

# RANSAC Recap

# RANSAC - Setup

- Given

- A dataset  $D = \{p_1, p_2, \dots, p_N\}$ 
  - Example 1: Line fitting:  $\{(x_1, y_1), \dots, (x_n, y_n)\}$
  - Example 2: Homography fitting:  $\{(\vec{Q}_1, \vec{q}_1), (\vec{Q}_2, \vec{q}_2), \dots, (\vec{Q}_N, \vec{q}_N)\}$
- A set of parameters  $\theta$  that need to be fitted
  - Line fitting:  $\theta = (m, b)$
  - Homography estimation  $\theta = H, \|h\| = 1$
- A cost function  $C(p, \theta)$ 
  - Line fitting:  $C((x, y), (m, b)) = \|y - (mx + b)\|^2$
  - Homography estimation  $C((\vec{Q}, \vec{q}), H) = E(H)$  (Reprojection error)
- A minimum number needed  $k$ 
  - Line fitting: 2
  - Homography estimation: 4

# RANSAC - Setup

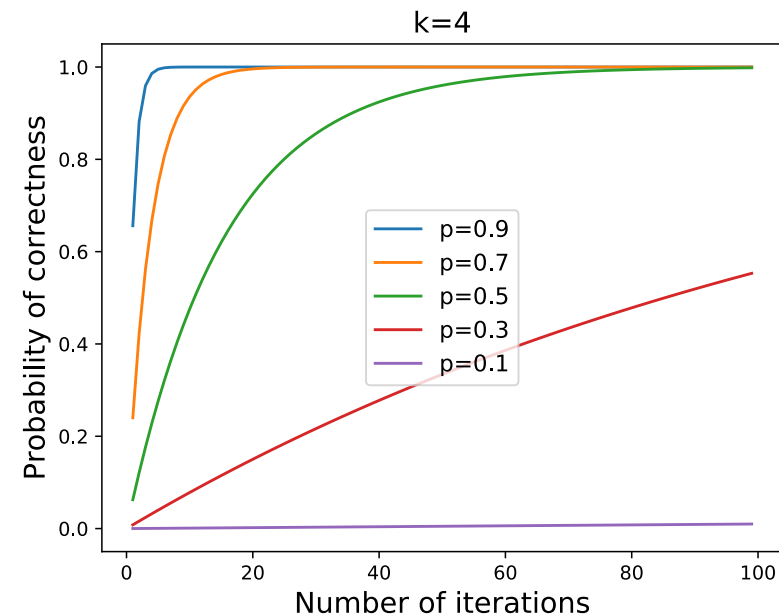

- Given
  - A dataset  $D = \{p_1, p_2, \dots, p_N\}$
  - A set of parameters  $\theta$  that need to be fitted
  - A cost function  $C(p, \theta)$
  - $k$
  - $\theta^* = \min_{\theta} \sum_i C(p_i, \theta)$ ?
  - Problem: outliers

# RANSAC - Algorithm

- Given:  $D = \{p_1, p_2, \dots, p_N\}$ ,  $C(\theta, p)$ ,  $k$
- $\theta_{best} \leftarrow \phi$ ,  $D_{inlier} \leftarrow \phi$
- For  $i = 1, \dots, S$ 
  - Sample  $k$  points
  - Minimize  $C$  for these  $k$  points to get  $\theta_{hyp}$
  - Compute the set of inliers:  $D_{hyp} = \{p \in D : C(\theta_{hyp}, p) < \delta\}$
  - If size of  $D_{hyp}$  is more than size of  $D_{inlier}$ 
    - $\theta_{best} \leftarrow \theta_{hyp}$
    - $D_{inlier} \leftarrow D_{hyp}$
- Minimize  $\theta$  over  $D_{inlier}$

# RANSAC: how many iterations do we need?


- $p$  = inlier fraction
- $k$  = minimum number of data points
- $S$  = iter
- $P = (1 - (1 - p^k)^S)$



# Binocular Stereo

# Binocular stereo

- General case: cameras can be arbitrary locations and orientations
- If we know where cameras are, we can shoot rays from corresponding pixels and intersect



# Binocular stereo : Triangulation

- Suppose we have two cameras
  - Calibrated: parameters known
- And a pair of corresponding pixels
- Find 3D location of point!






