



Feature-based Alignment

Chapter 6 R. Szelisky

Guido Gerig

CS 6320 Spring 2012

Slide Credits: Trevor Darrell, Berkeley (C280 CV Course), Steve Seitz, Kristen Grauman, Alyosha Efros, L. Lazebnik, Marc Pollefeys

Original Slides Prof. Trevor Darrel (08Alignment, 06LocalFeatures): please visit <http://www.eecs.berkeley.edu/~trevor/CS280Notes/>



Today: Alignment



(a)



(b)



(c)



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: $176^\circ \times 135^\circ$





Motion models

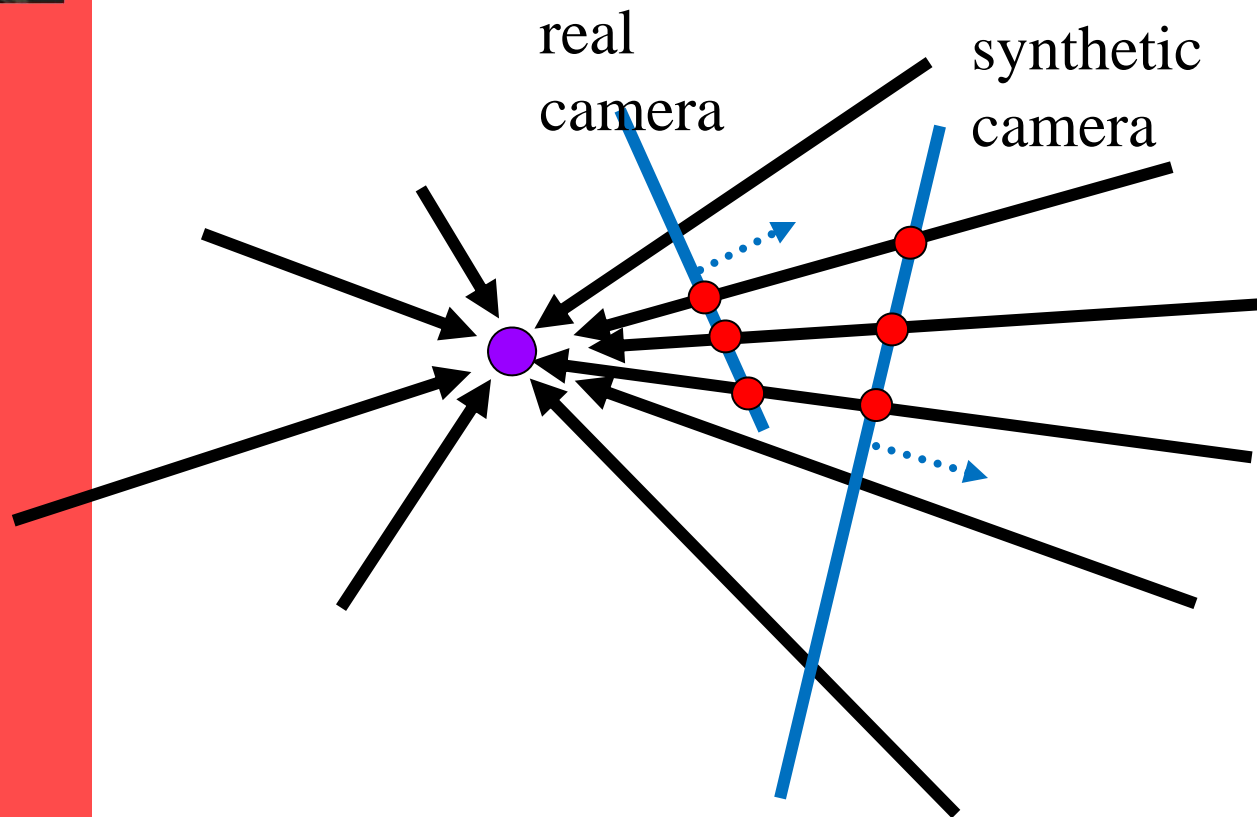


Motion models

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




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*



Motion models



Translation


Affine

Perspective

3D rotation




2 unknowns



6 unknowns



8 unknowns



3 unknowns



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.)



