Chapter 2
Passive 3D Imaging

Stephen Se and Nick Pears

Abstract   We describe passive, multiple-view 3D imaging systems that recover 3D information from scenes that are illuminated only with ambient lighting. Much of the material is concerned with using the geometry of stereo 3D imaging to formulate estimation problems. Firstly, we present an overview of the common techniques used to recover 3D information from camera images. Secondly, we discuss camera modeling and camera calibration as an essential introduction to the geometry of the imaging process and the estimation of geometric parameters. Thirdly, we focus on 3D recovery from multiple views, which can be obtained using multiple cameras at the same time (stereo), or a single moving camera at different times (structure from motion). Epipolar geometry and finding image correspondences associated with the same 3D scene point are two key aspects for such systems, since epipolar geometry establishes the relationship between two camera views, while depth information can be inferred from the correspondences. The details of both stereo and structure from motion, the two essential forms of multiple-view 3D reconstruction technique, are presented. Towards the end of the chapter, we present several real-world applications.

2.1 Introduction

Passive 3D imaging has been studied extensively for several decades and it is a core topic in many of the major computer vision conferences and journals. Essentially, a *passive* 3D imaging system, also known as a passive 3D vision system, is one in which we can recover 3D scene information, without that system having to project its own source of light or other source of electromagnetic radiation

S. Se (✉)
MDA Systems Ltd., 13800 Commerce Parkway, Richmond, BC V6V 2J3, Canada
e-mail: sse@mdacorporation.com

N. Pears
Department of Computer Science, University of York, Deramore Lane, York YO10 5GH, UK
e-mail: nick.pears@york.ac.uk

(EMR) onto that scene. By contrast, an *active* 3D imaging system has an EMR projection subsystem, which is commonly in the infra-red or visible wavelength region.

Several passive 3D information sources (cues) relate closely to human vision and other animal vision. For example, in stereo vision, fusing the images recorded by our two eyes and exploiting the difference between them gives us a sense of depth. The aim of this chapter is to present the fundamental principles of passive 3D imaging systems so that readers can understand their strengths and limitations, as well as how to implement a subset of such systems, namely those that exploit multiple views of the scene.

Passive, multiple-view 3D imaging originates from the mature field of photogrammetry and, more recently, from the younger field of computer vision. In contrast to photogrammetry, computer vision applications rely on fast, automatic techniques, sometimes at the expense of precision. Our focus is from the computer vision perspective.

A recurring theme of this chapter is that we consider some aspect of the geometry of 3D imaging and formulate a linear least squares estimation problem to estimate the associated geometric parameters. These estimates can then optionally be improved, depending on the speed and accuracy requirements of the application, using the linear estimate as an initialization for a non-linear least squares refinement. In contrast to the linear stage, this non-linear stage usually optimizes a cost function that has a well-defined geometric meaning.

Chapter Outline   We will start with an overview of various techniques for passive 3D imaging systems, including single-view and multiple-view approaches. However, the main body of this chapter is focused on 3D recovery from multiple views, which can be obtained using multiple cameras simultaneously (stereo) or a single moving camera (structure from motion). A good starting point to understand this subject matter is knowledge of the image formation process in a single camera and how to capture this process in a camera model. This modeling is presented in Sect. 2.3 and the following section describes camera calibration: the estimation of the parameters in the developed camera model. In order to understand how to search efficiently for left-right feature pairs that correspond to the same scene point in a stereo image pair (the *correspondence* problem), a good understanding of two-view geometry is required, which establishes the relationship between two camera views. Hence Sect. 2.5 details this geometry, known as *epipolar geometry*, and shows how it can be captured and used in linear (vector-matrix) form. Following this, we can begin to consider the correspondence problem and the first step is to simplify the search to be across the same horizontal scanlines in each image, by warping the stereo image pair in a process known as *rectification*. This is described in Sect. 2.6. The following section then focuses on the correspondence search itself and then Sect. 2.8 details the process of generating a 3D point cloud from a set of image correspondences.

With increasing computer processing power and decreasing camera prices, many real-world applications of passive 3D imaging systems have been emerging in re-

cent years. Thus, later in the chapter (Sect. 2.9), some recent applications involving such systems are presented. Several commercially available stereo vision systems will first be presented. We then describe 3D modeling systems that generate photo-realistic 3D models from image sequences, which have a wide range of applications. Later in this section, passive 3D imaging systems for mobile robot pose estimation and obstacle detection are described. Finally, multiple-view passive 3D imaging systems are compared to their counterpart within active 3D imaging systems. This acts as a bridge to Chap. 3, where such systems will be discussed in detail.

2.2 An Overview of Passive 3D Imaging Systems

Most cameras today use either a *Charge Coupled Device* (CCD) image sensor or a *Complementary Metal Oxide Semiconductor* (CMOS) sensor, both of which capture light and convert it into electrical signals. Typically, CCD sensors provide higher quality, lower noise images whereas CMOS sensors are less expensive, more compact and consume less power. However, these stereotypes are becoming less pronounced. The cameras employing such image sensors can be hand-held or mounted on different platforms such as *Unmanned Ground Vehicles* (UGVs), *Unmanned Aerial Vehicles* (UAVs) and optical satellites.

Passive 3D vision techniques can be categorized as follows: (i) Multiple view approaches, (ii) Single view approaches. We outline each of these in the following two subsections.

2.2.1 Multiple View Approaches

In multiple view approaches, the scene is observed from two or more viewpoints, by either multiple cameras at the same time (stereo) or a single moving camera at different times (structure from motion). From the gathered images, the system is to infer information on the 3D structure of the scene.

Stereo refers to multiple images taken simultaneously using two or more cameras, which are collectively called a stereo camera. For example, binocular stereo uses two viewpoints, trinocular stereo uses three viewpoints, or alternatively there may be many cameras distributed around the viewing sphere of an object. *Stereo* derives from the Greek word *stereos* meaning *solid*, thus implying a 3D form of visual information. In this chapter, we will use the term *stereo vision* to imply a binocular stereo system. At the top of Fig. 2.1, we show an outline of such a system.

If we can determine that imaged points in the left and right cameras correspond to the same scene point, then we can determine two directions (3D rays) along which the 3D point must lie. (The camera parameters required to convert the 2D image positions to 3D rays come from a camera calibration procedure.) Then, we can intersect the 3D rays to determine the 3D position of the scene point, in a process

Fig. 2.1 *Top*: Plan view of
the operation of a simple
stereo rig. Here the optical
axes of the two cameras are
parallel to form a rectilinear
rig. However, often the
cameras are rotated towards
each other (verged) to
increase the overlap in their
fields of view. *Center*:
A commercial stereo camera,
supplied by Videre Design
(figure courtesy of [59]),
containing SRI's Small
Vision System [26]. *Bottom*:
Left and *right* views of a
stereo pair (images courtesy
of [34])

known as *triangulation*. A scene point, **X**, is shown in Fig. 2.1 as the intersection
of two rays (colored black) and a nearer point is shown by the intersection of two
different rays (colored blue). Note that the difference between left and right image
positions, the *disparity*, is greater for the nearer scene point. Note also that the scene
surface colored red cannot be observed by the right camera, in which case no 3D
shape measurement can be made. This scene portion is sometimes referred to as a
missing part and is the result of self-occlusion. A final point to note is that, although
the real image sensor is behind the lens, it is common practice to envisage and use
a conceptual image position in front of the lens so that the image is the same orien-

tation as the scene (i.e. not inverted top to bottom and left to right) and this position is shown in the figure.

Despite the apparent simplicity of Fig. 2.1(top), a large part of this chapter is required to present the various aspects of stereo 3D imaging in detail, such as calibration, determining left-to-right image correspondences and dense 3D shape reconstruction. A typical commercial stereo camera, supplied by Videre Design.[1] is shown in the center of Fig. 2.1, although many computer vision researchers build their own stereo rigs, using off-the-shelf digital cameras and a slotted steel bar mounted on a tripod. Finally, at the bottom of Fig. 2.1, we show the left and right views of a typical stereo pair taken from the Middlebury webpage [34].

In contrast to stereo vision, *structure from motion* (SfM) refers to a single moving camera scenario, where image sequences are captured over a period of time. While stereo refers to fixed relative viewpoints with synchronized image capture, SfM refers to variable viewpoints with sequential image capture. For image sequences captured at a high frame rate, optical flow can be computed, which estimates the motion field from the image sequences, based on the spatial and temporal variations of the image brightness. Using the local brightness constancy alone, the problem is under-constrained as the number of variables is twice the number of measurements. Therefore, it is augmented with additional global smoothness constraints, so that the motion field can be estimated by minimizing an energy function [23, 29]. 3D motion of the camera and the scene structure can then be recovered from the motion field.

2.2.2 Single View Approaches

In contrast to these two multiple-view approaches, 3D shape can be inferred from a single viewpoint using information sources (cues) such as shading, texture and focus. Not surprisingly, these techniques are called *shape from shading*, *shape from texture* and *shape from focus* respectively.

Shading on a surface can provide information about local surface orientations and overall surface shape, as illustrated in Fig. 2.2, where the technique in [24] has been used. Shape from shading [22] uses the shades in a grayscale image to infer the shape of the surfaces, based on the reflectance map which links image intensity with surface orientation. After the surface normals have been recovered at each pixel, they can be integrated into a depth map using regularized surface fitting. The computations involved are considerably more complicated than for multiple-view approaches. Moreover, various assumptions, such as uniform albedo, reflectance and known light source directions, need to be made and there are open issues with convergence to a solution. The survey in [65] reviews various techniques and provides some comparative results. The approach can be enhanced when lights shining from different directions can be turned on and off separately. This technique is

---

[1]http://www.videredesign.com.

Fig. 2.2 Examples of synthetic shape from shading images (*left column*) and corresponding shape from shading reconstruction (*right column*)

known as *photometric stereo* [61] and it takes two or more images of the scene from the same viewpoint but under different illuminations in order to estimate the surface normals.

The foreshortening of regular patterns depends on how the surface slants away from the camera viewing direction and provides another cue on the local surface orientation. Shape from texture [17] estimates the shape of the observed surface from the distortion of the texture created by the imaging process, as illustrated in Fig. 2.3. Therefore, this approach works only for images with texture surfaces and assumes the presence of a regular pattern. Shape from shading is combined with shape from texture in [60] where the two techniques can complement each other. While the texture components provide information in textured region, shading helps in the uniform region to provide detailed information on the surface shape.

Shape from focus [37, 41] estimates depth using two input images captured from the same viewpoint but at different camera depths of field. The degree of blur is a strong cue for object depth as it increases as the object moves away from the camera's focusing distance. The relative depth of the scene can be constructed from

Fig. 2.3 Examples of synthetic shape from texture images (**a**, **c**) and corresponding surface normal estimates (**b**, **d**). Figure courtesy of [17]

the image blur where the amount of defocus can be estimated by averaging the squared gradient in a region.

Single view metrology [13] allows shape recovery from a single perspective view of a scene given some geometric information determined from the image. By exploiting scene constraints such as orthogonality and parallelism, a vanishing line and a vanishing point in a single image can be determined. Relative measurements of shape can then be computed, which can be upgraded to absolute metric measurements if the dimensions of a reference object in the scene are known.

While 3D recovery from a single view is possible, such methods are often not practical in terms of either robustness or speed or both. Therefore, the most commonly used approaches are based on multiple views, which is the focus of this chapter. The first step to understanding such approaches is to understand how to model the image formation process in the cameras of a stereo rig. Then we need to know how to estimate the parameters of this model. Thus camera modeling and camera calibration are discussed in the following two main sections.

2.3 Camera Modeling

A camera is a device in which the 3D scene is projected down onto a 2D image. The most commonly used projection in computer vision is 3D perspective projection. Figure 2.4 illustrates perspective projection based on the pinhole camera model, where **C** is the position of the pinhole, termed the camera center or the *center of projection*. Recall that, although the real image plane is behind the camera center, it is common practice to employ a virtual image plane in front of the camera, so that the image is conveniently at the same orientation as the scene.

Clearly, from this figure, the path of imaged light is modeled by a ray that passes from a 3D world point **X** through the camera center. The intersection of this ray with the image plane defines where the image, $\mathbf{x}_c$, of the 3D scene point, **X**, lies. We can reverse this process and say that, for some point on the image plane, its corresponding scene point must lie somewhere along the ray connecting the center of projection, **C**, and that imaged point, $\mathbf{x}_c$. We refer to this as *back-projecting* an image point to an infinite ray that extends out into the scene. Since we do not know

Fig. 2.4 Projection based on a pinhole camera model where a 3D object is projected onto the image plane. Note that, although the real image plane is behind the camera center, it is common practice to employ a virtual image plane in front of the camera, so that the image is conveniently at the same orientation as the scene

how far along the ray the 3D scene point lies, explicit depth information is lost in the imaging process. This is the main source of geometric ambiguity in a single image and is the reason why we refer to the recovery of the depth information from stereo and other cues as *3D reconstruction*.

Before we embark on our development of a mathematical camera model, we need to digress briefly and introduce the concept of homogeneous coordinates (also called projective coordinates), which is the natural coordinate system of analytic projective geometry and hence has wide utility in geometric computer vision.

2.3.1 Homogeneous Coordinates

We are all familiar with expressing the position of some point in a plane using a pair of coordinates as $[x, y]^T$. In general for such systems, n coordinates are used to describe points in an n-dimensional space, $\mathbb{R}^n$. In analytic projective geometry, which deals with algebraic theories of points and lines, such points and lines are typically described by *homogeneous coordinates*, where $n + 1$ coordinates are used to describe points in an n-dimensional space. For example, a general point in a plane is described as $\mathbf{x} = [x_1, x_2, x_3]^T$, and the general equation of a line is given by $\mathbf{l}^T \mathbf{x} = 0$ where $\mathbf{l} = [l_1, l_2, l_3]^T$ are the homogeneous coordinates of the line.[2] Since the right hand side of this equation for a line is zero, it is an homogeneous equation, and any non-zero multiple of the point $\lambda[x_1, x_2, x_3]^T$ is the same point, similarly any non-zero multiple of the line's coordinates is the same line. The symmetry in this

---

[2]You may wish to compare $\mathbf{l}^T \mathbf{x} = 0$ to two well-known parameterizations of a line in the (x, y) plane, namely: $ax + by + c = 0$ and $y = mx + c$ and, in each case, write down homogeneous coordinates for the point $\mathbf{x}$ and the line $\mathbf{l}$.

equation is indicative of the fact that points and lines can be exchanged in many theories of projective geometry; such theories are termed *dual* theories. For example, the cross product of two lines, expressed in homogeneous coordinates, yields their intersecting point, and the cross-product of a pair of points gives the line between them.

Note that we can easily convert from homogeneous to inhomogeneous coordinates, simply by dividing through by the third element, thus $[x_1, x_2, x_3]^T$ maps to $[\frac{x_1}{x_3}, \frac{x_2}{x_3}]^T$. A key point about homogeneous coordinates is that they allow the relevant transformations in the imaging process to be represented as linear mappings, which of course are expressed as matrix-vector equations. However, although the mapping between homogeneous world coordinates of a point and homogeneous image coordinates is linear, the mapping from homogeneous to inhomogeneous coordinates is non-linear, due to the required division.

The use of homogeneous coordinates fits well with the relationship between image points and their associated back-projected rays into the scene space. Imagine a mathematical (virtual) image plane at a distance of one metric unit in front of the center of projection, as shown in Fig. 2.4. With the camera center, $\mathbf{C}$, the homogeneous coordinates $[x, y, 1]^T$ define a 3D scene ray as $[\lambda x, \lambda y, \lambda]^T$, where λ is the unknown distance ($\lambda > 0$) along the ray. Thus there is an intuitive link between the depth ambiguity associated with the 3D scene point and the equivalence of homogeneous coordinates up to an arbitrary non-zero scale factor.

Extending the idea of thinking of homogeneous image points as 3D rays, consider the cross product of two homogeneous points. This gives a direction that is the normal of the plane that contains the two rays. The line between the two image points is the intersection of this plane with the image plane. The dual of this is that the cross product of two lines in the image plane gives the intersection of their associated planes. This is a direction orthogonal to the normals of both of these planes and is the direction of the ray that defines the point of intersection of the two lines in the image plane. Note that any point with its third homogeneous element zero defines a ray parallel to the image plane and hence meets it at infinity. Such a point is termed a *point at infinity* and there is an infinite set of these points $[x_1, x_2, 0]^T$ that lie on the *line at infinity* $[0, 0, 1]^T$; Finally, note that the 3-tuple $[0, 0, 0]^T$ has no meaning and is undefined. For further reading on homogeneous coordinates and projective geometry, please see [21] and [12].

