Multiple View Geometry

Lecture Set 01
Introduction

林惠勇
Huei-Yung Lin

lin@ee.ccu.edu.tw

Robot Vision Lab
Department of Electrical Engineering
National Chung Cheng University
Chiayi 621, Taiwan

Course Topics

- Introduction to Multiple View Geometry
- 2D Projective Geometry
- 3D Projective Geometry
- Estimation (2D Homography)
- Camera Models
- Camera Calibration
- More on Single View Geometry
- The Epipolar Geometry
- 3D Reconstruction
- Computing the Fundamental Matrix
- Structure Computation
- Planes and Homographies
- The Trifocal Tensor
- Three-View Computation

Useful Resource

- `MATLAB Functions for Multiple View Geometry`
- `Peter Kovesi's Matlab functions for Computer Vision`
- `Camera Calibration Toolbox for Matlab`

- Please print out the lecture notes and lecture slides.
- Note: This is a rather high level computer vision course, the students are expected to have introductory computer vision and image processing background.

Some Application Demos

- Markerless Motion Capture #1 (Full Body Tracking from Volumetric Data)
- Markerless Motion Capture #2 (Model Fiting to Multiple 2D Projections)
- Multi-Camera Surveillance
- Multi-Camera Recognition
- Augmented Reality #1
- Augmented Reality #2
- Navigation (Building Facade Reconstruction and Vehicle Detection)
- 3D Reconstruction from Photos
- License Plate Recognition
- Virtual Window #1
- Virtual Window #2

The Ubiquitous Projective Geometry

- What is *projective transformation*? Is it seen in your daily life?
  - A coin appears as an ellipse instead of a circle.
  - The blackboard looks like a quadrilateral instead of rectangle.
  - These are examples of “planar” projective transformation.

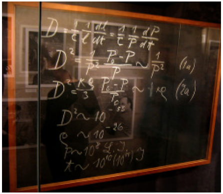

The Ubiquitous Projective Geometry

- What properties are preserved under projective transformation?
  - Shape? – No.
  - Lengths? – No.
  - Angles? – No.
  - Distances? – No.
  - Straightness? – Yes!
- Thus, a projective transformation can be defined by *a mapping that preserves straight lines*.
  - On a 2D plane, a set of points forms a straight line map to a set of points forms another (or identical) straight line.


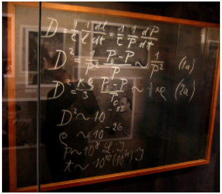

Projective Space & Euclidean Space

- Problem with the familiar Euclidean geometry:
  - Some exceptions are always required...
  - *Two lines do not always meet on a plane.* – It is called parallel.
  - *Two planes do not always intersect in 3-space.*
  - We can say that they meet at "infinity".
- By adding the points at infinity, the *Euclidean space* is transformed to the "*projective space*".
  - Two lines in the 2-space always meet.
  - Parallel lines intersect at the "ideal points".
- *Projective space can be thought as an extension of Euclidean space with ideal points.*
- The properties (such as distances, angles, etc.) of Euclidean space will be used for computational purposes.

Coordinates

- A point is represented by (x, y) in the Euclidean 2-space.
- In *homogeneous coordinates*, a 2-D point can be represented by $(x, y, 1)$ or (kx, ky, k), where $k \neq 0$.
  - There is a 1-1 mapping from (kx, ky, k) to (x, y) if $k \neq 0$. Just scale the first two entries by $1/k$.
  - The homogeneous representation $(0, 0, k)$ maps to the origin $(0, 0)$ if $k \neq 0$.
- Questions: What point in 2-space is represented by $(x, y, 0)$?
  - The points $(x, y, 0) \mapsto (x/0, y/0)$ which is the points at infinity.
  - This is how the points at infinity arise.
- The Euclidean space $\mathbb{R}^n$ can be extended to a projective space $\mathbb{P}^n$ by representing points as *homogeneous vectors*.
  - The points at infinity in 2-D projective space form a line – *the line at infinity*.
  - The points at infinity in 3-D projective space form a plane – *the plane at infinity*.

