

# *Image Stitching and Alignment*

# Multiple Images

- ❖ So far, algorithms deal with *a single, static* image
- ❖ In the real world, a static pattern is a rarity, continuous motion and change are the rule
- ❖ Human eyes are well-equipped to take advantage of motion or change in *an image sequence*
- ❖ *Stitching (Alignment) and Motion*
  - ❑ Stitching has a “global” model – all pixel movement can be explained by a simple mathematic model (far field, pure rotational, pure translation)
  - ❑ 2D motion field is a “local” model – pixels by themselves (similarity in a local neighborhood only)

# General Taxonomy

---

- ❖ Camera motion and the Scene is static
  - ❑ Driving, panorama
  - ❑ Near field (hard) vs. Far field (easy)
  - ❑ General camera motion (hard) vs. special camera motion (e.g., rotation only, easier)
  - ❑ General scene (hard) vs. special scene (planar, easier)
- ❖ Object motion and the camera is stationary
  - ❑ Surveillance
  - ❑ Background modeling and subtraction
- ❖ Both camera and object are moving
  - ❑ Sports video, driving, diving, etc.

# Alignment

- ❖ Homographies
- ❖ Rotational Panoramas
- ❖ RANSAC
- ❖ Global alignment
- ❖ Warping
- ❖ Blending



(a)



(b)








(c)

# Motivation: Recognition



# Motivation: medical image registration



# *Motivation: Mosaics*

---

## ❖ Getting the whole picture

❑ Consumer camera:  $50^\circ \times 35^\circ$





# *Motivation: Mosaics*


- ❖ Getting the whole picture
  - ❑ Consumer camera:  $50^\circ \times 35^\circ$
  - ❑ Human Vision:  $176^\circ \times 135^\circ$





# *Motivation: Mosaics*

- ❖ Getting the whole picture
  - ❑ Consumer camera:  $50^\circ \times 35^\circ$
  - ❑ Human Vision:  $176^\circ \times 135^\circ$



# *Motion models*


- ❖ What happens when we take two images with a camera and try to align them?
  - translation?
  - rotation?
  - scale?
  - affine?
  - perspective?



# Image Warping


- ❖ image filtering: change *range* of image

- ❖  $g(x) = h(f(x))$



- ❖ image warping: change *domain* of image


- ❖  $g(x) = f(h(x))$



# Image Warping


❖ image filtering: change *range* of image

$$\text{❖ } g(x) = h(f(x))$$



❖ image warping: change *domain* of image

$$\text{❖ } g(x) = f(h(x))$$



# *Parametric (global) warping*

❖ Examples of parametric warps:



translation



rotation



aspect



affine



perspective




cylindrical




# Image Warping

- ❖ Given a coordinate transform  $\mathbf{x}' = \mathbf{h}(\mathbf{x})$  and a source image  $f(\mathbf{x})$ , how do we compute a transformed image  $g(\mathbf{x}') = f(\mathbf{h}(\mathbf{x}))$ ?



# Forward Warping


- ❖ Send each pixel  $f(x)$  to its corresponding location  $x' = h(x)$  in  $g(x')$
- What if pixel lands “between” two pixels?






# Forward Warping

- ❖ Send each pixel  $f(x)$  to its corresponding location  $x' = h(x)$  in  $g(x')$
- What if pixel lands “between” two pixels?
- Answer: add “contribution” to several pixels, normalize later (*splatting*)




# Inverse Warping

- ❖ Get each pixel  $g(x')$  from its corresponding location  $x' = h(x)$  in  $f(x)$
- What if pixel comes from “between” two pixels?



# Inverse warping



Get each pixel  $g(x', y')$  from its corresponding location

$(x, y) = T^{-1}(x', y')$  in the first image

Q: what if pixel comes from “between” two pixels?


A: *Interpolate* color value from neighbors

- nearest neighbor, bilinear...



# Bilinear interpolation

Sampling at  $f(x, y)$ :



$$\begin{aligned} f(x, y) = & (1 - a)(1 - b) f[i, j] \\ & + a(1 - b) f[i + 1, j] \\ & + ab f[i + 1, j + 1] \\ & + (1 - a)b f[i, j + 1] \end{aligned}$$

# Interpolation

## ❖ Possible interpolation filters:

- nearest neighbor
- bilinear
- bicubic (interpolating)
- sinc / FIR

## ❖ Needed to prevent “jaggies” and “texture crawl”



# 2D coordinate transformations

- ❖ translation:  $\mathbf{x}' = \mathbf{x} + \mathbf{t}$   $\mathbf{x} = (x, y)$
- ❖ rotation:  $\mathbf{x}' = \mathbf{R} \mathbf{x} + \mathbf{t}$
- ❖ similarity:  $\mathbf{x}' = s \mathbf{R} \mathbf{x} + \mathbf{t}$
- ❖ affine:  $\mathbf{x}' = \mathbf{A} \mathbf{x} + \mathbf{t}$
- ❖ perspective:  $\underline{\mathbf{x}}' \cong \mathbf{H} \underline{\mathbf{x}}$   $\underline{\mathbf{x}} = (x, y, 1)$   
( $\underline{\mathbf{x}}$  is a *homogeneous* coordinate)
- ❖ These all form a nested *group* (closed w/ inv.)

# *Homogeneous Coordinates*

---

- ❖ consistent representation for all linear transform (including translation)
- ❖ can be concatenated & pre-computed

$$(x, y) \rightarrow (wx, wy, w), w \neq 0$$

$$(wx, wy, w) \rightarrow (wx / w, wy / w)$$



$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & T_x \\ 0 & 1 & T_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} S_x & 0 & 0 \\ 0 & S_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = (TRS) \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

# Basic 2D Transformations

## ❖ Basic 2D transformations as 3x3 matrices

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Translate

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Scale

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos \Theta & -\sin \Theta & 0 \\ \sin \Theta & \cos \Theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Rotate

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & sh_x & 0 \\ sh_y & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Shear

# 2D Affine Transformations

---

$$\begin{bmatrix} x' \\ y' \\ w \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

- ❖ Affine transformations are combinations of ...
  - ❑ Linear transformations, and
  - ❑ Translations
  
- ❖ Parallel lines remain parallel

# Projective Transformations


$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

## ❖ Projective transformations:

- ❑ Affine transformations, and

- ❑ Projective warps

## ❖ Parallel lines do not necessarily remain parallel




# *Fitting an affine transformation*



Affine model approximates perspective projection of planar objects.

# Fitting an affine transformation


- Assuming we know the correspondences, how do we get the transformation?



$$\begin{bmatrix} x'_i \\ y'_i \end{bmatrix} = \begin{bmatrix} m_1 & m_2 \\ m_3 & m_4 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$$


# Fitting an affine transformation

- Assuming we know the correspondences, how do we get the transformation?



$$\begin{bmatrix} x'_i \\ y'_i \end{bmatrix} = \begin{bmatrix} m_1 & m_2 \\ m_3 & m_4 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$$

^

$$\begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \phantom{m_1} \\ \phantom{m_2} \\ \phantom{m_3} \\ \phantom{m_4} \\ \phantom{t_1} \\ \phantom{t_2} \end{bmatrix}$$





# *Fitting an affine transformation*

$$\begin{bmatrix} & & \Lambda & & & & \\ x_i & y_i & 0 & 0 & 1 & 0 & \\ 0 & 0 & x_i & y_i & 0 & 1 & \\ & & \Lambda & & & & \end{bmatrix} \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \Lambda \\ x'_i \\ y'_i \\ \Lambda \end{bmatrix}$$

- How many matches (correspondence pairs) do we need to solve for the transformation parameters?
- Once we have solved for the parameters, how do we compute the coordinates of the corresponding point for ?  $(x_{new}, y_{new})$

# Panoramas



...



image from S. Seitz

Obtain a wider angle view by combining multiple images.

