

Image Alignment and Stitching: A Tutorial

Richard Szeliski¹

¹ *Microsoft Research, USA, szeliski@microsoft.com*

Abstract

This tutorial reviews image alignment and image stitching algorithms. Image alignment algorithms can discover the correspondence relationships among images with varying degrees of overlap. They are ideally suited for applications such as video stabilization, summarization, and the creation of panoramic mosaics. Image stitching algorithms take the alignment estimates produced by such registration algorithms and blend the images in a seamless manner, taking care to deal with potential problems such as blurring or ghosting caused by parallax and scene movement as well as varying image exposures. This tutorial reviews the basic motion models underlying alignment and stitching algorithms, describes effective direct (pixel-based) and feature-based alignment algorithms, and describes blending algorithms used to produce seamless mosaics. It ends with a discussion of open research problems in the area.

1

Introduction

Algorithms for aligning images and stitching them into seamless photo-mosaics are among the oldest and most widely used in computer vision. Frame-rate image alignment is used in every camcorder that has an “image stabilization” feature. Image stitching algorithms create the high-resolution photo-mosaics used to produce today’s digital maps and satellite photos. They also come bundled with most digital cameras currently being sold, and can be used to create beautiful ultra wide-angle panoramas.

An early example of a widely used image registration algorithm is the patch-based translational alignment (optical flow) technique developed by Lucas and Kanade [123]. Variants of this algorithm are used in almost all motion-compensated video compression schemes such as MPEG and H.263 [113]. Similar parametric motion estimation algorithms have found a wide variety of applications, including video summarization [20,93,111,203], video stabilization [81], and video compression [95,114]. More sophisticated image registration algorithms have also been developed for medical imaging and remote sensing – see [29,71,226] for some previous surveys of image registration techniques.

2 Introduction

In the photogrammetry community, more manually intensive methods based on surveyed *ground control points* or manually registered *tie points* have long been used to register aerial photos into large-scale photo-mosaics [181]. One of the key advances in this community was the development of *bundle adjustment* algorithms that could simultaneously solve for the locations of all of the camera positions, thus yielding globally consistent solutions [207]. One of the recurring problems in creating photo-mosaics is the elimination of visible seams, for which a variety of techniques have been developed over the years [1, 50, 135, 136, 148].

In film photography, special cameras were developed at the turn of the century to take ultra wide-angle panoramas, often by exposing the film through a vertical slit as the camera rotated on its axis [131]. In the mid-1990s, image alignment techniques were started being applied to the construction of wide-angle seamless panoramas from regular hand-held cameras [43, 124, 193, 194]. More recent work in this area has addressed the need to compute globally consistent alignments [167, 178, 199], the removal of “ghosts” due to parallax and object movement [1, 50, 178, 210], and dealing with varying exposures [1, 116, 124, 210]. (A collection of some of these papers can be found in [19].) These techniques have spawned a large number of commercial stitching products [43, 168], for which reviews and comparison can be found on the Web.

While most of the above techniques work by directly minimizing pixel-to-pixel dissimilarities, a different class of algorithms works by extracting a sparse set of *features* and then matching these to each other [7, 30, 35, 38, 129, 227]. Feature-based approaches have the advantage of being more robust against scene movement and are potentially faster, if implemented the right way. Their biggest advantage, however, is the ability to “recognize panoramas,” i.e., to automatically discover the adjacency (overlap) relationships among an unordered set of images, which makes them ideally suited for fully automated stitching of panoramas taken by casual users [30].

What, then, are the essential problems in image alignment and stitching? For image alignment, we must first determine the appropriate mathematical model relating pixel coordinates in one image to pixel coordinates in another. Section 2 reviews these basic *motion*

models. Next, we must somehow estimate the correct alignments relating various pairs (or collections) of images. Section 3 discusses how *direct* pixel-to-pixel comparisons combined with gradient descent (and other optimization techniques) can be used to estimate these parameters. Section 4 discusses how distinctive *features* can be found in each image and then efficiently matched to rapidly establish correspondences between pairs of images. When multiple images exist in a panorama, techniques must be developed to compute a globally consistent set of alignments and to efficiently discover which images overlap one another. These issues are discussed in Section 5.

For image stitching, we must first choose a final compositing surface onto which to warp and place all of the aligned images (Section 6). We also need to develop algorithms to seamlessly blend overlapping images, even in the presence of parallax, lens distortion, scene motion, and exposure differences (Section 6). In the last section of this survey, additional applications of image stitching and open research problems were discussed.

2

Motion Models

Before we can register and align images, we need to establish the mathematical relationships that map pixel coordinates from one image to another. A variety of such *parametric motion models* are possible, from simple 2D transforms, to planar perspective models, 3D camera rotations, lens distortions, and the mapping to nonplanar (e.g., cylindrical) surfaces [194].

To facilitate working with images at different resolutions, we adopt a variant of the *normalized device coordinates* used in computer graphics [145, 216]. For a typical (rectangular) image or video frame, we let the pixel coordinates range from $[-1, 1]$ along the longer axis, and $[-a, a]$ along the shorter, where a is the inverse of the *aspect ratio*, as shown in Figure 2.1.¹ For an image with width W and height H , the equations mapping integer pixel coordinates $\bar{\mathbf{x}} = (\bar{x}, \bar{y})$ to normalized device coordinates $\mathbf{x} = (x, y)$ are

$$x = \frac{2\bar{x} - W}{S} \quad \text{and} \quad y = \frac{2\bar{y} - H}{S}, \quad \text{where} \quad S = \max(W, H). \quad (2.1)$$

¹In computer graphics, it is usual to have both axes range from $[-1, 1]$, but this requires the use of two different focal lengths for the vertical and horizontal dimensions, and makes it more awkward to handle mixed portrait and landscape mode images.




Fig. 2.1 Mapping from pixel coordinates to normalized device coordinates.

Note that if we work with images in a *pyramid*, we need to halve the S value after each decimation step rather than recomputing it from $\max(W, H)$, since the (W, H) values may get rounded off or truncated in an unpredictable manner. *Note that for the rest of this paper, we use normalized device coordinates when referring to pixel coordinates.*

2.1 2D (planar) Motions

Having defined our coordinate system, we can now describe how coordinates are transformed. The simplest transformations occur in the 2D plane and are illustrated in Figure 2.2.

Translation 2D translations can be written as $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ or

$$\mathbf{x}' = [\mathbf{I} \ \mathbf{t}] \tilde{\mathbf{x}}, \tag{2.2}$$

where \mathbf{I} is the (2×2) identity matrix and $\tilde{\mathbf{x}} = (x, y, 1)$ is the *homogeneous* or *projective* 2D coordinate.




Fig. 2.2 Basic set of 2D planar transformations.

Rotation + Translation This transformation is also known as *2D rigid body motion* or the *2D Euclidean transformation* (since Euclidean distances are preserved). It can be written as $\mathbf{x}' = \mathbf{R}\mathbf{x} + \mathbf{t}$ or

$$\mathbf{x}' = [\mathbf{R} \quad \mathbf{t}] \tilde{\mathbf{x}}, \quad (2.3)$$

where

$$\mathbf{R} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \quad (2.4)$$

is an orthonormal rotation matrix with $\mathbf{R}\mathbf{R}^T = \mathbf{I}$ and $|\mathbf{R}| = 1$.

Scaled Rotation Also known as the *similarity transform*, this transform can be expressed as $\mathbf{x}' = s\mathbf{R}\mathbf{x} + \mathbf{t}$, where s is an arbitrary scale factor. It can also be written as

$$\mathbf{x}' = [s\mathbf{R} \quad \mathbf{t}] \tilde{\mathbf{x}} = \begin{bmatrix} a & -b & t_x \\ b & a & t_y \end{bmatrix} \tilde{\mathbf{x}}, \quad (2.5)$$

where we no longer require that $a^2 + b^2 = 1$. The similarity transform preserves angles between lines.

Affine The affine transform is written as $\mathbf{x}' = \mathbf{A}\tilde{\mathbf{x}}$, where \mathbf{A} is an arbitrary 2×3 matrix, i.e.,

$$\mathbf{x}' = \begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \end{bmatrix} \tilde{\mathbf{x}}. \quad (2.6)$$

Parallel lines remain parallel under affine transformations.

Projective This transform, also known as a *perspective transform* or *homography*, operates on homogeneous coordinates $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{x}}'$,






$$\tilde{\mathbf{x}}' \sim \tilde{\mathbf{H}}\tilde{\mathbf{x}}, \quad (2.7)$$

where \sim denotes equality up to scale and $\tilde{\mathbf{H}}$ is an arbitrary 3×3 matrix. Note that $\tilde{\mathbf{H}}$ is itself homogeneous, i.e., it is only defined up to a scale. The resulting homogeneous coordinate $\tilde{\mathbf{x}}'$ must be normalized in order to obtain an inhomogeneous result \mathbf{x}' , i.e.,

$$x' = \frac{h_{00}x + h_{01}y + h_{02}}{h_{20}x + h_{21}y + h_{22}} \quad \text{and} \quad y' = \frac{h_{10}x + h_{11}y + h_{12}}{h_{20}x + h_{21}y + h_{22}}. \quad (2.8)$$

Perspective transformations preserve straight lines.

Table 2.1 Hierarchy of 2D coordinate transformations. The 2×3 matrices are extended with a third $[\mathbf{0}^T \ 1]$ row to form a full 3×3 matrix for homogeneous coordinate transformations.

Name	Matrix	Number of d.o.f.	Preserves	Icon
Translation	$[\mathbf{I} \mathbf{t}]_{2 \times 3}$	2	Orientation + ...	
Rigid (Euclidean)	$[\mathbf{R} \mathbf{t}]_{2 \times 3}$	3	Lengths + ...	
Similarity	$[s\mathbf{R} \mathbf{t}]_{2 \times 3}$	4	Angles + ...	
Affine	$[\mathbf{A}]_{2 \times 3}$	6	Parallelism + ...	
Projective	$[\hat{\mathbf{H}}]_{3 \times 3}$	8	Straight lines	

Hierarchy of 2D Transformations The preceding set of transformations are illustrated in Figure 2.2 and summarized in Table 2.1. The easiest way to think of these is as a set of (potentially restricted) 3×3 matrices operating on 2D homogeneous coordinate vectors. Hartley and Zisserman [86] contains a more detailed description of the hierarchy of 2D planar transformations.

The above transformations form a nested set of *groups*, i.e., they are closed under composition and have an inverse that is a member of the same group. Each (simpler) group is a subset of the more complex group below it.

2.2 3D Transformations

A similar nested hierarchy exists for 3D coordinate transformations that can be denoted using 4×4 transformation matrices, with 3D equivalents to translation, rigid body (Euclidean) and affine transformations, and homographies (sometimes called collineations) [86].

The process of *central projection* maps 3D coordinates $\mathbf{p} = (X, Y, Z)$ to 2D coordinates $\mathbf{x} = (x, y, 1)$ through a *pinhole* at the camera origin onto a 2D projection plane a distance f along the z axis,

$$x = f \frac{X}{Z}, \quad y = f \frac{Y}{Z}, \quad (2.9)$$




Fig. 2.3 Central projection, showing the relationship between the 3D and 2D coordinates \mathbf{p} and \mathbf{x} , as well as the relationship between the focal length f and the field of view θ .

as shown in Figure 2.3. The relationship between the (unit-less) focal length f and the field of view θ is given by

$$f^{-1} = \tan \frac{\theta}{2} \quad \text{or} \quad \theta = 2 \tan^{-1} \frac{1}{f}. \quad (2.10)$$

To convert the focal length f to its more commonly used 35 mm equivalent, multiply the above number by 17.5 (the half-width of a 35 mm photo negative frame). To convert it to pixel coordinates, multiply it by $S/2$ (half-width for a landscape photo).

In the computer graphics literature, perspective projection is often written as a permutation matrix that permutes the last two elements of homogeneous 4-vector $\mathbf{p} = (X, Y, Z, 1)$,

$$\tilde{\mathbf{x}} \sim \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \mathbf{p}, \quad (2.11)$$

followed by a scaling and translation into screen and *z-buffer* coordinates.

In computer vision, it is traditional to drop the *z-buffer* values, since these cannot be sensed in an image and to write

$$\tilde{\mathbf{x}} \sim \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \mathbf{p} = [\mathbf{K} \mid \mathbf{0}] \mathbf{p}, \quad (2.12)$$

where $\mathbf{K} = \text{diag}(f, f, 1)$ is called the *intrinsic calibration* matrix.² This matrix can be replaced by a more general upper-triangular matrix \mathbf{K} that accounts for nonsquare pixels, skew, and a variable optic center location [86]. However, in practice, the simple focal length scaling used above provides high-quality results when stitching images from regular cameras.

In this paper, I prefer to use a 4×4 *projection matrix*, \mathbf{P} ,

$$\tilde{\mathbf{x}} \sim \left[\begin{array}{c|c} \mathbf{K} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{array} \right] \mathbf{p} = \mathbf{P}\mathbf{p}, \quad (2.13)$$

which maps the homogeneous 4-vector $\mathbf{p} = (X, Y, Z, 1)$ to a special kind of homogeneous *screen vector* $\tilde{\mathbf{x}} = (x, y, 1, d)$. This allows me to denote the upper-left 3×3 portion of the projection matrix \mathbf{P} as \mathbf{K} (making it compatible with the computer vision literature), while not dropping altogether the inverse screen depth information d (which is also sometimes called the *disparity* d [144]). This latter quantity is necessary to reason about mappings between images of a 3D scene, as described below.

What happens when we take two images of a 3D scene from different camera positions and/or orientations (Figure 2.4a)? A 3D point \mathbf{p} gets mapped to an image coordinate $\tilde{\mathbf{x}}_0$ in camera 0 through the




Fig. 2.4 A point is projected into two images: (a) relationship between the 3D point coordinate $(X, Y, Z, 1)$ and the 2D projected point $(x, y, 1, d)$; (b) planar homography induced by points all lying on a common plane $\hat{\mathbf{n}}_0 \cdot \mathbf{p} + c_0 = 0$.

²The last column of \mathbf{K} usually contains the optical center (c_x, c_y) , but this can be set to zero if we use normalized device coordinates.

combination of a 3D rigid-body (Euclidean) motion \mathbf{E}_0 ,

$$\mathbf{x}_0 = \begin{bmatrix} \mathbf{R}_0 & \mathbf{t}_0 \\ \mathbf{0}^T & 1 \end{bmatrix} \mathbf{p} = \mathbf{E}_0 \mathbf{p}, \quad (2.14)$$

and a perspective projection \mathbf{P}_0 ,

$$\tilde{\mathbf{x}}_0 \sim \mathbf{P}_0 \mathbf{E}_0 \mathbf{p}. \quad (2.15)$$

Assuming that we know the z -buffer value d_0 for a pixel in one image, we can map it back to the 3D coordinate \mathbf{p} using

$$\mathbf{p} \sim \mathbf{E}_0^{-1} \mathbf{P}_0^{-1} \tilde{\mathbf{x}}_0 \quad (2.16)$$

and then project it into another image yielding

$$\tilde{\mathbf{x}}_1 \sim \mathbf{P}_1 \mathbf{E}_1 \mathbf{p} = \mathbf{P}_1 \mathbf{E}_1 \mathbf{E}_0^{-1} \mathbf{P}_0^{-1} \tilde{\mathbf{x}}_0 = \mathbf{M}_{10} \tilde{\mathbf{x}}_0. \quad (2.17)$$

Unfortunately, we do not usually have access to the depth coordinates of pixels in a regular photographic image. However, for a *planar scene*, we can replace the last row of \mathbf{P}_0 in (2.13) with a general *plane equation*, $\hat{\mathbf{n}}_0 \cdot \mathbf{p} + c_0$ that maps points on the plane to $d_0 = 0$ values (Figure 2.4b). Then, if we set $d_0 = 0$, we can ignore the last column of \mathbf{M}_{10} in (2.17) and also its last row, since we do not care about the final z -buffer depth. The mapping equation (2.17) thus reduces to

$$\tilde{\mathbf{x}}_1 \sim \tilde{\mathbf{H}}_{10} \tilde{\mathbf{x}}_0, \quad (2.18)$$

where $\tilde{\mathbf{H}}_{10}$ is a general 3×3 homography matrix and $\tilde{\mathbf{x}}_1$ and $\tilde{\mathbf{x}}_0$ are now 2D homogeneous coordinates (i.e., 3-vectors) [193, 194].³ This justifies the use of the 8-parameter homography as a general alignment model for mosaics of planar scenes [124, 194].⁴

³For points off the reference plane, we get out-of-plane *parallax* motion, which is why this representation is often called the *plane plus parallax* representation [110, 165, 196].

⁴Note that for a single pair of images, the fact that a 3D plane is being viewed by a set of rigid cameras does not reduce the total number of degrees of freedom. However, for a large collection of images taken of a planar surface (e.g., a whiteboard) from a calibrated camera, we could reduce the number of degrees of freedom per image from 8 to 6 by assuming that the plane is at a canonical location (e.g., $z = 1$).




Fig. 2.5 Pure 3D camera rotation. The form of the homography (mapping) is particularly simple and depends only on the 3D rotation matrix and focal lengths.

Rotational Panoramas The more interesting case is when the camera undergoes pure rotation (Figure 2.5), which is equivalent to assuming all points are very far from the camera, i.e., on the *plane at infinity*. Setting $\mathbf{t}_0 = \mathbf{t}_1 = 0$, we get the simplified 3×3 homography

$$\tilde{\mathbf{H}}_{10} = \mathbf{K}_1 \mathbf{R}_1 \mathbf{R}_0^{-1} \mathbf{K}_0^{-1} = \mathbf{K}_1 \mathbf{R}_{10} \mathbf{K}_0^{-1}, \quad (2.19)$$

where $\mathbf{K}_k = \text{diag}(f_k, f_k, 1)$ is the simplified camera intrinsic matrix [194]. This can also be re-written as

$$\begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} \sim \begin{bmatrix} f_1 & & \\ & f_1 & \\ & & 1 \end{bmatrix} \mathbf{R}_{10} \begin{bmatrix} f_0^{-1} & & \\ & f_0^{-1} & \\ & & 1 \end{bmatrix} \begin{bmatrix} x_0 \\ y_0 \\ 1 \end{bmatrix} \quad (2.20)$$

or

$$\begin{bmatrix} x_1 \\ y_1 \\ f_1 \end{bmatrix} \sim \mathbf{R}_{10} \begin{bmatrix} x_0 \\ y_0 \\ f_0 \end{bmatrix}, \quad (2.21)$$

which reveals the simplicity of the mapping equations and makes all of the motion parameters explicit. Thus, instead of the general 8-parameter homography relating a pair of images, we get the 3-, 4-, or 5-parameter *3D rotation* motion models corresponding to the cases where the focal length f is known, fixed, or variable [199]. Estimating the 3D rotation matrix (and optionally, focal length) associated with

each image is intrinsically more stable than estimating a full 8-d.o.f. homography, which makes this the method of choice for large-scale image stitching algorithms [30, 178, 199].

Parametrizing 3D Rotations If we are going to represent panoramas using a combination of rotations and focal lengths, what is the best way to represent the rotations themselves? The choices include:

- the full 3×3 matrix \mathbf{R} , which has to be re-orthonormalized after each update;
- Euler angles (α, β, γ) , which are a bad idea as you cannot always move smoothly from one rotation to another;
- the axis/angle (or exponential twist) representation, which represents the rotation by an axis $\hat{\mathbf{n}}$ and a rotation angle θ , or the product of the two,

$$\vec{\omega} = \theta \hat{\mathbf{n}} = (\omega_x, \omega_y, \omega_z), \quad (2.22)$$

which has the minimal number of 3-parameters, but is still not unique;

- and unit quaternions, which represent rotations with unit 4-vectors,

$$\mathbf{q} = (x, y, z, w) = (\mathbf{v}, w) = \left(\sin \frac{\theta}{2} \hat{\mathbf{n}}, \cos \frac{\theta}{2} \right), \quad (2.23)$$

where $\hat{\mathbf{n}}$ and θ are the rotation axis and angle.

The rotation matrix corresponding to a rotation by θ around an axis $\hat{\mathbf{n}}$ is

$$\mathbf{R}(\hat{\mathbf{n}}, \theta) = \mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2, \quad (2.24)$$

which is known as *Rodriguez's formula* [5], and $[\hat{\mathbf{n}}]_{\times}$ is the matrix form of the cross-product operator,

$$[\hat{\mathbf{n}}]_{\times} = \begin{bmatrix} 0 & -\hat{n}_z & \hat{n}_y \\ \hat{n}_z & 0 & -\hat{n}_x \\ -\hat{n}_y & \hat{n}_x & 0 \end{bmatrix}. \quad (2.25)$$

For small (infinitesimal) rotations, the rotation reduces to

$$\mathbf{R}(\vec{\omega}) \approx \mathbf{I} + \theta[\hat{\mathbf{n}}]_{\times} = \mathbf{I} + [\vec{\omega}]_{\times}. \quad (2.26)$$

Using the trigonometric identities

$$\sin \theta = 2 \sin \frac{\theta}{2} \cos \frac{\theta}{2} = 2\|\mathbf{v}\|w$$

and

$$(1 - \cos \theta) = 2 \sin^2 \frac{\theta}{2} = 2\|\mathbf{v}\|^2,$$

Rodriguez's formula for a quaternion can be converted to

$$\begin{aligned} \mathbf{R}(\mathbf{q}) &= \mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2 \\ &= \mathbf{I} + 2w[\mathbf{v}]_{\times} + 2[\mathbf{v}]_{\times}^2. \end{aligned} \quad (2.27)$$

This suggests a quick way to rotate a vector by a quaternion using a series of cross products, scalings, and additions. From this, we can derive the commonly used formula for $\mathbf{R}(\mathbf{q})$ as a function of $\mathbf{q} = (x, y, z, w)$,

$$\mathbf{R}(\mathbf{q}) = \begin{bmatrix} 1 - 2(y^2 + z^2) & 2(xy - zw) & 2(xz + yw) \\ 2(xy + zw) & 1 - 2(x^2 + z^2) & 2(yz - xw) \\ 2(xz - yw) & 2(yz + xw) & 1 - 2(x^2 + y^2) \end{bmatrix}. \quad (2.28)$$

The diagonal terms can be made more symmetrical by replacing $1 - 2(y^2 + z^2)$ with $(x^2 + w^2 - y^2 - z^2)$, etc.

Between the axis/angle representation and quaternions, I generally prefer unit quaternions, because they possess a nice algebra that makes it easy to take products (compositions), ratios (change in rotation), and linear interpolations [174]. For example, the product of two quaternions $\mathbf{q}_0 = (\mathbf{v}_0, w_0)$ and $\mathbf{q}_1 = (\mathbf{v}_1, w_1)$ is given by

$$\mathbf{q}_2 = \mathbf{q}_0 \mathbf{q}_1 = (\mathbf{v}_0 \times \mathbf{v}_1 + w_0 \mathbf{v}_1 + w_1 \mathbf{v}_0, w_0 w_1 - \mathbf{v}_0 \cdot \mathbf{v}_1), \quad (2.29)$$

with the property that $\mathbf{R}(\mathbf{q}_2) = \mathbf{R}(\mathbf{q}_0)\mathbf{R}(\mathbf{q}_1)$. (Note that quaternion multiplication is *not* commutative, just as 3D rotations and matrix multiplications are not.) Taking the inverse of a quaternion is also easy: just flip the sign of \mathbf{v} or w (but not both!). However, when it comes time to update rotation estimates, I use an *incremental* form of the axis/angle representation (2.26), as described in Subsection 4.3.

2.3 Cylindrical and Spherical Coordinates

An alternative to using homographies or 3D motions to align images is to first warp the images into *cylindrical* coordinates and to then use a pure translational model to align them [43, 193]. Unfortunately, this only works if the images are all taken with a level camera or with a known tilt angle.

Assume for now that the camera is in its canonical position, i.e., its rotation matrix is the identity, $\mathbf{R} = \mathbf{I}$, so that the optic axis is aligned with the z axis and the y axis is aligned vertically. The 3D ray corresponding to an (x, y) pixel is therefore (x, y, f) .

We wish to project this image onto a *cylindrical surface* of unit radius [193]. Points on this surface are parameterized by an angle θ and a height h , with the 3D cylindrical coordinates corresponding to (θ, h) given by

$$(\sin\theta, h, \cos\theta) \propto (x, y, f), \quad (2.30)$$

as shown in Figure 2.6a. From this correspondence, we can compute the formula for the *warped* or *mapped* coordinates [199],

$$x' = s\theta = s \tan^{-1} \frac{x}{f}, \quad (2.31)$$

$$y' = sh = s \frac{y}{\sqrt{x^2 + f^2}}, \quad (2.32)$$

where s is an arbitrary scaling factor (sometimes called the *radius* of the cylinder) that can be set to $s = f$ to minimize the distortion (scaling)




Fig. 2.6 Projection from 3D to cylindrical and spherical coordinates.

near the center of the image.⁵ The inverse of this mapping equation is given by

$$x = f \tan \theta = f \tan \frac{x'}{s}, \quad (2.33)$$

$$y = h \sqrt{x^2 + f^2} = \frac{y'}{s} f \sqrt{1 + \tan^2 x'/s} = f \frac{y'}{s} \sec \frac{x'}{s}. \quad (2.34)$$

Images can also be projected onto a *spherical surface* [199], which is useful if the final panorama includes a full sphere or hemisphere of views, instead of just a cylindrical strip. In this case, the sphere is parameterized by two angles (θ, ϕ) , with 3D spherical coordinates given by

$$(\sin \theta \cos \phi, \sin \phi, \cos \theta \cos \phi) \propto (x, y, f), \quad (2.35)$$

as shown in Figure 2.6b. The correspondence between coordinates is now given by [199]

$$x' = s\theta = s \tan^{-1} \frac{x}{f}, \quad (2.36)$$

$$y' = s\phi = s \tan^{-1} \frac{y}{\sqrt{x^2 + f^2}}, \quad (2.37)$$

while the inverse is given by

$$x = f \tan \theta = f \tan \frac{x'}{s}, \quad (2.38)$$

$$\begin{aligned} y &= \sqrt{x^2 + f^2} \tan \phi = \tan \frac{y'}{s} f \sqrt{1 + \tan^2 x'/s} \\ &= f \tan \frac{y'}{s} \sec \frac{x'}{s}. \end{aligned} \quad (2.39)$$

Note that it may be simpler to generate a scaled (x, y, z) direction from (2.35) followed by a perspective division by z and a scaling by f .

