Lecture 12

Camera calibration

Administrative

- A3 will be out soon
  - calibration + segmentation due 5/19

Administrative

Recitation this friday

- Geometry
- Calibration

So far: 3D to 2D projections via a camera

3D object

3D to 2D transform
(camera)

2D image

2D image

2D to 2D transform
(image filters)

So far: The pinhole camera

For this course, we focus on the pinhole model.

- Similar to thin lens model in Physics: central rays are not deviated.
- Assumes lens camera in focus.
- Useful approximation but ignores important lens distortions.

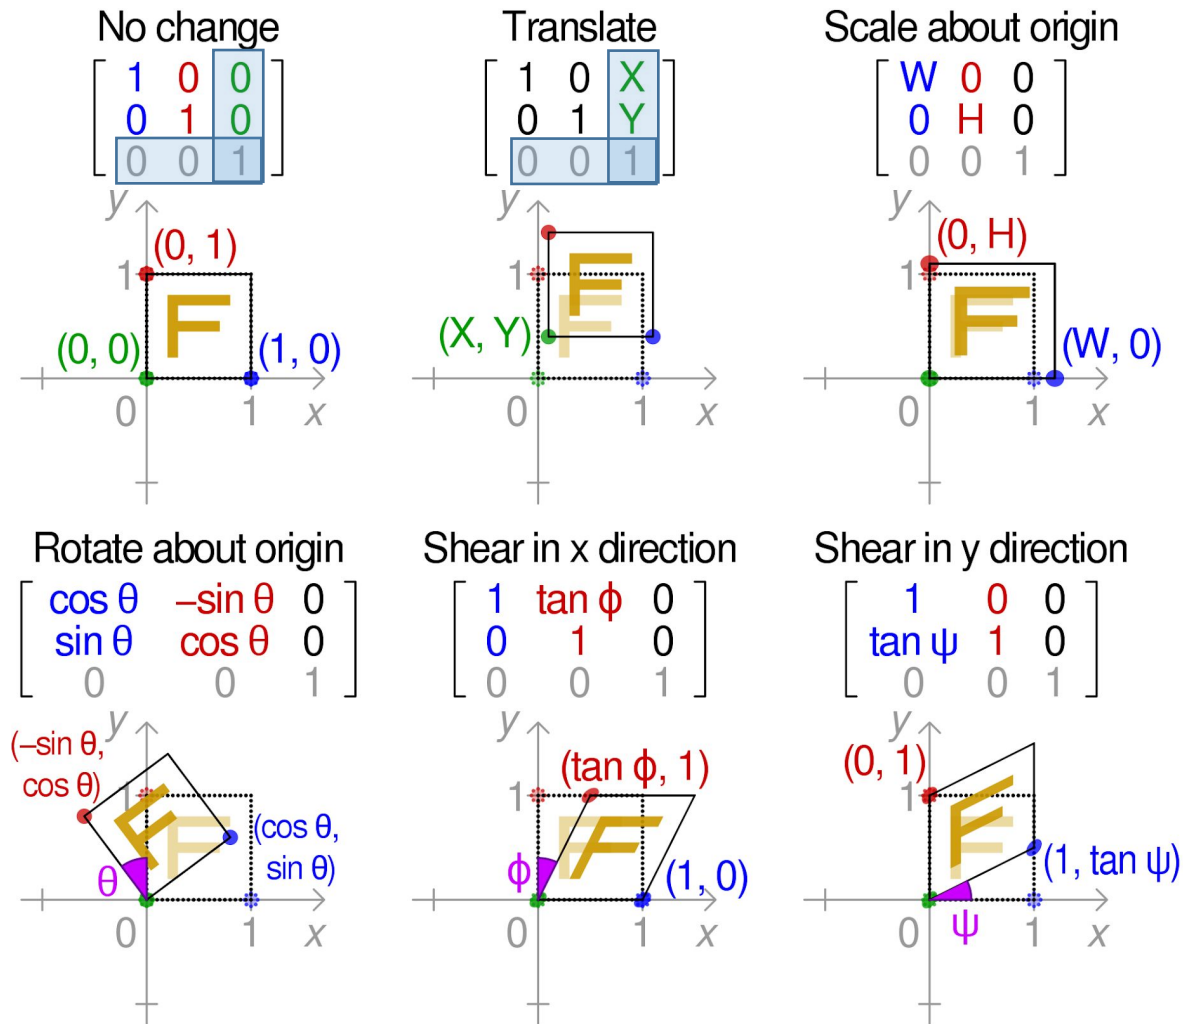# So far: 2D Transformations with homogeneous coordinates

No change

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

(0, 1)
(0, 0) (1, 0)

Translate

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

(X, Y)

Scale about origin

$$\begin{bmatrix} W & 0 & 0 \\ 0 & H & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

(0, H)
(W, 0)

Rotate about origin

$$\begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

(−sin θ, cos θ)
(cos θ, sin θ)
θ

Shear in x direction

$$\begin{bmatrix} 1 & \tan\phi & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

(tan φ, 1)
φ   (1, 0)

Shear in y direction

$$\begin{bmatrix} 1 & 0 & 0 \\ \tan\psi & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

(0, 1)
(1, tan ψ)
ψ

Figure: Wikipedia

So far: 2D Transformations with homogeneous coordinates

2D Projective Space: $\mathcal{P}^2 = \mathcal{R}^3 - (0,0,0)$ (same story in 3D with $\mathcal{P}^3$)

- heterogeneous → homogeneous

$$\begin{bmatrix} x \\ y \end{bmatrix} \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

- homogeneous → heterogeneous

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow \begin{bmatrix} x/w \\ y/w \end{bmatrix}$$

- points differing only by scale are *equivalent*: $(x, y, w) \sim \lambda (x, y, w)$

$$\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y}, \tilde{w}) = \tilde{w}(x, y, 1) = \tilde{w}\bar{\mathbf{x}}$$

So far: General pinhole camera matrix

What is the transformation $\mathbf{x} = \mathbf{P}\mathbf{X}$?