# How to stitch together a panorama?


## ❖ Basic Procedure

- ❑ Take a sequence of images from the same position
  - Rotate the camera about its optical center
- ❑ Compute transformation between second image and first
- ❑ Transform the second image to overlap with the first
- ❑ Blend the two together to create a mosaic
- ❑ (If there are more images, repeat)

## ❖ ...but **wait**, why should this work at all?

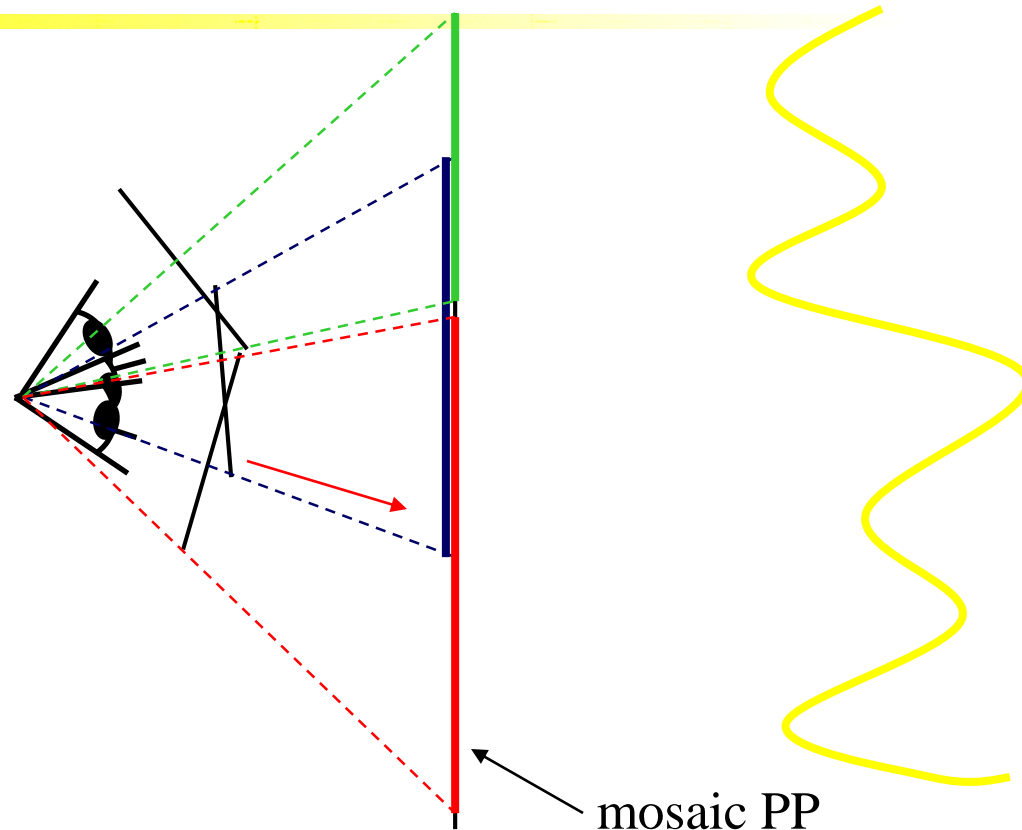
- ❑ What about the 3D geometry of the scene?
- ❑ Why aren't we using it?

# *Panoramas: generating synthetic views*




Can generate any synthetic camera view  
as long as it has **the same center of projection!**

# Image reprojection



- ❖ The mosaic has a natural interpretation in 3D
  - ❑ The images are reprojected onto a common plane
  - ❑ The mosaic is formed on this plane
  - ❑ Mosaic is a *synthetic wide-angle camera*

# Image reprojection



- ❖ The mosaic has a natural interpretation in 3D as a plane
- ❖ This is true even if the real scene is not planar as long as you have *the same focal point*

# *In reality*


---

- ❖ The scene is not planar
  - ❑ But if you are shooting panorama against far-away objects (e.g., from the south rim of the Grand Canyon against the north rim), the distance variation can be ignored
  - ❑ Panorama works best for far-field scene
- ❖ The rotation is about the person holding the camera, not the camera's focal center
  - ❑ If the scene is far away, such small deviation does not matter
- ❖ In fact, image stitching works well if you exercise some caution
- ❖ Why all phones these days have the panorama mode

# Homography

- ❖ How to relate two images from the same camera center?
  - how to map a pixel from PP1 to PP2?
- ❖ Think of it as a 2D **image warp** from one image to another.
- ❖ A projective transform is a mapping between any two PPs with the same center of projection
  - ❑ rectangle should map to arbitrary quadrilateral
  - ❑ parallel lines aren't
  - ❑ but must preserve straight lines
- ❖ called **Homography**

$$\begin{bmatrix} wx' \\ wy' \\ w \\ \mathbf{p}' \end{bmatrix} = \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \\ \mathbf{H} & & \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \\ \mathbf{p} \end{bmatrix}$$





# Homography

$(x, y)$



$$\begin{pmatrix} wx' / w & wy' / w \end{pmatrix} = (x', y')$$

To **apply** a given homography  $\mathbf{H}$


- Compute  $\mathbf{p}' = \mathbf{H}\mathbf{p}$  (regular matrix multiply)
- Convert  $\mathbf{p}'$  from homogeneous to image coordinates

$$\begin{bmatrix} wx' \\ wy' \\ w \end{bmatrix} = \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$\mathbf{p}' = \mathbf{H}\mathbf{p}$



# Homography



To **compute** the homography given pairs of corresponding points in the images, we need to set up an equation where the parameters of  $\mathbf{H}$  are the unknowns...

# Number of measurements required

- ❖ At least as many independent equations as degrees of freedom required
- ❖ Example:

$$\lambda \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \mathbf{H} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

The matrix  $\mathbf{H}$  is defined as:

$$\mathbf{H} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix}$$

2 independent equations / point  
8 degrees of freedom

$$4 \times 2 \geq 8$$



# Solving for homographies

$$\mathbf{p}' = \mathbf{H}\mathbf{p}$$
$$\begin{bmatrix} wx' \\ wy' \\ w \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

❖ Can set scale factor  $i=1$ . So, there are 8 unknowns.

❖ Set up a system of linear equations:

$$\text{❖ } \mathbf{A}\mathbf{h} = \mathbf{b}$$

❖ where vector of unknowns  $\mathbf{h} = [a,b,c,d,e,f,g,h]^T$

❖ Need at least 8 eqs, but the more the better...

❖ Solve for  $\mathbf{h}$ . If overconstrained, solve using least-squares:

$$\min \|\mathbf{A}\mathbf{h} - \mathbf{b}\|^2$$

❖ Work well if  $i$  is not close to 0 (not recommended!)

