

Computer Vision

A Modern Approach

David A. Forsyth

University of California at Berkeley

Jean Ponce

University of Illinois at Urbana-Champaign

=====*An Alan R. Apt Book*=====



Prentice Hall
Upper Saddle River, New Jersey 07458

Geometric Camera Calibration

This chapter addresses the problem of estimating the intrinsic and extrinsic parameters of a camera, a process known as *geometric camera calibration*. We assume throughout that the camera observes a set of features such as points or lines with known positions in some fixed world coordinate system (Figure 3.1): In this context, camera calibration can be modeled as an optimization process, where the discrepancy between the observed image features and their theoretical positions (as predicted by the perspective projection equations derived in chapter 2) is minimized with respect to the camera's intrinsic and extrinsic parameters.




Figure 3.1 Camera calibration setup: In this example, the calibration rig is formed by three grids drawn in orthogonal planes. Other patterns could be used as well, and they may involve lines or other geometric figures.

We start with an overview of *least-squares* techniques aimed at solving this type of optimization problems before presenting several linear and nonlinear approaches to calibration. Once a camera has been calibrated, it is possible to associate with any image point a well-defined ray passing through this point and the camera's optical center as well as perform accurate three-dimensional measurements from digitized pictures. An application to mobile robot localization is briefly discussed at the end of the chapter.

3.1 LEAST-SQUARES PARAMETER ESTIMATION

As already mentioned, calibrating a camera amounts to estimating the intrinsic and extrinsic parameters that minimize the mean-squared deviation from predicted to observed image features. This section introduces a class of optimization techniques, known as least-squares methods, for solving this kind of problem. They prove useful on several other occasions in the rest of this book.

3.1.1 Linear Least-Squares Methods

Let us first consider a system of p linear equations in q unknowns:

$$\begin{cases} u_{11}x_1 + u_{12}x_2 + \cdots + u_{1q}x_q = y_1 \\ u_{21}x_1 + u_{22}x_2 + \cdots + u_{2q}x_q = y_2 \\ \cdots \\ u_{p1}x_1 + u_{p2}x_2 + \cdots + u_{pq}x_q = y_p \end{cases} \iff \mathcal{U}\mathbf{x} = \mathbf{y}. \quad (3.1)$$

In this equation,

$$\mathcal{U} = \begin{pmatrix} u_{11} & u_{12} & \cdots & u_{1q} \\ u_{21} & u_{22} & \cdots & u_{2q} \\ \cdots & \cdots & \cdots & \cdots \\ u_{p1} & u_{p2} & \cdots & u_{pq} \end{pmatrix}, \quad \mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \cdots \\ x_q \end{pmatrix} \quad \text{and} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ \cdots \\ y_p \end{pmatrix}.$$

We know from linear algebra that (in general)

- when $p < q$, the set of solutions to this equation forms a $(q - p)$ -dimensional vector subspace of \mathbb{R}^q ;
- when $p = q$, there is a unique solution;
- when $p > q$, there is no solution.

This statement is true when the *rank* (i.e., the maximum number of independent rows or columns) of \mathcal{U} is maximal—that is, equal to $\min(p, q)$ (this is what we mean by *in general*). When the rank is smaller than $\min(p, q)$, the existence of solutions to Eq. (3.1) depends on the value of \mathbf{y} and whether it belongs to the *range* of \mathcal{U} (i.e., the subspace of \mathbb{R}^p spanned by its columns).

Normal Equations and the Pseudoinverse The rest of this section focuses on the overconstrained case $p > q$ and assumes that \mathcal{U} has maximal rank q . Since there is no exact solution in this case, we content ourselves with finding the vector \mathbf{x} that minimizes the error measure

$$E \stackrel{\text{def}}{=} \sum_{i=1}^p (u_{i1}x_1 + \cdots + u_{iq}x_q - y_i)^2 = |\mathcal{U}\mathbf{x} - \mathbf{y}|^2.$$

E is proportional to the mean-squared error associated with the equations, hence the name of least-squares methods given to techniques for minimizing it.

Now, we can write $\dot{E} = \mathbf{e} \cdot \mathbf{e}$, where $\mathbf{e} \stackrel{\text{def}}{=} \mathcal{U}\mathbf{x} - \mathbf{y}$. To find the vector \mathbf{x} minimizing E , we write that the derivatives of this error measure with respect to the coordinates x_i ($i = 1, \dots, q$) of \mathbf{x} must be zero—that is,

$$\frac{\partial E}{\partial x_i} = 2 \frac{\partial \mathbf{e}}{\partial x_i} \cdot \mathbf{e} = 0 \quad \text{for } i = 1, \dots, q.$$

But if the columns of \mathcal{U} are the vectors $\mathbf{c}_j = (u_{1j}, \dots, u_{mj})^T$ ($j = 1, \dots, q$), we have

$$\frac{\partial \mathbf{e}}{\partial x_i} = \frac{\partial}{\partial x_i} \left[\begin{pmatrix} \mathbf{c}_1 & \cdots & \mathbf{c}_q \end{pmatrix} \begin{pmatrix} x_1 \\ \cdots \\ x_q \end{pmatrix} - \mathbf{y} \right] = \frac{\partial}{\partial x_i} (x_i \mathbf{c}_1 + \cdots + x_q \mathbf{c}_q - \mathbf{y}) = \mathbf{c}_i.$$

In particular, writing that $\partial E / \partial x_i = 0$ implies that $\mathbf{c}_i^T (\mathcal{U}\mathbf{x} - \mathbf{y}) = 0$, and stacking the constraints associated with the q coordinates of \mathbf{x} yields the *normal equations* associated with our least-squares problem—that is,

$$\mathbf{0} = \begin{pmatrix} \mathbf{c}_1^T \\ \cdots \\ \mathbf{c}_q^T \end{pmatrix} (\mathcal{U}\mathbf{x} - \mathbf{y}) = \mathcal{U}^T (\mathcal{U}\mathbf{x} - \mathbf{y}) \iff \mathcal{U}^T \mathcal{U}\mathbf{x} = \mathcal{U}^T \mathbf{y}.$$

When \mathcal{U} has maximal rank q , the matrix $\mathcal{U}^T \mathcal{U}$ is easily shown to be invertible, and the solution of the normal equations is $\mathbf{x} = \mathcal{U}^\dagger \mathbf{y}$ with $\mathcal{U}^\dagger \stackrel{\text{def}}{=} [(\mathcal{U}^T \mathcal{U})^{-1} \mathcal{U}^T]$. The $q \times q$ matrix \mathcal{U}^\dagger is called the *pseudoinverse* of \mathcal{U} . It coincides with \mathcal{U}^{-1} when the matrix \mathcal{U} is square and nonsingular. Linear least-squares problems can be solved without explicitly computing the pseudoinverse, using, for example, QR decomposition or singular value decomposition (more on the latter in chapter 12), which are known to be better behaved numerically.

Homogeneous Systems and Eigenvalue Problems Let us now consider a variant of our original problem, where we have again a system of p linear equations in q unknowns, but the vector \mathbf{y} is zero—that is,

$$\begin{cases} u_{11}x_1 + u_{12}x_2 + \cdots + u_{1q}x_q = 0 \\ u_{21}x_1 + u_{22}x_2 + \cdots + u_{2q}x_q = 0 \\ \cdots \\ u_{p1}x_1 + u_{p2}x_2 + \cdots + u_{pq}x_q = 0 \end{cases} \iff \mathcal{U}\mathbf{x} = \mathbf{0}. \quad (3.2)$$

This is a *homogeneous* equation in \mathbf{x} (i.e., if \mathbf{x} is a solution, so is $\lambda \mathbf{x}$ for any $\lambda \neq 0$). When $p = q$ and the matrix \mathcal{U} is nonsingular, Eq. (3.2) admits as a unique solution $\mathbf{x} = \mathbf{0}$. Conversely, when $p \geq q$, nontrivial (i.e., nonzero) solutions may only exist when \mathcal{U} is singular with rank strictly smaller than q . In this context, minimizing $E = |\mathcal{U}\mathbf{x}|^2$ only makes sense when some additional constraint is imposed on \mathbf{x} since the value $\mathbf{x} = \mathbf{0}$ yields the zero global minimum of E . By homogeneity, we have $E(\lambda \mathbf{x}) = \lambda^2 E(\mathbf{x})$, and it is reasonable to choose the constraint $|\mathbf{x}|^2 = 1$, which avoids the trivial solution and forces the uniqueness of the result.