2.3.2 Perspective Projection Camera Model

We now return to the perspective projection (central projection) camera model and we note that it maps 3D world points in standard metric units into the pixel coordinates of an image sensor. It is convenient to think of this mapping as a cascade of three successive stages:

1. A 6 degree-of-freedom (DOF) rigid transformation consisting of a rotation, R (3 DOF), and translation, $\mathbf{t}$ (3 DOF), that maps points expressed in world coordinates to the same points expressed in camera centered coordinates.

2. A perspective projection from the 3D world to the 2D image plane.
3. A mapping from metric image coordinates to pixel coordinates.

We now discuss each of these projective mappings in turn.

2.3.2.1 Camera Modeling: The Coordinate Transformation

As shown in Fig. 2.4, the camera frame has its (X, Y) plane parallel to the image plane and Z is in the direction of the principal axis of the lens and encodes depth from the camera. Suppose that the camera center has *inhomogeneous* position $\tilde{\mathbf{C}}$ in the world frame[3] and the rotation of the camera frame is R_c relative to the world frame orientation. This means that we can express any *inhomogeneous* camera frame points as:

$$\tilde{\mathbf{X}}_c = \mathsf{R}_c^T (\tilde{\mathbf{X}} - \tilde{\mathbf{C}}) = \mathsf{R}\tilde{\mathbf{X}} + \mathbf{t}. \tag{2.1}$$

Here $\mathsf{R} = \mathsf{R}_c^T$ represents the rigid rotation and $\mathbf{t} = -\mathsf{R}_c^T \tilde{\mathbf{C}}$ represents the rigid translation that maps a scene point expressed in the world coordinate frame into a camera-centered coordinate frame. Equation (2.1) can be expressed as a projective mapping, namely one that is linear in homogeneous coordinates, to give:

$$\begin{bmatrix} X_c \\ Y_c \\ Z_c \\ 1 \end{bmatrix} = \begin{bmatrix} \mathsf{R} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}.$$

We denote P_r as the 4×4 homogeneous matrix representing the rigid coordinate transformation in the above equation.

2.3.2.2 Camera Modeling: Perspective Projection

Observing the similar triangles in the geometry of perspective imaging, we have

$$\frac{x_c}{f} = \frac{X_c}{Z_c}, \qquad \frac{y_c}{f} = \frac{Y_c}{Z_c}, \tag{2.2}$$

where (x_c, y_c) is the position (metric units) of a point in the camera's image plane and f is the distance (metric units) of the image plane to the camera center. (This is usually set to the focal length of the camera lens.) The two equations above can be written in linear form as:

$$Z_c \begin{bmatrix} x_c \\ y_c \\ 1 \end{bmatrix} = \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} X_c \\ Y_c \\ Z_c \\ 1 \end{bmatrix}.$$

---

[3]We use a tilde to differentiate n-tuple inhomogeneous coordinates from $(n + 1)$-tuple homogeneous coordinates.

We denote P_p as the 3×4 perspective projection matrix, defined by the value of f, in the above equation. If we consider an abstract image plane at $f = 1$, then points on this plane are termed *normalized image coordinates*[4] and from Eq. (2.2), these are given by

$$x_n = \frac{X_c}{Z_c}, \qquad y_n = \frac{Y_c}{Z_c}.$$

2.3.2.3 Camera Modeling: Image Sampling

Typically, the image on the image plane is sampled by an image sensor, such as a CCD or CMOS device, at the locations defined by an array of pixels. The final part of camera modeling defines how that array is positioned on the $[x_c, y_c]^T$ image plane, so that pixel coordinates can be generated. In general, pixels in an image sensor are not square and the number of pixels per unit distance varies between the x_c and y_c directions; we will call these scalings m_x and m_y. Note that pixel positions have their origin at the corner of the sensor and so the position of the principal point (where the principal axis intersects the image plane) is modeled with pixel coordinates $[x_0, y_0]^T$. Finally, many camera models also cater for any skew,[5] s, so that the mapping into pixels is given by:

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} m_x & s & x_0 \\ 0 & m_y & y_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_c \\ y_c \\ 1 \end{bmatrix}.$$

We denote P_c as the 3×3 projective matrix defined by the five parameters m_x, m_y, s, x_0 and y_0 in the above equation.

2.3.2.4 Camera Modeling: Concatenating the Projective Mappings

We can concatenate the three stages described in the three previous subsections to give

$$\lambda \mathbf{x} = \mathsf{P}_c \mathsf{P}_p \mathsf{P}_r \mathbf{X}$$

or simply

$$\lambda \mathbf{x} = \mathsf{P} \mathbf{X}, \tag{2.3}$$

where λ is non-zero and positive. We note the following points concerning the above equation

---

[4]We need to use a variety of image coordinate normalizations in this chapter. For simplicity, we will use the same subscript n, but it will be clear about how the normalization is achieved.

[5]Skew models a lack of orthogonality between the two image sensor sampling directions. For most imaging situations it is zero.

1. For any homogeneous image point scaled to $\lambda[x, \ y, \ 1]^T$, the scale λ is equal to the imaged point's depth in the camera centered frame ($\lambda = Z_c$).
2. Any non-zero scaling of the projection matrix $\lambda_P \mathsf{P}$ performs the same projection since, in Eq. (2.3), any non-zero scaling of homogeneous image coordinates is equivalent.
3. A camera with projection matrix P, or some non-zero scalar multiple of that, is informally referred to as *camera P* in the computer vision literature and, because of point 2 above, it is referred to as being defined *up to scale*.

The matrix P is a 3×4 projective camera matrix with the following structure:

$$\mathsf{P} = \mathsf{K}[\mathsf{R}|\mathbf{t}]. \tag{2.4}$$

The parameters within K are the camera's *intrinsic parameters*. These parameters are those combined from Sects. 2.3.2.2 and 2.3.2.3 above, so that:

$$\mathsf{K} = \begin{bmatrix} \alpha_x & s & x_0 \\ 0 & \alpha_y & y_0 \\ 0 & 0 & 1 \end{bmatrix},$$

where $\alpha_x = fm_x$ and $\alpha_y = fm_y$ represent the focal length in pixels in the x and y directions respectively. Together, the rotation and translation in Eq. (2.4) are termed the camera's *extrinsic parameters*. Since there are 5 DOF from intrinsic parameters and 6 DOF from extrinsic parameters, a camera projection matrix has only 11 DOF, not the full 12 of a general 3×4 matrix. This is also evident from the fact that we are dealing with homogeneous coordinates and so the overall scale of P does not matter.

By expanding Eq. (2.3), we have:

$$\lambda \underbrace{\begin{bmatrix} x \\ y \\ 1 \end{bmatrix}}_{\substack{\text{homogeneous} \\ \text{image} \\ \text{coordinates}}} = \underbrace{\begin{bmatrix} \alpha_x & s & x_0 \\ 0 & \alpha_y & y_0 \\ 0 & 0 & 1 \end{bmatrix}}_{\substack{\text{intrinsic} \\ \text{camera} \\ \text{parameters}}} \underbrace{\begin{bmatrix} r_{11} & r_{12} & r_{13} & t_x \\ r_{21} & r_{22} & r_{23} & t_y \\ r_{31} & r_{32} & r_{33} & t_z \end{bmatrix}}_{\substack{\text{extrinsic} \\ \text{camera} \\ \text{parameters}}} \underbrace{\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}}_{\substack{\text{homogeneous} \\ \text{world} \\ \text{coordinates}}}, \tag{2.5}$$

which indicates that both the intrinsic and extrinsic camera parameters are necessary to fully define a ray (metrically, not just in pixel units) in 3D space and hence make absolute measurements in multiple-view 3D reconstruction. Finally, we note that any non-zero scaling of scene homogeneous coordinates $[X, Y, Z, 1]^T$ in Eq. (2.5) gives the same image coordinates[6] which, for a single image, can be interpreted as ambiguity between the scene scale and the translation vector $\mathbf{t}$.

---

[6]The same homogeneous image coordinates *up to scale* or the same inhomogeneous image coordinates.

Fig. 2.5 Examples of radial distortion effects in lenses: (**a**) No distortion (**b**) Pincushion distortion (**c**) Barrel distortion (**d**) Fisheye distortion

2.3.3 Radial Distortion

Typical cameras have a lens distortion, which disrupts the assumed linear projective model. Thus a camera may not be accurately represented by the pinhole camera model that we have described, particularly if a low-cost lens or a wide field-of-view (short focal length) lens such as a fisheye lens is employed. Some examples of lens distortion effects are shown in Fig. 2.5. Note that the effect is non-linear and, if significant, it must be corrected so that the camera can again be modeled as a linear device. The estimation of the required distortion parameters to do this is often encompassed within a camera calibration procedure, which is described in Sect. 2.4. With reference to our previous three-stage development of a projective camera in Sect. 2.3.2, lens distortion occurs at the second stage, which is the 3D to 2D projection, and this distortion is sampled by the image sensor.

Detailed distortion models contain a large number of parameters that model both radial and tangential distortion [7]. However, radial distortion is the dominant factor and usually it is considered sufficiently accurate to model this distortion only, using a low-order polynomial such as:

$$\begin{bmatrix} x_{nd} \\ y_{nd} \end{bmatrix} = \begin{bmatrix} x_n \\ y_n \end{bmatrix} + \begin{bmatrix} x_n \\ y_n \end{bmatrix} (k_1 r^2 + k_2 r^4),$$

where $[x_n, y_n]^T$ is the undistorted image position (i.e. that obeys our linear projection model) in normalized coordinates, $[x_{nd}, y_{nd}]^T$ is the distorted image position in normalized coordinates, k_1 and k_2 are the unknown radial distortion parameters, and $r = \sqrt{x_n^2 + y_n^2}$. Assuming zero skew, we also have

$$\begin{bmatrix} x_d \\ y_d \end{bmatrix} = \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} (x - x_0) \\ (y - y_0) \end{bmatrix} (k_1 r^2 + k_2 r^4), \qquad (2.6)$$

where the distorted position $[x_d, y_d]^T$ is now expressed in pixel coordinates and $[x, y]^T$ are the usual pixel coordinates predicted by the linear pinhole model. Note that r is still defined in normalized image coordinates and so a non-unity aspect ratio ($m_x \neq m_y$) in the image sensor does not invalidate this equation. Also note that both Eq. (2.6) and Fig. 2.5 indicate that distortion increases away from the center of the image. In the barrel distortion, shown in Fig. 2.5(c), distortion correction requires that image points are moved slightly towards the center of the image, more so if they are near the edges of the image. Correction could be applied to the whole image, as

in dense stereo, or just a set of relevant features, such as extracted corner points. Clearly, the latter process is computationally cheaper.

Now that we have discussed the modeling of a camera's image formation process in detail, we now need to understand how to estimate the parameters within this model. This is the focus of the next section, which details camera calibration.

2.4 Camera Calibration

Camera calibration [8] is the process of finding the parameters of the camera that produced a given image of a scene. This includes both extrinsic parameters R, $\mathbf{t}$ and intrinsic parameters, comprising those within the matrix K and radial distortion parameters, k_1, k_2. Once the intrinsic and extrinsic camera parameters are known, we know the camera projection matrix P and, taking into account of any radial distortion present, we can back-project any image pixel to a 3D ray in space. Clearly, as the intrinsic camera calibration parameters are tied to the focal length, changing the zoom on the lens would make the calibration invalid. It is also worth noting that calibration is not always required. For example, we may be more interested in approximate shape, where we need to know what objects in a scene are co-planar, rather than their absolute 3D position measurements. However, for stereo systems at least, camera calibration is commonplace.

Generally, it is not possible for an end-user to get the required calibration information to the required accuracy from camera manufacturer's specifications and external measurement of the position of cameras in some frame. Hence some sort of camera calibration procedure is required, of which there are several different categories. The longest established of these is *photogrammetric calibration*, where calibration is performed using a scene object of precisely known physical dimensions. Typically, several images of a special 3D target, such as three orthogonal planes with calibration grids (chessboard patterns of black and white squares), are captured and precise known translations may be used [58]. Although this gives accurate calibration results, it lacks flexibility due to the need for precise scene knowledge.

At the other end of the spectrum is *self-calibration* (auto-calibration) [21, 35], where no calibration target is used. The correspondences across three images of the same rigid scene provide enough constraints to recover a set of camera parameters which allow 3D reconstruction up to a similarity transform. Although this approach is flexible, there are many parameters to estimate and reliable calibrations cannot always be obtained.

Between these two extremes are 'desktop' camera calibration approaches that use images of planar calibration grids, captured at several unknown positions and orientations (i.e. a single planar chessboard pattern is manually held at several random poses and calibration images are captured and stored). This gives a good compromise between the accuracy of photogrammetric calibration and the ease of use of self-calibration. A seminal example is given by Zhang [64].

Although there are a number of publicly available camera calibration packages on the web, such as the Caltech camera calibration toolbox for MATLAB [9] and in

the OpenCV computer vision library [40], a detailed study of at least one approach is essential to understand calibration in detail. We will use Zhang's work [64] as a seminal example and this approach consists of two main parts:

(1) A *basic* calibration that is based on linear least squares and hence has a closed-form solution. In the formulation of the linear problem, a set of 9 parameters needs to be estimated. These are rather complicated combinations of the camera's intrinsic parameters and the algebraic least squares minimization to determine them has no obvious geometric meaning. Once intrinsic parameters have been extracted from these estimated parameters, extrinsic parameters can be determined using the projective mapping (homography) associated with each calibration grid image.

(2) A *refined* calibration that is based on non-linear least squares and hence has an iterative solution. Here it is possible to formulate a least squares error between the observed (inhomogeneous) image positions of the calibration grid corners and the positions predicted by the current estimate of intrinsic and extrinsic camera parameters. This has a clear geometric interpretation, but the sum of squares function that we wish to minimize is non-linear in terms of the camera parameters. A standard approach to solving this kind of problem is the *Levenberg-Marquardt* (LM) algorithm, which employs gradient descent when it is far from a minimum and Gauss-Newton minimization when it gets close to a minimum. Since the LM algorithm is a very general procedure, it is straightforward to employ more complex camera models, such as those that include parameters for the radial distortion associated with the camera lens.

The iterative optimization in (2) above needs to be within the basin of convergence of the global minimum and so the linear method in (1) is used to determine an initial estimation of camera parameters. The raw data used as inputs to the process consists of the image corner positions, as detected by an automatic corner detector [18, 52], of all corners in all calibration images and the corresponding 2D world positions, $[X, \ Y]^T$, of the corners on the calibration grid. Typically, correspondences are established by manually clicking one or more detected image corners, and making a quick visual check that the imaged corners are matched correctly using overlaying graphics or text. A typical set of targets is shown in Fig. 2.6.

In the following four subsections we outline the theory and practice of camera calibration. The first subsection details the estimation of the planar projective mapping between a scene plane (calibration grid) and its image. The next two subsections closely follow Zhang [64] and detail the basic calibration and then the refined calibration, as outlined above. These subsections refer to the case of a single camera and so a final fourth subsection is used to describe the additional issues associated with the calibration of a stereo rig.

2.4.1 Estimation of a Scene-to-Image Planar Homography

A homography is a projective transformation (*projectivity*) that maps points to points and lines to lines. It is a highly useful imaging model when we view planar scenes,
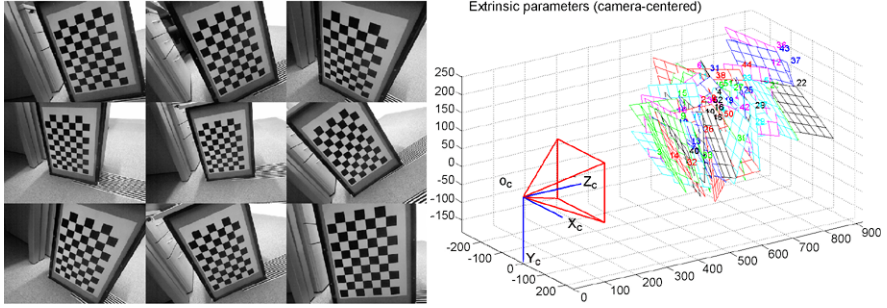
Fig. 2.6 *Left*: calibration targets used in a camera calibration process, image courtesy of Hao Sun. *Right*: after calibration, it is possible to determine the positions of the calibration planes using the estimated extrinsic parameters (Figure generated by the *Camera Calibration Toolbox for Matlab*, webpage maintained at Caltech by Jean-Yves Bouguet [9])

which is common in many computer vision processes, including the process of camera calibration.

