
Mirror-Based Extrinsic Camera Calibration

Joel A. Hesch¹, Anastasios I. Mourikis², and Stergios I. Roumeliotis¹

¹ Department of Computer Science and Engineering, University of Minnesota,
{joel, stergios}@cs.umn.edu

² Department of Electrical Engineering, University of California, Riverside,
mourikis@ee.ucr.edu

Abstract: This paper presents a method for determining the six degrees-of-freedom transformation between a camera and a base frame of interest. A planar mirror is maneuvered so as to allow the camera to observe the environment from several viewing angles. Points, whose coordinates in the base frame are known, are observed by the camera via their reflections in the mirror. Exploiting these measurements, we determine the camera-to-base transformation analytically, without assuming prior knowledge of the mirror motion or placement with respect to the camera. The computed solution is refined using a maximum likelihood estimator, to obtain high-accuracy estimates of the camera-to-base transformation and the mirror configuration for each image. We validate the accuracy and correctness of our method with simulations and real-world experiments.

1 Introduction

Cameras are utilized in a wide variety of applications ranging from surveillance and crowd monitoring, to vision-based robot localization. In order to obtain meaningful geometric information from a camera, two calibration procedures must be completed. The first is intrinsic calibration, that is, determining the internal camera parameters (e.g., focal length, principal point, and skew coefficients), which affect the image measurements. The second is extrinsic calibration, which is the process of computing the transformation between the camera and a base frame of reference. In a surveillance application, the base frame may be the room or building coordinate system, whereas on a mobile robot, the base frame could be the



Fig. 1. A camera equipped robot views itself in a mirror.

robot-body frame. Several works have addressed extrinsic calibration of a camera to another sensor (e.g., for odometry-to-camera [14], inertial measurement unit (IMU)-to-camera [15, 16], or laser scanner-to-camera [1, 23]). These exploit measurements from both sensors to determine their mutual transformation. However, very little attention has been devoted to determining the *camera-to-base* transformation, for a generic base frame.

In this paper, we deal exclusively with extrinsic camera calibration. Our objective is to determine the camera-to-base transformation from observations of points whose coordinates in the base frame are known. We consider the most limiting case, in which the known points do not lie within the camera’s field of view but can only be observed using a mirror (cf. Fig. 1). We maneuver a planar mirror in front of the camera to provide multiple views of the points. In our formulation, no prior information about the mirror motion or placement with respect to the camera is assumed. The configuration of the mirror and the camera-to-base transformation are treated as unknowns to be computed from the observations. The main contribution of this paper is an algorithm for determining the camera-to-base transformation *analytically*, which requires a minimum of 3 noncollinear points tracked in 3 images.

A direct approach to extrinsic camera calibration is to utilize all of the measurements in a maximum-likelihood estimator (MLE) for computing the unknown transformation [8]. This takes the form of a nonlinear least-squares problem, which seeks to iteratively minimize a nonconvex function of the unknown variables. While appealing for its ease of implementation, this method has a number of drawbacks. Without an accurate initial guess, the minimization process may take several iterations to converge, or even fail to find the correct solution. Additionally, the MLE provides no framework for studying the minimal measurement conditions required to compute a solution. For these reasons, we first determine the transformation analytically, and then employ an MLE to refine the computed solution. The rest of the paper is organized as follows. The proposed algorithm is discussed in Section 3, including presentations of the measurement model, the analytical solution, and the MLE utilized to refine the computed solution. This algorithm is tested both in simulation (Section 4) and experimentally (Section 5) to evaluate its accuracy and investigate its sensitivity to the system parameters. Section 6 discusses the conclusions of this work and suggests interesting directions of future research. Lastly, we comment on the extension of this work to robot-body cartography (Appendix B).

2 Related Work

We first review the related work, which falls into two categories: (i) hand-eye calibration, and (ii) catadioptric systems. Hand-eye calibration is the process of determining the six degrees-of-freedom (6 dof) transformation between a camera and a tool, which are both mounted on a robot manipulator [22, 2].

The hand-eye problem is solved by correlating the measurements of the camera and the encoders, which measure displacements of the robot joints. This process determines the pose of the camera with respect to the robot base. Subsequently, the camera-to-tool transformation is calculated by combining the estimated camera-to-robot-base transformation, and the robot-base-to-tool transformation, which is assumed to be known. This assumption necessitates the availability of precise technical drawings, and limits the applicability of these methods, since they cannot determine the camera-to-base transformation for a generic base frame.

We turn our attention to catadioptric systems, which are employed to perform synthetic multiple-view vision. Methods have been presented utilizing a single camera and planar [3, 4, 9], or conic mirrors [10]. Others accomplish stereo vision with reflections from free-form surfaces [24], and a trinocular mirror-based vision system also exists [19]. Additionally, stereo is achieved with a static camera and a moving mirror [12], or with a moving camera and two stationary spherical mirrors of known radii [17]. In contrast to the approaches above, we do not perform synthetic stereo, i.e., only a single observation of each point is available in each image.

Jang *et al.* demonstrate a moving planar mirror system for 3D scene reconstruction [11]. Exploiting a combination of fiducial points on the mirror and vanishing points in the reflections, they solve for the position of the mirror with respect to the camera. The 3D scene is determined based on synthetic stereo from multiple reflections. In contrast to this approach, we do not assume that the dimensions of the mirror, or its position with respect to the camera, are available. Finally, Kumar *et al.* determine the transformations between multiple cameras with non-overlapping fields of view, using mirror reflections of a calibration grid [13]. They require 5 views (per camera) of the calibration pattern to form a set of linear constraints which are solved for the unknown transformations. In contrast to [13], our method requires only 3 images, each containing observations of 3 known points, to determine the camera-to-base transformation analytically.

3 Computing the Transformation

In this section, we describe our approach for analytically determining the transformation between the camera frame, $\{C\}$, and a frame of interest, $\{B\}$, from observations of 3 points whose coordinates in $\{B\}$ are known. Frame $\{B\}$ is arbitrary and without loss of generality, we will refer to $\{B\}$ as the “base frame.” Example base frames vary by application, and may include (i) the robot-body frame, if the camera is mounted on a robot, (ii) the room or building frame, if the camera is utilized in a surveillance application, and (iii) the rig mount, if the camera is part of a stereo pair.