The error E can be rewritten as $|\mathcal{U}\mathbf{x}|^2 = \mathbf{x}^T (\mathcal{U}^T \mathcal{U}) \mathbf{x}$. The $q \times q$ matrix $\mathcal{U}^T \mathcal{U}$ is by construction symmetric positive semidefinite (i.e., its eigenvalues are all positive or zero), and it can be diagonalized in an orthonormal basis of eigenvectors \mathbf{e}_i ($i = 1, \dots, q$) associated with the eigenvalues $0 \leq \lambda_1 \leq \cdots \leq \lambda_q$. Thus we can write any unit vector as $\mathbf{x} = \mu_1 \mathbf{e}_1 + \cdots + \mu_q \mathbf{e}_q$ with $\mu_1^2 + \cdots + \mu_q^2 = 1$. In particular,

$$\begin{aligned} E(\mathbf{x}) - E(\mathbf{e}_1) &= \mathbf{x}^T (\mathcal{U}^T \mathcal{U}) \mathbf{x} - \mathbf{e}_1^T (\mathcal{U}^T \mathcal{U}) \mathbf{e}_1 = \lambda_1^2 \mu_1^2 + \cdots + \lambda_q^2 \mu_q^2 - \lambda_1^2 \\ &\geq \lambda_1^2 (\mu_1^2 + \cdots + \mu_q^2 - 1) = 0. \end{aligned}$$

It follows that the unit vector \mathbf{x} minimizing E is the eigenvector \mathbf{e}_1 associated with the minimum eigenvalue of $\mathcal{U}^T \mathcal{U}$, and the corresponding minimum value of E is λ_1^2 . Various methods are available for computing the eigenvectors and eigenvalues of a symmetric matrix, including Jacobi transformations and reduction to tridiagonal form followed by QR decomposition. Singular value decomposition can also be used to compute the eigenvectors and eigenvalues without actually constructing the matrix $\mathcal{U}^T \mathcal{U}$.

Before illustrating the use of homogeneous linear least-squares techniques with an example, let us pause for a minute to consider the slightly more general problem of minimizing $|\mathcal{U}\mathbf{x}|^2$ under the constraint $|\mathcal{V}\mathbf{x}|^2 = 1$, where \mathcal{V} is an $r \times q$ matrix (this reduces to homogeneous linear least squares when $\mathcal{V} = \text{Id}$). A vector \mathbf{x} and a scalar λ such that

$$\mathcal{U}^T \mathcal{U} \mathbf{x} = \lambda \mathcal{V}^T \mathcal{V} \mathbf{x}$$

are called a *generalized eigenvector* and the corresponding *generalized eigenvalue* of the $q \times q$ symmetric matrices $\mathcal{U}^T \mathcal{U}$ and $\mathcal{V}^T \mathcal{V}$. As shown in the exercises, the solution of our constrained optimization problem is precisely the unit generalized eigenvector associated with the minimum generalized eigenvalue (which is in this case guaranteed to be positive or zero by construction). As before, effective methods for computing the generalized eigenvectors and eigenvalues of a pair of symmetric matrices are available.

Example 3.1 Fitting a line to points in a plane.

Consider n points p_i ($i = 1, \dots, n$) in a plane, with coordinates (x_i, y_i) in some fixed coordinate system (Figure 3.2). What is the straight line that best fits these points? To answer this question, we must first quantify how well a line δ fits a set of points or, equivalently, define some error function E measuring the discrepancy between this line and the points. The best-fitting line can then be found by minimizing E .

A reasonable choice for the error function is the mean-squared distance between the points and the line (Figure 3.2). We saw in chapter 2 that the equation of a line with unit normal $\mathbf{n} = (a, b)^T$ lying at a distance d from the origin is $ax + by = d$. It is in fact easy to show that the perpendicular distance between a point with coordinates $(x, y)^T$ and this line is $|ax + by - d|$. We can therefore




Figure 3.2 The line that best fits n points in the plane can be defined as the line δ that minimizes the mean-squared perpendicular distance to these points (i.e., in this diagram, the mean-squared length of the short parallel line segments joining δ to the points).

use

$$E(a, b, d) = \sum_{i=1}^n (ax_i + by_i - d)^2$$

as our error measure, and the line-fitting problem reduces to the minimization of E with respect to a , b , and d under the constraint $a^2 + b^2 = 1$. Differentiating E with respect to d shows that, at a minimum of this function, we must have $0 = \partial E / \partial d = -2 \sum_{i=1}^n (ax_i + by_i - d)$, thus

$$d = a\bar{x} + b\bar{y}, \quad \text{where} \quad \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad \text{and} \quad \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i, \quad (3.3)$$

and the two scalars \bar{x} and \bar{y} are simply the coordinates of the center of mass of the input points. Substituting this expression for d in the definition of E yields

$$E = \sum_{i=1}^n [a(x_i - \bar{x}) + b(y_i - \bar{y})]^2 = |\mathcal{U}\mathbf{n}|^2 \quad \text{where} \quad \mathcal{U} = \begin{pmatrix} x_1 - \bar{x} & y_1 - \bar{y} \\ \dots & \dots \\ x_n - \bar{x} & y_n - \bar{y} \end{pmatrix},$$

and our original problem finally reduces to minimizing $|\mathcal{U}\mathbf{n}|^2$ with respect to \mathbf{n} under the constraint $|\mathbf{n}|^2 = 1$. We recognize a homogeneous linear least-squares problem, whose solution is the unit eigenvector associated with the minimum eigenvalue of the 2×2 matrix $\mathcal{U}^T \mathcal{U}$. Once a and b have been computed, the value of d is immediately obtained from Eq. (3.3). Note that $\mathcal{U}^T \mathcal{U}$ is easily shown to be equal to

$$\begin{pmatrix} \sum_{i=1}^n x_i^2 - n\bar{x}^2 & \sum_{i=1}^n x_i y_i - n\bar{x}\bar{y} \\ \sum_{i=1}^n x_i y_i - n\bar{x}\bar{y} & \sum_{i=1}^n y_i^2 - n\bar{y}^2 \end{pmatrix},$$

that is, the matrix of second moments of inertia of the points p_i . In fact, the line best fitting these points in the sense defined in this section is simply their axis of least inertia as defined in elementary mechanics.

3.1.2 Nonlinear Least-Squares Methods

Let us now consider a general system of p equations in q unknowns:

$$\begin{cases} f_1(x_1, x_2, \dots, x_q) = 0 \\ f_2(x_1, x_2, \dots, x_q) = 0 \\ \dots \\ f_p(x_1, x_2, \dots, x_q) = 0 \end{cases} \iff \mathbf{f}(\mathbf{x}) = \mathbf{0}. \quad (3.4)$$

Here, f_i denotes, for $i = 1, \dots, p$, a (possibly nonlinear) differentiable function from \mathbb{R}^p to \mathbb{R} , and we take $\mathbf{f} = (f_1, \dots, f_p)^T$ and $\mathbf{x} = (x_1, \dots, x_q)^T$. In general,

- when $p < q$, the solutions form a $(q - p)$ -dimensional *subset* of \mathbb{R}^q ;
- when $p = q$, there is a *finite set* of solutions;
- when $p > q$, there is no solution.

Let us emphasize the main differences with the linear case: In general, the dimension of the solution set is still $q - p$ in the underconstrained case, but this set does not form a vector space anymore. Its structure depends on the nature of the functions f_i . Likewise, there is usually a finite number of solutions instead of a unique one in the case $p = q$. A precise definition of

the general conditions that a family of functions f_i ($i = 1, \dots, p$) has to satisfy for the prior statement to be true is unfortunately beyond the scope of this book.