# Triangulation

- Suppose we have two cameras
  - Calibrated: parameters known
- And a pair of corresponding pixels
- Find 3D location of point!

$P^{(1)}$



$P^{(2)}$



# Triangulation

$$\begin{array}{c} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} \\ \\ \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} \end{array} \leftarrow \begin{array}{c} \vec{\mathbf{x}}_{img}^{(1)} \\ \\ \vec{\mathbf{x}}_{img}^{(2)} \end{array} \equiv \begin{array}{c} P^{(1)} \vec{\mathbf{x}}_w \\ \\ P^{(2)} \vec{\mathbf{x}}_w \end{array} \rightarrow \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

The diagram illustrates the triangulation process. On the left, two image points are shown as column vectors:  $\begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$  and  $\begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix}$ . Arrows point from these image points to a central set of equations. The first equation is  $\vec{\mathbf{x}}_{img}^{(1)} \equiv P^{(1)} \vec{\mathbf{x}}_w$ , and the second is  $\vec{\mathbf{x}}_{img}^{(2)} \equiv P^{(2)} \vec{\mathbf{x}}_w$ . From the right-hand side of these equations, two arrows point to a single column vector on the right:  $\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$ . This vector represents the world point  $\vec{\mathbf{x}}_w$  that is common to both camera observations.

# Triangulation

$$\vec{\mathbf{X}}_{img}^{(1)} \equiv P^{(1)} \vec{\mathbf{X}}_w$$

$$\lambda x_1 = P_{11}^{(1)} X + P_{12}^{(1)} Y + P_{13}^{(1)} Z + P_{14}^{(1)}$$

$$\lambda y_1 = P_{21}^{(1)} X + P_{22}^{(1)} Y + P_{23}^{(1)} Z + P_{24}^{(1)}$$

$$\lambda = P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}$$

$$(P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}) x_1 = P_{11}^{(1)} X + P_{12}^{(1)} Y + P_{13}^{(1)} Z + P_{14}^{(1)}$$
$$X(P_{31}^{(1)} x_1 - P_{11}^{(1)}) + Y(P_{32}^{(1)} x_1 - P_{12}^{(1)}) + Z(P_{33}^{(1)} x_1 - P_{13}^{(1)}) + (P_{34}^{(1)} x_1 - P_{14}^{(1)}) = 0$$

# Triangulation

$$\vec{\mathbf{x}}_{img}^{(1)} \equiv P^{(1)} \vec{\mathbf{x}}_w$$

$$X(P_{31}^{(1)}x_1 - P_{11}^{(1)}) + Y(P_{32}^{(1)}x_1 - P_{12}^{(1)}) + Z(P_{33}^{(1)}x_1 - P_{13}^{(1)}) + (P_{34}^{(1)}x_1 - P_{14}^{(1)}) = 0$$

$$X(P_{31}^{(1)}y_1 - P_{21}^{(1)}) + Y(P_{32}^{(1)}y_1 - P_{22}^{(1)}) + Z(P_{33}^{(1)}y_1 - P_{23}^{(1)}) + (P_{34}^{(1)}y_1 - P_{24}^{(1)}) = 0$$

- 1 image gives 2 equations
- Need 2 images!
- Solve linear equations to get 3D point location

# Linear vs non-linear optimization

$$\lambda x_1 = P_{11}^{(1)} X + P_{12}^{(1)} Y + P_{13}^{(1)} Z + P_{14}^{(1)}$$

$$\lambda y_1 = P_{21}^{(1)} X + P_{22}^{(1)} Y + P_{23}^{(1)} Z + P_{24}^{(1)}$$

$$\lambda = P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}$$

$$x_1 = \frac{P_{11}^{(1)} X + P_{12}^{(1)} Y + P_{13}^{(1)} Z + P_{14}^{(1)}}{P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}}$$

$$y_1 = \frac{P_{21}^{(1)} X + P_{22}^{(1)} Y + P_{23}^{(1)} Z + P_{24}^{(1)}}{P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}}$$

