

# Introduction to 3D Vision

- How do we obtain 3D image data?
- What can we do with it?

What can you determine about

1. the sizes of objects

2. the distances of objects from the camera?



What knowledge  
do you use to  
analyze this image?

What objects are shown in this image?

How can you estimate distance from the camera?

What feature changes with distance?



# 3D Shape from X


- shading
- silhouette
- texture

) mainly research

- stereo
- light striping
- motion

) used in practice

# Perspective Imaging Model: 1D




This is the axis of the real image plane.

$O$  is the center of projection.

This is the axis of the **front image plane**, which we use.

$$\frac{x_i}{f} = \frac{x_C}{z_C}$$

# Perspective in 2D (Simplified)



3D object point  
 $P = (x_c, y_c, z_c)$   
 $= (x_w, y_w, z_w)$

ray

$P' = (x_i, y_i, f)$

$y_i$

$F$

$f$

$x_i$

$y_c$

optical axis

$z_w = z_c$

$Z_c$

$x_c$

camera

$X_c$

$Y_c$

$$\frac{x_i}{f} = \frac{x_c}{z_c}$$

$$x_i = (f/z_c)x_c$$


$$y_i = (f/z_c)y_c$$

$$\frac{y_i}{f} = \frac{y_c}{z_c}$$

Here camera coordinates  
 equal world coordinates.

# 3D from Stereo

● 3D point




disparity: the difference in image location of the same 3D point when projected under perspective to two different cameras.

$$d = x_{\text{left}} - x_{\text{right}}$$

# Depth Perception from Stereo

## Simple Model: Parallel Optic Axes



$$\frac{z}{f} = \frac{x}{x_l}$$

$$\frac{z}{f} = \frac{x-b}{x_r}$$

$$\frac{z}{f} = \frac{y}{y_l} = \frac{y}{y_r}$$

y-axis is perpendicular to the page.



# Resultant Depth Calculation

For stereo cameras with parallel optical axes, focal length  $f$ , baseline  $b$ , corresponding image points  $(x_l, y_l)$  and  $(x_r, y_r)$  with disparity  $d$ :

$$z = f * b / (x_l - x_r) = f * b / d$$


$$x = x_l * z / f \quad \text{or} \quad b + x_r * z / f$$

$$y = y_l * z / f \quad \text{or} \quad y_r * z / f$$

This method of determining depth from disparity is called **triangulation**.



# Finding Correspondences

- If the correspondence is correct, triangulation works **VERY** well.
- But correspondence finding is not perfectly solved. for the general stereo problem.
- For some very specific applications, it can be solved for those specific kind of images, e.g. windshield of a car.





