

Minimal Cases for Computing the Generalized Relative Pose using Affine Correspondences

Banglei Guan¹, Ji Zhao*, Daniel Barath² and Friedrich Fraundorfer^{3,4}

¹College of Aerospace Science and Engineering, National University of Defense Technology, China

²Computer Vision and Geometry Group, Department of Computer Science, ETH Zürich

³Institute for Computer Graphics and Vision, Graz University of Technology, Austria

⁴Remote Sensing Technology Institute, German Aerospace Center, Germany

guanbanglei12@nudt.edu.cn zhaoji84@gmail.com dbarath@ethz.ch fraundorfer@icg.tugraz.at

Abstract

We propose three novel solvers for estimating the relative pose of a multi-camera system from affine correspondences (ACs). A new constraint is derived interpreting the relationship of ACs and the generalized camera model. Using the constraint, we demonstrate efficient solvers for two types of motions assumed. Considering that the cameras undergo planar motion, we propose a minimal solution using a single AC and a solver with two ACs to overcome the degenerate case. Also, we propose a minimal solution using two ACs with known vertical direction, e.g., from an IMU. Since the proposed methods require significantly fewer correspondences than state-of-the-art algorithms, they can be efficiently used within RANSAC for outlier removal and initial motion estimation. The solvers are tested both on synthetic data and on real-world scenes from the KITTI odometry benchmark. It is shown that the accuracy of the estimated poses is superior to the state-of-the-art techniques.

1. Introduction

Relative pose estimation from two views of a camera, or a multi-camera system is regarded as a fundamental problem in computer vision [22, 10, 45, 46, 51], which plays an important role in simultaneous localization and mapping (SLAM) and structure-from-motion (SfM). Thus, improving the accuracy, efficiency and robustness of relative pose estimation algorithms is always an important research topic [30, 50, 1, 16, 6, 12, 31]. Motivated by the fact that multi-camera systems are available in self-driving cars, micro aerial vehicles or AR headsets, this paper investigates the problem of estimating the relative pose of multi-camera systems from affine correspondences (ACs), see Fig. 1.

Since a multi-camera system contains multiple individ-




Figure 1. An affine correspondence in camera C_i between consecutive frames k and $k+1$. The local affine transformation A relates the infinitesimal patches around point correspondence $(\mathbf{x}_{ij}, \mathbf{x}'_{ij})$.

ual cameras connected by being fixed to a single rigid body, it has the advantage of large field-of-view and high accuracy [47, 14]. The main difference of a multi-camera system and a standard pinhole camera is the absence of a single projection center [41]. Due to the different camera model, the relative pose estimation problem of multi-camera systems [24] is different from the monocular cameras [39, 18], which results in different equations. In order to remove outlier matches, most of the state-of-the-art SLAM and SfM pipelines using a multi-camera system [21, 23] apply the relative pose estimation algorithms repeatedly in a robust estimation framework, e.g. the Random Sample Consensus (RANSAC) [13]. This outlier removal process has to be efficient, which directly affects the real-time performance of SLAM and SfM. The computational complexity and, thus, the processing time of the RANSAC procedure depends exponentially on the number of points required for the relative pose estimation of multi-camera system.

Therefore, exploring the minimal solutions for relative pose estimation of multi-camera system is of significant importance and has received sustained attention [24, 32, 10, 25, 33, 50, 28]. The idea of deriving minimal solu-

*Corresponding author.

tions for relative pose estimation of multi-camera systems ranges back to the work of Stewénius *et al.* with the 6-point method [24]. Then other classical works have been subsequently proposed, such as the 17-point linear method [32] and techniques based on iterative optimization [27]. The minimal number of necessary points can be further reduced by taking additional motion constraints into account [29] or exploiting the measurements from other sensors, like an inertial measurement unit (IMU) [30, 48, 49, 34, 36]. Typically, the assumption of planar motion or considering known vertical direction are common for self-driving cars and ground robots [9, 20, 18, 44, 31], which makes the outlier removal more efficient and numerically more stable.

All previously mentioned relative pose solvers estimate the pose parameters from a set of point correspondences (PCs), *e.g.*, coming from SIFT [35] or SURF [7] detectors. Due to containing more information about the underlying surface geometry than PCs, ACs enable to estimate the pose from fewer correspondences. In this paper, we focus on the relative pose estimation of a multi-camera system from ACs, instead of PCs. The contributions of this paper are:

- A new constraint that interprets the relationship of ACs and the generalized camera model is derived under general motion. This constraint can be easily generalized to special cases of multi-camera motion, *e.g.*, planar motion and known vertical direction.
- When the motion is planar (*i.e.*, the body to which the cameras are fixed moves on a plane; 3DOF), a single AC is sufficient to recover the planar motion of a multi-camera system. In order to deal with the degenerate case of the 1AC solver, we also propose a new method to estimate the relative pose from two ACs. The point-based solver [29] requires at least two PCs and requires the Ackermann motion model to hold.
- A third solver is proposed for the case when the vertical direction is known (4DOF), *e.g.*, from an IMU attached to the multi-camera system. We show that two ACs are enough to recover the relative pose. In contrast, the point-based solver requires four PCs [30, 48].

2. Related Work

Due to the absence of a single center of projection, the camera model of multi-camera systems is different from the standard pinhole camera. Pless proposed to express the light rays as Plücker lines and derived the generalized camera model which has become a standard representation for the multi-camera systems [41]. Stewénius *et al.* proposed the first minimal solution to estimate the relative pose of a multi-camera system from 6 PCs, which produces up to 64 solutions [24]. Li *et al.* provided several linear solvers to compute the relative pose, among which the most commonly used one requires 17 PCs [32]. Kneip *et al.* proposed

an iterative approach for the relative pose estimation based on eigenvalue minimization [27]. Ventura *et al.* used first-order approximation of the rotation to simplify the problem and estimated the relative pose from 6 PCs [50].

By considering additional motion constraints or using additional information provided by an IMU, the number of required PCs can be further reduced. Lee *et al.* presented a minimal solution with two PCs for the ego-motion estimation of a multi-camera system, which constrains the relative motion by the Ackermann motion model [29]. In addition, a variety of algorithms have been proposed when a common direction of the multi-camera system is known, *i.e.* an IMU provides the roll and pitch angles of the multi-camera system. The relative pose estimation with known vertical direction requires a minimum of 4 PCs [30, 48, 34].

Exploiting the additional affine parameters besides the image coordinates has been recently proposed for the relative pose estimation of monocular cameras, which reduces the number of required points significantly. Bentolila *et al.* estimated the fundamental matrix from three ACs [8]. Raposo *et al.* computed homography and essential matrix using two ACs [43]. Barath *et al.* derived the constraints between the local affine transformation and the essential matrix and recovered the essential matrix from two ACs [4]. Hajder *et al.* [20] and Guan *et al.* [18, 19] proposed several minimal solutions for relative pose from a single AC under the planar motion assumption or with knowledge of a vertical direction. The above mentioned works are only suitable for the monocular perspective camera. For multi-camera systems, Alyousefi *et al.* recently proposed a linear solver to estimate the relative pose using 6 ACs [2]. Guan *et al.* estimated the relative pose from 2 ACs by utilizing a first-order rotation approximation [17]. In this paper, we focus on the minimal number of ACs to estimate the relative pose of a multi-camera system. Table 1 shows a summary of the solvers, including the DOF of the motion, feature types and number of points required.