Cylindrical image stitching algorithms are most commonly used when the camera is known to be level and only rotating around its vertical axis [43]. Under these conditions, images at different rotations are

⁵The scale can also be set to a larger or smaller value for the final compositing surface, depending on the desired output panorama resolution – see Section 6.




Fig. 2.7 An example of a cylindrical panorama: (a) two cylindrically warped images related by a horizontal translation; (b) part of a cylindrical panorama composited from a sequence of images.

related by a pure horizontal translation.⁶ This makes it attractive as an initial class project in an introductory computer vision course, since the full complexity of the perspective alignment algorithm (Subsections 3.5 and 4.3) can be avoided. Figure 2.7 shows how two cylindrically warped images from a leveled rotational panorama are related by a pure translation [199].

Professional panoramic photographers sometimes also use a pan-tilt head that makes it easy to control the tilt and to stop at specific *detents* in the rotation angle. This not only ensures a uniform coverage of the visual field with a desired amount of image overlap, but also makes it possible to stitch the images using cylindrical or spherical coordinates and pure translations. In this case, pixel coordinates (x, y, f) must first be rotated using the known tilt and panning angles before being projected into cylindrical or spherical coordinates [43]. Having a roughly known panning angle also makes it easier to compute the alignment, since the rough relative positioning of all the input images is known ahead of time, enabling a reduced search range for alignment. Figure 2.8 shows a full 3D rotational panorama unwrapped onto the surface of a sphere [199].

One final coordinate mapping worth mentioning is the *polar* mapping where the north pole lies along the optic axis rather than the

⁶Small vertical tilts can sometimes be compensated for with vertical translations.




Fig. 2.8 An example of a spherical panorama constructed from 54 photographs.

vertical axis,

$$(\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi) = s(x, y, z). \quad (2.40)$$

In this case, the mapping equations become

$$x' = s\phi \cos\theta = s \frac{x}{r} \tan^{-1} \frac{r}{z}, \quad (2.41)$$

$$y' = s\phi \sin\theta = s \frac{y}{r} \tan^{-1} \frac{r}{z}, \quad (2.42)$$

where $r = \sqrt{x^2 + y^2}$ is the *radial distance* in the (x, y) plane and $s\phi$ plays a similar role in the (x', y') plane. This mapping provides an attractive visualization surface for certain kinds of wide-angle panoramas and is also a good model for the distortion induced by *fish-eye lenses*, as discussed in Subsection 2.4. Note how for small values of (x, y) , the mapping equations reduces to $x' \approx sx/z$, which suggests that s plays a role similar to the focal length f .

2.4 Lens Distortions

When images are taken with wide-angle lenses, it is often necessary to model *lens distortions* such as *radial distortion*. The radial distortion model says that coordinates in the observed images are displaced away (*barrel distortion*) or toward (*pincushion distortion*) the image center by an amount proportional to their radial distance (Figures 2.9a and b). The simplest radial distortion models use low-order polynomials, e.g.,

$$\begin{aligned} x' &= x(1 + \kappa_1 r^2 + \kappa_2 r^4), \\ y' &= y(1 + \kappa_1 r^2 + \kappa_2 r^4), \end{aligned} \quad (2.43)$$




Fig. 2.9 Examples of radial lens distortion: (a) barrel, (b) pincushion, and (c) fisheye. The fisheye image spans almost a complete 180° from side-to-side.

where $r^2 = x^2 + y^2$ and κ_1 and κ_2 are called the *radial distortion parameters* [28,181].⁷ More complex distortion models also include *tangential (decentering) distortions* [181], but these are usually not necessary for consumer-level stitching.

A variety of techniques can be used to estimate the radial distortion parameters for a given lens. One of the simplest and most useful is to take an image of a scene with a lot of straight lines, especially lines aligned with and near the edges of the image. The radial distortion parameters can then be adjusted until all of the lines in the image are straight, which is commonly called the *plumb line method* [28,57,101].

Another approach is to use several overlapping images and to combine the estimation of the radial distortion parameters together with the image alignment process. Sawhney and Kumar [167] use a hierarchy of motion models (translation, affine, and projective) in a coarse-to-fine strategy coupled with a quadratic radial distortion correction term. They use direct (intensity-based) minimization to compute the alignment. Stein [187] uses a feature-based approach combined with a general 3D motion model (and quadratic radial distortion), which requires more matches than a parallax-free rotational panorama but is potentially more general. More recent approaches sometimes simultaneously compute both the unknown intrinsic parameters and the radial

⁷Sometimes the relationship between x and x' is expressed the other way around, i.e., using primed (final) coordinates on the right-hand side.

distortion coefficients, which may include higher order terms or more complex rational or nonparametric forms [13, 45, 84, 185, 189, 201, 204].

Fisheye lenses require a different model than traditional polynomial models of radial distortion (Figure 2.9c). Instead, fisheye lenses behave, to a first approximation, as *equi-distance* projectors of angles away from the optic axis [222], which is the same as the *polar projection* described by Equations (2.40)–(2.42). Xiong and Turkowski [222] describe how this model can be extended with the addition of an extra quadratic correction in ϕ , and how the unknown parameters (center of projection, scaling factor s , etc...) can be estimated from a set of overlapping fisheye images using a direct (intensity-based) nonlinear minimization algorithm.

Even more general models of lens distortion exist. For example, one can represent any lens as a mapping of pixel to rays in space [39, 75, 76, 202], either represented as a dense mapping or using a sparser interpolated smooth function such as a spline [39, 70].

3

Direct (pixel-based) Alignment

Once we have chosen a suitable motion model to describe the alignment between a pair of images, we need to devise some method to estimate its parameters. One approach is to shift or warp the images relative to each other and to look at how much the pixels agree. Approaches that use pixel-to-pixel matching are often called *direct methods*, as opposed to the *feature-based methods* described in the next section.

To use a direct method, a suitable *error metric* must first be chosen to compare the images. Once this has been established, a suitable *search* technique must be devised. The simplest technique is to exhaustively try all possible alignments, i.e., to do a *full search*. In practice, this may be too slow, so *hierarchical* coarse-to-fine techniques based on image pyramids have been developed. Alternatively, Fourier transforms can be used to speed up the computation. To get sub-pixel precision in the alignment, *incremental* methods based on a Taylor series expansion of the image function are often used. These can also be applied to *parametric motion models*. Each of these techniques is described in more detail below.

3.1 Error Metrics

The simplest way to establish an alignment between two images is to shift one image relative to the other. Given a *template* image $I_0(\mathbf{x})$ sampled at discrete pixel locations $\{\mathbf{x}_i = (x_i, y_i)\}$, we wish to find where it is located in image $I_1(\mathbf{x})$. A least-squares solution to this problem is to find the minimum of the *sum of squared differences* (SSD) function

$$E_{\text{SSD}}(\mathbf{u}) = \sum_i [I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i)]^2 = \sum_i e_i^2, \quad (3.1)$$

where $\mathbf{u} = (u, v)$ is the *displacement*, and $e_i = I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i)$ is called the *residual error* (or the *displaced frame difference* in the video coding literature).¹ (We ignore for the moment the possibility that parts of I_0 may lie outside the boundaries of I_1 or be otherwise not visible.)

In general, the displacement \mathbf{u} can be fractional, so a suitable interpolation function must be applied to image $I_1(\mathbf{x})$. In practice, a bilinear interpolant is often used, but bi-cubic interpolation should yield slightly better results. Color images can be processed by summing differences across all three color channels, although it is also possible to first transform the images into a different color space or to only use the luminance (which is often done in video encoders).

Robust Error Metrics We can make the above error metric more robust to outliers by replacing the squared error terms with a robust function $\rho(e_i)$ [22, 78, 92, 188] to obtain

$$E_{\text{SRD}}(\mathbf{u}) = \sum_i \rho(I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i)) = \sum_i \rho(e_i). \quad (3.2)$$

The robust norm $\rho(e)$ is a function that grows less quickly than the quadratic penalty associated with least squares. One such function, sometimes used in motion estimation for video coding because of its speed, is the *sum of absolute differences* (SAD) metric, i.e.,

$$E_{\text{SAD}}(\mathbf{u}) = \sum_i |I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i)| = \sum_i |e_i|. \quad (3.3)$$

¹The usual justification for using least squares is that it is the optimal estimate with respect to Gaussian noise. See the discussion below on robust alternatives.

However, since this function is not differentiable at the origin, it is not well suited to gradient-descent approaches such as the ones presented in Subsection 3.4.

Instead, a smoothly varying function that is quadratic for small values but grows more slowly away from the origin is often used. Black and Rangarajan [24] discuss a variety of such functions, including the *Geman–McClure* function,

$$\rho_{\text{GM}}(x) = \frac{x^2}{1 + x^2/a^2}, \quad (3.4)$$

where a is a constant that can be thought of as an *outlier threshold*. An appropriate value for the threshold can itself be derived using robust statistics [78, 92, 163], e.g., by computing the *median of absolute differences*, $\text{MAD} = \text{med}_i |e_i|$, and multiplying by 1.4 to obtain a robust estimate of the standard deviation of the nonoutlier noise process [188].

Spatially Varying Weights The error metrics above ignore the fact that for a given alignment, some of the pixels being compared may lie outside the original image boundaries. Furthermore, we may want to partially or completely downweight the contributions of certain pixels. For example, we may want to selectively “erase” some parts of an image from consideration, e.g., when stitching a mosaic where unwanted foreground objects have been cut out. For applications such as background stabilization, we may want to downweight the middle part of the image, which often contains independently moving objects being tracked by the camera.

All of these tasks can be accomplished by associating a spatially varying per-pixel weight value with each of the two images being matched. The error metric then become the weighted (or *windowed*) SSD function,

$$E_{\text{WSSD}}(\mathbf{u}) = \sum_i w_0(\mathbf{x})w_1(\mathbf{x}_i + \mathbf{u})[I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i)]^2, \quad (3.5)$$

where the weighting functions w_0 and w_1 are zero outside the valid ranges of the images.

If a large range of potential motions is allowed, the above metric can have a bias toward smaller overlap solutions. To counteract this

bias, the windowed SSD score can be divided by the overlap area

$$A = \sum_i w_0(\mathbf{x})w_1(\mathbf{x}_i + \mathbf{u}) \quad (3.6)$$

to compute a *per-pixel* (or mean) squared pixel error. The square root of this quantity is the *root mean squared* intensity error

$$\text{RMS} = \sqrt{E_{\text{WSSD}}/A}, \quad (3.7)$$

often seen reported in comparative studies.

Bias and Gain (exposure differences) Often, the two images being aligned were not taken with the same exposure. A simple model of linear (affine) intensity variation between the two images is the *bias and gain* model,

$$I_1(\mathbf{x} + \mathbf{u}) = (1 + \alpha)I_0(\mathbf{x}) + \beta, \quad (3.8)$$

where β is the *bias* and α is the *gain* [12, 66, 68, 123]. The least squares formulation then becomes

$$\begin{aligned} E_{\text{BG}}(\mathbf{u}) &= \sum_i [I_1(\mathbf{x}_i + \mathbf{u}) - (1 + \alpha)I_0(\mathbf{x}_i) - \beta]^2 \\ &= \sum_i [\alpha I_0(\mathbf{x}_i) + \beta - e_i]^2. \end{aligned} \quad (3.9)$$

Rather than taking a simple squared difference between corresponding patches, it becomes necessary to perform a *linear regression*, which is somewhat more costly. Note that for color images, it may be necessary to estimate a different bias and gain for each color channel to compensate for the automatic *color correction* performed by some digital cameras.

A more general (spatially varying nonparametric) model of intensity variation, which is computed as part of the registration process, is presented in [97]. This can be useful for dealing with local variations such as the *vignetting* caused by wide-angle lenses. It is also possible to pre-process the images before comparing their values, e.g., by using band-pass filtered images [20, 34] or using other local transformations such as histograms or rank transforms [48, 225], or to maximize *mutual information* [108, 214].

Correlation An alternative to taking intensity differences is to perform *correlation*, i.e., to maximize the *product* (or *cross-correlation*) of the two aligned images,

$$E_{CC}(\mathbf{u}) = \sum_i I_0(\mathbf{x}_i) I_1(\mathbf{x}_i + \mathbf{u}). \quad (3.10)$$

At first glance, this may appear to make bias and gain modeling unnecessary, since the images will prefer to line up regardless of their relative scales and offsets. However, this is actually not true. If a very bright patch exists in $I_1(\mathbf{x})$, the maximum product may actually lie in that area.

For this reason, *normalized cross-correlation* is more commonly used,

$$E_{NCC}(\mathbf{u}) = \frac{\sum_i [I_0(\mathbf{x}_i) - \bar{I}_0] [I_1(\mathbf{x}_i + \mathbf{u}) - \bar{I}_1]}{\sqrt{\sum_i [I_0(\mathbf{x}_i) - \bar{I}_0]^2 [I_1(\mathbf{x}_i + \mathbf{u}) - \bar{I}_1]^2}}, \quad (3.11)$$

where

$$\bar{I}_0 = \frac{1}{N} \sum_i I_0(\mathbf{x}_i) \quad \text{and} \quad (3.12)$$

$$\bar{I}_1 = \frac{1}{N} \sum_i I_1(\mathbf{x}_i + \mathbf{u}) \quad (3.13)$$

are the *mean images* of the corresponding patches and N is the number of pixels in the patch. The normalized cross-correlation score is always guaranteed to be in the range $[-1, 1]$, which makes it easier to handle in some higher-level applications (such as deciding which patches truly match). Note, however, that the NCC score is undefined if either of the two patches has zero variance (and in fact, its performance degrades for noisy low-contrast regions).

3.2 Hierarchical Motion Estimation

Now that we have defined an alignment cost function to optimize, how do we find its minimum? The simplest solution is to do a *full search* over some range of shifts, using either integer or sub-pixel steps. This is often the approach used for *block matching* in *motion compensated*

video compression, where a range of possible motions (say ± 16 pixels) is explored.²

To accelerate this search process, *hierarchical motion estimation* is often used, where an image pyramid is first constructed, and a search over a smaller number of discrete pixels (corresponding to the same range of motion) is first performed at coarser levels [3, 20, 154]. The motion estimate from one level of the pyramid can then be used to initialize a smaller *local* search at the next finer level. While this is not guaranteed to produce the same result as full search, it usually works almost as well and is much faster.

More formally, let

$$I_k^{(l)}(\mathbf{x}_j) \leftarrow \tilde{I}_k^{(l-1)}(2\mathbf{x}_j) \quad (3.14)$$

be the *decimated* image at level l obtained by subsampling (*downsampling*) a smoothed (pre-filtered) version of the image at level $l - 1$. At the coarsest level, we search for the best displacement $\mathbf{u}^{(l)}$ that minimizes the difference between images $I_0^{(l)}$ and $I_1^{(l)}$. This is usually done using a full search over some range of displacements $\mathbf{u}^{(l)} \in 2^{-l}[-S, S]^2$ (where S is the desired *search range* at the finest (original) resolution level), optionally followed by the incremental refinement step described in Subsection 3.4.

Once a suitable motion vector has been estimated, it is used to *predict* a likely displacement

$$\hat{\mathbf{u}}^{(l-1)} \leftarrow 2\mathbf{u}^{(l)} \quad (3.15)$$

for the next finer level.³ The search over displacements is then repeated at the finer level over a much narrower range of displacements, say $\hat{\mathbf{u}}^{(l-1)} \pm 1$, again optionally combined with an incremental refinement step [3]. A nice description of the whole process, extended to parametric motion estimation (Subsection 3.5), can be found in [20].

²In stereo matching, an explicit search over all possible disparities (i.e., a *plane sweep*) is almost always performed, since the number of search hypotheses is much smaller due to the 1D nature of the potential displacements [170].

³This doubling of displacements is only necessary if displacements are defined in integer *pixel* coordinates, which is the usual case in the literature, e.g., [20]. If *normalized device coordinates* (Section 2) are used instead, the displacements (and search ranges) need not change from level to level, although the step sizes will need to be adjusted (to keep search steps of roughly one pixel).

3.3 Fourier-Based Alignment

When the search range corresponds to a significant fraction of the larger image (as is the case in image stitching), the hierarchical approach may not work that well, since it is often not possible to coarsen the representation too much before significant features get blurred away. In this case, a Fourier-based approach may be preferable.

Fourier-based alignment relies on the fact that the Fourier transform of a shifted signal has the same magnitude as the original signal but linearly varying phase, i.e.,

$$\mathcal{F}\{I_1(\mathbf{x} + \mathbf{u})\} = \mathcal{F}\{I_1(\mathbf{x})\} e^{-2\pi j \mathbf{u} \cdot \mathbf{f}} = \mathcal{I}_1(\mathbf{f}) e^{-2\pi j \mathbf{u} \cdot \mathbf{f}}, \quad (3.16)$$

where \mathbf{f} is the vector-valued frequency of the Fourier transform and we use calligraphic notation $\mathcal{I}_1(\mathbf{f}) = \mathcal{F}\{I_1(\mathbf{x})\}$ to denote the Fourier transform of a signal [146, p. 57].

Another useful property of Fourier transforms is that convolution in the spatial domain corresponds to multiplication in the Fourier domain [146, p. 58]. Thus, the Fourier transform of the cross-correlation function E_{CC} can be written as

$$\begin{aligned} \mathcal{F}\{E_{CC}(\mathbf{u})\} &= \mathcal{F}\left\{\sum_i I_0(\mathbf{x}_i) I_1(\mathbf{x}_i + \mathbf{u})\right\} \\ &= \mathcal{F}\{I_0(\mathbf{u}) \bar{*} I_1(\mathbf{u})\} = \mathcal{I}_0(\mathbf{f}) \mathcal{I}_1^*(\mathbf{f}), \end{aligned} \quad (3.17)$$

where

$$f(\mathbf{u}) \bar{*} g(\mathbf{u}) = \sum_i f(\mathbf{x}_i) g(\mathbf{x}_i + \mathbf{u}) \quad (3.18)$$

is the *correlation* function, i.e., the convolution of one signal with the reverse of the other, and $\mathcal{I}_1^*(\mathbf{f})$ is the *complex conjugate* of $\mathcal{I}_1(\mathbf{f})$. (This is because convolution is defined as the summation of one signal with the reverse of the other [146].)

Thus, to efficiently evaluate E_{CC} over the range of all possible values of \mathbf{u} , we take the Fourier transforms of both images $I_0(\mathbf{x})$ and $I_1(\mathbf{x})$, multiply both transforms together (after conjugating the second one), and take the inverse transform of the result. The Fast Fourier Transform algorithm can compute the transform of an $N \times M$ image

in $O(NM \log NM)$ operations [146]. This can be significantly faster than the $O(N^2M^2)$ operations required to do a full search when the full range of image overlaps is considered.

While Fourier-based convolution is often used to accelerate the computation of image correlations, it can also be used to accelerate the sum of squared differences function (and its variants) as well. Consider the SSD formula given in (3.1). Its Fourier transform can be written as

$$\begin{aligned} \mathcal{F}\{E_{\text{SSD}}(\mathbf{u})\} &= \mathcal{F}\left\{\sum_i [I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i)]^2\right\} \\ &= \delta(\mathbf{f}) \sum_i [I_0^2(\mathbf{x}_i) + I_1^2(\mathbf{x}_i)] - 2\mathcal{I}_0(\mathbf{f})\mathcal{I}_1^*(\mathbf{f}). \end{aligned} \quad (3.19)$$

Thus, the SSD function can be computed by taking twice the correlation function and subtracting it from the sum of the energies in the two images.

Windowed Correlation Unfortunately, the Fourier convolution theorem only applies when the summation over \mathbf{x}_i is performed over *all* the pixels in both images, using a circular shift of the image when accessing pixels outside the original boundaries. While this is acceptable for small shifts and comparably sized images, it makes no sense when the images overlap by a small amount or one image is a small subset of the other.

In that case, the cross-correlation function should be replaced with a *windowed* (weighted) cross-correlation function,

$$E_{\text{WCC}}(\mathbf{u}) = \sum_i w_0(\mathbf{x}_i)I_0(\mathbf{x}_i) w_1(\mathbf{x}_i + \mathbf{u})I_1(\mathbf{x}_i + \mathbf{u}), \quad (3.20)$$

$$= [w_0(\mathbf{x})I_0(\mathbf{x})] \star [w_1(\mathbf{x})I_1(\mathbf{x})], \quad (3.21)$$

where the weighting functions w_0 and w_1 are zero outside the valid ranges of the images, and both images are padded so that circular shifts return 0 values outside the original image boundaries.

An even more interesting case is the computation of the *weighted* SSD function introduced in (3.5),

$$\begin{aligned} E_{\text{WSSD}}(\mathbf{u}) &= \sum_i w_0(\mathbf{x})w_1(\mathbf{x}_i + \mathbf{u})[I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i)]^2 \\ &= w_0(\mathbf{x})\bar{*}[w_1(\mathbf{x})I_1^2(\mathbf{x})] + [w_0(\mathbf{x})I_0^2(\mathbf{x})]\bar{*}w_1(\mathbf{x}) \\ &\quad - 2[w_0(\mathbf{x})I_0(\mathbf{x})]\bar{*}[w_1(\mathbf{x})I_1(\mathbf{x})]. \end{aligned} \quad (3.22)$$

The Fourier transform of the resulting expression is therefore

$$\mathcal{F}\{E_{\text{WSSD}}(\mathbf{u})\} = \mathcal{W}_0(\mathbf{f})\mathcal{S}_1^*(\mathbf{f}) + \mathcal{S}_0(\mathbf{f})\mathcal{W}_1^*(\mathbf{f}) - 2\hat{\mathcal{I}}_0(\mathbf{f})\hat{\mathcal{I}}_1^*(\mathbf{f}), \quad (3.23)$$

where

$$\begin{aligned} \mathcal{W}_0 &= \mathcal{F}\{w_0(\mathbf{x})\}, & \mathcal{W}_1 &= \mathcal{F}\{w_1(\mathbf{x})\}, \\ \hat{\mathcal{I}}_0 &= \mathcal{F}\{w_0(\mathbf{x})I_0(\mathbf{x})\}, & \hat{\mathcal{I}}_1 &= \mathcal{F}\{w_1(\mathbf{x})I_1(\mathbf{x})\}, \\ \mathcal{S}_0 &= \mathcal{F}\{w_0(\mathbf{x})I_0^2(\mathbf{x})\}, & \text{and } \mathcal{S}_1 &= \mathcal{F}\{w_1(\mathbf{x})I_1^2(\mathbf{x})\} \end{aligned} \quad (3.24)$$

are the Fourier transforms of the weighting functions and the *weighted* original and squared image signals. Thus, for the cost of a few additional image multiplies and Fourier transforms, the correct windowed SSD function can be computed. (To my knowledge, I have not seen this formulation written down before, but I have been teaching it to students for several years now.)

The same kind of derivation can be applied to the bias-gain corrected sum of squared difference function E_{BG} . Again, Fourier transforms can be used to efficiently compute all the correlations needed to perform the linear regression in the bias and gain parameters to estimate the exposure-compensated difference for each potential shift.

Phase Correlation A variant of regular correlation (3.17) that is sometimes used for motion estimation is *phase correlation* [29, 109]. Here, the spectrum of the two signals being matched is *whitened* by dividing each per-frequency product in (3.17) by the magnitudes of the Fourier transforms,

$$\mathcal{F}\{E_{\text{PC}}(\mathbf{u})\} = \frac{\mathcal{I}_0(\mathbf{f})\mathcal{I}_1^*(\mathbf{f})}{\|\mathcal{I}_0(\mathbf{f})\|\|\mathcal{I}_1(\mathbf{f})\|}, \quad (3.25)$$

before taking the final inverse Fourier transform. In the case of noiseless signals with perfect (cyclic) shift, we have $I_1(\mathbf{x} + \mathbf{u}) = I_0(\mathbf{x})$, and hence from (3.16), we obtain

$$\begin{aligned}\mathcal{F}\{I_1(\mathbf{x} + \mathbf{u})\} &= \mathcal{I}_1(\mathbf{f})e^{-2\pi j\mathbf{u}\cdot\mathbf{f}} = \mathcal{I}_0(\mathbf{f}) \quad \text{and} \\ \mathcal{F}\{E_{\text{PC}}(\mathbf{u})\} &= e^{-2\pi j\mathbf{u}\cdot\mathbf{f}}.\end{aligned}\tag{3.26}$$

The output of phase correlation (under ideal conditions) is therefore a single spike (impulse) located at the correct value of \mathbf{u} , which (in principle) makes it easier to find the correct estimate.

