we find that

$$\boldsymbol{m}^T \boldsymbol{C}^2 \boldsymbol{m} = \sum_{\alpha,\beta,\gamma=1}^{3} m_\beta c_{\beta\alpha} c_{\alpha\gamma} m_\gamma = \sum_{\alpha,\beta,\gamma=1}^{3} c_{\beta\alpha} m_\beta m_\gamma c_{\alpha\gamma} = \sum_{\alpha=1}^{3} \boldsymbol{\theta}^T \boldsymbol{\rho}_\alpha{}^T \boldsymbol{m}\boldsymbol{m}^T \boldsymbol{\rho}_\alpha \boldsymbol{\theta}.$$

Likewise

$$\boldsymbol{m}^T\boldsymbol{C}\boldsymbol{W}\dot{\boldsymbol{m}} = \sum_{\alpha=1}^{3}\boldsymbol{\theta}^T{\boldsymbol{\rho}_\alpha}^T\boldsymbol{m}\dot{\boldsymbol{m}}^T\boldsymbol{\sigma}_\alpha\boldsymbol{\theta},$$

$$\dot{\boldsymbol{m}}^T\boldsymbol{W}^2\dot{\boldsymbol{m}} = \sum_{\alpha=1}^{3}\boldsymbol{\theta}^T{\boldsymbol{\sigma}_\alpha}^T\dot{\boldsymbol{m}}\dot{\boldsymbol{m}}^T\boldsymbol{\sigma}_\alpha\boldsymbol{\theta},$$

$$\boldsymbol{m}^T\boldsymbol{W}^2\boldsymbol{m} = \sum_{\alpha=1}^{3}\boldsymbol{\theta}^T{\boldsymbol{\sigma}_\alpha}^T\boldsymbol{m}\boldsymbol{m}^T\boldsymbol{\sigma}_\alpha\boldsymbol{\theta}.$$

Combining the last four identities with

$$\begin{aligned}\|2\boldsymbol{C}\boldsymbol{m}+\boldsymbol{W}\dot{\boldsymbol{m}}\|^2+\|\boldsymbol{W}\boldsymbol{m}\|^2&\\ =4\boldsymbol{m}^T\boldsymbol{C}^2\boldsymbol{m}+4\boldsymbol{m}^T\boldsymbol{C}\boldsymbol{W}\dot{\boldsymbol{m}}-\dot{\boldsymbol{m}}^T\boldsymbol{W}^2\dot{\boldsymbol{m}}-\boldsymbol{m}^T\boldsymbol{W}^2\boldsymbol{m},&\end{aligned}$$

we see that

$$\|2\boldsymbol{C}\boldsymbol{m}+\boldsymbol{W}\dot{\boldsymbol{m}}\|^2+\|\boldsymbol{W}\boldsymbol{m}\|^2=\boldsymbol{\theta}^T\boldsymbol{N}_{\boldsymbol{m},\dot{\boldsymbol{m}}}\boldsymbol{\theta}, \tag{10}$$

where

$$\begin{aligned}\boldsymbol{N}_{\boldsymbol{m},\dot{\boldsymbol{m}}}=4\sum_{\alpha=1}^{3}{\boldsymbol{\rho}_\alpha}^T\boldsymbol{m}\boldsymbol{m}^T\boldsymbol{\rho}_\alpha+4\sum_{\alpha=1}^{3}{\boldsymbol{\rho}_\alpha}^T\boldsymbol{m}\dot{\boldsymbol{m}}^T\boldsymbol{\sigma}_\alpha&\\ -\sum_{\alpha=1}^{3}{\boldsymbol{\sigma}_\alpha}^T\dot{\boldsymbol{m}}\dot{\boldsymbol{m}}^T\boldsymbol{\sigma}_\alpha-\sum_{\alpha=1}^{3}{\boldsymbol{\sigma}_\alpha}^T\boldsymbol{m}\boldsymbol{m}^T\boldsymbol{\sigma}_\alpha.&\end{aligned}$$