# 3 Main Matching Methods

1. Cross correlation using small windows.



dense

2. Symbolic feature matching, usually using segments/corners.



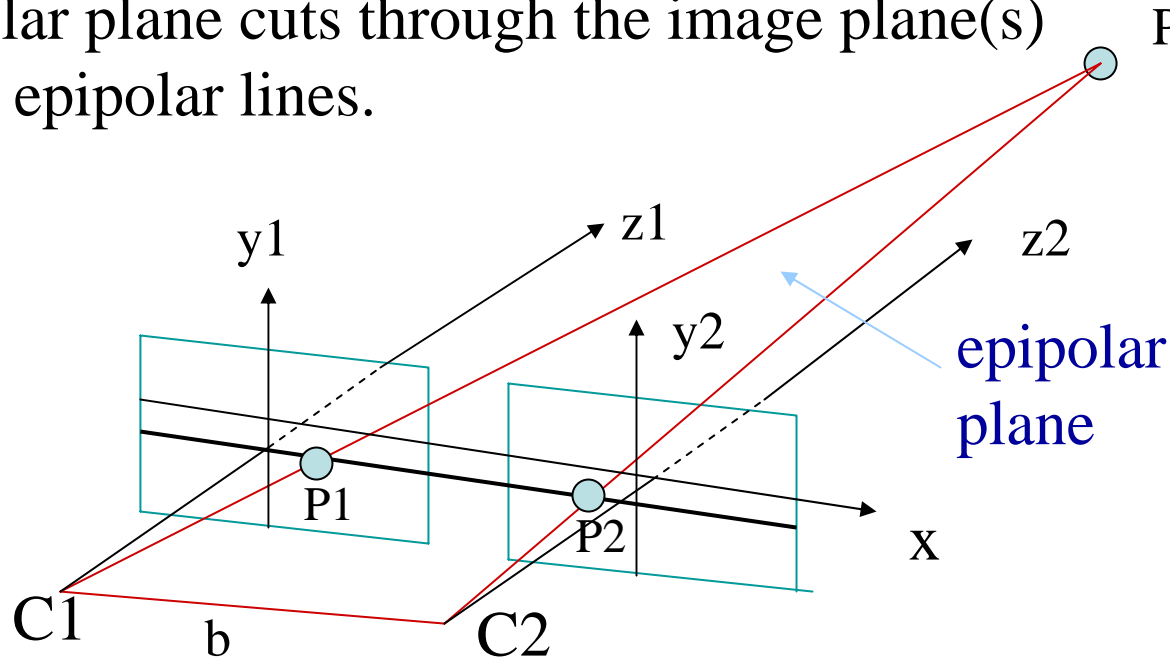
sparse

3. Use the newer interest operators, ie. SIFT.

sparse


# Epipolar Geometry Constraint: 1. Normal Pair of Images

The epipolar plane cuts through the image plane(s) forming 2 epipolar lines.




The match for  $P_1$  (or  $P_2$ ) in the other image, must lie on the same epipolar line.

# Epipolar Geometry: General Case



# Constraints

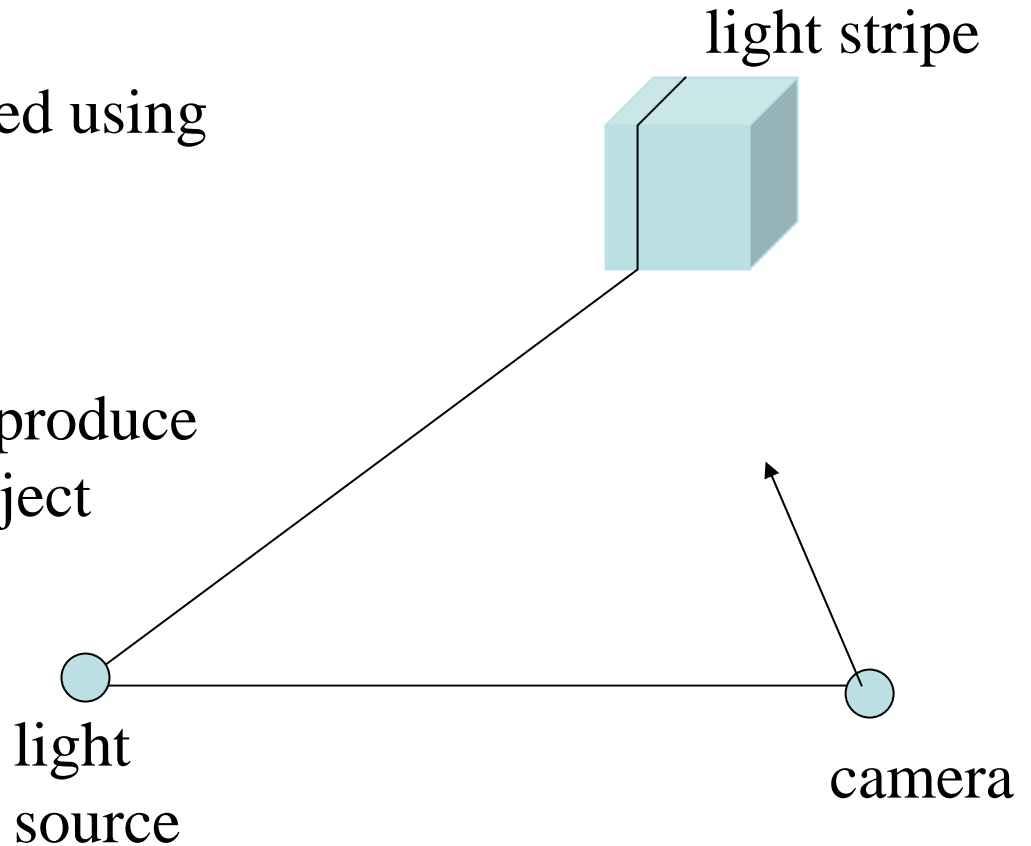
1. Epipolar Constraint:  
Matching points lie on corresponding epipolar lines.
2. Ordering Constraint:  
Usually in the same order across the lines.