Phase correlation has a reputation in some quarters of outperforming regular correlation, but this behavior depends on the characteristics of the signals and noise. If the original images are contaminated by noise in a narrow frequency band (e.g., low-frequency noise or peaked frequency “hum”), the whitening process effectively de-emphasizes the noise in these regions. However, if the original signals have very low signal-to-noise ratio at some frequencies (say, two blurry or low-textured images with lots of high-frequency noise), the whitening process can actually decrease performance. Recently, gradient cross-correlation has emerged as a promising alternative to phase correlation [4], although further systematic studies are probably warranted. Phase correlation has also been studied by Fleet and Jepson [61] as a method for estimating general optical flow and stereo disparity.

Rotations and Scale While Fourier-based alignment is mostly used to estimate translational shifts between images, it can, under certain limited conditions, also be used to estimate in-plane rotations and scales. Consider two images that are related *purely* by rotation, i.e.,

$$I_1(\hat{\mathbf{R}}\mathbf{x}) = I_0(\mathbf{x}).\tag{3.27}$$

If we re-sample the images into *polar coordinates*,

$$\begin{aligned}\tilde{I}_0(r, \theta) &= I_0(r \cos \theta, r \sin \theta) \quad \text{and} \\ \tilde{I}_1(r, \theta) &= I_1(r \cos \theta, r \sin \theta),\end{aligned}\tag{3.28}$$

we obtain

$$\tilde{I}_1(r, \theta + \hat{\theta}) = \tilde{I}_0(r, \theta).\tag{3.29}$$

The desired rotation can then be estimated using an FFT shift-based technique.

If the two images are also related by a scale,

$$I_1(e^{\hat{s}}\hat{\mathbf{R}}\mathbf{x}) = I_0(\mathbf{x}), \quad (3.30)$$

we can re-sample into *log-polar coordinates*,

$$\tilde{I}_0(s, \theta) = I_0(e^s \cos \theta, e^s \sin \theta) \quad \text{and} \quad \tilde{I}_1(s, \theta) = I_1(e^s \cos \theta, e^s \sin \theta), \quad (3.31)$$

to obtain

$$\tilde{I}_1(s + \hat{s}, \theta + \hat{\theta}) = \tilde{I}_0(s, \theta). \quad (3.32)$$

In this case, care must be taken to choose a suitable range of s values that reasonably samples the original image.

For images that are also translated by a small amount,

$$I_1(e^{\hat{s}}\hat{\mathbf{R}}\mathbf{x} + \mathbf{t}) = I_0(\mathbf{x}), \quad (3.33)$$

De Castro and Morandi [51] proposed an ingenious solution that uses several steps to estimate the unknown parameters. First, both images are converted to the Fourier domain, and only the magnitudes of the transformed images are retained. In principle, the Fourier magnitude images are insensitive to translations in the image plane (although the usual caveats about border effects apply). Next, the two magnitude images are aligned in rotation and scale using the polar or log-polar representations. Once rotation and scale are estimated, one of the images can be de-rotated and scaled, and a regular translational algorithm can be applied to estimate the translational shift.

Unfortunately, this trick only applies when the images have large overlap (small translational motion). For more general motion of patches or images, the parametric motion estimator described in Subsection 3.5 or the feature-based approaches described in Section 4 need to be used.

3.4 Incremental Refinement

The techniques described up till now can estimate translational alignment to the nearest pixel (or potentially fractional pixel if smaller search steps are used). In general, image stabilization and stitching




Fig. 3.1 Taylor series approximation of a function and the incremental computation of the optic flow correction amount. $\mathbf{J}_1(\mathbf{x}_i + \mathbf{u})$ is the image gradient at $(\mathbf{x}_i + \mathbf{u})$ and e_i is the current intensity difference.

applications require much higher accuracies to obtain acceptable results.

To obtain better *sub-pixel* estimates, we can use one of several techniques [205]. One possibility is to evaluate several discrete (integer or fractional) values of (u, v) around the best value found so far and to *interpolate* the matching score to find an analytic minimum.

A more commonly used approach, first proposed by Lucas and Kanade [123], is to do *gradient descent* on the SSD energy function (3.1), using a Taylor Series expansion of the image function (Figure 3.1),

$$\begin{aligned} E_{\text{LK-SSD}}(\mathbf{u} + \Delta\mathbf{u}) &= \sum_i [I_1(\mathbf{x}_i + \mathbf{u} + \Delta\mathbf{u}) - I_0(\mathbf{x}_i)]^2 \\ &\approx \sum_i [I_1(\mathbf{x}_i + \mathbf{u}) + \mathbf{J}_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} \\ &\quad - I_0(\mathbf{x}_i)]^2 \end{aligned} \quad (3.34)$$

$$= \sum_i [\mathbf{J}_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} + e_i]^2, \quad (3.35)$$

where

$$\mathbf{J}_1(\mathbf{x}_i + \mathbf{u}) = \nabla I_1(\mathbf{x}_i + \mathbf{u}) = \left(\frac{\partial I_1}{\partial x}, \frac{\partial I_1}{\partial y} \right) (\mathbf{x}_i + \mathbf{u}) \quad (3.36)$$

is the *image gradient* at $(\mathbf{x}_i + \mathbf{u})$ and

$$e_i = I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i), \quad (3.37)$$

first introduced in (3.1), is the current intensity error.⁴

⁴We follow the convention, commonly used in robotics and in [9], that derivatives with respect to (column) vectors result in row vectors, so that fewer transposes are needed in the formulas.

The above least squares problem can be minimized by solving the associated *normal equations* [69],

$$\mathbf{A}\Delta\mathbf{u} = \mathbf{b}, \quad (3.38)$$

where

$$\mathbf{A} = \sum_i \mathbf{J}_1^T(\mathbf{x}_i + \mathbf{u})\mathbf{J}_1(\mathbf{x}_i + \mathbf{u}) \quad (3.39)$$

and

$$\mathbf{b} = - \sum_i e_i \mathbf{J}_1^T(\mathbf{x}_i + \mathbf{u}) \quad (3.40)$$

are called the (Gauss–Newton approximation of the) *Hessian* and *gradient-weighted residual vector*, respectively.⁵ These matrices are also often written as

$$\mathbf{A} = \begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix} \quad \text{and} \quad \mathbf{b} = - \begin{bmatrix} \sum I_x I_t \\ \sum I_y I_t \end{bmatrix}, \quad (3.41)$$

where the subscripts in I_x and I_y denote spatial derivatives, and I_t is called the *temporal derivative*, which makes sense if we are computing instantaneous velocity in a video sequence.

The gradients required for $\mathbf{J}_1(\mathbf{x}_i + \mathbf{u})$ can be evaluated at the same time as the image warps required to estimate $I_1(\mathbf{x}_i + \mathbf{u})$, and in fact are often computed as a side-product of image interpolation. If efficiency is a concern, these gradients can be replaced by the gradients in the *template* image,

$$\mathbf{J}_1(\mathbf{x}_i + \mathbf{u}) \approx \mathbf{J}_0(\mathbf{x}), \quad (3.42)$$

since near the correct alignment, the template and displaced target images should look similar. This has the advantage of allowing the pre-computation of the Hessian and Jacobian images, which can result in significant computational savings [9,77]. A further reduction in computation can be obtained by writing the warped image $I_1(\mathbf{x}_i + \mathbf{u})$ used to compute e_i in (3.37) as a convolution of a sub-pixel interpolation filter

⁵The true Hessian is the full second derivative of the error function E , which may not be positive definite.

with the discrete samples in I_1 [149]. Pre-computing the inner product between the gradient field and shifted version of I_1 allows the iterative re-computation of e_i to be performed in constant time (independent of the number of pixels).

The effectiveness of the above incremental update rule relies on the quality of the Taylor series approximation. When far away from the true displacement (say 1–2 pixels), several iterations may be needed. (It is possible, however, to estimate a value for \mathbf{J}_1 using a least-squares fit to a series of larger displacements in order to increase the range of convergence [98].)

When started in the vicinity of the correct solution, only a few iterations usually suffice. A commonly used stopping criterion is to monitor the magnitude of the displacement correction $\|\mathbf{u}\|$ and to stop when it drops below a certain threshold (say $1/10$ th of a pixel). For larger motions, it is usual to combine the incremental update rule with a hierarchical coarse-to-fine search strategy, as described in Subsection 3.2.

Conditioning and Aperture Problems Sometimes, the inversion of the linear system (3.38) can be poorly conditioned because of lack of two-dimensional texture in the patch being aligned. A commonly occurring example of this is the *aperture problem*, first identified in some of the early papers on optic flow [91] and then studied more extensively by Anandan [3]. Consider an image patch that consists of a slanted edge moving to the right (Figure 3.2). Only the *normal*




Fig. 3.2 Aperture problems for different image patches: (a) stable (“corner-like”) flow; (b) classic aperture problem (barber-pole illusion); (c) textureless region.

component of the velocity (displacement) can be reliably recovered in this case. This manifests itself in (3.38) as a *rank-deficient* matrix \mathbf{A} , i.e., one whose smaller eigenvalue is very close to zero.⁶

When Equation (3.38) is solved, the component of the displacement along the edge is very poorly conditioned and can result in wild guesses under small noise perturbations. One way to mitigate this problem is to add a *prior* (soft constraint) on the expected range of motions [10, 72, 180]. This can be accomplished by adding a small value to the diagonal of \mathbf{A} , which essentially biases the solution toward smaller $\Delta\mathbf{u}$ values that still (mostly) minimize the squared error.

However, the pure Gaussian model assumed when using a simple (fixed) quadratic prior, as in [180], does not always hold in practice, e.g., because of aliasing along strong edges [206]. For this reason, it may be prudent to add some small fraction (say 5%) of the larger eigenvalue to the smaller one before doing the matrix inversion.

Uncertainty Modeling The reliability of a particular patch-based motion estimate can be captured more formally with an *uncertainty model*. The simplest such model is a *covariance matrix*, which captures the expected variance in the motion estimate in all possible directions. Under small amounts of additive Gaussian noise, it can be shown that the covariance matrix $\Sigma\mathbf{u}$ is proportional to the inverse of the Hessian \mathbf{A} ,

$$\Sigma\mathbf{u} = \sigma_n^2 \mathbf{A}^{-1}, \quad (3.43)$$

where σ_n^2 is the variance of the additive Gaussian noise [3, 128, 191]. For larger amounts of noise, the linearization performed by the Lucas–Kanade algorithm in (3.35) is only approximate, so the above quantity becomes the *Cramer–Rao lower bound* on the true covariance. Thus, the minimum and maximum eigenvalues of the Hessian \mathbf{A} can now be interpreted as the (scaled) inverse variances in the least-certain and most-certain directions of motion. (A more detailed analysis using a more realistic model of image noise can be found in [184].)

⁶The matrix \mathbf{A} is by construction always guaranteed to be symmetric positive semi-definite, i.e., it has real non-negative eigenvalues.

Bias and Gain, Weighting, and Robust Error Metrics The Lucas–Kanade update rule can also be applied to the bias–gain equation (3.9) to obtain

$$E_{\text{LK-BG}}(\mathbf{u} + \Delta\mathbf{u}) = \sum_i [\mathbf{J}_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} + e_i - \alpha I_0(\mathbf{x}_i) - \beta]^2 \quad (3.44)$$

[12,66,68,123]. The resulting 4×4 system of equations in can be solved to simultaneously estimate the translational displacement update $\Delta\mathbf{u}$ and the bias and gain parameters β and α .⁷

A similar formulation can be derived for images (templates) that have a *linear appearance variation*,

$$I_1(\mathbf{x} + \mathbf{u}) \approx I_0(\mathbf{x}) + \sum_j \lambda_j B_j(\mathbf{x}), \quad (3.45)$$

where the $B_j(\mathbf{x})$ are the *basis images* and the λ_j are the unknown coefficients [11, 12, 77]. Potential linear appearance variations include illumination changes [77] and small non-rigid deformations [23].

A weighted (windowed) version of the Lucas–Kanade algorithm is also possible,

$$E_{\text{LK-WSSD}}(\mathbf{u} + \Delta\mathbf{u}) = \sum_i w_0(\mathbf{x})w_1(\mathbf{x}_i + \mathbf{u})[\mathbf{J}_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} + e_i]^2. \quad (3.46)$$

Note that here, in deriving the Lucas–Kanade update from the original weighted SSD function (3.5), we have neglected taking the derivative of $w_1(\mathbf{x}_i + \mathbf{u})$ weighting function with respect to \mathbf{u} , which is usually acceptable in practice, especially if the weighting function is a binary mask with relatively few transitions.

Baker *et al.* [11] only use the $w_0(\mathbf{x})$ term, which is reasonable if the two images have the same extent and no (independent) cutouts in the overlap region. They also discuss the idea of making the weighting proportional to $\nabla I(\mathbf{x})$, which helps for very noisy images, where the gradient itself is noisy. Similar observation, formulated in terms of *total least squares* [212], have been made by other researchers studying optic flow (motion) estimation [6,139,217]. Lastly, Baker *et al.* [11] show how

⁷In practice, it may be possible to decouple the bias–gain and motion update parameters, i.e., to solve two independent 2×2 systems, which is a little faster.

evaluating (3.46) at just the *most reliable* (highest gradient) pixels does not significantly reduce performance for large enough images, even if only 5%–10% of the pixels are used. (This idea was originally proposed by Dellaert and Collins [53], who used a more sophisticated selection criterion.)

The Lucas–Kanade incremental refinement step can also be applied to the robust error metric introduced in Subsection 3.1,

$$E_{\text{LK-SRD}}(\mathbf{u} + \Delta\mathbf{u}) = \sum_i \rho(\mathbf{J}_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} + e_i). \quad (3.47)$$

We can take the derivative of this function w.r.t. \mathbf{u} and set it to 0,

$$\sum_i \psi(e_i) \frac{\partial e_i}{\partial \mathbf{u}} = \sum_i \psi(e_i) \mathbf{J}_1(\mathbf{x} + \mathbf{u}) = 0, \quad (3.48)$$

where $\Psi(e) = \rho'(e)$ is the derivative of ρ . If we introduce a weight function $w(e) = \Psi(e)/e$, we can write this as

$$\sum_i w(e_i) \mathbf{J}_1^T(\mathbf{x} + \mathbf{u}) [\mathbf{J}_1(\mathbf{x}_i + \mathbf{u})\Delta\mathbf{u} + e_i] = 0. \quad (3.49)$$

This results in the *Iteratively Re-weighted Least Squares* algorithm, which alternates between computing the weight functions $w(e_i)$ and solving the above weighted least squares problem [92, 188]. Alternative incremental robust least squares algorithms can be found in [11, 22, 24, 166] and textbooks and tutorials on robust statistics [78, 92, 163, 188].

3.5 Parametric Motion

Many image alignment tasks, for example image stitching with hand-held cameras, require the use of more sophisticated motion models, as described in Section 2. Since these models typically have more parameters than pure translation, a full search over the possible range of values is impractical. Instead, the incremental Lucas–Kanade algorithm can be generalized to parametric motion models and used in conjunction with a hierarchical search algorithm [9, 20, 66, 123, 157].

For parametric motion, instead of using a single constant translation vector \mathbf{u} , we use a spatially varying *motion field* or *correspondence map*, $\mathbf{x}'(\mathbf{x}; \mathbf{p})$, parameterized by a low-dimensional vector \mathbf{p} , where \mathbf{x}' can

Table 3.1 Jacobians of the 2D coordinate transformations.

Transform	Matrix	Parameters	Jacobian $\mathbf{J}_{\mathbf{x}'}$
Translation	$\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \end{bmatrix}$	(t_x, t_y)	$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$
Euclidean	$\begin{bmatrix} c_\theta & -s_\theta & t_x \\ s_\theta & c_\theta & t_y \end{bmatrix}$	(t_x, t_y, θ)	$\begin{bmatrix} 1 & 0 & -s_\theta x - c_\theta y \\ 0 & 1 & c_\theta x - s_\theta y \end{bmatrix}$
Similarity	$\begin{bmatrix} 1+a & -b & t_x \\ b & 1+a & t_y \end{bmatrix}$	(t_x, t_y, a, b)	$\begin{bmatrix} 1 & 0 & x & -y \\ 0 & 1 & y & x \end{bmatrix}$
Affine	$\begin{bmatrix} 1+a_{00} & a_{01} & t_x \\ a_{10} & 1+a_{11} & t_y \end{bmatrix}$	$(t_x, t_y, a_{00}, a_{01}, a_{10}, a_{11})$	$\begin{bmatrix} 1 & 0 & x & y & 0 & 0 \\ 0 & 1 & 0 & 0 & x & y \end{bmatrix}$
Projective	$\begin{bmatrix} 1+h_{00} & h_{01} & h_{02} \\ h_{10} & 1+h_{11} & h_{12} \\ h_{20} & h_{21} & 1 \end{bmatrix}$	(h_{00}, \dots, h_{21})	(See text)

be any of the motion models presented in Section 2. The parametric incremental motion update rule now becomes

$$E_{\text{LK-PM}}(\mathbf{p} + \Delta\mathbf{p}) = \sum_i [I_1(\mathbf{x}'(\mathbf{x}_i; \mathbf{p} + \Delta\mathbf{p})) - I_0(\mathbf{x}_i)]^2 \quad (3.50)$$

$$\approx \sum_i [I_1(\mathbf{x}'_i) + \mathbf{J}_1(\mathbf{x}'_i)\Delta\mathbf{p} - I_0(\mathbf{x}_i)]^2 \quad (3.51)$$

$$= \sum_i [\mathbf{J}_1(\mathbf{x}'_i)\Delta\mathbf{p} + e_i]^2, \quad (3.52)$$

where the Jacobian is now

$$\mathbf{J}_1(\mathbf{x}'_i) = \frac{\partial I_1}{\partial \mathbf{p}} = \nabla I_1(\mathbf{x}'_i) \frac{\partial \mathbf{x}'}{\partial \mathbf{p}}(\mathbf{x}_i), \quad (3.53)$$

i.e., the product of the image gradient ∇I_1 with the Jacobian of correspondence field, $\mathbf{J}_{\mathbf{x}'} = \partial \mathbf{x}' / \partial \mathbf{p}$.

Table 3.1 shows the motion Jacobians $\mathbf{J}_{\mathbf{x}'}$ for the 2D planar transformations introduced in Section 2.⁸ Note how I have re-parameterized the motion matrices so that they are always the identity at the origin $\mathbf{p} = 0$. This will become useful below, when we talk about the compositional and inverse compositional algorithms. (It also makes it easier to impose priors on the motions.)

⁸The derivatives of the 3D rotational motion model introduced in Subsection 2.2 are given in Subsection 4.3.

The derivatives in Table 3.1 are all fairly straightforward, except for the projective 2D motion (homography), which requires a per-pixel division to evaluate, c.f. (2.8), re-written here in its new parametric form as

$$x' = \frac{(1 + h_{00})x + h_{01}y + h_{02}}{h_{20}x + h_{21}y + 1} \quad \text{and} \quad y' = \frac{h_{10}x + (1 + h_{11})y + h_{12}}{h_{20}x + h_{21}y + 1}. \quad (3.54)$$

The Jacobian is therefore

$$\mathbf{J}_{x'} = \frac{\partial \mathbf{x}'}{\partial \mathbf{p}} = \frac{1}{D} \begin{bmatrix} x & y & 1 & 0 & 0 & 0 & -x'x & -x'y \\ 0 & 0 & 0 & x & y & 1 & -y'x & -y'y \end{bmatrix}, \quad (3.55)$$

where D is the denominator in (3.54), which depends on the current parameter settings (as do x' and y').

For parametric motion, the (Gauss–Newton) *Hessian* and *gradient-weighted residual vector* become

$$\mathbf{A} = \sum_i \mathbf{J}_{x'}^T(\mathbf{x}_i) [\nabla I_1^T(\mathbf{x}'_i) \nabla I_1(\mathbf{x}'_i)] \mathbf{J}_{x'}(\mathbf{x}_i) \quad (3.56)$$

and

$$\mathbf{b} = - \sum_i \mathbf{J}_{x'}^T(\mathbf{x}_i) [e_i \nabla I_1^T(\mathbf{x}'_i)]. \quad (3.57)$$

Note how the expressions inside the square brackets are the same ones evaluated for the simpler translational motion case (3.39)–(3.40).

Patch-Based Approximation The computation of the Hessian and residual vectors for parametric motion can be significantly more expensive than for the translational case. For parametric motion with n parameters and N pixels, the accumulation of \mathbf{A} and \mathbf{b} takes $O(n^2N)$ operations [9]. One way to reduce this by a significant amount is to divide the image up into smaller sub-blocks (patches) P_j and to only accumulate the simpler 2×2 quantities inside the square brackets at the pixel level [178],

$$\mathbf{A}_j = \sum_{i \in P_j} \nabla I_1^T(\mathbf{x}'_i) \nabla I_1(\mathbf{x}'_i), \quad (3.58)$$

$$\mathbf{b}_j = \sum_{i \in P_j} e_i \nabla I_1^T(\mathbf{x}'_i). \quad (3.59)$$

The full Hessian and residual can then be approximated as

$$\begin{aligned} \mathbf{A} &\approx \sum_j \mathbf{J}_{\mathbf{x}'}^T(\hat{\mathbf{x}}_j) \left[\sum_{i \in P_j} \nabla I_1^T(\mathbf{x}'_i) \nabla I_1(\mathbf{x}'_i) \right] \mathbf{J}_{\mathbf{x}'}(\hat{\mathbf{x}}_j) \\ &= \sum_j \mathbf{J}_{\mathbf{x}'}^T(\hat{\mathbf{x}}_j) \mathbf{A}_j \mathbf{J}_{\mathbf{x}'}(\hat{\mathbf{x}}_j) \end{aligned} \quad (3.60)$$

and

$$\mathbf{b} \approx - \sum_j \mathbf{J}_{\mathbf{x}'}^T(\hat{\mathbf{x}}_j) \left[\sum_{i \in P_j} e_i \nabla I_1^T(\mathbf{x}'_i) \right] = - \sum_j \mathbf{J}_{\mathbf{x}'}^T(\hat{\mathbf{x}}_j) \mathbf{b}_j, \quad (3.61)$$

where $\hat{\mathbf{x}}_j$ is the *center* of each patch P_j [178]. This is equivalent to replacing the true motion Jacobian with a piecewise-constant approximation. In practice, this works quite well. The relationship of this approximation to feature-based registration is discussed in Subsection 4.4.

Compositional Approach For a complex parametric motion such as a homography, the computation of the motion Jacobian becomes complicated, and may involve a per-pixel division. Szeliski and Shum [199] observed that this can be simplified by first warping the target image I_1 according to the current motion estimate $\mathbf{x}'(\mathbf{x}; \mathbf{p})$,

$$\tilde{I}_1(\mathbf{x}) = I_1(\mathbf{x}'(\mathbf{x}; \mathbf{p})), \quad (3.62)$$

and then comparing this *warped* image against the template $I_0(\mathbf{x})$,

$$\begin{aligned} E_{\text{LK-SS}}(\Delta \mathbf{p}) &= \sum_i [\tilde{I}_1(\tilde{\mathbf{x}}(\mathbf{x}_i; \Delta \mathbf{p})) - I_0(\mathbf{x}_i)]^2 \\ &\approx \sum_i [\tilde{J}_1(\mathbf{x}_i) \Delta \mathbf{p} + e_i]^2 \end{aligned} \quad (3.63)$$

$$= \sum_i [\nabla \tilde{I}_1(\mathbf{x}_i) \mathbf{J}_{\tilde{\mathbf{x}}}(\mathbf{x}_i) \Delta \mathbf{p} + e_i]^2. \quad (3.64)$$

Note that since the two images are assumed to be fairly similar, only an *incremental* parametric motion is required, i.e., the incremental motion can be evaluated around $\mathbf{p} = 0$, which can lead to considerable simplifications. For example, the Jacobian of the planar projective transform

(3.54) now becomes

$$\mathbf{J}_{\tilde{\mathbf{x}}} = \left. \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{p}} \right|_{\mathbf{p}=0} = \begin{bmatrix} x & y & 1 & 0 & 0 & 0 & -x^2 & -xy \\ 0 & 0 & 0 & x & y & 1 & -xy & -y^2 \end{bmatrix}. \quad (3.65)$$

Once the incremental motion $\tilde{\mathbf{x}}$ has been computed, it can be *prepended* to the previously estimated motion, which is easy to do for motions represented with transformation matrices, such as those given in Tables 2.1 and 3.1. Baker and Matthews [9] call this the *forward compositional* algorithm, since the target image is being re-warped, and the final motion estimates are being composed.

If the appearance of the warped and template images is similar enough, we can replace the gradient of $\tilde{I}_1(\mathbf{x})$ with the gradient of $I_0(\mathbf{x})$, as suggested previously in (3.42). This has potentially a big advantage in that it allows the pre-computation (and inversion) of the Hessian matrix \mathbf{A} given in (3.56). The residual vector \mathbf{b} (3.57) can also be partially pre-computed, i.e., the *steepest descent* images $\nabla I_0(\mathbf{x})\mathbf{J}_{\tilde{\mathbf{x}}}(\mathbf{x})$ can be pre-computed and stored for later multiplication with the $e(\mathbf{x}) = \tilde{I}_1(\mathbf{x}) - I_0(\mathbf{x})$ error images [9]. This idea was first suggested by Hager and Belhumeur [77] in what Baker and Matthews [9] call a *forward additive* scheme.

