


Transformations and Fitting

EECS 442 – David Fouhey

Fall 2019, University of Michigan

http://web.eecs.umich.edu/~fouhey/teaching/EECS442_F19/


Last Class



1. How do we find distinctive / easy to locate features? (*Harris/Laplacian of Gaussian*)
2. How do we describe the regions around them? (*Normalize window, use histogram of gradient orientations*)

Earlier I promised


Solving for a Transformation



3: Solve for transformation T (e.g. such that $\mathbf{p1} \equiv T \mathbf{p2}$) that fits the matches well

Before Anything Else, Remember

You, with your
gigantic brain, see:



The computer
sees:

097	097	097	097	097	097	097	097	097	096	097	097	096	096	096
100	100	100	100	100	100	101	101	102	101	100	100	100	100	099
105	105	105	105	105	105	105	103	102	102	101	103	104	104	105
109	109	109	109	109	110	107	118	145	132	120	112	106	103	
113	113	113	112	112	113	110	129	160	160	164	162	157	151	
118	117	118	123	119	118	112	125	142	134	135	139	139	175	
123	121	125	162	166	157	149	153	160	151	150	146	137	168	
127	127	125	168	147	117	139	135	126	147	147	149	156	160	
133	130	150	179	145	132	160	134	150	150	111	145	126	121	
138	134	179	185	141	090	166	117	120	153	111	153	114	126	
144	151	188	178	159	154	172	147	159	170	147	185	105	122	
152	157	184	183	142	127	141	133	137	141	131	147	144	147	
130	147	185	180	139	131	154	121	140	147	107	147	120	128	
035	102	194	175	149	140	179	128	146	168	096	163	101	125	

You should expect **noise** (not at quite the right pixel) and **outliers** (random matches)

Today

- How do we fit **models** (i.e., a parameteric representation of data that's smaller than the data) to data?
- How do we handle:
 - **Noise** – least squares / total least squares
 - **Outliers** – RANSAC (random sample consensus)
 - **Multiple models** – Hough Transform (can also make RANSAC handle this with some effort)

Working Example: Lines

- We'll handle lines as our **models** today since you should be familiar with them
- Next class will cover more complex models. I promise we'll eventually stitch images together
- You can apply today's techniques on next class's models

Model Fitting

Need three ingredients

Data: what data are we trying to explain with a model?

Model: what's the compressed, parametric form of the data?

Objective function: given a prediction, how do we evaluate how correct it is?


Example: Least-Squares

Fitting a line to data

Data: $(x_1, y_1), (x_2, y_2), \dots, (x_k, y_k)$

Model: $(m, b) y_i = mx_i + b$
Or $(\mathbf{w}) y_i = \mathbf{w}^T \mathbf{x}_i$

Objective function:
 $(y_i - \mathbf{w}^T \mathbf{x}_i)^2$



Least-Squares Setup

$$\sum_{i=1}^k (y_i - \mathbf{w}^T \mathbf{x}_i)^2 \quad \rightarrow \quad \|\mathbf{Y} - \mathbf{X}\mathbf{w}\|_2^2$$

$$\mathbf{Y} = \begin{bmatrix} y_1 \\ \vdots \\ y_k \end{bmatrix} \quad \mathbf{X} = \begin{bmatrix} x_1 & 1 \\ \vdots & 1 \\ x_k & 1 \end{bmatrix} \quad \mathbf{w} = \begin{bmatrix} m \\ b \end{bmatrix}$$

Note: I'm writing the most general form here since we'll do it in general and you can make it specific if you'd like.

Solving Least-Squares

$$\|Y - Xw\|_2^2$$

$$\frac{\partial}{\partial w} \|Y - Xw\|_2^2 = 2X^T Xw - 2X^T Y$$

Recall: derivative is 0 at a maximum / minimum. Same is true about gradients.

$$0 = 2X^T Xw - 2X^T Y$$

$$X^T Xw = X^T Y$$

$$w = (X^T X)^{-1} X^T Y$$

Aside: **0** is a vector of 0s. **1** is a vector of 1s.