# Structured Light

3D data can also be derived using


- a single camera
- a light source that can produce stripe(s) on the 3D object



# Structured Light 3D Computation

3D data can also be derived using



- a single camera
- a light source that can produce stripe(s) on the 3D object



$$\begin{array}{ccc}
 & b & \\
 [x \ y \ z] & = & \frac{\phantom{b}}{f \cot \theta} [x' \ y' \ f] \\
 \text{3D} & & \text{image}
 \end{array}$$





# Depth from Multiple Light Stripes



What are these objects?


# Our (former) System

## 4-camera light-striping stereo

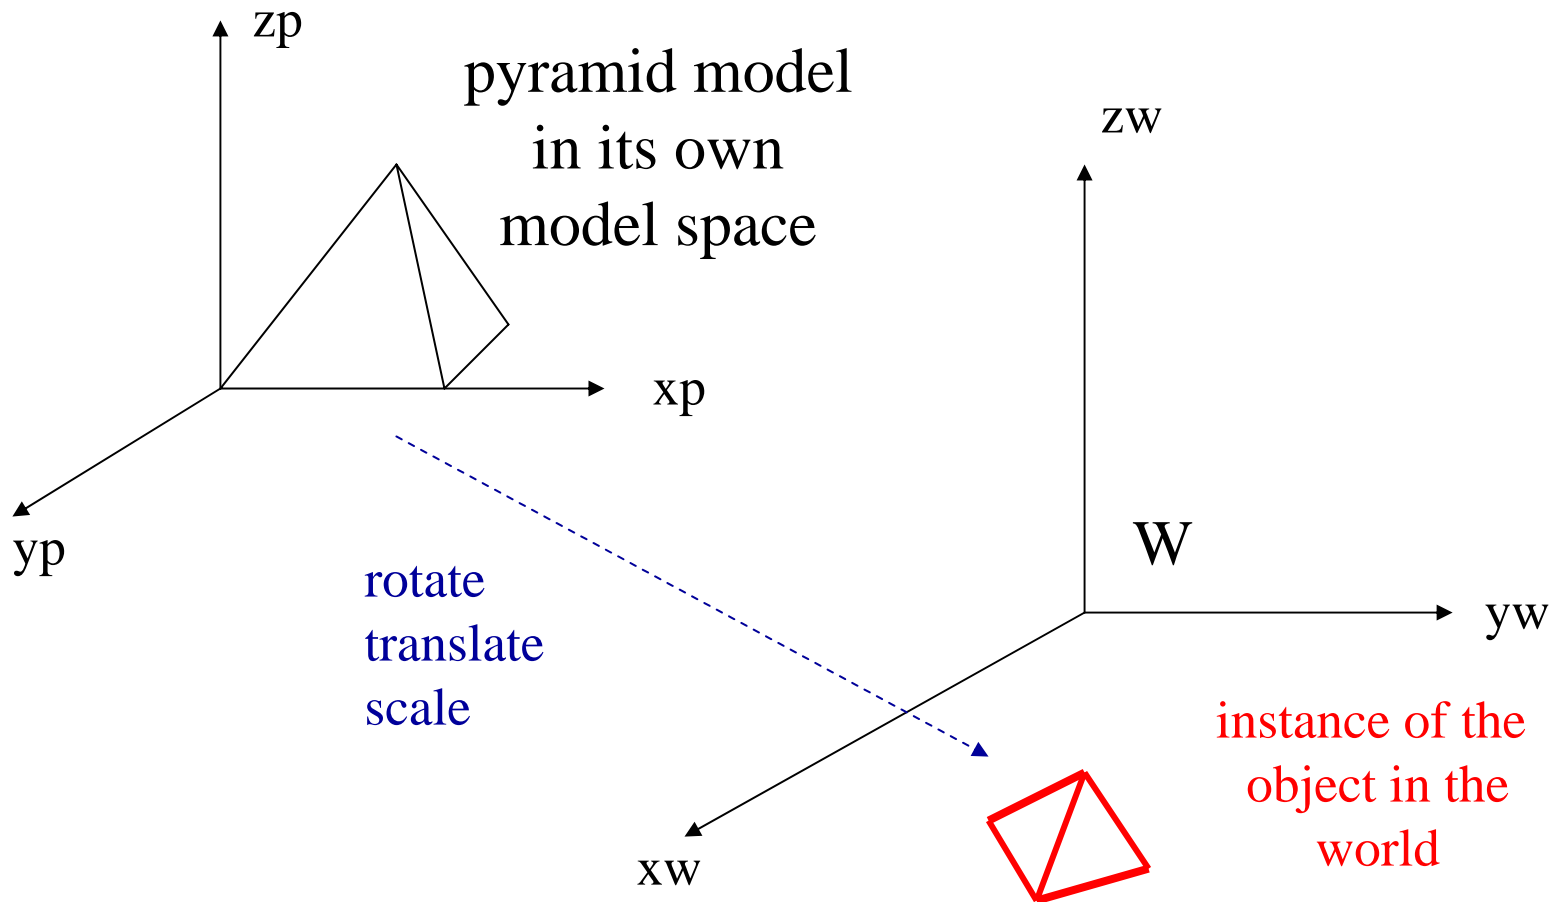


# Camera Model: Recall there are 5 Different Frames of Reference

- Object
- World
- Camera
- Real Image
- Pixel Image



# Rigid Body Transformations in 3D




# Translation and Scaling in 3D

$$\begin{bmatrix} {}^2P_x \\ {}^2P_y \\ {}^2P_z \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & x_0 \\ 0 & 1 & 0 & y_0 \\ 0 & 0 & 1 & z_0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} {}^1P_x \\ {}^1P_y \\ {}^1P_z \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} {}^2P_x \\ {}^2P_y \\ {}^2P_z \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & {}^2P_x \\ s_y & {}^2P_y \\ s_z & {}^2P_z \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & 0 & 0 \\ 0 & s_y & 0 & 0 \\ 0 & 0 & s_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} {}^1P_x \\ {}^1P_y \\ {}^1P_z \\ 1 \end{bmatrix}$$