Baker and Matthews [9] introduce one more variant they call the *inverse compositional* algorithm. Rather than (conceptually) re-warping the warped target image $\tilde{I}_1(\mathbf{x})$, they instead warp the template image $I_0(\mathbf{x})$ and minimize

$$\begin{aligned} E_{\text{LK-BM}}(\Delta \mathbf{p}) &= \sum_i [\tilde{I}_1(\mathbf{x}_i) - I_0(\tilde{\mathbf{x}}(\mathbf{x}_i; \Delta \mathbf{p}))]^2 \\ &\approx \sum_i [\nabla I_0(\mathbf{x}_i)\mathbf{J}_{\tilde{\mathbf{x}}}(\mathbf{x}_i)\Delta \mathbf{p} - e_i]^2. \end{aligned} \quad (3.66)$$

This is identical to the forward warped algorithm (3.64) with the gradients $\nabla \tilde{I}_1(\mathbf{x})$ replaced by the gradients $\nabla I_0(\mathbf{x})$, except for the sign of e_i . The resulting update $\Delta \mathbf{p}$ is the *negative* of the one computed by the modified (3.64), and hence the *inverse* of the incremental transformation must be prepended to the current transform. Because the inverse compositional algorithm has the potential of pre-computing the inverse Hessian and the steepest descent images, this makes it the preferred approach of those surveyed in [9]. Figure 3.3, taken from [11],

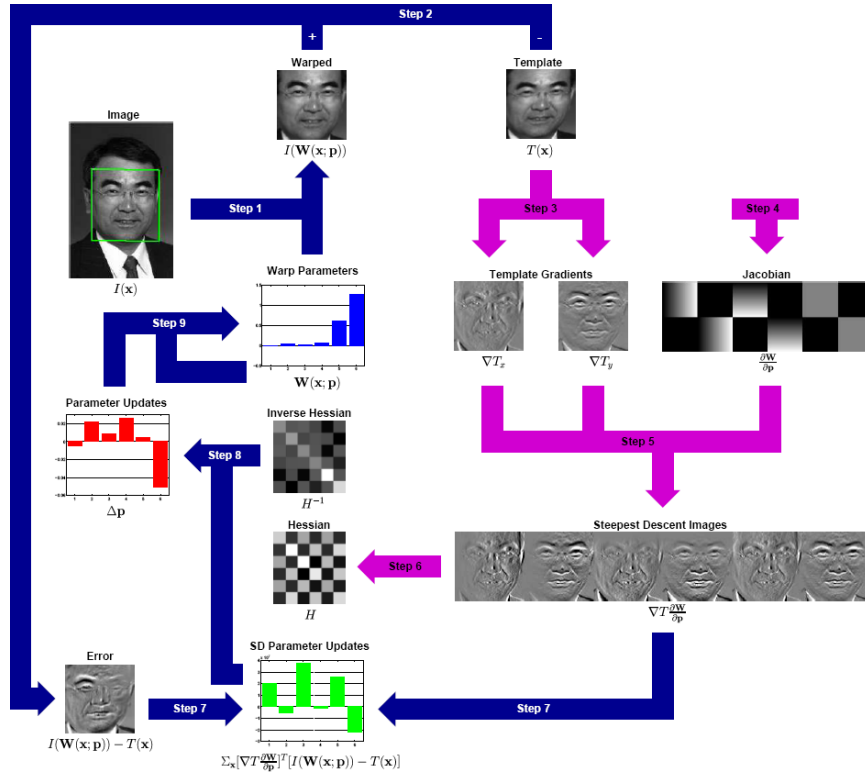


Fig. 3.3 A schematic overview of the inverse compositional algorithm (copied, with permission, from [11]). Steps 3–6 (light-color arrows) are performed once as a pre-computation. The main algorithm simply consists of iterating: image warping (Step 1), image differencing (Step 2), image dot products (Step 7), multiplication with the inverse of the Hessian (Step 8), and the update to the warp (Step 9). All of these steps can be performed efficiently.

beautifully shows all of the steps required to implement the inverse compositional algorithm.

Baker and Matthews [9] also discusses the advantage of using Gauss–Newton iteration (i.e., the first order expansion of the least squares, as above) vs. other approaches such as steepest descent and Levenberg–Marquardt. Subsequent parts of the series [10–12] discuss more advanced topics such as per-pixel weighting, pixel selection for efficiency, a more in-depth discussion of robust metrics and algorithms, linear appearance variations, and priors on parameters. They make for invaluable reading for anyone interested in implementing a highly tuned implementation of incremental image registration.

4

Feature-Based Registration

As I mentioned earlier, directly matching pixel intensities is just one possible approach to image registration. The other major approach is to first extract distinctive *features* from each image, to match these features to establish a global correspondence, and to then estimate the geometric transformation between the images. This kind of approach has been used since the early days of stereo matching [79, 80, 138] and has more recently gained popularity for image stitching applications [7, 30, 32, 35, 38, 129, 227].

In this section, I review methods for detecting distinctive points, for matching them, and for computing the image registration, including the 3D rotation model introduced in Subsection 2.2. I also discuss the relative advantages and disadvantages of direct and feature-based approaches.

4.1 Keypoint Detectors

As we saw in Subsection 3.4, the reliability of a motion estimate depends most critically on the size of the smallest eigenvalue of the image Hessian matrix, λ_0 [3]. This makes it a reasonable candidate

for finding points in the image that can be matched with high accuracy. (Older terminology in this field talked about “corner-like” features [138], but the modern usage is *keypoints*, *interest points*, or *salient points*.) Indeed, Shi and Tomasi [173] propose using this quantity to find *good features to track*, and then use a combination of translational and affine-based patch alignment to track such points through an image sequence.

Using a square patch with equal weighting may not be the best choice. Instead, a Gaussian weighting function can be used. Förstner [62] and Harris and Stephens [82] both proposed finding keypoints using such an approach. The Hessian and eigenvalue images can be efficiently evaluated using a sequence of filters and algebraic operations,

$$G_x(\mathbf{x}) = \frac{\partial}{\partial x} G_{\sigma_d}(\mathbf{x}) * I(\mathbf{x}), \quad (4.1)$$

$$G_y(\mathbf{x}) = \frac{\partial}{\partial y} G_{\sigma_d}(\mathbf{x}) * I(\mathbf{x}), \quad (4.2)$$

$$\mathbf{B}(\mathbf{x}) = \begin{bmatrix} G_x^2(\mathbf{x}) & G_x(\mathbf{x})G_y(\mathbf{x}) \\ G_x(\mathbf{x})G_y(\mathbf{x}) & G_y^2(\mathbf{x}) \end{bmatrix}, \quad (4.3)$$

$$\mathbf{A}(\mathbf{x}) = G_{\sigma_i}(\mathbf{x}) * \mathbf{B}(\mathbf{x}), \quad (4.4)$$

$$\lambda_{0,1}(\mathbf{x}) = \frac{a_{00} + a_{11} \mp \sqrt{(a_{00} - a_{11})^2 + a_{01}a_{10}}}{2}, \quad (4.5)$$

where G_{σ_d} is a noise-reducing pre-smoothing “derivative” filter of width σ_d , and G_{σ_i} is the integration filter whose scale σ_i controls the effective patch size. (The a_{ij} are the entries in the $\mathbf{A}(\mathbf{x})$ matrix, where I have dropped the (\mathbf{x}) for succinctness.) For example, Förstner [63] uses $\sigma_d = 0.7$ and $\sigma_i = 2$. Once the minimum eigenvalue image has been computed, local maxima can be found as potential keypoints.

The minimum eigenvalue is not the only quantity that can be used to find keypoints. A simpler quantity, proposed by Harris and Stephens [82] is

$$\det(\mathbf{A}) - \alpha \operatorname{tr}(\mathbf{A})^2 = \lambda_0\lambda_1 - \alpha(\lambda_0 + \lambda_1)^2 \quad (4.6)$$

with $\alpha = 0.06$. Triggs [206] suggest using the quantity

$$\lambda_0 - \alpha\lambda_1 \quad (4.7)$$

(say with $\alpha = 0.05$), which reduces the response at 1D edges, where aliasing errors sometimes affect the smaller eigenvalue. He also shows how the basic 2×2 Hessian can be extended to parametric motions to detect points that are also accurately localizable in scale and rotation. Brown *et al.* [32], on the other hand, use the harmonic mean,

$$\frac{\det \mathbf{A}}{\text{tr } \mathbf{A}} = \frac{\lambda_0 \lambda_1}{\lambda_0 + \lambda_1}, \quad (4.8)$$

which is a smoother function in the region where $\lambda_0 \approx \lambda_1$. Figure 4.1 shows isocontours of the various interest point operators (note that all the detectors require both eigenvalues to be large). Figure 4.2 shows the output of the multi-scale oriented patch detector of Brown *et al.* [32] at five different scales.

Schmid *et al.* [171] survey the vast literature on keypoint detection and perform some experimental comparisons to determine the *repeatability* of feature detectors, which is defined as the frequency with which keypoints detected in one image are found within $\epsilon = 1.5$ pixels of the corresponding location in a warped image. They also measure the *information content* available at each detected feature point, which they define as the entropy of a set of rotationally invariant local grayscale descriptors. Among the techniques they survey, they find that an *improved* version of the Harris operator with $\sigma_d = 1$ and $\sigma_i = 2$ works best.




Fig. 4.1 Isocontours of popular keypoint detection functions (taken from [31]). Each detector looks for points where the eigenvalues λ_0, λ_1 of $\mathbf{H} = \int_{\mathcal{R}} \nabla I \nabla I^T d\mathbf{x}$ are both large.




Fig. 4.2 Multi-scale Oriented Patches (MOPS) extracted at five pyramid levels (taken from [31]). The boxes show the feature orientation and the region from which the descriptor vectors are sampled.

More recently, feature detectors that are more invariant to scale [122, 132] and affine transformations [16, 99, 134, 169] have been proposed. These can be very useful when matching images that have different scales or aspects (e.g., for 3D object recognition). A simple way to achieve scale invariance is to look for scale-space maxima of Difference of Gaussian (DoG) [119, 122] or Harris corner [132, 206] detectors computed over a sub-octave pyramid, i.e., an image pyramid where the subsampling between adjacent levels is less than a factor of two. Lowe’s original paper [122] uses a half-octave ($\sqrt{2}$) pyramid, whereas Triggs [206] recommends using a quarter-octave ($\sqrt[4]{2}$). The area of feature point detectors and descriptors continues to be very active, with papers appearing every year at major computer vision conferences [17, 37, 106, 152, 160] – see the recent survey and comparison of affine region detectors by Mikolajczyk *et al.* [134].

Of course, keypoints are not the only kind of features that can be used for registering images. Zoghlami *et al.* [227] use line segments as well as point-like features to estimate homographies between pairs of images, whereas [14] use line segments with local correspondences along the edges to extract 3D structure and motion. Tuytelaars and Van Gool [209] use affine invariant regions to detect correspondences for wide baseline stereo matching. Matas *et al.* [127] detect *maximally stable*

regions, using an algorithm related to watershed detection, whereas Kadir *et al.* [100] detect salient regions where patch entropy and its rate of change with scale are locally maximal. (Corso and Hager [47] use a related technique to fit 2D oriented Gaussian kernels to homogeneous regions.) While these techniques can all be used to solve image registration problems, they will not be covered in more detail in this survey.

4.2 Feature Matching

After detecting the features (keypoints), we must *match* them, i.e., determine which features come from corresponding locations in different images. In some situations, e.g., for video sequences [173] or for stereo pairs that have been *rectified* [121, 170], the local motion around each feature point may be mostly translational. In this case, the error metrics introduced in Subsection 3.1 such as E_{SSD} or E_{NCC} can be used to directly compare the intensities in small patches around each feature point. (The comparative study by Mikolajczyk and Schmid [133] discussed below uses cross-correlation.) Because feature points may not be exactly located, a more accurate matching score can be computed by performing incremental motion refinement as described in Subsection 3.4, but this can be time consuming, and can sometimes even decrease performance [32].

If features are being tracked over longer image sequences, their appearance can undergo larger changes. In this case, it makes sense to compare appearances using an *affine* motion model. Shi and Tomasi [173] compare patches using a translational model between neighboring frames, and then use the location estimate produced by this step to initialize an affine registration between the patch in the current frame and the base frame where a feature was first detected. In fact, features are only detected infrequently, i.e., only in region where tracking has failed. In the usual case, an area around the current *predicted* location of the feature is searched with an incremental registration algorithm. This kind of algorithm is a *detect then track* approach, since detection occurs infrequently. It is appropriate for video sequences where the expected locations of feature points can be reasonably well predicted.

For larger motions, or for matching collections of images where the geometric relationship between them is unknown [30,169], a *detect then match* approach in which feature points are first detected in all images is more appropriate. Because the features can appear at different orientations or scales, a more *view invariant* kind of representation must be used. Mikolajczyk and Schmid [133] review some recently developed view-invariant local image descriptors and experimentally compare their performance.

The simplest method to compensate for in-plane rotations is to find a *dominant orientation* at each feature point location before sampling the patch or otherwise computing the descriptor. Brown *et al.* [32] use the direction of the average gradient orientation, computed within a small neighborhood of each feature point, whereas Lowe [122] (as well as Mikolajczyk and Schmid [133]) look for a peak in the local gradient orientation histogram. The descriptor can be made invariant to scale by only selecting feature points that are local maxima in scale space, as discussed in Subsection 4.1. Making the descriptors invariant to affine deformations (stretch, squash, and skew) is even harder [16, 169]. Mikolajczyk and Schmid [132] use the local second moment matrix around a feature point to define a canonical frame, whereas Corso and Hager [47] fit 2D oriented Gaussian kernels to homogeneous regions and store the weighted region statistics.

Among the local descriptors that Mikolajczyk and Schmid [133] compared, they found that David Lowe’s Scale Invariant Feature Transform (SIFT) [122] generally performed the best, followed by Freeman and Adelson’s steerable filters [64] and then cross-correlation (which could potentially be improved with an incremental refinement of location and pose – but see [32]). Differential invariants, whose descriptors are insensitive to changes in orientation by design, did not do as well.

SIFT features are computed by first estimating a local orientation using a histogram of the local gradient orientations, which is potentially more accurate than just the average orientation. Once the local frame has been established, gradients are copied into different orientation planes, and blurred re-sampled versions of these images are used as the features. This provides the descriptor with some insensitivity to small feature localization errors and geometric distortions [122].




Fig. 4.3 MOP descriptors are formed using an (8×8) sampling of bias/gain normalized intensity values, with a sample spacing of 5 pixels relative to the detection scale (taken from [31]). This low frequency sampling gives the features some robustness to keypoint location error, and is achieved by sampling at a higher pyramid level than the detection scale.

Steerable filters are combinations of derivative of Gaussian filters that permit the rapid computation of even and odd (symmetric and anti-symmetric) edge-like and corner-like features at all possible orientations [64]. Because they use reasonably broad Gaussians, they too are somewhat insensitive to localization and orientation errors.

For tasks that do not exhibit large amounts of foreshortening, such as image stitching, simple normalized intensity patches perform reasonably well and are simple to implement [32] (Figure 4.3). The field of feature descriptors continues to evolve rapidly, with newer results including a principal component analysis (PCA) of SIFT feature descriptors [105] and descriptors that include local color information [213].

Rapid Indexing and Matching The simplest way to find all corresponding feature points in an image pair is to compare all the features in one image against all the features in the other, using one of the local descriptors described above. Unfortunately, this is quadratic in the expected number of features, which makes it impractical for some applications.

More efficient matching algorithms can be devised using different kinds of *indexing schemes*, many of which are based on the idea of finding nearest neighbors in high-dimensional spaces. For example, Nene and Nayar [141] developed a technique they call *slicing* that uses a series of 1D binary searches to efficiently cull down a list of candidate points

that lie within a hypercube of the query point. They also provide a nice review of previous work in this area, including spatial data structures such as k -d trees [164]. Beis and Lowe [18] propose a Best-Bin-First (BBF) algorithm, which uses a modified search ordering for a k -d tree algorithm so that bins in feature space are searched in the order of their closest distance from query location. Shakhnarovich *et al.* [172] extend a previously developed technique called *locality sensitive hashing*, which uses unions of independently computed hashing functions, to be more sensitive to the distribution of points in parameter space, which they call *parameter-sensitive hashing*. Brown *et al.* [32] hash the first three (non-constant) Haar wavelets from an 8×8 image patch. Even more recently, Nister and Stewenius [143] use a *metric tree*, which consists of comparing feature descriptors to a small number of prototypes at each level in a hierarchy. Despite all of this promising work, the rapid computation of image feature correspondences is far from being a solved problem.

RANSAC and LMS Once an initial set of feature correspondences has been computed, we need to find a set that will produce a high-accuracy alignment. One possible approach is to simply compute a least squares estimate or to use a robustified (iteratively re-weighted) version of least squares, as discussed below (Subsection 4.3). However, in many cases, it is better to first find a good starting set of *inlier* correspondences, i.e., points that are all consistent with some particular motion estimate.¹

Two widely used solutions to this problem are called RANdom SAMple Consensus, or RANSAC for short [60] and *least median of squares* (LMS) [162]. Both techniques start by selecting (at random) a subset of k correspondences, which is then used to compute a motion estimate \mathbf{p} , as described in Subsection 4.3. The *residuals* of the full set of correspondences are then computed as

$$\mathbf{r}_i = \tilde{\mathbf{x}}'_i(\mathbf{x}_i; \mathbf{p}) - \hat{\mathbf{x}}'_i, \quad (4.9)$$

where $\tilde{\mathbf{x}}'_i$ are the *estimated* (mapped) locations, and $\hat{\mathbf{x}}'_i$ are the sensed (detected) feature point locations.

¹For direct estimation methods, hierarchical (coarse-to-fine) techniques are often used to lock onto the *dominant motion* in a scene [20, 21].

The RANSAC technique then counts the number of *inliers* that are within ϵ of their predicted location, i.e., whose $\|\mathbf{r}_i\| \leq \epsilon$. (The ϵ value is application dependent, but often is around 1–3 pixels.) Least median of squares finds the median value of the $\|\mathbf{r}_i\|$ values.

The random selection process is repeated S times, and the sample set with largest number of inliers (or with the smallest median residual) is kept as the final solution. Either the initial parameter guess \mathbf{p} or the full set of computed inliers is then passed on to the next data fitting stage. In a more recently developed version of RANSAC called PROSAC (PROgressive SAmple Consensus), random samples are initially added from the most “confident” matches, thereby speeding up the process of finding a (statistically) likely good set of inliers [44].

To ensure that the random sampling has a good chance of finding a true set of inliers, a sufficient number of trials S must be tried. Let p be the probability that any given correspondence is valid, and P be the total probability of success after S trials. The likelihood in one trial that all k random samples are inliers is p^k . Therefore, the likelihood that S such trials will all fail is

$$1 - P = (1 - p^k)^S, \quad (4.10)$$

and the required minimum number of trials is

$$S = \frac{\log(1 - P)}{\log(1 - p^k)}. \quad (4.11)$$

Stewart [188] gives the following examples of the required number of trials S to attain a 99% probability of success:

k	p	S
3	0.5	35
6	0.6	97
6	0.5	293

As you can see, the number of trials grows quickly with the number of sample points used. This provides a strong incentive to use the *minimum* number of sample points k possible for any given trial, which in practice is how RANSAC is normally used.

4.3 Geometric Registration

Once we have computed a set of matched feature point correspondences, the next step is to estimate the motion parameters \mathbf{p} that best register the two images. The usual way to do this is to use least squares, i.e., to minimize the sum of squared residuals given by (4.9),

$$E_{\text{LS}} = \sum_i \|\mathbf{r}_i\|^2 = \|\tilde{\mathbf{x}}'_i(\mathbf{x}_i; \mathbf{p}) - \hat{\mathbf{x}}'_i\|^2. \quad (4.12)$$

Many of the motion models presented in Subsection 2, i.e., translation, similarity, and affine, have a *linear* relationship between the motion and the unknown parameters \mathbf{p} .² In this case, a simple linear regression (least squares) using normal equations $\mathbf{A}\mathbf{p} = \mathbf{b}$ works well.

Uncertainty Weighting and Robust Regression The above least squares formulation assumes that all feature points are matched with the same accuracy. This is often not the case, since certain points may fall in more textured regions than others. If we associate a variance estimate σ_i^2 with each correspondence, we can minimize *weighted least squares* instead,

$$E_{\text{WLS}} = \sum_i \sigma_i^{-2} \|\mathbf{r}_i\|^2. \quad (4.13)$$

As discussed in Subsection 3.4, a covariance estimate for patch-based matching can be obtained by multiplying the inverse of the Hessian with the per-pixel noise estimate (3.43). Weighting each squared residual by the inverse covariance $\Sigma_i^{-1} = \sigma_n^{-2} \mathbf{A}_i$ (which is called the *information matrix*), we obtain

$$E_{\text{CWLS}} = \sum_i \|\mathbf{r}_i\|_{\Sigma_i^{-1}}^2 = \sum_i \mathbf{r}_i^T \Sigma_i^{-1} \mathbf{r}_i = \sum_i \sigma_n^{-2} \mathbf{r}_i^T \mathbf{A}_i \mathbf{r}_i, \quad (4.14)$$

where \mathbf{A}_i is the *patch Hessian* (3.58).

If there are outliers among the feature-based correspondences (and there almost always are), it is better to use a robust version of least

²2D Euclidean motion can be estimated with a linear algorithm by first estimating the cosine and sine entries independently, and then normalizing them so that their magnitude is 1.

squares, even if an initial RANSAC or MLS stage has been used to select plausible inliers. The robust least squares cost metric (analogous to (3.2)) is then

$$E_{\text{RLS}}(\mathbf{u}) = \sum_i \rho(\|\mathbf{r}_i\|_{\Sigma_i^{-1}}). \quad (4.15)$$

As before, a commonly used approach to minimize this quantity is to use iteratively re-weighted least squares, as described in Subsection 3.4.

Homography Update For *nonlinear* measurement equations such as the homography given in (3.54), re-written here as

$$\begin{aligned} \hat{x}' &= \frac{(1 + h_{00})x + h_{01}y + h_{02}}{h_{20}x + h_{21}y + 1} \quad \text{and} \\ \hat{y}' &= \frac{h_{10}x + (1 + h_{11})y + h_{12}}{h_{20}x + h_{21}y + 1}, \end{aligned} \quad (4.16)$$

an iterative solution is required to obtain accurate results. An initial guess for the eight unknowns $\{h_{00}, \dots, h_{21}\}$ can be obtained by multiplying both sides of the equations through by the denominator, which yields the linear set of equations,

$$\begin{bmatrix} \hat{x}' - x \\ \hat{y}' - y \end{bmatrix} = \begin{bmatrix} x & y & 1 & 0 & 0 & 0 & -\hat{x}'x & -\hat{x}'y \\ 0 & 0 & 0 & x & y & 1 & -\hat{y}'x & -\hat{y}'y \end{bmatrix} \begin{bmatrix} h_{00} \\ \vdots \\ h_{21} \end{bmatrix}. \quad (4.17)$$

However, this is not optimal from a statistical point of view, since the denominator can vary quite a bit from point to point.

One way to compensate for this is to *re-weight* each equation by the inverse of current estimate of the denominator, D ,

$$\frac{1}{D} \begin{bmatrix} \hat{x}' - x \\ \hat{y}' - y \end{bmatrix} = \frac{1}{D} \begin{bmatrix} x & y & 1 & 0 & 0 & 0 & -\hat{x}'x & -\hat{x}'y \\ 0 & 0 & 0 & x & y & 1 & -\hat{y}'x & -\hat{y}'y \end{bmatrix} \begin{bmatrix} h_{00} \\ \vdots \\ h_{21} \end{bmatrix}. \quad (4.18)$$

While this may at first seem to be the exact same set of equations as (4.17), because least squares is being used to solve the over-determined set of equations, the weightings *do* matter and produce a different set of normal equations that performs better in practice (with noisy data).

The most principled way to do the estimation, however, is to directly minimize the squared residual equations (4.9) using the Gauss–Newton approximation, i.e., performing a first-order Taylor series expansion in \mathbf{p} , which yields,

$$\hat{\mathbf{x}}'_i - \tilde{\mathbf{x}}'_i(\mathbf{x}_i; \mathbf{p}) = \mathbf{J}_{\mathbf{x}'_i} \Delta \mathbf{p} \quad (4.19)$$

or

$$\begin{bmatrix} \hat{x}' - \tilde{x}' \\ \hat{y}' - \tilde{y}' \end{bmatrix} = \frac{1}{D} \begin{bmatrix} x & y & 1 & 0 & 0 & 0 & -\tilde{x}'x & -\tilde{x}'y \\ 0 & 0 & 0 & x & y & 1 & -\tilde{y}'x & -\tilde{y}'y \end{bmatrix} \begin{bmatrix} \Delta h_{00} \\ \vdots \\ \Delta h_{21} \end{bmatrix}. \quad (4.20)$$

While this looks similar to (4.18), it differs in two important respects. First, the left-hand side consists of unweighted *prediction errors* rather than pixel displacements, and the solution vector is a *perturbation* to the parameter vector \mathbf{p} . Second, the quantities inside $\mathbf{J}_{\mathbf{x}'_i}$ involve *predicted* feature locations (\tilde{x}', \tilde{y}') instead of *sensed*, feature locations (\hat{x}', \hat{y}') . Both of these are subtle and yet they lead to an algorithm that, when combined with proper checking for downhill steps (as in the Levenberg–Marquardt algorithm), will converge to a minimum. (Iterating the (4.18) equations is not guaranteed to do so, since it is not minimizing a well-defined energy function.)

