

CAP 4453

Robot Vision

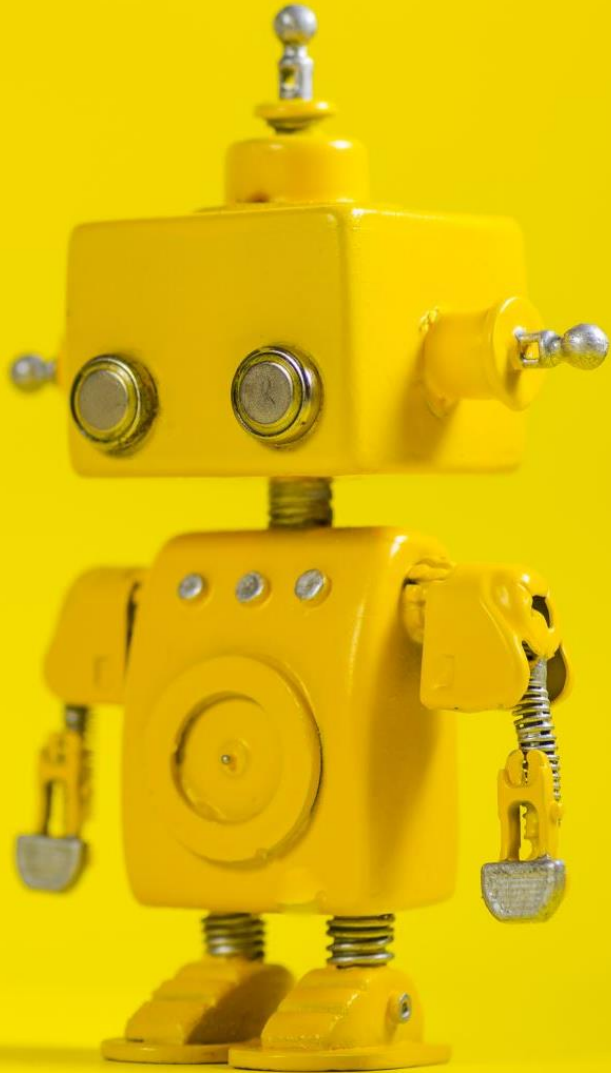
Dr. Gonzalo Vaca-Castaño
gonzalo.vacacastano@ucf.edu




Credits




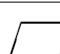

- Some slides comes directly from these sources:
 - Ioannis (Yannis) Gkioulekas (CMU)
 - Noah Snavely (Cornell)
 - Marco Zuliani

Short Review from last class



2D image transformations



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	

These transformations are a nested set of groups

- Closed under composition and inverse is a member



Projective transformations (aka homographies)

Projective transformations are combinations of


- affine transformations; and
- projective wraps


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

Properties of projective transformations:

- origin does not necessarily map to origin
- lines map to lines
- parallel lines do not necessarily map to parallel lines
- ratios are not necessarily preserved
- compositions of projective transforms are also projective transforms

8 DOF: vectors (and therefore matrices) are defined up to scale)





Robot Vision

10. Image warping II




Outline

- Linear algebra
- Image transformations
- 2D transformations.
- Projective geometry 101.
- Transformations in projective geometry.
- Classification of 2D transformations.
- **Determining unknown 2D transformations.**
- Determining unknown image warps.

Determining unknown transformations



Suppose we have two triangles: ABC and DEF.