Now, in view of (6) and (10),

$$\delta(\boldsymbol{\theta};\boldsymbol{m},\dot{\boldsymbol{m}})=\frac{\boldsymbol{\theta}^T\boldsymbol{M}_{\boldsymbol{m},\dot{\boldsymbol{m}}}\boldsymbol{\theta}}{\boldsymbol{\theta}^T\boldsymbol{N}_{\boldsymbol{m},\dot{\boldsymbol{m}}}\boldsymbol{\theta}}$$

implying that

$$J_2(\boldsymbol{\theta};\mathcal{S})=\sum_{i=1}^{n}\frac{\boldsymbol{\theta}^T\boldsymbol{M}_{\boldsymbol{m}_i,\dot{\boldsymbol{m}}_i}\boldsymbol{\theta}}{\boldsymbol{\theta}^T\boldsymbol{N}_{\boldsymbol{m}_i,\dot{\boldsymbol{m}}_i}\boldsymbol{\theta}}.$$

Hence, immediately,

$$[\nabla_{\boldsymbol{\theta}}J_2(\boldsymbol{\theta};\mathcal{S})]^T=2\boldsymbol{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}, \tag{11}$$

where

$$\boldsymbol{X}_{\boldsymbol{\theta}}=\sum_{i=1}^{n}\frac{\boldsymbol{M}_{\boldsymbol{m}_i,\dot{\boldsymbol{m}}_i}}{\boldsymbol{\theta}^T\boldsymbol{N}_{\boldsymbol{m}_i,\dot{\boldsymbol{m}}_i}\boldsymbol{\theta}}-\sum_{i=1}^{n}\frac{\boldsymbol{\theta}^T\boldsymbol{M}_{\boldsymbol{m}_i,\dot{\boldsymbol{m}}_i}\boldsymbol{\theta}}{(\boldsymbol{\theta}^T\boldsymbol{N}_{\boldsymbol{m}_i,\dot{\boldsymbol{m}}_i}\boldsymbol{\theta})^2}\boldsymbol{N}_{\boldsymbol{m}_i,\dot{\boldsymbol{m}}_i}. \tag{12}$$

Again the minimiser $\widehat{\boldsymbol{\theta}}$ satisfies $[\nabla_{\boldsymbol{\theta}} J_2(\widehat{\boldsymbol{\theta}}; \mathcal{S})]^T = 2\lambda\widehat{\boldsymbol{\theta}}$ for some Lagrange multiplier λ. Combining this with (11), we conclude that $\widehat{\boldsymbol{\theta}}$ is an eigenvector of $\boldsymbol{X}_{\widehat{\boldsymbol{\theta}}}$ with λ as the associated eigenvalue. Now

$$\widehat{\boldsymbol{\theta}}^T \boldsymbol{X}_{\widehat{\boldsymbol{\theta}}} \widehat{\boldsymbol{\theta}} = \lambda \widehat{\boldsymbol{\theta}}^T \widehat{\boldsymbol{\theta}} = \lambda .$$

On the other hand, recourse to (12) reveals that $\widehat{\boldsymbol{\theta}}^T \boldsymbol{X}_{\widehat{\boldsymbol{\theta}}} \widehat{\boldsymbol{\theta}} = 0$. Therefore $\lambda = 0$ and, consequently,

$$\boldsymbol{X}_{\widehat{\boldsymbol{\theta}}} \widehat{\boldsymbol{\theta}} = \boldsymbol{0} . \tag{13}$$

This is a non-linear constraint on $\widehat{\boldsymbol{\theta}}$ which one might hope to resolve by employing a method of successive approximations of some kind. The following scheme for solving (13) suggests itself naturally:

1. Compute $\widehat{\boldsymbol{\theta}}_0$ using least-square fitting based on J_1.
2. Assuming that $\widehat{\boldsymbol{\theta}}_{k-1}$ is known, compute the matrix $\boldsymbol{X}_{\widehat{\boldsymbol{\theta}}_{k-1}}$.
3. Compute a normalised eigenvector of $\boldsymbol{X}_{\widehat{\boldsymbol{\theta}}_{k-1}}$ corresponding to the smallest eigenvalue and take this eigenvector for $\widehat{\boldsymbol{\theta}}_k$.
4. If $\widehat{\boldsymbol{\theta}}_k$ is sufficiently close to $\widehat{\boldsymbol{\theta}}_{k-1}$, then terminate the procedure; otherwise increment k and return to Step 2.