We address the most limiting scenario in which the points are only visible through reflections in a planar mirror that is moved in front of the camera

Algorithm 1 Computing the Camera to Base-Frame Transformation

Input: Set of 3 points tracked in N_c images**Output:** The camera-to-base transformation $\{{}_B^C\mathbf{R}, {}^C\mathbf{p}_B\}$ **for each** image in N_c **do**

Convert to three point pose estimation problem (P3P)

 Solve P3P to obtain combined homogeneous/reflection transformation: $\{\mathbf{A}, \mathbf{b}\}_k$ **end for****for each** triplet of solutions $\{\mathbf{A}, \mathbf{b}\}_k, \{\mathbf{A}, \mathbf{b}\}_{k'}, \{\mathbf{A}, \mathbf{b}\}_{k''}$ **do**

Compute mirror configurations from (14)

 Compute camera-to-base rotation ${}_B^C\mathbf{R}$ from (23) Compute camera-to-base translation ${}^C\mathbf{p}_B$ from (16)**end for**Utilize clustering to select the correct solution $\{{}_B^C\mathbf{R}, {}^C\mathbf{p}_B\}$ Refine the solution using a maximum-likelihood estimator

to provide multiple views of the scene. We exploit these observations to compute the transformation between $\{B\}$ and $\{C\}$, without knowledge of the mirror’s placement or motion with respect to the camera (cf. Algorithm 1). In what follows, we present the measurement model and discuss its relation with the three point pose estimation problem (P3P). Additionally, we comment on the unobservable cases of the system under consideration, and present an analytical method to compute the unknown transformation from a minimum of 3 points observed in 3 images which differ by rotations about two axes. Lastly, we summarize a maximum-likelihood approach for refining the computed transformation, a detailed discussion of which is available in [8].

3.1 Measurement Model

First, we present the measurement model that describes each of the camera observations. To simplify the presentation, in this section we focus on the case of a single point, observed in a single image. Consider a point \mathbf{p} , whose position with respect to frame $\{B\}$, ${}^B\mathbf{p}$, is known¹. We seek to express the point \mathbf{p} in the camera reference frame $\{C\}$. From geometry (cf. Fig. 2) we have two constraint equations:

$${}^C\mathbf{p}' = {}^C\mathbf{p} + 2d_p {}^C\mathbf{n} \quad (1)$$

$$d_p = d - {}^C\mathbf{n}^T {}^C\mathbf{p} \quad (2)$$

where ${}^C\mathbf{p}'$ is the reflection of ${}^C\mathbf{p}$, ${}^C\mathbf{n}$ is the mirror normal vector expressed in the camera frame, d is the distance between the mirror and the camera, and

¹ Throughout this paper, ${}^X\mathbf{y}$ denotes a vector \mathbf{y} expressed with respect to frame $\{X\}$, ${}^X_W\mathbf{R}$ is the rotation matrix rotating vectors from frame $\{W\}$ to frame $\{X\}$, and ${}^X\mathbf{p}_W$ is the position of the origin of frame $\{W\}$, expressed with respect to frame $\{X\}$. \mathbf{I}_n is the $n \times n$ identity matrix, and $\mathbf{0}_{m \times n}$ is the $m \times n$ matrix of zeros.

d_p is the distance between the mirror and the known point (both distances are defined along the mirror normal vector). Note also that

$${}^C\mathbf{p} = {}^C_B\mathbf{R}^B\mathbf{p} + {}^C\mathbf{p}_B, \quad (3)$$

where ${}^C_B\mathbf{R}$ is the matrix which rotates vectors between frames $\{B\}$ and $\{C\}$, and ${}^C\mathbf{p}_B$ is the origin of $\{B\}$ with respect to $\{C\}$. We substitute (2) and (3) into (1), and rearrange the terms to obtain:

$$\begin{aligned} {}^C\mathbf{p}' &= {}^C\mathbf{p} + 2d_p {}^C\mathbf{n} \\ &= (\mathbf{I}_3 - 2{}^C\mathbf{n}{}^C\mathbf{n}^T) {}^C\mathbf{p} + 2d {}^C\mathbf{n} \\ &= (\mathbf{I}_3 - 2{}^C\mathbf{n}{}^C\mathbf{n}^T) ({}^C_B\mathbf{R}^B\mathbf{p} + {}^C\mathbf{p}_B) + 2d {}^C\mathbf{n} \end{aligned} \quad (4)$$

which can be written in terms of homogeneous coordinates as:

$$\begin{bmatrix} {}^C\mathbf{p}' \\ 1 \end{bmatrix} = \begin{bmatrix} (\mathbf{I}_3 - 2{}^C\mathbf{n}{}^C\mathbf{n}^T) & 2d {}^C\mathbf{n} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} {}^C_B\mathbf{R} & {}^C\mathbf{p}_B \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} {}^B\mathbf{p} \\ 1 \end{bmatrix} \quad (5)$$

This expression will be useful in the ensuing derivations. The reflection of \mathbf{p} is observed by the camera, and this measurement is described by the perspective projection model:

$$\mathbf{z} = \frac{1}{p_3} \begin{bmatrix} p_1 \\ p_2 \end{bmatrix} + \boldsymbol{\eta} = \mathbf{h}({}^C\mathbf{p}') + \boldsymbol{\eta}, \quad {}^C\mathbf{p}' = [p_1 \ p_2 \ p_3]^T \quad (6)$$

where $\boldsymbol{\eta}$ is the pixel noise, assumed to be zero-mean Gaussian with covariance matrix $\sigma_\eta^2 \mathbf{I}_2$. Equations (4) and (6) define the measurement model, which expresses the observed image coordinates, \mathbf{z} , of the point as a function of the *known* position vector, ${}^B\mathbf{p}$, the *unknown* transformation between the camera and base frame, $\{{}^C_B\mathbf{R}, {}^C\mathbf{p}_B\}$, and the *unknown* configuration of the mirror with respect to the camera $\{{}^C\mathbf{n}, d\}$. Note that the transformation between the mirror and camera frame has 6 dof, however, only 3 dof appear in the measurement equation. These are expressed by the vector $d {}^C\mathbf{n}$, which has 2 dof from the mirror normal, ${}^C\mathbf{n}$, and 1 dof from the camera-to-mirror distance, d . The remaining 3 dof, which correspond to rotations about ${}^C\mathbf{n}$ and to translations of the mirror-frame origin in the mirror plane, do not affect the measurements, and are unobservable.