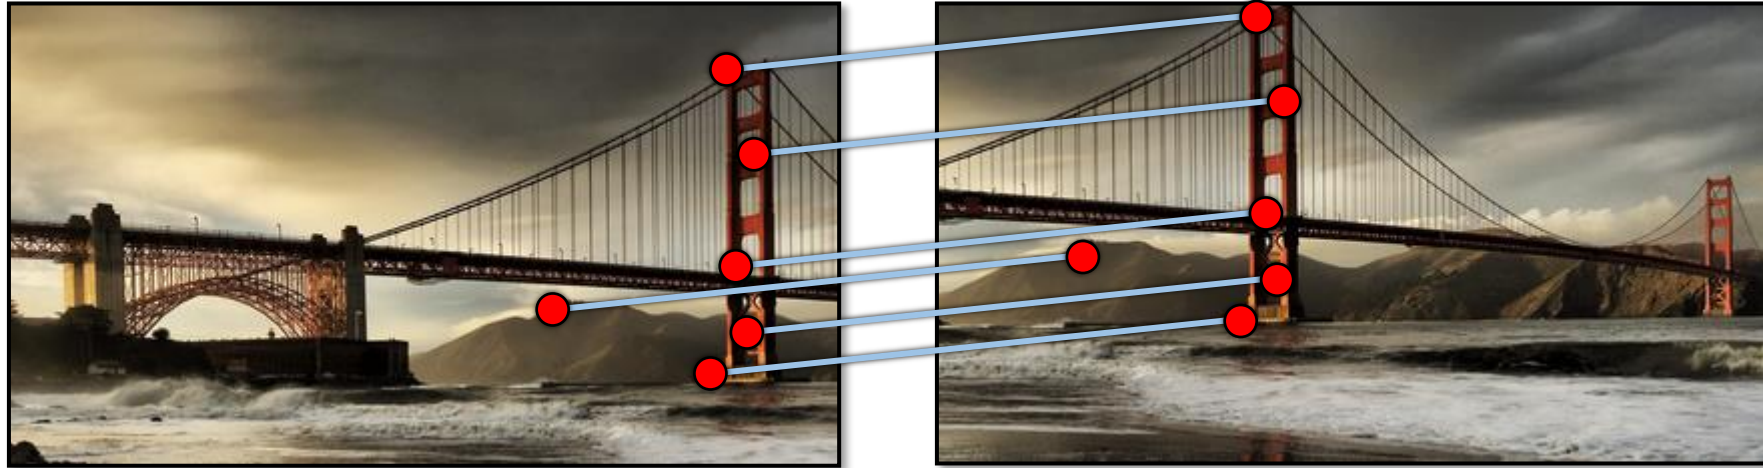
Determining unknown transformations

Suppose we have two triangles: ABC and DEF.

- What type of transformation will map A to D, B to E, and C to F?




Simple case: translations



(x_t, y_t)

How do we solve for (x_t, y_t) ?


Simple case: translations



Displacement of match $i = (\mathbf{x}'_i - \mathbf{x}_i, \mathbf{y}'_i - \mathbf{y}_i)$

$$(\mathbf{x}_t, \mathbf{y}_t) = \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}'_i - \mathbf{x}_i, \frac{1}{n} \sum_{i=1}^n \mathbf{y}'_i - \mathbf{y}_i \right)$$

Another view




$$\mathbf{x}_i + \mathbf{x}_t = \mathbf{x}'_i$$

$$\mathbf{y}_i + \mathbf{y}_t = \mathbf{y}'_i$$

- System of linear equations
 - What are the knowns? Unknowns?
 - How many unknowns? How many equations (per match)?

Another view



$$\mathbf{x}_i + \mathbf{x}_t = \mathbf{x}'_i$$

$$\mathbf{y}_i + \mathbf{y}_t = \mathbf{y}'_i$$

- Problem: more equations than unknowns
 - “Overdetermined” system of equations
 - We will find the *least squares* solution



Least squares formulation

- For each point $(\mathbf{x}_i, \mathbf{y}_i)$

$$\mathbf{x}_i + \mathbf{x}_t = \mathbf{x}'_i$$

$$\mathbf{y}_i + \mathbf{y}_t = \mathbf{y}'_i$$

- we define the *residuals* as

$$r_{\mathbf{x}_i}(\mathbf{x}_t) = (\mathbf{x}_i + \mathbf{x}_t) - \mathbf{x}'_i$$

$$r_{\mathbf{y}_i}(\mathbf{y}_t) = (\mathbf{y}_i + \mathbf{y}_t) - \mathbf{y}'_i$$