# Linear vs non-linear optimization

$$x_1 = \frac{P_{11}^{(1)} X + P_{12}^{(1)} Y + P_{13}^{(1)} Z + P_{14}^{(1)}}{P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}}$$

$$y_1 = \frac{P_{21}^{(1)} X + P_{22}^{(1)} Y + P_{23}^{(1)} Z + P_{24}^{(1)}}{P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}}$$

$$\left(x_1 - \frac{P_{11}^{(1)} X + P_{12}^{(1)} Y + P_{13}^{(1)} Z + P_{14}^{(1)}}{P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}}\right)^2 + \left(y_1 - \frac{P_{21}^{(1)} X + P_{22}^{(1)} Y + P_{23}^{(1)} Z + P_{24}^{(1)}}{P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}}\right)^2$$

Reprojection error

# Linear vs non-linear optimization


$$\left(x_1 - \frac{P_{11}^{(1)} X + P_{12}^{(1)} Y + P_{13}^{(1)} Z + P_{14}^{(1)}}{P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}}\right)^2 + \left(y_1 - \frac{P_{21}^{(1)} X + P_{22}^{(1)} Y + P_{23}^{(1)} Z + P_{24}^{(1)}}{P_{31}^{(1)} X + P_{32}^{(1)} Y + P_{33}^{(1)} Z + P_{34}^{(1)}}\right)^2$$

Reprojection error

- Reprojection error is the squared error between the true image coordinates of a point and the projected coordinates of hypothesized 3D point
- Actual error we care about
- Minimize total sum of reprojection error across all images
- *Non-linear optimization*

# Binocular stereo


- General case: cameras can be arbitrary locations and orientations






# Binocular stereo

- Special case: cameras are parallel to each other and translated along X axis




# Stereo with *rectified* cameras

- Special case: cameras are parallel to each other and translated along X axis



Stereo head



Kinect / depth cameras



# Stereo with “rectified cameras”



# Perspective projection in rectified cameras

- Without loss of generality, assume origin is at pinhole of 1<sup>st</sup> camera

$$\vec{\mathbf{x}}_{img}^{(1)} \equiv \begin{bmatrix} I & \mathbf{0} \end{bmatrix} \vec{\mathbf{x}}_w$$

$$\vec{\mathbf{x}}_{img}^{(2)} \equiv \begin{bmatrix} I & \mathbf{t} \end{bmatrix} \vec{\mathbf{x}}_w$$

$$\mathbf{t} = \begin{bmatrix} t_x \\ 0 \\ 0 \end{bmatrix}$$

# Perspective projection in rectified cameras

- Without loss of generality, assume origin is at pinhole of 1<sup>st</sup> camera

$$\vec{\mathbf{x}}_{img}^{(1)} \equiv \begin{bmatrix} I & \mathbf{0} \end{bmatrix} \vec{\mathbf{x}}_w$$

$$\vec{\mathbf{x}}_{img}^{(2)} \equiv \begin{bmatrix} I & \mathbf{t} \end{bmatrix} \vec{\mathbf{x}}_w$$

$$\mathbf{t} = \begin{bmatrix} t_x \\ 0 \\ 0 \end{bmatrix} \quad \vec{\mathbf{x}}_w = \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{x}_w \\ 1 \end{bmatrix}$$

# Perspective projection in rectified cameras

- Without loss of generality, assume origin is at pinhole of 1<sup>st</sup> camera

$$\vec{\mathbf{x}}_{img}^{(1)} \equiv [I \quad \mathbf{0}] \begin{bmatrix} \mathbf{x}_w \\ 1 \end{bmatrix} = \mathbf{x}_w = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

$$\vec{\mathbf{x}}_{img}^{(2)} \equiv [I \quad \mathbf{t}] \begin{bmatrix} \mathbf{x}_w \\ 1 \end{bmatrix} = \mathbf{x}_w + \mathbf{t} = \begin{bmatrix} X + t_x \\ Y \\ Z \end{bmatrix}$$

# Perspective projection in rectified cameras

- Without loss of generality, assume origin is at pinhole of 1<sup>st</sup> camera

$$\vec{\mathbf{x}}_{img}^{(1)} \equiv \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

$$\vec{\mathbf{x}}_{img}^{(2)} \equiv \begin{bmatrix} X + t_x \\ Y \\ Z \end{bmatrix}$$