Table 1. Relative pose solvers for multi-camera systems.

Solver	[32]	[27]	[24]	[2]	[30]	[48]	[34]	1AC plane	2AC plane	2AC vertical
Motion	6DOF				4DOF			3DOF		4DOF
Feature	PCs		ACs		PCs			ACs		
Point #	17	8	6		4			1	2	

3. Geometric Constraints from ACs

A multi-camera system is made up of individual cameras denoted by C_i , as shown in Fig. 1. Take an AC seen by same camera for an example. The geometric constraints can be easily generalized to the case that the AC is seen by different cameras. The extrinsic parameters of camera C_i expressed in a multi-camera reference frame are represented as $(\mathbf{R}_i, \mathbf{t}_i)$. For general motion, there is a 3DOF rotation and a 3DOF translation between two reference frames at time k and $k+1$. Rotation \mathbf{R} using Cayley parameterization

and translation \mathbf{t} can be written as:

$$\mathbf{R} = \frac{1}{1 + q_x^2 + q_y^2 + q_z^2} \cdot \begin{bmatrix} 1 + q_x^2 - q_y^2 - q_z^2 & 2q_xq_y - 2q_z & 2q_y + 2q_xq_z \\ 2q_xq_y + 2q_z & 1 - q_x^2 + q_y^2 - q_z^2 & 2q_yq_z - 2q_x \\ 2q_xq_z - 2q_y & 2q_x + 2q_yq_z & 1 - q_x^2 - q_y^2 + q_z^2 \end{bmatrix}, \quad (1)$$

$$\mathbf{t} = [t_x \quad t_y \quad t_z]^T, \quad (2)$$

where $[1, q_x, q_y, q_z]^T$ is a homogeneous quaternion vector. Note that 180 degree rotations are prohibited in Cayley parameterization, but this is a rare case for consecutive frames.

3.1. Generalized Camera Model

We give a brief description of generalized camera model (GCM) [41]. Let us denote an AC in camera C_i between consecutive frames k and $k + 1$ as $(\mathbf{x}_{ij}, \mathbf{x}'_{ij}, \mathbf{A})$, where \mathbf{x}_{ij} and \mathbf{x}'_{ij} are the normalized homogeneous image coordinates of feature point j and \mathbf{A} is a 2×2 local affine transformation. Indices i and j are the camera and point index, respectively. The related local affine transformation \mathbf{A} is a 2×2 linear transformation which relates the infinitesimal patches around \mathbf{x}_{ij} and \mathbf{x}'_{ij} [3]. The normalized homogeneous image coordinates $(\mathbf{p}_{ij}, \mathbf{p}'_{ij})$ expressed in the multi-camera reference frame are given as

$$\mathbf{p}_{ij} = \mathbf{R}_i \mathbf{x}_{ij}, \quad \mathbf{p}'_{ij} = \mathbf{R}_i \mathbf{x}'_{ij}. \quad (3)$$

The unit direction of rays $(\mathbf{u}_{ij}, \mathbf{u}'_{ij})$ expressed in the multi-camera reference frame are given as: $\mathbf{u}_{ij} = \mathbf{p}_{ij} / \|\mathbf{p}_{ij}\|$, $\mathbf{u}'_{ij} = \mathbf{p}'_{ij} / \|\mathbf{p}'_{ij}\|$. The 6-dimensional vector Plücker lines corresponding to the rays are denoted as $\mathbf{l}_{ij} = [\mathbf{u}_{ij}^T, (\mathbf{t}_i \times \mathbf{u}_{ij})^T]^T$, $\mathbf{l}'_{ij} = [\mathbf{u}'_{ij}^T, (\mathbf{t}_i \times \mathbf{u}'_{ij})^T]^T$. The generalized epipolar constraint is written as [41]

$$\mathbf{l}_{ij}^T \begin{bmatrix} [\mathbf{t}]_{\times} \mathbf{R}_i & \mathbf{R}_i \\ \mathbf{R}_i & \mathbf{0} \end{bmatrix} \mathbf{l}_{ij} = 0, \quad (4)$$

where \mathbf{l}_{ij}^T and \mathbf{l}_{ij} are Plücker lines between two consecutive frames at time k and $k + 1$.

3.2. Affine Transformation Constraint

We denote the transition matrix of camera coordinate system C_i between consecutive frames k and $k + 1$ as $(\mathbf{R}_{Ci}, \mathbf{t}_{Ci})$, which is represented as:

$$\begin{bmatrix} \mathbf{R}_{Ci} & \mathbf{t}_{Ci} \\ \mathbf{0} & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R}_i & \mathbf{t}_i \\ \mathbf{0} & 1 \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R}_i & \mathbf{t}_i \\ \mathbf{0} & 1 \end{bmatrix} \quad (5)$$

$$= \begin{bmatrix} \mathbf{R}_i^T \mathbf{R} \mathbf{R}_i & \mathbf{R}_i^T \mathbf{R} \mathbf{t}_i + \mathbf{R}_i^T \mathbf{t} - \mathbf{R}_i^T \mathbf{t}_i \\ \mathbf{0} & 1 \end{bmatrix}.$$

Essential matrix \mathbf{E} of two frames of camera C_i is given as:

$$\mathbf{E} = [\mathbf{t}_{Ci}]_{\times} \mathbf{R}_{Ci} = \mathbf{R}_i^T [\mathbf{R}_i \mathbf{t}_{Ci}]_{\times} \mathbf{R} \mathbf{R}_i, \quad (6)$$

where $[\mathbf{R}_i \mathbf{t}_{Ci}]_{\times} = \mathbf{R} [\mathbf{t}_i]_{\times} \mathbf{R}^T + [\mathbf{t}]_{\times} - [\mathbf{t}_i]_{\times}$. The relationship of essential matrix \mathbf{E} and local affine transformation \mathbf{A} is formulated as follows [4]:

$$(\mathbf{E}^T \mathbf{x}'_{ij})_{(1:2)} = -(\hat{\mathbf{A}}^T \mathbf{E} \mathbf{x}_{ij})_{(1:2)}, \quad (7)$$

where $\mathbf{n}_{ij} \triangleq \mathbf{E}^T \mathbf{x}'_{ij}$ and $\mathbf{n}'_{ij} \triangleq \mathbf{E} \mathbf{x}_{ij}$ denote the epipolar lines in their implicit form in frames of camera C_i at times k and $k + 1$. The subscript 1 and 2 represent the first and second equations of the equation system, respectively. $\hat{\mathbf{A}}$ is a 3×3 matrix: $\hat{\mathbf{A}} = [\mathbf{A} \quad \mathbf{0}; \mathbf{0} \quad 0]$. By substituting Eq. (6) into Eq. (7), we obtain:

$$(\mathbf{R}_i^T \mathbf{R}^T [\mathbf{R}_i \mathbf{t}_{Ci}]_{\times}^T \mathbf{R}_i \mathbf{x}'_{ij})_{(1:2)} = -(\hat{\mathbf{A}}^T \mathbf{R}_i^T [\mathbf{R}_i \mathbf{t}_{Ci}]_{\times} \mathbf{R} \mathbf{R}_i \mathbf{x}_{ij})_{(1:2)}. \quad (8)$$

Based on Eq. (3), the above equation is reformulated and expanded as follows:

$$(\mathbf{R}_i^T ([\mathbf{t}_i]_{\times} \mathbf{R}^T + \mathbf{R}^T [\mathbf{t}]_{\times} - \mathbf{R}^T [\mathbf{t}_i]_{\times}) \mathbf{p}'_{ij})_{(1:2)} = (\hat{\mathbf{A}}^T \mathbf{R}_i^T (\mathbf{R} [\mathbf{t}_i]_{\times} + [\mathbf{t}]_{\times} \mathbf{R} - [\mathbf{t}_i]_{\times} \mathbf{R}) \mathbf{p}_{ij})_{(1:2)}. \quad (9)$$

Eq. (9) interprets the new epipolar constraints which a local affine transformation implies on the i -th camera from a multi-camera system between two frames k and $k + 1$.