Derivation for the Curious

$$\begin{aligned}\|Y - Xw\|_2^2 &= (Y - Xw)^T (Y - Xw) \\ &= Y^T Y - 2w^T X^T Y + (Xw)^T Xw\end{aligned}$$

$$\frac{\partial}{\partial w} (Xw)^T (Xw) = 2 \left(\frac{\partial}{\partial w} Xw^T \right) Xw = 2X^T Xw$$

$$\begin{aligned}\frac{\partial}{\partial w} \|Y - Xw\|_2^2 &= 0 - 2X^T Y + 2X^T Xw \\ &= 2X^T Xw - 2X^T Y\end{aligned}$$

Two Solutions to Getting W

In One Go

Implicit form
(normal equations)

$$X^T X w = X^T Y$$

Explicit form
(don't do this)

$$w = (X^T X)^{-1} X^T Y$$

Iteratively

Recall: gradient is also
direction that makes
function go up the most.


What could we do?

$$w_0 = \mathbf{0}$$

$$w_{i+1} = w_i - \gamma \left(\frac{\partial}{\partial w} \|Y - Xw\|_2^2 \right)$$

What's The Problem?

- Vertical lines impossible!
- Not rotationally invariant: the line will change depending on orientation of points



Alternate Formulation

Recall: $ax + by + c = 0$

$$\mathbf{l}^T \mathbf{p} = 0$$


$$\mathbf{l} \equiv [a, b, c] \quad \mathbf{p} \equiv [x, y, 1]$$

Can always rescale \mathbf{l} .

Pick a, b, d such that

$$\|\mathbf{n}\|_2^2 = \|[a, b]\|_2^2 = 1$$

$$d = -c$$




Alternate Formulation

Now: $ax + by - d = 0$

$$\mathbf{n}^T [x, y] - d = 0$$

Point to line distance:

$$\frac{\mathbf{n}^T [x, y] - d}{\|\mathbf{n}\|_2} = \mathbf{n}^T [x, y] - d$$




Total Least-Squares

Fitting a line to data

Data: $(x_1, y_1), (x_2, y_2), \dots, (x_k, y_k)$

Model: $(\mathbf{n}, d), \|\mathbf{n}\|^2 = 1$
 $\mathbf{n}^T[x_i, y_i] - d = 0$

Objective function:
 $(\mathbf{n}^T[x_i, y_i] - d)^2$



Total Least Squares Setup

Figure out objective first, then figure out $\|n\|=1$

$$\sum_{i=1}^k (\mathbf{n}^T [x, y] - d)^2 \rightarrow \|X\mathbf{n} - \mathbf{1}d\|_2^2$$

$$\mathbf{X} = \begin{bmatrix} x_1 & y_1 \\ \vdots & \vdots \\ x_k & y_k \end{bmatrix} \quad \mathbf{1} = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \quad \mathbf{n} = \begin{bmatrix} a \\ b \end{bmatrix} \quad \boldsymbol{\mu} = \frac{1}{k} \mathbf{1}^T \mathbf{X}$$

The mean / center of mass of the points:
we'll use it later

Solving Total Least-Squares

$$\begin{aligned}\|X\mathbf{n} - \mathbf{1}d\|_2^2 &= (X\mathbf{n} - \mathbf{1}d)^T (X\mathbf{n} - \mathbf{1}d) \\ &= (X\mathbf{n})^T (X\mathbf{n}) - 2d\mathbf{1}^T X\mathbf{n} + d^2\mathbf{1}^T \mathbf{1}\end{aligned}$$