Least squares formulation

- Goal: minimize sum of squared residuals

$$C(\mathbf{x}_t, \mathbf{y}_t) = \sum_{i=1}^n \left(r_{\mathbf{x}_i}(\mathbf{x}_t)^2 + r_{\mathbf{y}_i}(\mathbf{y}_t)^2 \right)$$

- “Least squares” solution
- For translations, is equal to mean (average) displacement



Least squares formulation

- Can also write as a matrix equation

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \\ \vdots & \vdots \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} x'_1 - x_1 \\ y'_1 - y_1 \\ x'_2 - x_2 \\ y'_2 - y_2 \\ \vdots \\ x'_n - x_n \\ y'_n - y_n \end{bmatrix}$$

$$\mathbf{A}$$

$2n \times 2$

$$\mathbf{t}$$

2×1

$$=$$

$$\mathbf{b}$$

$2n \times 1$



Least squares

$$\mathbf{A}\mathbf{t} = \mathbf{b}$$

- Find \mathbf{t} that minimizes

$$\|\mathbf{A}\mathbf{t} - \mathbf{b}\|^2$$

- To solve, form the *normal equations*

$$\mathbf{A}^T \mathbf{A} \mathbf{t} = \mathbf{A}^T \mathbf{b}$$

$$\mathbf{t} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$$



Solving the linear system

Convert the system to a linear least-squares problem:

$$E_{LLS} = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$$

Expand the error:

$$E_{LLS} = \mathbf{x}^\top (\mathbf{A}^\top \mathbf{A}) \mathbf{x} - 2\mathbf{x}^\top (\mathbf{A}^\top \mathbf{b}) + \|\mathbf{b}\|^2$$

Minimize the error:


Set derivative to 0 $(\mathbf{A}^\top \mathbf{A})\mathbf{x} = \mathbf{A}^\top \mathbf{b}$

Solve for \mathbf{x} $\mathbf{x} = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{b}$ ←

In Python:


```
1 import numpy as np
2 x, resid, rank, s = np.linalg.lstsq(A, b)
3 x
```

Note: You almost never want to compute the inverse of a matrix.



Least Squares Error

$$E_{\text{LS}} = \sum_i \|\mathbf{f}(\mathbf{x}_i; \mathbf{p}) - \mathbf{x}'_i\|^2$$




Least Squares Error

$$E_{\text{LS}} = \sum_i \left\| \mathbf{f}(\mathbf{x}_i; \mathbf{p}) - \mathbf{x}'_i \right\|^2$$

What is this?

What is this?

What is this?



$$\|x\| = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2}$$

Least Squares Error


$$E_{\text{LS}} = \sum_i \left\| \mathbf{f}(\mathbf{x}_i; \mathbf{p}) - \mathbf{x}'_i \right\|^2$$

Euclidean
(L2) norm

squared!


predicted
location

measured
location



Least Squares Error

$$E_{\text{LS}} = \sum_i \underbrace{\|f(\mathbf{x}_i; \mathbf{p}) - \mathbf{x}'_i\|}_{\text{Residual (projection error)}}^2$$




Least Squares Error

$$E_{\text{LS}} = \sum_i \|\mathbf{f}(\mathbf{x}_i; \mathbf{p}) - \mathbf{x}'_i\|^2$$

What is the free variable?

What do we want to optimize?



Find parameters that minimize squared error

$$\hat{\mathbf{p}} = \arg \min_{\mathbf{p}} \sum_i \|\mathbf{f}(\mathbf{x}_i; \mathbf{p}) - \mathbf{x}'_i\|^2$$



General form of linear least squares

(**Warning:** change of notation. \mathbf{x} is a vector of parameters!)


$$\begin{aligned} E_{\text{LLS}} &= \sum_i |\mathbf{a}_i \mathbf{x} - \mathbf{b}_i|^2 \\ &= \|\mathbf{A} \mathbf{x} - \mathbf{b}\|^2 \quad (\text{matrix form}) \end{aligned}$$

Determining unknown transformations

Suppose we have two triangles: ABC and DEF.