3.2 Three Point Perspective Pose Estimation Problem

We now briefly review the three point perspective pose estimation problem (P3P) and discuss how it relates to our problem. The goal of P3P is to determine the 6 dof transformation, $\{{}^C_B\mathbf{R}, {}^C\mathbf{p}_B\}$, between a camera frame, $\{C\}$, and a base frame, $\{B\}$, given the known coordinates of 3 noncollinear points, ${}^B\mathbf{p}_i$, $i = 1 \dots 3$, in $\{B\}$, and their corresponding perspective projections, \mathbf{z}_i , $i = 1 \dots 3$, in $\{C\}$, defined as:²

² The indices in this paper are: i for points, j for images, and k for solutions.




Fig. 2. Observation of the point ${}^C\mathbf{p}'$ which is the reflection of ${}^C\mathbf{p}$. In this figure, the mirror plane is perpendicular to the page. Only the reflected point is in the camera's field of view; the real point is not observed directly by the camera.

$$\mathbf{z}_i = \frac{1}{p_{3i}} \begin{bmatrix} p_{1i} \\ p_{2i} \end{bmatrix}, \quad {}^C\mathbf{p}'_i = [p_{1i} \ p_{2i} \ p_{3i}]^T \quad (7)$$

$$\begin{bmatrix} {}^C\mathbf{p}'_i \\ 1 \end{bmatrix} = \begin{bmatrix} {}^C_B\mathbf{R} & {}^C\mathbf{p}_B \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} {}^B\mathbf{p}_i \\ 1 \end{bmatrix} \quad (8)$$

This problem has up to 4 pairs of solutions, where for each pair, there is one solution lying in front of the center of perspectivity and one behind it [5].

Equation (8) differs from (5) in that the former expresses a homogeneous transformation, while the latter describes a homogeneous transformation followed by a reflection. In other words, our scenario is equivalent to a P3P in which an “imaginary” camera with a left-handed reference frame lies behind the mirror and observes the true points (not the reflections). To bring (5) into a form similar to (8), we convert the imaginary camera to a right-handed system by premultiplying with a reflection across the y axis (although any axis can be chosen):

$$\begin{aligned} \begin{bmatrix} \check{C}\mathbf{p}' \\ 1 \end{bmatrix} &= \begin{bmatrix} (\mathbf{I}_3 - 2\mathbf{e}_2\mathbf{e}_2^T) & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} (\mathbf{I}_3 - 2{}^C\mathbf{n}{}^C\mathbf{n}^T) & 2d{}^C\mathbf{n} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} {}^C_B\mathbf{R} & {}^C\mathbf{p}_B \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} {}^B\mathbf{p} \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} \check{C}_B\mathbf{R} & \check{C}\mathbf{p}_B \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} {}^B\mathbf{p} \\ 1 \end{bmatrix} \end{aligned} \quad (9)$$

where $\mathbf{e}_2 = [0 \ 1 \ 0]^T$. The origin of $\{\check{C}\}$ coincides with that of $\{C\}$, their x and z axes are common, and their y axes lie in opposite directions. Note that this additional reflection can be implemented easily, by simply negating the sign of the y -coordinates of the image measurements.

Applying any method for solving the P3P problem to the modified one (cf. (9)), we obtain up to 4 solutions, in general, for the unknown transforma-

tion, $\{\check{C}_B \mathbf{R}, \check{C}_B \mathbf{p}_B\}$, between the body frame $\{B\}$ and the right-handed frame $\{\check{C}\}$ of the “imaginary” camera behind the mirror. In order to obtain the left-handed camera system, we “reflect back” each of the solutions:

$$\begin{bmatrix} {}^C \mathbf{p}' \\ 1 \end{bmatrix} = \begin{bmatrix} (\mathbf{I}_3 - 2\mathbf{e}_2\mathbf{e}_2^T) & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} \check{C}_B \mathbf{R} & \check{C}_B \mathbf{p}_B \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} {}^B \mathbf{p} \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} {}^B \mathbf{p} \\ 1 \end{bmatrix} \quad (10)$$

where the pair $\{\mathbf{A}, \mathbf{b}\}$, describes a reflection and a homogeneous transformation. Equating (5) and (10), we observe that:

$$\mathbf{A} = (\mathbf{I}_3 - 2^C \mathbf{n}^C \mathbf{n}^C{}^T) {}^C_B \mathbf{R} \quad (11)$$

$$\mathbf{b} = (\mathbf{I}_3 - 2^C \mathbf{n}^C \mathbf{n}^C{}^T) {}^C \mathbf{p}_B + 2d^C \mathbf{n} \quad (12)$$

To summarize, in order to exploit the similarity of our problem to the P3P, we execute the following steps: First, the y coordinates of the image measurements are negated. Then, the measurements are processed by a P3P algorithm to obtain $\{\check{C}_B \mathbf{R}, \check{C}_B \mathbf{p}_B\}$. Subsequently we employ (10), to obtain up to 4 solutions for \mathbf{A} and \mathbf{b} . In the next section, we describe our approach for recovering the unknowns, $\{{}^C_B \mathbf{R}, {}^C \mathbf{p}_B, {}^C \mathbf{n}, d\}$, from \mathbf{A} and \mathbf{b} using (11) and (12).

3.3 Solution from 3 points in 3 images

When only 2 points are observed, regardless of the number of images, there is not enough information to determine the transformation, since 3 noncollinear points are required to define the base frame of reference³. From 1 image with 3 points, there are not enough constraints to determine the unknowns (cf. (5) and (6)). As shown in [7], from 2 images with 3 points observed in each, the number of constraints equals the number of unknowns. Unfortunately in this case, rotations of the two mirror planes about the axis of their intersection are unobservable, and thus two images are not sufficient [7].