So far: General pinhole camera matrix

$$\mathbf{P} = \mathbf{K}[\mathbf{R}|\mathbf{t}] \qquad \text{where} \qquad \mathbf{t} = -\mathbf{RC}$$

$$\mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_1 & r_2 & r_3 & t_1 \\ r_4 & r_5 & r_6 & t_2 \\ r_7 & r_8 & r_9 & t_3 \end{bmatrix}$$

intrinsic parameters   extrinsic parameters

Today's agenda

- Properties of Perspective transformations
- Introduction to Camera Calibration
- Linear camera calibration method
- Calculating intrinsics and extrinsics
- Depth estimation

Today's agenda

- Properties of Perspective transformations
- Introduction to Camera Calibration
- Linear camera calibration method
- Calculating intrinsics and extrinsics
- Depth estimation

# Similar illusion as last lecture

The Ames room illusion

Actual position of Person A

Apparent position of person A

Actual and apparent position of person B

Apparent shape of room

Viewing peephole

Projective Geometry

Q. Who is taller?

The two blue lines are the same length

Slide credit: J. Hays

Projective Geometry

What is not preserved?

- Length
- Angles

Parallel?

Perpendicular?

Slide credit: J. Hays

Projective Geometry

What is preserved?

- Straight lines are still straight

Slide credit: J. Hays

# Projection of lines

Q. When is parallelism preserved?

Piero della Francesca, *Flagellation of Christ*, 1455-1460

Slide: Lazebnik

Projection of lines

Q. When is parallelism preserved?

When the parallel lines are also parallel to the image plane

Piero della Francesca, *Flagellation of Christ*, 1455-1460

Projection of lines

Patterns on
non-fronto-parallel planes
are distorted by a
homography

If parallel lines are no longer parallel, where do they intersect?

Projection of planes

What about patterns on *fronto-parallel planes*?

$$(x, y, z) \rightarrow \left(f\frac{x}{z}, f\frac{y}{z} \right)$$

- All points on a fronto-parallel plane are at a fixed depth z
- The pattern gets scaled by a factor of f / z, but angles and ratios of lengths/areas are preserved

Vanishing Points & Lines

image plane

vanishing point **v**

camera center

line in the scene

Figure: S Lazebnik

Vanishing Points & Lines

image plane

vanishing point **v**

camera center

line in the scene

Figure: S. Lazebnik

Vanishing Points & Lines

Image - Wikipedia

1-, 2-, 3-perspective

3D objects can have 1, 2, or 3 vanishing points depending on the camera location.

An image with multiple objects can have an arbitrary number of vanishing points.

vanishing points

one-point perspective   two-point perspective   three-point perspective

Vanishing lines for a house

Parallel lines in the world intersect in the image at a vanishing point

Vertical vanishing point (at infinity)

Vanishing line

Vanishing point

Vanishing point

Not all lines that intersect in 2D are parallel in 3D!

Slide from Efros, photo from Criminisi

Horizon: vanishing line of the ground plane (and planes parallel to it)

Parallel planes *in the world* intersect *in the image* at a **vanishing line**

Vanishing lines of planes

- The projection of parallel 3D planes intersect at a vanishing line

- How can we construct the vanishing line of a plane?

Vanishing lines of planes

- The projection of parallel 3D planes intersect at a vanishing line

- How can we construct the vanishing line of a plane?

camera center

plane in the scene

Vanishing lines of planes

Horizon: vanishing line of the ground plane

Q. Is this parachutist **higher** or **lower** than the person taking this picture?

camera center

ground plane

Image source: S. Seitz

Q. Let's try and answer these together with a neighbor

1. All lines that intersect in 2D are parallel in 3D
2. Lines in 3D always intersect each other in 2D at a single point
3. All non-parallel 3D planes in the real world intersect each other
4. 3D lines in the real world always intersect each other
5. All 3D lines intersect each other in the 2D image
6. All parallel lines in 3D meet at the same vanishing point
7. Non-intersecting lines in 3D meet at the same vanishing point
8. If a set of parallel 3D lines are also parallel to a particular plane, their vanishing point will lie on the vanishing line of the plane

1. All lines that intersect in 2D are parallel in 3D

Not true.

2. Lines in 3D always intersect each other in 2D at a single point

Lines in 3D are still lines in 2D.

Any two distinct lines meet in exactly one point (which may lie 'at infinity' if the lines are parallel in the Euclidean sense).

3. All non-parallel 3D planes in the real world intersect each other

True. They intersect at a line.

Parallel planes could potentially intersect at infinity (but don't have to)

The horizon line is where the ground planes intersect

4. 3D lines in the real world always intersect each other

5. All 3D lines intersect each other in the 2D image

6. All parallel lines in 3D meet at the same vanishing point

Q. how would you go about proving this?

6. All parallel lines in 3D meet at the same vanishing point

consider a simple camera with its camera z-axis aligned with the world z-axis:

For line $\ell_i(t) = \big(p_{ix} + t\,d_x,\ p_{iy} + t\,d_y,\ p_{iz} + t\,d_z\big)$,

its projected image in the 2D plane is:

$$\big(x_i(t),\ y_i(t)\big) \;=\; \pi\big(\ell_i(t)\big) \;=\; \left(\frac{p_{ix} + t\,d_x}{p_{iz} + t\,d_z},\ \frac{p_{iy} + t\,d_y}{p_{iz} + t\,d_z} \right),$$

6. All parallel lines in 3D meet at the same vanishing point

consider a simple camera with its camera z-axis aligned with the world z-axis:

For line $\ell_i(t) = \left(p_{ix} + t\, d_x,\ p_{iy} + t\, d_y,\ p_{iz} + t\, d_z \right)$,

its projected image in the 2D plane is:

$$\left(x_i(t),\ y_i(t) \right) \;=\; \pi(\ell_i(t)) \;=\; \left(\frac{p_{ix} + t\, d_x}{p_{iz} + t\, d_z},\ \frac{p_{iy} + t\, d_y}{p_{iz} + t\, d_z} \right),$$

as long as $p_{iz} + t\, d_z \neq 0.$

when dz = 0, the lines are parallel in both 3D and 2D and never intersect

VP₁   VP₂

6. All parallel lines in 3D meet at the same vanishing point

Since the lines are parallel in 3D, they intersection

As $t \to \infty$, the ratio

$$x_i(t) = \frac{p_{ix} + t\,d_x}{p_{iz} + t\,d_z} \quad \text{and} \quad y_i(t) = \frac{p_{iy} + t\,d_y}{p_{iz} + t\,d_z}$$

behaves like

$$x_i(t) \approx \frac{t\,d_x}{t\,d_z} = \frac{d_x}{d_z}, \quad y_i(t) \approx \frac{d_y}{d_z}.$$

In the limit $t \to \infty$, $\big(x_i(t),\, y_i(t)\big)$ approaches the **same** coordinate $\left(\frac{d_x}{d_z},\ \frac{d_y}{d_z} \right)$, regardless of

$\mathbf{p_i}$.

7. Non-intersecting lines in 3D meet at the same vanishing point

8. If a set of parallel 3D lines are also parallel to a particular plane, their vanishing point will lie on the vanishing line of the plane

Same proof as 6.

Properties of Vanishing Points & Lines

- The projections of parallel 3D lines intersect

  at a vanishing point

- The projection of parallel 3D planes intersect

  at a vanishing line

- Vanishing point <-> 3D direction of a line

- Vanishing line <-> 3D orientation of a surface

Vanishing Points are a key perspective tool:

from making realistic drawings, to measuring

3D from 2D, and even for camera calibration!

How do you calculate vanishing points?

Q. What do you need first?

How do you calculate vanishing points?

Lines!

Q. How do you get the lines?

How do you calculate vanishing points?

Lines!

First calculate edges using Canny Edge Detector.

Then use RANSAC or Hough transforms to get the lines.

Q. Once you have the lines, how do you find the vanishing points?

Vertical vanishing point (at infinity)

Vanishing line

Vanishing point

Vanishing point

How do you calculate vanishing points?

Lines!

First calculate edges using Canny Edge Detector.

Then use RANSAC or Hough transforms to get the lines.

Q. Once you have the lines, how do you find the vanishing points?

RANSAC again, but in pixel space to count the intersections!

What can you do with the vanishing points?

You can calculate camera intrinsics.

Let v_1, v_2 be the *homogeneous* coordinates two distinct vanishing points.

Calculating the lines from the vanishing points *(assuming we already know intrinsics K)*

You can calculate camera intrinsics.

Let v_1, v_2 be the *homogeneous* coordinates two distinct vanishing points.

Then the 3D directions of these lines—after undoing the intrinsics—are:

$$\mathbf{d}_1 \propto K^{-1}\mathbf{v}_1, \quad \mathbf{d}_2 \propto K^{-1}\mathbf{v}_2$$

Vertical vanishing point (at infinity)

Vanishing line

Vanishing point

Vanishing point

Calculating the lines from the vanishing points
(assuming we already know intrinsics K)

You can calculate camera intrinsics.

Let v_1, v_2 be the *homogeneous* coordinates two distinct vanishing points.

Then the 3D directions of these lines—after undoing the intrinsics—are:

$$\mathbf{d}_1 \propto K^{-1}\mathbf{v}_1, \quad \mathbf{d}_2 \propto K^{-1}\mathbf{v}_2$$

Because these directions are orthogonal in 3D, we have: $\mathbf{d}_1^\top \mathbf{d}_2 = 0$.

Calculating the lines from the vanishing points
(assuming we already know intrinsics K)

You can calculate camera intrinsics.

Let v_1, v_2 be the *homogeneous* coordinates two distinct vanishing points.

Then the 3D directions of these lines—after undoing the intrinsics—are:

$$\mathbf{d}_1 \propto K^{-1}\mathbf{v}_1, \quad \mathbf{d}_2 \propto K^{-1}\mathbf{v}_2$$

Because these directions are orthogonal in 3D, we have: $\mathbf{d}_1^\top \mathbf{d}_2 = 0.$

Substituting $d_i = K^{-1}v_i$, we get:

$$(K^{-1}\mathbf{v}_1)^\top (K^{-1}\mathbf{v}_2) = \mathbf{v}_1^\top (K^{-1})^\top (K^{-1}) \mathbf{v}_2 = 0$$

Calculating the lines from the vanishing points
(assuming we already know intrinsics K)

You can calculate camera intrinsics.

Let v_1, v_2 be the *homogeneous* coordinates two distinct vanishing points.

Then the 3D directions of these lines—after undoing the intrinsics—are:

$$\mathbf{d}_1 \propto K^{-1} \mathbf{v}_1, \quad \mathbf{d}_2 \propto K^{-1} \mathbf{v}_2$$

Because these directions are orthogonal in 3D, we have: $\mathbf{d}_1^\top \mathbf{d}_2 = 0$.

Substituting $d_i = K^{-1} v_i$, we get:

$$(K^{-1} \mathbf{v}_1)^\top (K^{-1} \mathbf{v}_2) = \mathbf{v}_1^\top (K^{-1})^\top (K^{-1}) \mathbf{v}_2 = 0$$

K has only 3 values:
$$\begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix}$$

Each pair of vanishing points gives us 1 equation. We need at least 3 pairs

Calculating the lines from the vanishing points *(assuming we already know intrinsics K)*

You can calculate camera intrinsics.