Homogeneity

- In classical Euclidean geometry all points are the same.
  - When coordinates are added, one point is picked out as the *origin*.
- *Euclidean transform*: A change of coordinates in which the axes of the coordinate frame are *shifted* and *rotated* to a different position.
- *Affine transformation*: Applying a linear transformation to $\mathbb{R}^n$, followed by a Euclidean translation.
- *The points at infinity remain at infinity for both Euclidean and affine transformation.*
  - The points at infinity are *preserved* by these transformations.
  - The points at infinity are special in both Euclidean and affine geometry.
- *The points at infinity are no different from other points in PG.*
  - The points with "zero final coordinate" in a homogeneous coordinate representation is just an accident of *the choice of coordinate frame.*
  - Those points can be transformed to arbitary other points (with "non-zero final coordinate") by a *projective transformation*.
  - The points at infinity are NOT preserved in projective space.

Projective Transformation

- A projective transformation of projective space $\mathbb{P}^n$ is represented by a linear transformation of homogeneous coordinates

$$\mathbf{X}' = \mathtt{H}_{(n+1)\times(n+1)}\mathbf{X}$$

- *In computer vision problems, projective space is used as a convenient way of representing the real 3-D world, by extending it to the 3-D projective space.*
- *Images are extended to be thought of as lying in the 2-D projective space.*
- Although we usually work with the projective spaces, we are aware that the line and plane at infinity are in some way special.
  - Usually we treat all points in projective space as equals when it suits us, and singling out the plane at infinity in space or the line at infinity in the image when that becomes necessary.

Projective Space

- The projective space is initially *homogeneous*, with no particular coordinate frame being preferred.

- There is no concept of parallelism of lines – Parallel lines (or planes in the three-dimensional case) are ones that meet at infinity.
- There is no concept of which points are at infinity – All points are created equal, parallelism is not a concept of projective geometry.

- In order for such a concept to make sense *to us in the familiar 3D world*, we need to pick out some particular line, and decide that this is the line at infinity.
  - This results in a situation where although all points are created equal, some are more equal than others.
  - This forms the *affine geometry*! (Later.)

2-D Projective Space

- The points at infinity in the plane show up in the image as the horizontal line.
  - Railway tracks show up in the image as lines meeting at the horizon.
- Points in the image lying above the horizon (the image of the sky) meet the plane at a point behind the camera.
- There's a 1-1 relationship between points in the image and points in the world plane:
  - *Points at infinity in the world plane correspond to a horizon line in the image.*
  - *Parallel lines in the world correspond to lines meeting at the horizon.*
- The world plane and its image are just alternative ways of viewing the geometry of a projective plane, *plus a distinguished line*.

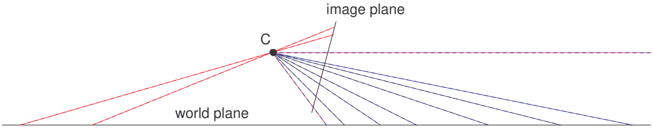

2-D Affine Geometry

- *Affine geometry*: The geometry of the projective plane and a distinguished line.
- *Affine transforamtion*: Any projective transformation that maps *the distinguished line* in one space to *the distinguished line* of the other space.
- By identifying a special line as the “line at infinity” we are able to define parallelism of straight lines in the plane. (How?)
- Summary:
    - By distinguishing a special line in a projective plane, we gain the concept of parallelism and with it *affine geometry*.
    - Affine geometry is seen as specialization of projective geometry, in which we single out a particular line and call it *the line at infinity*.
- How about Euclidean geometry?
    - Can we single out some special feature of the line or plane at infinity to make affine geometry become Euclidean geometry?
    - Yes! – with the concept of *circular points* and *absolute conic*.

Circle & Ellipse

- A circle is not a concept of affine geometry:
  - Stretching the plane preserves the line at infinity, but turns a circle into ellipse.
  - Thus, affine geometry does not distinguish between circles and ellipses.
  - However, circles and ellipses are distinct in Euclidean geometry.
- An ellipse or a circle is described by a second-degree equation.
  - Any two of them should generally intersect in four points.
  - This is not true for two distinct circles!
  - Two distinct circles intersect in no more than two points. Why?
  - The other two intersections are *complex points*!
- The equation for a circle in homogeneous coordinates (x, y, w) is of the form $(x/w - a)^2 + (y/w - b)^2 = r^2$
  - The points $(x, y, w)^\top = (1, \pm i, 0)^\top$ lie on every such circle. (Check!)
  - Thus, they lie on the intersection of any two circles.
  - Since their final coordinate is zero, these two points lie on *the line at infinity*.
  - They are called the *circular points* of the plane.