There is no general method for finding all the solutions of Eq. (3.4) when $p = q$ or for finding the global minimum of the least-squares error

$$E(\mathbf{x}) \stackrel{\text{def}}{=} |\mathbf{f}(\mathbf{x})|^2 = \sum_{i=1}^p f_i^2(\mathbf{x})$$

when $p > q$. Instead, we present next a number of iterative methods that linearize the problem in hope of finding at least one suitable solution. They all rely on a first-order Taylor expansion of the functions f_i in the neighborhood of a point \mathbf{x} :

$$f_i(\mathbf{x} + \delta\mathbf{x}) = f_i(\mathbf{x}) + \delta x_1 \frac{\partial f_i}{\partial x_1}(\mathbf{x}) + \dots + \delta x_q \frac{\partial f_i}{\partial x_q}(\mathbf{x}) + O(|\delta\mathbf{x}|^2) \approx f_i(\mathbf{x}) + \nabla f_i(\mathbf{x}) \cdot \delta\mathbf{x}.$$

Here, $\nabla f_i(\mathbf{x}) = (\partial f_i / \partial x_1, \dots, \partial f_i / \partial x_q)^T$ is the *gradient* of f_i at the point \mathbf{x} , and we have neglected the second-order term $O(|\delta\mathbf{x}|^2)$. It follows immediately that

$$\mathbf{f}(\mathbf{x} + \delta\mathbf{x}) \approx \mathbf{f}(\mathbf{x}) + \mathcal{J}_f(\mathbf{x})\delta\mathbf{x}, \quad (3.5)$$

where $\mathcal{J}_f(\mathbf{x})$ is the *Jacobian* of \mathbf{f} —that is, the $p \times q$ matrix

$$\mathcal{J}_f(\mathbf{x}) \stackrel{\text{def}}{=} \begin{pmatrix} \nabla f_1^T(\mathbf{x}) \\ \dots \\ \nabla f_p^T(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} \frac{\partial f_1}{\partial x_1}(\mathbf{x}) & \dots & \frac{\partial f_1}{\partial x_q}(\mathbf{x}) \\ \dots & \dots & \dots \\ \frac{\partial f_p}{\partial x_1}(\mathbf{x}) & \dots & \frac{\partial f_p}{\partial x_q}(\mathbf{x}) \end{pmatrix}.$$

Newton's Method: Square Systems of Nonlinear Equations As mentioned earlier, Eq. (3.4) admits (in general) a finite number of solutions when $p = q$. Although there is no general method for finding all of these solutions when \mathbf{f} is arbitrary, Eq. (3.5) can be used as the basis for a simple iterative algorithm for finding one of these solutions: Given some current estimate \mathbf{x} of the solution, the idea is to compute a perturbation $\delta\mathbf{x}$ of this estimate such that $\mathbf{f}(\mathbf{x} + \delta\mathbf{x}) \approx \mathbf{0}$, or, according to Eq. (3.5),

$$\mathcal{J}_f(\mathbf{x})\delta\mathbf{x} = -\mathbf{f}(\mathbf{x}).$$

When the Jacobian is nonsingular, $\delta\mathbf{x}$ is easily found as the solution of this $q \times q$ system of linear equations, and the process is repeated until convergence.

Newton's method converges rapidly once close to a solution: It has a *quadratic convergence rate* (i.e., the error at step $k + 1$ is proportional to the square of the error at step k). When started far from a solution, Newton's method as presented here may be unreliable. Various strategies can be used to improve its robustness, but their discussion is beyond the scope of this book.

Newton's Method: Overconstrained Systems of Nonlinear Equations When p is greater than q , we seek a local minimum of the least-squares error E . Newton's method can be adapted to this case by noting that such a minimum is a zero of the error's gradient. More precisely, we introduce $\mathbf{F}(\mathbf{x}) = \frac{1}{2}\nabla E(\mathbf{x})$ and use Newton's method to find the desired minimum as a solution of the $q \times q$ system of nonlinear equations $\mathbf{F}(\mathbf{x}) = \mathbf{0}$. Differentiating E shows that

$$\mathbf{F}(\mathbf{x}) = \mathcal{J}_f^T(\mathbf{x})\mathbf{f}(\mathbf{x}), \quad (3.6)$$

and differentiating this expression shows in turn that the Jacobian of \mathbf{F} is

$$\mathcal{J}_{\mathbf{F}}(\mathbf{x}) = \mathcal{J}_f^T(\mathbf{x})\mathcal{J}_f(\mathbf{x}) + \sum_{i=1}^p f_i(\mathbf{x})\mathcal{H}_{f_i}(\mathbf{x}). \quad (3.7)$$

In this equation, $\mathcal{H}_{f_i}(\mathbf{x})$ denotes the *Hessian* of f_i —that is, the $q \times q$ matrix of second derivatives

$$\mathcal{H}_{f_i}(\mathbf{x}) \stackrel{\text{def}}{=} \begin{pmatrix} \frac{\partial^2 f_i}{\partial x_1^2}(\mathbf{x}) & \cdots & \frac{\partial^2 f_i}{\partial x_1 \partial x_q}(\mathbf{x}) \\ \cdots & \cdots & \cdots \\ \frac{\partial^2 f_i}{\partial x_1 \partial x_q}(\mathbf{x}) & \cdots & \frac{\partial^2 f_i}{\partial x_q^2}(\mathbf{x}) \end{pmatrix}.$$

The term $\delta\mathbf{x}$ in Newton's method satisfies $\mathcal{J}_{\mathbf{F}}(\mathbf{x})\delta\mathbf{x} = -\mathbf{F}(\mathbf{x})$. Equivalently, combining Eqs. (3.6) and (3.7) shows that $\delta\mathbf{x}$ is the solution of

$$\left[\mathcal{J}_f^T(\mathbf{x})\mathcal{J}_f(\mathbf{x}) + \sum_{i=1}^p f_i(\mathbf{x})\mathcal{H}_{f_i}(\mathbf{x}) \right] \delta\mathbf{x} = -\mathcal{J}_f^T(\mathbf{x})\mathbf{f}(\mathbf{x}). \quad (3.8)$$

The Gauss–Newton and Levenberg–Marquardt Algorithms Newton's method requires computing the Hessians of the functions f_i , which may be difficult and/or expensive. We discuss here two other approaches to nonlinear least-squares that do not involve the Hessians. Let us first consider the Gauss–Newton algorithm: In this approach, we use again a first-order Taylor expansion of \mathbf{f} to minimize E , but this time we seek the value of $\delta\mathbf{x}$ that minimizes $E(\mathbf{x} + \delta\mathbf{x})$ for a given value of \mathbf{x} . Substituting Eq. (3.5) into Eq. (3.4) yields

$$E(\mathbf{x} + \delta\mathbf{x}) = |\mathbf{f}(\mathbf{x} + \delta\mathbf{x})|^2 \approx |\mathbf{f}(\mathbf{x}) + \mathcal{J}_f(\mathbf{x})\delta\mathbf{x}|^2.$$

At this point, we are back in the linear least-squares setting, and the adjustment $\delta\mathbf{x}$ can be computed as the solution of $\mathcal{J}_f^T(\mathbf{x})\delta\mathbf{x} = -\mathbf{f}(\mathbf{x})$ or, equivalently, according to the definition of the pseudoinverse,

$$\mathcal{J}_f^T(\mathbf{x})\mathcal{J}_f(\mathbf{x})\delta\mathbf{x} = -\mathcal{J}_f^T(\mathbf{x})\mathbf{f}(\mathbf{x}). \quad (3.9)$$

Comparing Eqs. (3.8) and (3.9), we see that the Gauss–Newton algorithm can be thought of as an approximation of Newton's method where the term involving the Hessians \mathcal{H}_{f_i} has been neglected. This is justified when the values of the functions f_i at a solution (the *residuals*) are small since the matrices \mathcal{H}_{f_i} are multiplied by these residuals in Eq. (3.8). In this case, the performance of the Gauss–Newton algorithm is comparable to that of Newton's method, with (nearly) quadratic convergence close to a solution. When the residuals at the solution are too large, however, it may converge slowly or not at all.