Suppose that we view a planar scene, then we can define the (X, Y) axes of the world coordinate system to be within the plane of the scene and hence $Z = 0$ everywhere. Equation (2.5) indicates that, as far as a planar scene is concerned, the imaging process can be reduced to:

$$\lambda \mathbf{x} = \mathsf{K}[\mathbf{r}_1 \quad \mathbf{r}_2 \quad \mathbf{t}][X, Y, 1]^T,$$

where $\mathbf{r}_1$ and $\mathbf{r}_2$ are the first and second columns of the rotation matrix R, hence:

$$\lambda \mathbf{x} = \mathsf{H}[X, Y, 1]^T, \qquad \mathsf{H} = \mathsf{K}[\mathbf{r}_1 \quad \mathbf{r}_2 \quad \mathbf{t}]. \tag{2.7}$$

The 3×3 matrix H is termed a planar homography, which is defined up to a scale factor,[7] and hence has eight degrees of freedom instead of nine.

By expanding the above equation, we have:

$$\lambda \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} X \\ Y \\ 1 \end{bmatrix}. \tag{2.8}$$

If we map homogeneous coordinates to inhomogeneous coordinates, by dividing through by λ, this gives:

$$x = \frac{h_{11}X + h_{12}Y + h_{13}}{h_{31}X + h_{32}Y + h_{33}} \tag{2.9}$$

$$y = \frac{h_{21}X + h_{22}Y + h_{23}}{h_{31}X + h_{32}Y + h_{33}}. \tag{2.10}$$

---

[7]Due to the scale equivalence of homogeneous coordinates.

From a set of four correspondences in a general position,[8] we can formulate a set of eight linear equations in the eight unknowns of a homography matrix. This is because each correspondence provides a pair of constraints of the form given in Eqs. (2.9) and (2.10).

Rearranging terms in four pairs of those equations allows us to formulate the homography estimation problem in the form:

$$\mathsf{A}\mathbf{h} = \mathbf{0}, \tag{2.11}$$

where A is an 8×9 data matrix derived from image and world coordinates of corresponding points and $\mathbf{h}$ is the 9-vector containing the elements of the homography matrix. Since A has rank 8, it has a 1-dimensional null space, which provides a non-trivial (non-zero vector) solution for Eq. (2.11). This can be determined from a Singular Value Decomposition (SVD) of the data matrix, which generates three matrices $(\mathsf{U}, \mathsf{D}, \mathsf{V})$ such that $\mathsf{A} = \mathsf{U}\mathsf{D}\mathsf{V}^T$. Here, D is a diagonal matrix of singular values and U, V are orthonormal matrices. Typically, SVD algorithms order the singular values in descending order down the diagonal of D and so the required solution, corresponding to a singular value of zero, is extracted as the last column of V. Due to the homogeneous form of Eq. (2.11), the solution is determined up to a non-zero scale factor, which is acceptable because H is only defined up to scale. Often a unit scale is chosen (i.e. $\|\mathbf{h}\| = 1$) and this scaling is returned automatically in the columns of V.

In general, a larger number of correspondences than the minimum will not exactly satisfy the same homography because of image noise. In this case, a least squares solution to $\mathbf{h}$ can be determined in an over-determined system of linear equations. We follow the same procedure as above but this time the data matrix is of size $2n \times 9$ where $n > 4$ is the number of correspondences. When we apply SVD, we still select the last column of V corresponding to the smallest singular value in D. (Note that, in this case, the smallest singular value will be non-zero.)

Data normalization prior to the application of SVD is *essential* to give stable estimates [21]. The basic idea is to translate and scale both image and world coordinates to avoid orders of magnitude difference between the columns of the data matrix. Image points are translated so that their centroid is at the origin and scaled to give a root-mean-squared (RMS) distance of $\sqrt{2}$ from that origin, so that the 'average' image point has coordinates of unity magnitude. Scene points should be normalized in a similar way except that they should be scaled to give an RMS distance of $\sqrt{3}$.

When using homogeneous coordinates, the normalizations can be applied using matrix operators N_i, N_s, such that new *normalized* coordinates are given as:

$$\mathbf{x}_n = \mathsf{N}_i\mathbf{x}, \qquad \mathbf{X}_n = \mathsf{N}_s\mathbf{X}$$

for the image points and scene points respectively. Suppose that the homography computed from normalized coordinates is $\tilde{\mathsf{H}}$, then the homography relating the original coordinates of the correspondences is given as

$$\mathsf{H} = \mathsf{N}_i^{-1}\tilde{\mathsf{H}}\mathsf{N}_s.$$

---

[8]No three points collinear.

2.4.2 Basic Calibration

From the known planar scene target and the resulting image, a scene-to-image planar homography can be estimated as described in the previous subsection. Suppose that we describe such a homography as a set of 3×1 column vectors, i.e. $\mathsf{H} = [\mathbf{h}_1 \, \mathbf{h}_2 \, \mathbf{h}_3]$, then comparing this to Eq. (2.7) we have:

$$\lambda_H \mathbf{h}_1 = \mathsf{K}\mathbf{r}_1, \qquad \lambda_H \mathbf{h}_2 = \mathsf{K}\mathbf{r}_2, \tag{2.12}$$

where λ_H is a scale factor, accounting for the particular scale of an estimated homography. Noting that the columns of the rotation matrix, $\mathbf{r}_1, \mathbf{r}_2$ are orthonormal:

$$\mathbf{r}_1^T \mathbf{r}_2 = \mathbf{h}_1^T \mathsf{K}^{-T} \mathsf{K}^{-1} \mathbf{h}_2 = 0, \tag{2.13}$$

$$\mathbf{r}_1^T \mathbf{r}_1 = \mathbf{r}_2^T \mathbf{r}_2 \Rightarrow \mathbf{h}_1^T \mathsf{K}^{-T} \mathsf{K}^{-1} \mathbf{h}_1 = \mathbf{h}_2^T \mathsf{K}^{-T} \mathsf{K}^{-1} \mathbf{h}_2. \tag{2.14}$$

These equations provide one constraint each on the intrinsic parameters.

We construct a symmetric matrix B such that

$$\mathsf{B} = \mathsf{K}^{-T} \mathsf{K}^{-1} = \begin{bmatrix} B_{11} & B_{12} & B_{13} \\ B_{12} & B_{22} & B_{23} \\ B_{13} & B_{23} & B_{33} \end{bmatrix}.$$

Let the ith column vector of H be $\mathbf{h}_i = [h_{1i}, h_{2i}, h_{3i}]^T$, we have:

$$\mathbf{h}_i^T \mathsf{B} \mathbf{h}_j = \mathbf{v}_{ij}^T \mathbf{b},$$

where

$$\mathbf{v}_{ij} = [h_{1i} h_{1j}, h_{1i} h_{2j} + h_{2i} h_{1j}, h_{2i} h_{2j}, h_{3i} h_{1j} + h_{1i} h_{3j}, h_{3i} h_{2j} + h_{2i} h_{3j}, h_{3i} h_{3j}]^T$$

and $\mathbf{b}$ is the vector containing six independent entries of the symmetric matrix B:

$$\mathbf{b} = [B_{11}, B_{12}, B_{22}, B_{13}, B_{23}, B_{33}]^T.$$

Therefore, the two constraints in Eqs. (2.13) and (2.14) can be rewritten as:

$$\begin{bmatrix} \mathbf{v}_{12}^T \\ (\mathbf{v}_{11} - \mathbf{v}_{22})^T \end{bmatrix} \mathbf{b} = \mathbf{0}.$$

If n images of the planar calibration grid are observed, n sets of these equations can be stacked into a matrix-vector equation as:

$$\mathsf{V}\mathbf{b} = \mathbf{0},$$

where V is a $2n \times 6$ matrix. Although a minimum of three planar views allows us to solve for $\mathbf{b}$, it is recommended to take more and form a least squares solution. In this case, the solution for $\mathbf{b}$ is the eigenvector of $\mathsf{V}^T \mathsf{V}$ associated with the smallest eigenvalue. Once $\mathbf{b}$ is estimated, we know the matrix B up to some unknown scale factor, λ_B, and all of the intrinsic camera parameters can be computed by expanding the right hand side of $\mathsf{B} = \lambda_B \mathsf{K}^{-T} \mathsf{K}^{-1}$ in terms of its individual elements. Although this is somewhat laborious, it is straightforward algebra of simultaneous equations, where five intrinsic camera parameters plus one unknown scale factor can be derived from the six parameters of the symmetric matrix B. Zhang [64] presents the solution:

$$y_0 = \frac{(B_{12}B_{13} - B_{11}B_{23})}{(B_{11}B_{22} - B_{12}^2)}$$

$$\lambda_B = B_{33} - \frac{[B_{13}^2 + y_0(B_{12}B_{13} - B_{11}B_{23})]}{B_{11}}$$

$$\alpha_x = \sqrt{\frac{\lambda_B}{B_{11}}}$$

$$\alpha_y = \sqrt{\frac{\lambda_B B_{11}}{(B_{11}B_{22} - B_{12}^2)}}$$

$$s = \frac{-B_{12}\alpha_x^2\alpha_y}{\lambda_B}$$

$$x_0 = \frac{s y_0}{\alpha_y} - \frac{B_{13}\alpha_x^2}{\lambda_B}.$$

Once K is known, the extrinsic camera parameters for each image can be computed using Eq. (2.12):

$$\mathbf{r}_1 = \lambda_H \mathsf{K}^{-1}\mathbf{h}_1$$

$$\mathbf{r}_2 = \lambda_H \mathsf{K}^{-1}\mathbf{h}_2$$

$$\mathbf{r}_3 = \mathbf{r}_1 \times \mathbf{r}_2$$

$$\mathbf{t} = \lambda_H \mathsf{K}^{-1}\mathbf{h}_3,$$

where

$$\lambda_H = \frac{1}{\|\mathsf{K}^{-1}\mathbf{h}_1\|} = \frac{1}{\|\mathsf{K}^{-1}\mathbf{h}_2\|}.$$

The vectors $\mathbf{r}_1$, $\mathbf{r}_2$ will not be exactly orthogonal and so the estimated rotation matrix does not exactly represent a rotation. Zhang [64] suggests performing SVD on the estimated rotation matrix so that $\mathsf{U}\mathsf{S}\mathsf{V}^T = \mathsf{R}$. Then the closest pure rotation matrix in terms of Frobenius norm to that estimated is given as $\mathsf{R}' = \mathsf{U}\mathsf{V}^T$.

2.4.3 Refined Calibration

After computation of the linear solution described above, it can be iteratively refined via a non-linear least squares minimization using the *Levenberg-Marquardt* (LM) algorithm. As previously mentioned, the camera parameters can be extended at this stage to include an estimation for the lens distortion parameters, to give us the following minimization:

$$\hat{\mathbf{p}} = \min_{\mathbf{p}} \left\{ \sum_{i=1}^{n} \sum_{j=1}^{m} \left\| \mathbf{x}_{i,j} - \hat{\mathbf{x}}_{i,j}(\mathsf{K}, k_1, k_2, \mathsf{R}_i, \mathbf{t}_i, \mathbf{X}_j) \right\|^2 \right\},$$

where $\mathbf{x}_{i,j}$ is the image of world point $\mathbf{X}_j$ in image i and $\hat{\mathbf{x}}_{i,j}$ is the predicted projection of the same world point according to Eq. (2.7) (using estimated intrinsic and extrinsic camera parameters) followed by radial distortion according to Eq. (2.6).

The vector $\mathbf{p}$ contains all of the free parameters within the planar projection (homography) function plus two radial distortion parameters k_1 and k_2 as described in Sect. 2.3.3. Initial estimates of these radial distortion parameters can be set to zero. LM iteratively updates all parameters according to the equation:

$$\mathbf{p}_{k+1} = \mathbf{p}_k + \delta\mathbf{p}_k$$
$$\delta\mathbf{p}_k = -\left(\mathsf{J}^T\mathsf{J} + \lambda_J \operatorname{diag}\left(\mathsf{J}^T\mathsf{J}\right)\right)^{-1}\mathsf{J}^T\mathbf{e},$$

where J is the Jacobian matrix containing the first derivatives of the residual $\mathbf{e}$ with respect to each of the camera parameters.

Thus computation of the Jacobian is central to LM minimization. This can be done either numerically or with a custom routine, if analytical expressions for the Jacobian entries are known. In the numerical approach, each parameter is incremented and the function to be minimized (the least squares error function in this case) is computed and divided by the increment, which should be the maximum of 10^{-6} and $10^{-4} \times |p_i|$, where p_i is some current parameter value [21]. In the case of providing a custom Jacobian function, the expressions are long and complicated in the case of camera calibration, and so the use of a symbolic mathematics package can help reduce human error in constructing the partial differentials.

Note that there are LM implementations available on many platforms, for example in MATLAB's optimization toolbox, or the C/C++ `levmar` package. A detailed discussion of iterative estimation methods including LM is given in Appendix 6 of Hartley and Zisserman's book [21].

2.4.4 Calibration of a Stereo Rig

It is common practice to choose the optical center of one camera to be the origin of a stereo camera's 3D coordinate system. (The midpoint of the stereo baseline, which connects the two optical centers is also occasionally used.) Then, the relative rigid location of cameras, $[\mathsf{R},\ \mathbf{t}]$, within this frame, along with both sets of intrinsic parameters, is required to generate a pair of projection matrices and hence a pair of 3D rays from corresponding image points that intersect at their common scene point.

The previous two subsections show how we can calculate the intrinsic parameters for any single camera. If we have a stereo pair, which is our primary interest, then we would compute a pair of intrinsic parameter matrices, one for the left camera and one for the right. In most cases, the two cameras are the same model and hence we would expect the two intrinsic parameter matrices to be very similar.

Also, we note that, for each chessboard position, two sets of extrinsic parameters, $[\mathsf{R}, \mathbf{t}]$, are generated, one for the left camera's position relative to the calibration plane and one for the right. Clearly, each left-right pair of extrinsic parameters

should have approximately[9] the same relationship, which is due to the fixed rigid rotation and translation of one camera relative to another in the stereo rig.

Once two sets of intrinsic parameters and *one* set of extrinsic parameters encoding the *relative* rigid pose of one camera relative to another has been computed, the results are often refined in a global stereo optimization procedure, again using the Levenberg-Marquardt approach. To reduce n sets of relative extrinsic parameters to one set, we could choose the set associated with the closest calibration plane or compute some form of robust average.

All parameter estimates, both intrinsic and extrinsic, can be improved if the LM optimization is now performed over a minimal set of parameters, since the extrinsic parameters are reduced from 12 (two rotations and two translations) to 6 (one rotation and one translation) per calibration grid location. This approach ensures global rigidity of the stereo rig going from left to right camera. An implementation of global stereo optimization to refine stereo camera parameters is given in the Caltech camera calibration toolbox for MATLAB [9].

2.5 Two-View Geometry

3D reconstruction from an image pair must solve two problems: the correspondence problem and the reconstruction problem.

- *Correspondence problem.* For a point $\mathbf{x}$ in the left image, which is the corresponding point $\mathbf{x}'$ in the right image, where $\mathbf{x}$ and $\mathbf{x}'$ are images of the same physical scene point $\mathbf{X}$?
- *Reconstruction problem.* Given two corresponding points $\mathbf{x}$ and $\mathbf{x}'$, how do we compute the 3D coordinates of scene point $\mathbf{X}$?

Of these problems the correspondence problem is significantly more difficult as it is a search problem whereas, for a stereo camera of known calibration, reconstruction to recover the 3D measurements is a simple geometric mechanism. Since we have sets of three unique points, $(\mathbf{x}, \mathbf{x}', \mathbf{X})$, this mechanism is called triangulation (not to be confused with surface mesh triangulation, described in Chap. 4).

This section is designed to give the reader a good general grounding in two-view geometry and estimation of the key two-view geometric relations that can be useful even when extrinsic or intrinsic camera calibration information is not available.[10] As long as the concept of epipolar geometry is well understood, the remaining main sections of this chapter can be followed easily.

---

[9] 'Approximately', because of noise in the imaged corner positions supplied to the calibration process.

[10] Extrinsic parameters are always not known in a structure from motion problem, they are part of what we are trying to solve for. Intrinsic parameters may or may not be known, depending on the application.

2.5.1 Epipolar Geometry

Epipolar geometry establishes the relationship between two camera views. When we have calibrated cameras and we are dealing with metric image coordinates, it is dependent only on the relative pose between the cameras. When we have uncalibrated cameras and we are dealing with pixel-based image coordinates, it is additionally dependent on the cameras' intrinsic parameters, however, it is independent of the scene.

Once the epipolar geometry is known, for any image point in one image, we know that its corresponding point (its match) in the other image, must lie on a line, which is known as the *epipolar line* associated with the original point. This epipolar constraint greatly reduces the correspondence problem from a 2D search over the whole image to a 1D search along the epipolar line only, and hence reduces computational cost and ambiguities.

The discussion here is limited to two-view geometry only. A similar constraint called the *trifocal tensor* is applicable for three views, but is outside the scope of this chapter. For further information on the trifocal tensor and n-view geometries, please refer to [21].