Because these directions are orthogonal

We will use this method in your assignments. We will explore another way in lecture to calculate camera intrinsics

lines - after undoing the intrinsics - are:

$$\mathbf{d}_1 \propto K^{-1}\mathbf{v}_1, \quad \mathbf{d}_2 \propto K^{-1}\mathbf{v}_2$$

K has only 3 values:

$$\begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix}$$

Each pair of vanishing points gives us 1 equation. We need at least 3 pairs

Today's agenda

- Properties of Perspective transformations
- **Introduction to Camera Calibration**
- Linear camera calibration method
- Calculating intrinsics and extrinsics
- Depth estimation

Recap: The Pinhole Camera Model

$$\tilde{\mathbf{x}}^{\mathbf{I}} \sim \mathbf{P}\tilde{\mathbf{X}}^{\mathbf{W}} \qquad \mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} [\mathbf{I} \mid \mathbf{0}] \begin{bmatrix} \mathbf{R} & -\mathbf{RC} \\ \mathbf{0} & 1 \end{bmatrix} = \mathbf{K}[\mathbf{R}|\mathbf{t}]$$

intrinsic parameters **K** (3 x 3): correspond to camera internals (image-to-image transformation)

perspective projection (3 x 4): maps 3D to 2D points (camera-to-image transformation)

extrinsic parameters (4 x 4): correspond to camera externals (world-to-camera transformation)

Camera Calibration & Pose Estimation

$$\tilde{\mathbf{x}}^{\mathbf{I}} \sim \mathbf{P}\tilde{\mathbf{X}}^{\mathbf{W}} \qquad \mathbf{P} = \mathbf{K}[\mathbf{R} \,|\mathbf{t}]$$

How can we estimate **P** and its components?

Camera Calibration: estimating *intrinsics* **K**

Pose Estimation: estimating *extrinsics* **[R | t]**

Camera calibration from 3D-2D correspondences

Given n points with ***known*** 3D coordinates X_i and known image projections x_i, estimate the camera parameters $\mathbf{P}$ such that $\tilde{\mathbf{x}}_i^{\text{I}} \sim \mathbf{P}\tilde{\mathbf{X}}_i^{\text{W}}$

Credits: Lazebnik

Camera calibration from 3D-2D correspondences

Given n points with **known** 3D coordinates $\boldsymbol{X}_i$ and known image projections $\boldsymbol{x}_i$, estimate the camera parameters $\mathbf{P}$ such that $\tilde{\mathbf{x}}_i^{I} \sim \mathbf{P}\tilde{\mathbf{X}}_i^{W}$

Known 2D image coords	Known 3D locations
880 214	312.747 309.140 30.086
43 203	305.796 311.649 30.356
270 197	307.694 312.358 30.418
886 347	310.149 307.186 29.298
745 302	311.937 310.105 29.216
943 128	311.202 307.572 30.682
476 590	307.106 306.876 28.660
419 214	309.317 312.490 30.230
317 335	307.435 310.151 29.318
783 521	308.253 306.300 28.881
235 427	306.650 309.301 28.905
665 429	308.069 306.831 29.189
655 362	309.671 308.834 29.029
427 333	308.255 309.955 29.267
412 415	307.546 308.613 28.963
746 351	311.036 309.206 28.913
434 415	307.518 308.175 29.069
525 234	309.950 311.262 29.990
716 308	312.160 310.772 29.080
602 187	311.988 312.709 30.514

Camera calibration from 3D-2D correspondences

Given n points with **known** 3D coordinates $\boldsymbol{X}_i$ and known image projections $\boldsymbol{x}_i$, estimate the camera parameters $\mathbf{P}$ such that $\tilde{\mathbf{x}}_i^{\mathrm{I}} \sim \mathbf{P}\tilde{\mathbf{X}}_i^{\mathrm{W}}$

Many good solutions for accurate 3D position from good fiducial markers: ArUco, AprilTags, …

Fig. 1. Example of augmented reality scene. (a) Input image containing a set of fiducial markers. (b) Markers automatically detected and used for camera pose estimation. (c) Augmented scene without considering user's occlusion. (d) Augmented scene considering occlusion.

Garrido-Jurado, S., Muñoz-Salinas, R., Madrid-Cuevas, F. J., & Marín-Jiménez, M. J. (2014). Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*

https://april.eecs.umich.edu/software/apriltag

Mapping between 3D point and image points

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

What are the knowns and unknowns?

Mapping between 3D point and image points

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

Heterogeneous coordinates

$$x' = \frac{\boldsymbol{p}_1^\top \boldsymbol{X}}{\boldsymbol{p}_3^\top \boldsymbol{X}} \qquad y' = \frac{\boldsymbol{p}_2^\top \boldsymbol{X}}{\boldsymbol{p}_3^\top \boldsymbol{X}}$$

(non-linear relationship between coordinates)

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \text{---} & \boldsymbol{p}_1^\top & \text{---} \\ \text{---} & \boldsymbol{p}_2^\top & \text{---} \\ \text{---} & \boldsymbol{p}_3^\top & \text{---} \end{bmatrix} \begin{bmatrix} | \\ \boldsymbol{X} \\ | \end{bmatrix}$$

How can we make these relations linear?

Today's agenda

- Properties of Perspective transformations
- Introduction to Camera Calibration
- Linear camera calibration method
- Calculating intrinsics and extrinsics
- Depth estimation

How can we make these relations linear?

$$x' = \frac{\boldsymbol{p}_1^\top \boldsymbol{X}}{\boldsymbol{p}_3^\top \boldsymbol{X}} \qquad\qquad y' = \frac{\boldsymbol{p}_2^\top \boldsymbol{X}}{\boldsymbol{p}_3^\top \boldsymbol{X}}$$

Make them linear with algebraic manipulation…

$$\boldsymbol{p}_1^\top \boldsymbol{X} - \boldsymbol{p}_3^\top \boldsymbol{X} x' = 0 \qquad\qquad \boldsymbol{p}_2^\top \boldsymbol{X} - \boldsymbol{p}_3^\top \boldsymbol{X} y' = 0$$

Now we can setup a system of linear equations with multiple corresponding points

Camera Calibration: Linear Method

$$p_i \equiv MX_i$$

Remember (from geometry): this implies MX_i & p_i are proportional/scaled copies of each other

$$p_i = \lambda MX_i, \lambda \neq 0$$

Remember (from homography fitting): this implies their cross product is **0**

$$p_i \times MX_i = 0$$

$$\boldsymbol{p}_1^\top \boldsymbol{X} - \boldsymbol{p}_3^\top \boldsymbol{X} x' = 0 \qquad \boldsymbol{p}_2^\top \boldsymbol{X} - \boldsymbol{p}_3^\top \boldsymbol{X} y' = 0$$

In matrix form …
$$\begin{bmatrix} \boldsymbol{X}^\top & \boldsymbol{0} & -x'\boldsymbol{X}^\top \\ \boldsymbol{0} & \boldsymbol{X}^\top & -y'\boldsymbol{X}^\top \end{bmatrix} \begin{bmatrix} \boldsymbol{p}_1 \\ \boldsymbol{p}_2 \\ \boldsymbol{p}_3 \end{bmatrix} = \boldsymbol{0}$$

Q1. How many equations does each correspondence give us?