# Direct Linear Transformation (DLT)

$$H = \begin{bmatrix} h^{1T} \\ h^{2T} \\ h^{3T} \end{bmatrix}$$

$$\mathbf{x}'_i \times \mathbf{Hx}_i = 0 \quad \mathbf{x}'_i = (x'_i, y'_i, w'_i)^T \quad \mathbf{Hx}_i = \begin{pmatrix} h^{1T} \mathbf{x}_i \\ h^{2T} \mathbf{x}_i \\ h^{3T} \mathbf{x}_i \end{pmatrix}$$

$$\mathbf{x}'_i \times \mathbf{Hx}_i = \begin{pmatrix} y'_i h^{3T} \mathbf{x}_i - w'_i h^{2T} \mathbf{x}_i \\ w'_i h^{1T} \mathbf{x}_i - x'_i h^{3T} \mathbf{x}_i \\ x'_i h^{2T} \mathbf{x}_i - y'_i h^{1T} \mathbf{x}_i \end{pmatrix}$$

$$\begin{bmatrix} 0^T & -w'_i \mathbf{x}_i^T & y'_i \mathbf{x}_i^T \\ w'_i \mathbf{x}_i^T & 0^T & -x'_i \mathbf{x}_i^T \\ -y'_i \mathbf{x}_i^T & x'_i \mathbf{x}_i^T & 0^T \end{bmatrix} \begin{pmatrix} h^1 \\ h^2 \\ h^3 \end{pmatrix} = 0$$

$$\mathbf{A}_i \mathbf{h} = 0$$



# Direct Linear Transformation (DLT)

❖ Equations are linear in  $\mathbf{h}$

$$A_i \mathbf{h} = 0$$

- Only 2 out of 3 are linearly independent (indeed, 2 eq/pt)

$$\begin{bmatrix} 0^T & -w'_i X_i^T & y'_i X_i^T \\ 0^T & -w'_i Y_i^T & y'_i X_i^T \\ w'_i X_i^T & 0^T & -x'_i X_i^T \\ w'_i Y_i^T & 0^T & -x'_i X_i^T \\ y'_i X_i^T & x'_i X_i^T & 0^T \end{bmatrix} \begin{pmatrix} h^1 \\ h^2 \\ h^3 \end{pmatrix} = 0$$

(only drop third row if  $w'_i \neq 0$ )

- Holds for any homogeneous representation, e.g.  $(x'_i, y'_i, 1)$

# Direct Linear Transformation (DLT)

## ❖ Solving for $\mathbf{H}$

$$\begin{bmatrix} A_1 \\ A_2 \\ A_3 \\ A_4 \end{bmatrix} \mathbf{h} = \mathbf{0}$$

size  $A$  is  $8 \times 9$  or  $12 \times 9$ , but rank 8

Trivial solution is  $\mathbf{h} = \mathbf{0}_9^T$  is not interesting

1-D null-space yields solution of interest

pick for example the one with  $\|\mathbf{h}\| = 1$



# Direct Linear Transformation (DLT)

- ❖ Over-determined solution

$$\begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_n \end{bmatrix} \mathbf{h} = \mathbf{0}$$

No exact solution because of inexact measurement  
i.e. "noise"

Find approximate solution

- Additional constraint needed to avoid 0, e.g.  $\|\mathbf{h}\| = 1$
- $A\mathbf{h} = \mathbf{0}$  not possible, so minimize  $\|A\mathbf{h}\|$





# *DLT algorithm*

## Objective

Given  $n \geq 4$  2D to 2D point correspondences  $\{x_i \leftrightarrow x_i'\}$ , determine the 2D homography matrix  $H$  such that  $x_i' = Hx_i$


## Algorithm

- (i) For each correspondence  $x_i \leftrightarrow x_i'$  compute  $A_i$ . Usually only two first rows needed.
- (ii) Assemble  $n$   $2 \times 9$  matrices  $A_i$  into a single  $2n \times 9$  matrix  $A$
- (iii) Obtain SVD of  $A$ . Solution for  $h$  is last column of  $V$
- (iv) Determine  $H$  from  $h$

# *changing camera center*

❖ Does it still work?