As shown in Fig. 2.7(a), the image points $\mathbf{x}$ and $\mathbf{x}'$, world point $\mathbf{X}$ and the camera centers are co-planar and this plane is called the epipolar plane, which is shaded in the figure. If we only know $\mathbf{x}$, how is the corresponding point $\mathbf{x}'$ constrained? The line $\mathbf{l}'$ is the intersection of the epipolar plane with the second image plane. $\mathbf{l}'$ is called the epipolar line, which is the image in the second view of the ray back-projected from $\mathbf{x}$. As the point $\mathbf{x}'$ lies on $\mathbf{l}'$, the correspondences search does not need to cover the entire image but can be restricted only to the line $\mathbf{l}'$. In fact, if any point on epipolar line $\mathbf{l}$ has a corresponding point in the second image, it must lie on epipolar line $\mathbf{l}'$ and vice-versa. Thus $\mathbf{l}$ and $\mathbf{l}'$ are called *conjugate* epipolar lines.

The epipole is the point of intersection of the line joining the camera centers with the image plane. The epipole $\mathbf{e}$ is the projection of the second camera center on the first image, while the epipole $\mathbf{e}'$ is the projection of the first camera center on the second image.

In essence, two-view epipolar geometry describes the intersection of the image planes with the pencil of planes having the baseline as the pencil axis, as illustrated in Fig. 2.7(b). Note that the baseline is the line joining the two camera centers.[11] All epipolar lines intersect at the epipole of the respective image to give a pencil of epipolar lines in each image. Note that the epipoles are not necessarily within the boundaries of the image. A special case is when the cameras are oriented in the same direction and they are separated by a translation parallel to both image planes. In this case, the epipoles are at infinity and the epipolar lines are parallel. Furthermore, if the translation is in the X direction only and the cameras have the same intrinsic parameters, the conjugate epipolar lines lie on the same image rows. This is an ideal set up when we search for correspondences between the two images. However, we

---

[11]The length of the baseline is the magnitude of the extrinsic translation vector, $\mathbf{t}$.

Fig. 2.7 (**a**) The epipolar geometry establishes the relationship between the two camera views. (**b**) The epipolar planes rotate around the baseline and all epipolar lines intersect at the epipole

may prefer some camera vergence to improve the field-of-view overlap between the two cameras and, in this case, the images need to be warped so that the epipolar lines become horizontal again. This *rectification* process is discussed later in the chapter.

The epipolar constraint can be represented algebraically by a 3×3 matrix called the *fundamental matrix* (F), when we are dealing with raw pixel coordinates, and by the *essential matrix* (E) when the intrinsic parameters of the cameras are known and we are dealing with metrically expressed coordinates (e.g. millimeters) in the image plane.

### 2.5.2 Essential and Fundamental Matrices

Both the essential and fundamental matrices derive from a simple co-planarity constraint. For simplicity it is best to look at the epipolar relation using the essential matrix first and then adapt it using the camera intrinsic parameters to obtain a relation for pixel-based image coordinates, which involves the fundamental matrix.

Referring to Fig. 2.8, we have a world point $\mathbf{X}$ that projects to points $\mathbf{x}_c$ and $\mathbf{x}_c'$ in the image planes. These image plane points are expressed as 3-vectors, so that they are effectively the 3D positions of the imaged points expressed metrically in their own camera frame, hence the subscript c. (Note also that they can be regarded as normalized homogeneous image coordinates, with the scale set to the focal length, f, although any non-zero scale would suffice.) We know that the three vectors $\mathbf{Cx}_c$, $\mathbf{C}'\mathbf{x}_c'$ and $\mathbf{t}$ are co-planar, so we can choose one of the two camera frames to express this co-planarity, using the scalar triple product. If we choose the right frame (primed), then we must rotate vector $\mathbf{Cx}_c$ using rotation matrix R, to give:

$$\mathbf{x}_c'^T (\mathbf{t} \times \mathbf{R}\mathbf{x}_c) = 0.$$

Expressing the cross product with $\mathbf{t}$ by the multiplication with the skew-symmetric matrix $[\mathbf{t}]_x$, we have:

$$\mathbf{x}_c'[\mathbf{t}]_x \mathbf{R}\mathbf{x}_c = 0,$$

Fig. 2.8 The essential matrix $E = [\mathbf{t}]_x R$ encodes the epipolar geometry. It is used to relate the correspondences $\mathbf{x}_c$ and $\mathbf{x}'_c$ between two images, when these image locations are expressed in metric units. If pixel-based coordinates are used (for example, if intrinsic camera parameters are unknown) epipolar geometry is encoded by the fundamental matrix F

where

$$[\mathbf{t}]_x = \begin{bmatrix} 0 & -t_z & t_y \\ t_z & 0 & -t_x \\ -t_y & t_x & 0 \end{bmatrix}$$

and thus we have:

$$E = [\mathbf{t}]_x R \tag{2.15}$$

and

$$\mathbf{x}_c'^T E \mathbf{x}_c = 0. \tag{2.16}$$

Thus the essential matrix encapsulates only extrinsic parameters, namely, the rotation and translation associated with the *relative* pose of the two cameras. The implication of this is that, in applications where R and $\mathbf{t}$ have not been computed in a calibration procedure, they may be recoverable from an estimate of E, which will be discussed further in Sect. 2.8.2 in the context of structure from motion.

In many practical situations, we also need to deal with uncalibrated cameras where the intrinsic parameters are unknown (i.e. the mapping between metric image coordinates and raw pixel values is unknown). The shifting and scaling operations required for this conversion can be encapsulated in matrices K and K', as follows:

$$\mathbf{x} = K\mathbf{x}_c, \qquad \mathbf{x}' = K'\mathbf{x}'_c,$$

where K and K' are the 3×3 matrices containing the intrinsic camera parameters for the two cameras. Inserting these relations into Eq. (2.16) gives:

$$\mathbf{x}'^T K'^{-T} E K^{-1} \mathbf{x} = 0$$
$$\mathbf{x}'^T F \mathbf{x} = 0$$

thus

$$F = K'^{-T} E K^{-1} = K'^{-T} [\mathbf{t}]_x R K^{-1}$$

and we can see that the fundamental matrix encapsulates both intrinsic and extrinsic parameters. The interpretation of the epipolar constraint given by the fundamental matrix, is that, if points $\mathbf{x}$ and $\mathbf{x}'$ correspond, then $\mathbf{x}'$ must lie on the epipolar line given by $\mathbf{l}' = F\mathbf{x}$ and therefore the dot product between $\mathbf{x}'$ and $F\mathbf{x}$ is zero.

Some key properties of the fundamental matrix are summarized below:

- If F is the fundamental matrix between camera P and camera P', then F^T is the fundamental matrix between camera P' and camera P.
- F is a projective mapping taking a point to a line. If $\mathbf{l}$ and $\mathbf{l}'$ are corresponding (i.e. conjugate) epipolar lines, then any point $\mathbf{x}$ on $\mathbf{l}$ maps to the same line $\mathbf{l}'$. Hence, there is no inverse mapping (zero determinant, rank 2).
- F has seven degrees of freedom. While a 3×3 homogeneous matrix has eight independent ratios, there is also an additional constraint that the determinant of F is zero (F is rank 2), which further removes one degree of freedom.
- For any point $\mathbf{x}$ in the first image, the corresponding epipolar line in the second image is $\mathbf{l}' = F\mathbf{x}$. Similarly, $\mathbf{l} = F^T\mathbf{x}'$ represents the epipolar line in the first image corresponding to $\mathbf{x}'$ in the second image.
- The epipoles are determined as the left and right nullspaces of the fundamental matrix. This is evident, since each epipole is on every epipolar line in their respective image. This is written as $\mathbf{e}'^T\mathbf{l}' = \mathbf{e}'^T F\mathbf{x} = 0 \ \forall \mathbf{x}$, hence $\mathbf{e}'^T F = 0$. Similarly $\mathbf{l}^T\mathbf{e} = \mathbf{x}'^T F\mathbf{e} = 0 \ \forall \mathbf{x}'$, hence $F\mathbf{e} = 0$.
- The SVD (*Singular Value Decomposition*) of F is given as $F = U \operatorname{diag}(\sigma_1, \sigma_2, 0) V^T$ where $U = [\mathbf{u}_1, \mathbf{u}_2, \mathbf{e}']$, $V = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{e}]$. Thus finding the column in V that corresponds to the zero singular value gives a simple method of computation of the epipoles from the fundamental matrix.
- For cameras with some vergence (epipoles not at infinity) to give camera projection matrices: $P = K[I|0]$ and $P' = K'[R|\mathbf{t}]$, then we have: $F = K'^{-T}[\mathbf{t}]_\times RK^{-1} = [K'\mathbf{t}]_\times K'RK^{-1} = K'^{-T}RK^T[KR^T\mathbf{t}]_\times$ [21].

2.5.3 The Fundamental Matrix for Pure Translation

If the two identical cameras ($K = K'$) are separated by a pure translation ($R = I$), the fundamental matrix has a simple form, which can be shown to be [21]:

$$F = [K\mathbf{t}]_x = \left[\mathbf{e}'\right]_x = \begin{bmatrix} 0 & -e'_z & e'_y \\ e'_z & 0 & -e'_x \\ -e'_y & e'_x & 0 \end{bmatrix}.$$

In this case, the epipoles are at the same location in both images. If the translation is parallel to the image plane, the epipoles are at infinity with $e_z = e'_z = 0$ and the epipolar lines are parallel in both images. When discussing rectilinear stereo rigs and rectification later, we will be particularly interested in the case when the translation is parallel to the camera's x-axis, in which case the epipolar lines are

parallel and horizontal and thus correspond to image scan (raster) lines. In this case $\mathbf{e}' = [1, 0, 0]^T$ and the fundamental matrix is:

$$F = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}$$

and hence the relationship between corresponding points $\mathbf{x}$ and $\mathbf{x}'$ is given by $\mathbf{x}'^T F \mathbf{x} = 0$ which reduces to $y = y'$.

2.5.4 Computation of the Fundamental Matrix

As the fundamental matrix is expressed in terms of corresponding image points, F can be computed from image correspondences alone. No camera calibration information is needed and pixel coordinates are used directly. Note that there are degenerate cases in the estimation of F. These occur in two common and well-known instances: (i) when the relative pose between the two views can be described by a pure rotation and (ii) when the scene is planar. For now we consider scenarios where such degeneracies do not occur and we return to them later.

By expanding $\mathbf{x}'^T F \mathbf{x} = 0$ where $\mathbf{x} = [x, y, 1]^T$ and $\mathbf{x}' = [x', y', 1]^T$ and

$$F = \begin{bmatrix} f_{11} & f_{12} & f_{13} \\ f_{21} & f_{22} & f_{23} \\ f_{31} & f_{32} & f_{33} \end{bmatrix}$$

we obtain:

$$x'xf_{11} + x'yf_{12} + x'f_{13} + y'xf_{21} + y'yf_{22} + y'f_{23} + xf_{31} + yf_{32} + f_{33} = 0.$$

As each feature correspondence provides one equation, for n correspondences, we get the following set of linear equations:

$$\begin{bmatrix} x_1'x_1 & x_1'y_1 & x_1' & y_1'x_1 & y_1'y_1 & y_1' & x_1 & y_1 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_n'x_n & x_n'y_n & x_n' & y_n'x_n & y_n'y_n & y_n' & x_n & y_n & 1 \end{bmatrix} \begin{bmatrix} f_{11} \\ f_{12} \\ f_{13} \\ f_{21} \\ f_{22} \\ f_{23} \\ f_{31} \\ f_{32} \\ f_{33} \end{bmatrix} = \mathbf{0} \qquad (2.17)$$

or more compactly,

$$A\mathbf{f} = \mathbf{0},$$

where A is termed the data matrix and $\mathbf{f}$ is the vector of unknown elements of F.

The eight-point algorithm[12] [27] can be used as a very simple method to solve for F linearly using eight correspondences. As this is a homogeneous set of equations, $\mathbf{f}$ can only be determined up to a scale factor. With eight correspondences, Eq. (2.17) can be solved by linear methods, where the solution is the nullspace of A. (This can be found from the column in V that corresponds to the *zero* singular value in D in the singular value decomposition $A = UDV^T$.) However, a solution with a minimal set of correspondences is often inaccurate, particularly if the correspondences are not well spread over the images, or they may not provide enough strong constraints if some of them are near-collinear or co-planar. It is preferable to use more than eight correspondences, then the least squares solution for $\mathbf{f}$ is given by the singular vector corresponding to the *smallest* singular value of A.

Note that this approach is similar to that for determining the homography matrix, discussed earlier in Sect. 2.4.1. As with that approach, it is essential to normalize the pixel coordinates of each image before applying SVD [19, 21], using a mean-centering translation and a scaling so that the RMS distance of the points to the origin is $\sqrt{2}$. When using homogeneous coordinates, this normalization can be applied using matrix operators N, N', such that new *normalized* image coordinates are given as $\mathbf{x}_n = N\mathbf{x}$, $\mathbf{x}'_n = N'\mathbf{x}'$.

In general the solution for F_n (the subscript n now denotes that we have based the estimate on normalized image coordinates) will not have zero determinant (its rank will be 3 and not 2), which means that the epipolar lines will not intersect at a single point. In order to enforce this, we can apply SVD a second time, this time to the initially estimated fundamental matrix so that $F_n = UDV^T$. We then set the smallest singular value (in the third row and third column of D) to zero to produce matrix D' and update the estimate of the fundamental matrix as $F_n = UD'V^T$.

Of course, the estimate of F_n maps points to epipolar lines in the normalized image space. If we wish to search for correspondences within the original image space, we need to de-normalize the fundamental matrix estimate as $F = N'^T F_n N$.

Typically, there are many correspondences between a pair of images, including mostly inliers but also some outliers. This is inevitable, since matching is a local search and ambiguous matches exist, which will be discussed further in Sect. 2.7. Various robust methods for estimating the fundamental matrix, which address the highly corrupting effect of outliers, are compared in [55]. In order to compute F from these correspondences automatically, a common method is to use a robust statistics technique called Random Sample Consensus (RANSAC) [16], which we now outline:

1. Extract features in both images, for example, from a corner detector [18].
2. Perform feature matching between images (usually over a local area neighborhood) to obtain a set of potential matches or *putative correspondences*.
3. Repeat the following steps N times:

---

[12]There are several other approaches, such as the seven-point algorithm.

Table 2.1 Number of samples required to get at least one good sample with 99 % probability for various sample size s and outlier fraction ε

Sample size s	$\varepsilon = 10\,\%$	$\varepsilon = 20\,\%$	$\varepsilon = 30\,\%$	$\varepsilon = 40\,\%$	$\varepsilon = 50\,\%$
4	5	9	17	34	72
5	6	12	26	57	146
6	7	16	37	97	293
7	8	20	54	163	588
8	9	26	78	272	1177

- Select eight putative correspondences randomly.
- Compute F using these eight points, as described above.
- Find the number of inliers[13] that support F.

4. Find the F with the highest number of inliers (largest support) among the N trials.
5. Use this F to look for additional matches outside the search range used for the original set of putative correspondences.
6. Re-compute a least squares estimate of F using all inliers.

Note that re-computing F in the final step may change the set of inliers, as the epipolar lines are adjusted. Thus, a possible refinement is to iterate computation of a linear least squares estimate of F and its inliers, until a stable set of inliers is achieved or some maximum number of iterations is reached. The refinement achieved is often considered to be not worth the additional computational expense if processing time is considered important or if the estimate of F is to be used as the starting point for more advanced iterative non-linear refinement techniques, described later.

In the RANSAC approach, N is the number of trials (putative F computations) needed to get at least one good sample with a high probability (e.g. 99 %). How large should N be? The probability p of getting a good sample is given by:

$$p = 1 - \left(1 - (1 - \varepsilon)^s\right)^N,$$

where ε is the fraction of outliers (incorrect feature correspondences) and s is the number of correspondences selected for each trial. The above equation can be re-arranged as:

$$N = \frac{\log(1 - p)}{\log(1 - (1 - \varepsilon)^s)}. \tag{2.18}$$

The number of samples required for various sample size and outlier fraction based on Eq. (2.18) are shown in Table 2.1. It can be seen that the number of samples gets higher as the outlier fraction increases.

By repeatedly selecting a group of correspondences, the inlier support would be high for a correct hypothesis in which all the correspondences within the sample

---

[13] An inlier is a putative correspondence that lies within some threshold of its expected position predicted by F. In other words image points must lie within a threshold from their epipolar lines generated by F.

size, s, are correct. This allows the robust removal of outliers and the computation of F using inliers only. As the fraction of outliers may not be known in advance, an adaptive RANSAC method can be used where the number of outliers at each iteration is used to re-compute the total number of iterations required.