The above formula is analogous to the *additive* algorithm for direct registration since the change to the *full* transformation is being computed in Subsection 3.5. If we prepend an incremental homography to the current homography instead, i.e., we use a *compositional* algorithm, we get $D = 1$ (since $\mathbf{p} = 0$) and the above formula simplifies to

$$\begin{bmatrix} \hat{x}' - x \\ \hat{y}' - y \end{bmatrix} = \begin{bmatrix} x & y & 1 & 0 & 0 & 0 & -x^2 & -xy \\ 0 & 0 & 0 & x & y & 1 & -xy & -y^2 \end{bmatrix} \begin{bmatrix} \Delta h_{00} \\ \vdots \\ \Delta h_{21} \end{bmatrix}, \quad (4.21)$$

where I have replaced (\tilde{x}', \tilde{y}') with (x, y) for conciseness. (Notice how this results in the same Jacobian as (3.65).)

Rotational Panorama Update As described in Subsection 2.2, representing the alignment of images in a panorama using a collection

of rotation matrices and focal lengths results in a much more stable estimation problem than directly using homographies [194, 199]. Given this representation, how do we update the rotation matrices to best align two overlapping images?

Recall from (2.18) to (2.19) that the equations relating two views can be written as

$$\tilde{\mathbf{x}}_1 \sim \tilde{\mathbf{H}}_{10} \tilde{\mathbf{x}}_0 \quad \text{with} \quad \tilde{\mathbf{H}}_{10} = \mathbf{K}_1 \mathbf{R}_{10} \mathbf{K}_0^{-1}, \quad (4.22)$$

where $\mathbf{K}_k = \text{diag}(f_k, f_k, 1)$ is the calibration matrix and $\mathbf{R}_{10} = \mathbf{R}_1 \mathbf{R}_0^{-1}$ is rotation *between* the two views. The best way to update \mathbf{R}_{10} is to prepend an *incremental* rotation matrix $\mathbf{R}(\vec{\omega})$ to the current estimate \mathbf{R}_{10} [178, 199],

$$\begin{aligned} \tilde{\mathbf{H}}(\vec{\omega}) &= \mathbf{K}_1 \mathbf{R}(\vec{\omega}) \mathbf{R}_{10} \mathbf{K}_0^{-1} \\ &= [\mathbf{K}_1 \mathbf{R}(\vec{\omega}) \mathbf{K}_1^{-1}] [\mathbf{K}_1 \mathbf{R}_{10} \mathbf{K}_0^{-1}] = \mathbf{D} \tilde{\mathbf{H}}_{10}. \end{aligned} \quad (4.23)$$

Note that here I have written the update rule in the *compositional* form, where the incremental update \mathbf{D} is *prepended* to the current homography $\tilde{\mathbf{H}}_{10}$. Using the small-angle approximation to $\mathbf{R}(\vec{\omega})$ given in (2.26), we can write the incremental update matrix as

$$\begin{aligned} \mathbf{D} &= \mathbf{K}_1 \mathbf{R}(\vec{\omega}) \mathbf{K}_1^{-1} \approx \mathbf{K}_1 (\mathbf{I} + [\vec{\omega}]_{\times}) \mathbf{K}_1^{-1} \\ &= \begin{bmatrix} 1 & -\omega_z & f_1 \omega_y \\ \omega_z & 1 & -f_1 \omega_x \\ -\omega_y/f_1 & \omega_x/f_1 & 1 \end{bmatrix}. \end{aligned} \quad (4.24)$$

Notice how there is now a nice one-to-one correspondence between the entries in the \mathbf{D} matrix and the h_{00}, \dots, h_{21} parameters used in Table 3.1 and (4.16), i.e.,

$$\begin{aligned} &(h_{00}, h_{01}, h_{02}, h_{10}, h_{11}, h_{12}, h_{20}, h_{21}) \\ &= (0, -\omega_z, f_1 \omega_y, \omega_z, 0, -f_1 \omega_x, -\omega_y/f_1, \omega_x/f_1). \end{aligned} \quad (4.25)$$

We can therefore apply the chain rule to (4.21) and (4.25) to obtain

$$\begin{bmatrix} \hat{x}' - x \\ \hat{y}' - y \end{bmatrix} = \begin{bmatrix} -xy/f_1 & f_1 + x^2/f_1 & -y \\ -(f_1 + y^2/f_1) & xy/f_1 & x \end{bmatrix} \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix}, \quad (4.26)$$




Fig. 4.4 Four images taken with a hand-held model registered using a 3D rotation motion model (from [199]). Notice how the homographies, rather than being arbitrary, have a well-defined keystone shape whose width increases away from the origin.

which give us the linearized update equations needed to estimate $\vec{\omega} = (\omega_x, \omega_y, \omega_z)$.³ Notice that this update rule depends on the focal length f_1 of the *target* view, and is independent of the focal length f_0 of the *template* view. This is because the compositional algorithm essentially makes small perturbations to the target. Once the incremental rotation vector $\vec{\omega}$ has been computed, the \mathbf{R}_1 rotation matrix can be updated using $\mathbf{R}_1 \leftarrow \mathbf{R}(\vec{\omega})\mathbf{R}_1$.

The formulas for updating the focal length estimates are a little more involved, and are given in [178]. I will not repeat them here, since an alternative update rule, based on minimizing the difference between back-projected 3D rays, will be given in Subsection 5.1. Figure 4.4 shows the alignment of four images under the 3D rotation motion model.

Focal Length Initialization To initialize the 3D rotation model, we need to simultaneously estimate the focal length(s) of the camera(s) and an initial guess for a rotation matrix. This can be obtained directly

³This is the same as the rotational component of instantaneous rigid flow [20] and the same as the update equations given in [178, 199].

from a homography-based (planar perspective) alignment $\tilde{\mathbf{H}}_{10}$, using the formulas first presented in [199].

Using the simplified form of the calibration matrices $\mathbf{K}_k = \text{diag}(f_k, f_k, 1)$ first used in (2.19), we can re-write (4.22) as

$$\mathbf{R}_{10} \sim \mathbf{K}_1^{-1} \tilde{\mathbf{H}}_{10} \mathbf{K}_0 \sim \begin{bmatrix} h_{00} & h_{01} & f_0^{-1} h_{02} \\ h_{10} & h_{11} & f_0^{-1} h_{12} \\ f_1 h_{20} & f_1 h_{21} & f_0^{-1} f_1 h_{22} \end{bmatrix}, \quad (4.27)$$

where the h_{ij} are the elements of $\tilde{\mathbf{H}}_{10}$.

Using the orthonormality properties of the rotation matrix \mathbf{R}_{10} and the fact that the right-hand side of (4.27) is known only up to a scale, we obtain

$$h_{00}^2 + h_{01}^2 + f_0^{-2} h_{02}^2 = h_{10}^2 + h_{11}^2 + f_0^{-2} h_{12}^2 \quad (4.28)$$

and

$$h_{00} h_{10} + h_{01} h_{11} + f_0^{-2} h_{02} h_{12} = 0. \quad (4.29)$$

From this, we can compute estimates for f_0 of

$$f_0^2 = \frac{h_{12}^2 - h_{02}^2}{h_{00}^2 + h_{01}^2 - h_{10}^2 - h_{11}^2} \quad \text{if } h_{00}^2 + h_{01}^2 \neq h_{10}^2 + h_{11}^2 \quad (4.30)$$

or

$$f_0^2 = -\frac{h_{02} h_{12}}{h_{00} h_{10} + h_{01} h_{11}} \quad \text{if } h_{00} h_{10} \neq -h_{01} h_{11}. \quad (4.31)$$

(Note that the equations given in [199] are erroneous; the correct equations can be found in [178].) If neither of these conditions holds, we can also take the dot products between the first (or second) row and the third one. Similar result can be obtained for f_1 as well by analyzing the columns of $\tilde{\mathbf{H}}_{10}$. If the focal length is the same for both images, we can take the geometric mean of f_0 and f_1 as the estimated focal length $f = \sqrt{f_1 f_0}$. When multiple estimates of f are available, e.g., from different homographies, the median value can be used as the final estimate.




Fig. 4.5 Gap closing: (a) a gap is visible when the focal length is wrong ($f = 510$); (b) no gap is visible for the correct focal length ($f = 468$).

Gap Closing The techniques presented in this section can be used to estimate a series of rotation matrices and focal lengths, which can be chained together to create large panoramas. Unfortunately, because of accumulated errors, this approach will rarely produce a closed 360° panorama. Instead, there will invariably be either a gap or an overlap (Figure 4.5).

We can solve this problem by matching the first image in the sequence with the last one. The difference between the two rotation matrix estimates associated with this frame indicates the amount of misregistration. This error can be distributed evenly across the whole sequence by taking the quotient of the two quaternions associated with these rotations and dividing this “error quaternion” by the number of images in the sequence (assuming relatively constant inter-frame rotations). We can also update the estimated focal length based on the amount of misregistration. To do this, we first convert the error quaternion into a *gap angle*, θ_g . We then update the focal length using the equation $f' = f(1 - \theta_g/360^\circ)$.

Figure 4.5a shows the end of registered image sequence and the first image. There is a big gap between the last image and the first which are in fact the same image. The gap is 32° because the wrong estimate of focal length ($f = 510$) was used. Figure 4.5b shows the registration

after closing the gap with the correct focal length ($f = 468$). Notice that both mosaics show very little visual misregistration (except at the gap), yet Figure 4.5a has been computed using a focal length which has 9% error. Related approaches have been developed by [83,103,130,186] to solve the focal length estimation problem using pure panning motion and cylindrical images.

Unfortunately, this particular gap-closing heuristic only works for the kind of “one-dimensional” panorama where the camera is continuously turning in the same direction. In Section 5, I describe a different approach to removing gaps and overlaps that works for arbitrary camera motions.

4.4 Direct vs. Feature-Based Alignment

Given that there exist these two alternative approaches to aligning images, which is preferable?

I used to be firmly in the direct matching camp [94]. Early feature-based methods seemed to get confused in regions that were either too textured or not textured enough. The features would often be distributed unevenly over the images, thereby failing to match image pairs that should have been aligned. Furthermore, establishing correspondences relied on simple cross-correlation between patches surrounding the feature points, which did not work well when the images were rotated or had foreshortening due to homographies.

Today, feature detection and matching schemes are remarkably robust, and can even be used for known object recognition from widely separated views [122]. Features not only respond to regions of high “corneriness” [62,82], but also to “blob-like” regions [122], as well as uniform areas [209]. Furthermore, because they operate in scale space and use a dominant orientation (or orientation invariant descriptors), they can match images that differ in scale, orientation, and even foreshortening. My own recent experience in working with feature-based approaches is that if the features are well distributed over the image and the descriptors reasonably designed for repeatability, enough correspondences to permit image stitching can usually be found [32].

The other major reason I used to prefer direct methods was that they make optimal use of the information available in image alignment, since they measure the contribution of *every* pixel in the image. Furthermore, assuming a Gaussian noise model (or a robustified version of it), they properly weight the contribution of different pixels, e.g., by emphasizing the contribution of high-gradient pixels. (See Baker *et al.* [11], who suggest that adding even more weight at strong gradients is preferable because of noise in the gradient estimates.) One could argue that for a blurry image with only slowly varying gradients, a direct approach will find an alignment, whereas a feature detector will fail to find anything. However, such images rarely occur in practice in consumer imaging, and the use of scale-space features means that some features can be found at lower resolutions.

The biggest disadvantage of direct techniques is that they have a limited range of convergence. Even though they can be used in a hierarchical (coarse-to-fine) estimation framework, in practice it is hard to use more than two or three levels of a pyramid before important details start to be blurred away. For matching sequential frames in a video, the direct approach can usually be made to work. However, for matching partially overlapping images in photo-based panoramas, they fail too often to be useful. Our older systems for image stitching [194,199] relied on Fourier-based correlation of cylindrical images and motion prediction to automatically align images, but had to be corrected by hand for more complex sequences. Our newer system [31,32] uses features and has a good success rate at automatically stitching panoramas without any user intervention.

Is there no rôle then for direct registration? I believe there is. Once a pair of images has been aligned with a feature-based approach, we can warp the two images to a common reference frame and re-compute a more accurate estimate using patch-based alignment. Notice how there is a close correspondence between the patch-based approximation to direct alignment given in (3.60)–(3.61) and the inverse covariance weighted feature-based least squares error metric (4.14).

In fact, if we divide the template images up into patches and place an imaginary “feature point” at the center of each patch, the two approaches return exactly the same answer (assuming that the correct

correspondences are found in each case). However, for this approach to succeed, we still have to deal with “outliers,” i.e., regions that do not fit the selected motion model due to either parallax (Subsection 5.2) or moving objects (Subsection 6.2). While a feature-based approach may make it somewhat easier to reason about outliers (features can be classified as inliers or outliers), the patch-based approach, since it establishes correspondences more densely, is potentially more useful for removing local mis-registration (parallax), as we discuss in Subsection 5.2.

5

Global Registration

So far, I have discussed how to register pairs of images using both direct and feature-based methods using a variety of motion models. In most applications, we are given more than a single pair of images to register. The goal is then to find a *globally consistent* set of alignment parameters that minimize the mis-registration between all pairs of images [46, 167, 178, 199]. To do this, we need to extend the pairwise matching criteria (3.1), (3.51), and (4.12) to a global energy function that involves all of the per-image pose parameters (Subsection 5.1). Once we have computed the global alignment, we often need to perform *local adjustments* such as *parallax removal* to reduce double images and blurring due to local mis-registrations (Subsection 5.2). Finally, if we are given an unordered set of images to register, we need to discover which images go together to form one or more panoramas. This process of *panorama recognition* is described in Subsection 5.3.

5.1 Bundle Adjustment

One way to register a large number of images is to add new images to the panorama one at a time, aligning the most recent image with

the previous ones already in the collection [199], and discovering, if necessary, which images it overlaps [167]. In the case of 360° panoramas, accumulated error may lead to the presence of a *gap* (or excessive overlap) between the two ends of the panorama, which can be fixed by stretching the alignment of all the images using a process called *gap closing* [199]. However, a better alternative is to simultaneously align all the images together using a least squares framework to correctly distribute any mis-registration errors.

The process of simultaneously adjusting pose parameters for a large collection of overlapping images is called *bundle adjustment* in the photogrammetry community [207]. In computer vision, it was first applied to the general structure from motion problem [197] and then later specialized for panoramic image stitching [46, 167, 178].

In this section, I formulate the problem of global alignment using a feature-based approach, since this results in a simpler system. An equivalent direct approach can be obtained either by dividing images into patches and creating a virtual feature correspondence for each one (as discussed in Subsection 4.4 and [178]), or by replacing the per-feature error metrics with per-pixel metrics.

Consider the feature-based alignment problem given in (4.12), i.e.,

$$E_{\text{pairwise-LS}} = \sum_i \|\mathbf{r}_i\|^2 = \|\tilde{\mathbf{x}}'_i(\mathbf{x}_i; \mathbf{p}) - \hat{\mathbf{x}}'_i\|^2. \quad (5.1)$$

For multi-image alignment, instead of having a single collection of pairwise feature correspondences, $\{(\mathbf{x}_i, \hat{\mathbf{x}}'_i)\}$, we have a collection of n features, with the location of the i th feature point in the j th image denoted by \mathbf{x}_{ij} and its scalar confidence (inverse variance) denoted by c_{ij} .¹ Each image also has some associated *pose* parameters.

In this section, I assume that this pose consists of a rotation matrix \mathbf{R}_j and a focal length f_j , although formulations in terms of homographies are also possible [167, 176]. The equation mapping a 3D point \mathbf{x}_i into a point \mathbf{x}_{ij} in frame j can be re-written from (2.15) to (2.19) as

$$\tilde{\mathbf{x}}_{ij} \sim \mathbf{K}_j \mathbf{R}_j \mathbf{x}_i \quad \text{and} \quad \mathbf{x}_i \sim \mathbf{R}_j^{-1} \mathbf{K}_j^{-1} \tilde{\mathbf{x}}_{ij}, \quad (5.2)$$

¹ Features that not seen in image j have $c_{ij} = 0$. We can also use 2×2 inverse covariance matrices Σ_{ij}^{-1} in place of c_{ij} , as shown in (4.14).

where $\mathbf{K}_j = \text{diag}(f_j, f_j, 1)$ is the simplified form of the calibration matrix. The motion mapping a point \mathbf{x}_{ij} from frame j into a point \mathbf{x}_{ik} in frame k is similarly given by

$$\tilde{\mathbf{x}}_{ik} \sim \tilde{\mathbf{H}}_{kj} \tilde{\mathbf{x}}_{ij} = \mathbf{K}_k \mathbf{R}_k \mathbf{R}_j^{-1} \mathbf{K}_j^{-1} \tilde{\mathbf{x}}_{ij}. \quad (5.3)$$

Given an initial set of $\{(\mathbf{R}_j, f_j)\}$ estimates obtained from chaining pairwise alignments, how do we refine these estimates?

One approach is to directly extend the pairwise energy $E_{\text{pairwise-LS}}$ (5.1) to a multiview formulation,

$$E_{\text{all-pairs-2D}} = \sum_i \sum_{jk} c_{ij} c_{ik} \|\tilde{\mathbf{x}}_{ik}(\hat{\mathbf{x}}_{ij}; \mathbf{R}_j, f_j, \mathbf{R}_k, f_k) - \hat{\mathbf{x}}_{ik}\|^2, \quad (5.4)$$

where the $\tilde{\mathbf{x}}_{ik}$ function is the *predicted* location of feature i in frame k given by (5.3), $\hat{\mathbf{x}}_{ij}$ is the *observed* location, and the “2D” in the subscript indicates that an image-plane error is being minimized [176]. Note that since $\tilde{\mathbf{x}}_{ik}$ depends on the $\hat{\mathbf{x}}_{ij}$ observed value, we actually have an *errors-in-variable* problem, which in principle requires more sophisticated techniques than least squares to solve. However, in practice, if we have enough features, we can directly minimize the above quantity using regular nonlinear least squares and obtain an accurate multi-frame alignment.²

While this approach works well in practice, it suffers from two potential disadvantages. First, since a summation is taken over all pairs with corresponding features, features that are observed many times get over-weighted in the final solution. (In effect, a feature observed m times gets counted $\binom{m}{2}$ times instead of m times.) Second, the derivatives of $\tilde{\mathbf{x}}_{ik}$ w.r.t. the $\{(\mathbf{R}_j, f_j)\}$ are a little cumbersome, although using the incremental correction to \mathbf{R}_j introduced in Subsection 2.2 makes this more tractable.

An alternative way to formulate the optimization is to use true bundle adjustment, i.e., to solve not only for the pose parameters $\{(\mathbf{R}_j, f_j)\}$ but also for the 3D point positions $\{\mathbf{x}_i\}$,

$$E_{\text{BA-2D}} = \sum_i \sum_j c_{ij} \|\tilde{\mathbf{x}}_{ij}(\mathbf{x}_i; \mathbf{R}_j, f_j) - \hat{\mathbf{x}}_{ij}\|^2, \quad (5.5)$$

²While there exists an overall pose ambiguity in the solution, i.e., all the \mathbf{R}_j can be post-multiplied by an arbitrary rotation \mathbf{R}_g , a well-conditioned nonlinear least squares algorithm such as Levenberg–Marquardt will handle this degeneracy without trouble.

where $\tilde{\mathbf{x}}_{ij}(\mathbf{x}_i; \mathbf{R}_j, f_j)$ is given by (5.2). The disadvantage of full bundle adjustment is that there are more variables to solve for, so both, each iteration and the overall convergence may be slower. (Imagine how the 3D points need to “shift” each time some rotation matrices are updated.) However, the computational complexity of each linearized Gauss–Newton step can be reduced using sparse matrix techniques [85, 197, 207].

An alternative formulation is to minimize the error in 3D projected ray directions [178], i.e.,

$$E_{\text{BA-3D}} = \sum_i \sum_j c_{ij} \|\tilde{\mathbf{x}}_i(\hat{\mathbf{x}}_{ij}; \mathbf{R}_j, f_j) - \mathbf{x}_i\|^2, \quad (5.6)$$

where $\tilde{\mathbf{x}}_i(\mathbf{x}_{ij}; \mathbf{R}_j, f_j)$ is given by the second half of (5.2). This in itself has no particular advantage over (5.5). In fact, since errors are being minimized in 3D ray space, there is a bias toward estimating longer focal lengths, since the angles between rays become smaller as f increases.

However, if we eliminate the 3D rays \mathbf{x}_i , we can derive a pairwise energy formulated in 3D ray space [178],

$$E_{\text{all-pairs-3D}} = \sum_i \sum_{jk} c_{ij} c_{ik} \|\tilde{\mathbf{x}}_i(\hat{\mathbf{x}}_{ij}; \mathbf{R}_j, f_j) - \tilde{\mathbf{x}}_i(\hat{\mathbf{x}}_{ik}; \mathbf{R}_k, f_k)\|^2. \quad (5.7)$$

This results in the simplest set of update equations [178], since the f_k can be folded into the creation of the homogeneous coordinate vector as in (2.21). Thus, even though this formula over-weights features that occur more frequently, it is the method used both by Shum and Szeliski [178] and in our current work [32]. To reduce the bias toward longer focal lengths, I multiply each residual (3D error) by $\sqrt{f_j f_k}$, which is similar to projecting the 3D rays into a “virtual camera” of intermediate focal length, and which seems to work well in practice.

Up Vector Selection As mentioned above, there exists a global ambiguity in the pose of the 3D cameras computed by the above methods. While this may not appear to matter, people have a preference for the final stitched image being “upright” rather than twisted or tilted. More concretely, people are used to seeing photographs displayed so

that the vertical (gravity) axis points straight up in the image. Consider how you usually shoot photographs: while you may pan and tilt the camera any way, you usually keep vertical scene lines parallel to the vertical edge of the image. In other words, the horizontal edge of your camera (its x -axis) usually stays parallel to the ground plane (perpendicular to the world gravity direction).

Mathematically, this constraint on the rotation matrices can be expressed as follows. Recall from (5.2) that the $3\text{D} \rightarrow 2\text{D}$ projection is given by

$$\tilde{\mathbf{x}}_{ik} \sim \mathbf{K}_k \mathbf{R}_k \mathbf{x}_i. \quad (5.8)$$

We wish to post-multiply each rotation matrix \mathbf{R}_k by a global rotation \mathbf{R}_g such that the projection of the global y -axis, $\hat{\mathbf{j}} = (0, 1, 0)$ is perpendicular to the image x -axis, $\hat{\mathbf{i}} = (1, 0, 0)$.³

This constraint can be written as

$$\hat{\mathbf{i}}^T \mathbf{R}_k \mathbf{R}_g \hat{\mathbf{j}} = 0 \quad (5.9)$$

(note that the scaling by the calibration matrix is irrelevant here). This is equivalent to requiring that the first row of \mathbf{R}_k , $\mathbf{r}_{k0} = \hat{\mathbf{i}}^T \mathbf{R}_k$ be perpendicular to the second column of \mathbf{R}_g , $\mathbf{r}_{g1} = \mathbf{R}_g \hat{\mathbf{j}}$. This set of constraints (one per input image) can be written as a least squares problem,

$$\mathbf{r}_{g1} = \arg \min_r \sum_k (\mathbf{r}^T \mathbf{r}_{k0})^2 = \arg \min_r \mathbf{r}^T \left[\sum_k \mathbf{r}_{k0} \mathbf{r}_{k0}^T \right] \mathbf{r}. \quad (5.10)$$

Thus, \mathbf{r}_{g1} is the smallest eigenvector of the *scatter* or *moment* matrix spanned by the individual camera rotation x -vectors, which should generally be of the form $(c, 0, s)$ when the cameras are upright.

To fully specify the \mathbf{R}_g global rotation, we need to specify one additional constraint. This is related to the *view selection* problem discussed in Subsection 6.1. One simple heuristic is to prefer the average z -axis

³Note that here we use the convention common in computer graphics that the vertical world axis corresponds to y . This is a natural choice if we wish the rotation matrix associated with a “regular” image taken horizontally to be the identity, rather than a 90° rotation around the x -axis.

of the individual rotation matrices, $\bar{\mathbf{k}} = \sum_k \hat{\mathbf{k}}^T \mathbf{R}_k$ to be close to the world z -axis, $\mathbf{r}_{g2} = \mathbf{R}_g \hat{\mathbf{k}}$. We can therefore compute the full rotation matrix \mathbf{R}_g in three steps:

- (1) $\mathbf{r}_{g1} = \min$ eigenvector $(\sum_k \mathbf{r}_{k0} \mathbf{r}_{k0}^T)$;
- (2) $\mathbf{r}_{g0} = \mathcal{N}((\sum_k \mathbf{r}_{k2}) \times \mathbf{r}_{g1})$;
- (3) $\mathbf{r}_{g2} = \mathbf{r}_{g0} \times \mathbf{r}_{g1}$,

where $\mathcal{N}(\mathbf{v}) = \mathbf{v}/\|\mathbf{v}\|$ normalizes a vector \mathbf{v} .

5.2 Parallax Removal

Once we have optimized the global orientations and focal lengths of our cameras, we may find that the images are still not perfectly aligned, i.e., the resulting stitched image looks blurry or ghosted in some places. This can be caused by a variety of factors, including unmodeled radial distortion, 3D parallax (failure to rotate the camera around its optical center), small scene motions such as waving tree branches, and large-scale scene motions such as people moving in and out of pictures.