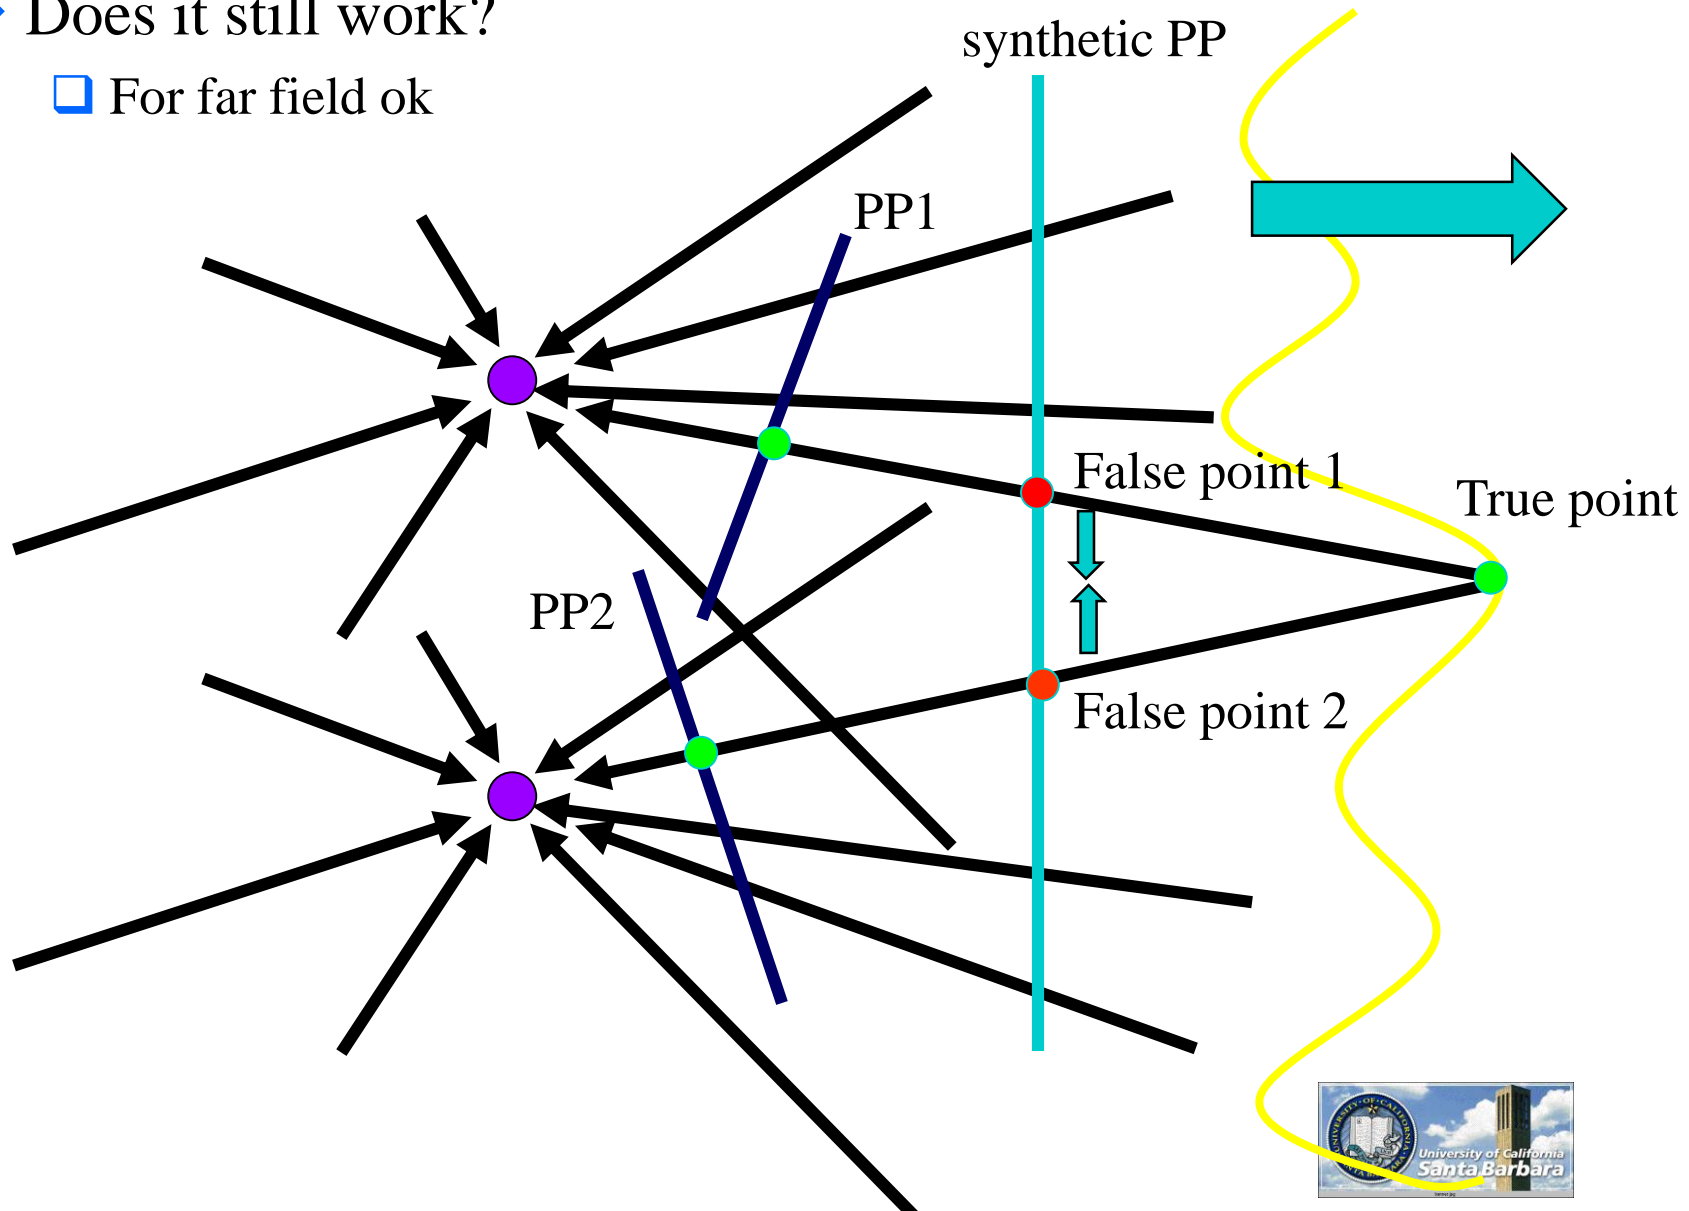
❑ Not for near field




# *changing camera center*

❖ Does it still work?

☐ For far field ok



# *Planar scene (or far away)*



- ❖ PP3 is a projection plane of both centers of projection, so we are OK!
- ❖ This is how big aerial photographs are made








200 ft  
100 m



# Outliers


- ❖ **Outliers** can hurt the quality of our parameter estimates, e.g.,
  - ❑ an erroneous pair of matching points from two images
  - ❑ an edge point that is noise, or doesn't belong to the line we are fitting.



# *Example: least squares line fitting*


---

- ❖ Assuming all the points that belong to a particular line are known






# *Outliers affect least squares fit*



# Outliers affect least squares fit



# RANSAC

---

- ❖ RANdOm Sample Consensus
- ❖ Approach: we want to avoid the impact of outliers, so let's look for “inliers”, and use those only.
- ❖ Intuition: if an outlier is chosen to compute the current fit, then the resulting line won't have much support from rest of the points.

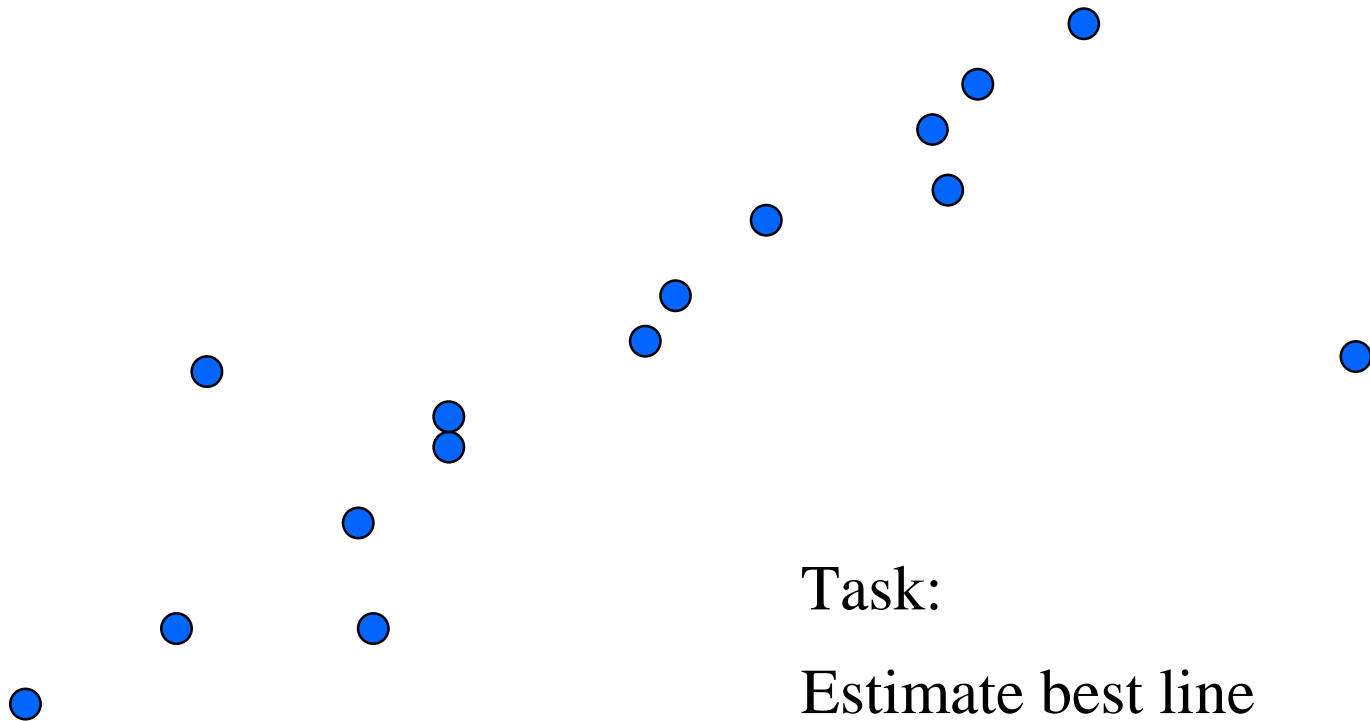
# RANSAC

## ❖ RANSAC loop:

1. Randomly select a *seed group* of points on which to base transformation estimate (e.g., a group of matches)
  2. Compute transformation from seed group
  3. Find *inliers* to this transformation
  4. If the number of inliers is sufficiently large, re-compute least-squares estimate of transformation on all of the inliers
- ❖ Keep the transformation with the largest number of inliers