For each AC $(\mathbf{x}_{ij}, \mathbf{x}'_{ij}, \mathbf{A})$, we get three polynomials based on Eqs. (4) and (9), see supplementary material for details. Motivated by scenarios like self-driving cars, ground robots or AR headsets, we investigate relevant special cases of multi-camera motion, *i.e.*, planar motion and motion with known vertical direction, see Fig. 2. We show that two special cases can be efficiently solved with ACs.

4. Relative Pose under Planar Motion




Figure 2. Special cases of multi-camera motion: (a) Planar motion in top-view. There are three unknowns: yaw angle θ , translation direction ϕ and translation distance ρ . (b) Motion with known vertical direction. There are four unknowns: a Y-axis rotation \mathbf{R}_y and 3D translation $\hat{\mathbf{t}} = [\hat{t}_x, \hat{t}_y, \hat{t}_z]^T$.

When assuming that the body, to which the camera system is rigidly fixed, moves on a planar surface (as visualized in Fig. 2(a)), there are only a Y-axis rotation and 2D translation between the reference frames k and $k + 1$. Similar to Eqs. (1) and (2), the rotation $\mathbf{R} = \mathbf{R}_y$ and the translation \mathbf{t}

from frame k to $k + 1$ is written as:

$$\mathbf{R}_y = \frac{1}{1 + q_y^2} \begin{bmatrix} 1 - q_y^2 & 0 & -2q_y \\ 0 & 1 + q_y^2 & 0 \\ 2q_y & 0 & 1 - q_y^2 \end{bmatrix}, \quad (10)$$

$$\mathbf{t} = [t_x \quad 0 \quad t_z]^T.$$

where $q_y = \tan(\frac{\theta}{2})$, $t_x = \rho \sin(\phi)$, $t_z = -\rho \cos(\phi)$, ρ is the distance between two multi-camera reference frames.

4.1. Solution by Reduction to a Single Polynomial

By substituting Eq. (10) into Eqs. (4) and (9), we get an equation system of three polynomials for three unknowns q_y , t_x and t_z . Since an AC generally provides three independent constraints for relative pose, a single AC is sufficient to recover the planar motion of a multi-camera system. After separating q_y from t_x , t_z , the three independent constraints from an AC form matrix equation:

$$\frac{1}{1 + q_y^2} \underbrace{\begin{bmatrix} M_{11} & M_{12} & M_{13} \\ M_{21} & M_{22} & M_{23} \\ M_{31} & M_{32} & M_{33} \end{bmatrix}}_{\mathbf{M}(q_y)} \begin{bmatrix} t_x \\ t_z \\ 1 \end{bmatrix} = \mathbf{0}, \quad (11)$$

where M_{ij} ($i, j \in [1, 3]$) is an element of coefficient matrix $\mathbf{M}(q_y)$ and are formed by the polynomial coefficients and one unknown variable q_y , see supplementary material for details. Since $\mathbf{M}(q_y)$ is a square matrix, Eq. (11) has a non-trivial solution only if the determinant of $\mathbf{M}(q_y)/(1 + q_y^2)$ is zero. The expansion of $\det(\mathbf{M}(q_y)/(1 + q_y^2)) = 0$ gives a 4-degree univariate polynomial as follows:

$$\text{quot}(\sum_{i=0}^6 w_i q_y^i, q_y^2 + 1) = 0, \quad (12)$$

where $\text{quot}(a, b)$ means calculating the quotient of a divided by b , w_0, \dots, w_6 are formed by a Plücker line correspondence and an affine transformation between the corresponding feature points.

Note that the coefficients are divided by $q_y^2 + 1$, which reduces the polynomial degree and improves the efficiency of the solution. The univariate polynomial Eq. (12) leads to an explicit analytic solution with a maximum of 4 real roots. Once the solutions for q_y are found, the remaining unknowns t_x and t_z are solved by substituting q_y into $\mathbf{M}(q_y)$ and solving the linear system via calculating its null vector. Finally, the rotation matrix \mathbf{R}_y is recovered from Eq. (10).

However, we proved that the solver using a single AC has a degenerate case, *i.e.*, the distances between the motion plane and optical centers of the cameras being equal, see supplementary material for details. This degenerate case might happen in the self-driving scenario, which would lead to that both the translation direction and the translation scale cannot be calculated using one AC. To overcome this issue,

two ACs are used to estimate the relative pose. For example, the first and second constraints of the first AC, and the first constraint of the second AC are also stacked into three equations in three unknowns, just as Eq. (11). The solution procedure remains the same, except that the code for constructing the coefficient matrix $\mathbf{M}(q_y)$ is replaced.

5. Relative Pose with Known Vertical Direction

In this section a minimal solution using two ACs is proposed for relative motion estimation for multi-camera systems with known vertical direction, see Fig. 2(b). In this case, an IMU is coupled with the multi-camera system and the relative rotation between the IMU and the reference frame is known. The IMU provides the known roll and pitch angles for the reference frame. So the reference frame can be aligned with the measured vertical direction, such that the X-Z-plane of the aligned reference frame is parallel to the ground plane and the Y-axis is parallel to the vertical direction. Rotation \mathbf{R}_{imu} for aligning the reference frame to the aligned reference frame is written as:

$$\mathbf{R}_{\text{imu}} = \mathbf{R}_p \mathbf{R}_r$$

$$= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta_p) & \sin(\theta_p) \\ 0 & -\sin(\theta_p) & \cos(\theta_p) \end{bmatrix} \begin{bmatrix} \cos(\theta_r) & \sin(\theta_r) & 0 \\ -\sin(\theta_r) & \cos(\theta_r) & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

where θ_r and θ_p are roll and pitch angles provided by the coupled IMU, respectively. Thus, there are only a Y-axis rotation $\mathbf{R} = \mathbf{R}_y$ and 3D translation $\tilde{\mathbf{t}} = \mathbf{R}'_{\text{imu}} \mathbf{t} = [\tilde{t}_x, \tilde{t}_y, \tilde{t}_z]^T$ to be estimated between the aligned multi-camera reference frames at time k and $k + 1$. In this section, we show that the geometric constraints in Section 3 can be generalized to the multi-camera motion with known vertical direction.