As the fundamental matrix has only seven degrees of freedom, a minimum of seven correspondences are required to compute F. When there are only seven correspondences, $\det(F) = 0$ constraint also needs to be imposed, resulting in a cubic equation to solve and hence may produce up to three solutions and all three must be tested for support. The advantage of using seven correspondences is that fewer trials are required to achieve the same probability of getting a good sample, as illustrated in Table 2.1.

Fundamental matrix refinement techniques are often based on the Levenberg-Marquardt algorithm, such that some non-linear cost function is minimized. For example a geometric cost function can be formulated as the sum of the squared distances between image points and the epipolar lines generated from their associated corresponding points and the estimate of F. This is averaged over both points in a correspondence and over all corresponding points (i.e. all those that agree with the estimate of F). The minimization can be expressed as:

$$F = \min_{F}\left(\frac{1}{N}\sum_{i=1}^{N}\left(d\left(\mathbf{x}'_i, F\mathbf{x}_i\right)^2 + d\left(\mathbf{x}_i, F^T\mathbf{x}'_i\right)^2\right)\right),$$

where $d(\mathbf{x}, \mathbf{l})$ is the distance of a point $\mathbf{x}$ to a line $\mathbf{l}$, expressed in pixels. For more details of this and other non-linear refinement schemes, the reader is referred to [21].

2.5.5 Two Views Separated by a Pure Rotation

If two views are separated by a pure rotation around the camera center, the baseline is zero, the epipolar plane is not defined and a useful fundamental matrix cannot be computed. In this case, the back-projected rays from each camera cannot form a triangulation to compute depth. This lack of depth information is intuitive because, under rotation, all points in the same direction move across the image in the same way, regardless of their depth. Furthermore, if the translation magnitude is small, the epipolar geometry is close to this degeneracy and computation of the fundamental matrix will be highly unstable.

In order to model the geometry of correspondences between two rotated views, a *homography*, described by a 3×3 matrix H, should be estimated instead. As described earlier, a homography is a projective transformation (*projectivity*) that maps points to points and lines to lines. For two identical cameras ($K = K'$), the scene-to-image projections are:

$$\mathbf{x} = K[I|\mathbf{0}]\mathbf{X}, \qquad \mathbf{x}' = K[R|\mathbf{0}]\mathbf{X}$$

hence

$$\mathbf{x}' = KRK^{-1}\mathbf{x} = H\mathbf{x}. \tag{2.19}$$

Fig. 2.9 The homography
induced by a plane π, where
a point **x** in the first image
can be transferred to the point
x$'$ in the second image

We can think of this homography as a mapping of image coordinates onto normalized coordinates (centered on the principal point at a unit metric distance from the camera). These points are rotated and then multiplying by K generates the image coordinates on the focal plane of the second, rotated camera.

2.5.6 Two Views of a Planar Scene

A homography should also be estimated for planar scenes where correspondences cannot uniquely define the epipolar geometry and hence the fundamental matrix. Similar to Eq. (2.7), the *2D-to-2D* projection of the world plane π in Fig. 2.9 to the left and right images are given by:

$$\lambda_x \mathbf{x} = \mathsf{H}_x \mathbf{X}, \qquad \lambda_{x'} \mathbf{x}' = \mathsf{H}_{x'} \mathbf{X},$$

where $\mathsf{H}_x, \mathsf{H}_{x'}$ are 3×3 homography matrices (homographies) and $\mathbf{x}, \mathbf{x}'$ are homogeneous image coordinates. The planar homographies form a group and hence we can form a composite homography as $\mathsf{H} = \mathsf{H}_{x'} \mathsf{H}_x^{-1}$ and it is straightforward to show that:

$$\lambda \mathbf{x}' = \mathsf{H}\mathbf{x}.$$

Figure 2.9 illustrates this mapping from **x** to **x**$'$ and we say that a homography is induced by the plane π. Homography estimation follows the same approach as was described in Sect. 2.4.1 for a scene-to-image planar homography (replacing **X** with **x** and **x** with **x**$'$ in Eqs. (2.8) to (2.10)).

Note that a minimum of four correspondences (no three points collinear in either image) are required because, for the homography, each correspondence generates a pair of constraints. Larger numbers of correspondences allow a least squares solution to an over-determined system of linear equations. Again suitable normalizations are required before SVD is applied to determine the homography.

A RANSAC-based technique can also be used to handle outliers, similar to the fundamental matrix estimation method described in Sect. 2.5.4. By repeatedly selecting the minimal set of four correspondences randomly to compute H and counting the number of inliers, the H with the largest number of inliers can be chosen.

Additional matches that are not in the original set of putative correspondences can be obtained using the best H. Then, H can be re-computed using all supporting matches in a linear least squares minimization using SVD.

Finally we note that, as in the case of the fundamental matrix, a non-linear optimization can be applied to refine the homography solution, if required by the application. The interested reader is referred to [21] for the details of the geometric cost function to be minimized.

2.6 Rectification

Typically, in a stereo rig, the cameras are horizontally displaced and rotated towards each other by an equal amount (verged), in order to overlap their fields of view. In this case, epipolar lines lie at a variety of angles across the two images, complicating the search for correspondences. In contrast, if these cameras had their principal axes parallel to each other (no vergence) and the two cameras had identical intrinsic parameters, conjugate (corresponding) epipolar lines would lie along the same horizontal scanline in each image, as observed in Sect. 2.5.3. This configuration is known as a standard rectilinear stereo rig. Clearly it is desirable to retain the improved stereo viewing volume associated with verged cameras and yet have the simplicity of correspondence search associated with a rectilinear rig.

To achieve this we can warp or *rectify* the raw images associated with the verged system such that corresponding epipolar lines become collinear and lie on the same scanline. A second advantage is that the equations for 3D reconstruction are very simply related to image disparity after image rectification, since they correspond to those of a simple rectilinear stereo rig. This triangulation computation is described later in the chapter.

Rectification can be achieved either with camera calibration information, for example in a typical stereo application, or without calibration information, for example in a typical structure from motion application. We discuss the calibrated case in the following subsection and give a brief mention of uncalibrated approaches in Sect. 2.6.2.

2.6.1 Rectification with Calibration Information

Here we assume a calibrated stereo rig, where we know both the intrinsic and the extrinsic parameters. Knowing this calibration information gives a simple rectification approach, where we find an image mapping that generates, from the original images, a pair of images that would have been obtained from a rectilinear rig. Of course, the field of view of each image is still bound by the real original cameras, and so the rectified images tend to be a different shape than the originals (e.g. slightly trapezoidal in a verged stereo rig).

Depending on the lenses used and the required accuracy of the application, it may be considered necessary to correct for radial distortion, using estimated parameters k_1 and k_2 from the calibration. To do the correction, we employ Eq. (2.6) in order to compute the unknown, undistorted pixel coordinates, $[x, y]^T$, from the known distorted coordinates, $[x_d, y_d]^T$. Of course, an iterative solution is required for this non-linear equation and the undistorted pixel coordinates can be initialized to the distorted coordinates at the start of this process.

Assuming some vergence, we wish to map the image points onto a pair of (virtual) image planes that are parallel to the baseline and in the same plane. Thus we can use the homography structure in Eq. (2.19) that warps images between a pair of rotated views. Given that we already know the intrinsic camera parameters, we need to determine the rotation matrices associated with the rectification of the left and right views. We will assume that the origin of the stereo system is at the optical center of the left camera and calibration information gives $[\mathsf{R}, \mathbf{t}]$ to define the rigid position of the right camera relative to this. To get the rotation matrix that we need to apply to image points of the left camera, we define the rectifying rotation matrix as:

$$\mathsf{R}_{rect} = \begin{bmatrix} \mathbf{r}_1^T \\ \mathbf{r}_2^T \\ \mathbf{r}_3^T \end{bmatrix},$$

where $\mathbf{r}_i$, $i = 1 \ldots 3$ are a set of mutually orthogonal unit vectors. The first of these is in the direction of the epipole or, equivalently, the direction of the translation to the right camera, $\mathbf{t}$. (This ensures that epipolar lines will be horizontal in the rectified image.) Hence the unit vector that we require is:

$$\mathbf{r}_1 = \frac{\mathbf{t}}{\|\mathbf{t}\|}.$$

The second vector $\mathbf{r}_2$ is orthogonal to the first and obtained as the cross product of $\mathbf{t}$ and the original left optical axis $[0, 0, 1]^T$ followed by a normalization to unit length to give:

$$\mathbf{r}_2 = \frac{1}{\sqrt{t_x^2 + t_y^2}} [-t_y, t_x, 0]^T.$$

The third vector is mutually orthogonal to the first two and so is computed using the cross product as $\mathbf{r}_3 = \mathbf{r}_1 \times \mathbf{r}_2$.

Given that the real right camera is rotated relative to the real left camera, we need to apply a rotation RR_{rect} to the image points of the right camera. Hence, applying homographies to left and right image points, using the form of Eq. (2.19), we have:

$$\mathbf{x}_{rect} = \mathsf{K}\mathsf{R}_{rect}\mathsf{K}^{-1}\mathbf{x}$$
$$\mathbf{x}'_{rect} = \mathsf{K}'\mathsf{R}\mathsf{R}_{rect}\mathsf{K}'^{-1}\mathbf{x}',$$

where K and K' are the 3×3 matrices containing the intrinsic camera parameters for the left and right cameras respectively. Note that, even with the same make and

Fig. 2.10 An image pair before rectification (**a**) and after rectification (**b**). The overlay shows that the corresponding *left* and *right* features lie on the same image row after rectification. Figure courtesy of [43]

model of camera, we may find that the focal lengths associated with K and K′ are slightly different. Thus we need to scale one rectified image by the ratio of focal lengths in order to place them on the same focal plane.

As the rectified coordinates are, in general, not integer, *resampling* using some form of interpolation is required. The rectification is often implemented in reverse, so that the pixel values in the new image plane can be computed as a bilinear interpolation of the four closest pixels values in the old image plane. Rectified images give a very simple triangulation reconstruction procedure, which is described later in Sect. 2.8.1.2.

2.6.2 Rectification Without Calibration Information

When calibration information is not available, rectification can be achieved using an estimate of the fundamental matrix, which is computed from correspondences within the raw image data. A common approach is to compute a pair of rectifying homographies for the left and right images [20, 33] so that the fundamental matrix associated with the rectified images is the same form as that for a standard rectilinear rig and the 'new cameras' have the same intrinsic camera parameters. Since such rectifying homographies map the epipoles to infinity ($[1, 0, 0]^T$), this approach fails when the epipole lies within the image. This situation is common in structure from motion problems, when the camera translates in the direction of its Z-axis. Several authors have tackled this problem by directly resampling the original images along their epipolar lines, which are specified by an estimated fundamental matrix. For example, the image is reparameterized using polar coordinates around the epipoles to reduce the search ambiguity to half epipolar lines [42, 43]. Figure 2.10 shows an example of an image pair before and after rectification for this scheme, where the corresponding left and right features lie on the same image row afterwards. Specialized rectifications exist, for example [10] which allows image matching over large forward translations of the camera although, in this scheme, rotations are not catered for.

2.7 Finding Correspondences

Finding correspondences is an essential step for 3D reconstruction from multiple views. The correspondence problem can be viewed as a search problem, which asks, given a pixel in the left image, which is the corresponding pixel in the right image? Of course there is something of a circular dependency here. We need to find correspondences to determine the epipolar geometry, yet we need the epipolar geometry to find (denser) correspondences in an efficient manner. The RANSAC sampling approach described earlier, showed us how to break into this loop. Once we have the epipolar geometry constraint, the search space is reduced from a 2D search to the epipolar line only.

The following assumptions underpin most methods for finding correspondences in image pairs. These assumptions hold when the distance of the world point from the cameras is much larger than the baseline.

- Most scene points are visible from both viewpoints.
- Corresponding image regions are similar.

Two questions are involved: what is a suitable image element to match and what is a good similarity measure to adopt? There are two main classes of correspondence algorithms: correlation-based and feature-based methods. Correlation-based methods recover dense correspondences where the element to match is an image window centered on some pixel and the similarity measure is the correlation between the windows. Feature-based methods typically establish sparse correspondences where the element to match is an image feature and the similarity measure is the distance between descriptors of the image features.

2.7.1 Correlation-Based Methods

If the element to match is only a single image pixel, ambiguous matches exist. Therefore, windows are used for matching in correlation-based methods and the similarity criterion is a measure of the correlation between the two windows. A larger window gives larger image context which can reduce the probability of ambiguities, but this has its own problems which will be discussed in Sect. 2.8.1.1. The selected correspondence is given by the window that maximizes a similarity criterion or minimizes a dissimilarity criterion within a search range. Once a match is found, the offset between the two windows can be computed, which is called the *disparity* from which the depth can be recovered. Some commonly used criteria for correlation-based methods are described next.

Based on the rectified images in Fig. 2.11, we define the window function, where m, an odd integer, is the image window size so that:

$$W_m(x, y) = \left\{ (u, v) \middle| x - \frac{(m-1)}{2} \leq u \leq x + \frac{(m-1)}{2}, \right.$$
$$\left. y - \frac{(m-1)}{2} \leq v \leq y + \frac{(m-1)}{2} \right\}. \tag{2.20}$$

Fig. 2.11 Correlation-based methods look for the matching image window between the *left* and *right* rectified images. An *m* by *m* window centering at the pixel is used for correlation (Raw image pair courtesy of the Middlebury Stereo Vision Page [34], originally sourced from Tsukuba University)

The dissimilarity can be measured by the *Sum of Squared Differences* (SSD) cost for instance, which is the intensity difference as a function of disparity d:

$$SSD(x, y, d) = \sum_{(u,v) \in W_m(x,y)} \left[I_l(u, v) - I_r(u - d, v) \right]^2,$$

where I_l and I_r refer to the intensities of the left and right images respectively.

If two image windows correspond to the same world object, the pixel values of the windows should be similar and hence the SSD value would be relatively small. As shown in Fig. 2.11, for each pixel in the left image, correlation-based methods would compare the SSD measure for pixels within a search range along the corresponding epipolar line in the right image. The disparity value that gives the lowest SSD value indicates the best match.

A slight variation of SSD is the *Sum of Absolute Differences* (SAD) where the absolute values of the differences are added instead of the squared values:

$$SAD(x, y, d) = \sum_{(u,v) \in W_m(x,y)} \left| I_l(u, v) - I_r(u - d, v) \right|.$$

This cost measure is less computationally expensive as it avoids the multiplication operation required for SSD. On the other hand, the SSD cost function penalizes the large intensity difference more due to the squaring operation.

The intensities between the two image windows may vary due to illumination changes and non-Lambertian reflection. Even if the two images are captured at the same time by two cameras with identical models, non-Lambertian reflection and differences in the gain and sensitivity can cause variation in the intensity. In these cases, SSD or SAD may not give a low value even for the correct matches. For these reasons, it is a good idea to normalize the pixels in each window. A first level of normalization would be to ensure that the intensities in each window are zero-mean. A second level of normalization would be to scale the zero-mean intensities so that they either have the same range or, preferably, unit variance. This can be

achieved by dividing each pixel intensity by the standard deviation of window pixel intensities, after the zero mean operation, i.e. normalized pixel intensities are given as:

$$I_n = \frac{I - \bar{I}}{\sigma_I},$$

where $\bar{I}$ is the mean intensity and σ_I is the standard deviation of window intensities. While SSD measures the dissimilarity and hence the smaller the better, *Normalized Cross-Correlation* (NCC) measures the similarity and hence, the larger the better. Again, the pixel values in the image window are normalized first by subtracting the average intensity of the window so that only the relative variation would be correlated. The NCC measure is computed as follows:

$$NCC(x, y, d) = \frac{\sum_{(u,v) \in W_m(x,y)} (I_l(u, v) - \overline{I_l})(I_r(u - d, v) - \overline{I_r})}{\sqrt{\sum_{(u,v) \in W_m(x,y)} (I_l(u, v) - \overline{I_l})^2 (I_r(u - d, v) - \overline{I_r})^2}},$$

where

$$\overline{I_l} = \frac{1}{m^2} \sum_{(u,v) \in W_m(x,y)} I_l(u, v), \qquad \overline{I_r} = \frac{1}{m^2} \sum_{(u,v) \in W_m(x,y)} I_r(u, v).$$

2.7.2 Feature-Based Methods

Rather than matching each pixel, feature-based methods only search for correspondences to a sparse set of features, such as those located by a repeatable, well-localized interest point detector (e.g. a corner detector). Apart from locating the features, feature extraction algorithms also compute some sort of feature descriptors for their representation, which can be used for the similarity criterion. The correct correspondence is given by the most similar feature pair, the one with the minimum distance between the feature descriptors.

Stable features are preferred in feature-based methods to facilitate matching between images. Typical examples of image features are edge points, lines and corners. For example, a feature descriptor for a line could contain the length, the orientation, coordinates of the midpoint or the average contrast along the edge line. A problem with linear features is that the matching can be poorly localized along the length of a line particularly if a linear feature is fragmented (imagine a smaller fragment from the left image sliding along a larger fragment from the right image). This is known as the aperture problem, referring to the fact that a *local* match 'looks through' a small aperture.