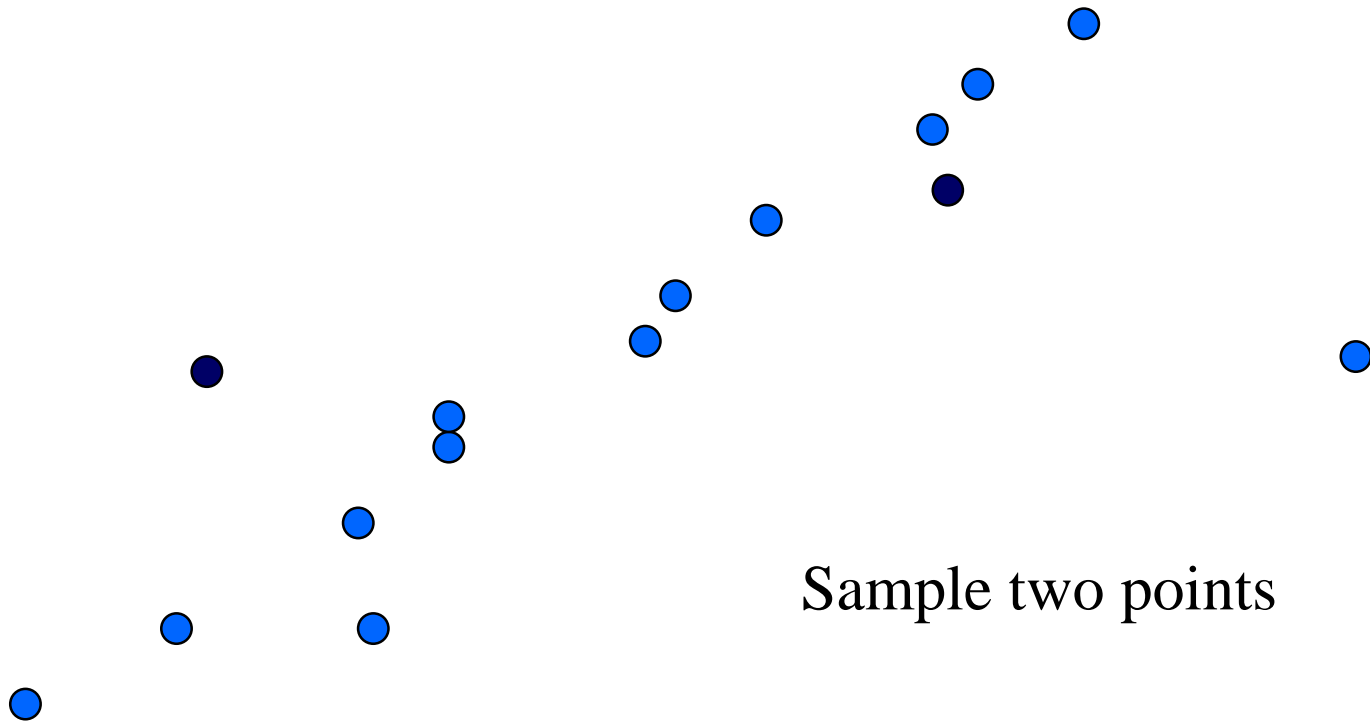
# *RANSAC Line Fitting Example*

---



# *RANSAC Line Fitting Example*


---



Sample two points

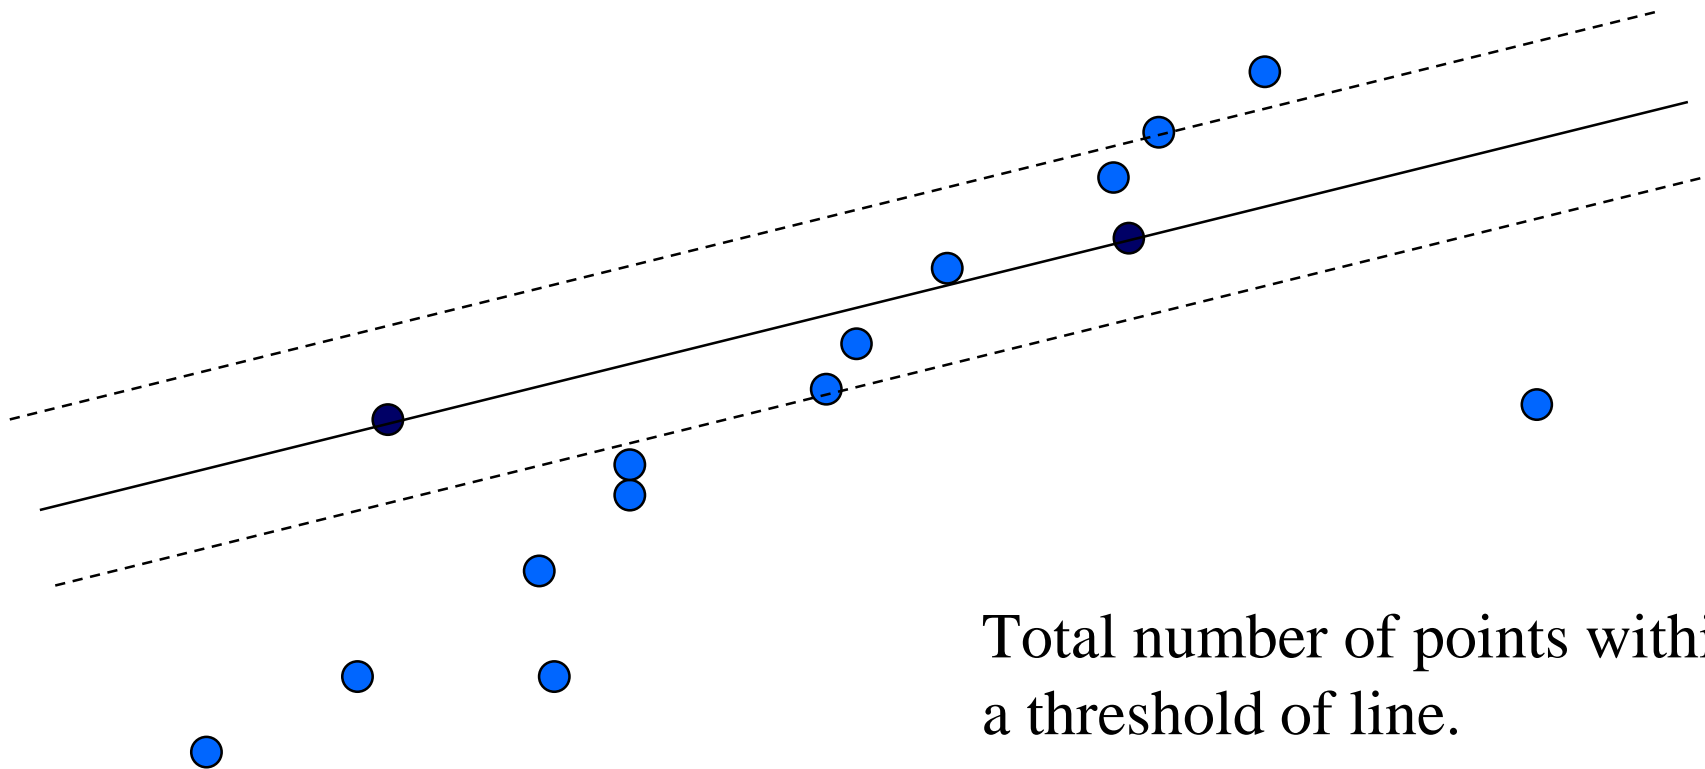
# *RANSAC Line Fitting Example*

---



Fit Line

# *RANSAC Line Fitting Example*

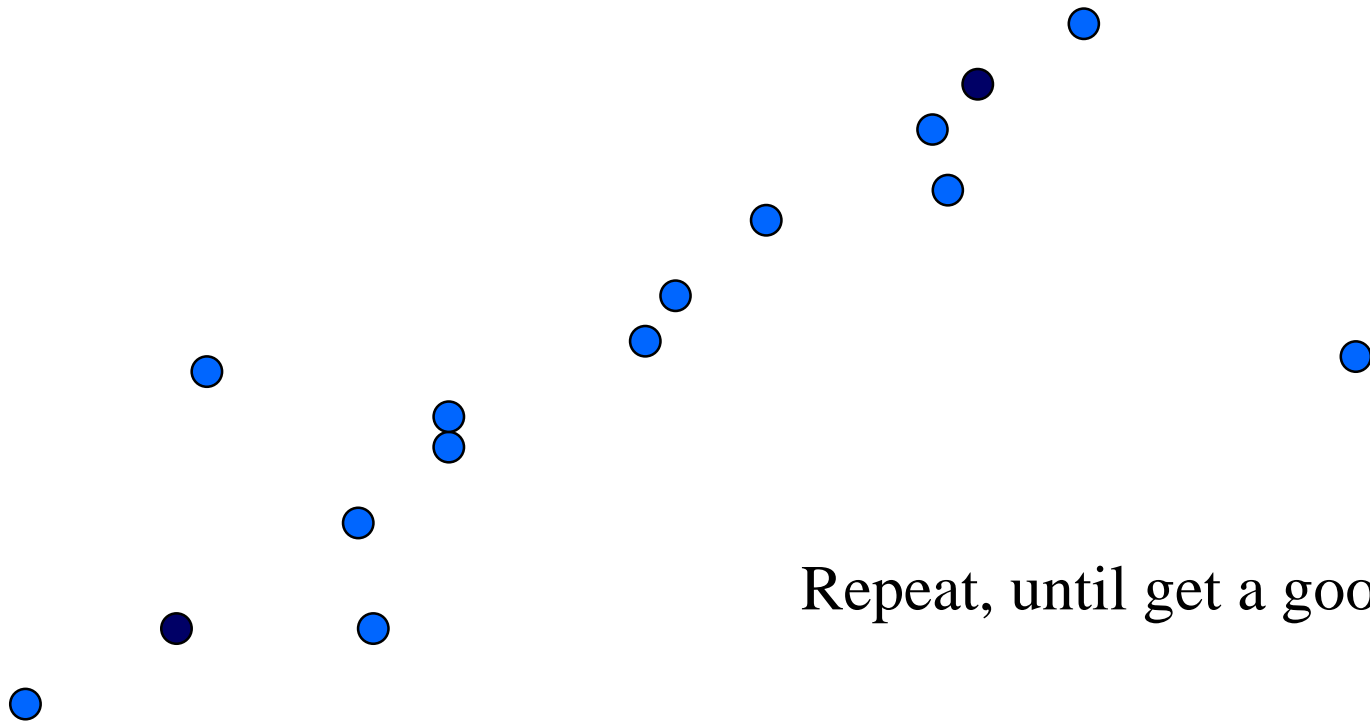


Total number of points within  
a threshold of line.



# *RANSAC Line Fitting Example*

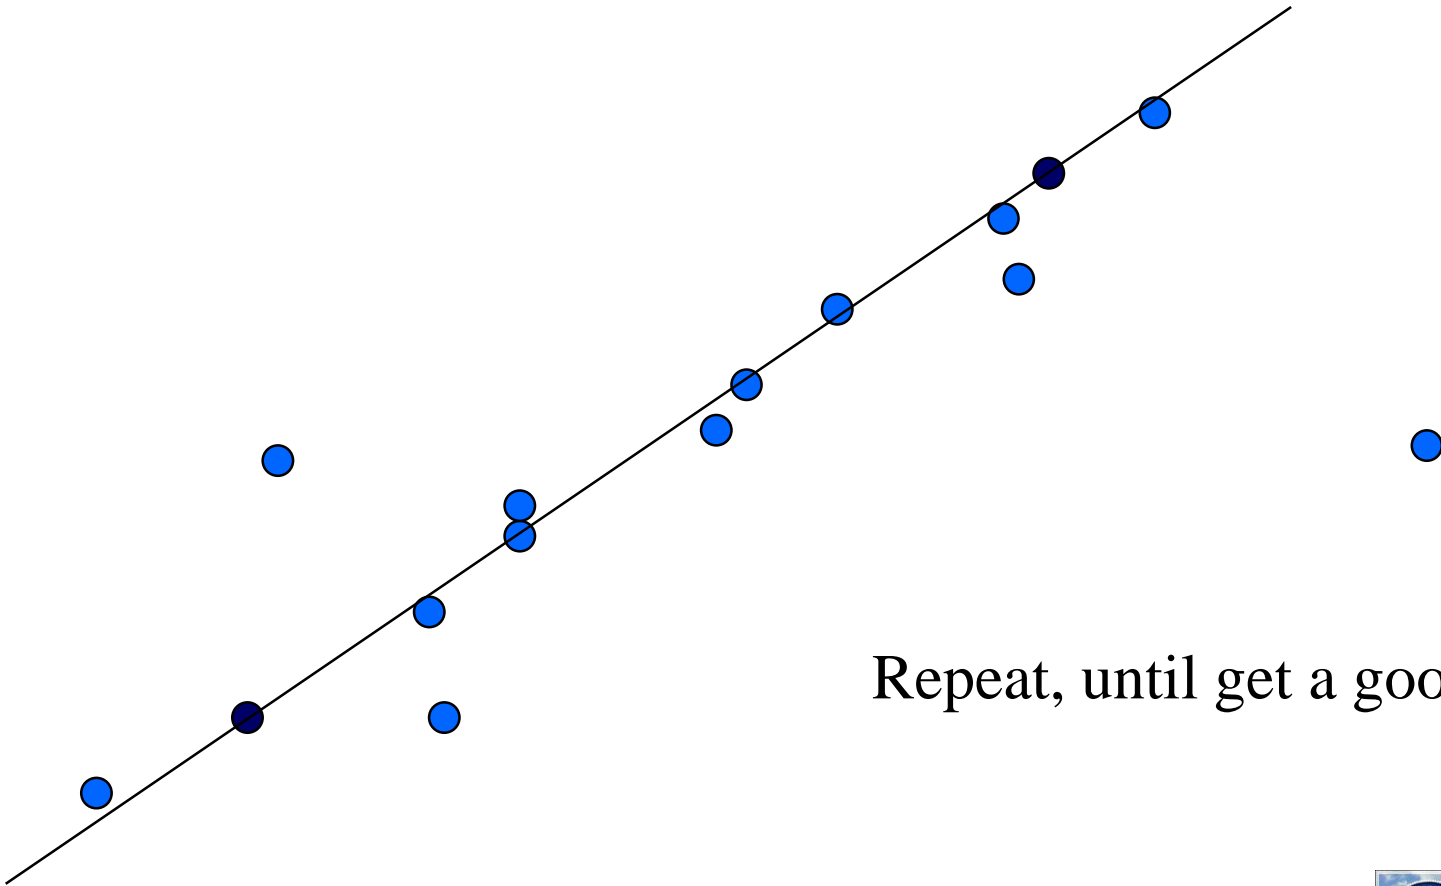
---



Repeat, until get a good result

# *RANSAC Line Fitting Example*


---



Repeat, until get a good result

# *RANSAC Line Fitting Example*

---



Repeat, until get a good result

# How Many Trials?

- ❖ Well, theoretically it is  $C(n,p)$  to find all possible  $p$ -tuples
- ❖ Very expensive

$$1 - (1 - (1 - \varepsilon)^p)^m$$

$\varepsilon$  : fraction of bad data

$(1 - \varepsilon)$  : fraction of good data

$(1 - \varepsilon)^p$  : all  $p$  samples are good

$1 - (1 - \varepsilon)^p$  : at least one sample is bad

$(1 - (1 - \varepsilon)^p)^m$  : got bad data in all  $m$  tries

$1 - (1 - (1 - \varepsilon)^p)^m$  : got at least one good  $p$  set in  $m$  tries



# How Many Trials (cont.)

- ❖ Make sure the probability is high (e.g. >95%)
- ❖ given  $p$  and epsilon, calculate  $m$

$p$	5%	10	20	25	30	40	50
		%	%	%	%	%	%
1	1	2	2	3	3	4	5
2	2	2	3	4	5	7	11
3	2	3	5	6	8	13	23
4	2	3	6	8	11	22	47
5	3	4	8	12	17	38	95


