


Lecture 12: Transformations and Fitting II

Administrative

HW2 Due Tomorrow, 2/19 at 11:59pm

HW3 is released, due a week from
Friday, 2/28 at 11:59pm

So Far



1. How do we find distinctive / easy to locate features? (*Harris/Laplacian of Gaussian*)
2. How do we describe the regions around them? (*histogram of gradients*)
3. How do we match features? (L2 distance)
4. How do we handle outliers? (RANSAC)

Today

As promised: warping one image to another



Why Mosaic?

- Compact Camera FOV = $50 \times 35^\circ$



Slide credit: Brown & Lowe

Why Mosaic?

- Compact Camera FOV = $50 \times 35^\circ$
- Human FOV = $200 \times 135^\circ$



Slide credit: Brown & Lowe

Why Mosaic?

- Compact Camera FOV = $50 \times 35^\circ$
- Human FOV = $200 \times 135^\circ$
- Panoramic Mosaic = $360 \times 180^\circ$



Slide credit: Brown & Lowe

Why Bother with the Math?





Slide credit: A. Efros

Homework 1 Style



Translation only via alignment



Slide credit: A. Efros

More Sophisticated Result



Slide credit: A. Efros

Today

Categories of Transformations

Fitting Transformations

Applying Transformations

Blending Images

Today

Categories of Transformations

Fitting Transformations

Applying Transformations

Blending Images

Image Transformations

Image filtering: change range of image

$$g(x) = T(f(x))$$





Image warping: change **domain** of image

$$g(x) = f(T(x))$$



Slide credit: A. Efros

Image Transformations

Image filtering: change range of image

$$g(x) = T(f(x))$$





Image warping: change **domain** of image

$$g(x) = f(T(x))$$



Slide credit: A. Efros

Parametric (Global) Warping

Examples of parametric warps



translation



rotation



aspect



affine



perspective



cylindrical

Slide credit: A. Efros

Parametric (Global) Warping


T is a coordinate changing machine

$$\mathbf{p}' = T(\mathbf{p})$$

Note: T is the same for all points, has relatively few parameters, and does **not** depend on image content



$$\mathbf{p} = (x, y)$$



$$\mathbf{p}' = (x', y')$$

Slide credit: A. Efros

Parametric (Global) Warping


Today we'll deal with linear warps

$$\mathbf{p}' \equiv T\mathbf{p}$$

T: matrix; \mathbf{p} , \mathbf{p}' : 2D points. Start with normal points and $=$, then do homogeneous coords and \equiv



$$\mathbf{p} = (x, y)$$




$$\mathbf{p}' = (x', y')$$

Slide credit: A. Efros

Scaling

Scaling multiplies each component (x,y) by a scalar.
Uniform scaling is the same for all components.


Note the corner goes from $(1,1)$ to $(2,2)$



Slide credit: A. Efros

Scaling

Non-uniform scaling multiplies each component by a different scalar.



Slide credit: A. Efros

Scaling

What does T look like?

$$x' = ax$$

$$y' = by$$

Let's convert to a matrix:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \underbrace{\begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}} \begin{bmatrix} x \\ y \end{bmatrix}$$

scaling matrix S

What's the inverse of S?

2D Rotation



Rotation Matrix

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

But wait! Aren't sin/cos non-linear?

x' is a linear combination/function of x, y

x' is not a linear function of θ

What's the inverse of R_θ ? $I = R_\theta^T R_\theta$

Things you can do with 2x2

Identity / No Transformation



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


Shear



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


Slide credit: A. Efros

Things you can do with 2x2



2D Mirror About Y-Axis

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



2D Mirror About X,Y

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

Slide credit: A. Efros


What is Preserved?



3D lines project to 2D lines so lines are preserved

Projections of parallel 3D lines are not necessarily parallel, so not parallelism

Distant objects are smaller so size is not preserved



2x2: What is Preserved

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = T \begin{bmatrix} x \\ y \end{bmatrix}$$

After multiplication by T (irrespective of T)


- Origin is origin: **$\mathbf{0} = T\mathbf{0}$**
 - Lines are lines
- Parallel lines are parallel

Things You Can't Do With 2x2

What about translation?

$$x' = x + t_x, y' = y + t_y$$

How do we fix it?




Homogenous Coordinates Again

What about translation?

$$x' = x + t_x, y' = y + t_y$$

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



Slide credit: A. Efros

Representing 2D Transformations

How do we represent a 2D transformation?