When Eq. (3.9) is replaced by

$$[\mathcal{J}_f^T(\mathbf{x})\mathcal{J}_f(\mathbf{x}) + \mu\text{Id}]\delta\mathbf{x} = -\mathcal{J}_f^T(\mathbf{x})\mathbf{f}(\mathbf{x}), \quad (3.10)$$

where the parameter μ is allowed to vary at each iteration, we obtain the Levenberg–Marquardt algorithm, popular in computer vision circles. This is another variant of Newton's method where the term involving the Hessians is this time approximated by a multiple of the identity matrix. The Levenberg–Marquardt algorithm has convergence properties comparable to its Gauss–Newton cousin, but it is more robust: For example, unlike that algorithm, it can be used when the Jacobian \mathcal{J}_f does not have maximal rank and its pseudoinverse does not exist.

3.2 A LINEAR APPROACH TO CAMERA CALIBRATION

It is now time to go back to geometric camera calibration. We assume in this section that a calibration rig is observed by a camera and that the image positions (u_i, v_i) of n points P_i ($i = 1, \dots, n$) with known homogeneous coordinate vectors \mathbf{P}_i have been found in a picture of the rig, either automatically or by hand. We decompose the calibration process into (a) the computation of the perspective projection matrix \mathcal{M} associated with the camera in this coordinate system, followed by (b) the estimation of the intrinsic and extrinsic parameters of the camera from this matrix. Degenerate point configurations for which the first step of this process may fail are identified in Section 3.2.3. As shown shortly, writing that the points p_i are the perspective images of the points P_i imposes a set of n linear constraints on the 11 independent coefficients of the corresponding projection matrix. When $n > 11$, these equations generally do not admit a common root, but the techniques introduced in Section 3.1.1 can be used to effectively construct their solution in the least-squares sense.

3.2.1 Estimation of the Projection Matrix

Let us assume that our camera has nonzero skew. According to Theorem 1 from chapter 2, the matrix \mathcal{M} is not singular, but otherwise arbitrary. Clearing the denominators in the perspective projection Eq. (2.16) yields

$$\begin{cases} (\mathbf{m}_1 - u_i \mathbf{m}_3) \cdot \mathbf{P} = 0, \\ (\mathbf{m}_2 - v_i \mathbf{m}_3) \cdot \mathbf{P} = 0. \end{cases}$$

Collecting the constraints associated with our n points yields a system of $2n$ homogeneous linear equations in the twelve coefficients of the matrix \mathcal{M} —namely, $\mathcal{P}\mathbf{m} = 0$, where

$$\mathcal{P} \stackrel{\text{def}}{=} \begin{pmatrix} \mathbf{P}_1^T & \mathbf{0}^T & -u_1 \mathbf{P}_1^T \\ \mathbf{0}^T & \mathbf{P}_1^T & -v_1 \mathbf{P}_1^T \\ \dots & \dots & \dots \\ \mathbf{P}_n^T & \mathbf{0}^T & -u_n \mathbf{P}_n^T \\ \mathbf{0}^T & \mathbf{P}_n^T & -v_n \mathbf{P}_n^T \end{pmatrix} \quad \text{and} \quad \mathbf{m} \stackrel{\text{def}}{=} \begin{pmatrix} m_1 \\ m_2 \\ m_3 \end{pmatrix} = 0.$$

When $n \geq 6$, homogeneous linear least-squares can be used to compute the value of the unit vector \mathbf{m} (hence the matrix \mathcal{M}) that minimizes $|\mathcal{P}\mathbf{m}|^2$ as the solution of an eigenvalue problem.

3.2.2 Estimation of the Intrinsic and Extrinsic Parameters

Once the projection matrix \mathcal{M} has been estimated, its expression in terms of the camera intrinsic and extrinsic parameters (Eq. [2.17] in chapter 2) can be used to recover these parameters as follows: We write as before $\mathcal{M} = (\mathcal{A} \ \mathbf{b})$, with \mathbf{a}_1^T , \mathbf{a}_2^T , and \mathbf{a}_3^T denoting the rows of \mathcal{A} . We obtain

$$\rho(\mathcal{A} \ \mathbf{b}) = \mathcal{K}(\mathcal{R} \ \mathbf{t}) \iff \rho \begin{pmatrix} \mathbf{a}_1^T \\ \mathbf{a}_2^T \\ \mathbf{a}_3^T \end{pmatrix} = \begin{pmatrix} \alpha \mathbf{r}_1^T - \alpha \cot \theta \mathbf{r}_2^T + u_0 \mathbf{r}_3^T \\ \frac{\beta}{\sin \theta} \mathbf{r}_2^T + v_0 \mathbf{r}_3^T \\ \mathbf{r}_3^T \end{pmatrix},$$

where ρ is an unknown scale factor introduced here to account for the fact that the recovered matrix \mathcal{M} has unit Frobenius form since $|\mathcal{M}| = |\mathbf{m}| = 1$.

In particular, using the fact that the rows of a rotation matrix have unit length and are perpendicular to each other yields immediately

$$\begin{cases} \rho = \varepsilon/|\mathbf{a}_3|, \\ \mathbf{r}_3 = \rho\mathbf{a}_3, \\ u_0 = \rho^2(\mathbf{a}_1 \cdot \mathbf{a}_3), \\ v_0 = \rho^2(\mathbf{a}_2 \cdot \mathbf{a}_3), \end{cases} \quad \text{where } \varepsilon = \mp 1. \quad (3.11)$$

Since θ is always in the neighborhood of $\pi/2$ with a positive sine, we have

$$\begin{cases} \rho^2(\mathbf{a}_1 \times \mathbf{a}_3) = -\alpha\mathbf{r}_2 - \alpha \cot\theta\mathbf{r}_1, \\ \rho^2(\mathbf{a}_2 \times \mathbf{a}_3) = \frac{\beta}{\sin\theta}\mathbf{r}_1, \end{cases} \quad \text{and} \quad \begin{cases} \rho^2|\mathbf{a}_1 \times \mathbf{a}_3| = \frac{|\alpha|}{\sin\theta}, \\ \rho^2|\mathbf{a}_2 \times \mathbf{a}_3| = \frac{|\beta|}{\sin\theta}. \end{cases} \quad (3.12)$$

Thus,

$$\begin{cases} \cos\theta = -\frac{(\mathbf{a}_1 \times \mathbf{a}_3) \cdot (\mathbf{a}_2 \times \mathbf{a}_3)}{|\mathbf{a}_1 \times \mathbf{a}_3||\mathbf{a}_2 \times \mathbf{a}_3|}, \\ \alpha = \rho^2|\mathbf{a}_1 \times \mathbf{a}_3|\sin\theta, \\ \beta = \rho^2|\mathbf{a}_2 \times \mathbf{a}_3|\sin\theta, \end{cases} \quad (3.13)$$

since the sign of the magnification parameters α and β is normally known in advance and can be taken to be positive.

We can now compute \mathbf{r}_1 and \mathbf{r}_2 from the second part of Eq. (3.12) as

$$\begin{cases} \mathbf{r}_1 = \frac{\rho^2 \sin\theta}{\beta}(\mathbf{a}_2 \times \mathbf{a}_3) = \frac{1}{|\mathbf{a}_2 \times \mathbf{a}_3|}(\mathbf{a}_2 \times \mathbf{a}_3), \\ \mathbf{r}_2 = \mathbf{r}_3 \times \mathbf{r}_1. \end{cases} \quad (3.14)$$

Note that there are two possible choices for the matrix \mathcal{R} depending on the value of ε . The translation parameters can now be recovered by writing $\mathcal{K}\mathbf{t} = \rho\mathbf{b}$, and hence $\mathbf{t} = \rho\mathcal{K}^{-1}\mathbf{b}$. In practical situations, the sign of t_z is often known in advance (this corresponds to knowing whether the origin of the world coordinate system is in front or behind the camera), which allows the choice of a unique solution for the calibration parameters.

3.2.3 Degenerate Point Configurations

We now examine the *degenerate configurations* of the points P_i ($i = 1, \dots, n$) that may cause the failure of the camera calibration process. We focus on the (ideal) case where the data points \mathbf{p}_i ($i = 1, \dots, n$) can be measured with zero error, and we identify the *nullspace* of the matrix \mathcal{P} (i.e., the subspace of \mathbb{R}^{12} formed by the vectors \mathbf{l} such that $\mathcal{P}\mathbf{l} = \mathbf{0}$).