Euclidean Geometry

- Although the two circular points are complex, they satisfy a pair of real equations:

$$\begin{cases} x^2 + y^2 = 0 \\ w = 0 \end{cases}$$

- In some sense, they are *fixed* in Euclidean geometry.
    - We may define a circle as a conic that passes through the two circular points.
    - Concepts such as angle and length ratios may be defined in terms of these circular points. (Later!)
- *Euclidean geometry arises from projective geometry by singling out first a line at infinity and subsequently, two points called circular points lying on this line.*
- In the standard Euclidean coordinate system, the circular points have the coordinates $(1, \pm i, 0)^\top$.
    - In assigning a Euclidean structure to a projective plane, however, we may designate *any* line and any two (complex) points on that line as being the line at infinity and the circular points. (Why?)

3-D Euclidean Geometry

- Recap: The Euclidean plane is defined in terms of the projective plane by specifying a line at infinity and a pair of circular points.
- For the 3-D Euclidean space:
  - Two spheres intersect in a circle – a two-degree curve.
  - But two general ellipsoids intersect in a four-degree curve.
  - Similar to the 2-D case, in homogeneous coordinates $(X, Y, Z, T)^\top$ all spheres intersect the plane at infinity in a curve with the equations:

$$\begin{cases} X^2 + Y^2 + Z^2 = 0 \\ T = 0 \end{cases}$$

  - This is a second-degree curve (a conic) lying on the plane at infinity, and consisting only of complex points.
  - It is known as the *absolute conic*.
- We may consider 3D Euclidean space to be derived from projective space by singling out a particular plane as *the plane at infinity* and specifying a particular conic lying in this plane to be *the absolute conic*.

Camera Projection

- The usual way of modeling the process of image formation is by *central projection*.
  - In applying projective geometry to the imaging process, the world can be modeled as a 3-D projective space, equal to $\mathbb{R}^3$ along with points at infinity.
  - Similarly the model for the image is the 2-D projective plane $\mathbb{P}^2$.
  - Central projection is simply a map from $\mathbb{P}^3$ to $\mathbb{P}^2$.

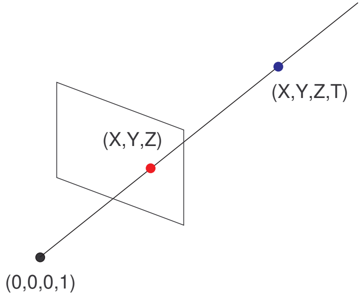

Camera Projection

- Consider points in $\mathbb{P}^3$ written in terms of homogeneous coordinates $(X, Y, Z, T)^\top$ and let the center of projection (COP) be the origin $(0, 0, 0, 1)^\top$
  - The set of all points $(X, Y, Z, T)^\top$ for fixed X, Y and Z, but varying T form a single ray passing through the COP, and hence all mapping to the same point.
  - The final coordinate of $(X, Y, Z, T)^\top$ is irrelevant to where the point is imaged.
  - The image point is the point in $\mathbb{P}^2$ with homogeneous coordinates $(X, Y, Z)^\top$.
  - The mapping may be represented by a mapping of 3-D homogeneous coordinates, represented by a 3×4 matrix $\mathtt{P} = [\ \mathtt{I}_{3\times 3} \mid \mathbf{0}_3\ ]$ (How?).

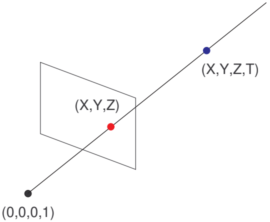

Camera Projection

- The matrix $\mathtt{P}$ is known as the *camera matrix*.
  - The action of a projective camera on a point in space may be expressed in terms of a linear mapping of homogeneous coordinates as

$$\begin{pmatrix} x \\ y \\ w \end{pmatrix} = \mathtt{P}_{3\times 4} \begin{pmatrix} X \\ Y \\ Z \\ T \end{pmatrix}$$

  - If all the points lie on a plane then the linear mapping reduces to a *projective transformation*:

$$\begin{pmatrix} x \\ y \\ w \end{pmatrix} = \mathtt{H}_{3\times 3} \begin{pmatrix} X \\ Y \\ T \end{pmatrix}$$