From 3 images with 3 points in each, there are 18 scalar measurements (cf. (6)) and 15 unknowns; 6 from $\{{}^C_B \mathbf{R}, {}^C \mathbf{p}_B\}$, and 3 for each mirror configuration $\{\mathbf{n}_j, d_j\}$ for $j = 1 \dots 3$ (cf. (5)). This is an overdetermined system that is nonlinear in the unknown variables. Using P3P as an intermediate step, and momentarily ignoring multiple solutions, we have constraints of the form (11), (12) for each image.

For each pair of images, $j, j' \in \{1 \dots 3\}$, let the unit vector $\mathbf{m}_{jj'}$, be perpendicular to both \mathbf{n}_j , and $\mathbf{n}_{j'}$ (i.e., $\mathbf{n}_j^T \mathbf{m}_{jj'} = \mathbf{n}_{j'}^T \mathbf{m}_{jj'} = 0$). Clearly, $\mathbf{m}_{jj'} = \alpha \mathbf{n}_j \times \mathbf{n}_{j'}$, where α is a scaling constant to ensure unit length. Then, the following relationship holds (cf. (11)):

$$\mathbf{A}_j \mathbf{A}_{j'}^T \mathbf{m}_{jj'} = (\mathbf{I}_3 - 2^C \mathbf{n}_j^C \mathbf{n}_j^C{}^T) (\mathbf{I}_3 - 2^C \mathbf{n}_{j'}^C \mathbf{n}_{j'}^C{}^T) \mathbf{m}_{jj'} = \mathbf{m}_{jj'} \quad (13)$$

³ In the case of 2 points, or 3 or more collinear points, rotations about the line that the points lie on will not be observable.

Thus, by solving for the eigenvector corresponding to the unit eigenvalue of $\mathbf{A}_j \mathbf{A}_j^T$, we determine $\mathbf{m}_{jj'}$ up to sign (it can be shown that $\mathbf{A}_j \mathbf{A}_j^T$ is a special orthogonal matrix with two complex conjugate eigenvalues, and one eigenvalue equal to one). Employing the properties of the cross product, we obtain⁴

$$\mathbf{n}_1 = \frac{\mathbf{m}_{13} \times \mathbf{m}_{12}}{\|\mathbf{m}_{13} \times \mathbf{m}_{12}\|}, \quad \mathbf{n}_2 = \frac{\mathbf{m}_{21} \times \mathbf{m}_{23}}{\|\mathbf{m}_{21} \times \mathbf{m}_{23}\|}, \quad \mathbf{n}_3 = \frac{\mathbf{m}_{13} \times \mathbf{m}_{23}}{\|\mathbf{m}_{13} \times \mathbf{m}_{23}\|} \quad (14)$$

Once we have determined the unit vectors corresponding to the 3 mirror planes, the rotation matrix, ${}^C_B \mathbf{R}$, can be computed independently from 3 sets of equations:

$${}^C_B \mathbf{R}_j = (\mathbf{I} - 2\mathbf{n}_j \mathbf{n}_j^T) \mathbf{A}_j, \quad j = 1 \dots 3 \quad (15)$$

In order to utilize all the available information, and to minimize numerical errors, we want to compute an ‘‘average’’ ${}^C_B \mathbf{R}$ from these 3 sets of equations. However, employing the arithmetic mean is inappropriate since the property of orthonormality is not maintained. We address this issue with the procedure described in Appendix A.

Once the rotation, ${}^C_B \mathbf{R}$, and the mirror normal vectors, \mathbf{n}_j , $j = 1 \dots 3$, are determined, the remaining unknowns $\{{}^C \mathbf{p}_B, d_1, d_2, d_3\}$ appear linearly in the constraint equations (cf. (12))

$$\begin{bmatrix} (\mathbf{I} - 2\mathbf{n}_1 \mathbf{n}_1^T) & 2\mathbf{n}_1 & \mathbf{0} & \mathbf{0} \\ (\mathbf{I} - 2\mathbf{n}_2 \mathbf{n}_2^T) & \mathbf{0} & 2\mathbf{n}_2 & \mathbf{0} \\ (\mathbf{I} - 2\mathbf{n}_3 \mathbf{n}_3^T) & \mathbf{0} & \mathbf{0} & 2\mathbf{n}_3 \end{bmatrix} \begin{bmatrix} {}^C \mathbf{p}_B \\ d_1 \\ d_2 \\ d_3 \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \mathbf{b}_3 \end{bmatrix} \Leftrightarrow \mathbf{D} \mathbf{x} = \mathbf{c} \quad (16)$$

where \mathbf{D} is a 9×6 known matrix, \mathbf{c} is a 9×1 known vector, and \mathbf{x} is the 6×1 vector of unknowns. The least-squares solution for \mathbf{x} in this linear system is $\mathbf{x} = \mathbf{D}^\dagger \mathbf{c}$, where \mathbf{D}^\dagger denotes the Moore-Penrose generalized inverse of \mathbf{D} . From (14), (15), (23), and (16) the mirror configurations as well as the camera-to-base transformation are computed.

Up to this point, we assumed that the P3P solution was unique for each image. In general, however, there may be up to 4 solutions per image. Recall that 3 images are required to compute the camera-to-base transformation analytically, hence, there are up to 64 solutions for $\{{}^C_B \mathbf{R}, {}^C \mathbf{p}_B, d_1, d_2, d_3, \mathbf{n}_1, \mathbf{n}_2, \mathbf{n}_3\}$, arising from the $4 \times 4 \times 4$ possible combinations of P3P solutions. If the measurements were noiseless, we would expect that only one of these solutions would yield zero reprojection error (i.e., would satisfy all constraints exactly). This is because the problem at hand is over-constrained (18 constraints for 15 unknowns), and we expect to have a unique solution. In the presence of pixel noise, none of the solutions will satisfy the measurements exactly, thus, we choose the solution with the minimum reprojection error.