# Rotation in 3D is about an axis



rotation by angle  $\theta$   
about the x axis

$$\begin{pmatrix} P_{x'} \\ P_{y'} \\ P_{z'} \\ 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta & 0 \\ 0 & \sin \theta & \cos \theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} P_x \\ P_y \\ P_z \\ 1 \end{pmatrix}$$

# Rotation about Arbitrary Axis



One translation and two rotations to line it up with a major axis. Now rotate it about that axis. Then apply the reverse transformations (R2, R1, T) to move it back.

$$\begin{pmatrix} P_x' \\ P_y' \\ P_z' \\ 1 \end{pmatrix} = \begin{pmatrix} r_{11} & r_{12} & r_{13} & t_x \\ r_{21} & r_{22} & r_{23} & t_y \\ r_{31} & r_{32} & r_{33} & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} P_x \\ P_y \\ P_z \\ 1 \end{pmatrix}$$

# The Camera Model

How do we get an image point IP from a world point P?

$$\begin{pmatrix} s \text{ Ipr} \\ s \text{ Ipc} \\ s \end{pmatrix} = \begin{pmatrix} c11 & c12 & c13 & c14 \\ c21 & c22 & c23 & c24 \\ c31 & c32 & c33 & 1 \end{pmatrix} \begin{pmatrix} P_x \\ P_y \\ P_z \\ 1 \end{pmatrix}$$

image  
point

camera matrix **C**

world  
point

What's in C?