- What type of transformation will map A to D, B to E, and C to F?
- How do we determine the unknown parameters?

Important: We will see a different procedure for dealing with homographies!



Affine transform:
uniform scaling + shearing
+ rotation + translation


$$\begin{bmatrix} a_1 & a_2 & a_3 \\ a_4 & a_5 & a_6 \\ 0 & 0 & 1 \end{bmatrix}$$

How many degrees of freedom do we have?

Determining unknown transformations

Suppose we have two triangles: ABC and DEF.

- What type of transformation will map A to D, B to E, and C to F?
- How do we determine the unknown parameters?



unknowns \rightarrow

$$\mathbf{x}' = \mathbf{M}\mathbf{x}$$


\leftarrow point correspondences

- One point correspondence gives how many equations?
- How many point correspondences do we need?

Determining unknown transformations

Suppose we have two triangles: ABC and DEF.

- What type of transformation will map A to D, B to E, and C to F?
- How do we determine the unknown parameters?



unknowns

$$\mathbf{x}' = \mathbf{M}\mathbf{x}$$

point correspondences

How do we solve this for \mathbf{M} ?

Affine transformations

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



- How many unknowns?
- How many equations per match?
- How many matches do we need?



Affine transformations

- Residuals:

$$r_{x_i}(a, b, c, d, e, f) = (ax_i + by_i + c) - x'_i$$

$$r_{y_i}(a, b, c, d, e, f) = (dx_i + ey_i + f) - y'_i$$

- Cost function:

$$C(a, b, c, d, e, f) = \sum_{i=1}^n (r_{x_i}(a, b, c, d, e, f)^2 + r_{y_i}(a, b, c, d, e, f)^2)$$



Affine transformations

- Matrix form

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & y_1 & 1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_2 & y_2 & 1 \\ \vdots & & & & & \\ x_n & y_n & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_n & y_n & 1 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = \begin{bmatrix} x'_1 \\ y'_1 \\ x'_2 \\ y'_2 \\ \vdots \\ x'_n \\ y'_n \end{bmatrix}$$

A
 $2n \times 6$

t
 6×1

=

b
 $2n \times 1$



Determining unknown transformations

Affine transformation:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} p_1 & p_2 & p_3 \\ p_4 & p_5 & p_6 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Why can we drop the last line?

Vectorize transformation parameters:

$$\begin{bmatrix} x' \\ y' \\ x' \\ y' \\ \vdots \\ x' \\ y' \end{bmatrix} = \begin{bmatrix} x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \\ x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \\ \vdots & & & \vdots & & \\ x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \end{bmatrix} \begin{bmatrix} p_1 \\ p_2 \\ p_3 \\ p_4 \\ p_5 \\ p_6 \end{bmatrix}$$

Stack equations from point correspondences:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \end{bmatrix} \begin{bmatrix} p_1 \\ p_2 \\ p_3 \\ p_4 \\ p_5 \\ p_6 \end{bmatrix}$$



b

A

x

Notation in system form:



General form of linear least squares

(Warning: change of notation. \mathbf{x} is a vector of parameters!)

$$\begin{aligned} E_{\text{LLS}} &= \sum_i |\mathbf{a}_i \mathbf{x} - \mathbf{b}_i|^2 \\ &= \|\mathbf{A} \mathbf{x} - \mathbf{b}\|^2 \quad (\text{matrix form}) \end{aligned}$$

This function is quadratic.

How do you find the root of a quadratic?