Moreover, when N_c images are available, there are $N_s = \binom{N_c}{3}$ analytically computed transformations. However, some of these may be inaccurate as a

⁴ For the remainder of the paper, we drop the superscript ‘ C ’ from \mathbf{n}_j , $j = 1 \dots 3$.

result of degenerate sets of measurements (e.g., when 3 images are all taken from similar viewing angles). In order to identify the correct solution, we employ spectral clustering to determine the largest set of similar solutions [18]. Specifically, we adopt the unit-quaternion representation of rotation [21], ${}^C\bar{q}_B$, which corresponds to ${}^C_B\mathbf{R}$, and denote each solution as $\{{}^C\bar{q}_B^{(k)}, {}^C\mathbf{p}_B^{(k)}\}$ for $k = 1 \dots N_s$. To carry out spectral clustering, we define an affinity matrix, \mathbf{L} , in which each element is the Mahalanobis distance between a pair of solutions, indexed by k and k' :

$$\mathbf{L}_{kk'} = [\delta\boldsymbol{\theta}_{kk'}^T \ \delta\mathbf{p}_{kk'}^T] \left[(\mathbf{H}_k^T \boldsymbol{\Gamma}^{-1} \mathbf{H}_k)^{-1} + (\mathbf{H}_{k'}^T \boldsymbol{\Gamma}^{-1} \mathbf{H}_{k'})^{-1} \right]^{-1} \begin{bmatrix} \delta\boldsymbol{\theta}_{kk'} \\ \delta\mathbf{p}_{kk'} \end{bmatrix} \quad (17)$$

where $\delta\boldsymbol{\theta}_{kk'}$ is the quaternion error-angle vector between ${}^C\bar{q}_B^{(k)}$ and ${}^C\bar{q}_B^{(k')}$ [21], and $\delta\mathbf{p}_{kk'} = {}^C\mathbf{p}_B^{(k)} - {}^C\mathbf{p}_B^{(k')}$ is the difference between the translation vectors. The matrices \mathbf{H}_k and $\mathbf{H}_{k'}$ are the measurement Jacobians with respect to the transformation [7], and $\boldsymbol{\Gamma} = \sigma_\eta^2 \mathbf{I}_2$ is the covariance of the pixel noise. For brevity, we do not discuss spectral clustering further here, but refer the reader to [18] for details. We compute the transformation, $\{{}^C_B\mathbf{R}, {}^C\mathbf{p}_B\}$, from the largest spectral cluster. The rotation, ${}^C_B\mathbf{R}$, is determined from (23) using all the quaternions in the cluster (cf. Appendix A), and the translation, ${}^C\mathbf{p}_B$, is computed as the arithmetic mean of the translations in the cluster.

3.4 Refining the Solution

Due to the presence of pixel noise, and the fact that noise was not accounted for in the analytical solution, the result of the procedure presented in Sections 3.2-3.3 may be coarse (cf. Section 4). Hence, we employ an MLE to refine our analytically computed estimate. We now proceed with an overview of the MLE for determining the unknown transformation between the camera and base frame. Let the vector of all unknown parameters be denoted by \mathbf{x} . This vector comprises the unknown transformation, as well as $\{{}^C\mathbf{n}_j, d_j\}$, $j = 1 \dots N_c$, that describe each mirror configuration:

$$\mathbf{x} = [{}^C\mathbf{p}_B^T \ {}^C\bar{q}_B^T \ {}^C\mathbf{n}_1^T \ d_1 \ \dots \ {}^C\mathbf{n}_{N_c}^T \ d_{N_c}]^T \quad (18)$$

Under the assumption of Gaussian pixel noise, the likelihood of the measurements is given by:

$$\begin{aligned} L(\mathcal{Z}; \mathbf{x}) &= \prod_{i=1}^{N_p} \prod_{j=1}^{N_c} p(\mathbf{z}_{ij}; \mathbf{x}) = \prod_{i=1}^{N_p} \prod_{j=1}^{N_c} \frac{1}{2\pi\sigma_\eta^2} \exp \left[-\frac{(\mathbf{z}_{ij} - \mathbf{h}({}^C_j\mathbf{p}'_i))^T (\mathbf{z}_{ij} - \mathbf{h}({}^C_j\mathbf{p}'_i))}{2\sigma_\eta^2} \right] \\ &= \prod_{i=1}^{N_p} \prod_{j=1}^{N_c} \frac{1}{2\pi\sigma_\eta^2} \exp \left[-\frac{(\mathbf{z}_{ij} - \mathbf{h}_{ij}(\mathbf{x}))^T (\mathbf{z}_{ij} - \mathbf{h}_{ij}(\mathbf{x}))}{2\sigma_\eta^2} \right] \end{aligned}$$

where the dependence on \mathbf{x} is explicitly shown (cf. (5), (6)), and N_p is the total number of points observed in each of the N_c images. Maximizing the likelihood is equivalent to minimizing its negative logarithm, or minimizing the cost function:

$$J(\mathbf{x}) = \sum_{i=1}^{N_p} \sum_{j=1}^{N_c} (\mathbf{z}_{ij} - \mathbf{h}_{ij}(\mathbf{x}))^T (\mathbf{z}_{ij} - \mathbf{h}_{ij}(\mathbf{x})) \quad (19)$$

This is a nonlinear least-squares problem, and thus we employ Gauss-Newton iterative minimization for estimating \mathbf{x} . At each iteration, indexed by ℓ , the estimate is changed by:

$$\delta \mathbf{x}^{(\ell)} = \left(\sum_{i,j} \mathbf{H}_{ij}^{(\ell)T} \mathbf{H}_{ij}^{(\ell)} \right)^{-1} \left(\sum_{i,j} \mathbf{H}_{ij}^{(\ell)T} (\mathbf{z}_{ij} - \mathbf{h}_{ij}(\mathbf{x}^{(\ell)})) \right)$$

where $\mathbf{H}_{ij}^{(\ell)}$ is the Jacobian of the measurement \mathbf{z}_{ij} with respect to \mathbf{x} , evaluated at the current iterate, $\mathbf{x}^{(\ell)}$. Note that the analytically computed solution from Sections 3.2-3.3 is utilized as the initial iterate, $\mathbf{x}^{(0)}$, of the MLE. Since the MLE is not a main contribution of this work, we limit our discussion here, but refer the reader to [8] for more details.