Let \mathbf{l} be such a vector. Introducing the vectors formed by successive quadruples of its coordinates (i.e., $\boldsymbol{\lambda} = (l_1, l_2, l_3, l_4)^T$, $\boldsymbol{\mu} = (l_5, l_6, l_7, l_8)^T$, and $\boldsymbol{\nu} = (l_9, l_{10}, l_{11}, l_{12})^T$) allows us to write

$$\mathbf{0} = \mathcal{P}\mathbf{l} = \begin{pmatrix} \mathbf{P}_1^T & \mathbf{0}^T & -u_1\mathbf{P}_1^T \\ \mathbf{0}^T & \mathbf{P}_1^T & -v_1\mathbf{P}_1^T \\ \dots & \dots & \dots \\ \mathbf{P}_n^T & \mathbf{0}^T & -u_n\mathbf{P}_n^T \\ \mathbf{0}^T & \mathbf{P}_n^T & -v_n\mathbf{P}_n^T \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda} \\ \boldsymbol{\mu} \\ \boldsymbol{\nu} \end{pmatrix} = \begin{pmatrix} \mathbf{P}_1^T\boldsymbol{\lambda} - u_1\mathbf{P}_1^T\boldsymbol{\nu} \\ \mathbf{P}_1^T\boldsymbol{\mu} - v_1\mathbf{P}_1^T\boldsymbol{\nu} \\ \dots \\ \mathbf{P}_n^T\boldsymbol{\lambda} - u_n\mathbf{P}_n^T\boldsymbol{\nu} \\ \mathbf{P}_n^T\boldsymbol{\mu} - v_n\mathbf{P}_n^T\boldsymbol{\nu} \end{pmatrix}. \quad (3.15)$$

Combining Eqs. (2.16) and (3.15) yields

$$\begin{cases} P_i^T \lambda - \frac{m_1^T P_i}{m_3^T P_i} P_i^T \nu = 0, \\ P_i^T \mu - \frac{m_2^T P_i}{m_3^T P_i} P_i^T \nu = 0, \end{cases} \quad \text{for } i = 1, \dots, n.$$

After clearing the denominators and rearranging the terms, we finally obtain

$$\begin{cases} P_i^T (\lambda m_3^T - m_1 \nu^T) P_i = 0, \\ P_i^T (\mu m_3^T - m_2 \nu^T) P_i = 0, \end{cases} \quad \text{for } i = 1, \dots, n. \quad (3.16)$$

As expected, the vector l associated with $\lambda = m_1$, $\mu = m_2$, and $\nu = m_3$ is a solution of these equations. Are there other solutions?

Let us first consider the case where the points P_i ($i = 1, \dots, n$) all lie in some plane Π —that is, according to Eq. (2.2), $\Pi \cdot P_i = 0$ for some 4-vector Π . Clearly, choosing (λ, μ, ν) equal to $(\Pi, \mathbf{0}, \mathbf{0})$, $(\mathbf{0}, \Pi, \mathbf{0})$, or $(\mathbf{0}, \mathbf{0}, \Pi)$, or any linear combination of these vectors yields a solution of Eq. (3.16). In other words, the nullspace of \mathcal{P} contains the four-dimensional vector space spanned by these vectors and m . In practice, this means that the fiducial points P_i should not all lie in the same plane.

In general, for a given nonzero value of the vector l , the points P_i that satisfy Eq. (3.16) must lie on the curve where the two quadric surfaces defined by the corresponding equations intersect. A closer look at Eq. (3.16) reveals that the straight line where the planes defined by $m_3 \cdot P = 0$ and $\nu \cdot P = 0$ intersect lies on both quadrics. It can be shown that the intersection curve of these two surfaces consists of this line and a *twisted cubic* curve Γ passing through the origin. A twisted cubic is entirely determined by six points lying on it, and it follows that seven points chosen at random do not fall on Γ . Since this curve passes through the origin, choosing $n \geq 6$ random points generally guarantees that the matrix \mathcal{P} has rank 11 and the projection matrix can be recovered in a unique fashion.

3.3 TAKING RADIAL DISTORTION INTO ACCOUNT

We have assumed so far that our camera is equipped with a perfect lens. As shown in chapter 1, real lenses suffer from a number of aberrations. In this section, we show how to account for *radial distortion*, a type of aberration that depends on the distance separating the optical axis from the point of interest. We assume that the image center is known so that we can take $u_0 = v_0 = 0$ and model the projection process as

$$p = \frac{1}{z} \begin{pmatrix} 1/\lambda & 0 & 0 \\ 0 & 1/\lambda & 0 \\ 0 & 0 & 1 \end{pmatrix} \mathcal{M}P, \quad (3.17)$$

where λ is a polynomial function of the squared distance d^2 between the image center and the image point p . In most applications, it is sufficient to use a low-degree polynomial (e.g., $\lambda = 1 + \sum_{p=1}^q \kappa_p d^{2p}$, with $q \leq 3$) and the *distortion coefficients* κ_p ($p = 1, \dots, q$) are normally assumed to be small. Note that d^2 is naturally expressed in terms of the *normalized* image coordinates of the point p (i.e., $d^2 = \hat{u}^2 + \hat{v}^2$). Substituting $u_0 = 0$ and $v_0 = 0$ in Eq. (2.13) allows us, after some algebraic manipulation, to rewrite d^2 as a function of u and v instead—namely,

$$d^2 = \frac{u^2}{\alpha^2} + \frac{v^2}{\beta^2} + 2 \frac{uv}{\alpha\beta} \cos \theta. \quad (3.18)$$

Using Eq. (3.18) to write λ as an explicit function of u and v in Eq. (3.17) yields highly nonlinear constraints on the $q + 11$ camera parameters. Although these parameters in principle can all be found using the general *nonlinear least-squares* techniques introduced in the next section, we prefer here a two-stage approach tailored to the calibration problem: Eliminating λ from Eq. (3.17) first allows us to use *linear* least squares to estimate nine of the camera parameters. The $q+2$ remaining ones are then computed from Eqs. (3.17) and (3.18) by a simple nonlinear process.

3.3.1 Estimation of the Projection Matrix

Geometrically, radial distortion changes the distance between the image center and the image point p , but it does not affect the direction of the vector joining these two points. This is the *radial alignment constraint* introduced by Tsai (1987a), and it can be expressed algebraically by writing

$$\lambda \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \frac{\mathbf{m}_1 \cdot \mathbf{P}}{\mathbf{m}_3 \cdot \mathbf{P}} \\ \frac{\mathbf{m}_2 \cdot \mathbf{P}}{\mathbf{m}_3 \cdot \mathbf{P}} \end{pmatrix} \implies v(\mathbf{m}_1 \cdot \mathbf{P}) - u(\mathbf{m}_2 \cdot \mathbf{P}) = 0. \quad (3.19)$$

Given n fiducial points, we obtain n linear equations in the eight coefficients of the vectors \mathbf{m}_1 and \mathbf{m}_2 —namely,

$$\mathcal{Q}\mathbf{n} = 0, \quad \text{where} \quad \mathcal{Q} \stackrel{\text{def}}{=} \begin{pmatrix} v_1 \mathbf{P}_1^T & -u_1 \mathbf{P}_1^T \\ \dots & \dots \\ v_n \mathbf{P}_n^T & -u_n \mathbf{P}_n^T \end{pmatrix} \quad \text{and} \quad \mathbf{n} = \begin{pmatrix} \mathbf{m}_1 \\ \mathbf{m}_2 \end{pmatrix}. \quad (3.20)$$

Note the similarity with the previous case. When $n \geq 8$, this system of equations is in general overconstrained, and a solution with unit norm can be found using linear least squares.