First solve for d at optimum (set to 0)

$$\frac{\partial}{\partial d} \|X\mathbf{n} - \mathbf{1}d\|_2^2 = 0 - 2\mathbf{1}^T X\mathbf{n} + 2dk$$

$$0 = -2\mathbf{1}^T X\mathbf{n} + 2dk \longrightarrow 0 = -\mathbf{1}^T X\mathbf{n} + dk$$

$$\longrightarrow d = \frac{1}{k} \mathbf{1}^T X\mathbf{n} = \mu n$$

Solving Total Least-Squares

$$\begin{aligned}\|X\mathbf{n} - \mathbf{1}d\|_2^2 &= \|X\mathbf{n} - \mathbf{1}\mu\mathbf{n}\|_2^2 & d &= \mu\mathbf{n} \\ &= \|(X - \mathbf{1}\mu) \mathbf{n}\|_2^2\end{aligned}$$

Objective is then:

$$\arg \min_{\|\mathbf{n}\|=1} \|(X - \mathbf{1}\mu) \mathbf{n}\|_2^2$$

Homogeneous Least Squares

$\arg \min_{\|v\|_2=1} \|Av\|_2^2 \rightarrow$ Eigenvector corresponding to smallest eigenvalue of $A^T A$

Why do we need $\|v\|^2 = 1$ or some other constraint?

Applying it in our case:

$$\mathbf{n} = \text{smallest_eigenvec}((\mathbf{X} - \mathbf{1}\boldsymbol{\mu})^T (\mathbf{X} - \mathbf{1}\boldsymbol{\mu}))$$

Note: technically homogeneous only refers to $\|Av\|=0$ but it's common shorthand in computer vision to refer to the specific problem of $\|v\|=1$

Details For ML-People


Matrix we take the eigenvector of looks like:

$$(X - \mathbf{1}\mu)^T(X - \mathbf{1}\mu) = \begin{bmatrix} \sum_i (x_i - \mu_x)^2 & \sum_i (x_i - \mu_x)(y_i - \mu_y) \\ \sum_i (x_i - \mu_x)(y_i - \mu_y) & \sum_i (y_i - \mu_y)^2 \end{bmatrix}$$


This is a scatter matrix or scalar multiple of the covariance matrix. We're doing PCA, but taking the least principal component to get the normal.

Note: If you don't know PCA, just ignore this slide; it's to help build connections to people with a background in data science/ML.


Running Least-Squares




Running Least-Squares



Ruining Least Squares



Ruining Least Squares



Ruining Least Squares

Way to think of it #1:

$$\|Y - XW\|_2^2$$

$100^2 \gg 10^2$: least-squares prefers having no large errors, even if the model is useless overall

Way to think of it #2:

$$W = \underline{(X^T X)^{-1} X^T Y}$$

Weights are a linear transformation of the output variable: can manipulate W by manipulating Y .

