CSE 252B: Computer Vision II

Lecturer: Serge Belongie
Scribes: Jia Mao, Andrew Rabinovich

LECTURE 9
Affine and Euclidean Reconstruction

9.1. Stratified reconstruction

Recall that in 3D reconstruction from two uncalibrated views, we can only obtain the structure of the scene up to a projective transformation, denoted $\boldsymbol{X}_p$.

The true structure $\boldsymbol{X}$ is related to $\boldsymbol{X}_p$ by

$$\boldsymbol{X}_p = H_p H_a g_e \boldsymbol{X}$$

where $\boldsymbol{X}_e = g_e \boldsymbol{X}$ differs from $\boldsymbol{X}$ by a Euclidean transformation $g_e = (R_e, T_e)$, $\boldsymbol{X}_a = H_a g_e \boldsymbol{X}_e$ differs from $\boldsymbol{X}_e$ by a general affine transformation, and $\boldsymbol{X}_p$ differs from $\boldsymbol{X}_a$ by a general projective transformation. Figure 1 illustrates the different structures obtained at each stratum of the reconstruction process.

[1]Department of Computer Science and Engineering, University of California, San Diego.

April 26, 2004

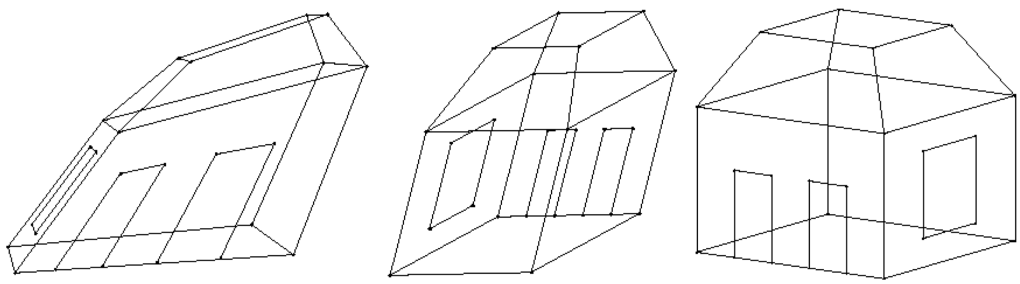

Figure 1. Projective structure $\boldsymbol{X}_p$, affine structure $\boldsymbol{X}_a$, and Euclidean structure $\boldsymbol{X}_e$ obtained in different stages of reconstruction (MaSKS fig6.10)

In this lecture, we will describe the stratified reconstruction approach in 3D. Namely, first recover the affine structure $\boldsymbol{X}_a$ by identifying the matrix H_p^{-1}, and then recover the Euclidean structure $\boldsymbol{X}_e$ by identifying the matrix H_a^{-1}. Moreover, we will see that one can go directly from projective to Euclidean structure if we know the "ground truth" coordinates of 5 points in the scene.

Analogous to the stratified reconstruction in 2D, the tools for us to do the affine and Euclidean upgrades in 2D and 3D are listed in Table 9.1.

	Affine upgrade	Euclidean upgrade
2D	line at infinity $\boldsymbol{\ell}_\infty = (0,0,1)^\top$	circular points $\boldsymbol{I}, \boldsymbol{J} = (1, \pm i, 0)^\top$
3D	plane at infinity $\boldsymbol{\pi}_\infty = (0,0,0,1)^\top$	absolute conic Ω_∞

Table 1. Useful entities for 2D vs. 3D stratified reconstruction

9.2. Affine upgrade

The line at infinity $\boldsymbol{\ell}_\infty$ is used to do affine upgrade in the 2D reconstruction because it stays fixed under an affine transformation. In the 3D case, the plane at infinity $\boldsymbol{\pi}_\infty$ is the corresponding entity that allows us to achieve the same goal.

9.2.1. Plane at infinity $\boldsymbol{\pi}_\infty$

Parallelism is preserved in the affine structure. The plane at infinity $\boldsymbol{\pi}_\infty$ enables us to identify parallelism. In particular,

- Two planes are parallel iff the line of intersection is on $\boldsymbol{\pi}_\infty$.
- Two lines are parallel iff the point of intersection is on $\boldsymbol{\pi}_\infty$.