Observe that, on account of (9) and (12),

$$\boldsymbol{X}_{\boldsymbol{\theta}} = \boldsymbol{D}_{\boldsymbol{\theta}} - \boldsymbol{E}_{\boldsymbol{\theta}} ,$$

where

$$\boldsymbol{E}_{\boldsymbol{\theta}} = \sum_{i=1}^{n} \frac{\boldsymbol{\theta}^T \boldsymbol{M}_{\boldsymbol{m}_i, \dot{\boldsymbol{m}}_i} \boldsymbol{\theta}}{(\boldsymbol{\theta}^T \boldsymbol{N}_{\boldsymbol{m}_i, \dot{\boldsymbol{m}}_i} \boldsymbol{\theta})^2} \boldsymbol{N}_{\boldsymbol{m}_i, \dot{\boldsymbol{m}}_i} .$$

Therefore $\boldsymbol{X}_{\boldsymbol{\theta}}$ can be viewed a modification of $\boldsymbol{D}_{\boldsymbol{\theta}}$. Accordingly, the estimator embodied by the above algorithm can be viewed as a modification of the iteratively reweighted least squares estimator from the previous subsection.

Note that other more refined schemes for solving (13) may readily be developed. One such scheme may be obtained by linearising the left-hand side of (13) to incorporate the matrix-valued derivative of the mapping $\boldsymbol{\theta} \mapsto \boldsymbol{X}_{\boldsymbol{\theta}}$.

3.6 Enforcing the cubic constraint

An estimate $\widehat{\boldsymbol{\theta}}$ obtained by any of the methods considered above may fail to satisfy equation (2). A corrective procedure for modifying estimates to accommodate this constraint is therefore needed.

Given a (normalised) estimate $\widehat{\boldsymbol{\theta}} = \boldsymbol{\theta}_{\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}}$, possibly failing to satisfy (2), define $\widehat{\boldsymbol{\theta}}_\rho = \boldsymbol{\theta}_{\widehat{\boldsymbol{C}}_\rho, \widehat{\boldsymbol{W}}_\rho}$ by

$$\widehat{\boldsymbol{C}}_\rho = \frac{\widehat{\boldsymbol{C}} - \boldsymbol{P}\widehat{\boldsymbol{C}}\boldsymbol{P}}{\|\boldsymbol{\theta}_{\widehat{\boldsymbol{C}} - \boldsymbol{P}\widehat{\boldsymbol{C}}\boldsymbol{P}, \widehat{\boldsymbol{W}}}\|} \quad \text{and} \quad \widehat{\boldsymbol{W}}_\rho = \frac{\widehat{\boldsymbol{W}}}{\|\boldsymbol{\theta}_{\widehat{\boldsymbol{C}} - \boldsymbol{P}\widehat{\boldsymbol{C}}\boldsymbol{P}, \widehat{\boldsymbol{W}}}\|},$$

Fig. 1. Image sequence of a calibration grid.

where

$$\boldsymbol{P} = \boldsymbol{I} + \|\widehat{\boldsymbol{w}}\|^{-2}\widehat{\boldsymbol{W}}^2.$$

Note that the vector $\widehat{\boldsymbol{\theta}}_\rho$ comes automatically normalised. It is easily verified that if $\widehat{\boldsymbol{\theta}}$ satisfies (2), then $\boldsymbol{P}\widehat{\boldsymbol{C}}\boldsymbol{P} = \boldsymbol{0}$ and hence $\widehat{\boldsymbol{\theta}} = \widehat{\boldsymbol{\theta}}_\rho$. Since $\boldsymbol{P}\widehat{\boldsymbol{w}} = \widehat{\boldsymbol{w}}$ and $\widehat{\boldsymbol{w}}^T\boldsymbol{P} = \widehat{\boldsymbol{w}}^T$, it follows that $\widehat{\boldsymbol{w}}^T\widehat{\boldsymbol{C}}_\rho\widehat{\boldsymbol{w}} = 0$, which in turn immediately implies that ${\widehat{\boldsymbol{w}}_\rho}^T\widehat{\boldsymbol{C}}_\rho\widehat{\boldsymbol{w}}_\rho = 0$. Thus passing from $\widehat{\boldsymbol{\theta}}$ to $\widehat{\boldsymbol{\theta}}_\rho$ gives a sought-after modification procedure.

3.7 Robust estimation

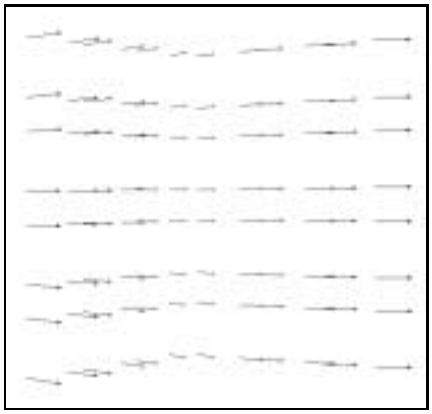

Fig. 2. Optical flow for calibration grid.