3.3.2 Estimation of the Intrinsic and Extrinsic Parameters

Once \mathbf{m}_1 and \mathbf{m}_2 have been estimated, we can define as before the corresponding values of \mathbf{a}_1 , \mathbf{a}_2 and write

$$\rho \begin{pmatrix} \mathbf{a}_1^T \\ \mathbf{a}_2^T \end{pmatrix} = \begin{pmatrix} \alpha \mathbf{r}_1^T - \alpha \cot \theta \mathbf{r}_2^T + u_0 \mathbf{r}_3^T \\ \frac{\beta}{\sin \theta} \mathbf{r}_2^T + v_0 \mathbf{r}_3^T \end{pmatrix}.$$

Calculating the norm and dot product of the vectors \mathbf{a}_1 and \mathbf{a}_2 immediately yields the aspect ratio and the skew of the camera as

$$\frac{\beta}{\alpha} = \frac{|\mathbf{a}_2|}{|\mathbf{a}_1|} \quad \text{and} \quad \cos \theta = -\frac{\mathbf{a}_1 \cdot \mathbf{a}_2}{|\mathbf{a}_1||\mathbf{a}_2|}. \quad (3.21)$$

Using the fact that \mathbf{r}_2^T is the second row of a rotation matrix, and thus has unit norm, now yields

$$\alpha = \varepsilon \rho |\mathbf{a}_1| \sin \theta \quad \text{and} \quad \beta = \varepsilon \rho |\mathbf{a}_2| \sin \theta, \quad (3.22)$$

where, as before, $\varepsilon = \mp 1$. After some simple algebraic manipulation, we obtain

$$\begin{cases} \mathbf{r}_1 = \frac{\varepsilon}{\sin \theta} \left(\frac{1}{|\mathbf{a}_1|} \mathbf{a}_1 + \frac{\cos \theta}{|\mathbf{a}_2|} \mathbf{a}_2 \right), \\ \mathbf{r}_2 = \frac{\varepsilon}{|\mathbf{a}_2|} \mathbf{a}_2. \end{cases}$$

Using these equations and $\mathbf{r}_3 = \mathbf{r}_1 \times \mathbf{r}_2$ allows us to recover the rotation matrix \mathcal{R} up to a twofold ambiguity. Two of the translation parameters can also be recovered by writing

$$\begin{pmatrix} \alpha t_x - \alpha \cot \theta t_y \\ \frac{\beta}{\sin \theta} t_y \end{pmatrix} = \rho \begin{pmatrix} b_1 \\ b_2 \end{pmatrix},$$

where b_1 and b_2 are the first two coordinates of the vector \mathbf{b} , which in turn allows us to compute these parameters t_x and t_y as

$$\begin{cases} t_x = \frac{\varepsilon}{\sin \theta} \left(\frac{b_1}{|\mathbf{a}_1|} + \frac{b_2 \cos \theta}{|\mathbf{a}_2|} \right), \\ t_y = \frac{\varepsilon b_2}{|\mathbf{a}_2|}. \end{cases}$$

Without further constraints, it is impossible to recover t_z and the absolute scale of the magnification parameters, or equivalently, the value of ρ , from the values of \mathbf{m}_1 and \mathbf{m}_2 only. To estimate these parameters, it is necessary to go back to the original projection equations: We rewrite the left side of Eq. (3.19) as

$$\begin{cases} (\mathbf{m}_1 - \lambda u \mathbf{m}_3) \cdot \mathbf{P} = 0, \\ (\mathbf{m}_2 - \lambda v \mathbf{m}_3) \cdot \mathbf{P} = 0, \end{cases} \quad (3.23)$$

Here \mathbf{m}_1 and \mathbf{m}_2 are known and, according to Eq. (2.17), $\mathbf{m}_3^T = (\mathbf{r}_3^T \quad t_z)$, where \mathbf{r}_3 is also known. Now, combining the expression for d^2 given in Eq. (3.18) with the expressions for α , β , and $\cos \theta$ given in Eqs. (3.21) and (3.22) yields

$$d^2 = \frac{1}{\rho^2} \frac{|u \mathbf{a}_2 - v \mathbf{a}_1|^2}{|\mathbf{a}_1 \times \mathbf{a}_2|^2},$$

and substituting this value in Eq. (3.23) yields a nonlinear equation in ρ , t_z , and the distortion parameters κ_p ($p = 1, \dots, q$). Given enough data points, the nonlinear least-squares techniques that have been presented in Section 3.1.2 can be used to solve for these parameters. These methods are iterative and require initial guesses for all unknowns. Here, a reasonable estimate for ρ and t_z can be found using linear least squares by first assuming that $\lambda = 1$. Likewise, zero values are reasonable initial guesses for the distortion parameters. As before, the twofold ambiguity can be resolved when the sign of t_z is known in advance.

3.3.3 Degenerate Point Configurations

Let us determine the degenerate point configurations for which the vectors \mathbf{m}_1 and \mathbf{m}_2 cannot be uniquely determined. Given a vector \mathbf{l} in the nullspace, we define the vectors $\boldsymbol{\lambda} = (l_1, l_2, l_3, l_4)^T$ and $\boldsymbol{\mu} = (l_5, l_6, l_7, l_8)^T$ and write

$$\mathbf{0} = \mathcal{Q} \mathbf{l} = \begin{pmatrix} v_1 \mathbf{P}_1^T & -u_1 \mathbf{P}_1^T \\ \dots & \dots \\ v_n \mathbf{P}_n^T & -u_n \mathbf{P}_n^T \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda} \\ \boldsymbol{\mu} \end{pmatrix} = \begin{pmatrix} v_1 \mathbf{P}_1^T \boldsymbol{\lambda} - u_1 \mathbf{P}_1^T \boldsymbol{\mu} \\ \dots \\ v_n \mathbf{P}_n^T \boldsymbol{\lambda} - u_n \mathbf{P}_n^T \boldsymbol{\mu} \end{pmatrix}.$$

Taking into account the values of u_i and v_i , rearranging the terms, and clearing the denominators yields

$$\mathbf{P}_i^T (\mathbf{m}_2 \boldsymbol{\lambda}^T - \mathbf{m}_1 \boldsymbol{\mu}^T) \mathbf{P}_i = 0 \quad \text{for } i = 1, \dots, n. \quad (3.24)$$

The vector l associated with $\boldsymbol{\lambda} = \mathbf{m}_1$ and $\boldsymbol{\mu} = \mathbf{m}_2$ is of course a solution of these equations (in the noise-free case; i.e., when all image positions are exact). When the points P_i ($i = 1, \dots, n$) all lie in some plane Π or, equivalently, $\Pi \cdot P_i = 0$ for some 4-vector Π , we can choose $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ equal to $(\Pi, \mathbf{0})$, $(\mathbf{0}, \Pi)$, or any linear combination of these two vectors, and construct a solution of Eq. (3.24). The nullspace of \mathcal{Q} contains the three-dimensional vector space spanned by these vectors and l . Thus, as before, points that all lie in the same plane cannot be used in this calibration method.

More generally, for a given value of $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$, the points P_i form a degenerate configuration when they lie on the quadric surface defined by Eq. (3.24). Note that this surface contains the four straight lines defined by $\boldsymbol{\lambda} \cdot \mathbf{P} = \boldsymbol{\mu} \cdot \mathbf{P} = 0$, $\boldsymbol{\lambda} \cdot \mathbf{P} = \mathbf{m}_1 \cdot \mathbf{P} = 0$, $\boldsymbol{\mu} \cdot \mathbf{P} = \mathbf{m}_2 \cdot \mathbf{P} = 0$ and $\mathbf{m}_1 \cdot \mathbf{P} = \mathbf{m}_2 \cdot \mathbf{P} = 0$. Therefore, it must consist of two planes or be a cone, hyperboloid of one sheet, or hyperbolic paraboloid. In any case, for a large enough number of points in general position, our least-squares problem admits a unique solution.

3.4 ANALYTICAL PHOTOGRAMMETRY

The techniques presented so far ignore some of the constraints associated with the calibration process. For example, the camera skew was assumed to be arbitrary instead of (very close to) zero in Section 3.2. We present in this section a nonlinear approach to camera calibration that takes into account *all* the relevant constraints. This approach is borrowed from *photogrammetry*—an engineering field whose aim is to recover quantitative geometric information from one or several pictures, with applications in cartography, military intelligence, city planning, and so on. For many years, photogrammetry relied on a combination of geometric, optical, and mechanical methods to recover three-dimensional information from pictures, but the advent of computers in the 1950s has made a purely computational approach to this problem feasible. This is the domain of *analytical photogrammetry*, where the intrinsic parameters of a camera define its *interior orientation*, and the extrinsic parameters define its *exterior orientation*.