4 Simulations

In this section, we study the accuracy of the analytically computed camera-to-base transformation (cf. Sections 3.2-3.3). In particular, we investigate how the accuracy is affected by the following parameters: (i) pixel noise, (ii) number of images, (iii) distance from camera to mirror, and (iv) range of the mirror’s angular motion. We consider a “standard” case, in which 3 points placed at the corners of a right triangle with sides measuring $20 \times 20 \times 20\sqrt{2}$ cm are observed in 200 images, while a mirror placed at a distance of 0.5 m is rotated by 30° in two directions. Then, we vary each of the aforementioned parameters individually, in order to examine its effect on the solution accuracy. In Fig. 3, we plot the average RMS error for the position and attitude. These averages are computed over 10 trials. Some key observations are the following:

- Increasing the camera’s pixel noise decreases the accuracy of the computed solution. In the extreme case when the camera measurements become substantially noisy, e.g., $\sigma = 2$ pixels, the average RMS error is 1° in attitude and 15 cm in position.
- As expected, increasing the number of images results in higher accuracy. However, the improvement follows a “law of diminishing return,” i.e., when a large number of images is already available, the impact of recording more observations is smaller.

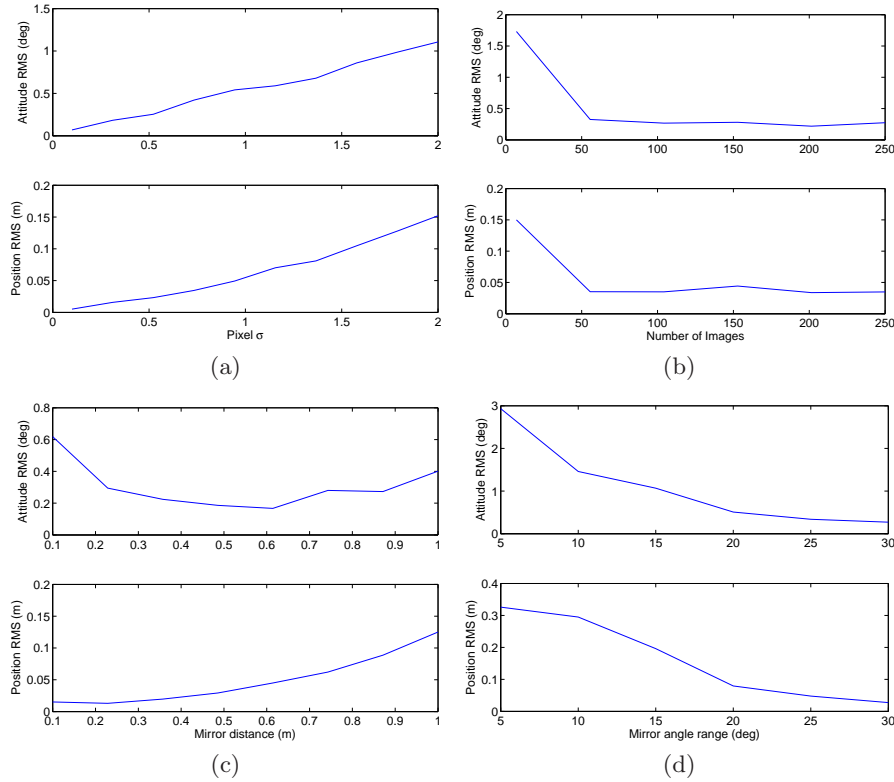


Fig. 3. Average RMS error over 10 trials for attitude and position plotted versus: (a) pixel noise, (b) number of images, (c) mirror distance, and (d) range of mirror rotation.

- Changing the distance from the mirror to the camera has a significant effect on the position accuracy. When the mirror is at a distance of 1 m, the average RMS error for position is approximately 13 cm. The magnitude of this error suggests that the mirror distance should be kept small. Additionally, it highlights the need to refine our analytically computed transformation with an MLE. As we show in [8], the accuracy of the MLE is approximately 5 times better in attitude, and 10 times better in position compared to the analytical solution.
- Increasing the range of the mirror’s angular motion results in improved accuracy. The effect on the attained accuracy is significant, and thus every effort should be made to move the mirror in the widest range of motion allowed by the camera’s field of view.

As a final remark, we note that using the analytical solution as an initial guess for the MLE helps the latter converge to the correct minimum 100% of the time. Moreover, the number of iterations required by the MLE was 7




Fig. 4. (a) Observation of a point on the robot reflected in the mirror, (b) Image with 3 fiducial points, captured during experimentation.

(compared to 18 when using a naive initial guess), on average, which shows that the availability of a precise analytical solution improves the speed and robustness of the overall estimation process.

5 Experiments

The method described in the preceding sections was employed for computing the transformation between a camera and a base frame attached on the robot-body. For this purpose, 3 fiducial points were placed in known positions on the robot as shown in Fig. 4(b). The origin of frame $\{B\}$ coincides with the top-left fiducial point, while both frames $\{B\}$ and $\{C\}$ are right-handed systems with the axes of $\{B\}$ approximately aligned with those of $\{C\}$. These fiducial points were tracked using the KLT algorithm [20] in 1000 images, recorded by a Firewire camera with resolution of 1024×768 pixels.

A planar mirror was maneuvered in different spatial configurations, and in distances varying between 30 and 50 cm from the camera, in order to generate a wide range of views. All the measurements were processed to compute the transformation analytically: ${}^C\mathbf{p}_B = [-14.13 \ -10.25 \ -13.89]^T$ cm, and ${}^C\mathbf{q}_B = [-0.0401 \ -0.0017 \ -0.0145 \ 0.9991]^T$. This initial solution was refined using the MLE described in Section 3.4, to obtain a better estimate for the transformation between the two frames of interest. The Gauss-Newton minimization converged after 8 iterations, to the following solution for the transformation: ${}^C\mathbf{p}_B = [-14.80 \ -15.96 \ -14.95]^T$ cm, and ${}^C\mathbf{q}_B = [0.0045 \ 0.0774 \ 0.0389 \ 0.9962]^T$. The corresponding 3σ accuracy bounds are $[1.1 \ 1.6 \ 5.0]$ mm for the position, and $[0.2419 \ 0.2313 \ 0.0665]$ degrees for the orientation estimates. We point out that the estimates agree with our best guess from manual measurement. We believe that the attained accuracy (given

by the 3σ bounds from the MLE) is sufficiently high for most practical applications.