5.1. Generalized Camera Model

Let us denote the rotation matrices from the roll and pitch angles of the two corresponding multi-camera reference frames at time k and $k + 1$ as \mathbf{R}_{imu} and \mathbf{R}'_{imu} . The relative rotation between two reference frames is

$$\mathbf{R} = (\mathbf{R}'_{\text{imu}})^T \mathbf{R}_{\text{imu}}. \quad (13)$$

We substitute Eq. (13) into Eq. (4) yields:

$$\underbrace{\left(\begin{bmatrix} \mathbf{R}'_{\text{imu}} & \mathbf{0} \\ \mathbf{0} & \mathbf{R}'_{\text{imu}} \end{bmatrix} \mathbf{l}_{ij}' \right)^T}_{\tilde{\mathbf{l}}_{ij}'} \begin{bmatrix} [\tilde{\mathbf{t}}]_{\times} \mathbf{R}_y & \mathbf{R}_y \\ \mathbf{R}_y & \mathbf{0} \end{bmatrix} \quad (14)$$

$$\underbrace{\left(\begin{bmatrix} \mathbf{R}_{\text{imu}} & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{\text{imu}} \end{bmatrix} \mathbf{l}_{ij} \right)}_{\tilde{\mathbf{l}}_{ij}} = 0,$$

where $\tilde{\mathbf{l}}_{ij} \leftrightarrow \tilde{\mathbf{l}}_{ij}'$ are the corresponding Plücker lines expressed in the aligned multi-camera reference frame.

5.2. Affine Transformation Constraint

In this case, the transition matrix of the camera coordinate system C_i between consecutive frames k and $k + 1$ is represented as follows:

$$\begin{bmatrix} \mathbf{R}_{C_i} & \mathbf{t}_{C_i} \\ \mathbf{0} & 1 \end{bmatrix} = \left(\begin{bmatrix} \mathbf{R}'_{imu} & \mathbf{0} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R}_i & \mathbf{t}_i \\ \mathbf{0} & 1 \end{bmatrix} \right)^{-1} \begin{bmatrix} \mathbf{R}_y & \tilde{\mathbf{t}} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R}'_{imu} & \mathbf{0} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R}_i & \mathbf{t}_i \\ \mathbf{0} & 1 \end{bmatrix}. \quad (15)$$

we denote that

$$\begin{bmatrix} \tilde{\mathbf{R}}_{imu} & \tilde{\mathbf{t}}_{imu} \\ \mathbf{0} & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R}_{imu} & \mathbf{0} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R}_i & \mathbf{t}_i \\ \mathbf{0} & 1 \end{bmatrix}, \quad (16)$$

$$\begin{bmatrix} \tilde{\mathbf{R}}'_{imu} & \tilde{\mathbf{t}}'_{imu} \\ \mathbf{0} & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R}'_{imu} & \mathbf{0} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R}_i & \mathbf{t}_i \\ \mathbf{0} & 1 \end{bmatrix}.$$

By substituting Eq. (16) into Eq. (15), we obtain

$$\begin{bmatrix} \mathbf{R}_{C_i} & \mathbf{t}_{C_i} \\ \mathbf{0} & 1 \end{bmatrix} = \begin{bmatrix} (\tilde{\mathbf{R}}'_{imu})^T \mathbf{R}_y \tilde{\mathbf{R}}_{imu} & (\tilde{\mathbf{R}}'_{imu})^T (\mathbf{R}_y \tilde{\mathbf{t}}_{imu} + \tilde{\mathbf{t}} - \tilde{\mathbf{t}}'_{imu}) \\ \mathbf{0} & 1 \end{bmatrix}. \quad (17)$$

Essential matrix \mathbf{E} between the two frames is given as

$$\mathbf{E} = [\mathbf{t}_{C_i}]_{\times} \mathbf{R}_{C_i} = (\tilde{\mathbf{R}}'_{imu})^T [\tilde{\mathbf{R}}_{imu} \mathbf{t}_{C_i}]_{\times} \mathbf{R}_y \tilde{\mathbf{R}}_{imu}, \quad (18)$$

where $[\tilde{\mathbf{R}}_{imu} \mathbf{t}_{C_i}]_{\times} = \mathbf{R}_y [\tilde{\mathbf{t}}_{imu}]_{\times} \mathbf{R}_y^T + [\tilde{\mathbf{t}}]_{\times} - [\tilde{\mathbf{t}}'_{imu}]_{\times}$. By substituting Eq. (18) into Eq. (7), we obtain

$$\begin{aligned} & (\tilde{\mathbf{R}}_{imu}^T \mathbf{R}_y^T [\tilde{\mathbf{R}}'_{imu} \mathbf{t}_{C_i}]_{\times}^T \tilde{\mathbf{R}}_{imu} \mathbf{x}_{ij})_{(1:2)} = \\ & - (\hat{\mathbf{A}}^T (\tilde{\mathbf{R}}'_{imu})^T [\tilde{\mathbf{R}}_{imu} \mathbf{t}_{C_i}]_{\times} \mathbf{R}_y \tilde{\mathbf{R}}_{imu} \mathbf{x}_{ij})_{(1:2)}. \end{aligned} \quad (19)$$

We denote the normalized homogeneous image coordinates expressed in the aligned multi-camera reference frame as $(\tilde{\mathbf{p}}_{ij}, \tilde{\mathbf{p}}'_{ij})$, which are given as

$$\tilde{\mathbf{p}}_{ij} = \tilde{\mathbf{R}}_{imu} \mathbf{x}_{ij}, \quad \tilde{\mathbf{p}}'_{ij} = \tilde{\mathbf{R}}'_{imu} \mathbf{x}'_{ij}. \quad (20)$$

Based on the above equation, Eq. (19) is rewritten as

$$\begin{aligned} & (\tilde{\mathbf{R}}_{imu}^T ([\tilde{\mathbf{t}}_{imu}]_{\times} \mathbf{R}_y^T + \mathbf{R}_y^T [\tilde{\mathbf{t}}]_{\times} - \mathbf{R}_y^T [\tilde{\mathbf{t}}'_{imu}]_{\times}) \tilde{\mathbf{p}}'_{ij})_{(1:2)} = \\ & (\hat{\mathbf{A}}^T (\tilde{\mathbf{R}}'_{imu})^T (\mathbf{R}_y [\tilde{\mathbf{t}}_{imu}]_{\times} + [\tilde{\mathbf{t}}]_{\times} \mathbf{R}_y - [\tilde{\mathbf{t}}'_{imu}]_{\times} \mathbf{R}_y) \tilde{\mathbf{p}}_{ij})_{(1:2)} \end{aligned} \quad (21)$$

5.3. Solution by Reduction to a Single Polynomial

Based on Eqs. (14) and (21), we get an equation system of three polynomials for four unknowns q_y , \tilde{t}_x , \tilde{t}_y and \tilde{t}_z . Recall that there are three independent constraints provided by one AC. Thus, one more equation is required which can be taken from a second AC. In principle, one arbitrary equation can be chosen from Eqs. (14) and (21), for example,

three constraints of the first AC, and the first constraint of the second AC are stacked into 4 equations in 4 unknowns:

$$\frac{1}{1 + q_y^2} \underbrace{\begin{bmatrix} \tilde{M}_{11} & \tilde{M}_{12} & \tilde{M}_{13} & \tilde{M}_{14} \\ \tilde{M}_{21} & \tilde{M}_{22} & \tilde{M}_{23} & \tilde{M}_{24} \\ \tilde{M}_{31} & \tilde{M}_{32} & \tilde{M}_{33} & \tilde{M}_{34} \\ \tilde{M}_{41} & \tilde{M}_{42} & \tilde{M}_{43} & \tilde{M}_{44} \end{bmatrix}}_{\tilde{\mathbf{M}}(q_y)} \begin{bmatrix} \tilde{t}_x \\ \tilde{t}_y \\ \tilde{t}_z \\ 1 \end{bmatrix} = \mathbf{0}, \quad (22)$$

where the elements \tilde{M}_{ij} ($i = 1, \dots, 4; j = 1, \dots, 4$) of the coefficient matrix $\tilde{\mathbf{M}}(q_y)$ are formed by the polynomial coefficients and one unknown variable q_y , see supplementary material for details. Since $\tilde{\mathbf{M}}(q_y)/(1 + q_y^2)$ is a square matrix, Eq. (22) has a non-trivial solution only if the $\det(\tilde{\mathbf{M}}(q_y)/(1 + q_y^2)) = 0$. The expansion of the determinant equation gives a 6-degree univariate polynomial:

$$\text{quot}(\sum_{i=0}^8 w_i q_y^i, q_y^2 + 1) = 0, \quad (23)$$

where $\tilde{w}_0, \dots, \tilde{w}_8$ are formed by two Plücker line correspondences and two affine transformations between the corresponding feature points.