Common Fixes

Replace Least-Squares objective

Let $E = Y - XW$

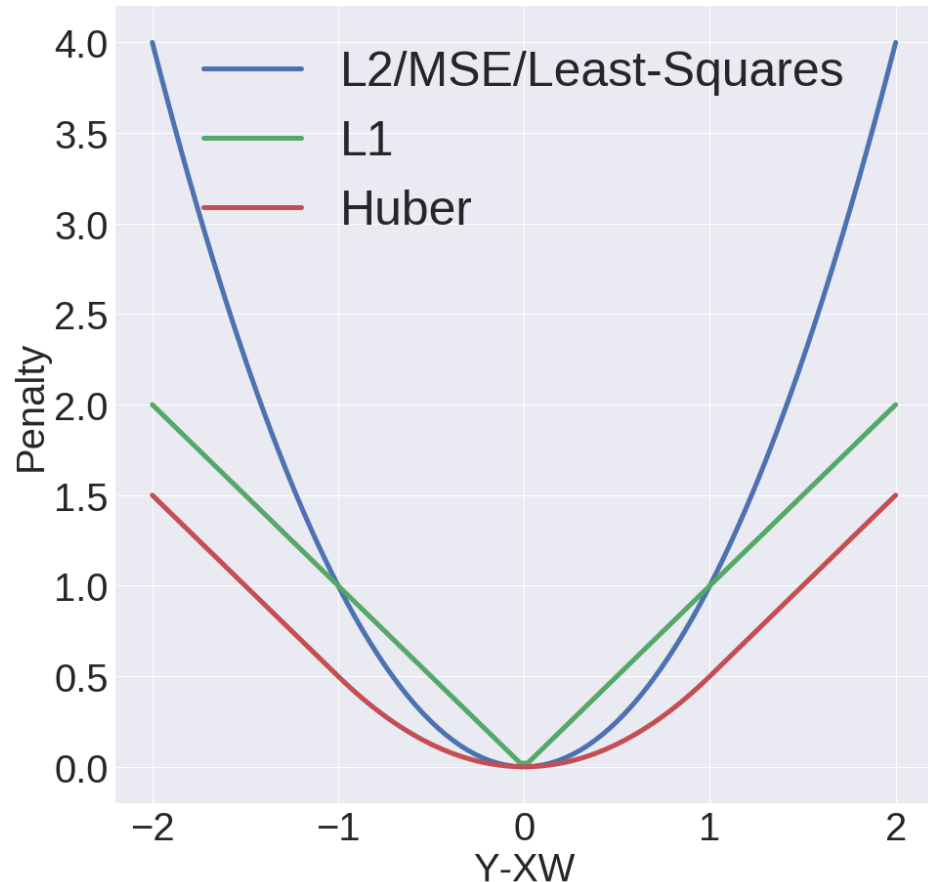
LS/L2/MSE: E_i^2

L1: $|E_i|$

Huber:

$|E_i| \leq \delta$: $\frac{1}{2}E_i^2$

$|E_i| > \delta$: $\delta(|E_i| - \frac{\delta}{2})$




Issues with Common Fixes


- Usually complicated to optimize:
 - Often no closed form solution
 - Typically not something you could write yourself
 - Sometimes not convex (no global optimum)
- Not simple to extend more complex objectives to things like total-least squares
- Typically don't handle a ton of outliers (e.g., 80% outliers)

Outliers in Computer Vision


Single outlier:
rare




Many outliers:
common



Ruining Least Squares Continued



Ruining Least Squares Continued



A Simple, Yet Clever Idea

- *What we really want*: model explains **many** points “**well**”
- *Least Squares*: model makes as few big mistakes as possible over the entire dataset
- *New objective*: find model for which error is “small” for as many data points as possible
- *Method*: RANSAC (**RA**ndom **SA**mple **C**onsensus)

RANSAC For Lines

```
bestLine, bestCount = None, -1
```

```
for trial in range(numTrials):
```

```
    subset = pickPairOfPoints(data)
```

```
    line = totalLeastSquares(subset)
```

```
    E = linePointDistance(data, line)
```


```
    inliers = E < threshold
```

```
    if #inliers > bestCount:
```


```
        bestLine, bestCount = line, #inliers
```

Running RANSAC

Trial
#1




Best
Model:




Best
Count:
-1

Running RANSAC

Trial
#1




Best
Model:




Best
Count:
-1

Running RANSAC

Trial
#1



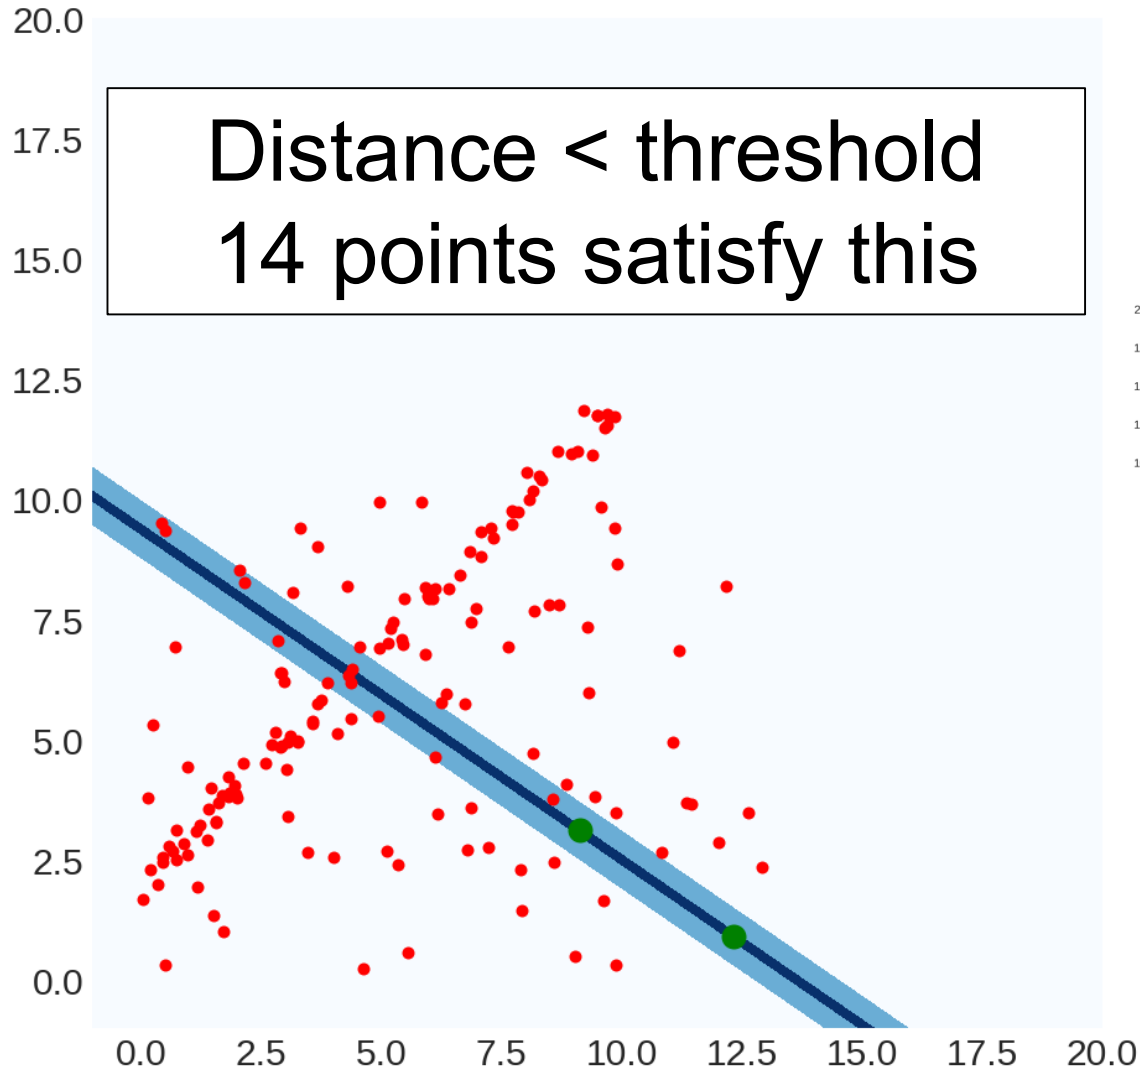
Best
Model:




Best
Count:
-1

Running RANSAC

Trial
#1




Best
Model:




Best
Count:
-1

Running RANSAC

Trial
#1




Best
Model:




Best
Count:
14

Running RANSAC

Trial
#2




Best
Model:




Best
Count:
14

Running RANSAC

Trial
#2




Best
Model:




Best
Count:
22

Running RANSAC

Trial
#3




Best
Model:




Best
Count:
22

Running RANSAC

Trial
#3




Best
Model:




Best
Count:
22

Running RANSAC

Trial
#9




Best
Model:




Best
Count:
22

Running RANSAC

Trial
#9




Best
Model:




Best
Count:
76

Running RANSAC

Trial
#9




Best
Model:




Best
Count:
76

Running RANSAC

Trial
#100




Best
Model:



Best
Count:
85

Running RANSAC



RANSAC In General

best, bestCount = None, -1

for trial in range(**NUM_TRIALS**):

 subset = pickSubset(data, **SUBSET_SIZE**)

 model = fitModel(subset)

 E = computeError(data, line)

 inliers = E < **THRESHOLD**

 if #(inliers) > bestCount:

 best, bestCount = model, #(inliers)

(often refit on the inliers for best model)

Parameters – Num Trials

r is the fraction of outliers (e.g., 80%)
Suppose we pick s points (e.g., 2)
we run RANSAC N times (e.g., 500)

What's the probability of picking a sample set with no outliers?

$$\approx (1 - r)^s \quad (4\%)$$

What's the probability of picking a sample set with any outliers?

$$1 - (1 - r)^s \quad (96\%)$$

Parameters – Num Trials

r is the fraction of outliers (e.g., 80%)
Suppose we pick s points (e.g., 2)
we run RANSAC N times (e.g., 500)

What's the probability of picking a sample set with any outliers?

$$1 - (1 - r)^s \quad \textbf{(96\%)}$$

What's the probability of picking only sample sets with outliers?

$$(1 - (1 - r)^s)^N \quad \textbf{(10^{-7}\% N=500)}$$

$$\textbf{(13\% N=50)}$$

What's the probability of picking any set with inliers?

$$1 - (1 - (1 - r)^s)^N$$

Parameters – Num Trials



P(\$157M Jackpot):
1 / 302,575,350

Death by
vending
machine



P(Death):
 $\approx 1 / 112,000,000$

RANSAC fails to fit a
line with 80% outliers
after trying only 500
times

P(Failure):
1 / 731,784,961

Parameters – Subset Size

- Always the smallest possible set for fitting the model.
- Minimum number for lines: 2 data points
- Minimum number of planes: **how many?**

- **Why intuitively?**
- You'll find out more precisely in homework 3.

Parameters – Threshold

- Common sense; there's no magical threshold

RANSAC Pros and Cons


Pros

1. Ridiculously simple
2. Ridiculously effective
3. Works in general

Cons


1. Have to tune parameters
2. No theory (so can't derive parameters via theory)
3. Not magic, especially with lots of outliers

Hough Transform



Hough Transform


1. Discretize space of parametric models
2. Each pixel votes for all compatible models
3. Find models compatible with many pixels



P.V.C. Hough, *Machine Analysis of Bubble Chamber Pictures*, Proc. Int. Conf. High Energy Accelerators and Instrumentation, 1959

Hough Transform


Line in image = point in parameter space



Hough Transform

Point in image = line in parameter space


All lines through the point: $b = x_0 m + y_0$



Hough Transform


Point in image = line in parameter space

All lines through the point: $b = x_1 m + y_1$




Hough Transform

Line through two points in image = intersection of two lines in parameter space (i.e., solutions to both equations)



Hough Transform

Line through two points in image = intersection of two lines in parameter space (i.e., solutions to both equations)



Hough Transform

- *Recall*: m, b space is awful
- $ax+by+c=0$ is better, but *unbounded*
- Trick: write lines using angle + offset (normally a mediocre way, but makes things bounded)