Solving the linear system

Convert the system to a linear least-squares problem:

$$E_{LLS} = \|\mathbf{Ax} - \mathbf{b}\|^2$$

Expand the error:

$$E_{LLS} = \mathbf{x}^\top (\mathbf{A}^\top \mathbf{A}) \mathbf{x} - 2\mathbf{x}^\top (\mathbf{A}^\top \mathbf{b}) + \|\mathbf{b}\|^2$$

Minimize the error:

Set derivative to 0 $(\mathbf{A}^\top \mathbf{A}) \mathbf{x} = \mathbf{A}^\top \mathbf{b}$

Solve for \mathbf{x} $\mathbf{x} = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{b}$ ←

In Python:

```
1 import numpy as np
2 x, resid, rank, s = np.linalg.lstsq(A,b)
3 x
```

Note: You almost never want to compute the inverse of a matrix.




Linear least squares estimation only works when the transform function is ?



Linear least squares estimation only works when the transform function is **linear! (duh)**

Also doesn't deal well with outliers (next class !!!)

Homographies




To unwarped (rectify) an image

- solve for homography \mathbf{H} given \mathbf{p} and \mathbf{p}'
- solve equations of the form: $w\mathbf{p}' = \mathbf{H}\mathbf{p}$
 - linear in unknowns: w and coefficients of \mathbf{H}
 - \mathbf{H} is defined up to an arbitrary scale factor
 - how many points are necessary to solve for \mathbf{H} ?


Create point correspondences

Given a set of matched feature points $\{p_i, p'_i\}$ find the best estimate of H such that

$$P' = H \cdot P$$



original image



target image

How many correspondences do we need?



Determining the homography matrix

Write out linear equation for each correspondence:

$$P' = H \cdot P \quad \text{or} \quad \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \alpha \begin{bmatrix} h_1 & h_2 & h_3 \\ h_4 & h_5 & h_6 \\ h_7 & h_8 & h_9 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$



Determining the homography matrix

Write out linear equation for each correspondence:

$$P' = H \cdot P \quad \text{or} \quad \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \alpha \begin{bmatrix} h_1 & h_2 & h_3 \\ h_4 & h_5 & h_6 \\ h_7 & h_8 & h_9 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Expand matrix multiplication:

$$x' = \alpha(h_1x + h_2y + h_3)$$

$$y' = \alpha(h_4x + h_5y + h_6)$$

$$1 = \alpha(h_7x + h_8y + h_9)$$



Determining the homography matrix

Write out linear equation for each correspondence:

$$P' = H \cdot P \quad \text{or} \quad \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \alpha \begin{bmatrix} h_1 & h_2 & h_3 \\ h_4 & h_5 & h_6 \\ h_7 & h_8 & h_9 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Expand matrix multiplication:

$$x' = \alpha(h_1x + h_2y + h_3)$$

$$y' = \alpha(h_4x + h_5y + h_6)$$

$$1 = \alpha(h_7x + h_8y + h_9)$$

Divide out unknown scale factor:

$$x'(h_7x + h_8y + h_9) = (h_1x + h_2y + h_3)$$

$$y'(h_7x + h_8y + h_9) = (h_4x + h_5y + h_6)$$

*How do you
rearrange terms
to make it a
linear system?*



$$x'(h_7x + h_8y + h_9) = (h_1x + h_2y + h_3)$$

$$y'(h_7x + h_8y + h_9) = (h_4x + h_5y + h_6)$$

Just rearrange the terms

$$h_7xx' + h_8yx' + h_9x' - h_1x - h_2y - h_3 = 0$$

$$h_7xy' + h_8yy' + h_9y' - h_4x - h_5y - h_6 = 0$$



Solving for homographies

$$x'_i(h_{20}x_i + h_{21}y_i + h_{22}) = h_{00}x_i + h_{01}y_i + h_{02}$$

$$y'_i(h_{20}x_i + h_{21}y_i + h_{22}) = h_{10}x_i + h_{11}y_i + h_{12}$$

$$\begin{bmatrix} x_i & y_i & 1 & 0 & 0 & 0 & -x'_i x_i & -x'_i y_i & -x'_i \\ 0 & 0 & 0 & x_i & y_i & 1 & -y'_i x_i & -y'_i y_i & -y'_i \end{bmatrix} \begin{bmatrix} h_{00} \\ h_{01} \\ h_{02} \\ h_{10} \\ h_{11} \\ h_{12} \\ h_{20} \\ h_{21} \\ h_{22} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$