6 Conclusions and Future Work

In this paper, we propose a method for computing the 6 dof transformation between a camera and a base frame of reference. A mirror is maneuvered in front of the camera, to provide observations of known points from different viewing angles and distances. These measurements are utilized to analytically compute the camera-to-base transformation, and the solution is refined using a maximum-likelihood estimator, which produces estimates for the camera-to-base transformation, as well as for the mirror configuration in each image. The approach was validated both in simulation and using real-world experimentation. One of the key advantages of the proposed method is its ease of use; it only requires a mirror, and it provides a solution with as little as 3 points viewed in 3 images. When more information is available, it can be incorporated to produce a more accurate estimate of the transformation.

In our future work, we will investigate the feasibility of mirror-based robot-body cartography (briefly discussed in Appendix B). Furthermore, we plan to extend this method to the case where the coordinates of the points in the base frame are not known *a priori*, but are estimated along with the camera-to-base transformation and the mirror configurations.

Appendix A

In this section, we describe the procedure we employ for computing an “average rotation,” given N_q rotation estimates $\bar{q}_j, j = 1 \dots N_q$. We adopt the quaternion notation from [21] and denote the quaternion of rotation arising from the j th set of equations as \bar{q}_j , which corresponds to ${}^C_B\mathbf{R}_j$ (cf. (15)). Assuming that \bar{q} is the optimal estimate, and employing the small error-angle approximation, we write the following expression for the error in each of the \bar{q}_j s:

$$\bar{q}_j \otimes \bar{q}^{-1} \simeq \begin{bmatrix} \mathbf{k}_j \delta\theta_j \\ 1 \end{bmatrix}, \quad j = 1 \dots N_q \quad (20)$$

where \otimes denotes quaternion multiplication, \mathbf{k}_j is the unit-vector axis of rotation, and $\delta\theta_j$ is the error angle between the two quaternions. Rewriting this last expression as a matrix-vector multiplication [21], yields

$$\mathcal{L}(\bar{q}_j) \bar{q}^{-1} = \begin{bmatrix} \mathbf{k}_j \delta\theta_j \\ 1 \end{bmatrix}, \quad j = 1 \dots N_q \quad (21)$$

where $\mathcal{L}(\bar{q}_j)$, is the left-side quaternion multiplication matrix parameterized by \bar{q}_j . Projecting this relation, to keep only the error components, we obtain:

$$\mathbf{P}\mathcal{L}(\bar{q}_j)\bar{q}^{-1} = \mathbf{k}_j\delta\theta_j, \quad j = 1 \dots N_q \quad (22)$$

where $\mathbf{P} = [\mathbf{I}_3 \mathbf{0}_{3 \times 1}]$. Stacking these relations, we have

$$\begin{bmatrix} \mathbf{P}\mathcal{L}(\bar{q}_1) \\ \vdots \\ \mathbf{P}\mathcal{L}(\bar{q}_{N_q}) \end{bmatrix} \bar{q}^{-1} = \begin{bmatrix} \mathbf{k}_1\delta\theta_1 \\ \vdots \\ \mathbf{k}_{N_q}\delta\theta_{N_q} \end{bmatrix} \quad (23)$$

Our goal is to find the \bar{q}^{-1} that minimizes the norm of the right-hand side. This occurs when $\bar{q}^{-1} = \mathbf{v}(\sigma_{min})$, i.e., we select \bar{q}^{-1} to be the right singular vector corresponding to the minimum singular value of the $3N_q \times 4$ matrix multiplying \bar{q}^{-1} in (23). After finding \bar{q}^{-1} by SVD, we compute the optimal estimate for the rotational matrix ${}^C_B\mathbf{R} = \mathbf{R}(\bar{q})$, which is the rotational matrix parameterized by the quaternion \bar{q} .

Appendix B

We hereafter address the problem of mirror-based robot-body cartography. That is, building a 3D map of the robot body utilizing mirror reflections. We assume that in addition to the 3 points which are known in the robot body frame, we observe another point, \mathbf{p}_u , which is unknown in the body frame. From one image, we have (cf. (4)):

$$s^C \mathbf{p}'_{u_0} = (\mathbf{I}_3 - 2^C \mathbf{n} \mathbf{n}^T) {}^C_B \mathbf{R}^B \mathbf{p}_u + (\mathbf{I}_3 - 2^C \mathbf{n} \mathbf{n}^T) {}^C \mathbf{p}_B + 2d^C \mathbf{n} \quad (24)$$

where s is an unknown scale factor and ${}^C \mathbf{p}'_{u_0}$ is the unit vector along the direction of ${}^C \mathbf{p}'_u$. Premultiplying both sides by the reflection matrix yields

$$s (\mathbf{I}_3 - 2^C \mathbf{n} \mathbf{n}^T) {}^C \mathbf{p}'_{u_0} = {}^C_B \mathbf{R}^B \mathbf{p}_u + {}^C \mathbf{p}_B - 2d^C \mathbf{n} \quad (25)$$

We assume that the transformation from $\{B\}$ to $\{C\}$, as well as the mirror configuration have been determined using the method outlined in this paper. Hence, the quantities $\{{}^C_B \mathbf{R}, {}^C \mathbf{p}_B, d, {}^C \mathbf{n}\}$ are known and ${}^C \mathbf{p}'_{u_0}$ is measured (through the image reflection), while the quantities $\{s, {}^B \mathbf{p}_u\}$ are unknown. From a single image, there are 3 constraints (cf. (25)) and 4 unknowns; hence, we can constrain ${}^B \mathbf{p}_u$ to lie on a line parameterized by s . If the point is observed in 2 consecutive images, then we will have 6 constraints and 5 unknowns, 3 corresponding to the unknown point's coordinates and 2 to the unknown scale factors. In this case, we expect that ${}^B \mathbf{p}_u$ can be determined uniquely.