Cameras as Points

- In a central projection, points in $\mathbb{P}^3$ are mapped to points in $\mathbb{P}^2$, all points in a ray passing through the COP projecting to the same point in an image.
- We can think of the ray through the COP as representing the image point.
- The set of all image points is the same as the set of rays through the camera center.
- If we represent the ray from $(0,0,0,1)^\top$ through the point $(X,Y,Z,T)^\top$ by $(X,Y,Z)^\top$, then for any constant k, the ray $k(X,Y,Z)^\top$ represents the same ray.
  - The rays themselves are represented by homogeneous coordinates.
  - They make up a 2-dimensional space of rays.
  - The set of rays themselves may be thought of as a representation of the image space $\mathbb{P}^2$.
- In this representation of the image, *only the camera center is important*.
  - The camera center alone determines the set of rays forming the image.
  - *Two images taken from "the same point" in space are projective equivalent.*

Cameras as Points

- Only when we start to *measure* points in an image, that a particular coordinate frame for the image needs to be specified.
  - A particular camera matrix might also be needed.

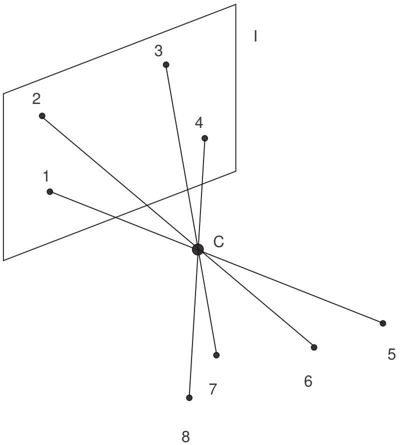

Cameras as Points

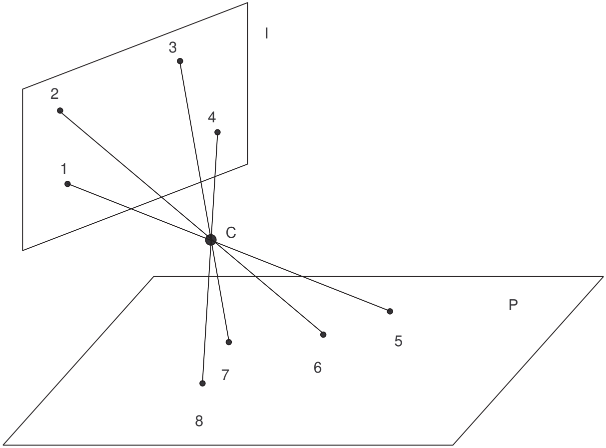

Cameras as Points

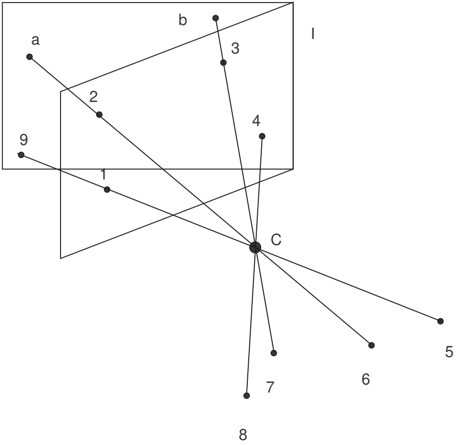

Cameras as Points

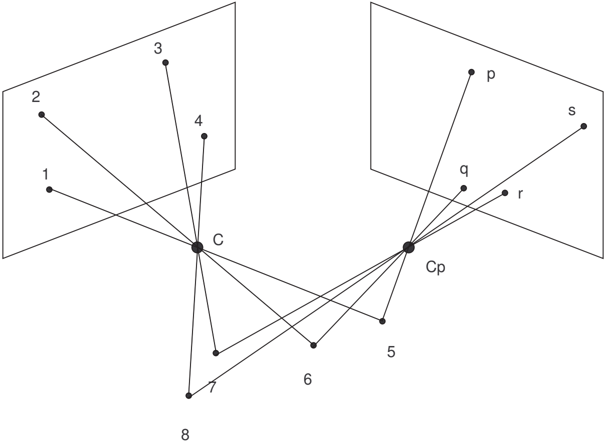

Calibrated Cameras

- The Euclidean geometry of the 3-D world is determined by specifying
  - a particular plane in $\mathbb{P}^3$ as being the plane at infinity;
  - a specific conic Ω in that plane as beging the absolute conic.
- For a camera *not* located on the plane at infinity, the plane at infinity in the world maps one-to-one onto the image plane.
  - Because any point in the image defines a ray in space that meets the plane at infinity in a single point.
- Thus, the plane at infinity does not tell us anything new about the image.
- The absolute conic (in the plane at infinity) must project to a conic in the image.
- The resulting image curve is called the *Image of the Absolute Conic*, or IAC.
- If the location of IAC is known in an image, then we say the camera is *calibrated*.