Solving for homographies

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & -x'_1 x_1 & -x'_1 y_1 & -x'_1 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -y'_1 x_1 & -y'_1 y_1 & -y'_1 \\ & & & & & \vdots & & & \\ x_n & y_n & 1 & 0 & 0 & 0 & -x'_n x_n & -x'_n y_n & -x'_n \\ 0 & 0 & 0 & x_n & y_n & 1 & -y'_n x_n & -y'_n y_n & -y'_n \end{bmatrix} \begin{bmatrix} h_{00} \\ h_{01} \\ h_{02} \\ h_{10} \\ h_{11} \\ h_{12} \\ h_{20} \\ h_{21} \\ h_{22} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

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

$2n \times 9$ 9 2n

Defines a least squares problem: minimize $\|\mathbf{A}\mathbf{h} - \mathbf{0}\|^2$

- Since \mathbf{h} is only defined up to scale, solve for unit vector $\hat{\mathbf{h}}$
- Solution: $\hat{\mathbf{h}}$ = eigenvector of $\mathbf{A}^T \mathbf{A}$ with smallest eigenvalue
- Works with 4 or more points



Reminder: Determining affine transformations

Affine transformation:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} p_1 & p_2 & p_3 \\ p_4 & p_5 & p_6 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Vectorize transformation parameters:

$$\begin{bmatrix} x' \\ y' \\ x' \\ y' \\ \vdots \\ x' \\ y' \end{bmatrix} = \begin{bmatrix} x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \\ x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \\ \vdots & & & \vdots & & \\ x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \end{bmatrix} \begin{bmatrix} p_1 \\ p_2 \\ p_3 \\ p_4 \\ p_5 \\ p_6 \end{bmatrix}$$

Stack equations from point correspondences:

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

$$\underbrace{\begin{bmatrix} x' \\ y' \end{bmatrix}}_{\mathbf{b}} = \underbrace{\begin{bmatrix} x & y & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & y & 1 \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} p_1 \\ p_2 \\ p_3 \\ p_4 \\ p_5 \\ p_6 \end{bmatrix}}_{\mathbf{x}}$$

Notation in system form:

$$\boxed{Ax = b}$$



Reminder: Determining affine transformations

Convert the system to a linear least-squares problem:

$$E_{LLS} = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$$

Expand the error:

$$E_{LLS} = \mathbf{x}^\top (\mathbf{A}^\top \mathbf{A}) \mathbf{x} - 2\mathbf{x}^\top (\mathbf{A}^\top \mathbf{b}) + \|\mathbf{b}\|^2$$

Minimize the error:

Set derivative to 0 $(\mathbf{A}^\top \mathbf{A})\mathbf{x} = \mathbf{A}^\top \mathbf{b}$

Solve for \mathbf{x} $\mathbf{x} = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{b}$ ←

In Python:

```
1 import numpy as np
2 x, resid, rank, s = np.linalg.lstsq(A,b)
3 x
```

Note: You almost never want to compute the inverse of a matrix.

Determining the homography matrix

Stack together constraints from multiple point correspondences:

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

$$\begin{bmatrix} -x & -y & -1 & 0 & 0 & 0 & xx' & yx' & x' \\ 0 & 0 & 0 & -x & -y & -1 & xy' & yy' & y' \\ -x & -y & -1 & 0 & 0 & 0 & xx' & yx' & x' \\ 0 & 0 & 0 & -x & -y & -1 & xy' & yy' & y' \\ \vdots & & & & & & & & \\ -x & -y & -1 & 0 & 0 & 0 & xx' & yx' & x' \\ 0 & 0 & 0 & -x & -y & -1 & xy' & yy' & y' \end{bmatrix} \begin{bmatrix} h_1 \\ h_2 \\ h_3 \\ h_4 \\ h_5 \\ h_6 \\ h_7 \\ h_8 \\ h_9 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

Homogeneous linear least squares problem

- How do we solve this?