This univariate polynomial leads to a maximum of 6 solutions. Equation (23) can be efficiently solved by the companion matrix method [11] or Sturm bracketing method [39]. Once q_y has been obtained, the rotation matrix \mathbf{R}_y is recovered from Eq. (10). For the relative pose between two multi-camera reference frames at time k and $k + 1$, the rotation matrix \mathbf{R} is recovered from Eq. (13) and the translation is computed by $\mathbf{t} = (\mathbf{R}'_{imu})^T \tilde{\mathbf{t}}$.

6. Experiments

In this section, we conduct extensive experiments on both synthetic and real-world data to evaluate the performance of the proposed methods. Our solvers are compared with state-of-the-art techniques.

For relative pose estimation under planar motion, the solvers using one AC and two ACs proposed in Section 4 are referred to as 1AC plane method and 2AC plane method, respectively. The accuracy of 1AC plane and 2AC plane are compared with the methods 17pt-Li [32], 8pt-Kneip [27], 6pt-Stewénus [24] and 6AC-Ventura [2].

For relative pose estimation with known vertical direction, the solver proposed in Section 5 is referred to as 2AC vertical method. We compare the accuracy of 2AC vertical with the methods 17pt-Li [32], 8pt-Kneip [27], 6pt-Stewénus [24], 4pt-Lee [30], 4pt-Sweeney [48], 4pt-Liu [34] and 6AC-Ventura [2].

A single run of the proposed solvers 1AC plane, 2AC plane and 2AC vertical take 3.6, 3.6 and 17.8 μs in C++, respectively. Due to space limitations, the

efficiency comparison and stability study are provided in the supplementary material. In the experiments, all the solvers are integrated within RANSAC to reject outliers. For the point-based solvers, only the point coordinates of ACs are used. The relative pose which produces the highest number of inliers is chosen. The confidence of RANSAC is set to 0.99 and an inlier threshold angle is set to 0.1° by following the definition in OpenGV [26]. We also show the feasibility of our methods on the KITTI dataset [15]. This experiment demonstrates that our methods are well suited for visual odometry in road driving scenarios.

6.1. Experiments on Synthetic Data

We made a simulated 2-camera rig system by following the KITTI autonomous driving platform. The baseline length between two simulated cameras is set to 1 meter and the cameras are installed at different altitude. The multi-camera reference frame is set at the center of the camera rig and the translation between two multi-camera reference frames is 3 meters. The resolution of the cameras is 640×480 pixels and the focal lengths are 400 pixels. The principal points are set to the image center (320, 240).

The synthetic scene is composed of a ground plane and 50 random planes. All 3D planes are randomly generated within the range of -5 to 5 meters (along axes X and Y), and 10 to 20 meters (Z-axis direction), that are expressed in the respective axis of the multi-camera reference frame. We choose 50 ACs from the ground plane and an AC from each random plane randomly, thus, having a total of 100 ACs. For each AC, a random 3D point from a plane is reprojected onto two cameras to get the image point pair. The corresponding affine transformation is obtained by the following procedure. First, four points are chosen as the vertices of a square in view 1, where the center of the square is the point coordinates of AC. The side length of the square is set as 20 or 40 pixels. A larger side length causes smaller noise of affine transformation. Second, the four corresponding points in view 2 are calculated by the ground truth homography. Third, four sampled point pairs are contaminated by Gaussian noise, which is similar to the noise added to the coordinates of image point pair. Fourth, the noisy homography matrix is estimated using the four sampled point pairs. The noisy affine transformation is the first-order approximation of the noisy homography matrix. This procedure enables an indirect but geometrically interpretable way of adding noise to the affine transformation [5].

A total of 1000 trials are carried out in the synthetic experiment. In each test, 100 ACs are generated randomly. The ACs for the methods are selected randomly and the error is measured on the relative pose which produces the most inliers within the RANSAC scheme. This also allows us to select the best candidate from multiple solutions by counting their inliers in a RANSAC-like procedure. The

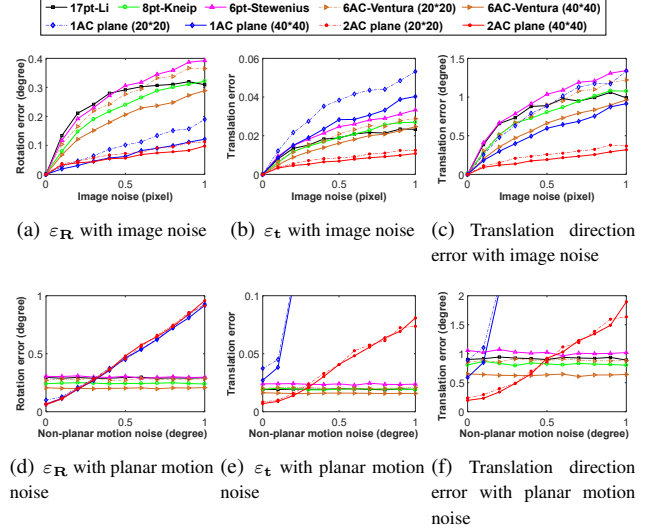


Figure 3. Rotation and translation error under planar motion. (a–c): varying image noise under perfect planar motion. (d–f): varying planar motion noise and fixed 1.0 pixel std. image noise.

median of errors are used to assess the rotation and translation error. The rotation error is computed as the angular difference between the ground truth rotation and the estimated rotation: $\varepsilon_R = \arccos((\text{trace}(\mathbf{R}_{gt}\mathbf{R}^T) - 1)/2)$, where \mathbf{R}_{gt} and \mathbf{R} are the ground truth and estimated rotation matrices. Following the definition in [42, 30], the translation error is defined as: $\varepsilon_t = 2 \|(\mathbf{t}_{gt} - \mathbf{t})\| / (\|\mathbf{t}_{gt}\| + \|\mathbf{t}\|)$, where \mathbf{t}_{gt} and \mathbf{t} are the ground truth and estimated translations.

6.1.1 Planar Motion Estimation

In this scenario, the planar motion of the multi-camera system is described by (θ, ϕ) , see Fig. 2(a). The magnitudes of both angles ranges from -10° to 10° . We use Gaussian image noise with a standard deviation ranging from 0 to 1 pixel. Fig. 3(a–c) shows the performance of the proposed 1AC plane and 2AC plane methods against image noise. Since the noise magnitude of affine transformation is influenced by the support region of sampled points, the AC-based methods have better performance with larger support region at the same magnitude of image noise. It can be seen that 2AC plane performs better than the other compared methods under perfect planar motion, even though the size of the square is 20 pixels. The 1AC plane method performs better than the PC-based methods and the 6AC-Ventura method in rotation estimation, but has worse performance in translation estimation. As shown in Fig. 3(c), we plot the translation direction error as an additional evaluation. It is interesting to see that when the side length of the square is 40 pixels, the 1AC plane method performs better than the PC-based methods and the 6AC-Ventura method in translation direction estimation.