Let's pick scaling

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} \equiv \begin{bmatrix} s_x & 0 & a \\ 0 & s_y & b \\ d & e & f \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

What's

a	b	d	e	f
0	0	0	0	1

Affine Transformations

Affine: *linear transformation plus translation*



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

Will the last coordinate always be 1?

In general (without homogeneous coordinates)

$$\mathbf{x}' = \mathbf{Ax} + \mathbf{b}$$

Composing Transforms

We can combine transformations via matrix multiplication.

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} \equiv \underbrace{\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}}_{T(t_x, t_y)} \underbrace{\begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{R(\theta)} \underbrace{\begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{S(s_x, s_y)} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

Does order matter?

Affine: What is Preserved

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} \equiv \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \equiv \mathbf{T} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

After multiplication by T (irrespective of T)

- ~~• Origin is origin: $0 = T0$~~
- Lines are lines
- Parallel lines are parallel

Perspective Transformations

Set bottom row to not $[0,0,1]$

Called a perspective/projective transformation or a *homography*



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

How many degrees of freedom?

How Many Degrees of Freedom?

Recall: can always scale by non-zero value

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

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} \equiv \frac{1}{i} \begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} \equiv \frac{1}{i} \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix} \equiv \begin{bmatrix} a/i & b/i & c/i \\ d/i & e/i & f/i \\ g/i & h/i & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

Homography can always be re-scaled by $\lambda \neq 0$

Perspective: What is Preserved


$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} \equiv \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \equiv \mathbf{T} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

After multiplication by T (irrespective of T)

- ~~• Origin is origin: $0 = T0$~~
- Lines are lines
- ~~• Parallel lines are parallel~~
- ~~• Ratios between distances~~

Transformation Families

In general: transformations are a nested set of groups



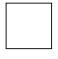

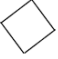
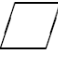

Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	2	orientation + ...	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	3	lengths + ...	
similarity	$\begin{bmatrix} s\mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	4	angles + ...	
affine	$\begin{bmatrix} \mathbf{A} \end{bmatrix}_{2 \times 3}$	6	parallelism + ...	
projective	$\begin{bmatrix} \tilde{\mathbf{H}} \end{bmatrix}_{3 \times 3}$	8	straight lines	

Diagram credit: R. Szeliski

What Can Homography Do?

Homography example 1: any two views
of a *planar* surface



Figure Credit: S. Lazebnik

What Can Homography Do?

Homography example 2: any images from two cameras sharing a camera center





Figure Credit: S. Lazebnik

What Can Homography Do?



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

What Can Homography Do?

Homography sort of example “3”: far away scene that can be approximated by a plane




Figure credit: Brown & Lowe

Fun With Homography

Original image

St. Petersburg
photo by A. Tikhonov




Virtual camera rotations



Slide Credit: A. Efros

Fun With Homography




The floor (enlarged)

Slide from A. Criminisi



**Automatically
rectified floor**

Fun With Homography




From Martin Kemp *The Science of Art*
(*manual reconstruction*)

Slide from A. Criminisi

Automatic rectification



Fun With Homography



St. Lucy Altarpiece, D. Veneziano

Slide from A. Criminisi

What is the (complicated) shape of the floor pattern?