As a consequence, point-based features that are well-localized in two mutually orthogonal directions, have been preferred by researchers and practitioners in the field of computer vision. For example, the Harris corner detector [18] extracts points

Fig. 2.12 Wide baseline matching between two images with SIFT. The size and orientation of the squares correspond to the scale and orientation of the matching SIFT features

that differ as much as possible from neighboring points. This is achieved by looking for high curvatures in two mutually orthogonal directions, as the gradient is ill-defined in the neighborhood of corners. The corner strength or the grayscale values in a window region around each corner could be used as the descriptor. Another corner detector SUSAN [52] detects features based on the size, centroid and second moments of the local areas. As it does not compute image derivatives, it is robust to noise and does not require image smoothing.

Wide baseline matching refers to the situation where the two camera views differ considerably. Here, matching has to operate successfully over more difficult conditions, since there are larger geometric and photometric variations between the images.

In recent years, many interest point detection algorithms have been proposed that are scale invariant and viewpoint invariant to a certain extent which facilitates wide baseline matching. An interest point refers to an image feature that is stable under local and global perturbation and the local image structure is rich in terms of local image contents. These features are often described by a distinctive feature descriptor which is used as the similarity criterion. They can be used even when epipolar geometry is not yet known, as such distinctive descriptors allow correspondences to be searched over the whole image relatively efficiently.

For example, the *Scale Invariant Feature Transform* (SIFT) [28] and the *Speeded-Up Robust Feature* (SURF) [2] are two popular features which were developed for image feature generation in object recognition applications. The SIFT feature is described by a local image vector with 128 elements, which is invariant to image translation, scaling, rotation and partially invariant to illumination changes and affine or 3D projections.

Figure 2.12 shows an example of matching SIFT features across large baseline and viewpoint variation. It can be seen that most matches are correct, thanks to the invariance and discriminative nature of SIFT features.

Table 2.2 Different types of 3D reconstruction

| A priori knowledge	3D reconstruction
Intrinsic and extrinsic parameters	Absolute 3D reconstruction
Intrinsic parameters only	Metric 3D reconstruction (up to a scale factor)
No information	Projective 3D reconstruction

2.8 3D Reconstruction

Different types of 3D reconstruction can be obtained based on the amount of *a priori* knowledge available, as illustrated in Table 2.2. The simplest method to recover 3D information is stereo where the intrinsic and extrinsic parameters are known and the *absolute metric* 3D reconstruction can be obtained. This means we can determine the actual dimensions of structures, such as: *height of door* $= 1.93$ m.

For structure from motion, if no such prior information is available, only a projective 3D reconstruction can be obtained. This means that 3D structure is known only up to an arbitrary projective transformation so we know, for example, how many planar faces the object has and what point features are collinear, but we do not know anything about the scene dimensions and angular measurements within the scene. If intrinsic parameters are available, the projective 3D reconstruction can be upgraded to a metric reconstruction, where the 3D reconstruction is known up to a scale factor (i.e. a scaled version of the original scene). There is more detail to this hierarchy of reconstruction than we can present here (for example *affine* 3D reconstruction lies between the metric and projective reconstructions) and we refer the interested reader to [21].

2.8.1 Stereo

Stereo vision refers to the ability to infer information on the 3D structure and distance of a scene from two or more images taken from different viewpoints. The disparities of all the image points form the *disparity map*, which can be displayed as an image. If the stereo system is calibrated, the disparity map can be converted to a 3D point cloud representing the scene.

The discussion here focuses on binocular stereo for two image views only. Please refer to [51] for a survey of multiple-view stereo methods that reconstruct a complete 3D model instead of just a single disparity map, which generates *range image* information only. In such a 3D imaging scenario, there is at most one depth per image plane point, rear facing surfaces and other self-occlusions are not imaged and the data is sometimes referred to as 2.5D.

Fig. 2.13 A sample disparity map (**b**) obtained from the left image (**a**) and the right image (**c**). The disparity value for the pixel highlighted in *red* in the disparity map corresponds to the length of the line linking the matching features in the right image. Figure courtesy of [43]

2.8.1.1 Dense Stereo Matching

The aim of dense stereo matching is to compute disparity values for all the image points from which a dense 3D point cloud can be obtained. Correlation-based methods provide dense correspondences while feature-based methods only provide sparse correspondences. Dense stereo matching is more challenging than sparse correspondences as textureless regions do not provide information to distinguish the correct matches from the incorrect ones. The quality of correlation-based matching results depends highly on the amount of texture available in the images and the illumination conditions.

Figure 2.13 shows a sample disparity map after dense stereo matching. The disparity map is shown in the middle with disparity values encoded in grayscale level. The brighter pixels refer to larger disparities which mean the object is closer. For example, the ground pixels are brighter than the building pixels. An example of correspondences is highlighted in red in the figure. The pixel itself and the matching pixel are marked and linked on the right image. The length of the line corresponds to the disparity value highlighted in the disparity map.

Comparing image windows between two images could be ambiguous. Various matching constraints can be applied to help reduce the ambiguity, such as:

- Epipolar constraint
- Ordering constraint
- Uniqueness constraint
- Disparity range constraint

The epipolar constraint reduces the search from 2D to the epipolar line only, as has been described in Sect. 2.5. The ordering constraint means that if pixel b is to the right of a in the left image, then the correct correspondences a' and b' must also follow the same order (i.e. b' is to the right of a' in the right image). This constraint fails if there is occlusion.

The uniqueness constraint means that each pixel has at most one corresponding pixel. In general, there is a one-to-one correspondence for each pixel, but there is none in the case of occlusion or noisy pixels.

Fig. 2.14 The effect of window size on correlation-based methods: (**a**) input images (**b**) disparity map for a small correlation window (**c**) disparity map for a large correlation window (Raw image pair courtesy of the Middlebury Stereo Vision Page [34], originally sourced from Tsukuba University)

The disparity range constraint limits the disparity search range according to the prior information of the expected scene. Maximum disparity sets how close the object can be while the minimum disparity sets how far the object can be. Zero disparity refers to objects at infinity.

One important parameter for these correlation-based methods is the window size m in Eq. (2.20). While using a larger window size provides more intensity variation and hence more context for matching, this may cause problems around the occlusion area and at object boundaries, particularly for wide baseline matching.

Figure 2.14 shows the effect of window size on the resulting disparity map. The disparity map in the middle is for a window size of 3×3. It can be seen that, while it captures details well, it is very noisy, as the smaller window provides less information for matching. The disparity map on the right is for a window size of 15×15. It can be seen that while it looks very clean, the boundaries are not well-defined. Moreover, the use of a larger window size also increases the processing time as more pixels need to be correlated. The best window size is a trade-off between these two effects and is dependent on the level of fine detail in the scene.

For local methods, disparity computation at a given point depends on the intensity value within a local window only. The best matching window is indicated by the lowest dissimilarity measure or the highest similarity measure which uses information in the local region only. As pixels in an image are correlated (they may belong to the same object for instance), global methods could improve the stereo matching quality by making use of information outside the local window region.

Global methods perform optimization across the image and are often formulated as an energy minimization problem. Dynamic programming approaches [3, 5, 11] compute the minimum-cost path through the matrix of all pair-wise matching costs between two corresponding scanlines so that the best set of matches that satisfy the ordering constraint can be obtained. Dynamic programming utilizes information along each scanline independently, therefore, it may generate results that are not consistent across scanlines.

Graph cuts [6, 25] is one of the current state-of-the-art optimization techniques. These approaches make use of information across the whole image and produce high quality disparity maps. There is a trade-off between stereo matching quality

and the processing time. Global methods such as graph cuts, max flow [45], and belief propagation [53, 54] produce better disparity maps than local methods but they are very computationally intensive.

Apart from the algorithm itself, the processing time also depends on the image resolution, the window size and the disparity search range. The higher the image resolution, the more pixels need to be processed to produce the disparity map. The similarity measure needs to correlate more pixels for a larger window size. The disparity search range affects how many such measures need to be computed in order to find the correct match.

Hierarchical stereo matching methods have been proposed by down-sampling the original image into a pyramid [4, 44]. Dense stereo matching is first performed on the lowest resolution image and disparity ranges can be propagated back to the finer resolution image afterwards. This coarse-to-fine hierarchical approach allows fast computation to deal with a large disparity range, as a narrower disparity range can be used for the original image. Moreover, the more precise disparity search range helps to obtain better matches in the low texture areas.

The Middlebury webpage [34] provides standard datasets with ground truth information for researchers to benchmark their algorithms so that the performance of various algorithms can be evaluated and compared. A wide spectrum of dense stereo matching algorithms have been benchmarked, as illustrated in Fig. 2.15 [46]. Researchers can submit results of new algorithms which are ranked based on various metrics, such as RMS error between computed disparity map and ground truth map, percentage of bad matching pixels and so on. It can be observed from Fig. 2.15 that it is very difficult to understand algorithmic performance by qualitative inspection of disparity maps and the quantitative measures presented in [46] are required.

2.8.1.2 Triangulation

When the corresponding left and right image points are known, two rays from the camera centers through the left and right image points can be back-projected. The two rays and the stereo baseline lie on a plane (the epipolar plane) and form a triangle, hence the reconstruction is termed 'triangulation'. Here we describe triangulation for a rectilinear arrangement of two views or, equivalently, two rectified views.

After image rectification, the stereo geometry becomes quite simple as shown in Fig. 2.16, which shows the top-down view of a stereo system composed of two pinhole cameras. The necessary parameters, such as baseline and focal length, are obtained from the original stereo calibration. The following two equations can be obtained based on the geometry:

$$x'_c = f\frac{X}{Z}$$
$$x_c = f\frac{X+B}{Z},$$

True disparities	19 – Belief propagation	11 – GC + occlusions	20 – Layered stereo
10 – Graph cuts	*4 – Graph cuts	13 – Genetic algorithm	6 – Max flow
12 – Compact windows	9 – Cooperative alg.	15 – Stochastic diffusion	*2 – Dynamic progr.
14 – Realtime SAD	*3 – Scanline opt.	7 – Pixel-to-pixel stereo	*1 – SSD+MF

Fig. 2.15 Comparative disparity maps for the top fifteen dense stereo matching algorithms in [46] in decreasing order of performance. The *top left* disparity map is the ground truth. Performance here is measured as the percentage of bad matching pixels in regions where there are no occlusions. This varies from 1.15 % in algorithm 19 to 5.23 % in algorithm 1. Algorithms marked with a ∗ were implemented by the authors of [46], who present a wider range of algorithms in their publication. Figure courtesy of [46]

Fig. 2.16 The stereo geometry becomes quite simple after image rectification. The world coordinate frame is arbitrarily centered on the right camera. B is the stereo baseline and f is the focal length. Disparity is given by $d = x_c - x_c'$

where x'_c and x_c are the corresponding horizontal image coordinates (in metric units) in the right and left images respectively, f is the focal length and B is the baseline distance.

Disparity d is defined as the difference in horizontal image coordinates between the corresponding left and right image points, given by:

$$d = x_c - x'_c = \frac{fB}{Z}.$$

Therefore,

$$Z = \frac{fB}{d},$$

$$X = \frac{Zx'_c}{f}, \qquad Y = \frac{Zy'_c}{f}, \tag{2.21}$$

where y'_c is the vertical image coordinates in the right image.

This shows that the 3D world point can be computed once disparity is available: $(x'_c, y'_c, d) \mapsto (X, Y, Z)$. Disparity maps can be converted into depth maps using these equations to generate a 3D point cloud. It can be seen that triangulation is straightforward compared to the earlier stages of computing the two-view relations and finding correspondences.

Stereo matches are found by seeking the minimum of some cost functions across the disparity search range. This computes a set of disparity estimates in some discretized space, typically integer disparities, which may not be accurate enough for 3D recovery. 3D reconstruction using such quantized disparity maps leads to many thin layers of the scene. Interpolation can be applied to obtain sub-pixel disparity accuracy, such as fitting a curve to the SSD values for the neighboring pixels to find the peak of the curve, which provides more accurate 3D world coordinates.

By taking the derivatives of Eq. (2.21), the standard deviation of depth is given by:

$$\Delta Z = \frac{Z^2}{Bf} \Delta d,$$

where Δd is the standard deviation of the disparity. This equation shows that the depth uncertainty increases quadratically with depth. Therefore, stereo systems typically are operated within a limited range. If the object is far away, the depth estimation becomes more uncertain. The depth error can be reduced by increasing the baseline, focal length or image resolution. However, each of these has detrimental effects. For example, increasing the baseline makes matching harder and causes viewed objects to self-occlude, increasing the focal length reduces the depth of field, and increasing image resolution increases processing time and data bandwidth requirements. Thus, we can see that design of stereo cameras typically involves a range of performance trade-offs, where trade-offs are selected according to the application requirements.

Figure 2.17 compares the depth uncertainty for three stereo configuration assuming a disparity standard deviation of 0.1 pixel. A stereo camera with higher

Fig. 2.17 A plot illustrating the stereo uncertainty with regard to image resolution and baseline distance. A larger baseline and higher resolution provide better accuracy, but each of these has other costs

resolution (dashed line) provides better accuracy than the one with lower resolution (dotted line). A stereo camera with a wider baseline (solid line) provides better accuracy than the one with a shorter baseline (dashed line).

A quick and simple method to evaluate the accuracy of 3D reconstruction, is to place a highly textured planar target at various depths from the sensor, fit a least squares plane to the measurements and measure the residual RMS error. In many cases, this gives us a good measure of depth repeatability, unless there are significant systematic errors, for example from inaccurate calibration of stereo camera parameters. In this case, more sophisticated processes and ground truth measurement equipment are required. Capturing images of a target of known size and shape at various depths, such as a textured cube, can indicate how reconstruction performs when measuring in all three spatial dimensions.

2.8.2 Structure from Motion

Structure from motion (SfM) is the simultaneous recovery of 3D structure and camera relative pose (position and orientation) from image correspondences and it refers to the situation where images are captured by a moving camera. There are three sub-problems in structure from motion.

- Correspondence: which elements of an image frame correspond to which elements of the next frame.

- Ego-motion and reconstruction: determination of camera motion (sometimes called ego-motion) and structure of the observed world.
- Segmentation: extraction of regions corresponding to one or more moving objects.

The third sub-problem is a relatively recent problem in structure from motion, where some objects in the scene may have moved between frames. For dynamic scenes, features belonging to moving objects could be identified and removed as outliers. Alternatively one could consider an environment to contain an unknown number (n) of independently moving objects and a static environment as $n + 1$ SfM sub-problems, each having their own F matrix. However, for the following discussion, we assume that the scene is static, without any moving objects.

By matching features between frames, we obtain at least eight correspondences from which the fundamental matrix can be recovered as described in Sect. 2.5.4. Without camera calibration parameters, only the projective reconstruction can be obtained where orthogonal lines in the world may not be reconstructed as orthogonal. While this may be useful by itself, most practical applications require at least metric reconstruction where the reconstructed 3D model is a scaled version of the real scene.

Metric reconstruction requires camera intrinsic parameters which can be estimated from the images themselves using self-calibration (auto-calibration) techniques [21, 35] developed in recent years. Such methods exploit some prior information of the scene itself such as parallel lines, vanishing points and so on. For better accuracy and more robustness, the camera intrinsic parameters can be obtained with a calibration procedure using a known calibration grid, as discussed in Sect. 2.4.

Once the camera intrinsic parameters are known, the essential matrix E can be computed from the fundamental matrix. According to Eq. (2.15), the motion can be recovered from E, where $\mathbf{t}$ is determined up to a scale factor only (since we can multiply Eq. (2.16) by an arbitrary non-zero scale factor). The physical insight into this is that the same image disparity between a pair of views can occur for a point close to the camera positions and a point n-times the distance away with n-times the translation. Effectively we have scaled similar triangles in the triangulation-based reconstruction process.

SVD can be applied to extract $\mathbf{t}$ and R from E as follows [21]. Application of SVD gives the factorization $E = UDV^T$. By defining:

$$W = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \qquad Z = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

the solution is given by:

$$R = UWV^T \quad \text{or} \quad UW^TV^T$$
$$\mathbf{t} = \pm\mathbf{u}_3,$$

where $\mathbf{u}_3$ is the third column of matrix U. With two possible choices of R and $\mathbf{t}$, there are four possible solutions. Testing with a single point to determine if it is in

front of both cameras is sufficient to decide among the four different solutions. For further details, please refer to [21].