Each of these problems can be treated with a different approach. Radial distortion can be estimated (potentially before the camera's first use) using one of the techniques discussed in Subsection 2.4. For example, the *plumb line method* [28, 57, 101] adjusts radial distortion parameters until slightly curved lines become straight, while mosaic-based approaches adjust them until mis-registration is reduced in image overlap areas [167, 187].

3D parallax can be attacked by doing a full 3D bundle adjustment, i.e., replacing the projection equation (5.2) used in (5.5) with (2.15), which models camera translations. The 3D positions of the matched features points and cameras can then be simultaneously recovered, although this can be significantly more expensive than parallax-free image registration. Once the 3D structure has been recovered, the scene could (in theory) be projected to a single (central) viewpoint that contains no parallax. However, in order to do this, dense *stereo* correspondence needs to be performed [111, 170, 198], which may not be possible if the images only contain partial overlap. In that case, it may be necessary to correct for parallax only in the overlap areas, which

can be accomplished using a *Multi-Perspective Plane Sweep* (MPPS) algorithm [102, 211].

When the motion in the scene is very large, i.e., when objects appear and disappear completely, a sensible solution is to simply *select* pixels from only one image at a time as the source for the final composite [1, 50, 136], as discussed in Subsection 6.2. However, when the motion is reasonably small (on the order of a few pixels), general 2D motion estimation (optic flow) can be used to perform an appropriate correction before blending using a process called *local alignment* [104, 178]. This same process can also be used to compensate for radial distortion and 3D parallax, although it uses a weaker motion model than explicitly modeling the source of error, and may therefore fail more often or introduce unwanted distortions.

The local alignment technique introduced by Shum and Szeliski [178] starts with the global bundle adjustment (5.7) used to optimize the camera poses. Once these have been estimated, the *desired* location of a 3D point \mathbf{x}_i can be estimated as the *average* of the back-projected 3D locations,

$$\bar{\mathbf{x}}_i \sim \sum_j c_{ij} \tilde{\mathbf{x}}_i(\hat{\mathbf{x}}_{ij}; \mathbf{R}_j, f_j), \quad (5.11)$$

which can be projected into each image j to obtain a *target location* $\bar{\mathbf{x}}_{ij}$. The difference between the target locations $\bar{\mathbf{x}}_{ij}$ and the original features \mathbf{x}_{ij} provide a set of local motion estimates

$$\mathbf{u}_{ij} = \bar{\mathbf{x}}_{ij} - \mathbf{x}_{ij}, \quad (5.12)$$


which can be interpolated to form a dense correction field $\mathbf{u}_j(\mathbf{x}_j)$. In their system, Shum and Szeliski [178] use an *inverse warping* algorithm where the sparse \mathbf{u}_{ij} values are placed at the new target locations $\bar{\mathbf{x}}_{ij}$, interpolated using bilinear kernel functions [142] and then added to the original pixel coordinates when computing the warped (corrected) image. To get a reasonably dense set of features to interpolate, Shum and Szeliski [178] place a feature point at the center of each patch (the patch size controls the smoothness in the local alignment stage), rather than relying on features extracted using an interest operator.

An alternative approach to motion-based de-ghosting was proposed by Kang *et al.* [104], who estimated dense optical flow between each input image and a central *reference* image. The accuracy of the flow vector is checked using a photo-consistency measure before a given warped pixel is considered valid and therefore used to compute a high dynamic range radiance estimate, which is the goal of their overall algorithm. The requirement for having a reference image makes their approach less applicable to general image mosaicing, although an extension to this case could certainly be envisaged.

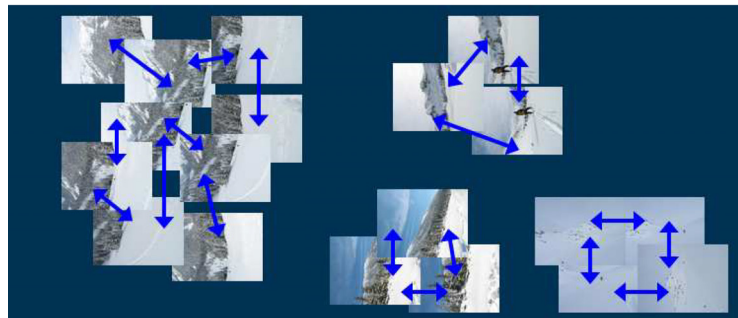
5.3 Recognizing Panoramas

The final piece needed to perform fully automated image stitching is a technique to recognize which images actually go together, that Brown and Lowe [30] call *recognizing panoramas*. If the user takes images in sequence so that each image overlaps its predecessor and also specifies the first and last images to be stitched, bundle adjustment combined with the process of *topology inference* can be used to automatically assemble a panorama [167]. However, users often jump around when taking panoramas, e.g., they may start a new row on top of a previous one, or jump back to take a repeated shot, or create 360° panoramas where end-to-end overlaps need to be discovered. Furthermore, the ability to discover multiple panoramas taken by a user over an extended period of time can be a big convenience.


To recognize panoramas, Brown and Lowe [30] first find all pairwise image overlaps using a feature-based method and then find connected components in the overlap graph to “recognize” individual panoramas (Figure 5.1). The feature-based matching stage first extracts SIFT feature locations and feature descriptors [122] from all the input images and then places these in an indexing structure, as described in Subsection 4.2. For each image pair under consideration, the nearest matching neighbor is found for each feature in the first image, using the indexing structure to rapidly find candidates, and then comparing feature descriptors to find the best match. RANSAC is then used to find a set of *inlier* matches, using a pair of matches to hypothesize a similarity motion model that is then used to count the number of inliers.



(a)



(b)



(c)

Fig. 5.1 Recognizing panoramas using our new algorithm [31]: (a) input images with pairwise matches; (b) images grouped into connected components (panoramas); and (c) individual panoramas registered and blended into stitched composites.

In practice, the most difficult part of getting a fully automated stitching algorithm to work is deciding which pairs of images actually correspond to the same parts of the scene. Repeated structures such as windows (Figure 5.2) can lead to false matches when using a




Fig. 5.2 Matching errors [31]: accidental matching of several features can lead to matches between pairs of images that do not actually overlap.




Fig. 5.3 Validation of image matches by direct pixel error comparison can fail when the scene contains moving objects.

feature-based approach. One way to mitigate this problem is to perform a direct pixel-based comparison between the registered images to determine if they actually are different views of the same scene. Unfortunately, this heuristic may fail if there are moving objects in the scene

(Figure 5.3). While there is no magic bullet for this problem short of full scene understanding, further improvements can likely be made by applying domain-specific heuristics such as priors on typical camera motions as well as machine learning techniques applied to the problem of match validation.

6

Compositing

Once we have registered all of the input images with respect to each other, we need to decide how to produce the final stitched (mosaic) image. This involves selecting a final compositing surface (flat, cylindrical, spherical, etc.) and view (reference image). It also involves selecting which pixels contribute to the final composite and how to optimally blend these pixels to minimize visible seams, blur, and ghosting.

In this section, I review techniques that address these problems, namely compositing surface parameterization, pixel/seam selection, blending, and exposure compensation. My emphasis is on fully *automated* approaches to the problem. Since the creation of high-quality panoramas and composites is as much an *artistic* endeavor as a computational one, various interactive tools have been developed to assist this process, e.g., [1, 117, 161]. I will not cover these in this article, except where they provide automated solutions to our problems.

6.1 Choosing a Compositing Surface

The first choice to be made is how to represent the final image. If only a few images are stitched together, a natural approach is to select one of

the images as the *reference* and to then warp all of the other images into the reference coordinate system. The resulting composite is sometimes called a *flat* panorama, since the projection onto the final surface is still a perspective projection, and hence straight lines remain straight (which is often a desirable attribute).

For larger fields of view, however, we cannot maintain a flat representation without excessively stretching pixels near the border of the image. (In practice, flat panoramas start to look severely distorted once the field of view exceeds 90° or so.) The usual choice for compositing larger panoramas is to use a cylindrical [43, 193] or spherical [199] projection, as described in Subsection 2.3. In fact, any surface used for *environment mapping* in computer graphics can be used, including a *cube map* that represents the full viewing sphere with the six square faces of a cube [73, 199]. Cartographers have also developed a number of alternative methods for representing the globe [33].

The choice of parameterization is somewhat application dependent, and involves a tradeoff between keeping the local appearance undistorted (e.g., keeping straight lines straight) and providing a reasonably uniform sampling of the environment. Automatically making this selection and smoothly transitioning between representations based on the extent of the panorama is an interesting topic for future research.

View Selection Once we have chosen the output parameterization, we still need to determine which part of the scene will be *centered* in the final view. As mentioned above, for a flat composite, we can choose one of the images as a reference. Often, a reasonable choice is the one that is geometrically most central. For example, for rotational panoramas represented as a collection of 3D rotation matrices, we can choose the image whose z -axis is closest to the average z -axis (assuming a reasonable field of view). Alternatively, we can use the average z -axis (or quaternion, but this is trickier) to define the reference rotation matrix.

For larger (e.g., cylindrical or spherical) panoramas, we can still use the same heuristic if a subset of the viewing sphere has been imaged. In the case of full 360° panoramas, a better choice might be to choose the middle image from the sequence of inputs, or sometimes the first image,

assuming this contains the object of greatest interest. In all of these cases, having the user control the final view is often highly desirable. If the “up vector” computation described in Subsection 5.1 is working correctly, this can be as simple as panning over the image or setting a vertical “center line” for the final panorama.

Coordinate Transformations Once we have selected the parameterization and reference view, we still need to compute the mappings between the input and output pixels coordinates.

If the final compositing surface is flat (e.g., a single plane or the face of a cube map) and the input images have no radial distortion, the coordinate transformation is the simple homography described by (2.19). This kind of warping can be performed in graphics hardware by appropriately setting texture mapping coordinates and rendering a single quadrilateral.

If the final composite surface has some other analytic form (e.g., cylindrical or spherical), we need to convert every pixel in the final panorama into a viewing ray (3D point) and then map it back into each image according to the projection (and optionally radial distortion) equations. This process can be made more efficient by pre-computing some lookup tables, e.g., the partial trigonometric functions needed to map cylindrical or spherical coordinates to 3D coordinates and/or the radial distortion field at each pixel. It is also possible to accelerate this process by computing exact pixel mappings on a coarser grid and then interpolating these values.

When the final compositing surface is a texture-mapped polyhedron, a slightly more sophisticated algorithm must be used. Not only do the 3D and texture map coordinates have to be properly handled, but a small amount of *overdraw* outside of the triangle footprints in the texture map is necessary, to ensure that the texture pixels being interpolated during 3D rendering have valid values [199].

Sampling Issues While the above computations can yield the correct (fractional) pixel addresses in each input image, we still need to pay attention to sampling issues. For example, if the final panorama has a lower resolution than the input images, pre-filtering the input

images is necessary to avoid aliasing. These issues have been extensively studied in both the image processing and computer graphics communities. The basic problem is to compute the appropriate pre-filter, which depends on the distance (and arrangement) between neighboring samples in a source image. Various approximate solutions, such as MIP mapping [219] or elliptically weighted Gaussian averaging [74] have been developed in the graphics community. For highest visual quality, a higher order (e.g., cubic) interpolator combined with a spatially adaptive pre-filter may be necessary [215]. Under certain conditions, it may also be possible to produce images with a higher resolution than the input images using a process called *super-resolution* (Section 7).

6.2 Pixel Selection and Weighting

Once the source pixels have been mapped onto the final composite surface, we must still decide how to blend them in order to create an attractive looking panorama. If all of the images are in perfect registration and identically exposed, this is an easy problem (any pixel or combination will do). However, for real images, visible seams (due to exposure differences), blurring (due to mis-registration), or ghosting (due to moving objects) can occur.

Creating clean, pleasing looking panoramas involves both deciding which pixels to use and how to weight or blend them. The distinction between these two stages is a little fluid, since per-pixel weighting can be thought of as a combination of selection and blending. In this section, I discuss spatially varying weighting, pixel selection (seam placement), and then more sophisticated blending.

Feathering and Center-Weighting The simplest way to create a final composite is to simply take an *average* value at each pixel,

$$C(\mathbf{x}) = \sum_k w_k(\mathbf{x}) \tilde{I}_k(\mathbf{x}) / \sum_k w_k(\mathbf{x}), \quad (6.1)$$

where $\tilde{I}_k(\mathbf{x})$ are the *warped* (re-sampled) images and $w_k(\mathbf{x})$ is 1 at valid pixels and 0 elsewhere. On computer graphics hardware, this kind of



Fig. 6.1 Final composites computed by a variety of algorithms: (a) average, (b) median, (c) feathered average, (d) p -norm $p = 10$, (e) Vornoi, (f) weighted ROD vertex cover with feathering, (g) graph cut seams with Poisson blending, (h) and with pyramid blending.

summation can be performed in an *accumulation buffer* (using the A channel as the weight).

Simple averaging usually does not work very well, since exposure differences, mis-registrations, and scene movement are all very visible (Figure 6.1a). If rapidly moving objects are the only problem, taking a *median* filter (which is a kind of pixel selection operator) can often be used to remove them [93] (Figure 6.1b). Conversely, center-weighting

(discussed below) and *minimum likelihood* selection [1] can sometimes be used to retain multiple copies of a moving object (Figure 6.4).

A better approach to averaging is to weight pixels near the center of the image more heavily and to down-weight pixels near the edges. When an image has some cutout regions, down-weighting pixels near the edges of both cutouts and edges is preferable. This can be done by computing a *distance map* or *grassfire transform*,

$$w_k(\mathbf{x}) = \left\| \arg \min_{\mathbf{y}} \{ \|\mathbf{y}\| \mid \tilde{I}_k(\mathbf{x} + \mathbf{y}) \text{ is invalid} \} \right\|, \quad (6.2)$$

where each valid pixel is tagged with its Euclidean distance to the nearest invalid pixel. The Euclidean distance map can be efficiently computed using a two-pass raster algorithm [26,49]. Weighted averaging with a distance map is often called *feathering* [42,199,210] and does a reasonable job of blending over exposure differences. However, blurring and ghosting can still be problems (Figure 6.1c). Note that weighted averaging is *not* the same as compositing the individual images with the classic *over* operation [25,153], even when using the weight values (normalized to sum up to one) as *alpha* (translucency) channels. This is because the over operation attenuates the values from more distant surfaces, and hence is not equivalent to a direct sum.

One way to improve feathering is to raise the distance map values to some large power, i.e., to use $w_k^p(\mathbf{x})$ in (6.1). The weighted averages then become dominated by the larger values, i.e., they act somewhat like a *p-norm*. The resulting composite can often provide a reasonable tradeoff between visible exposure differences and blur (Figure 6.1d).

In the limit as $p \rightarrow \infty$, only the pixel with the maximum weight gets selected,

$$C(\mathbf{x}) = \tilde{I}_{l(\mathbf{x})}(\mathbf{x}), \quad (6.3)$$

where

$$l = \arg \max_k w_k(\mathbf{x}) \quad (6.4)$$

is the *label assignment* or *pixel selection* function that selects the image to be used at each pixel. This hard pixel selection process produces a visibility mask-sensitive variant of the familiar *Vornoi diagram*, which

assigns each pixel to the nearest image center in the set [150, 220]. The resulting composite, while useful for artistic guidance and in high-overlap panoramas (*manifold mosaics*) tends to have very hard edges with noticeable seams when the exposures vary (Figure 6.1e).

Xiong and Turkowski [223] use this Voronoi idea (local maximum of the grassfire transform) to select seams for Laplacian pyramid blending (which is discussed below). However, since the seam selection is performed sequentially as new images are added in, some artifacts can occur.

Optimal Seam Selection Computing the Voronoi diagram is one way to select the *seams* between regions where different images contribute to the final composite. However, Voronoi images totally ignore the local image structure underlying the seam.

A better approach is to place the seams in regions where the images agree, so that transitions from one source to another are not visible. In this way, the algorithm avoids “cutting through” moving objects where a seam would look unnatural [50]. For a pair of images, this process can be formulated as a simple dynamic program starting from one (short) edge of the overlap region and ending at the other [50, 56, 135, 136].

When multiple images are being composited, the dynamic program idea does not readily generalize. (For square texture tiles being composited sequentially, Efros and Freeman [56] run a dynamic program along each of the four tile sides.)

To overcome this problem, Uyttendaele *et al.* [210] observed that for well-registered images, moving objects produce the most visible artifacts, namely translucent looking *ghosts*. Their system therefore decides which objects to keep and which ones to erase. First, the algorithm compares all overlapping input image pairs to determine *regions of difference* (RODs) where the images disagree. Next, a graph is constructed with the RODs as vertices and edges representing ROD pairs that overlap in the final composite (Figure 6.2). Since the presence of an edge indicates an area of disagreement, vertices (regions) must be removed from the final composite until no edge spans a pair of remaining vertices. The smallest such set can be computed using a *vertex cover* algorithm. Since several such covers may exist, a *weighted vertex cover*




Fig. 6.2 Computation of regions of differences (RODs): (a) three overlapping images with a moving face; (b) corresponding RODs; and (c) graph of coincident RODs. (Taken from [210]).

is used instead, where the vertex weights are computed by summing the feather weights in the ROD [210]. The algorithm therefore prefers removing regions that are near the edge of the image, which reduces the likelihood that partially visible objects will appear in the final composite. (It is also possible to infer which object in a region of difference is the foreground object by the “edginess” (pixel differences) across the ROD boundary, which should be higher when an object is present [87].) Once the desired excess regions of difference have been removed, the final composite can be created using a feathered blend (Figure 6.1f).

A different approach to pixel selection and seam placement was recently proposed by Agarwala *et al.* [1]. Their system computes the label assignment that optimizes the sum of two objective functions. The first is a per-pixel *image objective* that determines which pixels are likely to produce good composites,

$$\mathcal{C}_D = \sum_{\mathbf{x}} D_{l(\mathbf{x})}(\mathbf{x}), \quad (6.5)$$

where $D_{l(\mathbf{x})}(\mathbf{x})$ is the *data penalty* associated with choosing image l at pixel \mathbf{x} . In their system, users can select which pixels to use by “painting” over an image with the desired object or appearance, which sets $D(\mathbf{x}, l)$ to a large value for all labels l other than the one selected by the user (Figure 6.3). Alternatively, automated selection criteria can be used, such as *maximum likelihood* that prefers pixels that occur repeatedly (for object removal), or *minimum likelihood* for objects that occur infrequently (for greatest object retention). Using a more traditional center-weighted data term tends to favor objects that are centered in the input images (Figure 6.4).



Fig. 6.3 From a set of five source images (of which four are shown on the left), Photomontage quickly creates a composite family portrait in which everyone is smiling and looking at the camera (right). Users simply flip through the stack and coarsely draw strokes using the designated source image objective over the people they wish to add to the composite. The user-applied strokes and computed regions are color-coded by the borders of the source images on the left (middle). (Copied, with permission, from [1]).



Fig. 6.4 Set of five photos tracking a snowboarder's jump stitched together into a seamless composite. Because the algorithm prefers pixels near the center of the image, multiple copies of the boarder are retained.

The second term is a *seam objective* that penalizes differences in labelings between adjacent images,

$$\mathcal{C}_S = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{N}} S_{l(\mathbf{x}), l(\mathbf{y})}(\mathbf{x}, \mathbf{y}), \quad (6.6)$$

where $S_{l(\mathbf{x}), l(\mathbf{y})}(\mathbf{x}, \mathbf{y})$ is the image-dependent *interaction penalty* or *seam cost* of placing a seam between pixels \mathbf{x} and \mathbf{y} , and \mathcal{N} is the set of N_4 neighboring pixels. For example, the simple color-based seam penalty used in [1, 112] can be written as

$$S_{l(\mathbf{x}), l(\mathbf{y})}(\mathbf{x}, \mathbf{y}) = \|\tilde{I}_{l(\mathbf{x})}(\mathbf{x}) - \tilde{I}_{l(\mathbf{y})}(\mathbf{x})\| + \|\tilde{I}_{l(\mathbf{x})}(\mathbf{y}) - \tilde{I}_{l(\mathbf{y})}(\mathbf{y})\|. \quad (6.7)$$

More sophisticated seam penalties can also look at image gradients or the presence of image edges [1]. Seam penalties are widely used in other computer vision applications such as stereo matching [27] to give the labeling function its *coherence* or *smoothness*. An alternative approach, which places seams along strong consistent edges in overlapping images using a watershed computation has recently been developed by Soille [182].

The sum of the two objective functions is often called the *Markov Random Field* (MRF) energy, since it arises as the negative log-likelihood of an MRF distribution [67]. For general energy functions, finding the minimum can be NP-hard [27]. However, a variety of approximate optimization techniques have been developed over the years, including *simulated annealing* [67], graph cuts [27], and loopy belief propagation [190, 200]. Both Kwatra *et al.* [112] and Agarwala *et al.* [1] use graph cuts, which involves cycling through a set of simpler α -*expansion* re-labelings, each of which can be solved with a graph cut (max-flow) polynomial-time algorithm [27].

For the result shown in Figure 6.1g, Agarwala *et al.* [1] use a large data penalty for invalid pixels and 0 for valid pixels. Notice how the seam placement algorithm avoids regions of differences, including those that border the image and which might result in cut off objects. Graph cuts [1] and vertex cover [210] often produce similar looking results, although the former is significantly slower since it optimizes over all pixels, while the latter is more sensitive to the thresholds used to determine regions of difference.

6.3 Blending

Once the seams have been placed and unwanted object removed, we still need to blend the images to compensate for exposure differences and other mis-alignments. The spatially varying weighting (feathering) previously discussed can often be used to accomplish this. However, it is difficult in practice to achieve a pleasing balance between smoothing out low-frequency exposure variations and retaining sharp enough transitions to prevent blurring (although using a high exponent does help).

Laplacian Pyramid Blending An attractive solution to this problem was developed by Burt and Adelson [34]. Instead of using a single transition width, a frequency-adaptive width is used by creating a band-pass (Laplacian) pyramid and making the transition widths a function of the pyramid level. The process operates as follows.

First, each warped image is converted into a band-pass (Laplacian) pyramid, which involves smoothing each level with a $1/16(1, 4, 6, 4, 1)$ binomial kernel, subsampling the smoothed image by a factor of 2, and subtracting the reconstructed (low-pass) image from the original. This creates a reversible, overcomplete representation of the image signal. Invalid and edge pixels are filled with neighboring values to make this process well defined.

Next, the *mask* (valid pixel) image associated with each source image is converted into a low-pass (Gaussian) pyramid. These blurred and subsampled masks become the weights used to perform a per-level feathered blend of the band-pass source images.

Finally, the composite image is reconstructed by interpolating and summing all of the pyramid levels (band-pass images). The result of applying this pyramid blending is shown in Figure 6.1i.

Gradient Domain Blending An alternative approach to multi-band image blending is to perform the operations in the *gradient domain*. Reconstructing images from their gradient fields has a long history in computer vision [89], starting originally with work in brightness constancy [88], shape from shading [90], and photometric stereo [221]. More recently, related ideas have been used for reconstructing images from their edges [58], removing shadows from images [218], separating reflections from a single image [115], and *tone mapping* high dynamic range images by reducing the magnitude of image edges (gradients) [59].

Pérez *et al.* [151] showed how gradient domain reconstruction can be used to do seamless object insertion in image editing applications. Rather than copying pixels, the *gradients* of the new image fragment are copied instead. The actual pixel values for the copied area are then computed by solving a *Poisson equation* that locally matches the gradients while obeying the fixed *Dirichlet* (exact matching) conditions at the seam boundary. Pérez *et al.* [151] show that this is

equivalent to computing an additive *membrane* interpolant of the mismatch between the source and destination images along the boundary. (The membrane interpolant is known to have nicer interpolation properties for arbitrary-shaped constraints than frequency-domain interpolants [142].) In earlier work, Peleg [148] also proposed adding a smooth function to force a consistency along the seam curve.

Agarwala *et al.* [1] extended this idea to a multi-source formulation, where it no longer makes sense to talk of a destination image whose exact pixel values must be matched at the seam. Instead, *each* source image contributes its own gradient field, and the Poisson equation is solved using *Neumann* boundary conditions, i.e., dropping any equations that involve pixels outside the boundary of the image.

Rather than solving the Poisson partial differential equations, Agarwala *et al.* [1] directly minimize *variational problem*,

$$\min_{C(\mathbf{x})} \|\nabla C(\mathbf{x}) - \nabla \tilde{I}_{l(\mathbf{x})}(\mathbf{x})\|^2. \quad (6.8)$$

The discretized form of this equation is a set of gradient constraint equations

$$C(\mathbf{x} + \hat{\mathbf{i}}) - C(\mathbf{x}) = \tilde{I}_{l(\mathbf{x})}(\mathbf{x} + \hat{\mathbf{i}}) - \tilde{I}_{l(\mathbf{x})}(\mathbf{x}) \quad \text{and} \quad (6.9)$$

$$C(\mathbf{x} + \hat{\mathbf{j}}) - C(\mathbf{x}) = \tilde{I}_{l(\mathbf{x})}(\mathbf{x} + \hat{\mathbf{j}}) - \tilde{I}_{l(\mathbf{x})}(\mathbf{x}), \quad (6.10)$$

where $\hat{\mathbf{i}} = (1, 0)$ and $\hat{\mathbf{j}} = (0, 1)$ are unit vectors in the x and y directions.¹ They then solve the associated sparse least squares problem. Since this system of equations is only defined up to an additive constraint, Agarwala *et al.* [1] ask the user to select the value of one pixel. In practice, a better choice might be to weakly bias the solution toward reproducing the original color values.