Automatically rectified floor

Fun With Homography



**Automatic
rectification**



**From Martin Kemp, *The Science of Art*
(*manual reconstruction*)**

Slide from A. Criminisi

Today

Categories of Transformations

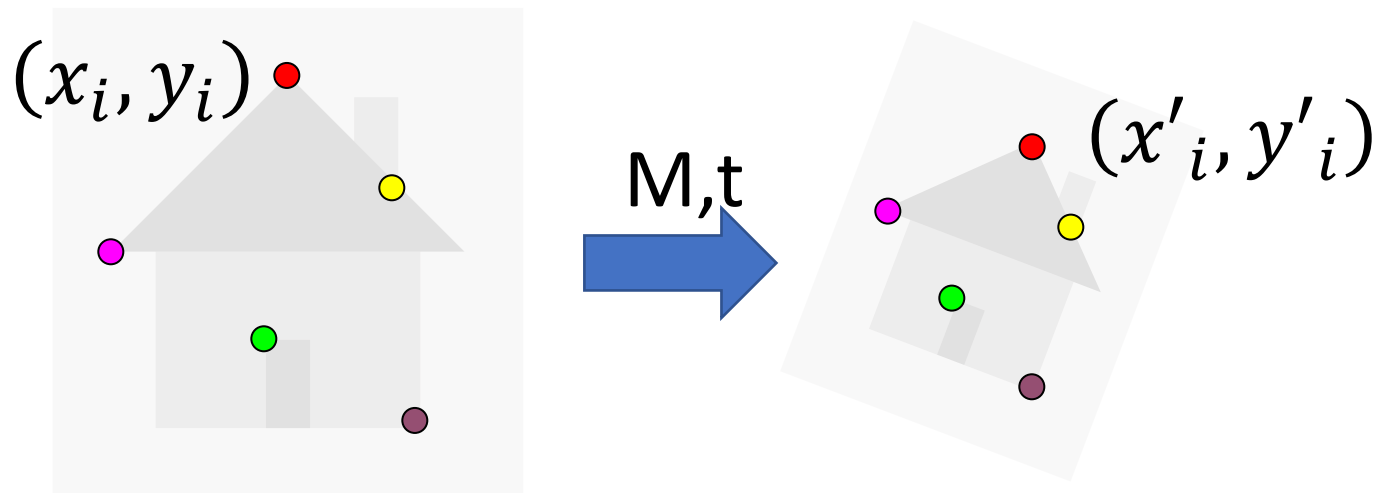
Fitting Transformations

Applying Transformations

Blending Images

Fitting Transformations

Setup: have pairs of correspondences



$$\begin{bmatrix} x'_i \\ y'_i \end{bmatrix} = \mathbf{M} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \mathbf{t}$$

Slide Credit: S. Lazebnik


Fitting Transformations: Affine

Affine Transformation: M, t

Data: (x_i, y_i, x'_i, y'_i) for
 $i=1, \dots, k$

Model:
 $[x'_i, y'_i] = \mathbf{M}[x_i, y_i] + \mathbf{t}$

Objective function:
 $\| [x'_i, y'_i] - \mathbf{M}[x_i, y_i] + \mathbf{t} \|^2$



Fitting Transformations: Affine

Given correspondences: $\mathbf{p}' = [x'_i, y'_i]$, $\mathbf{p} = [x_i, y_i]$

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

Set up two equations per point

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

Solve with least squares!

Fitting Transformations: Affine

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

2 equations per point, 6 unknowns

How many points do we need?

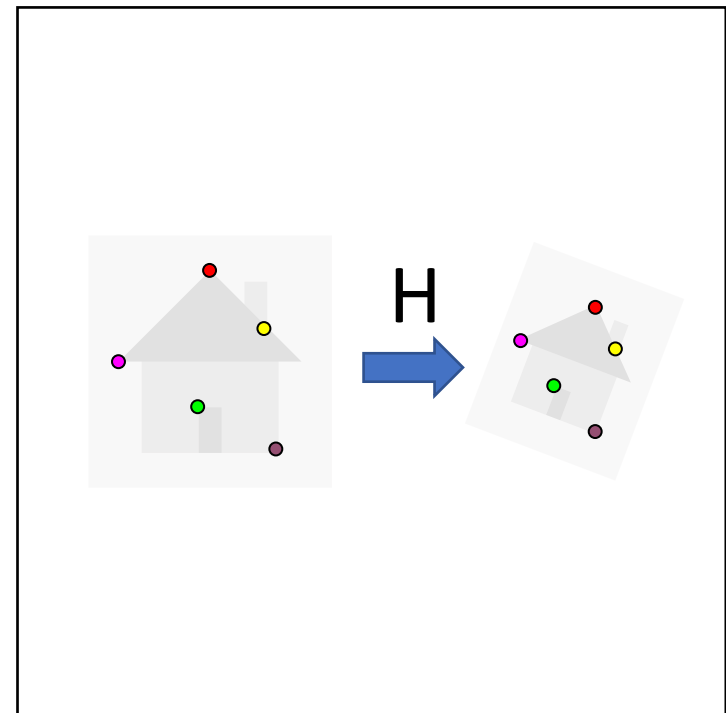
Fitting Transformations: Homography

Homography: H

Data: (x_i, y_i, x'_i, y'_i) for
 $i=1, \dots, k$

Model:
 $[x'_i, y'_i, 1] \equiv \mathbf{H}[x_i, y_i, 1]$

Objective function:
It's complicated



Fitting Transformations: Homography

Want:
$$\begin{bmatrix} x'_i \\ y'_i \\ w'_i \end{bmatrix} \equiv \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} x_i \\ y_i \\ w_i \end{bmatrix}$$

$$\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} x_i \\ y_i \\ w_i \end{bmatrix} \equiv \mathbf{H} \mathbf{x}_i \equiv \begin{bmatrix} \mathbf{h}_1^T \\ \mathbf{h}_2^T \\ \mathbf{h}_3^T \end{bmatrix} \mathbf{x}_i \equiv \begin{bmatrix} \mathbf{h}_1^T \mathbf{x}_i \\ \mathbf{h}_2^T \mathbf{x}_i \\ \mathbf{h}_3^T \mathbf{x}_i \end{bmatrix}$$