# Perspective projection in rectified cameras

- Without loss of generality, assume origin is at pinhole of 1<sup>st</sup> camera

$$\begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} \equiv \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

$$\begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} \equiv \begin{bmatrix} X + t_x \\ Y \\ Z \end{bmatrix}$$

# Perspective projection in rectified cameras

- Without loss of generality, assume origin is at pinhole of 1<sup>st</sup> camera

$$\begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} \equiv \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

$$\begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} \equiv \begin{bmatrix} X + t_x \\ Y \\ Z \end{bmatrix}$$

# Perspective projection in rectified cameras

- Without loss of generality, assume origin is at pinhole of 1<sup>st</sup> camera

$$\begin{bmatrix} \lambda x_1 \\ \lambda y_1 \\ \lambda \end{bmatrix} = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

$$\begin{bmatrix} \lambda x_2 \\ \lambda y_2 \\ \lambda \end{bmatrix} = \begin{bmatrix} X + t_x \\ Y \\ Z \end{bmatrix}$$

# Perspective projection in rectified cameras

- Without loss of generality, assume origin is at pinhole of 1<sup>st</sup> camera

X coordinate differs by  $t_x/Z$

$$x_1 = \frac{X}{Z} \qquad x_2 = \frac{X + t_x}{Z}$$

$$y_1 = \frac{Y}{Z} \qquad y_2 = \frac{Y}{Z}$$

Y coordinate is the same!

# Perspective projection in rectified cameras

- X coordinate differs by  $t_x/Z$
- That is, difference in X coordinate is *inversely proportional to depth*
- Difference in X coordinate is called *disparity*
- Translation between cameras ( $t_x$ ) is called *baseline*
  
- *disparity = baseline / depth*

# The disparity image


- For pixel  $(x,y)$  in one image, only need to know disparity to get correspondence
- Create an image with color at  $(x,y) = \text{disparity}$



right image





left image



disparity

# Perspective projection in rectified cameras



- For rectified cameras, correspondence problem is easier
- Only requires searching along a particular *row*.

# NCC - Normalized Cross Correlation


- Lighting and color change pixel intensities
- Example: increase brightness / contrast
- $I' = \alpha I + \beta$
- Subtract patch mean: invariance to  $\beta$
- Divide by norm of vector: invariance to  $\alpha$
- $x' = x - \langle x \rangle$
- $x'' = \frac{x'}{\|x'\|}$
- *similarity* =  $x'' \cdot y''$



Why not SIFT?



# Cross-correlation of neighborhood




regions A, B, write as vectors  $\mathbf{a}$ ,  $\mathbf{b}$

translate so that mean is zero

$$\mathbf{a} \rightarrow \mathbf{a} - \langle \mathbf{a} \rangle, \quad \mathbf{b} \rightarrow \mathbf{b} - \langle \mathbf{b} \rangle$$


$$\text{cross correlation} = \frac{\mathbf{a} \cdot \mathbf{b}}{|\mathbf{a}| |\mathbf{b}|}$$

Invariant to  $I \rightarrow \alpha I + \beta$




left image band

right image band



cross  
correlation




target region



left image band



right image band



cross  
correlation

# The NCC cost volume

- Consider  $M \times N$  image
- Suppose there are  $D$  possible disparities.
- For every pixel,  $D$  possible scores
- Can be written as an  $M \times N \times D$  array
- To get disparity, take max along 3<sup>rd</sup> axis

# Computing the NCC volume

1. For every pixel  $(x, y)$ 
  1. For every disparity  $d$ 
    1. Get normalized patch from image 1 at  $(x, y)$
    2. Get normalized patch from image 2 at  $(x + d, y)$
    3. Compute NCC

# Computing the NCC volume


## 1. For every disparity $d$

### 1. For every pixel $(x, y)$

1. Get normalized patch from image 1 at  $(x, y)$
2. Get normalized patch from image 2 at  $(x + d, y)$
3. Compute NCC

Assume all pixels lie at same disparity  $d$  (i.e., lie on same plane) and compute cost for each

**Plane sweep stereo**



NCC volume



Disparity