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







Image alignment




- Two broad approaches:
 - Direct (pixel-based) alignment
 - Search for alignment where most pixels agree
 - Feature-based alignment
 - Search for alignment where *extracted features* agree
 - Can be verified using pixel-based alignment



Fitting an affine transformation

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


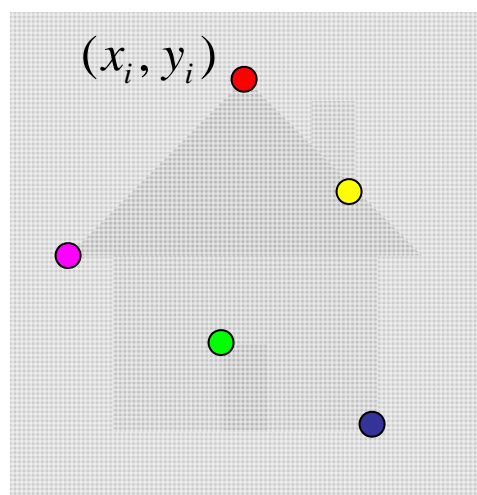


$$\begin{bmatrix} x'_i \\ y'_i \end{bmatrix} = \begin{bmatrix} m_1 & m_2 \\ m_3 & m_4 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$$



Fitting an affine transformation

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



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



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



Fitting an affine transformation

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

- 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

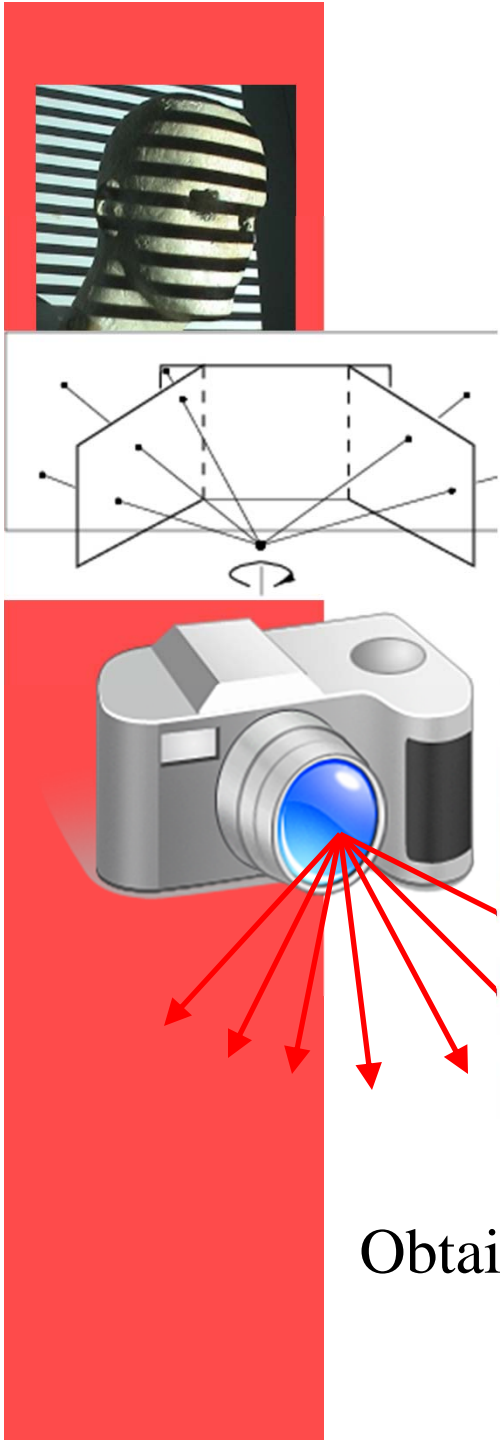




image from S. Seitz

Obtain a wider angle view by combining multiple images.