Once **t** (up to scale) and R have been extracted from E, the sparse scene structure can be recovered by computing the intersection between the back-projected rays. In general, due to measurement noise, these will not intersect in 3D space. The simplest solution is to compute the mid-point of the shortest perpendicular line between the two rays. However, a refined solution is to choose a reconstructed scene point **X**, such that it minimizes the sum of square errors between the actual image positions and their positions predicted by their respective camera projection matrices. The scene structure is only determined up to a scale factor but in some applications this could be constrained, for example, if some measurement is known in the scene, or the translation can be estimated from the wheel odometry of a mobile robot. In summary, this method first estimates the intrinsic camera parameters (or uses an existing calibration) after which the extrinsic camera parameters are recovered. Both the intrinsic and extrinsic camera parameters are then used to compute the scene structure.

Alternatively, *bundle adjustment*[14] offers a more accurate method that simultaneously optimizes the 3D structure and the 6-DOF camera pose (extrinsic camera parameters) for each view in an image sequence [57]. Sometimes the intrinsic camera parameters are also refined in the procedure. This is a batch process that iteratively refines the camera parameters and the 3D structure in order to minimize the sum of the reprojection errors. (A reprojection error is the Euclidean distance between an image feature and its reprojection into the image plane after computing the 3D world coordinate and the camera pose associated with that image point.) Since a specific reprojection error is only dependent on its own scene point and own viewpoint, the structure of the equations is sparse. Thus, even though bundle adjustment is thought to be fairly computationally expensive, exploitation of sparse linear algebra algorithms can significantly mitigate this. Such procedures are referred to as *sparse bundle adjustment*.

Using consecutive video frames gives poor 3D accuracy due to the very short baseline. An image pair formed by a larger time increment would provide better 3D information. However, if the time increment is too large, the camera could have moved significantly and it would be harder to establish correct correspondences. One possible solution to this is to track features over several short baseline frames using a small, local area-based search, before computing 3D from a pair of frames tracked over a significantly longer baseline.

2.9 Passive Multiple-View 3D Imaging Systems

Examples of passive multiple-view 3D imaging systems and their applications will now be presented, including stereo cameras, 3D modeling and mobile robot naviga-

---

[14]Bundle adjustment methods appeared several decades ago in the photogrammetry literature and are now used widely in the computer vision community.

tion. 3D modeling systems generate photo-realistic 3D models from sequences of images and have a wide range of applications. For mobile robot applications, passive multiple-view 3D imaging systems are used for localization, building maps and obstacle avoidance.

2.9.1 Stereo Cameras

Stereo cameras can be custom-built by mounting two individual cameras on a rigid platform separated by a fixed baseline. However, it is important that, for non-static scenes or for mobile platforms, the two cameras are synchronized so that they capture images at the same time. In order to obtain absolute 3D information, as discussed earlier in Table 2.2, the stereo camera needs to be calibrated to recover the intrinsic and extrinsic parameters. It is also critical that the relative camera pose does not change over time, otherwise, re-calibration would be required.

Commercial off-the-shelf (COTS) stereo vision systems have been emerging in recent years. These cameras often have a fixed baseline and are pre-calibrated by the vendor. Typically, they are nicely packaged and convenient to use and an example was given earlier in Fig. 2.1. The Point Grey Research Bumblebee camera[15] is another example, which comes pre-calibrated and an application programming interface (API) is provided to configure the camera and grab images, as well as rectify the images and perform dense stereo matching.

It is desirable to obtain disparity maps in real-time in many applications, for example obstacle detection for mobile robots. Hardware-accelerated correlation-based stereo systems are now commercially available, which can offer a high update rate required for mobile robot navigation, as well as to free up the processor for other tasks.

The Tyzx DeepSea G2 stereo vision system[16] provides real-time embedded 3D vision processing without the use of separate computer. The custom image processing chip (an *Application-Specific Integrated Circuit* or ASIC), a *Field Programmable Gate Array* (FPGA) and an embedded PowerPC are all enclosed in the self-contained camera package. Different baselines and lens options are available. Real-time 3D depth data can be obtained via an Ethernet connection. Figure 2.18 shows that the Tyzx system is used on a rugged military *Unmanned Ground Vehicle* (UGV) for obstacle detection [62].

Videre Design [59] offers fixed baseline and variable baseline stereo cameras, as well as a stereo camera with onboard processing. Their *stereo on a chip* (STOC) camera performs stereo processing onboard the camera and these are available with different fixed baselines. The fixed baseline cameras are pre-calibrated at the factory while the variable baseline cameras can be field-calibrated, offering flexibility for different range requirements.

---

[15]http://www.ptgrey.com/products/stereo.asp.

[16]http://www.tyzx.com/products/DeepSeaG2.html.

Fig. 2.18 A military UGV (Unmanned Ground Vehicle) equipped with the Tyzx DeepSea G2 stereo vision system [62]. Image courtesy of iRobot Corporation

Dense stereo matching can be highly parallelized, therefore such algorithms are highly suitable to run on graphics processing units (GPUs) to free up the CPU for other tasks. GPUs have a parallel throughput architecture that supports executing many concurrent threads, providing immense speed-up for highly parallelized algorithms. A dense stereo matching algorithm has been implemented on a commodity graphics card [63] to perform several hundred millions of disparity evaluations per second. This corresponds to 20 Hz for 512×512 image resolution with 32 disparity search range, therefore real-time performance can be achieved without the use of specialized hardware.

2.9.2 3D Modeling

The creation of photo-realistic 3D models of observed scenes has been an active research topic for many years. Such 3D models are very useful for both visualization and measurements in various applications such as planetary rovers, defense, mining, forensics, archeology and virtual reality.

Pollefeys et al. [43] and Nister [38] presented systems which create surface models from a sequence of images taken with a hand-held video camera. The camera motion is recovered by matching corner features in the image sequence. Dense stereo matching is carried out between the frames. The input images are used as surface texture to produce photo-realistic 3D models. These monocular approaches only output a scaled version of the original scene, but can be scaled with some prior information. Moreover, it requires a long processing time.

The objective of the DARPA Urbanscape project [36] is to develop a real-time data collection and processing system for the automatic geo-registered 3D reconstruction of urban scenes from video data. Multiple video streams as well as *Global Positioning System* (GPS) and *Inertial Navigation System* (INS) measurements are collected to reconstruct photo-realistic 3D models and place them in geo-registered coordinates. An example of a large-scale 3D reconstruction is shown in Fig. 2.19.

Fig. 2.19 An example of 3D modeling of urban scene from the Urbanscape project. Figure courtesy of [36]


Fig. 2.20 The user points the stereo camera freely at the scene of interest (*left*) and the photo-realistic 3D model of the scene is generated (*right*). Figure adapted from [47]

A stereo-camera based 3D vision system is capable of quickly generating calibrated photo-realistic 3D models of unknown environments. *Instant Scene Modeler* (iSM) can process stereo image sequences captured by an unconstrained hand-held stereo camera [47]. Dense stereo matching is performed to obtain 3D point clouds from each stereo pair. 3D point clouds from each stereo pair are merged together to obtain a color 3D point cloud. Furthermore, a surface triangular mesh is generated from the point cloud. This is followed by texture mapping, which involves mapping image textures to the mesh triangles. As adjacent triangles in the mesh may use different texture images, seamlines may appear unless texture blending is performed. The resulting photo-realistic 3D models can be visualized from different views and absolute measurements can be performed on the models. Figure 2.20 shows the user pointing the hand-held COTS stereo camera to freely scan the scene and the resulting photo-realistic 3D model, which is a textured triangular mesh.

For autonomous vehicles and planetary rovers, the creation of 3D terrain models of the environment is useful for visualization and path planning [1]. Moreover, the 3D modeling process achieves significant data compression, allowing the transfer of data as compact surface models instead of raw images. This is beneficial for plane-

Fig. 2.21 First image of a sequence captured by an autonomous rover in a desert in Nevada (*left*). Terrain model generated with virtual rover model inserted (*right*). Resulting terrain model and rover trajectory (*bottom*). Figure courtesy of [1]


Fig. 2.22 Mars Exploration Rover stereo image processing (*left*) and the reconstructed color 3D point cloud (*right*), with a virtual rover model inserted. Figure courtesy of [31]

tary rover exploration due to the limited bandwidth available. Figure 2.21 shows a photo-realistic 3D model created from a moving autonomous vehicle that traveled over 40 m in a desert in Nevada.

One of the key technologies required for planetary rover navigation is the ability to sense the nearby 3D terrain. Stereo cameras are suitable for planetary exploration thanks to their low power and low mass requirements and the lack of moving parts. The NASA *Mars Exploration Rovers* (MERs), named *Opportunity* and *Spirit*, both use passive stereo image processing to measure geometric information about the environment [31]. This is done by matching and triangulating pixels from a pair of rectified stereo images to generate a 3D point cloud. Figure 2.22 shows an example of the stereo images captured and the color 3D point cloud generated which represents the imaged terrain.

Fig. 2.23 3D model of a mock crime scene obtained with a hand-held stereo camera. Figure courtesy of [48]

Fig. 2.24 Underground mine 3D model (*left*) and consecutive 3D models as the mine advances (*right*). The *red* and *blue* lines on the *left* are geological features annotated by geologists to help with the ore body modeling. Figure courtesy of [48]

Documenting crime scenes is a tedious process that requires the investigators to record vast amounts of data by using video, still cameras and measuring devices, and by taking samples and recording observations. With passive 3D imaging systems, 3D models of the crime scene can be created quickly without much disturbance to the crime scene. The police can also perform additional measurements using the 3D model after the crime scene is released. The 3D model can potentially be shown in court so that the judge and the jury can understand the crime scene better. Figure 2.23 shows a 3D reconstruction of a mock crime scene generated from a hand-held stereo sequence within minutes after acquisition [48].

Photo-realistic 3D models are useful for survey and geology in underground mining. The mine map can be updated after each daily drill/blast/ore removal cycle to minimize any deviation from the plan. In addition, the 3D models can also allow the mining companies to monitor how much ore is taken at each blast. Figure 2.24 shows a photo-realistic 3D model of an underground mine face annotated with ge-

Fig. 2.25 3D reconstruction of a building on the ground using video (*left*) and using infra-red video (*right*) captured by an UAV (Unmanned Aerial Vehicle). Figure courtesy of [50]

ological features and consecutive 3D models of a mine tunnel created as the mine advances [48].

Airborne surveillance and reconnaissance are essential for successful military missions. *Unmanned Aerial Vehicles* (UAVs) are becoming the platform of choice for such surveillance operations and video cameras are among the most common sensors onboard UAVs. Photo-realistic 3D models can be generated from UAV video data to provide situational awareness as it is easier to understand the scene by visualizing it in 3D. The 3D model can be viewed from different perspectives and allow distance measurements and line-of-sight analysis. Figure 2.25 shows a 3D reconstruction of a building on the ground using video and infra-red video captured by an UAV [50]. The photo-realistic 3D models are geo-referenced and can be visualized in 3D *Geographical Information System* (GIS) viewers such as Google Earth.

2.9.3 Mobile Robot Localization and Mapping

Mobile robot localization and mapping is the process of simultaneously tracking the position of a mobile robot relative to its environment and building a map of the environment. Accurate localization is a prerequisite for building a good map and having an accurate map is essential for good localization. Therefore, *Simultaneous Localization and Mapping* (SLAM) is a critical underlying capability for successful mobile robot applications. To achieve a SLAM capability, high resolution passive vision systems can capture images in milliseconds, hence they are suitable for moving platforms such as mobile robots.

Stereo vision systems are commonly used on mobile robots, as they can measure the full six degrees of freedom (DOF) of the change in robot pose. This is known as visual odometry. By matching visual landmarks between frames to recover the robot motion, visual odometry is not affected by wheel slip and hence is more accurate than the wheel-based odometry. For outdoor robots with GPS receivers, visual odometry can also augment the GPS to provide better accuracy, and it is also valuable in environments where GPS signals are not available.
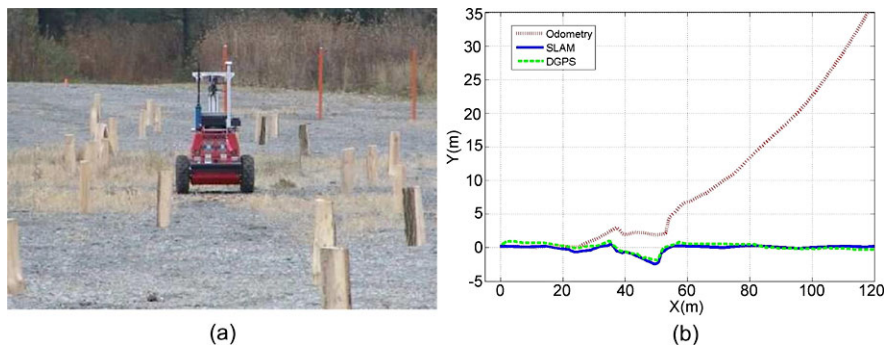
Fig. 2.26 (**a**) Autonomous rover on a gravel test site with obstacles (**b**) Comparison of the estimated path by SLAM, wheel odometry and DGPS (Differential GPS). Figure courtesy of [1]

Unlike in 3D modeling where correlation-based dense stereo matching is typically performed, feature-based matching is sufficient for visual odometry and SLAM; indeed, it is preferable for real-time robotics applications, as it is computationally less expensive. Such features are used for localization and a feature map is built at the same time.

The MERs *Opportunity* and *Spirit* are equipped with visual odometry capability [32]. An update to the rover's pose is computed by tracking the motion of autonomously-selected terrain features between two pairs of stereo images. It has demonstrated good performance and successfully detected slip ratios as high as 125 % even while driving on slopes as high as 31 degrees.

As SIFT features [28] are invariant to image translation, scaling, rotation, and fairly robust to illumination changes and affine or even mild projective deformation, they are suitable landmarks for robust SLAM. When the mobile robot moves around in an environment, landmarks are observed over time but from different angles, distances or under different illumination. SIFT features are extracted and matched between the stereo images to obtain 3D SIFT landmarks which are used for indoor SLAM [49] and for outdoor SLAM [1]. Figure 2.26 shows a field trial of an autonomous vehicle at a gravel test site with obstacles and a comparison of rover localization results. It can be seen that the vision-based SLAM trajectory is much better than the wheel odometry and matches well with the *Differential GPS* (DGPS).

Monocular visual SLAM applications have been emerging in recent years and these only require a single camera. The results are up to a scale factor, but can be scaled with some prior information. MonoSLAM [14] is a real-time algorithm which can recover the 3D trajectory of a monocular camera, moving rapidly through a previously unknown scene. The SLAM methodology is applied to the vision domain of a single camera, thereby achieving real-time and drift-free performance not offered by other structure from motion approaches.

Apart from localization, passive 3D imaging systems can also be used for obstacle/hazard detection in mobile robotics. Stereo cameras are often used as they can recover the 3D information without moving the robot. Figure 2.27 shows the stereo

Fig. 2.27 Examples of hazard detection using stereo images: a truck (*left*) and a person (*right*)

images and the hazard maps for a truck and a person respectively. Correlation-based matching is performed to generate a dense 3D point cloud. Clusters of point cloud that are above the ground plane are considered as hazards.

2.10 Passive Versus Active 3D Imaging Systems

Before concluding, we briefly compare passive multiple-view 3D imaging systems and their active imaging counterpart, as a bridge between this and the following chapter. Passive systems do not emit any illumination and only perceive the ambient light reflected from the scene. Typically this is reflected sunlight when outdoors, or the light reflected from standard room lighting when indoors. On the other hand, active systems include their own source of illumination, which has two main benefits:

- 3D structure can be determined in smooth, textureless regions. For passive stereo, it would be difficult to extract features and correspondences in such circumstances.
- The correspondence problem either disappears, for example a single spot of light may be projected at any one time, or is greatly simplified by controlling the structure of the projected light.

The geometric principle of determining depth from a light (or other EMR) projector (e.g. laser) and a camera is identical to the passive binocular stereo situation. The physical difference is that, instead of using triangulation applied to a pair of back-projected rays, we apply triangulation to the axis of the projected light and a single back-projected ray.

Compared with active approaches, passive systems are more computationally intensive as the 3D data is computed from processing the images and matching image features. Moreover, the depth data could be noisier as it relies on the natural texture in the scene and ambient lighting condition. Unlike active scanning systems such as laser scanners, cameras could capture complete images in milliseconds, hence they can be used as mobile sensors or operate in dynamic environments. The cost, size, mass and power requirements of cameras are generally lower than those of active sensors.

2.11 Concluding Remarks

One of the key challenges for 3D vision researchers is to develop algorithms to recover accurate 3D information robustly under a wide range of illumination conditions which can be done by humans so effortlessly. While 3D passive vision algorithms have been maturing over the years, this is still an active topic in the research community and at major computer vision conferences. Many algorithms perform reasonably well with test data but there are still challenges to handle scenes with uncontrolled illumination. Other open issues include efficient global dense stereo matching, multi-image matching and fully automated accurate 3D reconstruction from images.