Points on $\boldsymbol{\pi}_\infty$ are ideal points of the form $\boldsymbol{X} = (X, Y, Z, 0)^\top$ satisfying the equation

$$(0, 0, 0, 1)^\top \boldsymbol{X} = 0$$

If we can identify the image of $\boldsymbol{\pi}_\infty$ (denoted by $(\boldsymbol{v}^\top, v_4)^\top$), then we can do the affine upgrade using

$$H_p^{-1} = \begin{bmatrix} I & \mathbf{0} \\ \boldsymbol{v}^\top & v_4 \end{bmatrix}$$

such that $\boldsymbol{X}_a = H_p^{-1}\boldsymbol{X}_p$ and $\Pi_{ia} = \Pi_{ip} H_p$, $i = 1, 2$.

As with any plane, we need 3 points to specify $\boldsymbol{\pi}_\infty$, which we can get from three vanishing points, assuming they are measurable.

9.2.2. Vanishing points

Perspective projection renders intersections of parallel lines to finite points on the image plane, called the *vanishing points*, as illustrated in Figure 2. These three vanishing points define $\boldsymbol{\pi}_\infty$. Any two parallel lines in the 3D projective reconstruction will, in theory, intersect at the same vanishing point. However, intersecting lines in 3D is problematic in the presence of noise.

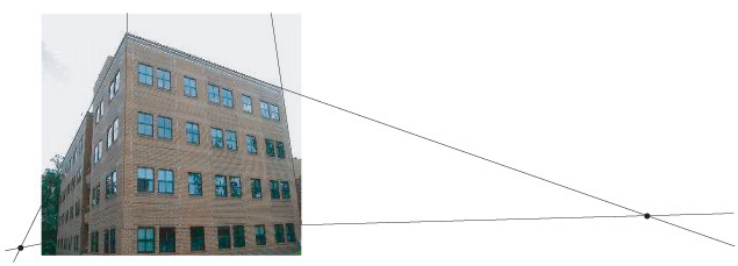

Figure 2. Parallel lines and the associated vanishing points. (MaSKS Figure 6.12) Two vanishing points are shown; the third is off the page far above the picture.

As an alternative, we can find the 3D coordinates of a vanishing point purely from 2D data in the two image planes. H&Z suggest the following algorithm:

(1) Find vanishing point $\boldsymbol{x}'_1$ in image 1 from a pair of imaged parallel lines.
(2) Compute epipolar line in image 2: $\boldsymbol{\ell}_2 = F\boldsymbol{x}'_1$. (The point in the second image corresponding to $\boldsymbol{x}'_1$ must lie on this line.)
(3) Intersect $\boldsymbol{\ell}_2$ with a corresponding parallel line in image 2, call it $\boldsymbol{x}'_2$.
(4) Find projective depth of $(\boldsymbol{x}'_1, \boldsymbol{x}'_2)$.

Note that step (2) ensures that the backprojected rays through $\boldsymbol{x}'_1$ and $\boldsymbol{x}'_2$ in 3D will intersect.

9.2.3. $\boldsymbol{\pi}_\infty$ for affine upgrade

(MaSKS Example 6.5) Now that we have three vanishing points $\boldsymbol{X}_p^j$, $j = 1, 2, 3$, we can then solve for $(\boldsymbol{v}^\top, v_4)^\top$ from the linear system

$$(v_1, v_2, v_3, v_4)\boldsymbol{X}_p^j = 0 \qquad j = 1, 2, 3$$

which then fully specifies $H_p^{-1} = \begin{bmatrix} I & \mathbf{0} \\ \boldsymbol{v}^\top & v_4 \end{bmatrix}$ as desired.

Other alternative methods also exist for affine upgrade, with more of them coming out each year. MaSKS §6.4.3 gives two examples, one of which is to exploit the case of pure translation (i.e. $R = I$), the other is to exploit the equal modulus constraint.

9.3. Direct upgrade from $\boldsymbol{X}_p$ to $\boldsymbol{X}_e$ using ground-truth points

An example of direct upgrade from projective to Euclidean structure in the 2D case is the four-point algorithm (a.k.a. DLT). It utilizes four ground-truth correspondences of which no three in each image are collinear to recover the planar homography matrix.

Similarly, in 3D, H has size 4×4, which has $16-1$(for scale)=15 degrees of freedom. Therefore, 5 ground-truth points suffice as long as they are in general position (i.e., no 4 points are coplanar). An example of this direct upgrade is shown in Figure 3. The analogous five-point algorithm is left as an exercise.