Reconstruction from Two Views (1/2)

- In the two-view case we consider a set of correspondences $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ in two images.
- There exist some camera matrices, $\mathtt{P}$ and $\mathtt{P}'$ and a set of 3-D points $\mathbf{X}_i$ such that $\mathtt{P}\mathbf{X}_i = \mathbf{x}_i$ and $\mathtt{P}'\mathbf{X}_i = \mathbf{x}'_i$. That is, the point $\mathbf{X}_i$ projects to the two given data points.
  - $\mathtt{P}$, $\mathtt{P}'$ and $\mathbf{X}_i$ are unknown. – It is our task to determine them.
- The reconstruction is possible at best up to a *similarity transformation* of the world.
- If the camera calibration information are *not* available, the reconstruction is possible at best up to a projective transformation – *a projective reconstruction*.

Reconstruction from Two Views (2/2)

- The basic tool in the reconstruction from two views is the *fundamental matrix*.
  - It arises from the coplanarity of camera centers, image points and space point.
  - It satisfies $\mathbf{x}_i'^{\top} \mathtt{F} \mathbf{x}_i = 0$ where $\mathtt{F}$ is a 3×3 matrix of rank 2.
  - It can be computed from a set of point correspondences.
- A pair of camera matrices $\mathtt{P}$ and $\mathtt{P}'$ uniquely determine a fundamental matrix $\mathtt{F}$.
- The fundamental matrix determines the pair of camera matrices, up to a 3-D projective ambiguity.
- The fundamental matrix encapsulates the complete projective geometry of the pair of cameras, and is unchanged by projective transformation of 3-D.

Transfer (1/2)

- A useful application of projective geometry is that of *transfer*:
  - Given the position of a point in one (or more) image(s), determine where it will appear in all other images of the set.
- To do this, we must first establish the relationship between the cameras using (for instance) a set of auxiliary point correspondences.
- Suppose the point is identified in two views (at $\mathbf{x}$ and $\mathbf{x}'$) and we wish to know its position $\mathbf{x}''$ in a third, then this may be computed by the following steps:
  (i) Compute the camera metrices of the three views $\mathtt{P}, \mathtt{P}', \mathtt{P}''$ from other point correspondences $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i \leftrightarrow \mathbf{x}''_i$.
  (ii) Triangulate the 3D point $\mathbf{X}$ from $\mathbf{x}$ and $\mathbf{x}'$ using $\mathtt{P}$ and $\mathtt{P}'$.
  (iii) Project the 3D point into the third view as $\mathbf{x}'' = \mathtt{P}''\mathbf{X}$.
- This procedure only requires projective information.

Transfer (2/2)

- Suppose the camera rotates about its center *or* that all the scene points of interest lie on a plane.
- Then the appropriate multiple view relations are the planar projective transformations between the images.
- In this case, a point seen in just one image can be transferred to any other image.

Euclidean Reconstruction

- Projective reconstruction is insufficient for many purposes, such as application to computer graphics.
- In order to obtain a reconstruction of the model in which objects have their correct (Euclidean) shape, it is necessary to determine the calibration of the cameras.
- Determining the Euclidean structure of the world is equivalent to specifying the plane at infinity and the absolute conic.
- Since the absolute conic lies in a plane, the plane at infinity, it is enough to find the absolute conic in space.
- The knowledge of the camera calibration is equivalent to being able to determine the Euclidean structure of the scene.

Applications

- 3D Graphical Models
  - 3D Reconstruction from Video #1
  - 3D Reconstruction from Video #2
  - 3D Reconstruction from Video #3
- Video Augmentation
  - Markerless Augmented Reality
  - Augmentation with Real-Time SLAM with a Single Camera
- Historical Site Reconstruction