We also evaluate the accuracy of the proposed methods 1AC plane and 2AC plane for increasing planar motion noise. To test such noise, we added a small randomly generated X-axis, Z-axis rotation and a YZ-plane translation [9] to the motion of the multi-camera system. The magnitude of non-planar motion noise ranges from 0° to 1° and the standard deviation of the image noise is set to 0.5 pixel. Figures 3(d–f) show the performance of the proposed 1AC plane method and 2AC plane method against planar motion noise. Methods 17pt-Li, 8pt-Kneip, 6pt-Stewenius and 6AC-Ventura deal with the 6DOF motion case and, thus they are not affected by the noise in the planarity assumption. It can be seen that the rotation accuracy of the 2AC plane method performs better than comparative methods when the planar motion noise is less than 0.2° . Since the estimation accuracy of translation direction of the 2AC plane method in Fig. 3(f) performs satisfactory, the main reason for poor performance of translation estimation is that the metric scale estimation is sensitive to the planar motion noise. In comparison with the 2AC plane method, the 1AC plane method has similar performance in rotation estimation, but performs poorly in translation estimation. The translation accuracy decreases significantly with the increase of the planar motion noise.

Both the 1AC plane and 2AC plane methods have a significant computational advantage over comparative methods, because the efficient solver for 4-degree polynomial equation takes only about $3.6 \mu s$. Moreover, since only a single AC is required, the 1AC plane method has the advantage of detecting a correct inlier set efficiently, which can then be used for accurate motion estimation with non-linear optimization. See supplementary material for details.

6.1.2 Motion with Known Vertical Direction

In this set of experiments, the translation direction of two multi-camera reference frames is chosen to produce either forward, sideways or random motions. The second reference frame is rotated around three axes randomly with angles ranging from -10° to 10° . Assuming known roll and pitch angles, the multi-camera reference frame is aligned with the vertical direction. Due to space limitations, we only show the results for random motion. Other results are in the supplementary material. Figs. 4(a) and (d) show the performance of 2AC vertical against image noise with perfect IMU data. The proposed method is robust to image noise and performs better than the other methods.

Figs. 4(b,e) and (c,f) show the performance of 2AC vertical against IMU noise in the random motion case, while the standard deviation of the image noise is fixed at 0.5 pixel. Note that the methods 17pt-Li, 8pt-Kneip, 6pt-Stewenius and 6AC-Ventura are not influenced by IMU noise, because these methods do not

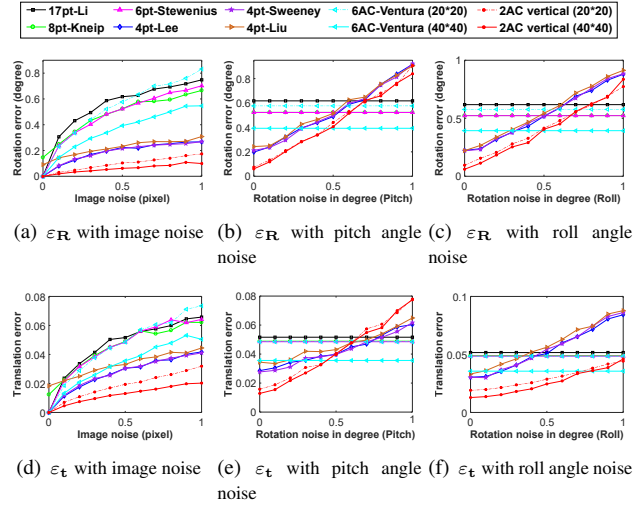


Figure 4. Rotation and translation error under random motion with known vertical direction. Upper row: rotation error. Bottom row: translation error. (a,d): varying image noise. (b,e) and (c,f): varying IMU angle noise and fixed 1.0 pixel std. image noise.

use the known vertical direction as a prior. The methods 4pt-Lee, 4pt-Sweeney and 4pt-Liu use the known vertical direction as a prior. It is interesting to see that the proposed method outperforms the comparative methods in the random motion case, even though the IMU noise is around 0.4° . The results under forward and sideways motion also demonstrate that the 2AC vertical method performs basically better than all comparative methods against image noise and provides comparable accuracy for increasing IMU noise.

6.2. Experiments on Real Data

We test the performance of our methods on the KITTI dataset [15] that consists of successive video frames from a forward facing stereo camera. The ground truth pose is provided from the built-in GPS/IMU units. We ignore the overlap in their fields of view and treat it as a general multi-camera system. The sequences labeled from 0 to 10, which have ground truth, are used for the evaluation. Therefore, the methods were tested on a total of 23000 image pairs. The ACs between consecutive frames in each camera are established by applying the ASIFT [38] detector. The extraction of ACs can also be sped up by MSER [37], GPU acceleration, or approximating ACs from SIFT features for subsequent video frames. The ACs across the two cameras are not matched and the metric scale is not estimated as the movement between consecutive frames is small. Besides, integrating the acceleration over time from an IMU is more suitable for recovering the scale [40]. All the solvers have been integrated into a RANSAC scheme.

Sequence	17pt-Li [32]		8pt-Kneip [27]		6pt-St. [24]		4pt-Lee [30]		4pt-Sw. [48]		4pt-Liu [34]		6AC-Ven. [2]		2AC plane		2AC vertical	
	ϵ_R	ϵ_t	ϵ_R	ϵ_t	ϵ_R	ϵ_t	ϵ_R	ϵ_t	ϵ_R	ϵ_t	ϵ_R	ϵ_t	ϵ_R	ϵ_t	ϵ_R	ϵ_t	ϵ_R	ϵ_t
00 (4541 images)	0.139	2.412	0.130	2.400	0.229	4.007	0.065	2.469	0.050	2.190	0.066	2.519	0.142	2.499	0.280	2.243	0.031	1.738
01 (1101 images)	0.158	5.231	0.171	4.102	0.762	41.19	0.137	4.782	0.125	11.91	0.105	3.781	0.146	3.654	0.168	2.486	0.025	1.428
02 (4661 images)	0.123	1.740	0.126	1.739	0.186	2.508	0.057	1.825	0.044	1.579	0.057	1.821	0.121	1.702	0.213	1.975	0.030	1.558
03 (801 images)	0.115	2.744	0.108	2.805	0.265	6.191	0.064	3.116	0.069	3.712	0.062	3.258	0.113	2.731	0.238	1.849	0.037	1.888
04 (271 images)	0.099	1.560	0.116	1.746	0.202	3.619	0.050	1.564	0.051	1.708	0.045	1.635	0.100	1.725	0.116	1.768	0.020	1.228
05 (2761 images)	0.119	2.289	0.112	2.281	0.199	4.155	0.054	2.337	0.052	2.544	0.056	2.406	0.116	2.273	0.185	2.354	0.022	1.532
06 (1101 images)	0.116	2.071	0.118	1.862	0.168	2.739	0.053	1.757	0.092	2.721	0.056	1.760	0.115	1.956	0.137	2.247	0.023	1.303
07 (1101 images)	0.119	3.002	0.112	3.029	0.245	6.397	0.058	2.810	0.065	4.554	0.054	3.048	0.137	2.892	0.173	2.902	0.023	1.820
08 (4071 images)	0.116	2.386	0.111	2.349	0.196	3.909	0.051	2.433	0.046	2.422	0.053	2.457	0.108	2.344	0.203	2.569	0.024	1.911
09 (1591 images)	0.133	1.977	0.125	1.806	0.179	2.592	0.056	1.838	0.046	1.656	0.058	1.793	0.124	1.876	0.189	1.997	0.027	1.440
10 (1201 images)	0.127	1.889	0.115	1.893	0.201	2.781	0.052	1.932	0.040	1.658	0.058	1.888	0.203	2.057	0.223	2.296	0.025	1.586

Table 2. Rotation and translation error on KITTI sequences (unit: degree).