Q2. How many correspondences do we need to solve for P?

$$\boldsymbol{p}_1^\top \boldsymbol{X} - \boldsymbol{p}_3^\top \boldsymbol{X} x' = 0 \qquad\qquad \boldsymbol{p}_2^\top \boldsymbol{X} - \boldsymbol{p}_3^\top \boldsymbol{X} y' = 0$$

In matrix form for 1 corresponding point …

$$\begin{bmatrix} \boldsymbol{X}^\top & \boldsymbol{0} & -x'\boldsymbol{X}^\top \\ \boldsymbol{0} & \boldsymbol{X}^\top & -y'\boldsymbol{X}^\top \end{bmatrix} \begin{bmatrix} \boldsymbol{p}_1 \\ \boldsymbol{p}_2 \\ \boldsymbol{p}_3 \end{bmatrix} = \boldsymbol{0}$$

For N points …

$$\begin{bmatrix} \boldsymbol{X}_1^T & 0 & -x_1'\boldsymbol{X}_1^T \\ \boldsymbol{0} & \boldsymbol{X}_1^T & -y_1'\boldsymbol{X}_1^T \\ \vdots & \vdots & \vdots \\ \boldsymbol{X}_N^T & 0 & -x_N'\boldsymbol{X}_N^T \\ \boldsymbol{0} & \boldsymbol{X}_N^T & -y_N'\boldsymbol{X}_N^T \end{bmatrix} \begin{bmatrix} \boldsymbol{p_1} \\ \boldsymbol{p_2} \\ \boldsymbol{p_3} \end{bmatrix} = 0$$

How do we solve this system?

A few things to look out for

- N points should not be co-planar. Otherwise, the rows will not be independent
- Usually any measurements you make in 3D space will be noisy. So you need more than the minimum number of points!
- P has 12 values but we only need to find 11 since everything is scaled

For N points …

$$\begin{bmatrix} \boldsymbol{X}_1^T & \boldsymbol{0} & -x_1' \boldsymbol{X}_1^T \\ \boldsymbol{0} & \boldsymbol{X}_1^T & -y_1' \boldsymbol{X}_1^T \\ \vdots & \vdots & \vdots \\ \boldsymbol{X}_N^T & \boldsymbol{0} & -x_N' \boldsymbol{X}_N^T \\ \boldsymbol{0} & \boldsymbol{X}_N^T & -y_N' \boldsymbol{X}_N^T \end{bmatrix} \begin{bmatrix} \boldsymbol{p}_1 \\ \boldsymbol{p}_2 \\ \boldsymbol{p}_3 \end{bmatrix} = 0$$

How do we solve this system?

Solve for camera matrix via total least squares

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \|\mathbf{A}\boldsymbol{x}\|^2 \text{ subject to } \|\boldsymbol{x}\|^2 = 1$$

$$\mathbf{A} = \begin{bmatrix} \boldsymbol{X}_1^T & \boldsymbol{0} & -x_1' \boldsymbol{X}_1^T \\ \boldsymbol{0} & \boldsymbol{X}_1^T & -y_1' \boldsymbol{X}_1^T \\ \vdots & \vdots & \vdots \\ \boldsymbol{X}_N^T & \boldsymbol{0} & -x_N' \boldsymbol{X}_N^T \\ \boldsymbol{0} & \boldsymbol{X}_N^T & -y_N' \boldsymbol{X}_N^T \end{bmatrix} \qquad \boldsymbol{x} = \begin{bmatrix} \boldsymbol{p}_1 \\ \boldsymbol{p}_2 \\ \boldsymbol{p}_3 \end{bmatrix}$$

Singular Value Decomposition!

Singular Value Decomposition (SVD)

- Represents any matrix **A** as a product of three matrices: $\mathbf{U\Sigma V^T}$
- Python command:
    - [U,**Σ**,V]= numpy.linalg.svd(A)

$$\overset{U}{\begin{bmatrix} -.40 & .916 \\ .916 & .40 \end{bmatrix}} \times \overset{\Sigma}{\begin{bmatrix} 5.39 & 0 \\ 0 & 3.154 \end{bmatrix}} \times \overset{V^T}{\begin{bmatrix} -.05 & .999 \\ .999 & .05 \end{bmatrix}} = \overset{A}{\begin{bmatrix} 3 & -2 \\ 1 & 5 \end{bmatrix}}$$

Singular Value Decomposition (SVD)

- Beyond 2x2 matrices:
  - In general, if **A** is *m* x *n*, then **U** will be *m* x *m*, **Σ** will be *m* x *n*, and **V^T** will be *n* x *n*.
  - (Note the dimensions work out to produce *m* x *n* after multiplication)

$$
\overset{U}{\begin{bmatrix} -.40 & .916 \\ .916 & .40 \end{bmatrix}} \times \overset{\Sigma}{\begin{bmatrix} 5.39 & 0 \\ 0 & 3.154 \end{bmatrix}} \times \overset{V^T}{\begin{bmatrix} -.05 & .999 \\ .999 & .05 \end{bmatrix}} = \overset{A}{\begin{bmatrix} 3 & -2 \\ 1 & 5 \end{bmatrix}}
$$

Singular Value Decomposition (SVD)