# Best Practice

---

- ❖ Randomized selection can completely remove outliers
- ❖ “plutocratic”
- ❖ Results are based on a small set of features
- ❖ LS is most fair, everyone get an equal say
- ❖ “democratic”
- ❖ But can be seriously influenced by bad data
- ❖ Use randomized algorithm to remove outliers
- ❖ Use LS for final “polishing” of results (using all “good” data)
- ❖ Allow up to 50% outliers theoretically




# *RANSAC example: Translation*






# *RANSAC example: Translation*

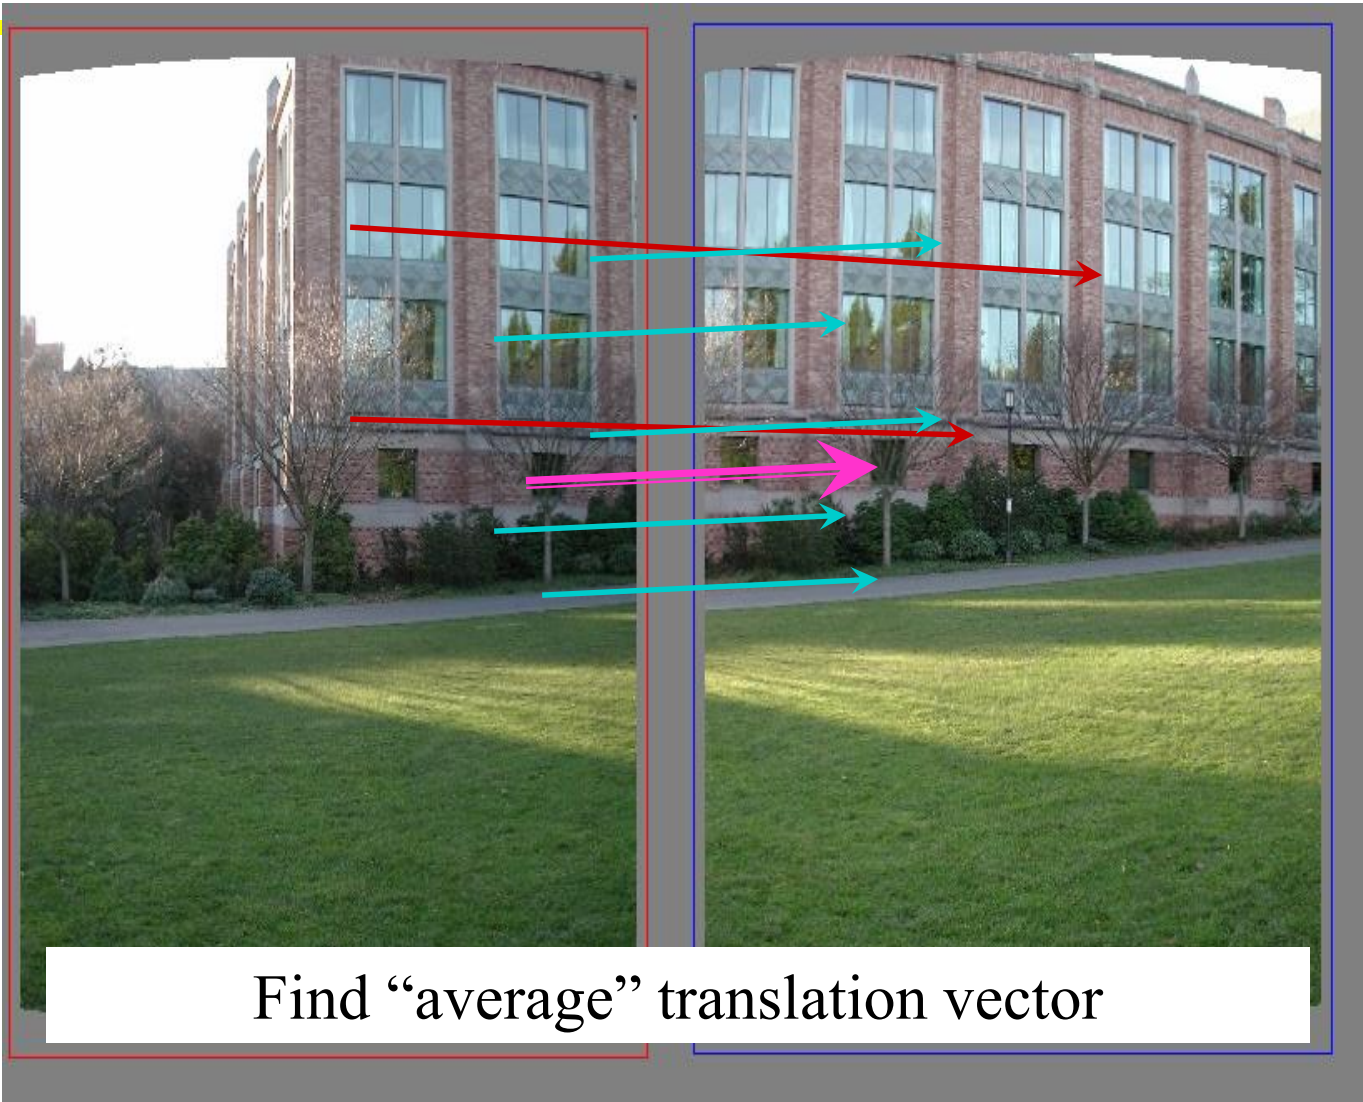




# *RANSAC example: Translation*



# *RANSAC example: Translation*







# *Feature-based alignment outline*




# *Feature-based alignment outline*



- Extract features


# *Feature-based alignment outline*



- Extract features
- Compute *putative matches*




# *Feature-based alignment outline*



- Extract features
- Compute *putative matches*
- Loop:
  - *Hypothesize transformation  $T$*  (small group of putative matches that are related by  $T$ )

# *Feature-based alignment outline*



- Extract features
- Compute *putative matches*
- Loop:
  - ❑ *Hypothesize* transformation  $T$  (small group of putative matches that are related by  $T$ )
  - ❑ *Verify* transformation (search for other matches consistent with  $T$ )

# *Feature-based alignment outline*




- Extract features
- Compute *putative matches*
- Loop:
  - ❑ *Hypothesize* transformation  $T$  (small group of putative matches that are related by  $T$ )
  - ❑ *Verify* transformation (search for other matches consistent with  $T$ )



# *Panoramas*


- ❖ What if you want a 360° field of view?



mosaic Projection Cylinder

# *Cylindrical panoramas*

---



## ❖ Steps

- Project each image onto a cylinder (warp)
- Estimate motion (a pure translation now)
- Blend
- Optional: project it back (unwarp)
- Output the resulting mosaic

# Cylindrical Panoramas

- ❖ Map image to cylindrical or spherical coordinates
  - ❑ need *known* focal length
  - ❑ Work only if a single tilt (e.g., camera on tripod)




Image 384x300



$f = 180$  (pixels)




$f = 280$



# Determining the focal length

1. Initialize from homography  $H$   
(see text or [SzSh'97])
2. Use camera's EXIF tags (approx.)
3. Use a tape measure
4. Try and error 😊



# *Practical Methods for F*

---


- ❖ Use program jhead  
(<http://www.sentex.net/~mwandel/jhead/>)
- ❖ Mac, Windows, and Linux
- ❖ Sample outputs

```
File name      : 0805-153933.jpg
File size     : 463023 bytes
File date     : 2001:08:12 21:02:04
Camera make   : Canon
Camera model  : Canon PowerShot S100
Date/Time    : 2001:08:05 15:39:33
Resolution   : 1600 x 1200
Flash used   : No
Focal length : 5.4mm (35mm equivalent: 36mm)
CCD Width    : 5.23mm
Exposure time: 0.100 s (1/10)
Aperture     : f/2.8
Focus Dist.  : 1.18m
Metering Mode: center weight
Jpeg process : Baseline
```