Methods	17pt-Li [32]	8pt-Kneip [27]	6pt-St. [24]	4pt-Lee [30]	4pt-Sw. [48]	4pt-Liu [34]	6AC-Ven. [2]	2AC plane	2AC vertical
Mean time	52.82	10.36	79.76	0.85	0.63	0.45	6.83	0.07	0.09
Standard deviation	2.62	1.59	4.52	0.093	0.057	0.058	0.61	0.0071	0.0086

Table 3. Runtime of RANSAC averaged over KITTI sequences combined with different solvers (unit: s).

The proposed methods 2AC plane and 2AC vertical are compared against 17pt-Li [32], 8pt-Kneip [27], 6pt-Stewénus [24], 4pt-Lee [30], 4pt-Sweeney [48], 4pt-Liu [34] and 6AC-Ventura [2]. Since the KITTI dataset is captured by a stereo rig with both cameras having the same altitude, that is a degenerate case for the 1AC plane method, it is not performed in the experiment. For the 2AC plane method, the results are also compared to the ground truth of the 6DOF relative pose, even though this method only estimates two angles (θ , ϕ) with the plane motion assumption. For the 2AC vertical method, the roll and pitch angles obtained from the GPS/IMU units are used to align the multi-camera reference frame with the vertical direction [44, 18, 31]. To ensure the fairness of the experiment, the roll and pitch angles are also provided for the methods 4pt-Lee [30], 4pt-Sweeney [48] and 4pt-Liu [34], which use the known vertical direction as a prior. Table 2 shows the results of the rotation and translation estimation. The median error for each individual sequence is used to evaluate the estimation accuracy. The runtime of RANSAC averaged over KITTI sequences combined with different solvers is shown in Table 3. The reported runtimes include the robust relative pose estimation without feature extraction, *i.e.*, recovering the relative pose by RANSAC combined with a minimal solver.

The proposed 2AC vertical method offers the best overall performance among all the methods. The 6pt-Stewénus method performs poorly on sequence 01, because this sequence is a highway with few tractable close objects, and this method always fails to select the best candidate from multiple solutions under forward motion in the RANSAC scheme. Besides, it is interesting to see that

the translation accuracy of the 2AC plane method basically outperforms the 6pt-Stewénus method, even though the planar motion assumption does not fit the KITTI dataset well. To visualize the comparison results, the estimated trajectory for sequence 00 is plotted in the supplementary material. Due to the benefits of computational efficiency, both the 2AC plane method and the 2AC vertical method are quite suitable for finding a correct inlier set, which is then used for accurate motion estimation in visual odometry.

7. Conclusion

By exploiting the geometric constraints which interprets the relationship of ACs and the generalized camera model, we have proposed three solutions for the relative pose estimation of a multi-camera system. Under the planar motion assumption, we present two solvers to recover the relative pose of a multi-camera system, including a minimal solver using a single AC and a solver based on two ACs. In addition, a minimal solution with two ACs is proposed to solve for the relative pose of the multi-camera system with known vertical direction. Both planar motion and known vertical direction assumptions are realistic in autonomous driving scenes. We evaluate the proposed solvers on synthetic data and real image sequence datasets. The experimental results clearly showed that the proposed methods provide better efficiency and accuracy for relative pose estimation in comparison to state-of-the-art methods.

Acknowledgments

This work has been partially funded by the National Natural Science Foundation of China (11902349, 11727804).

References

- [1] Sameer Agarwal, Hon-Leung Lee, Bernd Sturmfels, and Rekha R. Thomas. On the existence of epipolar matrices. *International Journal of Computer Vision*, 121(3):403–415, 2017. [1](#)
- [2] Khaled Alyousefi and Jonathan Ventura. Multi-camera motion estimation with affine correspondences. In *International Conference on Image Analysis and Recognition*, pages 417–431, 2020. [2](#), [5](#), [8](#)
- [3] Daniel Barath. Five-point fundamental matrix estimation for uncalibrated cameras. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 235–243, 2018. [3](#)
- [4] Daniel Barath and Levente Hajder. Efficient recovery of essential matrix from two affine correspondences. *IEEE Transactions on Image Processing*, 27(11):5328–5337, 2018. [2](#), [3](#)
- [5] Daniel Barath and Zuzana Kukelova. Homography from two orientation-and scale-covariant features. In *IEEE International Conference on Computer Vision*, pages 1091–1099, 2019. [6](#)
- [6] Daniel Barath, Michal Polic, Wolfgang Förstner, Torsten Sattler, Tomas Pajdla, and Zuzana Kukelova. Making affine correspondences work in camera geometry computation. In *European Conference on Computer Vision*, 2020. [1](#)
- [7] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (SURF). *Computer Vision and Image Understanding*, 110(3):346–359, 2008. [2](#)
- [8] Jacob Bentolila and Joseph M Francos. Conic epipolar constraints from affine correspondences. *Computer Vision and Image Understanding*, 122:105–114, 2014. [2](#)
- [9] Sunglok Choi and Jong-Hwan Kim. Fast and reliable minimal relative pose estimation under planar motion. *Image and Vision Computing*, 69:103–112, 2018. [2](#), [7](#)
- [10] Brian Clipp, Jae-Hak Kim, Jan-Michael Frahm, Marc Pollefeys, and Richard Hartley. Robust 6dof motion estimation for non-overlapping, multi-camera systems. In *IEEE Workshop on Applications of Computer Vision*, pages 1–8. IEEE, 2008. [1](#)
- [11] David Cox, John Little, and Donal O’Shea. *Ideals, varieties, and algorithms: An introduction to computational algebraic geometry and commutative algebra*. Springer Science & Business Media, 2013. [5](#)
- [12] Iván Eichhardt and Daniel Barath. Relative pose from deep learned depth and a single affine correspondence. In *European Conference on Computer Vision*, 2020. [1](#)
- [13] Martin A Fischler and Robert C Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. [1](#)
- [14] Victor Fragoso, Joseph DeGol, and Gang Hua. gdl*: Generalized pose-and-scale estimation given scale and gravity priors. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2210–2219, 2020. [1](#)
- [15] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. [6](#), [7](#)
- [16] Banglei Guan, Pascal Vasseur, Cédric Demonceaux, and Friedrich Fraundorfer. Visual odometry using a homography formulation with decoupled rotation and translation estimation using minimal solutions. In *IEEE International Conference on Robotics and Automation*, pages 2320–2327, 2018. [1](#)
- [17] Banglei Guan, Ji Zhao, Daniel Barath, and Friedrich Fraundorfer. Efficient recovery of multi-camera motion from two affine correspondences. In *IEEE International Conference on Robotics and Automation*, pages 1–6, 2021. [2](#)
- [18] Banglei Guan, Ji Zhao, Zhang Li, Fang Sun, and Friedrich Fraundorfer. Minimal solutions for relative pose with a single affine correspondence. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1929–1938, 2020. [1](#), [2](#), [8](#)
- [19] Banglei Guan, Ji Zhao, Zhang Li, Fang Sun, and Friedrich Fraundorfer. Relative pose estimation with a single affine correspondence. *IEEE Transactions on Cybernetics*, pages 1–12, 2021. [2](#)
- [20] Levente Hajder and Daniel Barath. Relative planar motion for vehicle-mounted cameras from a single affine correspondence. In *IEEE International Conference on Robotics and Automation*, pages 8651–8657, 2020. [2](#)
- [21] Christian Häne, Lionel Heng, Gim Hee Lee, Friedrich Fraundorfer, Paul Furgale, Torsten Sattler, and Marc Pollefeys. 3D visual perception for self-driving cars using a multi-camera system: Calibration, mapping, localization, and obstacle detection. *Image and Vision Computing*, 68:14–27, 2017. [1](#)
- [22] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge University Press, 2003. [1](#)
- [23] Lionel Heng, Benjamin Choi, Zhaopeng Cui, Marcel Gepert, Sixing Hu, Benson Kuan, Peidong Liu, Rang Nguyen, Ye Chuan Yeo, Andreas Geiger, Gim Hee Lee, Marc Pollefeys, and Torsten Sattler. Project AutoVision: Localization and 3D scene perception for an autonomous vehicle with a multi-camera system. In *IEEE International Conference on Robotics and Automation*, pages 4695–4702, 2019. [1](#)
- [24] Stewénus Henrik, Oskarsson Magnus, Kalle Aström, and David Nistér. Solutions to minimal generalized relative pose problems. In *Workshop on Omnidirectional Vision in conjunction with ICCV*, pages 1–8, 2005. [1](#), [2](#), [5](#), [8](#)
- [25] Jae-Hak Kim, Hongdong Li, and Richard Hartley. Motion estimation for nonoverlapping multicamera rigs: Linear algebraic and L_∞ geometric solutions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1044–1059, 2009. [1](#)
- [26] Laurent Kneip and Paul Furgale. OpenGV: A unified and generalized approach to real-time calibrated geometric vision. In *IEEE International Conference on Robotics and Automation*, pages 12034–12043, 2014. [6](#)
- [27] Laurent Kneip and Hongdong Li. Efficient computation of relative pose for multi-camera systems. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 446–453, 2014. [2](#), [5](#), [8](#)
- [28] Laurent Kneip, Chris Sweeney, and Richard Hartley. The generalized relative pose and scale problem: View-graph fu-