Recall: $\mathbf{a} \equiv \mathbf{b} \rightarrow \mathbf{a} = \lambda \mathbf{b} \rightarrow \mathbf{a} \times \mathbf{b} = \mathbf{0}$

Fitting Transformations: Homography

Want:

$$\begin{bmatrix} x'_i \\ y'_i \\ w'_i \end{bmatrix} \equiv \begin{bmatrix} \mathbf{h}_1^T \mathbf{x}_i \\ \mathbf{h}_2^T \mathbf{x}_i \\ \mathbf{h}_3^T \mathbf{x}_i \end{bmatrix} \iff \begin{bmatrix} x'_i \\ y'_i \\ w'_i \end{bmatrix} \times \begin{bmatrix} \mathbf{h}_1^T \mathbf{x}_i \\ \mathbf{h}_2^T \mathbf{x}_i \\ \mathbf{h}_3^T \mathbf{x}_i \end{bmatrix} = \mathbf{0}$$

Cross-product

$$\begin{bmatrix} y'_i \mathbf{h}_3^T \mathbf{x}_i - w'_i \mathbf{h}_2^T \mathbf{x}_i \\ w'_i \mathbf{h}_1^T \mathbf{x}_i - x'_i \mathbf{h}_3^T \mathbf{x}_i \\ x'_i \mathbf{h}_2^T \mathbf{x}_i - y'_i \mathbf{h}_1^T \mathbf{x}_i \end{bmatrix} = \mathbf{0}$$

Re-arrange
and put 0s in

$$\begin{bmatrix} \mathbf{h}_1^T \mathbf{0} - w'_i \mathbf{h}_2^T \mathbf{x}_i + y'_i \mathbf{h}_3^T \mathbf{x}_i \\ w'_i \mathbf{h}_1^T \mathbf{x}_i + \mathbf{h}_2^T \mathbf{0} - x'_i \mathbf{h}_3^T \mathbf{x}_i \\ -y'_i \mathbf{h}_1^T \mathbf{x}_i + x'_i \mathbf{h}_2^T \mathbf{x}_i + \mathbf{h}_3^T \mathbf{0} \end{bmatrix} = \mathbf{0}$$

Fitting Transformations: Homography

Equation

$$\begin{bmatrix} \mathbf{h}_1^T \mathbf{0} - w'_i \mathbf{h}_2^T \mathbf{x}_i + y'_i \mathbf{h}_3^T \mathbf{x}_i \\ w'_i \mathbf{h}_1^T \mathbf{x}_i + \mathbf{h}_2^T \mathbf{0} - x'_i \mathbf{h}_3^T \mathbf{x}_i \\ -y'_i \mathbf{h}_1^T \mathbf{x}_i + x'_i \mathbf{h}_2^T \mathbf{x}_i + \mathbf{h}_3^T \mathbf{0} \end{bmatrix} = \mathbf{0}$$

Pull out h

$$\begin{bmatrix} \mathbf{0}^T & -w'_i \mathbf{x}_i^T & y'_i \mathbf{x}_i^T \\ w'_i \mathbf{x}_i^T & \mathbf{0}^T & -x'_i \mathbf{x}_i^T \\ -y'_i \mathbf{x}_i^T & x'_i \mathbf{x}_i^T & \mathbf{0}^T \end{bmatrix} \begin{bmatrix} \mathbf{h}_1 \\ \mathbf{h}_2 \\ \mathbf{h}_3 \end{bmatrix} = \mathbf{0}$$

Only two linearly independent equations

$$\frac{x'_i}{w'_i} [0 \quad -w'_i \quad y'_i] + \frac{y'_i}{w'_i} [w'_i \quad 0 \quad -x'_i] + [-y'_i \quad x'_i \quad 0] = \mathbf{0}$$