In this setting, we assume once again that we observe n fiducial points P_i ($i = 1, \dots, n$) whose positions in some world coordinate system are known, and we minimize the mean-squared distance between the measured positions (u_i, v_i) of their images and the positions $(\tilde{u}_i, \tilde{v}_i)$ predicted by the perspective projection equation with respect to a vector of camera parameters $\boldsymbol{\xi} = (\xi_1, \dots, \xi_q)^T$ ($q \geq 11$) that may include various distortion coefficients in addition to the usual intrinsic and extrinsic parameters. The least-squares error can be written as

$$E(\boldsymbol{\xi}) = \sum_{i=1}^n [(\tilde{u}_i(\boldsymbol{\xi}) - u_i)^2 + (\tilde{v}_i(\boldsymbol{\xi}) - v_i)^2],$$

where

$$\tilde{u}_i(\boldsymbol{\xi}) \stackrel{\text{def}}{=} \frac{\mathbf{m}_1(\boldsymbol{\xi}) \cdot \mathbf{P}_i}{\mathbf{m}_3(\boldsymbol{\xi}) \cdot \mathbf{P}_i} \quad \text{and} \quad \tilde{v}_i(\boldsymbol{\xi}) \stackrel{\text{def}}{=} \frac{\mathbf{m}_2(\boldsymbol{\xi}) \cdot \mathbf{P}_i}{\mathbf{m}_3(\boldsymbol{\xi}) \cdot \mathbf{P}_i}.$$

Contrary to the cases studied so far, the dependency of each error term on the unknown parameters $\boldsymbol{\xi}$ is not linear. Instead, it involves a combination of polynomial and trigonometric functions, and minimizing the overall error measure involves the use of the nonlinear least squares algorithms discussed in Section 3.1.2. To follow the notation introduced in that section,

we rewrite our error function as

$$E(\boldsymbol{\xi}) = |\mathbf{f}(\boldsymbol{\xi})|^2 = \sum_{j=1}^{2n} f_j^2(\boldsymbol{\xi}), \text{ where } \begin{cases} f_{2i-1}(\boldsymbol{\xi}) = \tilde{u}_i(\boldsymbol{\xi}) - u_i \\ f_{2i}(\boldsymbol{\xi}) = \tilde{v}_i(\boldsymbol{\xi}) - v_i \end{cases} \text{ for } i = 1, \dots, n.$$

The Gauss–Newton and Levenberg–Marquard techniques described in Section 3.1.2 require the gradient of the functions f_j , and Newton’s method requires both their gradient and Hessian. Here we only calculate the gradient or, equivalently, the Jacobian of \mathbf{f} . Let us drop the $\boldsymbol{\xi}$ argument for conciseness and define $\tilde{x}_i = \mathbf{m}_1 \cdot \mathbf{P}_i$, $\tilde{y}_i = \mathbf{m}_2 \cdot \mathbf{P}_i$ and $\tilde{z}_i = \mathbf{m}_3 \cdot \mathbf{P}_i$ ($i = 1, \dots, n$), so $\tilde{u}_i = \tilde{x}_i/\tilde{z}_i$ and $\tilde{v}_i = \tilde{y}_i/\tilde{z}_i$. We have

$$\begin{cases} \frac{\partial f_{2i-1}}{\partial \xi_j} = \frac{\partial \tilde{u}_i}{\partial \xi_j} = \frac{1}{z_i} \frac{\partial \tilde{x}_i}{\partial \xi_j} - \frac{\tilde{x}_i}{\tilde{z}_i^2} \frac{\partial \tilde{z}_i}{\partial \xi_j} = \frac{1}{\tilde{z}_i} \left(\frac{\partial}{\partial \xi_j} (\mathbf{m}_1 \cdot \mathbf{P}_i) - \tilde{u}_i \frac{\partial}{\partial \xi_j} (\mathbf{m}_3 \cdot \mathbf{P}_i) \right), \\ \frac{\partial f_{2i}}{\partial \xi_j} = \frac{\partial \tilde{v}_i}{\partial \xi_j} = \frac{1}{\tilde{z}_i} \frac{\partial \tilde{y}_i}{\partial \xi_j} - \frac{\tilde{y}_i}{\tilde{z}_i^2} \frac{\partial \tilde{z}_i}{\partial \xi_j} = \frac{1}{\tilde{z}_i} \left(\frac{\partial}{\partial \xi_j} (\mathbf{m}_2 \cdot \mathbf{P}_i) - \tilde{v}_i \frac{\partial}{\partial \xi_j} (\mathbf{m}_3 \cdot \mathbf{P}_i) \right), \end{cases}$$

which is easily rewritten as

$$\begin{pmatrix} \frac{\partial f_{2i-1}}{\partial \xi_j} \\ \frac{\partial f_{2i}}{\partial \xi_j} \end{pmatrix} = \frac{1}{\tilde{z}_i} \begin{pmatrix} \mathbf{P}_i^T & \mathbf{0}^T & -\tilde{u}_i \mathbf{P}_i^T \\ \mathbf{0}^T & \mathbf{P}_i^T & -\tilde{v}_i \mathbf{P}_i^T \end{pmatrix} \mathcal{J}\mathbf{m},$$

where \mathbf{m} is as before the vector of \mathbb{R}^{12} associated with \mathcal{M} , and $\mathcal{J}\mathbf{m}$ denotes its Jacobian with respect to $\boldsymbol{\xi}$. We finally obtain the Jacobian of \mathbf{f} as

$$\mathcal{J}\mathbf{f} = \begin{pmatrix} \frac{1}{\tilde{z}_1} \mathbf{P}_1^T & \mathbf{0}^T & -\frac{\tilde{u}_1}{\tilde{z}_1} \mathbf{P}_1^T \\ \mathbf{0}^T & \frac{1}{\tilde{z}_1} \mathbf{P}_1^T & -\frac{\tilde{v}_1}{\tilde{z}_1} \mathbf{P}_1^T \\ \dots & \dots & \dots \\ \frac{1}{\tilde{z}_n} \mathbf{P}_n^T & \mathbf{0}^T & -\frac{\tilde{u}_n}{\tilde{z}_n} \mathbf{P}_n^T \\ \mathbf{0}^T & \frac{1}{\tilde{z}_n} \mathbf{P}_n^T & -\frac{\tilde{v}_n}{\tilde{z}_n} \mathbf{P}_n^T \end{pmatrix} \mathcal{J}\mathbf{m}.$$

In this expression, \tilde{u}_i , \tilde{v}_i , \tilde{z}_i , and \mathbf{P}_i depend on the point considered, but $\mathcal{J}\mathbf{m}$ only depends on the intrinsic and extrinsic parameters of the camera. Note that this method requires an explicit parameterization of the matrix \mathcal{R} . Such a parameterization in terms of three elementary rotations about coordinate axes was mentioned in chapter 2. Many other parameterizations can be used as well (see Exercises and chapter 21).

3.5 AN APPLICATION: MOBILE ROBOT LOCALIZATION

The calibration methods presented in this chapter can be used in a variety of applications, from metrology to stereo vision and object localization in robotic tasks. Here we briefly describe the nonlinear approach to camera calibration proposed by Devy *et al.* (1997) and its application to mobile robot localization. Unlike the techniques discussed so far, this method uses several images (up to 20 in the experiments presented here) of a planar rectangular grid to calibrate a




Figure 3.3 Calibration experiments. Left: One of the 20 input pictures; note the strong radial distortion. Note the mobile robot at the top of the photograph with its characteristic LED pattern. Right: Average and maximum reprojections errors (in pixels) in the 20 images. *Calibration and localization software courtesy of Michel Devy. Experiments courtesy of Fred Rothganger.*

static camera (Figure 3.3). One of these pictures is taken with the grid lying on the ground, and it is used to define the world coordinate system. After a rough manual localization, the corners of the grid are found in each picture with a precision of 1/10 pixel using a parametric model of the gray-level surface in the neighborhood of a corner.