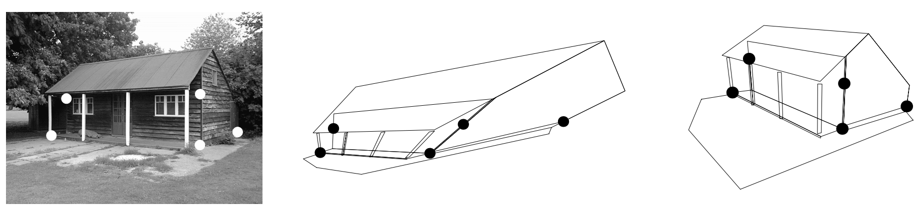

Figure 3. Direct upgrade from projective to Euclidean using 5 ground-truth points, which are in general position (i.e., no 4 points are coplanar). (H&Z Figure 9.6).

9.4. The Absolute Conic

In the absence of sufficient ground truth points, a versatile alternative is to use the *absolute conic* Ω_∞ or its dual, the *absolute dual quadric* Q_∞^*.

What is the absolute conic? The absolute conic lives on the plane at infinity $\boldsymbol{\pi}_\infty$. It cannot be seen – it can only be inferred. It is a point conic (as opposed to a line conic); points of the form $\boldsymbol{X} = (X, Y, Z, W)^\top$ on Ω_∞ satisfy:

$$X^2 + Y^2 + Z^2 = 0 \tag{9.1}$$
$$W = 0 \tag{9.2}$$

Thus for ideal points (i.e. points with $W = 0$) on $\boldsymbol{\pi}_\infty$, the defining equation is:

$$\begin{bmatrix} X & Y & Z \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = 0$$

From this we see $\Omega_\infty = I$. This is an imaginary conic.

9.4.1. Some Properties of the Absolute Conic

(1) All spheres intersect the plane at infinity in the absolute conic.
(2) All circles intersect the absolute conic in two points.

Suppose a circle lives in a plane $\boldsymbol{\pi}$. Then $\boldsymbol{\pi}$ intersects $\boldsymbol{\pi}_\infty$ in a line, and that line intersects Ω_∞ in two points. These two points are the circular points of the plane $\boldsymbol{\pi}$.

9.4.2. Absolute Conic and the Calibration Matrix

Hartley has shown how to find the image of the absolute conic (IAC) Ω'_∞ and factor it to get K, the calibration matrix. In particular, $\Omega'_\infty = S = (KK^\top)^{-1}$. This is the metric of the uncalibrated space.

Consider the relationship between points on the plane at infinity and the image plane. We can write these points as $\boldsymbol{X}_\infty = (\boldsymbol{d}^\top, 0)^\top$, and when imaged by a general camera, we get

$$\lambda \boldsymbol{x} = \Pi \boldsymbol{X}_\infty = KR[I, \boldsymbol{T}] \begin{bmatrix} \boldsymbol{d} \\ 0 \end{bmatrix} = KR\boldsymbol{d}$$

which does not depend on $\boldsymbol{T}$. It only depends on K (the calibration matrix), and R (the orientation w.r.t. the world frame), which will drop out.

You can observe this phenomenon by looking out the window of a moving car at a distant point, e.g. on a mountain or on the moon. While objects near the car appear to move rapidly, these distant points appear fixed on your retina.

Recall that the absolute conic lives on $\boldsymbol{\pi}_\infty$ and we know how to compute its image under the homography H. Under the mapping $\boldsymbol{x}'_2 = H\boldsymbol{x}'_1$, the conic $C \mapsto H^{-\top} C H^{-1}$, so

$$\Omega_\infty = I \mapsto (KR)^{-\top} I (KR)^{-1} = K^{-\top} R R^\top K^{-1} = (KK^\top)^{-1}$$

We thus have the result that the image of the absolute conic (IAC) is the conic

$$\Omega'_\infty = S = (KK^\top)^{-1}$$

A consequence of this is that if we can identify Ω'_∞, then we can get K! H&Z show in Example 7.17 (8.18 in the 2nd edition) how to do this using an image containing three squares.

9.4.3. Lyapunov Map

Underlying our search for the Euclidean upgrade is the following question: given $C = KRK^{-1}$, how much does C reveal about $S = (KK^\top)^{-1}$? One can express the relationship between C and S using a Lyapunov (or Sylvester) equation:

$$S^{-1} - CS^{-1}C^\top = 0$$

This is a matrix equation; the right hand side is a zero matrix of size 3×3. Hartley shows how two such matrix equations for rotations around different axes can be used to estimate C and recover K. We will discuss this in the next lecture.