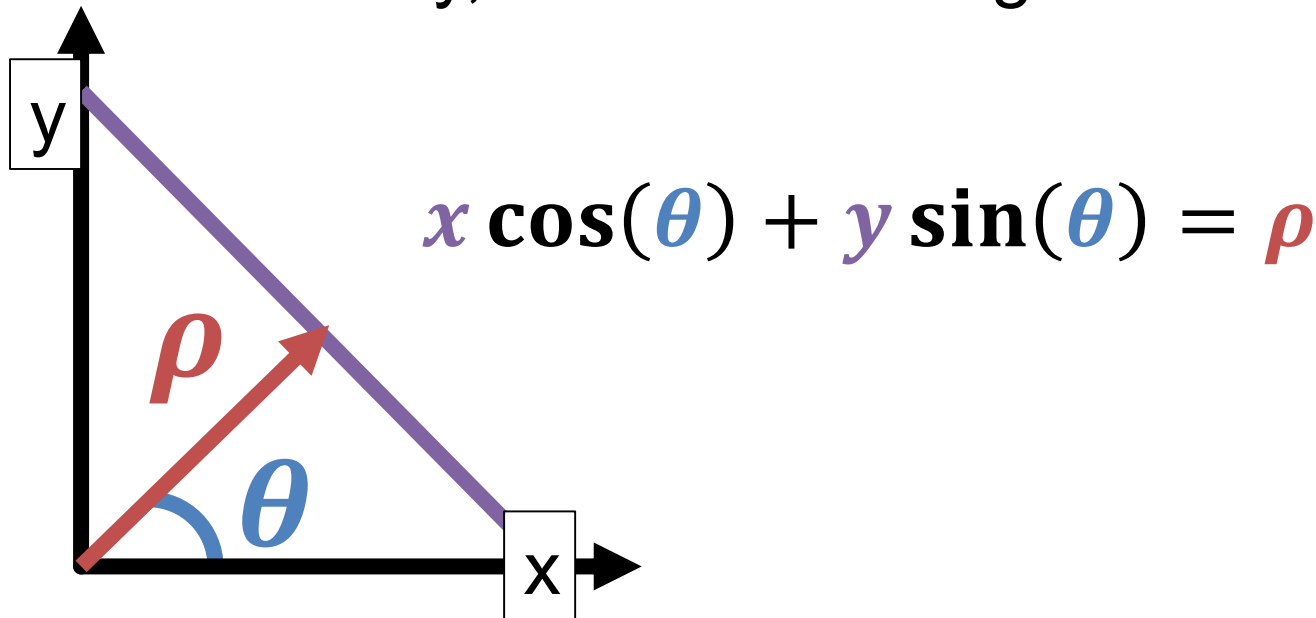


Diagram is remake of S. Seitz Slides; these are illustrative and values may not be real

Hough Transform Algorithm

Remember: $x \cos(\theta) + y \sin(\theta) = \rho$

Accumulator $H = \text{zeros}(?, ?)$

For x, y in `detected_points`:


For θ in `range(0, 180, ?)`:

$$\rho = x \cos(\theta) + y \sin(\theta)$$

$$H[\theta, \rho] += 1$$

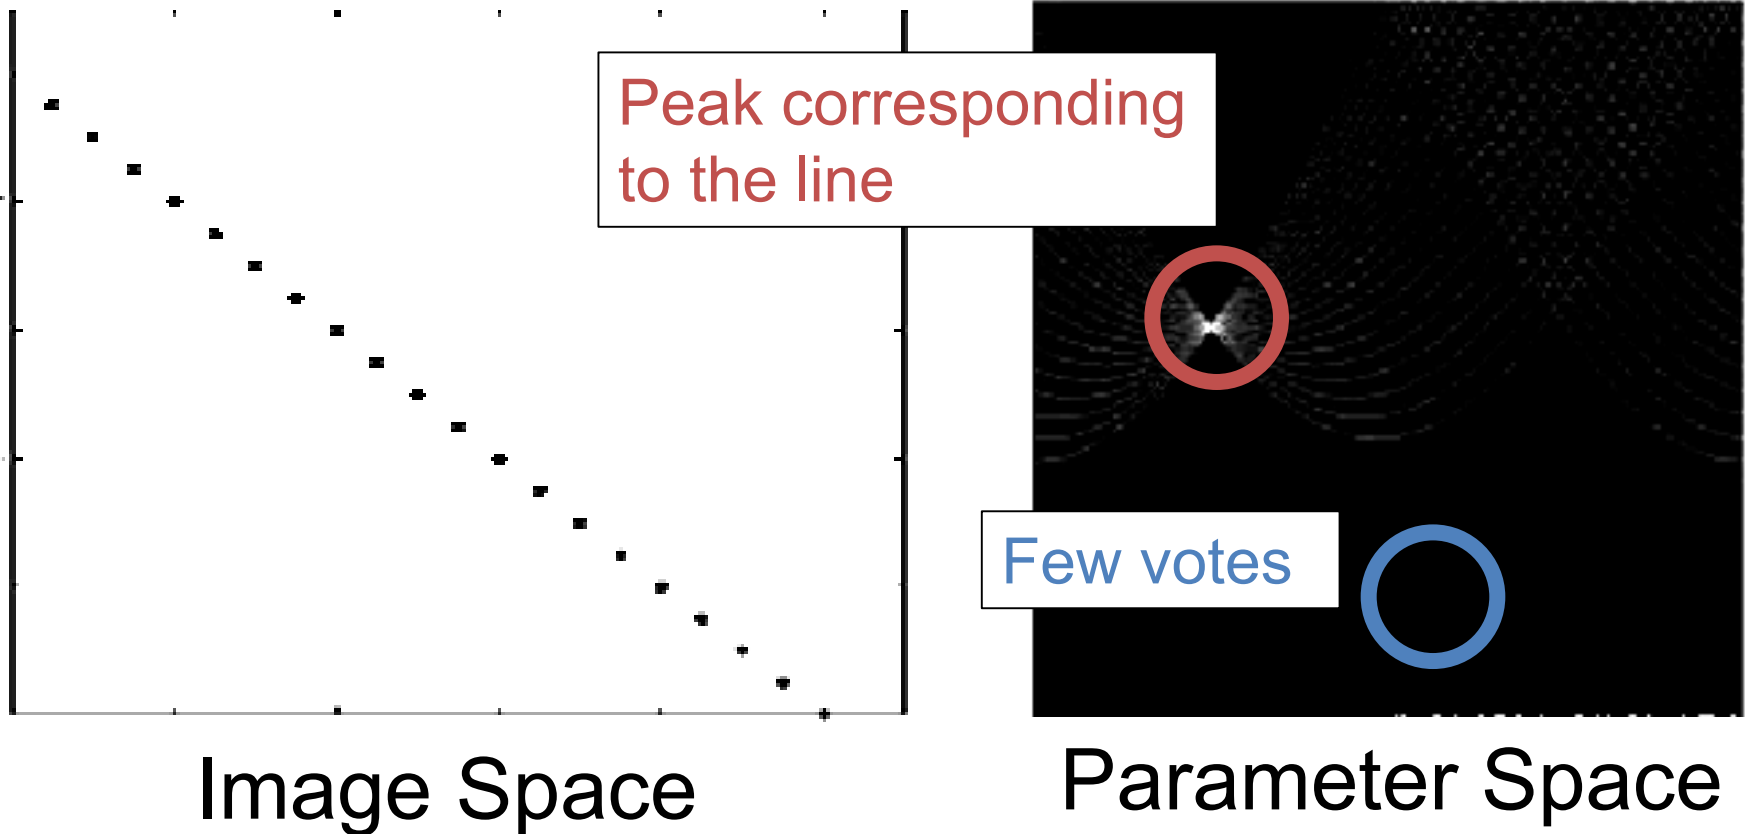
#any local maxima (θ, ρ) of H is a line

#of the form $\rho = x \cos(\theta) + y \sin(\theta)$



Example

Points (x,y) \rightarrow sinusoids



Hough Transform Pros / Cons

Pros

1. Handles **multiple** models
2. Some robustness to noise
3. In principle, general

Cons

1. Have to bin ALL parameters: exponential in #params
2. Have to parameterize your space nicely
3. Details really, really important (a working version requires a lot more than what I showed you)

Next Time

- What happens with fitting more complex transformations?