The imaging geometry is modeled as in Section 3.3, with three radial distortion coefficients and zero skew. The calibration algorithm recovers a single set of intrinsic parameters, and one set of extrinsic parameters per input image. An initial guess for the intrinsic parameters can be obtained from information supplied by the camera and frame-grabber manufacturers. An initial guess for the extrinsic parameters can be obtained for each image using a variant of Tsai's (1987a) algorithm. Briefly, the projection matrix is estimated via linear least squares by choosing the z coordinate axis of the world reference frame perpendicular to the calibration grid. Accordingly, Eq. (3.20) now becomes

$$\mathcal{Q}'\mathbf{n}' = 0, \quad \text{where} \quad \mathcal{Q}' = \begin{pmatrix} v_1x_1 & v_1y_1 & v_1 & -u_1x_1 & -u_1y_1 & -u_1 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ v_nx_n & v_ny_n & v_n & -u_nx_n & -u_ny_n & -u_n \end{pmatrix}$$

and $\mathbf{n}' = (m_{11}, m_{12}, m_{14}, m_{21}, m_{22}, m_{24})^T$.

Note that explicitly imposing that $z_i = 0$ avoids the degeneracies encountered by the previously discussed algorithms for coplanar points. Once \mathbf{n}' is known, since the intrinsic parameters are assumed to be (roughly) known, it is a simple matter to compute the extrinsic parameters (see Exercises). Once initial guesses for both the intrinsic and extrinsic parameters are available, non-linear optimization (in this case, the Levenberg–Marquardt algorithm) can be used to minimize as usual the mean-squared distance between predicted and observed image features.

Figure 3.3 shows the result of experiments conducted with a 576×768 camera equipped with a 4.5 mm lens and plots the errors found when reprojecting the corners of the calibration grid model into the 20 images. Once the camera has been calibrated, it can be used to monitor the position and orientation of mobile robots in the coordinate system attached to the ground reference image. Each robot carries an array of infrared LEDs forming a distinctive pattern (Figure 3.3). During localization experiments, the camera is equipped with an infrared filter that effectively blocks out all incoming light except for that from the LEDs. Each robot is identified

using a simple pattern matching algorithm, and its position and orientation are deduced from the image position of the LEDs and the camera parameters. With the camera mounted 4 m above the ground, typical localization errors within the entire field of view of the camera are below 2 cm in position and 1° in orientation, with maximum errors that may reach 5 cm and 5° .

3.6 NOTES

The linear calibration technique described in Section 3.2 is detailed in Faugeras (1993). Its variant that takes radial distortion into account is adapted from Tsai (1987a). Haralick and Shapiro (1992) present a concise introduction to analytical photogrammetry. The *Manual of Photogrammetry* is of course the gold standard, and newcomers (such as the authors of this book) will probably find the ingenious mechanisms and rigorous methods described in its various editions fascinating (Thompson *et al.*, 1966, Slama *et al.*, 1980). We come back to photogrammetry in the context of multiple images in chapter 10. General optimization techniques are discussed in Luenberger (1985), Bertsekas (1995), and Heath (2002), for example. An excellent survey and discussion of least-squares methods in the context of analytical photogrammetry can be found in Triggs *et al.* (2000). The output of least-squares methods admits a statistical interpretation in maximum-likelihood terms when the coordinates of the data points are modeled as random variables obeying a normal distribution. We come back to this interpretation in chapter 15, where we also revisit the problem of fitting a straight line to a set of points in the plane.

PROBLEMS

- 3.1. Show that the vector \mathbf{x} that minimizes $|\mathcal{U}\mathbf{x}|^2$ under the constraint $|\mathcal{V}\mathbf{x}|^2 = 1$ is the unit generalized eigenvector associated with the minimum generalized eigenvalue of the symmetric matrices $\mathcal{U}^T\mathcal{U}$ and $\mathcal{V}^T\mathcal{V}$.

Hint: This problem is equivalent to minimizing (without constraints) the error $E = |\mathcal{U}\mathbf{x}|^2/|\mathcal{V}\mathbf{x}|^2$ with respect to \mathbf{x} .

- 3.2. Show that the 2×2 matrix $\mathcal{U}^T\mathcal{U}$ involved in the line-fitting example from Section 3.1.1 is the matrix of second moments of inertia of the points p_i ($i = 1, \dots, n$).
- 3.3. Extend the line-fitting method presented in Section 3.1.1 to the problem of fitting a plane to n points in \mathbb{E}^3 .
- 3.4. Derive an expression for the Hessian of the functions $f_{2i-1}(\boldsymbol{\xi}) = \tilde{u}_i(\boldsymbol{\xi}) - u_i$ and $f_{2i}(\boldsymbol{\xi}) = \tilde{v}_i(\boldsymbol{\xi}) - v_i$ ($i = 1, \dots, n$) introduced in Section 3.4.
- 3.5. Euler angles. Show that the rotation obtained by first rotating about the z axis of some coordinate frame by an angle α , then rotating about the y axis of the new coordinate frame by an angle β and finally rotating about the z axis of the resulting frame by an angle γ can be represented in the original coordinate system by

$$\begin{pmatrix} \cos \alpha \cos \beta \cos \gamma - \sin \alpha \sin \gamma & -\cos \alpha \cos \beta \sin \gamma - \sin \alpha \cos \gamma & \cos \alpha \sin \beta \\ \sin \alpha \cos \beta \cos \gamma + \cos \alpha \sin \gamma & -\sin \alpha \cos \beta \sin \gamma + \cos \alpha \cos \gamma & \sin \alpha \sin \beta \\ -\sin \beta \cos \gamma & \sin \beta \sin \gamma & \cos \beta \end{pmatrix}.$$

- 3.6. The Rodrigues formula. Consider a rotation \mathcal{R} of angle θ about the axis \mathbf{u} (a unit vector). Show that $\mathcal{R}\mathbf{x} = \cos \theta \mathbf{x} + \sin \theta \mathbf{u} \times \mathbf{x} + (1 - \cos \theta)(\mathbf{u} \cdot \mathbf{x})\mathbf{u}$.

Hint: A rotation does not change the projection of a vector \mathbf{x} onto the direction \mathbf{u} of its axis and applies a planar rotation of angle θ to the projection of \mathbf{x} into the plane orthogonal to \mathbf{u} .

3.7. Use the Rodrigues formula to show that the matrix associated with \mathcal{R} is

$$\begin{pmatrix} u^2(1-c) + c & uv(1-c) - ws & uw(1-c) + vs \\ uv(1-c) + ws & v^2(1-c) + c & vw(1-c) - us \\ uw(1-c) - vs & vw(1-c) + us & w^2(1-c) + c \end{pmatrix}.$$

where $c = \cos \theta$ and $s = \sin \theta$.

3.8. Assuming that the intrinsic parameters of a camera are known, show how to compute its extrinsic parameters once the vector \mathbf{n}' defined in Section 3.5 is known.

Hint: Use the fact that the rows of a rotation matrix have unit norm.

3.9. Assume that n fiducial lines with known Plücker coordinates are observed by a camera.

(a) Show that the line projection matrix $\tilde{\mathcal{M}}$ introduced in the exercises of chapter 2 can be recovered using linear least squares when $n \geq 9$.

(b) Show that once $\tilde{\mathcal{M}}$ is known, the projection matrix \mathcal{M} can also be recovered using linear least squares.

Hint: Consider the rows \mathbf{m}_i of \mathcal{M} as the coordinate vectors of three planes Π_i and the rows $\tilde{\mathbf{m}}_i$ of $\tilde{\mathcal{M}}$ as the coordinate vectors of three lines, and use the incidence relationships between these planes and these lines to derive linear constraints on the vectors \mathbf{m}_i .

Programming Assignments

3.10. Use linear least-squares to fit a plane to n points $(x_i, y_i, z_i)^T$ ($i = 1, \dots, n$) in \mathbb{R}^3 .

3.11. Use linear least-squares to fit a conic section defined by $ax^2 + bxy + cy^2 + dx + ey + f = 0$ to n points $(x_i, y_i)^T$ ($i = 1, \dots, n$) in \mathbb{R}^2 .

3.12. Implement the linear calibration algorithm presented in Section 3.2.

3.13. Implement the calibration algorithm that takes into account radial distortion and that was presented in Section 3.3.

3.14. Implement the nonlinear calibration algorithm from Section 3.4.