Typically, a data set comprises two subsets: a large, dominant subset of valid data or *inliers*, and a relatively small subset of *outliers* or *contaminants*. Least squares minimisation is by nature global and hence vulnerable to distortion by outliers. To obtain robust estimates, outliers have to be detected and rejected. To identify the outliers, we employ Rousseeuw's *least median of squares regression* (LMS), as presented in [21].

The LMS technique is applied to a collection of samples, each being a seven-element subset of $\mathcal{S}$. The size of samples is such that it is minimal to allow an estimate of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ to be determined from a single sample. The collection is assumed to be large enough to comprise a sample containing no outlying data. Ideally, the estimator should consider the set of all seven-element samples. In practice, to make the search computationally feasible, the sample space is reduced to a family of q randomly chosen samples.[1] The number q is determined as follows. Assume that the proportion of outliers in $\mathcal{S}$ does not exceed ϵ,

[1] It is important that such a family be chosen so that points forming samples are evenly spread throughout the whole image.

where $0 < \epsilon < 1$. The probability P that a family of q samples contains at least one sample that is outlier-free is then approximately given by

$$P = 1 - (1 - (1 - \epsilon)^7)^q.$$

Consequently,

$$q = \left[\frac{\log(1 - P)}{\log(1 - (1 - \epsilon)^7)} \right],$$

where $[x]$ denotes the integral part of x. We exploit this formula by assuming that $P = 0.95$ and $\epsilon = 0.5$[2]. If after a first run the LMS estimator finds out that the level of contamination is essentially different from the initial value 0.5, we re-run the LMS procedure taking the estimated contamination level for an updated value of ϵ.

With q fixed, the LMS estimate of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ is obtained by executing the following steps:

1. Using a random selection scheme[3], choose a family $\mathcal{S}_0$ consisting of q subsets of $\mathcal{S}$, each subset containing seven elements.
2. For each $s \in \mathcal{S}_0$, compute three estimates $\widehat{\boldsymbol{\theta}}_{s,k}$ ($k \in \{1, 2, 3\}$) by using the seven-point algorithm from Subsection 3.1 (all three estimates may coincide).
3. For each $(s, k) \in \mathcal{S}_0 \times \{1, 2, 3\}$, determine the median
$$M_{s,k} = \underset{i=1,\ldots,n}{\text{med}} \, \delta(\widehat{\boldsymbol{\theta}}_{s,k}; \boldsymbol{m}_i, \dot{\boldsymbol{m}}_i)^2.$$
4. Letting $(s_q, k_q) \in \mathcal{S}_0 \times \{1, 2, 3\}$ be such that
$$M_{s_q,k_q} = \min_{(s,k) \in \mathcal{S}_0 \times \{1,2,3\}} M_{s,k},$$
take $\widehat{\boldsymbol{\theta}}_{s_q,k_q}$ for the LMS estimate of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$.

With the LMS estimate at hand, outliers within $\mathcal{S}$ are identified by using the following procedure:

1. Take $\hat{\sigma} = 1.4826 \left(1 + \frac{5}{n-7}\right) \sqrt{M_{s_q,k_q}}$ for the *robust standard deviation* of the measurements.
2. Declare $[\boldsymbol{m}_i^T, \dot{\boldsymbol{m}}_i^T]^T$ to be an outlier if and only if $\delta(\widehat{\boldsymbol{\theta}}_{s_q,k_q}; \boldsymbol{m}_i, \dot{\boldsymbol{m}}_i) > 2.5\hat{\sigma}$.

Finally, an overall algorithm for generating a robust estimate of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$ can be formulated as follows:

1. Extract outliers from $\mathcal{S}$ by using the above procedure.
2. Compute an estimate $\widehat{\boldsymbol{\theta}}$ by applying to the remaining elements of $\mathcal{S}$ (a refinement of) the iterative algorithm from Subsection 3.5 to solve (13).
3. Modify $\widehat{\boldsymbol{\theta}}$ to $\widehat{\boldsymbol{\theta}}_\rho$ by using the procedure from Subsection 3.6.

[2] 0.5 is roughly the *breakdown point* of the LMS estimator. The breakdown point of an estimator is the smallest proportion of the data that can have an arbitrarily large effect on its value.

[3] Such a scheme should be sufficiently regular to avoid clustered outcomes.

Fig. 3. Grid reconstruction from various views.