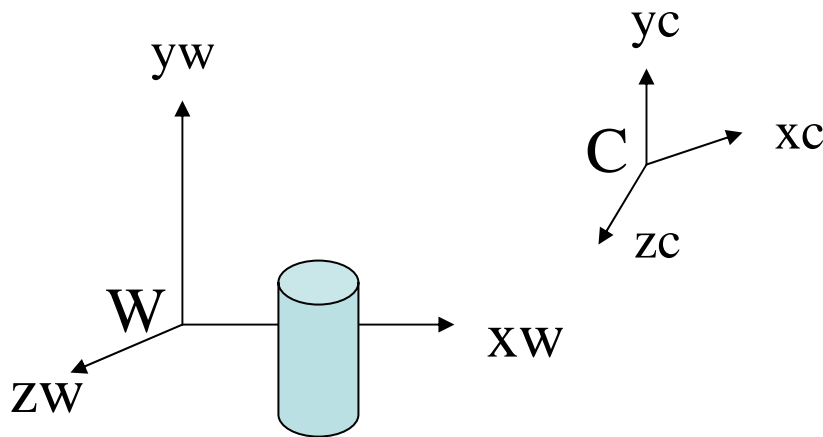
The camera model handles the rigid body transformation from world coordinates to camera coordinates plus the perspective transformation to image coordinates.

$$\begin{array}{l}
 1. \quad \mathbf{CP} = \mathbf{T R WP} \\
 2. \quad \mathbf{IP} = \pi(f) \mathbf{CP}
 \end{array}$$

$$\begin{array}{ccc}
 \begin{pmatrix} s \text{ Ip}_x \\ s \text{ Ip}_y \\ s \end{pmatrix} = & \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/f & 1 \end{pmatrix} & \begin{pmatrix} \text{CP}_x \\ \text{CP}_y \\ \text{CP}_z \\ 1 \end{pmatrix} \\
 \text{image} & \text{perspective} & \text{3D point in} \\
 \text{point} & \text{transformation} & \text{camera} \\
 & & \text{coordinates}
 \end{array}$$

# Camera Calibration

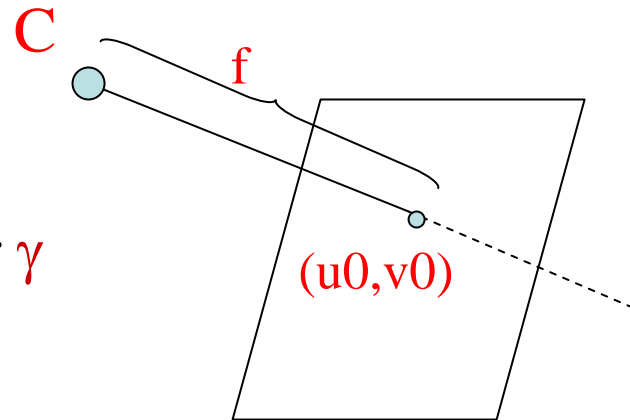
- In order work in 3D, we need to know the parameters of the particular camera setup.
- Solving for the camera parameters is called calibration.