In fact, this problem is analogous to “triangulation” of a point from two image views ([6], ch. 12). It is solvable when the origin of the camera frame

is different for the two views. In our case, this corresponds to the quantity $d^C \mathbf{n}$ changing. Thus, it suffices to either change the distance to the mirror, or the mirror's orientation with respect to the camera. We expect that the location of every unknown point on the robot-body, which is visible in the mirror reflections, can be determined in the body frame of reference, given that it can be reliably tracked in at least 2 images taken from different views.

Acknowledgments

This work was supported by the University of Minnesota (DTC), and the National Science Foundation (EIA-0324864, IIS-0643680, IIS-0811946). Anastasios Mourikis was supported by the UMN Doctoral Dissertation Fellowship. Joel Hesch was supported by the NIH Neuro-physical-computational Sciences Fellowship.

References

1. X. Brun and F. Goulette. Modeling and calibration of coupled fish-eye CCD camera and laser range scanner for outdoor environment reconstruction. In *Proc. of the International Conference on 3D Digital Imaging and Modeling*, pages 320–327, Montreal, QC, Canada, Aug. 21–23, 2007.
2. K. Daniilidis. Hand-eye calibration using dual quaternions. *International Journal of Robotics Research*, 18(3):286–298, 1999.
3. J. Gluckman and S. K. Nayar. Planar catadioptric stereo: Geometry and calibration. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 22–28, Ft. Collins, CO, June 23–25, 1999.
4. A. A. Goshtasby and W. A. Gruver. Design of a single-lens stereo camera system. *Pattern Recognition*, 26(6):923–937, June 1993.
5. R. M. Haralick, C.-N. Lee, K. Ottenberg, and M. Nölle. Review and analysis of solutions of the three point perspective pose estimation problem. *International Journal of Computer Vision*, 13(3):331–356, 1994.
6. R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.
7. J. A. Hesch, A. I. Mourikis, and S. I. Roumeliotis. Camera to robot-body calibration using planar mirror reflections. Technical Report 2008-001, University of Minnesota, Dept. of Comp. Sci. & Eng., MARS Lab, July 2008.
8. J. A. Hesch, A. I. Mourikis, and S. I. Roumeliotis. Determining the camera to robot-body transformation from planar mirror reflections. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3865–3871, Nice, France, Sept. 22–26, 2008.
9. M. Inaba, T. Hara, and H. Inoue. A stereo viewer based on a single camera with view-control mechanisms. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1857–1865, Yokohama, Japan, July 26–30, 1993.

10. G. Jang, S. Kim, and I. Kweon. Single camera catadioptric stereo system. In *Proc. of the 6th Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras (OMNIVIS/ICCV)*, Beijing, China, Oct. 21, 2005.
11. K. H. Jang, D. H. Lee, and S. K. Jung. A moving planar mirror based approach for cultural reconstruction. *Computer Animation and Virtual Worlds*, 15(3-4):415–423, July 2004.
12. M. Kanbara, N. Ukita, M. Kidode, and N. Yokoya. 3D scene reconstruction from reflection images in a spherical mirror. In *Proc. of the IEEE 18th International Conference on Pattern Recognition*, pages 874–897, Hong Kong, China, Aug. 20–24, 2006.
13. R. K. Kumar, A. Ilie, J.-M. Frahm, and M. Pollefeys. Simple calibration of non-overlapping cameras with a mirror. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Anchorage, AK, June 2008.
14. A. Martinelli and R. Siegwart. Observability properties and optimal trajectories for on-line odometry self-calibration. In *Proc. of the IEEE Conference on Decision and Control*, pages 3065–3070, San Diego, CA, Dec. 13–15, 2006.
15. F. M. Mirzaei and S. I. Roumeliotis. A Kalman filter-based algorithm for IMU-camera calibration. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2427–2434, San Diego, CA, Oct. 29–Nov. 2, 2007.
16. F. M. Mirzaei and S. I. Roumeliotis. A Kalman filter-based algorithm for IMU-camera calibration: Observability analysis and performance evaluation. *IEEE Transactions on Robotics*, 2008. (in press).
17. S. K. Nayar. Sphero: Determining depth using two specular spheres and a single camera. In *Proc. of the SPIE conference on Optics, Illumination, and Image Sensing for Machine Vision*, volume 1005, pages 245–254, Nov. 1988.
18. A. Y. Ng, M. I. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *Neural Information Processing Systems Conference*, volume 2, pages 849–856, British Columbia, Canada, Dec. 3–8, 2002.
19. B. K. Ramsgaard, I. Balslev, and J. Arnsparng. Mirror-based trinocular systems in robot-vision. In *Proc. of the IEEE International Conference on Pattern Recognition*, pages 499–502, Barcelona, Spain, Sept. 3–7, 2000.
20. J. Shi and C. Tomasi. Good features to track. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 593–600, Washington, DC, June 27–July 2, 1994.
21. N. Trawny and S. I. Roumeliotis. Indirect Kalman filter for 3D attitude estimation. Technical Report 2005-002, University of Minnesota, Dept. of Comp. Sci. & Eng., MARS Lab, Mar. 2005.
22. R. Y. Tsai and R. K. Lenz. A new technique for fully autonomous and efficient 3D robotics hand/eye calibration. *IEEE Transactions on Robotics and Automation*, 5(3):345–358, June 1989.
23. S. Wasielewski and O. Strauss. Calibration of a multi-sensor system laser rangefinder/camera. In *Proc. of the Intelligent Vehicles Symposium*, pages 472–477, Detroit, MI, Sept. 25–26, 1995.
24. A. Würz-Wessel and F. K. Stein. Calibration of a free-form surface mirror in a stereo vision system. In *Proc. of the IEEE Intelligent Vehicle Symposium*, volume 2, pages 471–476, Versailles, France, June 17–21, 2002.