Fitting Transformations: Homography

N points \rightarrow

$$\begin{array}{c}
 \leftarrow \boxed{9} \rightarrow \\
 \uparrow \boxed{2n} \downarrow \\
 \begin{bmatrix}
 \mathbf{0}^T & -w'_1 \mathbf{x}_1^T & y'_1 \mathbf{x}_1^T \\
 w'_1 \mathbf{x}_1^T & \mathbf{0}^T & -x'_1 \mathbf{x}_1^T \\
 \vdots & & \\
 \mathbf{0}^T & -w'_n \mathbf{x}_n^T & y'_n \mathbf{x}_n^T \\
 w'_n \mathbf{x}_n^T & \mathbf{0}^T & -x'_n \mathbf{x}_n^T
 \end{bmatrix}
 \begin{bmatrix}
 h_1 \\
 h_2 \\
 h_3
 \end{bmatrix}
 = \mathbf{0}
 \end{array}$$

$Ah = \mathbf{0}$

If h is up to scale, what do we use from last time?

$$h^* = \arg \min_{\|h\|=1} \|Ah\|^2 \rightarrow \text{Eigenvector of } A^T A \text{ with smallest eigenvalue}$$

Fitting Transforms: Small Detail

$\|Ah\|^2$ doesn't measure model fit (it's called an *algebraic error* that's mainly just convenient to minimize)

Really want *geometric error*:

$$\sum_{i=1}^k \|[x'_i, y'_i] - T([x_i, y_i])\|^2 + \|[x_i, y_i] - T^{-1}([x'_i, y'_i])\|^2$$

Fitting Transformations: Small Detail

Solution: initialize with algebraic ($\min ||Ah||$), optimize with geometric using standard non-linear optimizer

**In RANSAC, we always take just enough points to fit.
Why might this not make a big difference when fitting
a model with RANSAC?**

Today


Categories of Transformations

Fitting Transformations

Applying Transformations


Blending Images

Image Warping




Given a coordinate transform $(x', y') = T(x, y)$ and a source image $f(x, y)$, how do we compute a transformed image $g(x', y') = f(T(x, y))$?

Forward Warping





Send the value at each pixel (x, y) to
the new pixel $(x', y') = T([x, y])$

Forward Warping




If you don't hit an exact pixel, give the value to each of the neighboring pixels ("splatting").

Forward Warping



Suppose $T(x,y)$ scales by a factor of 3.
HmMMM.


Backward Warping



Find out where each pixel $g(x', y')$ should get its value from, and steal it.

Note: requires ability to invert T

Backward Warping



If you don't hit an exact pixel, figure out how to take it from the neighbors.

Today

Categories of Transformations

Fitting Transformations

Applying Transformations

Blending Images

Blending Images


Warped
Input 1

I_1




Warped
Input 2

I_2



α



$\alpha I_1 +$
 $(1-\alpha)I_2$



Slide Credit: A. Efros

Simple Approach: Two-Band Blending

- Brown & Lowe, 2003
 - Break up each image into high frequency + low frequency
 - Linearly blend low-frequency information
 - No blending for high-frequency: at each pixel take from one image or the other



Figure Credit: Brown & Lowe

Simple Approach: Two-Band Blending




Low frequency ($\lambda > 2$ pixels)



High frequency ($\lambda < 2$ pixels)

Linear Blending



2-band Blending



Today

Categories of Transformations

Fitting Transformations

Applying Transformations

Blending Images

Putting It All Together

How do you make a panorama?

Step 1: Find “features” to match

Step 2: Describe Features


Step 3: Match by Nearest Neighbor

Step 4: Fit H via RANSAC

Step 5: Blend Images

Putting It All Together: Step 1


Find corners/blobs




- (Multi-scale) Harris; or
- Laplacian of Gaussian

Putting It All Together: Step 2

Describe Regions Near Features




$$x_q \in R^{128}$$



Build histogram of
gradient orientations
(SIFT)

Putting It All Together: Step 3

Match Features Based On Region




Sort by distance to: x_q $\|x_q - x_1\| < \|x_q - x_2\| < \|x_q - x_3\|$

Accept match if: $\|x_q - x_1\| / \|x_q - x_2\|$

Nearest neighbor is far closer than 2nd nearest neighbor

Putting It All Together: Step 4

Fit transformation H via RANSAC



for trial in range(Ntrials):

Pick sample

Fit model

Check if more inliers

Re-fit model with most inliers

$$\arg \min_{\|h\|=1} \|Ah\|^2$$



Putting It All Together: Step 5

Warp images together



Resample images with inverse warping
and blend

So far:
Filtering and Matching

Next up:
Recognition
Linear Models
Neural Networks