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

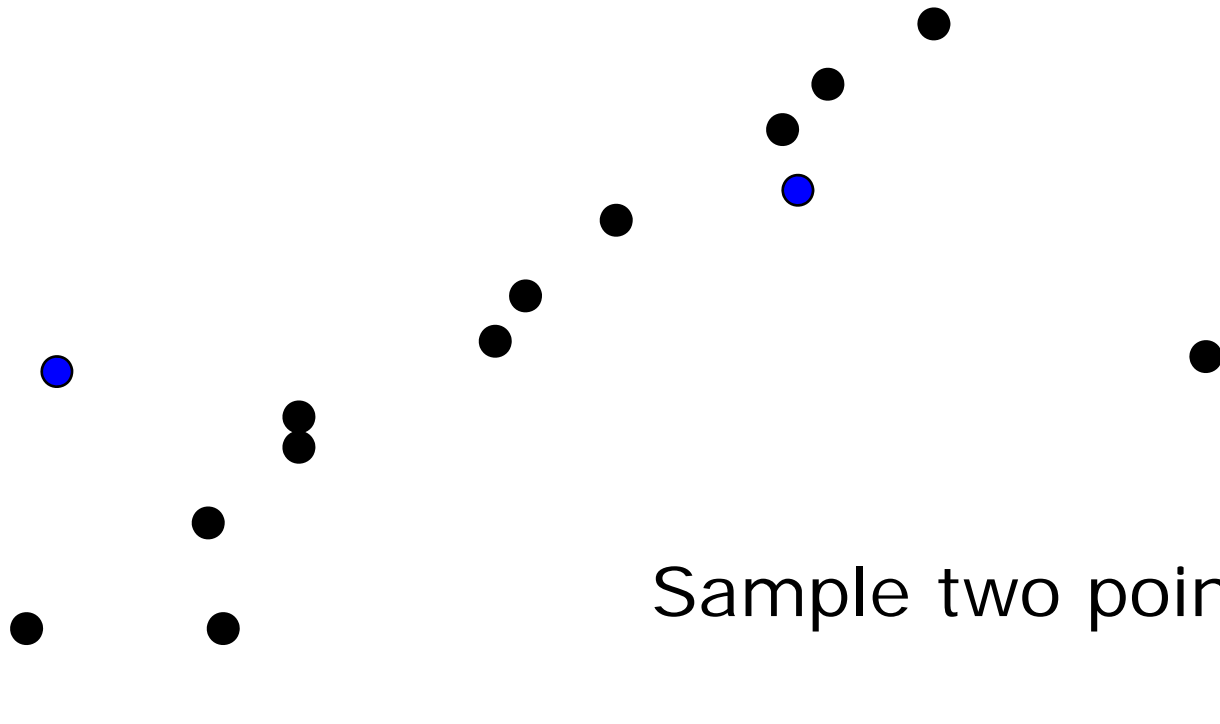


RANSAC Line Fitting Example



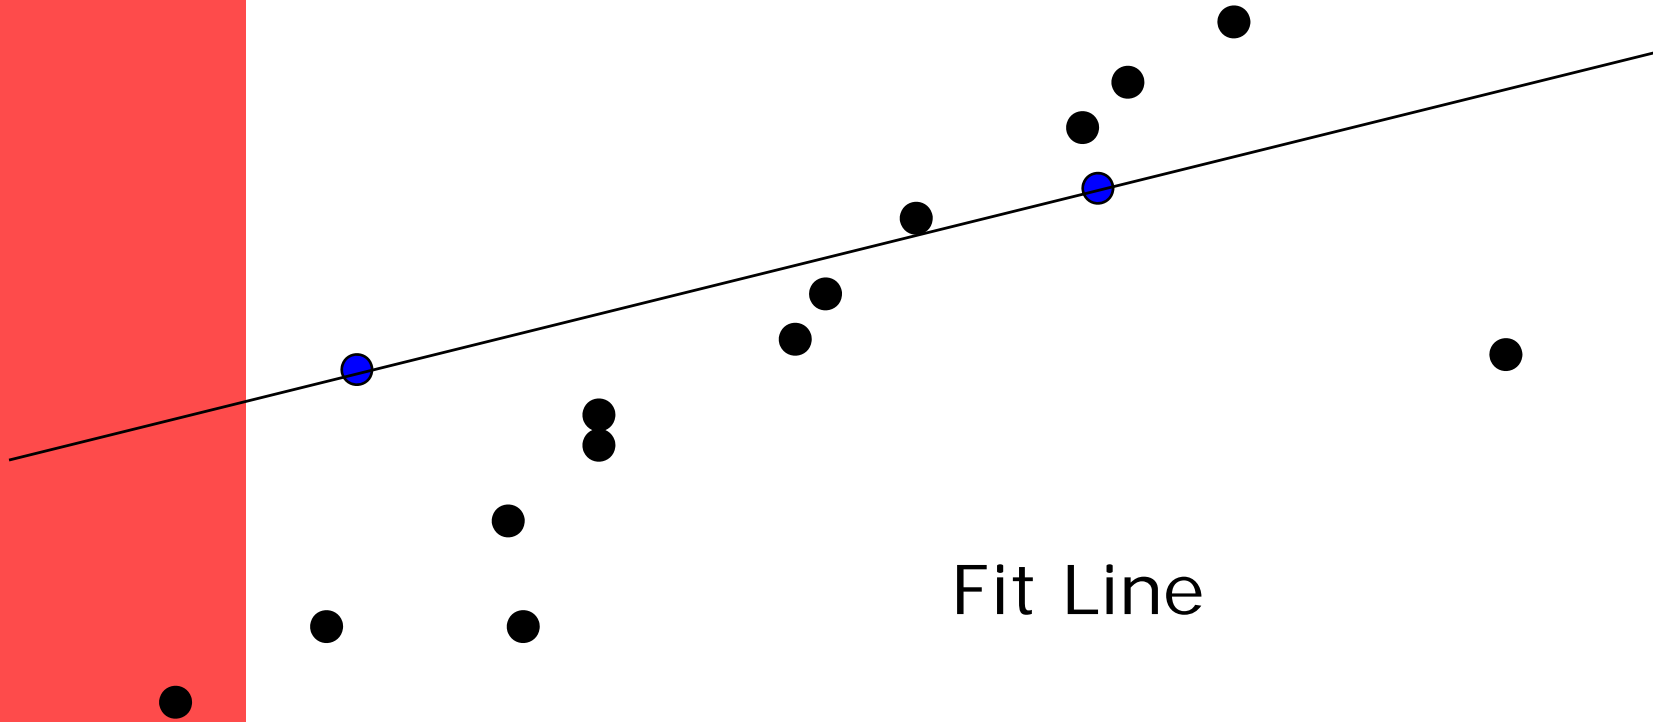


RANSAC Line Fitting Example



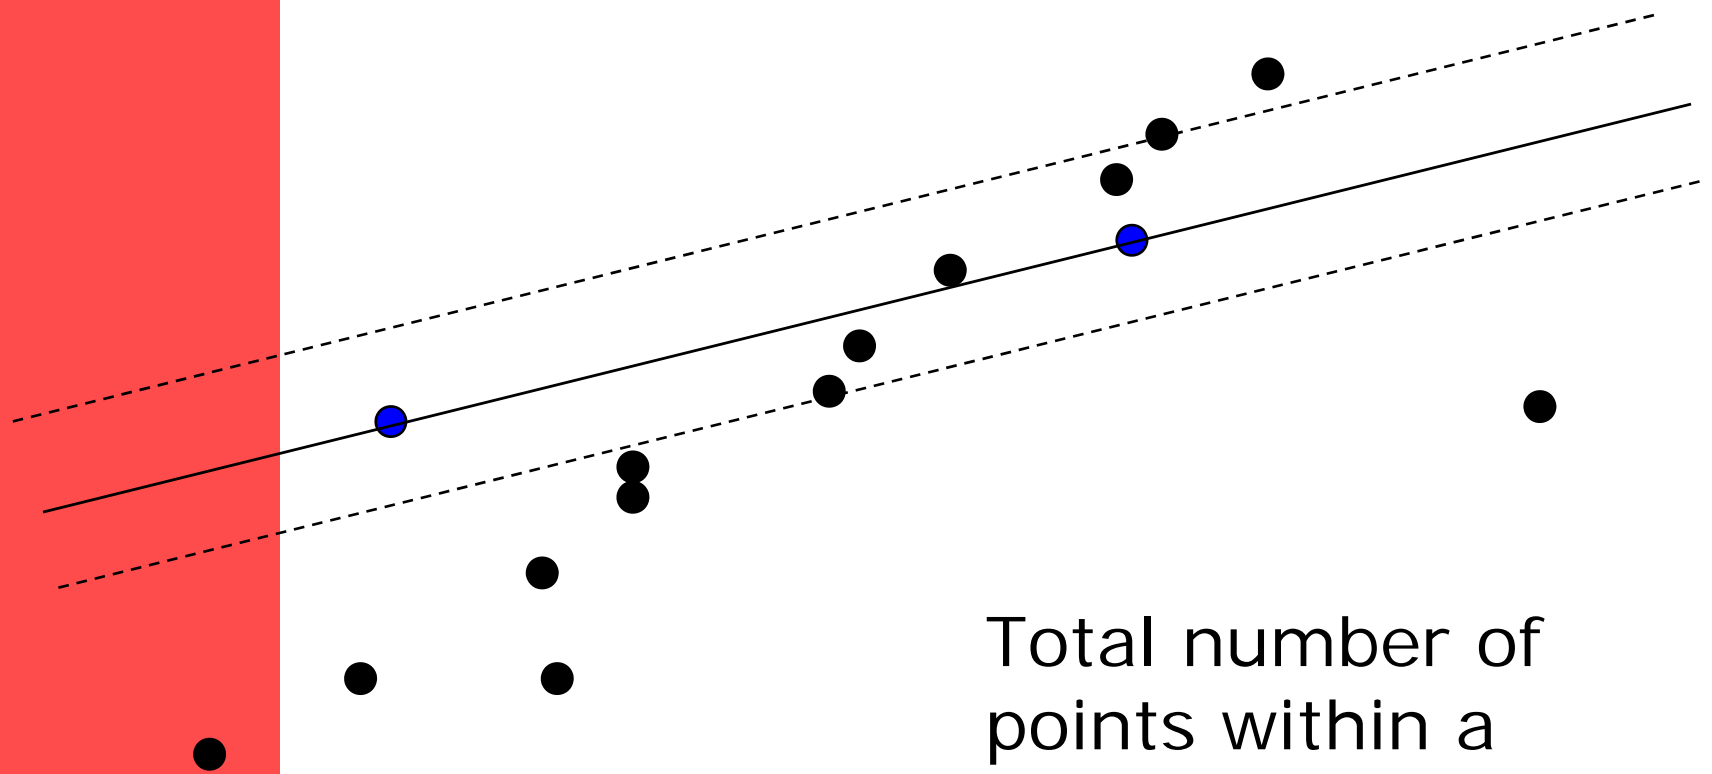


RANSAC Line Fitting Example





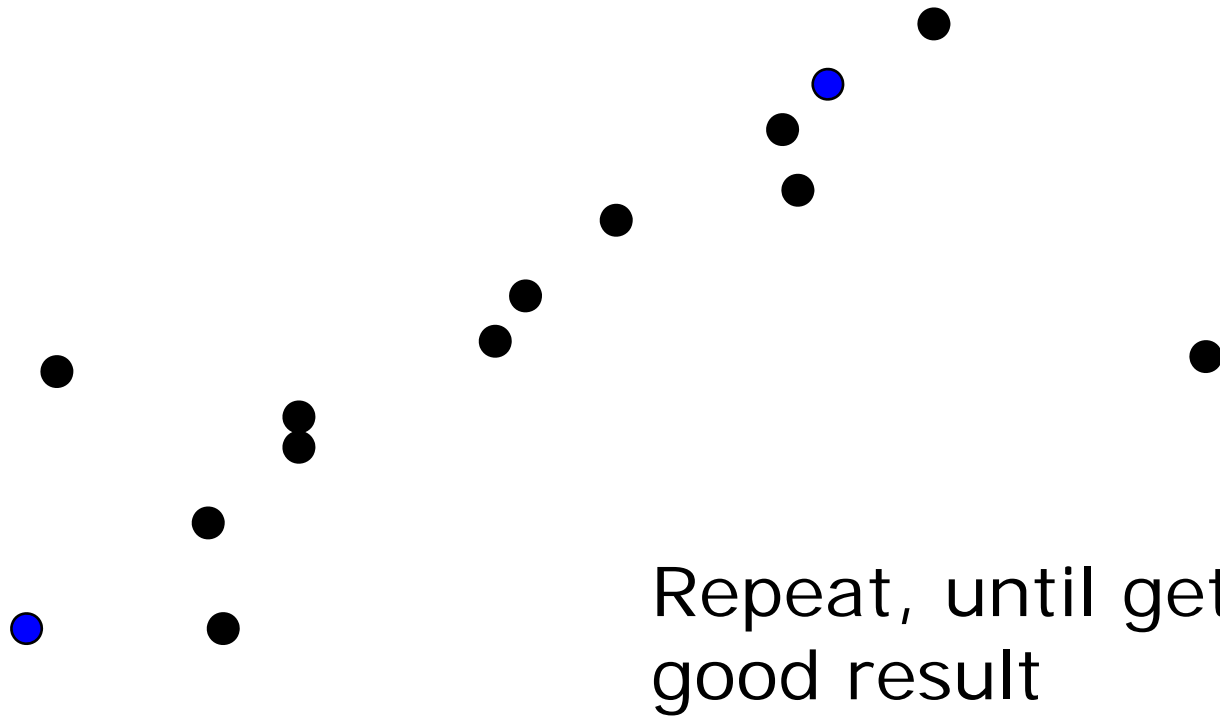
RANSAC Line Fitting Example



Total number of
points within a
threshold of line.




RANSAC Line Fitting Example






RANSAC Line Fitting Example



Repeat, until get a good result




RANSAC Line Fitting Example



Repeat, until get a good result




RANSAC example: Translation






RANSAC example: Translation



Select *one* match, count *inliers*




RANSAC example: Translation



Select *one* match, count *inliers*



RANSAC example: Translation




Find “average” translation vector

Towards large-scale mosaics...

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)

RANSAC motion model





Towards large-scale mosaics...

Probabilistic Feature Matching



Probabilistic model for verification





Other types of mosaics



- Can mosaic onto *any* surface if you know the geometry
 - See NASA's [Visible Earth project](#) for some stunning earth mosaics
 - <http://earthobservatory.nasa.gov/Newsroom/BlueMarble/>

Final thought: What is a "panorama"?



- Tracking a subject
- Repeated (best) shots
- Multiple exposures
- "Infer" what photographer wants?





Next time: 6.2 Pose Estimation

- 6.2 Pose Estimation
- Chapter 7: Structure from Motion