- sion via 2D-2D registration. In *IEEE Winter Conference on Applications of Computer Vision*, pages 1–9, 2016. [1](#)
- [29] Gim Hee Lee, Friedrich Fraundorfer, and Marc Pollefeys. Motion estimation for self-driving cars with a generalized camera. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2746–2753, 2013. [2](#)
- [30] Gim Hee Lee, Marc Pollefeys, and Friedrich Fraundorfer. Relative pose estimation for a multi-camera system with known vertical direction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 540–547, 2014. [1](#), [2](#), [5](#), [6](#), [8](#)
- [31] Bo Li, Evgeniy Martyshev, and Gim Hee Lee. Relative pose estimation of calibrated cameras with known SE(3) invariants. In *European Conference on Computer Vision*, 2020. [1](#), [2](#), [8](#)
- [32] Hongdong Li, Richard Hartley, and Jae-hak Kim. A linear approach to motion estimation using generalized camera models. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. [1](#), [2](#), [5](#), [8](#)
- [33] John Lim, Nick Barnes, and Hongdong Li. Estimating relative camera motion from the antipodal-epipolar constraint. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(10):1907–1914, 2010. [1](#)
- [34] Liu Liu, Hongdong Li, Yuchao Dai, and Quan Pan. Robust and efficient relative pose with a multi-camera system for autonomous driving in highly dynamic environments. *IEEE Transactions on Intelligent Transportation Systems*, 19(8):2432–2444, 2017. [2](#), [5](#), [8](#)
- [35] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. [2](#)
- [36] Evgeniy Martyshev and Bo Li. Efficient relative pose estimation for cameras and generalized cameras in case of known relative rotation angle. *Journal of Mathematical Imaging and Vision*, (10), 2020. [2](#)
- [37] Jiri Matas, Ondrej Chum, Martin Urban, and Tomás Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing*, 22(10):761–767, 2004. [7](#)
- [38] Jean-Michel Morel and Guoshen Yu. ASIFT: A new framework for fully affine invariant image comparison. *SIAM Journal on Imaging Sciences*, 2(2):438–469, 2009. [7](#)
- [39] David Nistér. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):756–777, 2004. [1](#), [5](#)
- [40] Gabriel Ntzi, Stephan Weiss, Davide Scaramuzza, and Roland Siegwart. Fusion of IMU and vision for absolute scale estimation in monocular SLAM. *Journal of intelligent & robotic systems*, 61(1-4):287–299, 2011. [7](#)
- [41] Robert Pless. Using many cameras as one. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages II–587, 2003. [1](#), [2](#), [3](#)
- [42] Long Quan and Zhongdan Lan. Linear n-point camera pose determination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):774–780, 1999. [6](#)
- [43] Carolina Raposo and Joao P Barreto. Theory and practice of structure-from-motion using affine correspondences. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5470–5478, 2016. [2](#)
- [44] Olivier Saurer, Pascal Vasseur, Rémi Boutteau, Cédric Demonceaux, Marc Pollefeys, and Friedrich Fraundorfer. Homography based egomotion estimation with a common direction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(2):327–341, 2016. [2](#), [8](#)
- [45] Davide Scaramuzza and Friedrich Fraundorfer. Visual odometry: The first 30 years and fundamentals. *IEEE Robotics & Automation Magazine*, 18(4):80–92, 2011. [1](#)
- [46] Johannes L Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. [1](#)
- [47] Chris Sweeney, John Flynn, Benjamin Nuernberger, Matthew Turk, and Tobias Höllerer. Efficient computation of absolute pose for gravity-aware augmented reality. In *IEEE International Symposium on Mixed and Augmented Reality*, pages 19–24, 2015. [1](#)
- [48] Chris Sweeney, John Flynn, and Matthew Turk. Solving for relative pose with a partially known rotation is a quadratic eigenvalue problem. In *IEEE International Conference on 3D Vision*, pages 483–490, 2014. [2](#), [5](#), [8](#)
- [49] Chris Sweeney, Laurent Kneip, Tobias Hollerer, and Matthew Turk. Computing similarity transformations from only image correspondences. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3305–3313, 2015. [2](#)
- [50] Jonathan Ventura, Clemens Arth, and Vincent Lepetit. An efficient minimal solution for multi-camera motion. In *IEEE International Conference on Computer Vision*, pages 747–755, 2015. [1](#), [2](#)
- [51] Ji Zhao, Wanting Xu, and Laurent Kneip. A certifiably globally optimal solution to generalized essential matrix estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12034–12043, 2020. [1](#)