- **U** and **V** are always **rotation** matrices.
  - Each column is a unit vector.
- **Σ** is a diagonal matrix
  - The number of nonzero entries = rank of **A**
  - The algorithm always sorts the entries high to low

$$\overset{U}{\begin{bmatrix} -.40 & .916 \\ .916 & .40 \end{bmatrix}} \times \overset{\Sigma}{\begin{bmatrix} 5.39 & 0 \\ 0 & 3.154 \end{bmatrix}} \times \overset{V^T}{\begin{bmatrix} -.05 & .999 \\ .999 & .05 \end{bmatrix}} = \overset{A}{\begin{bmatrix} 3 & -2 \\ 1 & 5 \end{bmatrix}}$$

Solve for camera matrix via total least squares

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \|\mathbf{A}\boldsymbol{x}\|^2 \text{ subject to } \|\boldsymbol{x}\|^2 = 1$$

$$\mathbf{A} = \begin{bmatrix} \boldsymbol{X}_1^T & \mathbf{0} & -x_1'\boldsymbol{X}_1^T \\ \mathbf{0} & \boldsymbol{X}_1^T & -y_1'\boldsymbol{X}_1^T \\ \vdots & \vdots & \vdots \\ \boldsymbol{X}_N^T & \mathbf{0} & -x_N'\boldsymbol{X}_N^T \\ \mathbf{0} & \boldsymbol{X}_N^T & -y_N'\boldsymbol{X}_N^T \end{bmatrix} \qquad \boldsymbol{x} = \begin{bmatrix} \boldsymbol{p}_1 \\ \boldsymbol{p}_2 \\ \boldsymbol{p}_3 \end{bmatrix}$$

$$\mathbf{A} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$$

Solution **x** is the column of **V** corresponding to the smallest singular value of A

Why is it the column of V with the smallest eigenvalue?

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \|\mathbf{A}\boldsymbol{x}\|^2 \text{ subject to } \|\boldsymbol{x}\|^2 = 1$$

Is equivalent to:

$$\mathbf{x} = \arg\min_{\mathbf{x}} \mathbf{x}^T(\mathbf{A}^T\mathbf{A})\mathbf{x} \qquad \text{such that} \qquad \|\mathbf{x}\|^2 = 1.$$

Why is it the column of V with the smallest eigenvalue?

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \|\mathbf{A}\boldsymbol{x}\|^2 \text{ subject to } \|\boldsymbol{x}\|^2 = 1$$

Is equivalent to:

$$\mathbf{x} = \arg\min_{\mathbf{x}} \mathbf{x}^T(\mathbf{A}^T\mathbf{A})\mathbf{x} \qquad \text{such that} \qquad \|\mathbf{x}\|^2 = 1.$$

$$
\begin{aligned}
A^TA &= (U\Sigma V)^T(U\Sigma V) \\
&= (V^T\Sigma U^T)(U\Sigma V) \\
&= V^T\Sigma(U^TU)\Sigma V \\
&= V^T\Sigma I\Sigma V \\
&= V^T\Sigma^2 V
\end{aligned}
$$

Why is it the column of V with the smallest eigenvalue?

$$\hat{x} = \arg\min_{x} \|\mathbf{A}x\|^2 \text{ subject to } \|x\|^2 = 1$$

Is equivalent to:

$$\mathbf{x} = \arg\min_{\mathbf{x}} \mathbf{x}^T (\mathbf{A}^T \mathbf{A})\mathbf{x} \qquad \text{such that} \qquad \|\mathbf{x}\|^2 = 1.$$

$A^T A = (U\Sigma V)^T (U\Sigma V)$
$\quad = (V^T \Sigma U^T)(U\Sigma V)$
$\quad = V^T \Sigma (U^T U)\Sigma V$
$\quad = V^T \Sigma I \Sigma V$
$\quad = V^T \Sigma^2 V$

$$\mathbf{A}^T \mathbf{A} \mathbf{v}_k = \sigma_k^2 \mathbf{v}_k,$$

So, $A^T A v_k$ is proportional to the vector's eigenvalue. So the vector corresponding to the smallest eigenvalue is what we want

Solve for camera matrix via total least squares

$$\hat{x} = \underset{x}{\arg\min} \|\mathbf{A}x\|^2 \text{ subject to } \|x\|^2 = 1$$

$$\mathbf{A} = \begin{bmatrix} X_1^T & 0 & -x_1'X_1^T \\ 0 & X_1^T & -y_1'X_1^T \\ \vdots & \vdots & \vdots \\ X_N^T & 0 & -x_N'X_N^T \\ 0 & X_N^T & -y_N'X_N^T \end{bmatrix} \qquad x = \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix}$$

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$$

Equivalently, solution **x** is the Eigenvector corresponding to the smallest Eigenvalue of **A**

A.2.1 Total least squares

In some problems, e.g., when performing geometric line fitting in 2D images or 3D plane fitting to point cloud data, instead of having measurement error along one particular axis, the measured points have uncertainty in all directions, which is known as the *errors-in-variables* model (Van Huffel and Lemmerling 2002; Matei and Meer 2006). In this case, it makes more sense to minimize a set of homogeneous squared errors of the form

$$E_{\text{TLS}} = \sum_i (\mathbf{a}_i \mathbf{x})^2 = \|\mathbf{A}\mathbf{x}\|^2, \tag{A.35}$$

which is known as *total least squares* (TLS) (Van Huffel and Vandewalle 1991; Björck 1996; Golub and Van Loan 1996; Van Huffel and Lemmerling 2002).

The above error metric has a trivial minimum solution at $\mathbf{x} = 0$ and is, in fact, homogeneous in $\mathbf{x}$. For this reason, we augment this minimization problem with the requirement that $\|\mathbf{x}\|^2 = 1$. which results in the eigenvalue problem

$$\mathbf{x} = \arg \min_{\mathbf{x}} \mathbf{x}^T (\mathbf{A}^T \mathbf{A}) \mathbf{x} \qquad \text{such that} \qquad \|\mathbf{x}\|^2 = 1. \tag{A.36}$$

The value of $\mathbf{x}$ that minimizes this constrained problem is the eigenvector associated with the smallest eigenvalue of $\mathbf{A}^T \mathbf{A}$. This is the same as the last right singular vector of $\mathbf{A}$, because

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \tag{A.37}$$

$$\mathbf{A}^T \mathbf{A} = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^T, \tag{A.38}$$

$$\mathbf{A}^T \mathbf{A}\mathbf{v}_k = \sigma_k^2 \mathbf{v}_k, \tag{A.39}$$

which is minimized by selecting the smallest σ_k value.