4 Experimental results

In order to assess the applicability of the approach, two tests with real image sequences were performed. The three images shown in Figure 1 were captured by a Phillips CCD camera with a 12.5-mm lens. Corners were localised to sub-pixel accuracy with the use of a corner detector, correspondences between the images were obtained, and the optical flow depicted in Figure 2 was computed by exploiting these correspondences (thus no use was made of intensity-based methods of flow computation). The LMS procedure was used to eliminate outliers, and a general solver was then employed to resolve (13), which resulted in an estimate of $\boldsymbol{\pi}(\boldsymbol{C}, \boldsymbol{W})$. Finally, closed-form expressions were employed to self-calibrate the system. With self-calibration completed, the pleasing reconstruction displayed in Figure 3 was finally obtained.

The images in Figure 4 were captured using a Nikon-N90/Kodak-DCS420 camera. The process outlined above produced the reconstruction depicted in Figure 5. To convey clearly the patterns of the calibration grid, reconstructed points in 3-space have been connected by line segments. Note that coplanarity of three dimensional points on the wall and table is well preserved, as is the orthogonality of groups of points on the different faces of the pillar. This simple reconstruction is again visually pleasing and suggests that the approach holds promise.

Acknowledgments

This research was in part funded by the Australian Research Council and the Cooperative Research Center for Sensor Signal and Information Processing.

References

1. M. Armstrong, A. Zisserman, and R. Hartley, *Self-calibration from image triplets*, In Buxton and Cipolla [8], pp. 3–16.
2. K. Åström and A. Heyden, *Continuous time matching constraints for image streams*, International Journal of Computer Vision, to appear.

Fig. 4. Extract from an office image sequence.

3. ———, *Multilinear forms in the infinitesimal-time case*, Proceedings, CVPR '96, IEEE Computer Society Conference on Computer Vision and Pattern Recognition (San Francisco, CA, June 18–20, 1996), IEEE Computer Society Press, Los Alamitos, CA, 1996, pp. 833–838.
4. J. L. Barron and R. Eagleson, *Recursive estimation of time-varying motion and structure parameters*, Pattern Recognition **29** (1996), no. 5, 797–818.
5. P. A. Beardsley, A. Zisserman, and D. W. Murray, *Sequential updating of projective and affine structure from motion*, International Journal of Computer Vision **23** (1997), no. 3, 235–259.
6. F. Bookstein, *Fitting conic sections to scattered data*, Computer Vision, Graphics, and Image Processing **9** (1979), no. 1, 56–71.
7. M. J. Brooks, W. Chojnacki, and L. Baumela, *Determining the egomotion of an uncalibrated camera from instantaneous optical flow*, Journal of the Optical Society of America A **14** (1997), no. 10, 2670–2677.
8. B. Buxton and R. Cipolla (eds.), *Computer Vision—ECCV '96*, Lecture Notes in Computer Science, vol. 1064, Fourth European Conference on Computer Vision, Cambridge, UK, April 14–18, 1996, Springer, Berlin, 1996.
9. O. D. Faugeras, Q. T. Luong, and S. J. Maybank, *Camera self-calibration: Theory and experiments*, Computer Vision—ECCV '92 (Second European Conference on Computer Vision, Santa Margherita Ligure, Italy, May 19–22, 1992) (G. Sandini, ed.), Springer, Berlin, 1992, pp. 321–334.
10. N. C. Gupta and L. N. Kanal, *3-D motion estimation from motion field*, Artificial Intelligence **78** (1995), no. 1-2, 45–86.
11. D. J. Heeger and A. D. Jepson, *Subspace methods for recovering rigid motion I: algorithm and implementation*, International Journal of Computer Vision **7** (1992), no. 2, 95–117.
12. K. Kanatani, *3-D interpretation of optical flow by renormalization*, International Journal of Computer Vision **11** (1993), no. 3, 267–282.
13. ———, *Statistical Optimization for Geometric Computation: Theory and Practice*, Elsevier, Amsterdam, 1996.
14. Q.-T. Luong and O. D. Faugeras, *The fundamental matrix: theory, algorithms, and stability analysis*, International Journal of Computer Vision **17** (1996), no. 1, 43–75.
15. ———, *Self-calibration of a moving camera from point correspondences and fundamental matrices*, International Journal of Computer Vision **22** (1997), no. 3, 261–289.
16. S. J. Maybank, *The angular velocity associated with the optical flowfield arising from motion through a rigid environment*, Proceedings of the Royal Society of London Ser. A **401** (1985), 317–326.