To accelerate the solution of this sparse linear system, [59] use multi-grid, whereas [1] use hierarchical basis preconditioned conjugate gradient descent [192, 195]. The resulting seam blending work very well in practice (Figure 6.1h), although care must be taken when copying large gradient values near seams so that a “double edge” is not introduced.

¹At seam locations, the right-hand side is replaced by the average of the gradients in the two source images.

Copying gradients directly from the source images after seam placement is just one approach to gradient domain blending. The paper by Levin *et al.* [116] examines several different variants on this approach, which they call *Gradient-domain Image STitching* (GIST). The techniques they examine include feathering (blending) the gradients from the source images, as well as using an L1 norm in performing the reconstruction of the image from the gradient field, rather than using an L2 norm as in (6.8). Their preferred technique is the L1 optimization of a feathered (blended) cost function on the original image gradients (which they call GIST1- l_1). Since L1 optimization using linear programming can be slow, they develop a faster iterative median-based algorithm in a multigrid framework. Visual comparisons between their preferred approach and what they call *optimal seam on the gradients* (which is equivalent to Agarwala *et al.* [1]’s approach) show similar results, while significantly improving on pyramid blending and feathering algorithms.

Exposure Compensation Pyramid and gradient domain blending can do a good job of compensating for moderate amounts of exposure differences between images. However, when the exposure differences become large, alternative approaches may be necessary.

Uyttendaele *et al.* [210] iteratively estimate a local correction between each source image and a blended composite. First, a block-based quadratic transfer function is fit between each source image and an initial feathered composite. Next, transfer functions are averaged with their neighbors to get a smoother mapping, and per-pixel transfer functions are computed by *splining* (interpolating) between neighboring block values. Once each source image has been smoothly adjusted, a new feathered composite is computed, and the process is repeated (typically three times). The results in [210] demonstrate that this does a better job of exposure compensation than simple feathering, and can handle local variations in exposure due to effects like lens vignetting.

High Dynamic Range Imaging A more principled approach to exposure compensation is to estimate a single *high dynamic range* (HDR) radiance map from the differently exposed images [52, 125,

137, 158]. Most techniques assume that the input images were taken with a fixed camera whose pixel values

$$I_k(\mathbf{x}) = f(c_k R(\mathbf{x}); \mathbf{p}) \quad (6.11)$$

are the result of applying a parameterized *radiometric transfer function* $f(R, \mathbf{p})$ to scaled radiance values $c_k R(\mathbf{x})$. The exposure values c_k are either known (by experimental setup, or from a camera's EXIF tags), or are computed as part of the fitting process.

The form of the parametric function differs from paper to paper. Mann and Picard [125] use a three-parameter $f(R) = \alpha + \beta R^\gamma$ function, Debevec and Malik [52] use a thin-plate cubic spline, while Mitsunaga and Nayar [137] use a low-order ($N \leq 10$) polynomial for the *inverse* of the transfer function.

To blend the estimated (noisy) radiance values into a final composite, Mann and Picard [125] use a hat function (accentuating mid-tone pixels), Debevec and Malik [52] use the derivative of the response function, while Mitsunaga and Nayar [137] optimize the signal-to-noise ratio (SNR), which emphasizes both higher pixel values and larger gradients in the transfer function.

Once a radiance map has been computed, it is usually necessary to display it on a lower gamut (i.e., 8-bit) screen or printer. A variety of *tone mapping* techniques have been developed for this purpose, which involve either computing spatially varying transfer functions or reducing image gradients to fit the available dynamic range [54, 59, 120, 159].

Unfortunately, casually acquired images may not be perfectly registered and may contain moving objects. Kang *et al.* [104] present an algorithm that combines global registration with local motion estimation (optic flow) to accurately align the images before blending their radiance estimates (Figure 6.5). Since the images may have widely different exposures, care must be taken when producing the motion estimates, which must themselves be checked for consistency to avoid the creation of ghosts and object fragments.

Even this approach, however, may not work when the camera is simultaneously undergoing large panning motions and exposure changes, which is a common occurrence in casually acquired panoramas. Under such conditions, different parts of the image may be seen at




Fig. 6.5 Merging multiple exposures to create a high dynamic range composite: (a–c) three different exposures; (d) merging the exposures using classic algorithms (note the ghosting due to the horse’s head movement); and (e) merging the exposures with motion compensation [104].

one or more exposures. Devising a method to blend all of these different sources while avoiding sharp transitions and dealing with scene motion is a challenging task that has recently been tackled by first finding a consensus mosaic and then selectively computing radiances in under- and over-exposed regions [55].

In the long term, the need to compute high dynamic range images from multiple exposures may be eliminated by advances in camera sensor technology [104, 140, 208, 224]. However, the need to blend such images and to tone map them to a pleasing final result will likely remain.

7

Extensions and Open Issues

In this paper, the basics of image alignment and stitching, concentrating on techniques for registering partially overlapping images and blending them to create seamless panoramas have been surveyed. A large number of additional techniques have been developed for solving related problems such as increasing the resolution of images by taking multiple displaced pictures (*super-resolution*), stitching videos together to create dynamic panoramas, and stitching videos and images in the presence of large amounts of parallax.

Perhaps, the most common question that comes up in relation to image stitching is the following. “Why can’t you just take multiple images of the same scene with sub-pixel displacements and produce an image with a higher effective resolution?” Indeed, this problem has been studied for a long time and is generally known as *multiple image super-resolution*.¹ Examples of papers that have addressed this issue include [35, 36, 40, 41, 96, 107]. (See [8] for a recent paper with lots of additional references and experimental comparisons.) The general idea is that different images of the same scene taken from slightly different

¹One can also increase the resolution of a single image using various kinds of non-linear or example-based interpolation techniques [8, 65].

positions (i.e., where the pixels do not sample exactly the same rays in space) contain more information than a single image. However, this is only true if the imager actually *aliases* the original signal, e.g., if the silicon sensor integrates over a finite area and the optics do not cut off all the frequencies above the Nyquist frequency. Motion estimation also needs to be very accurate for this to work, so that in practice, an increase in resolution greater than $2\times$ is difficult to achieve [8].

Another popular topic is video stitching [15, 93, 126, 166, 183, 203]. While this problem is in many ways a straightforward generalization of multiple-image stitching, the potential presence of large amounts of independent motion, camera zoom, and the desire to visualize dynamic events impose additional challenges. For example, moving foreground objects can often be removed using *median filtering*. Alternatively, foreground objects can be extracted into a separate layer [166] and later composited back into the stitched panoramas, sometimes as multiple instances to give the impressions of a “Chronophotograph” [126] and sometimes as video overlays [93]. Videos can also be used to create animated *panoramic video textures* in which different portions of a panoramic scene are animated with independently moving video loops [2, 156].

Video can also provide an interesting source of content for creating panoramas taken from moving cameras. While this invalidates the usual assumption of a single point of view (optical center), interesting results can still be obtained. For example, the VideoBrush system [168] uses thin strips taken from the center of the image to create a panorama taken from a horizontally moving camera. This idea can be generalized to other camera motions and compositing surfaces using the concept of mosaics on adaptive manifold [150]. Related ideas have been used to create panoramic matte paintings for multi-plane cell animation [220], for creating stitched images of scenes with parallax [111], and as 3D representations of more complex scenes using *multiple-center-of-projection images* [155].

Another interesting variant on video-based panoramas are *concentric mosaics* [175]. Here, rather than trying to produce a single panoramic image, the complete original video is kept and used to re-synthesize novel views (from different camera origins) using ray

remapping (light field rendering), thus endowing the panorama with a sense of 3D depth. The same data set can also be used to explicitly reconstruct the depth using multi-baseline stereo [118, 147, 177, 179].

Open Issues While image stitching is by now a fairly mature field with a variety of commercial products, there remain a large number of challenges and open extensions. One of these is to increase the reliability of fully automated stitching algorithms. As discussed in Subsection 5.3 and illustrated in Figures 5.2 and 5.3, it is difficult to simultaneously avoid matching spurious features or repeated patterns while also being tolerant to large outliers such as moving people. Advances in semantic scene understanding could help resolve some of these problems, as well as better machine learning techniques for feature matching and validation.

The problem of parallax has also not been adequately solved. For small amounts of parallax, the deghosting techniques described in Subsections 5.2 and 6.2 can often adequately disguise these effects through local warping and careful seam selection. For high-overlap panoramas, *concentric mosaics* [175], panoramas with parallax [118] and careful seam selection (with potential user guidance) [1] can be used. The most challenging case is limited overlap panoramas with large parallax, since the depth estimates needed to compensate for the parallax are only available in the overlap regions [102, 211].

References

- [1] A. Agarwala *et al.*, “Interactive digital photomontage,” *ACM Transactions on Graphics*, vol. 23, no. 3, pp. 292–300, August 2004.
- [2] A. Agarwala *et al.*, “Panoramic video textures,” *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 821–827, August 2005.
- [3] P. Anandan, “A computational framework and an algorithm for the measurement of visual motion,” *International Journal of Computer Vision*, vol. 2, no. 3, pp. 283–310, January 1989.
- [4] V. Argyriou and T. Vlachos, “Estimation of sub-pixel motion using gradient cross-correlation,” *Electronic Letters*, vol. 39, no. 13, pp. 980–982, June 2003.
- [5] N. Ayache, *Vision Stéréoscopique et Perception Multisensorielle*. Paris: InterEditions, 1989.
- [6] A. Bab Hadiashar and D. Suter, “Robust total least squares based optic flow computation,” in *Asian Conference on Computer Vision (ACCV’98)*, (Hong Kong), pp. 566–573, ACM, January 1998.
- [7] F. Badra, A. Qumsieh, and G. Dudek, “Rotation and zooming in image mosaicing,” in *IEEE Workshop on Applications of Computer Vision (WACV’98)*, (Princeton), pp. 50–55, October 1998. IEEE Computer Society.
- [8] S. Baker and T. Kanade, “Limits on super-resolution and how to break them,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 9, pp. 1167–1183, September 2002.
- [9] S. Baker and I. Matthews, “Lucas-Kanade 20 years on: A unifying framework: Part 1: The quantity approximated, the warp update rule, and the gradient descent approximation,” *International Journal of Computer Vision*, vol. 56, no. 3, pp. 221–255, March 2004.

- [10] S. Baker *et al.*, “Lucas-kanade 20 years on: A unifying framework: Part 4,” The Robotics Institute, Carnegie Mellon University, February 2004. Technical Report CMU-RI-TR-04-14.
- [11] S. Baker *et al.*, “Lucas-kanade 20 years on: A unifying framework: Part 2,” The Robotics Institute, Carnegie Mellon University, February 2003. Technical Report CMU-RI-TR-03-01.
- [12] S. Baker *et al.*, “Lucas-kanade 20 years on: A unifying framework: Part 3,” The Robotics Institute, Carnegie Mellon University, November 2003. Technical Report CMU-RI-TR-03-35.
- [13] J. P. Barreto and K. Daniilidis, “Fundamental matrix for cameras with radial distortion,” in *Tenth International Conference on Computer Vision (ICCV 2005)*, (Beijing, China), vol. 1, pp. 625–632, October 2005.
- [14] A. Bartoli, M. Coquerelle, and P. Sturm, “A framework for pencil-of-points structure-from-motion,” in *Eighth European Conference on Computer Vision (ECCV 2004)*, (Prague), vol. II, pp. 28–40, Springer-Verlag, May 2004.
- [15] P. Baudisch *et al.*, “Panoramic viewfinder: Providing a real-time preview to help users avoid flaws in panoramic pictures,” in *OZCHI 2005*, (Canberra, Australia), November 2005.
- [16] A. Baumberg, “Reliable feature matching across widely separated views,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2000)*, (Hilton Head Island), vol. 1, pp. 774–781, June 2000.
- [17] H. Bay, T. Tuytelaars, and L. V. Gool, “Surf: Speeded up robust features,” in *Computer Vision – ECCV 2006*, (A. Leonardis, H. Bischof, and A. Pinz, eds.), vol. 3951, pp. 404–417, of *Lecture Notes in Computer Science*, Springer, 2006.
- [18] J. S. Beis and D. G. Lowe, “Shape indexing using approximate nearest-neighbour search in high-dimensional spaces,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’97)*, (San Juan, Puerto Rico), pp. 1000–1006, June 1997.
- [19] R. Benosman and S. B. Kang, eds., *Panoramic Vision: Sensors, Theory, and Applications*. New York: Springer, 2001.
- [20] J. R. Bergen, P. Anandan, K. J. Hanna, and R. Hingorani, “Hierarchical model-based motion estimation,” in *Second European Conference on Computer Vision (ECCV’92)*, (Santa Margherita Ligure, Italy), pp. 237–252, Springer-Verlag, May 1992.
- [21] J. R. Bergen, P. J. Burt, R. Hingorani, and S. Peleg, “A three-frame algorithm for estimating two-component image motion,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 9, pp. 886–896, September 1992.
- [22] M. J. Black and P. Anandan, “The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields,” *Computer Vision and Image Understanding*, vol. 63, no. 1, pp. 75–104, 1996.
- [23] M. J. Black and A. D. Jepson, “EigenTracking: Robust matching and tracking of articulated objects using a view-based representation,” *International Journal of Computer Vision*, vol. 26, no. 1, pp. 63–84, January 1998.

- [24] M. J. Black and A. Rangarajan, "On the unification of line processes, outlier rejection, and robust statistics with applications in early vision," *International Journal of Computer Vision*, vol. 19, no. 1, pp. 57–91, 1996.
- [25] J. F. Blinn, "Jim Blinn's corner: Compositing, part 1: Theory," *IEEE Computer Graphics and Applications*, vol. 14, no. 5, pp. 83–87, September 1994.
- [26] G. Borgefors, "Distance transformations in digital images," *Computer Vision, Graphics and Image Processing*, vol. 34, no. 3, pp. 227–248, June 1986.
- [27] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1222–1239, November 2001.
- [28] D. C. Brown, "Close-range camera calibration," *Photogrammetric Engineering*, vol. 37, no. 8, pp. 855–866, 1971.
- [29] L. G. Brown, "A survey of image registration techniques," *Computing Surveys*, vol. 24, no. 4, pp. 325–376, December 1992.
- [30] M. Brown and D. Lowe, "Recognizing panoramas," in *Ninth International Conference on Computer Vision (ICCV'03)*, (Nice, France), pp. 1218–1225, October 2003.
- [31] M. Brown, R. Szeliski, and S. Winder, "Multi-image matching using multi-scale oriented patches," Technical Report MSR-TR-2004-133, Microsoft Research, December 2004.
- [32] M. Brown, R. Szeliski, and S. Winder, "Multi-image matching using multi-scale oriented patches," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2005)*, (San Diego, CA), vol. I, pp. 510–517, June 2005.
- [33] L. M. Bugayevskiy and J. P. Snyder, *Map Projections: A Reference Manual*. CRC Press, 1995.
- [34] P. J. Burt and E. H. Adelson, "A multiresolution spline with applications to image mosaics," *ACM Transactions on Graphics*, vol. 2, no. 4, pp. 217–236, October 1983.
- [35] D. Capel and A. Zisserman, "Automated mosaicing with super-resolution zoom," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'98)*, (Santa Barbara), pp. 885–891, June 1998.
- [36] D. Capel and A. Zisserman, "Super-resolution enhancement of text image sequences," in *Fifteenth International Conference on Pattern Recognition (ICPR'2000)*, (Barcelona, Spain), vol. I, pp. 600–605, IEEE Computer Society Press, September 2000.
- [37] G. Carneiro and A. Jepson, "The distinctiveness, detectability, and robustness of local image features," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2005)*, (San Diego, CA), vol. II, pp. 296–301, June 2005.
- [38] T. J. Cham and R. Cipolla, "A statistical framework for long-range feature matching in uncalibrated image mosaicing," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'98)*, (Santa Barbara), pp. 442–447, June 1998.
- [39] G. Champleboux *et al.*, "Accurate calibration of cameras and range imaging sensors, the NPBS method," in *IEEE International Conference on Robotics*

- and Automation*, (Nice, France), pp. 1552–1558, IEEE Computer Society Press, May 1992.
- [40] S. Chaudhuri, *Super-Resolution Imaging*. Springer, 2001.
 - [41] P. Cheeseman, B. Kanefsky, R. Hanson, and J. Stutz, “Super-resolved surface reconstruction from multiple images,” Technical Report FIA-93-02, NASA Ames Research Center, Artificial Intelligence Branch, January 1993.
 - [42] C.-Y. Chen and R. Klette, “Image stitching – comparisons and new techniques,” in *Computer Analysis of Images and Patterns (CAIP’99)*, (Ljubljana), pp. 615–622, Springer-Verlag, September 1999.
 - [43] S. E. Chen, “QuickTime VR – an image-based approach to virtual environment navigation,” *Computer Graphics (SIGGRAPH’95)*, pp. 29–38, August 1995.
 - [44] O. Chum and J. Matas, “Matching with prosac – progressive sample consensus,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2005)*, (San Diego, CA), vol. I, pp. 220–226, June 2005.
 - [45] D. Claus and A. Fitzgibbon, “A rational function lens distortion model for general cameras,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2005)*, (San Diego, CA), vol. I, pp. 213–219, June 2005.
 - [46] S. Coorg and S. Teller, “Spherical mosaics with quaternions and dense correlation,” *International Journal of Computer Vision*, vol. 37, no. 3, pp. 259–273, June 2000.
 - [47] J. Corso and G. Hager, “Coherent regions for concise and stable image description,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2005)*, (San Diego, CA), vol. II, pp. 184–190, June 2005.
 - [48] I. J. Cox, S. Roy, and S. L. Hingorani, “Dynamic histogram warping of image pairs for constant image brightness,” in *IEEE International Conference on Image Processing (ICIP’95)*, vol. 2, pp. 366–369, IEEE Computer Society, 1995.
 - [49] P. E. Danielsson, “Euclidean distance mapping,” *Computer Graphics and Image Processing*, vol. 14, no. 3, pp. 227–248, November 1980.
 - [50] J. Davis, “Mosaics of scenes with moving objects,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’98)*, (Santa Barbara), pp. 354–360, June 1998.
 - [51] E. De Castro and C. Morandi, “Registration of translated and rotated images using finite fourier transforms,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 5, pp. 700–703, September 1987.
 - [52] P. E. Debevec and J. Malik, “Recovering high dynamic range radiance maps from photographs,” in *Proceedings of SIGGRAPH 97*, Los Angeles, California, pp. 369–378, August 1997. ISBN 0-89791-896-7.
 - [53] F. Dellaert and R. Collins, “Fast image-based tracking by selective pixel integration,” in *ICCV Workshop on Frame-Rate Vision*, pp. 1–22, 1999.
 - [54] F. Durand and J. Dorsey, “Fast bilateral filtering for the display of high-dynamic-range images,” *ACM Transactions on Graphics (TOG)*, vol. 21, no. 3, pp. 257–266, 2002.

- [55] A. Eden, M. Uyttendaele, and R. Szeliski, "Seamless image stitching of scenes with large motions and exposure differences," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2006)*, (New York, NY), vol. 3, pp. 2498–2505, June 2006.
- [56] A. A. Efros and W. T. Freeman, "Image quilting for texture synthesis and transfer," in *SIGGRAPH 2001, Computer Graphics Proceedings*, (E. Fiume, ed.), pp. 341–346, ACM Press/ACM SIGGRAPH, 2001.
- [57] M. El Melegy and A. Farag, "Nonmetric lens distortion calibration: Closed-form solutions, robust estimation and model selection," in *Ninth International Conference on Computer Vision (ICCV 2003)*, (Nice, France), pp. 554–559, October 2003.
- [58] J. H. Elder and R. M. Goldgerg, "Image editing in the contour domain," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 3, pp. 291–296, March 2001.
- [59] R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," *ACM Transactions on Graphics (TOG)*, vol. 21, no. 3, pp. 249–256, 2002.
- [60] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, June 1981.
- [61] D. Fleet and A. Jepson, "Computation of component image velocity from local phase information," *International Journal of Computer Vision*, vol. 5, no. 1, pp. 77–104, January 1990.
- [62] W. Förstner, "A feature-based correspondence algorithm for image matching," *International Archives on Photogrammetry and Remote Sensing*, vol. 26, no. 3, pp. 150–166, 1986.
- [63] W. Förstner, "A framework for low level feature extraction," in *Third European Conference on Computer Vision (ECCV'94)*, (Stockholm, Sweden), pp. 383–394, Springer-Verlag, May 1994.
- [64] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, September 1991.
- [65] W. T. Freeman, T. R. Jones, and E. C. Pasztor, "Example-based super-resolution," *IEEE Computer Graphics and Applications*, vol. 22, no. 2, pp. 56–65, March–April 2002.
- [66] C.-S. Fuh and P. Maragos, "Motion displacement estimation using an affine model for image matching," *Optical Engineering*, vol. 30, no. 7, pp. 881–887, July 1991.
- [67] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, no. 6, pp. 721–741, November 1984.
- [68] M. A. Gennert, "Brightness-based stereo matching," in *Second International Conference on Computer Vision (ICCV'88)*, (Tampa), pp. 139–143, IEEE Computer Society Press, December 1988.
- [69] G. Golub and C. F. Van Loan, *Matrix Computation*. Baltimore and London: The John Hopkins University Press, Third Edition, 1996.

- [70] A. Goshtasby, "Correction of image deformation from lens distortion using bezier patches," *Computer Vision, Graphics, and Image Processing*, vol. 47, no. 4, pp. 385–394, November 1989.
- [71] A. Goshtasby, *2-D and 3-D Image Registration*. New York: Wiley, 2005.
- [72] V. M. Govindu, "Revisiting the brightness constraint: Probabilistic formulation and algorithms," in *Computer Vision – ECCV 2006*, (A. Leonardis, H. Bischof, and A. Pinz, eds.), vol. 3953, pp. 177–188, of *Lecture Notes in Computer Science*, Springer, 2006.
- [73] N. Greene, "Environment mapping and other applications of world projections," *IEEE Computer Graphics and Applications*, vol. 6, no. 11, pp. 21–29, November 1986.
- [74] N. Greene and P. Heckbert, "Creating raster Omnimax images from multiple perspective views using the elliptical weighted average filter," *IEEE Computer Graphics and Applications*, vol. 6, no. 6, pp. 21–27, June 1986.
- [75] K. D. Gremban, C. E. Thorpe, and T. Kanade, "Geometric camera calibration using systems of linear equations," in *IEEE International Conference on Robotics and Automation*, (Philadelphia), pp. 562–567, IEEE Computer Society Press, April 1988.
- [76] M. D. Grossberg and S. K. Nayar, "A general imaging model and a method for finding its parameters," in *Eighth International Conference on Computer Vision (ICCV 2001)*, (Vancouver, Canada), vol. II, pp. 108–115, July 2001.
- [77] G. D. Hager and P. N. Belhumeur, "Efficient region tracking with parametric models of geometry and illumination," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 10, pp. 1025–1039, October 1998.
- [78] F. R. Hampel *et al.*, *Robust Statistics: The Approach Based on Influence Functions*. New York: Wiley, 1986.
- [79] M. J. Hannah, *Computer Matching of Areas in Stereo Images*. PhD thesis, Stanford University, 1974.
- [80] M. J. Hannah, "Test results from SRI's stereo system," in *Image Understanding Workshop*, (Cambridge, Massachusetts), pp. 740–744, Morgan Kaufmann Publishers, April 1988.
- [81] M. Hansen, P. Anandan, K. Dana, G. van der Wal, and P. Burt, "Real-time scene stabilization and mosaic construction," in *IEEE Workshop on Applications of Computer Vision (WACV'94)*, (Sarasota), pp. 54–62, December 1994. IEEE Computer Society.
- [82] C. Harris and M. J. Stephens, "A combined corner and edge detector," in *Alvey Vision Conference*, pp. 147–152, 1988.
- [83] R. I. Hartley, "Self-calibration from multiple views of a rotating camera," in *Third European Conference on Computer Vision (ECCV'94)*, (Stockholm, Sweden), vol. 1, pp. 471–478, Springer-Verlag, May 1994.
- [84] R. I. Hartley and S. B. Kang, "Parameter-free radial distortion correction with centre of distortion estimation," in *Tenth International Conference on Computer Vision (ICCV 2005)*, (Beijing, China), vol. 2, pp. 1834–1841, October 2005.
- [85] R. I. Hartley and A. Zisserman, *Multiple View Geometry*. Cambridge, UK: Cambridge University Press, September 2000.