# Calculating $F$

- ❖ With image resolution (width x height), CCD width and  $f$ 
  - ❑  $f \cdot (\text{width} / \text{CCD width})$  or  $5.4 \cdot (1600 / 5.23) = 1652$  (pixels)
- ❖ With equivalent  $f$  (35mm film is 36mmx24mm)
  - ❑  $(\text{equivalent } f) \cdot (\text{width} / 36)$  or  $36 \cdot (1600 / 36) = 1600$  (pixels)
- ❖ If you don't have the above (more often than not), guess!
  - ❑ No zoom  $f \sim$  (picture width in pixels)
  - ❑ 2x zoom  $f \sim 2 \cdot$  (picture width in pixels)

# Cylindrical projection



Map 3D point  $(X, Y, Z)$  onto cylinder



$$(\hat{x}, \hat{y}, \hat{z}) = \frac{1}{\sqrt{X^2 + Z^2}}(X, Y, Z)$$

- Convert to cylindrical coordinates
- Convert to cylindrical image coordinates

$$(\sin\theta, h, \cos\theta) = (\hat{x}, \hat{y}, \hat{z})$$


$$(\tilde{x}, \tilde{y}) = (s\theta, sh) + (\tilde{x}_c, \tilde{y}_c)$$

– s defines size of the final image



# Cylindrical warping

❖ Given focal length  $f$  and image center  $(x_c, y_c)$



$$\theta = (x_{cyl} - x_c) / f$$

$$h = (y_{cyl} - y_c) / f$$

$$\hat{x} = \sin \theta$$

$$\hat{y} = h$$

$$\hat{z} = \cos \theta$$


$$x = f \hat{x} / \hat{z} + x_c$$

$$y = f \hat{y} / \hat{z} + y_c$$



# Spherical warping

❖ Given focal length  $f$  and image center  $(x_c, y_c)$



$$\theta = (x_{cyl} - x_c) / f$$

$$\varphi = (y_{cyl} - y_c) / f$$

$$\hat{x} = \sin \theta \cos \varphi$$

$$\hat{y} = \sin \varphi$$


$$\hat{z} = \cos \theta \cos \varphi$$

$$x = f \hat{x} / \hat{z} + x_c$$

$$y = f \hat{y} / \hat{z} + y_c$$

# 3D rotation

❖ Rotate image before placing on unrolled sphere



$$\theta = (x_{cyl} - x_c) / f$$

$$\phi = (y_{cyl} - y_c) / f$$

$$\hat{x} = \sin \theta \cos \phi$$

$$\hat{y} = \sin \phi$$


$$\hat{z} = \cos \theta \cos \phi$$

$$x = f \underline{\hat{x}} / \underline{\hat{z}} + x_c$$

$$y = f \underline{\hat{y}} / \underline{\hat{z}} + y_c$$

# Radial distortion

- ❖ Correct for “bending” in wide field of view lenses



$$\hat{r}^2 = \hat{x}^2 + \hat{y}^2$$

$$\hat{x}' = \hat{x} / (1 + \kappa_1 \hat{r}^2 + \kappa_2 \hat{r}^4)$$

$$\hat{y}' = \hat{y} / (1 + \kappa_1 \hat{r}^2 + \kappa_2 \hat{r}^4)$$

$$x = f \hat{x}' / \hat{z} + x_c$$

$$y = f \hat{y}' / \hat{z} + y_c$$

# Fisheye lens

- ❖ Extreme “bending” in ultra-wide fields of view



$$\hat{r}^2 = \hat{x}^2 + \hat{y}^2$$


$$(\cos \theta \sin \phi, \sin \theta \sin \phi, \cos \phi) = s (x, y, z)$$

Equations become

$$x' = s \phi \cos \theta = s \frac{x}{r} \tan^{-1} \frac{r}{z},$$
$$y' = s \phi \sin \theta = s \frac{y}{r} \tan^{-1} \frac{r}{z},$$


# Image Stitching

1. Align the images over each other
  - camera pan  $\leftrightarrow$  translation on cylinder
2. Blend the images together




# *Assembling the panorama*

---



❖ Stitch pairs together, blend, then crop

# *Problem: Drift*




## ❖ Error accumulation

- ❑ small (vertical) errors accumulate over time
- ❑ apply correction so that  $\text{sum} = 0$  (for  $360^\circ$  pan.)



# Problem: Drift



## ❖ Solution

- ❑ add another copy of first image at the end
- ❑ this gives a constraint:  $y_n = y_1$
- ❑ there are a bunch of ways to solve this problem
  - add displacement of  $(y_1 - y_n)/(n - 1)$  to each image after the first
  - compute a global warp:  $y' = y + ax$
  - run a big optimization problem, incorporating this constraint
    - best solution, but more complicated
    - known as “bundle adjustment”

*Full-view (360° spherical)*  
*panoramas*



# Full-view Panorama



+



+



+



+



# *Texture Mapped Model*

---



# *Global alignment*

---

- Register *all* pairwise overlapping images
- Use a 3D rotation model (one  $R$  per image)
- Use direct alignment (patch centers) or feature based
- *Infer* overlaps based on previous matches (incremental)
- Optionally *discover* which images overlap other images using feature selection (RANSAC)



# Bundle adjustment formulations

All pairs optimization: Confidence / uncertainty of point  $i$  in image  $j$

$$E_{\text{all-pairs-2D}} = \sum_i \sum_{jk} c_{ij} c_{ik} \left\| \tilde{\mathbf{x}}_{ik}(\hat{\mathbf{x}}_{ij}; \mathbf{R}_j, f_j, \mathbf{R}_k, f_k) - \hat{\mathbf{x}}_{ik} \right\|^2, \quad (9.29)$$

Map 2D point  $i$  in image  $j$  to 2D point in image  $k$

Full bundle adjustment, using 3-D point positions  $\{\mathbf{x}_i\}$

$$E_{\text{BA-2D}} = \sum_i \sum_j c_{ij} \left\| \tilde{\mathbf{x}}_{ij}(\mathbf{x}_i; \mathbf{R}_j, f_j) - \hat{\mathbf{x}}_{ij} \right\|^2, \quad (9.30)$$

Map 3D point  $i$  in to 2D point in image  $i$

Bundle adjustment using 3-D ray:

$$E_{\text{BA-3D}} = \sum_i \sum_j c_{ij} \left\| \tilde{\mathbf{x}}_i(\hat{\mathbf{x}}_{ij}; \mathbf{R}_j, f_j) - \mathbf{x}_i \right\|^2, \quad (9.31)$$

3-D ray from point  $i$

All-pairs 3-D ray formulation:

$$E_{\text{all-pairs-3D}} = \sum_i \sum_{jk} c_{ij} c_{ik} \left\| \tilde{\mathbf{x}}_i(\hat{\mathbf{x}}_{ij}; \mathbf{R}_j, f_j) - \tilde{\mathbf{x}}_i(\hat{\mathbf{x}}_{ik}; \mathbf{R}_k, f_k) \right\|^2. \quad (9.32)$$

3-D ray from points  $i$  and  $j$

Projected point  $\rightarrow$

$$\tilde{\mathbf{x}}_{ij} \sim \mathbf{K}_j \mathbf{R}_j \mathbf{x}_i \quad \text{and} \quad \mathbf{x}_i \sim \mathbf{R}_j^{-1} \mathbf{K}_j^{-1} \tilde{\mathbf{x}}_{ij}, \quad \leftarrow$$