We calculated the camera matrix!

Now we have:

$$\mathbf{P} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix}$$

Today's agenda

- Properties of Perspective transformations
- Introduction to Camera Calibration
- Linear camera calibration method
- **Calculating intrinsics and extrinsics**
- Depth estimation

Recap: The Pinhole Camera Model

$$\tilde{\mathbf{x}}^{I} \sim \mathbf{P}\tilde{\mathbf{X}}^{W} \qquad \mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} [\mathbf{I} \mid \mathbf{0}] \begin{bmatrix} \mathbf{R} & -\mathbf{RC} \\ \mathbf{0} & 1 \end{bmatrix} = \mathbf{K}[\mathbf{R}|\mathbf{t}]$$

intrinsic parameters **K** (3 x 3): correspond to camera internals (image-to-image transformation)

perspective projection (3 x 4): maps 3D to 2D points (camera-to-image transformation)

extrinsic parameters (4 x 4): correspond to camera externals (world-to-camera transformation)

Almost there …
$$\mathbf{P} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix}$$

Want to get $\qquad \mathbf{P} = \mathbf{K}\left[\mathbf{R} \mid -\mathbf{RC}\right]$

How can we calculate the intrinsic and extrinsic
parameters from the projection matrix?

Decomposition of the Camera Matrix

$$\mathbf{P} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \sim \mathbf{K}\,[\mathbf{R} \mid -\mathbf{RC}]$$

Q1. Is there a way we can get rid of **C**?

Decomposition of the Camera Matrix

$$\mathbf{P} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \sim \mathbf{K}\,[\mathbf{R} \mid -\mathbf{RC}]$$

$$\overline{\mathbf{P}} = \begin{bmatrix} p_1 & p_2 & p_3 \\ p_5 & p_6 & p_7 \\ p_9 & p_{10} & p_{11} \end{bmatrix} \sim \mathbf{KR}$$

Q2. Is there a way we can get rid of **R**?

Decomposition of the Camera Matrix

$$\mathbf{P} = \begin{bmatrix} p_1 & p_2 & p_3 & | & p_4 \\ p_5 & p_6 & p_7 & | & p_8 \\ p_9 & p_{10} & p_{11} & | & p_{12} \end{bmatrix} \sim \mathbf{K} \left[\mathbf{R} \mid -\mathbf{RC} \right]$$

$$\bar{\mathbf{P}} = \begin{bmatrix} p_1 & p_2 & p_3 \\ p_5 & p_6 & p_7 \\ p_9 & p_{10} & p_{11} \end{bmatrix} \sim \mathbf{KR}$$

$$\bar{\mathbf{P}}\bar{\mathbf{P}}^\top \sim \mathbf{K}\mathbf{R}\mathbf{R}^\top\mathbf{K}^\top$$
$$\bar{\mathbf{P}}\bar{\mathbf{P}}^\top \sim \mathbf{K}\mathbf{K}^\top \qquad (\mathbf{R}\mathbf{R}^\top = \mathbf{I})$$

Q. Do you remember a **property of K** that could help us solve this?

Decomposition of the Camera Matrix

$$\mathbf{P} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \sim \mathbf{K}\,[\mathbf{R} \mid -\mathbf{RC}]$$

$$\bar{\mathbf{P}} = \begin{bmatrix} p_1 & p_2 & p_3 \\ p_5 & p_6 & p_7 \\ p_9 & p_{10} & p_{11} \end{bmatrix} \sim \mathbf{KR}$$

$$\bar{\mathbf{P}}\bar{\mathbf{P}}^\top \sim \mathbf{KRR}^\top\mathbf{K}^\top$$

$$\bar{\mathbf{P}}\bar{\mathbf{P}}^\top \sim \mathbf{KK}^\top \qquad (\mathbf{RR}^\top = \mathbf{I})$$

K is upper triangular and positive definite!
$$\mathbf{K} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix}$$

Decomposition of the Camera Matrix

$$\mathbf{P} = \left[\begin{array}{c|c} \bar{\mathbf{P}} & \begin{matrix} p_4 \\ p_8 \\ p_{12} \end{matrix} \end{array} \right] \sim \mathbf{K}\,[\mathbf{R} \mid -\mathbf{RC}]$$

$\bar{\mathbf{P}}^{\top}\bar{\mathbf{P}} \sim \mathbf{K}^{\top}\mathbf{K}$ with $\mathbf{K}$ upper triangular p.d.

Obtain $\mathbf{K}$ by Cholesky decomposition of $\bar{\mathbf{P}}^{\top}\bar{\mathbf{P}} = \mathbf{LL}^{\top}$

$$\mathbf{K} \sim \mathbf{L}^{\top}$$

Scalar factor: fixed to $1/L_{3,3}$ ($K_{3,3} = 1$)

Once $\mathbf{K}$ is known, we can compute $\mathbf{R} \sim \mathbf{K}^{-1}\bar{\mathbf{P}}$

Everything is calculated up to a scaling factor. So we need to rescale

Decomposition of the Camera Matrix

$$\mathbf{P} = \left[\begin{array}{c|c} \overline{\mathbf{P}} & \begin{array}{c} p_4 \\ p_8 \\ p_{12} \end{array} \end{array} \right] \sim \mathbf{K} \left[\mathbf{R} \mid -\mathbf{RC} \right]$$

$\overline{\mathbf{P}}^\top \overline{\mathbf{P}} \sim \mathbf{K}^\top \mathbf{K}$ with $\mathbf{K}$ upper triangular p.d.