- [86] R. I. Hartley and A. Zisserman, *Multiple View Geometry*. Cambridge, UK: Cambridge University Press, March 2004.
- [87] C. Herley, "Automatic occlusion removal from minimum number of images," in *International Conference on Image Processing (ICIP 2005)*, (Genova), vol. II, pp. 1046–1049–16, September 2005.
- [88] B. K. P. Horn, "Determining lightness from an image," *Computer Graphics and Image Processing*, vol. 3, no. 1, pp. 277–299, December 1974.
- [89] B. K. P. Horn, *Robot Vision*. Cambridge, Massachusetts: MIT Press, 1986.
- [90] B. K. P. Horn and M. J. Brooks, *Shape from Shading*. Cambridge, Massachusetts: MIT Press, 1989.
- [91] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185–203, 1981.
- [92] P. J. Huber, *Robust Statistics*. New York: John Wiley & Sons, 1981.
- [93] M. Irani and P. Anandan, "Video indexing based on mosaic representations," *Proceedings of the IEEE*, vol. 86, no. 5, pp. 905–921, May 1998.
- [94] M. Irani and P. Anandan, "About direct methods," in *International Workshop on Vision Algorithms*, (Kerkyra, Greece), pp. 267–277, Springer, September 1999.
- [95] M. Irani, S. Hsu, and P. Anandan, "Video compression using mosaic representations," *Signal Processing: Image Communication*, vol. 7, pp. 529–552, 1995.
- [96] M. Irani and S. Peleg, "Improving resolution by image registration," *Graphical Models and Image Processing*, vol. 53, no. 3, pp. 231–239, May 1991.
- [97] J. Jia and C.-K. Tang, "Image registration with global and local luminance alignment," in *Ninth International Conference on Computer Vision (ICCV 2003)*, (Nice, France), pp. 156–163, October 2003.
- [98] F. Jurie and M. Dhome, "Hyperplane approximation for template matching," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 7, pp. 996–1000, July 2002.
- [99] T. Kadir and M. Brady, "Saliency, scale and image description," *International Journal of Computer Vision*, vol. 45, no. 2, pp. 83–105, November 2001.
- [100] T. Kadir, A. Zisserman, and M. Brady, "An affine invariant salient region detector," in *Eighth European Conference on Computer Vision (ECCV 2004)*, (Prague), vol. I, pp. 228–241, Springer-Verlag, May 2004.
- [101] S. B. Kang, "Radial distortion snakes," *IEICE Transactions on Information and Systems*, vol. E84-D, no. 12, pp. 1603–1611, December 2001.
- [102] S. B. Kang, R. Szeliski, and M. Uyttendaele, "Seamless stitching using multi-perspective plane sweep," Technical Report MSR-TR-2004-48, Microsoft Research, June 2004.
- [103] S. B. Kang and R. Weiss, "Characterization of errors in compositing panoramic images," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'97)*, (San Juan, Puerto Rico), pp. 103–109, June 1997.
- [104] S. B. Kang *et al.*, "High dynamic range video," *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 319–325, July 2003.

- [105] Y. Ke and R. Sukthankar, "PCA-SIFT: A more distinctive representation for local image descriptors," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2004)*, (Washington, DC), vol. II, pp. 506–513, June 2004.
- [106] C. Kenney, M. Zuliani, and B. Manjunath, "An axiomatic approach to corner detection," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2005)*, (San Diego, CA), vol. I, pp. 191–197, June 2005.
- [107] D. Keren, S. Peleg, and R. Brada, "Image sequence enhancement using sub-pixel displacements," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'88)*, (Ann Arbor, Michigan), pp. 742–746, IEEE Computer Society Press, June 1988.
- [108] J. Kim, V. Kolmogorov, and R. Zabih, "Visual correspondence using energy minimization and mutual information," in *Ninth International Conference on Computer Vision (ICCV 2003)*, (Nice, France), pp. 1033–1040, October 2003.
- [109] C. D. Kuglin and D. C. Hines, "The phase correlation image alignment method," in *IEEE 1975 Conference on Cybernetics and Society*, (New York), pp. 163–165, September 1975.
- [110] R. Kumar, P. Anandan, and K. Hanna, "Direct recovery of shape from multiple views: A parallax based approach," in *Twelfth International Conference on Pattern Recognition (ICPR'94)*, (Jerusalem, Israel), vol. A, pp. 685–688, IEEE Computer Society Press, October 1994.
- [111] R. Kumar, P. Anandan, M. Irani, J. Bergen, and K. Hanna, "Representation of scenes from collections of images," in *IEEE Workshop on Representations of Visual Scenes*, (Cambridge, Massachusetts), pp. 10–17, June 1995.
- [112] V. Kwatra *et al.*, "Graphcut textures: Image and video synthesis using graph cuts," *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 277–286, July 2003.
- [113] D. Le Gall, "MPEG: A video compression standard for multimedia applications," *Communications of the ACM*, vol. 34, no. 4, pp. 44–58, April 1991.
- [114] M.-C. Lee *et al.*, "A layered video object coding system using sprite and affine motion model," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 7, no. 1, pp. 130–145, February 1997.
- [115] A. Levin, A. Zomet, and Y. Weiss, "Separating reflections from a single image using local features," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2004)*, (Washington, DC), vol. I, pp. 306–313, June 2004.
- [116] A. Levin, A. Zomet, S. Peleg, and Y. Weiss, "Seamless image stitching in the gradient domain," in *Eighth European Conference on Computer Vision (ECCV 2004)*, (Prague), vol. IV, pp. 377–389, Springer-Verlag, May 2004.
- [117] Y. Li *et al.*, "Lazy snapping," *ACM Transactions on Graphics*, vol. 23, no. 3, pp. 303–308, August 2004.
- [118] Y. Li *et al.*, "Stereo reconstruction from multiperspective panoramas," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 1, pp. 44–62, January 2004.

- [119] T. Lindeberg, "Scale-space for discrete signals," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 3, pp. 234–254, March 1990.
- [120] D. Lischinski, Z. Farbman, M. Uytendaele, and R. Szeliski, "Interactive local adjustment of tonal values," *ACM Transactions on Graphics*, vol. 25, no. 3, pp. 646–653, August 2006.
- [121] C. Loop and Z. Zhang, "Computing rectifying homographies for stereo vision," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'99)*, (Fort Collins), vol. I, pp. 125–131, June 1999.
- [122] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, November 2004.
- [123] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application in stereo vision," in *Seventh International Joint Conference on Artificial Intelligence (IJCAI-81)*, (Vancouver), pp. 674–679, 1981.
- [124] S. Mann and R. W. Picard, "Virtual bellows: Constructing high-quality images from video," in *First IEEE International Conference on Image Processing (ICIP-94)*, (Austin), vol. I, pp. 363–367, November 1994.
- [125] S. Mann and R. W. Picard, "On being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures," in *IS&T's 48th Annual Conference*, (Washington, D. C.), pp. 422–428, Society for Imaging Science and Technology, May 1995.
- [126] M. Massey and W. Bender, "Salient stills: Process and practice," *IBM Systems Journal*, vol. 35, nos. 3&4, pp. 557–573, 1996.
- [127] J. Matas *et al.*, "Robust wide baseline stereo from maximally stable extremal regions," *Image and Vision Computing*, vol. 22, no. 10, pp. 761–767, September 2004.
- [128] L. H. Matthies, R. Szeliski, and T. Kanade, "Kalman filter-based algorithms for estimating depth from image sequences," *International Journal of Computer Vision*, vol. 3, pp. 209–236, 1989.
- [129] P. F. McLauchlan and A. Jaenicke, "Image mosaicing using sequential bundle adjustment," *Image and Vision Computing*, vol. 20, nos. 9–10, pp. 751–759, August 2002.
- [130] L. McMillan and G. Bishop, "Plenoptic modeling: An image-based rendering system," *Computer Graphics (SIGGRAPH'95)*, pp. 39–46, August 1995.
- [131] J. Meehan, *Panoramic Photography*. Watson-Guptill, 1990.
- [132] K. Mikolajczyk and C. Schmid, "Scale & affine invariant interest point detectors," *International Journal of Computer Vision*, vol. 60, no. 1, pp. 63–86, October 2004.
- [133] K. Mikolajczyk and C. Schmid, "A performance evaluation of local descriptors," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 10, pp. 1615–1630, October 2005.
- [134] K. Mikolajczyk *et al.*, "A comparison of affine region detectors," *International Journal of Computer Vision*, vol. 65, nos. 1–2, pp. 43–72, November 2005.
- [135] D. L. Milgram, "Computer methods for creating photomosaics," *IEEE Transactions on Computers*, vol. C-24, no. 11, pp. 1113–1119, November 1975.

- [136] D. L. Milgram, “Adaptive techniques for photomosaicking,” *IEEE Transactions on Computers*, vol. C-26, no. 11, pp. 1175–1180, November 1977.
- [137] T. Mitsunaga and S. K. Nayar, “Radiometric self calibration,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’99)*, (Fort Collins), vol. 1, pp. 374–380, June 1999.
- [138] H. Moravec, “The stanford cart and the cmu rover,” in *Proceedings of the IEEE*, vol. 71, no. 7, pp. 872–884, July 1983.
- [139] M. Mühlich and R. Mester, “The role of total least squares in motion analysis,” in *Fifth European Conference on Computer Vision (ECCV’98)*, (Freiburg, Germany), vol. II, pp. 305–321, Springer-Verlag, June 1998.
- [140] S. K. Nayar and T. Mitsunaga, “High dynamic range imaging: Spatially varying pixel exposures,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2000)*, (Hilton Head, Island), vol. 1, pp. 472–479, June 2000.
- [141] S. Nene and S. K. Nayar, “A simple algorithm for nearest neighbor search in high dimensions,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 9, pp. 989–1003, September 1997.
- [142] G. M. Nielson, “Scattered data modeling,” *IEEE Computer Graphics and Applications*, vol. 13, no. 1, pp. 60–70, January 1993.
- [143] D. Nister and H. Stewenius, “Scalable recognition with a vocabulary tree,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2006)*, (New York City, NY), vol. 3, pp. 2161–2168, June 2006.
- [144] M. Okutomi and T. Kanade, “A multiple baseline stereo,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 4, pp. 353–363, April 1993.
- [145] OpenGL-ARB, *OpenGL Reference Manual: The Official Reference Document to OpenGL, Version 1.1.*, Second Edition, Addison-Wesley, Reading, MA, 1997.
- [146] A. V. Oppenheim, R. W. Schaffer, and J. R. Buck, *Discrete-Time Signal Processing*, Second Edition, Pearson Education, 1999.
- [147] R. Peleg, M. Ben-Ezra, and Y. Pritch, “Omnistereo: Panoramic stereo imaging,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 3, pp. 279–290, March 2001.
- [148] S. Peleg, “Elimination of seams from photomosaics,” *Computer Vision, Graphics, and Image Processing*, vol. 16, pp. 1206–1210, 1981.
- [149] S. Peleg and A. Rav Acha, “Lucas-Kanade without iterative warping,” in *International Conference on Image Processing (ICIP-2006)*, (Atlanta), pp. 1097–1100, October 2006.
- [150] S. Peleg *et al.*, “Mosaicing on adaptive manifolds,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1144–1154, October 2000.
- [151] P. Pérez, M. Gangnet, and A. Blake, “Poisson image editing,” *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 313–318, July 2003.
- [152] B. Platel, E. Balmachnova, L. M. J. Florack, and B. M. ter Haar Romeny, “Top-points as interest points for image matching,” in *Computer Vision –*

- ECCV 2006*, (A. Leonardis, H. Bischof, and A. Pinz, eds.), vol. 3951, pp. 418–429, of *Lecture Notes in Computer Science*, Springer, 2006.
- [153] T. Porter and T. Duff, “Compositing digital images,” *Computer Graphics (SIGGRAPH’84)*, vol. 18, no. 3, pp. 253–259, July 1984.
- [154] L. H. Quam, “Hierarchical warp stereo,” in *Image Understanding Workshop*, (New Orleans), pp. 149–155, Science Applications International Corporation, December 1984.
- [155] P. Rademacher and G. Bishop, “Multiple-center-of-projection images,” in *Computer Graphics Proceedings, Annual Conference Series*, pp. 199–206, Proc. SIGGRAPH’98 (Orlando), ACM SIGGRAPH, July 1998.
- [156] A. Rav Acha, Y. Pritch, D. Lischinski, and S. Peleg, “Dynamosaics: Video mosaics with non-chronological time,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2005)*, (San Diego, CA), vol. I, pp. 58–65, June 2005.
- [157] J. Rehg and A. Witkin, “Visual tracking with deformation models,” in *IEEE International Conference on Robotics and Automation*, (Sacramento), pp. 844–850, IEEE Computer Society Press, April 1991.
- [158] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann, 2005.
- [159] E. Reinhard *et al.*, “Photographic tone reproduction for digital images,” *ACM Transactions on Graphics (TOG)*, vol. 21, no. 3, pp. 267–276, 2002.
- [160] E. Rosten and T. Drummond, “Machine learning for high-speed corner detection,” in *Computer Vision – ECCV 2006*, (A. Leonardis, H. Bischof, and A. Pinz, eds.), vol. 3951, pp. 430–443, of *Lecture Notes in Computer Science*, Springer, 2006.
- [161] C. Rother, V. Kolmogorov, and A. Blake, ““GrabCut” – interactive foreground extraction using iterated graph cuts,” *ACM Transactions on Graphics*, vol. 23, no. 3, pp. 309–314, August 2004.
- [162] P. J. Rousseeuw, “Least median of squares regression,” *Journal of the American Statistical Association*, vol. 79, pp. 871–880, 1984.
- [163] P. J. Rousseeuw and A. M. Leroy, *Robust Regression and Outlier Detection*. New York: Wiley, 1987.
- [164] H. Samet, *The Design and Analysis of Spatial Data Structures*. Reading, Massachusetts: Addison-Wesley, 1989.
- [165] H. S. Sawhney, “3D geometry from planar parallax,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’94)*, (Seattle), pp. 929–934, IEEE Computer Society, June 1994.
- [166] H. S. Sawhney and S. Ayer, “Compact representation of videos through dominant multiple motion estimation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 8, pp. 814–830, August 1996.
- [167] H. S. Sawhney and R. Kumar, “True multi-image alignment and its application to mosaicing and lens distortion correction,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 3, pp. 235–243, March 1999.

- [168] H. S. Sawhney *et al.*, “Videobrush: Experiences with consumer video mosaicing,” in *IEEE Workshop on Applications of Computer Vision (WACV’98)*, (Princeton), pp. 56–62, IEEE Computer Society, October 1998.
- [169] F. Schaffalitzky and A. Zisserman, “Multi-view matching for unordered image sets, or “How do I organize my holiday snaps?”,” in *Seventh European Conference on Computer Vision (ECCV 2002)*, (Copenhagen), vol. I, pp. 414–431, Springer-Verlag, May 2002.
- [170] D. Scharstein and R. Szeliski, “A taxonomy and evaluation of dense two-frame stereo correspondence algorithms,” *International Journal of Computer Vision*, vol. 47, no. 1, pp. 7–42, May 2002.
- [171] C. Schmid, R. Mohr, and C. Bauckhage, “Evaluation of interest point detectors,” *International Journal of Computer Vision*, vol. 37, no. 2, pp. 151–172, June 2000.
- [172] G. Shakhnarovich, P. Viola, and T. Darrell, “Fast pose estimation with parameter sensitive hashing,” in *Ninth International Conference on Computer Vision (ICCV 2003)*, (Nice, France), vol. 2, pp. 750–757, October 2003.
- [173] J. Shi and C. Tomasi, “Good features to track,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’94)*, (Seattle), pp. 593–600, IEEE Computer Society, June 1994.
- [174] K. Shoemake, “Animating rotation with quaternion curves,” *Computer Graphics (SIGGRAPH’85)*, vol. 19, no. 3, pp. 245–254, July 1985.
- [175] H.-Y. Shum and L.-W. He, “Rendering with concentric mosaics,” in *SIGGRAPH’99*, (Los Angeles), pp. 299–306, ACM SIGGRAPH, August 1999.
- [176] H.-Y. Shum and R. Szeliski, “Panoramic image mosaicing,” Technical Report MSR-TR-97-23, Microsoft Research, September 1997.
- [177] H.-Y. Shum and R. Szeliski, “Stereo reconstruction from multiperspective panoramas,” in *Seventh International Conference on Computer Vision (ICCV’99)*, (Kerkyra, Greece), pp. 14–21, September 1999.
- [178] H.-Y. Shum and R. Szeliski, “Construction of panoramic mosaics with global and local alignment,” *International Journal of Computer Vision*, vol. 36, no. 2, pp. 101–130, February 2000, Erratum published July 2002, vol. 48, no. 2, pp. 151–152, 2000.
- [179] H.-Y. Shum *et al.*, “Omnivergenet stereo,” in *Seventh International Conference on Computer Vision (ICCV’99)*, (Greece), pp. 22–29, September 1999.
- [180] E. P. Simoncelli, E. H. Adelson, and D. J. Heeger, “Probability distributions of optic flow,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’91)*, (Maui, Hawaii), pp. 310–315, IEEE Computer Society Press, June 1991.
- [181] C. C. Slama, ed., *Manual of Photogrammetry*. Fourth Edition, Falls Church, Virginia, 1980. American Society of Photogrammetry.
- [182] P. Soille, “Morphological image compositing,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 5, pp. 673–683, May 2006.
- [183] D. Steedly *et al.*, “Efficiently registering video into panoramic mosaics,” in *Tenth International Conference on Computer Vision (ICCV 2005)*, (Beijing, China), pp. 1300–1307, October 2005.

- [184] R. Steele and C. Jaynes, "Feature uncertainty arising from covariant image noise," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2005)*, (San Diego, CA), vol. I, pp. 1063–1070, June 2005.
- [185] R. M. Steele and C. Jaynes, "Overconstrained linear estimation of radial distortion and multi-view geometry," in *Computer Vision – ECCV 2006*, (A. Leonardis, H. Bischof, and A. Pinz, eds.), vol. 3951, pp. 253–264, of *Lecture Notes in Computer Science*, Springer, 2006.
- [186] G. Stein, "Accurate internal camera calibration using rotation, with analysis of sources of error," in *Fifth International Conference on Computer Vision (ICCV'95)*, (Cambridge, Massachusetts), pp. 230–236, June 1995.
- [187] G. Stein, "Lens distortion calibration using point correspondences," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'97)*, (San Juan, Puerto Rico), pp. 602–608, June 1997.
- [188] C. V. Stewart, "Robust parameter estimation in computer vision," *SIAM Reviews*, vol. 41, no. 3, pp. 513–537, September 1999.
- [189] P. Sturm, "Multi-view geometry for general camera models," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2005)*, (San Diego, CA), vol. I, pp. 206–212, June 2005.
- [190] J. Sun, N. Zheng, and H. Shum, "Stereo matching using belief propagation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 7, pp. 787–800, July 2003.
- [191] R. Szeliski, *Bayesian Modeling of Uncertainty in Low-Level Vision*. Boston: Kluwer Academic Publishers, 1989.
- [192] R. Szeliski, "Fast surface interpolation using hierarchical basis functions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 6, pp. 513–528, June 1990.
- [193] R. Szeliski, "Image mosaicing for tele-reality applications," in *IEEE Workshop on Applications of Computer Vision (WACV'94)*, (Sarasota), pp. 44–53, IEEE Computer Society, December 1994.
- [194] R. Szeliski, "Video mosaics for virtual environments," *IEEE Computer Graphics and Applications*, vol. 16, no. 2, pp. 22–30, March 1996.
- [195] R. Szeliski, "Locally adapted hierarchical basis preconditioning," *ACM Transactions on Graphics*, vol. 25, no. 3, pp. 1135–1143, August 2006.
- [196] R. Szeliski and J. Coughlan, "Hierarchical spline-based image registration," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'94)*, (Seattle), pp. 194–201, IEEE Computer Society, June 1994.
- [197] R. Szeliski and S. B. Kang, "Recovering 3D shape and motion from image streams using nonlinear least squares," *Journal of Visual Communication and Image Representation*, vol. 5, no. 1, pp. 10–28, March 1994.
- [198] R. Szeliski and S. B. Kang, "Direct methods for visual scene reconstruction," in *IEEE Workshop on Representations of Visual Scenes*, (Cambridge, Massachusetts), pp. 26–33, June 1995.
- [199] R. Szeliski and H.-Y. Shum, "Creating full view panoramic image mosaics and texture-mapped models," in *Computer Graphics (SIGGRAPH'97 Proceedings)*, pp. 251–258, August 1997.

- [200] M. F. Tappen and W. T. Freeman, “Comparison of graph cuts with belief propagation for stereo, using identical MRF parameters,” in *Ninth International Conference on Computer Vision (ICCV 2003)*, (Nice, France), pp. 900–907, October 2003.
- [201] J.-P. Tardif, P. Sturm, and S. Roy, “Self-calibration of a general radially symmetric distortion model,” in *Computer Vision – ECCV 2006*, (A. Leonardis, H. Bischof, and A. Pinz, eds.), vol. 3954, pp. 186–199, of *Lecture Notes in Computer Science*, Springer, 2006.
- [202] J.-P. Tardif *et al.*, “Self-calibration of a general radially symmetric distortion model,” in *Seventh European Conference on Computer Vision (ECCV 2002)*, (Graz), vol. 4, pp. 186–199, Springer-Verlag, May 2006.
- [203] L. Teodosio and W. Bender, “Salient video stills: Content and context preserved,” in *ACM Multimedia 93*, (Anaheim, California), pp. 39–46, August 1993.
- [204] S. Thirithala and M. Pollefeys, “The radial trifocal tensor: A tool for calibrating the radial distortion of wide-angle cameras,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2005)*, (San Diego, CA), vol. I, pp. 321–328, June 2005.
- [205] Q. Tian and M. N. Huhns, “Algorithms for subpixel registration,” *Computer Vision, Graphics, and Image Processing*, vol. 35, pp. 220–233, 1986.
- [206] B. Triggs, “Detecting keypoints with stable position, orientation, and scale under illumination changes,” in *Eighth European Conference on Computer Vision (ECCV 2004)*, (Prague), vol. IV, pp. 100–113, Springer-Verlag, May 2004.
- [207] B. Triggs *et al.*, “International Workshop on Vision Algorithms,” in *Bundle adjustment – a modern synthesis*, (Kerkyra, Greece), pp. 298–372, Springer, September 1999.
- [208] J. Tumblin, A. Agrawal, and R. Raskar, “Why i want a gradient camera,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2005)*, (San Diego, CA), vol. I, pp. 103–110, June 2005.
- [209] T. Tuytelaars and L. Van Gool, “Matching widely separated views based on affine invariant regions,” *International Journal of Computer Vision*, vol. 59, no. 1, pp. 61–85, August 2004.
- [210] M. Uyttendaele, A. Eden, and R. Szeliski, “Eliminating ghosting and exposure artifacts in image mosaics,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’2001)*, (Kauai, Hawaii), vol. II, pp. 509–516, December 2001.
- [211] M. Uyttendaele *et al.*, “Image-based interactive exploration of real-world environments,” *IEEE Computer Graphics and Applications*, vol. 24, no. 3, May/June 2004.
- [212] S. v. Huffel and J. Vandewalle, “The Total Least Squares Problem: Computational Aspects and Analysis,” in *Society for Industrial and Applied Mathematics*, (Philadelphia), 1991.
- [213] J. van de Weijer and C. Schmid, “Coloring local feature extraction,” in *Computer Vision – ECCV 2006*, (A. Leonardis, H. Bischof, and A. Pinz, eds.), vol. 3952, pp. 334–348, of *Lecture Notes in Computer Science*, Springer, 2006.

- [214] P. Viola and W. Wells III, "Alignment by maximization of mutual information," in *Fifth International Conference on Computer Vision (ICCV'95)*, (Cambridge, Massachusetts), pp. 16–23, June 1995.
- [215] L. Wang, S. B. Kang, R. Szeliski, and H.-Y. Shum, "Optimal texture map reconstruction from multiple views," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2001)*, (Kauai, Hawaii), vol. I, pp. 347–354, December 2001.
- [216] A. Watt, *3D Computer Graphics*. Addison-Wesley, Third Edition, 1995.
- [217] J. Weber and J. Malik, "Robust computation of optical flow in a multi-scale differential framework," *International Journal of Computer Vision*, vol. 14, no. 1, pp. 67–81, January 1995.
- [218] Y. Weiss, "Deriving intrinsic images from image sequences," in *Eighth International Conference on Computer Vision (ICCV 2001)*, (Vancouver, Canada), vol. II, pp. 7–14, July 2001.
- [219] L. Williams, "Pyramidal parametrics," *Computer Graphics*, vol. 17, no. 3, pp. 1–11, July 1983.
- [220] D. N. Wood *et al.*, "Multiperspective panoramas for cell animation," in *Computer Graphics Proceedings, Annual Conference Series*, (Los Angeles), pp. 243–250, Proc. SIGGRAPH'97, ACM SIGGRAPH, August 1997.
- [221] R. J. Woodham, "Analysing images of curved surfaces," *Artificial Intelligence*, vol. 17, pp. 117–140, 1981.
- [222] Y. Xiong and K. Turkowski, "Creating image-based VR using a self-calibrating fisheye lens," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'97)*, (San Juan, Puerto Rico), pp. 237–243, June 1997.
- [223] Y. Xiong and K. Turkowski, "Registration, calibration and blending in creating high quality panoramas," in *IEEE Workshop on Applications of Computer Vision (WACV'98)*, (Princeton), pp. 69–74, IEEE Computer Society, October 1998.
- [224] D. Yang *et al.*, "A 640x512 CMOS image sensor with ultra-wide dynamic range floating-point pixel level ADC," *IEEE Journal of Solid State Circuits*, vol. 34, no. 12, pp. 1821–1834, December 1999.
- [225] R. Zabih and J. Woodfill, "Non-parametric local transforms for computing visual correspondence," in *Third European Conference on Computer Vision (ECCV'94)*, (Stockholm, Sweden), vol. II, pp. 151–158, Springer-Verlag, May 1994.
- [226] B. Zitov'aa and J. Flusser, "Image registration methods: A survey," *Image and Vision Computing*, vol. 21, pp. 997–1000, October 2003.
- [227] I. Zoghlami, O. Faugeras, and R. Deriche, "Using geometric corners to build a 2D mosaic from a set of images," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'97)*, (San Juan, Puerto Rico), pp. 420–425, June 1997.