- **intrinsic** parameters are of the camera device
- **extrinsic** parameters are where the camera sits in the world

# Intrinsic Parameters

- principal point  $(u_0, v_0)$
- scale factors  $(dx, dy)$
- aspect ratio distortion factor  $\gamma$
- focal length  $f$
- lens distortion factor  $\kappa$   
(models radial lens distortion)



# Extrinsic Parameters

- translation parameters

$$t = [t_x \ t_y \ t_z]$$

- rotation matrix

$$R = \begin{pmatrix} r_{11} & r_{12} & r_{13} & 0 \\ r_{21} & r_{22} & r_{23} & 0 \\ r_{31} & r_{32} & r_{33} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Are there really  
nine parameters?

# Calibration Object

The idea is to snap images at different depths and get a lot of **2D-3D point correspondences**.




# The Tsai Procedure

- The Tsai procedure was developed by Roger Tsai at IBM Research and is most widely used.
- Several images are taken of the calibration object yielding point correspondences at different distances.
- Tsai's algorithm requires  $n > 5$  correspondences

$$\{(x_i, y_i, z_i), (u_i, v_i) \mid i = 1, \dots, n\}$$

between (real) image points and 3D points.

# Tsai's Geometric Setup




# Tsai's Procedure

- Given  $n$  point correspondences  $((x_i, y_i, z_i), (u_i, v_i))$
- Estimates
  - 9 rotation matrix values
  - 3 translation matrix values
  - focal length
  - lens distortion factor
- By solving several systems of equations



We use them for general stereo.



For a correspondence  $(r1, c1)$  in image 1 to  $(r2, c2)$  in image 2:

1. Both cameras were calibrated. Both camera matrices are then known. From the two camera equations we get

4 linear equations in 3 unknowns.

$$r1 = (b11 - b31*r1)\mathbf{x} + (b12 - b32*r1)\mathbf{y} + (b13 - b33*r1)\mathbf{z}$$

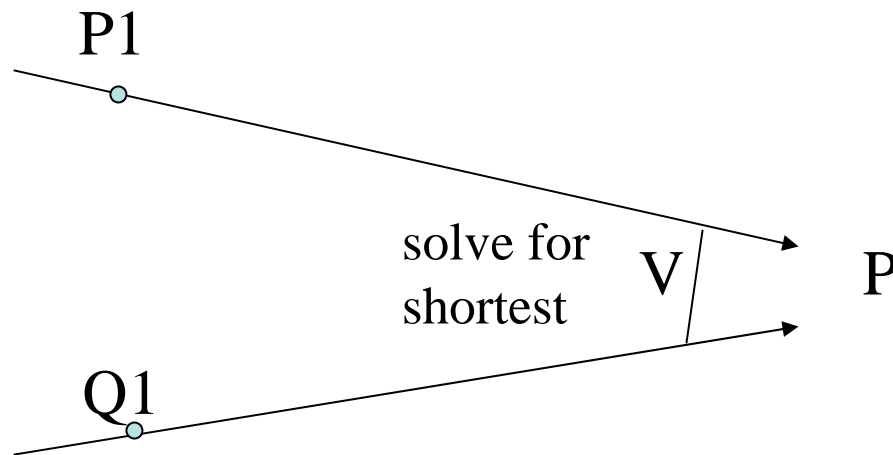
$$c1 = (b21 - b31*c1)\mathbf{x} + (b22 - b32*c1)\mathbf{y} + (b23 - b33*c1)\mathbf{z}$$

$$r2 = (c11 - c31*r2)\mathbf{x} + (c12 - c32*r2)\mathbf{y} + (c13 - c33*r2)\mathbf{z}$$

$$c2 = (c21 - c31*c2)\mathbf{x} + (c22 - c32*c2)\mathbf{y} + (c23 - c33*c2)\mathbf{z}$$

Direct solution uses 3 equations, won't give reliable results.


# Solve by computing the closest approach of the two skew rays.



Instead, we solve for the shortest line segment connecting the two rays and let  $P$  be its midpoint.

If the rays intersected perfectly in 3D, the intersection would be  $P$ .

# Application: Kari Pulli's Reconstruction of 3D Objects from light-stripping stereo.



# Application: Zhenrong Qian's 3D Blood Vessel Reconstruction from Visible Human Data