Obtain $\mathbf{K}$ by Cholesky decomposition of $\overline{\mathbf{P}}^\top \overline{\mathbf{P}} = \mathbf{L}\mathbf{L}^\top$

$$\mathbf{K} \sim \mathbf{L}^\top$$

Scalar factor: fixed to $1/L_{3,3}$ ($K_{3,3} = 1$)

Once $\mathbf{K}$ is known, we can compute $\mathbf{R} \sim \mathbf{K}^{-1} \overline{\mathbf{P}}$

$\mathbf{R}$ is a rotation matrix: $|\mathbf{R}| = 1$!

Scaling R

Let $\mathbf{R} = \lambda \mathbf{K}^{-1} \overline{\mathbf{P}}$

$$|\mathbf{R}| = 1 \Rightarrow \lambda = |\mathbf{K}^{-1} \overline{\mathbf{P}}|^{-1/3}$$

If A is an $n \times n$ matrix and λ is a scalar, then

$$\det(\lambda A) = \lambda^n \det(A).$$

Calculating C

$$\mathbf{P} = \begin{bmatrix} & \bar{\mathbf{P}} & & \begin{matrix} p_4 \\ p_8 \\ p_{12} \end{matrix} \end{bmatrix} \sim \mathbf{K}\,[\mathbf{R} \mid -\mathbf{RC}]$$

Finally, easy to know the camera center: $\mathbf{C} = -\boldsymbol{\lambda}^{-1}\,\mathbf{R}^{\top}\mathbf{K}^{-1}[p_4\ \ p_8\ \ p_{12}]^{\top}$

Linear Camera Calibration

- Advantages:
  - Simple to formulate
  - Analytical solution
- Disadvantages:
  - Doesn't model radial distortion (non-linear!)
  - Hard to impose constraints (e.g., known f)
  - Doesn't minimize the correct error function: the reprojection error in 2D is what we truly care about!
- Hence why **non-linear** methods are preferred in practice.
- They can reuse the linear method we just saw!

(out of scope for class) Non-Linear Camera Calibration

- Write down objective function in terms of intrinsic and extrinsic parameters, as sum of squared distances between **measured** 2D points x_i and **estimated** projections of corresponding 3D points:

$$\sum_i \|\text{proj}(\mathbf{K}[\mathbf{R} \mid \mathbf{t}]\mathbf{X_i}; \boldsymbol{\kappa}) - x_i\|_2^2$$

- Can include radial distortion (cf. Szeliski 2.1.5 & 11.1.4) or other parameters $\boldsymbol{\kappa}$ in the projection model (non-linear in the parameters!)

- Can include constraints such as known focal length, orthogonality, visibility of points, or even known $\mathbf{K}$ ("extrinsics calibration")

- Minimize error using standard non-linear optimization techniques (traditionally Levenberg-Marquardt, cf. Szeliski A.3, 8.1.3, 11.1.4)

- Iterative non-linear optimization is sensitive to initialization: use the output of the linear method we just saw!

Today's agenda

- Properties of Perspective transformations
- Introduction to Camera Calibration
- Linear camera calibration method
- Calculating intrinsics and extrinsics
- Depth estimation

# Estimating depth

b is the translation from camera at location L and camera at R

We want to estimate z

We know (u_L, v_L) and (u_R, v_R)

left camera

$P_O(x, y, z)$

$(0,0,0)$

BASELINE

$P_L(u_L, v_L)$

$(b,0,0)$

right camera

$P_R(u_R, v_R)$

Estimating depth

We can estimate
the intrinsics of
the camera from
the previous
sections

So, we know:

focal length: f_x, f_y

translation: p_x, p_y

$$K = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix}$$


$P_o(x, y, z)$

left camera

$(0,0,0)$

$P_L(u_L, v_L)$

BASELINE

$(b,0,0)$

right camera

$P_R(u_R, v_R)$

Estimating depth

$$u_L = f\frac{x}{z} + p_x$$

$$v_L = f\frac{y}{z} + p_y$$

P₀(x, y, z)

left camera

(0,0,0)

BASELINE

P_L(u_L, v_L)

(b,0,0)

right camera

P_R(u_R, v_R)

$$K = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix}$$

$$u_R = f\frac{x-b}{z} + p_x$$

$$v_R = f\frac{y}{z} + p_y$$

Estimating depth

$$u_L = f\frac{x}{z} + p_x$$

$$\boxed{z = \frac{fb}{u_L - u_R}}$$

$P_o(x, y, z)$

left camera

$(0,0,0)$

BASELINE

$P_L(u_L, v_L)$

$(b,0,0)$

right camera

$P_R(u_R, v_R)$

$$u_R = f\frac{x - b}{z} + p_x$$

Estimating depth

$$z = \frac{fb}{u_L - u_R}$$

Estimating pose [R|t] (extrinsics)

We have already done it together with estimating K!

What if we know K already?

Estimating [R|t] only = pose estimation, a.k.a. extrinsincs calibration (and previous linear method is called the "Direct Linear Transform")

Estimating pose [R|t] (extrinsics)

Other linear algorithm: PnP (Perspective-n-Point)

Commonly used solution available in standard libraries like OpenCV

Minimal form: P3P (3 noise-free non-colinear correspondences)

Main idea: same angle between rays of 2 2D points and 2 3D points

In practice: use $n \geq 4$ correspondences + RANSAC

More details:

Quan, Long; Lan, Zhong-Dan (1999). "Linear N-Point Camera Pose Determination" (PDF). *IEEE TPAMI*.
Lepetit, V.; Moreno-Noguer, M.; Fua, P. (2009). "EPnP: An Accurate O(n) Solution to the PnP Problem". *IJCV*

Figure: R. Szeliski

Estimating pose [R|t] (extrinsics)

Non-linear method:

- Minimize reprojection error as function of [R|t]

- More accurate and flexible (e.g., using constraints)

- Can be robustified and easy to implement via transformation decomposition and backpropagation (yes, like in Deep Learning!)

Figure via Szeliski

cf. Szeliski 11.2.2 for more details

Today's agenda

- Properties of Perspective transformations
- Introduction to Camera Calibration
- Linear camera calibration method
- Calculating intrinsics and extrinsics
- Depth estimation

Next lecture

Recognition