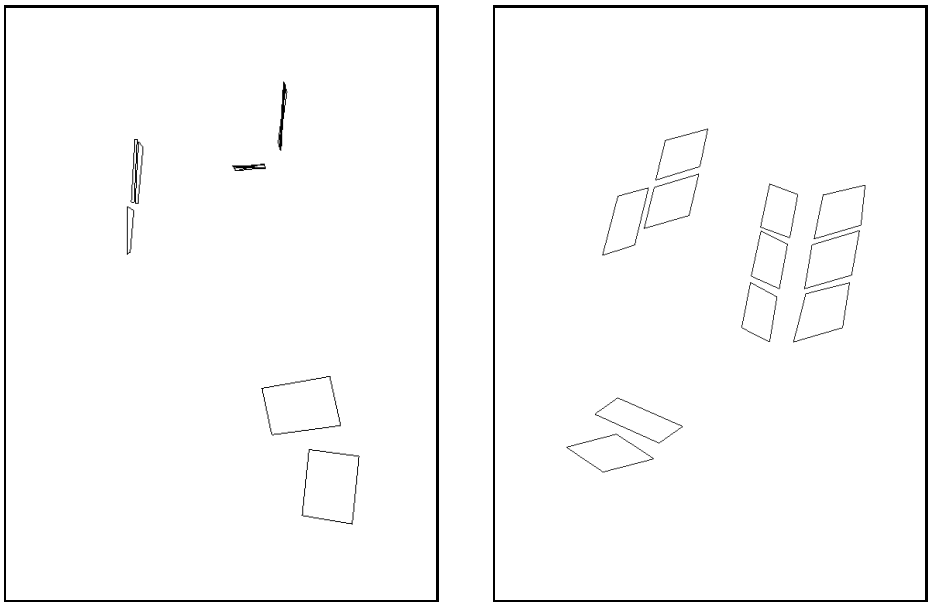

Fig. 5. Office reconstruction viewed from overhead and from right-side.

17. S. J. Maybank and O. D. Faugeras, *A theory of self-calibration of a moving camera*, International Journal of Computer Vision **8** (1992), no. 2, 123–151.
18. N. Ohta and K. Kanatani, *Optimal structure-from-motion algorithm for optical flow*, IEICE Transactions on Information and Systems **E78-D** (1995), no. 12, 1559–1566.
19. M. Pollefeys, M. Van Gool, and M. Proesmans, *Euclidean 3D reconstruction from image sequences with variable focal length*, In Buxton and Cipolla [8], pp. 31–42.
20. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, Cambridge University Press, Cambridge, 1995.
21. P. J. Rousseeuw and A. M. Leroy, *Robust Regression and Outlier Regression*, John Wiley & Sons, New York, 1987.
22. P. D. Sampson, *Fitting conics sections to 'very scattered' data: An iterative refinement of the Bookstein algorithm*, Computer Graphics and Image Processing **18** (1982), no. 1, 97–108.
23. S. Soatto and P. Perona, *Recursive 3D visual motion estimation using subspace constraints*, International Journal of Computer Vision **22** (1997), no. 3, 235–259.
24. P. H. S. Torr and D. W. Murray, *The development and comparison of robust methods for estimating the fundamental matrix*, International Journal of Computer Vision **24** (1997), no. 3, 271–300.
25. T. Viéville and O. D. Faugeras, *Motion analysis with a camera with unknown, and possibly varying intrinsic parameters*, Proceedings of the Fifth International Conference on Computer Vision (Cambridge, MA, June 1995), IEEE Computer Society Press, Los Alamitos, CA, 1995, pp. 750–756.
26. T. Viéville, O. D. Faugeras, and Q.-T. Luong, *Motion of points and lines in the uncalibrated case*, International Journal of Computer Vision **17** (1996), no. 1, 7–41.
27. T. Viéville, C. Zeller, and L. Robert, *Using collineations to compute motion and structure in an uncalibrated image sequence*, International Journal of Computer Vision **20** (1996), no. 3, 213–242.

28. J. Weng, T. S. Huang, and N. Ahuja, *Motion and structure from two perspective views: algorithms, error analysis, and error estimation*, IEEE Transactions on Pattern Analysis and Machine Intelligence **11** (1989), no. 5, 451–476.
29. Z. Zhang, *Parameter estimation techniques: a tutorial with application to conic fitting*, Image and Vision Computing **15** (1997), no. 1, 57–76.