Passive 3D imaging systems are becoming more prevalent as cameras are getting cheaper and computers are fast enough to handle the intensive processing requirements. Thanks to hardware acceleration and GPUs, real-time applications are more common, leading to a growing number of real-world applications.

After working through this chapter, you should be able to:

- Explain the fundamental concepts and challenges of passive 3D imaging systems.
- Explain the principles of epipolar geometry.
- Solve the correspondence problem by correlation-based and feature-based techniques (using off-the-shelf feature extractors).
- Estimate the fundamental matrix from correspondences.
- Perform dense stereo matching and compute a 3D point cloud.
- Explain the principles of structure from motion.
- Provide example applications of passive 3D imaging systems.

2.12 Further Reading

Two-view geometry is studied extensively in [21], which also covers the equivalent of epipolar geometry for three or more images. The eight-point algorithm was proposed in [19] to compute the fundamental matrix, while the five-point algorithm was proposed in [39] for calibrated cameras. Reference [57] provides a good tutorial and survey on bundle adjustment, which is also covered in textbooks [15, 21] and a recent survey article [35].

Surveys such as [46] serve as a guide to the extensive literature on stereo imaging. Structure from motion is extensively covered in review articles such as [35]. A step-by-step guide to 3D modeling from images is described in detail in [30]. Non-rigid structure from motion for dynamic scenes is discussed in [56].

Multiple-view 3D vision continues to be a highly active research topic and some of the major computer vision conferences include: the International Conference on Computer Vision (ICCV), IEEE Conference on Computer Vision and Pattern Recognition (CVPR) and the European Conference on Computer Vision (ECCV). Some of the relevant major journals include: International Journal of Computer

Vision (IJCV), IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) and Image and Vision Computing (IVC).

The International Society for Photogrammetry and Remote Sensing (ISPRS) proceedings and archives provide extensive literature on photogrammetry and related topics.

The following web sites provide comprehensive on-line resources for computer vision including 3D passive vision topics and are being updated regularly.

- CVonline (http://homepages.inf.ed.ac.uk/rbf/CVonline/) provides an on-line compendium of computer vision.
- VisionBib.Com (http://www.visionbib.com) contains annotated bibliography on a wide range of computer vision topics, as well as references to available datasets.
- Computer Vision online (http://www.computervisiononline.com) is a portal with links to software, hardware and datasets.
- OpenCV (http://opencv.willowgarage.com) is an open-source computer vision library.

2.13 Questions

1. What are the differences between passive and active 3D vision systems?
2. Name two approaches to recover 3D from single images and two approaches to recover 3D from multiple images.
3. What is the epipolar constraint and how can you use it to speed up the search for correspondences?
4. What are the differences between essential and fundamental matrices?
5. What is the purpose of rectification?
6. What are the differences between correlation-based and feature-based methods for finding correspondences?
7. What are the differences between local and global methods for dense stereo matching?
8. What are the differences between stereo and structure from motion?
9. What are the factors that affect the accuracy of stereo vision systems?

2.14 Exercises

Experimenting with stereo imaging requires that you have two images of a scene from slightly different viewpoints, with a good overlap between the views, and a significant number of well distributed corner features that can be matched. You will also need a corner detector. There are many stereo image pairs and corner detector implementations available on the web [40]. Of course, you can collect your own images either with a pre-packaged stereo camera or with a pair of standard digital cameras. The following programming exercises should be implemented in a language of your choice.

1. *Fundamental matrix with manual correspondences*. Run a corner detector on the image pair. Use a point-and-click GUI to manually label around 20 well distributed correspondences. Compute the fundamental matrix and plot the conjugate pair of epipolar lines on the images for each correspondence. Experiment with different numbers and combinations of correspondences, using a minimum of eight in the eight-point algorithm. Observe and comment on the sensitivity of the epipolar lines with respect to the set of correspondences chosen.

2. *Fundamental matrix estimation with outlier removal*. Add 4 incorrect corner correspondences to your list of 20 correct ones. Observe the effect on the computed fundamental matrix and the associated (corrupted) epipolar lines. Augment your implementation of fundamental matrix estimation with the RANSAC algorithm. Use a graphical overlay on your images to show that RANSAC has correctly identified the outliers, and verify that the fundamental matrix and its associated epipolar lines can now be computed without the corrupting effect of the outliers.

3. *Automatic feature correspondences*. Implement a function to *automatically* match corners between two images according to the *Sum of Squared Differences* (SSD) measure. Also, implement a function for the *Normalized Cross-Correlation* (NCC) measure. Compare the matching results with test images of similar brightness and also of different brightness.

4. *Fundamental matrix from automatic correspondences*. Use your fundamental matrix computation (with RANSAC) with the automatic feature correspondences. Determine the positions of the epipoles and, again, plot the epipolar lines.

The following additional exercises require the use of a stereo rig, which could be a pre-packaged stereo pair or a home-made rig with a pair of standard digital cameras. The cameras should have a small amount of vergence to overlap their fields of view.

5. *Calibration*. Create your own calibration target by printing off a chessboard pattern and pasting it to a flat piece of wood. Use a point-and-click GUI to semi-automate the corner correspondences between the calibration target and a set of captured calibration images. Implement a camera calibration procedure for a stereo pair to determine the intrinsic and extrinsic parameters of the stereo rig. If you have less time available you may choose to use some of the calibration libraries available on the web [9, 40].

6. *Rectification*. Compute an image warping (homography) to apply to each image in the stereo image pair, such that conjugate epipolar lines are horizontal (parallel to the x-axis) and have the same y-coordinate. Plot a set of epipolar lines to check that this rectification is correct.

7. *Dense stereo matching*. Implement a function to perform local dense stereo matching between left and right rectified images, using NCC as the similarity measure, and hence generate a disparity map for the stereo pair. Capture stereo images for a selection of scenes with varying amounts of texture within them and at varying distances from the cameras, and compare their disparity maps.

8. *3D reconstruction*. Implement a function to perform a 3D reconstruction from your disparity maps and camera calibration information. Use a graphics tool to visualize the reconstructions. Comment on the performance of the reconstructions for different scenes and for different distances from the stereo rig.

References

1. Barfoot, T., Se, S., Jasiobedzki, P.: Vision-based localization and terrain modelling for planetary rovers. In: Howard, A., Tunstel, E. (eds.) Intelligence for Space Robotics, pp. 71–92. TSI Press, Albuquerque (2006)
2. Bay, H., Ess, A., Tuytelaars, T., Van Gool, L.: SURF: speeded up robust features. Comput. Vis. Image Underst. **110**(3), 346–359 (2008)
3. Belhumeur, P.N.: A Bayesian approach to binocular stereopsis. Int. J. Comput. Vis. **19**(3), 237–260 (1996)
4. Bergen, J.R., Anandan, P., Hanna, K.J., Hinogorani, R.: Hierarchical model-based motion estimation. In: European Conference on Computer Vision (ECCV), Italy, pp. 237–252 (1992)
5. Birchfield, S., Tomasi, C.: Depth discontinuities by pixel-to-pixel stereo. Int. J. Comput. Vis. **35**(3), 269–293 (1999)
6. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE Trans. Pattern Anal. Mach. Intell. **23**(11), 1222–1239 (2001)
7. Brown, D.C.: Decentering distortion of lenses. Photogramm. Eng. **32**(3), 444–462 (1966)
8. Brown, D.C.: Close-range camera calibration. Photogramm. Eng. **37**(8), 855–866 (1971)
9. Camera Calibration Toolbox for MATLAB: http://www.vision.caltech.edu/bouguetj/calib_doc/. Accessed 20th October 2011
10. Chen, Z., Pears, N.E., Liang, B.: Monocular obstacle detection using Reciprocal-Polar rectification. Image Vis. Comput. **24**(12), 1301–1312 (2006)
11. Cox, I.J., Hingorani, S.L., Rao, S.B., Maggs, B.M.: A maximum likelihood stereo algorithm. Comput. Vis. Image Underst. **63**(3), 542–567 (1996)
12. Coxeter, H.S.M.: Projective Geometry, 2nd edn. Springer, Berlin (2003)
13. Criminisi, A., Reid, I., Zisserman, A.: Single view metrology. Int. J. Comput. Vis. **40**(2), 123–148 (2000)
14. Davison, A., Reid, I., Molton, N., Stasse, O.: MonoSLAM: real-time single camera SLAM. IEEE Trans. Pattern Anal. Mach. Intell. **29**(6), 1052–1067 (2007)
15. Faugeras, O., Luong, Q.T.: The Geometry of Multiple Images. MIT Press, Cambridge (2001)
16. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM **24**(6), 381–395 (1981)
17. Garding, J.: Shape from texture for smooth curved surfaces in perspective projection. J. Math. Imaging Vis. **2**, 329–352 (1992)
18. Harris, C., Stephens, M.J.: A combined corner and edge detector. In: Alvey Vision Conference, pp. 147–152 (1988)
19. Harley, R.I.: In defense of the 8-point algorithm. IEEE Trans. Pattern Anal. Mach. Intell. **19**(6), 580–593 (1997)
20. Hartley, R.I.: Theory and practice of projective rectification. Int. J. Comput. Vis. **35**(2), 115–127 (1999)
21. Hartley, R.I., Zisserman, A.: Multiple View Geometry in Computer Vision, 2nd edn. Cambridge University Press, Cambridge (2004)
22. Horn, B.K.P., Brooks, M.J. (eds.): Shape from Shading. MIT Press, Cambridge (1989)
23. Horn, B.K.P., Schunck, B.G.: Determining optical flow. Artif. Intell. **17**, 185–203 (1981)
24. Huang, R., P, S.W.A.: A shape-from-shading framework for satisfying data-closeness and structure-preserving smoothness constraints. In: Proceedings of the British Machine Vision Conference (2009)
25. Kolmogorov, V., Zabih, R.: Computing visual correspondence with occlusions using graph cuts. In: International Conference on Computer Vision (ICCV), Vancouver, pp. 508–515 (2001)
26. Konolige, K.: Small vision system: hardware and implementation. In: Proc. Int. Symp. on Robotics Research, Hayama, Japan, pp. 111–116 (1997)
27. Longuet-Higgins, H.C.: A computer algorithm for re-constructing a scene from two projections. Nature **293**, 133–135 (1981)

28. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. Int. J. Comput. Vis. **60**(2), 91–110 (2004)
29. Lucas, B.D., Kanade, T.: An interactive image registration technique with an application in stereo vision. In: International Joint Conference on Artificial Intelligence (IJCAI), Vancouver, pp. 674–679 (1981)
30. Ma, Y., Soatto, S., Kosecka, J., Sastry, S.S.: An Invitation to 3-D Vision. Springer, New York (2003)
31. Maimone, M., Biesiadecki, J., Tunstel, E., Cheng, Y., Leger, C.: Surface navigation and mobility intelligence on the Mars Exploration Rovers. In: Howard, A., Tunstel, E. (eds.) Intelligence for Space Robotics, pp. 45–69. TSI Press, Albuquerque (2006)
32. Maimone, M., Cheng, Y., Matthies, L.: Two years of visual odometry on the mars exploration rovers. J. Field Robot. **24**(3), 169–186 (2007)
33. Mallon, J., Whelan, P.F.: Projective rectification from the fundamental matrix. In: Image and Vision Computing, pp. 643–650 (2005)
34. The Middlebury stereo vision page: http://vision.middlebury.edu/stereo/. Accessed 16th November 2011
35. Moons, T., Van Gool, L., Vergauwen, M.: 3D reconstruction from multiple images. Found. Trends Comput. Graph. Vis. **4**(4), 287–404 (2010)
36. Mordohai, P., et al.: Real-time video-based reconstruction of urban environments. In: International Workshop on 3D Virtual Reconstruction and Visualization of Complex Architectures (3D-ARCH), Zurich, Switzerland (2007)
37. Nayar, S.K., Nakagawa, Y.: Shape from focus. IEEE Trans. Pattern Anal. Mach. Intell. **16**(8), 824–831 (1994)
38. Nister, D.: Automatic passive recovery of 3D from images and video. In: International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT), Thessaloniki, Greece, pp. 438–445 (2004)
39. Nister, D.: An efficient solution to the five-point relative pose problem. IEEE Trans. Pattern Anal. Mach. Intell. **26**(6), 756–770 (2004)
40. Open source computer vision library: http://opencv.willowgarage.com/wiki/. Accessed 20th October 2011
41. Pentland, A.P.: A new sense for depth of field. IEEE Trans. Pattern Anal. Mach. Intell. **9**(4), 523–531 (1987)
42. Pollefeys, M., Koch, R., Van Gool, L.: A simple and efficient rectification method for general motion. In: International Conference on Compute Vision (ICCV), Kerkyra, Greece, pp. 496–501 (1999)
43. Pollefeys, M., Van Gool, L., Vergauwen, M., Verbiest, F., Cornelis, K., Tops, J., Koch, R.: Visual modeling with a hand-held camera. Int. J. Comput. Vis. **59**(3), 207–232 (2004)
44. Quam, L.H.: Hierarchical warp stereo. In: Image Understanding Workshop, New Orleans, pp. 149–155 (1984)
45. Roy, S., Cox, I.J.: A maximum-flow formulation of the N-camera stereo correspondence problem. In: International Conference on Computer Vision (ICCV), Bombay, pp. 492–499 (1998)
46. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. Int. J. Comput. Vis. **47**(1/2/3), 7–42 (2002)
47. Se, S., Jasiobedzki, P.: Photo-realistic 3D model reconstruction. In: IEEE International Conference on Robotics and Automation, Orlando, Florida, pp. 3076–3082 (2006)
48. Se, S., Jasiobedzki, P.: Stereo-vision based 3D modeling and localization for unmanned vehicles. Int. J. Intell. Control Syst. **13**(1), 47–58 (2008)
49. Se, S., Lowe, D., Little, J.: Mobile robot localization and mapping with uncertainty using scale-invariant visual landmarks. Int. J. Robot. Res. **21**(8), 735–758 (2002)
50. Se, S., Firoozfam, P., Goldstein, N., Dutkiewicz, M., Pace, P.: Automated UAV-based video exploitation for mapping and surveillance. In: International Society for Photogrammetry and Remote Sensing (ISPRS) Commission I Symposium, Calgary (2010)
51. Seitz, S., Curless, B., Diebel, J., Scharstein, D., Szeliski, R.: A comparison and evaluation of multi-view stereo reconstruction algorithms. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), New York, pp. 519–526 (2006)

52. Smith, S.M., Brady, J.M.: SUSAN—a new approach to low level image processing. Int. J. Comput. Vis. **23**(1), 45–78 (1997)
53. Sun, J., Zheng, N., Shum, H.: Stereo matching using belief propagation. IEEE Trans. Pattern Anal. Mach. Intell. **25**(7), 787–800 (2003)
54. Tappen, M.F., Freeman, W.T.: Comparison of graph cuts with belief propagation for stereo, using identical MRF parameters. In: International Conference on Computer Vision (ICCV), Nice, France, pp. 900–907 (2003)
55. Torr, P.H.S., Murray, D.: The development and comparison of robust methods for estimating the fundamental matrix. Int. J. Comput. Vis. **24**(3), 271–300 (1997)
56. Torresani, L., Hertzmann, A., Bregler, C.: Non-rigid structure-from-motion: estimating shape and motion with hierarchical priors. IEEE Trans. Pattern Anal. Mach. Intell. **30**(5), 878–892 (2008)
57. Triggs, B., McLauchlan, P.F., Hartley, R.I., Fitzigibbon, A.W.: Bundle adjustment—a modern synthesis. In: International Workshop on Vision Algorithms, Kerkyra, Greece, pp. 298–372 (1999)
58. Tsai, R.Y.: A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. IEEE J. Robot. Autom. **3**(4), 323–344 (1987)
59. Videre Design: http://www.videredesign.com/. Accessed 16th November 2011
60. White, R., Forsyth, D.A.: Combining cues: shape from shading and texture. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), New York, pp. 1809–1816 (2006)
61. Woodham, R.J.: Analysing images of curved surfaces. Artif. Intell. **17**, 117–140 (1981)
62. Yamauchi, B.: Autonomous urban reconnaissance using man-portable UGVs. In: Unmanned Systems Technology. Orlando, Florida, SPIE, vol. 6230 (2006)
63. Yang, R., Pollefeys, M.: A versatile stereo implementation on commodity graphics hardware. Real-Time Imaging **11**(1), 7–18 (2005)
64. Zhang, Z.: A flexible new technique for camera calibration. IEEE Trans. Pattern Anal. Mach. Intell. **22**(11), 1330–1334 (2000)
65. Zhang, R., Tsai, P.S., Cryer, J.E., Shah, M.: Shaping from shading: a survey. IEEE Trans. Pattern Anal. Mach. Intell. **21**(8), 690–706 (1999)