Determining the homography matrix

Stack together constraints from multiple point correspondences:

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

$$\begin{bmatrix} -x & -y & -1 & 0 & 0 & 0 & xx' & yx' & x' \\ 0 & 0 & 0 & -x & -y & -1 & xy' & yy' & y' \\ -x & -y & -1 & 0 & 0 & 0 & xx' & yx' & x' \\ 0 & 0 & 0 & -x & -y & -1 & xy' & yy' & y' \\ \vdots & & & & & & & & \\ -x & -y & -1 & 0 & 0 & 0 & xx' & yx' & x' \\ 0 & 0 & 0 & -x & -y & -1 & xy' & yy' & y' \end{bmatrix} \begin{bmatrix} h_1 \\ h_2 \\ h_3 \\ h_4 \\ h_5 \\ h_6 \\ h_7 \\ h_8 \\ h_9 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

Homogeneous linear least squares problem

- Solve with SVD



Recap: Two Common Optimization Problems

Problem statement

$$\text{minimize } \|\mathbf{Ax} - \mathbf{b}\|^2$$

least squares solution to $\mathbf{Ax} = \mathbf{b}$

Solution

$$\mathbf{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$$

```
import numpy as np
x, resid, rank, s = np.linalg.lstsq(A, b)
```

Problem statement

$$\text{minimize } \mathbf{x}^T \mathbf{A}^T \mathbf{A} \mathbf{x} \text{ s.t. } \mathbf{x}^T \mathbf{x} = 1$$

non - trivial lsq solution to $\mathbf{Ax} = 0$

Solution

$$[\mathbf{v}, \lambda] = \text{eig}(\mathbf{A}^T \mathbf{A})$$

$$\lambda_1 < \lambda_{2..n} : \mathbf{x} = \mathbf{v}_1$$

Derivation using Least squares

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

The sum squared error can be written as:

$$f(\mathbf{h}) = \frac{1}{2} (A\mathbf{h} - \mathbf{0})^T (A\mathbf{h} - \mathbf{0})$$

$$f(\mathbf{h}) = \frac{1}{2} (A\mathbf{h})^T (A\mathbf{h})$$

$$f(\mathbf{h}) = \frac{1}{2} \mathbf{h}^T A^T A \mathbf{h}.$$

Taking the derivative of f with respect to \mathbf{h} and setting the result to zero,

$$\begin{aligned} \frac{d}{d\mathbf{h}} f = 0 &= \frac{1}{2} (A^T A + (A^T A)^T) \mathbf{h} \\ 0 &= A^T A \mathbf{h}. \end{aligned}$$

\mathbf{h} should equal the eigenvector of $B = A^T A$ that has an eigenvalue of zero

$$B\vec{h} = \lambda\vec{h}$$

(or, in the presence of noise the eigenvalue closest to zero)



Singular Value Decomposition

$$\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$$

ortho-normal diagonal ortho-normal

$n \times m$ $n \times n$ $n \times m$ $m \times m$

unit norm constraint

$$= \sum_{i=1}^9 \sigma_i \mathbf{u}_i \mathbf{v}_i^T$$

$n \times 1$ $1 \times m$

Singular Value Decomposition

$$A = U \Sigma V^{-1} \quad \Sigma = \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \cdot & \\ & & & \sigma_N \end{bmatrix}$$

U, V = orthogonal matrix

$$\sigma_i = \sqrt{\lambda_i}$$

σ = singular value

λ = eigenvalue of $A^t A$


General form of total least squares

(Warning: change of notation. \mathbf{x} is a vector of parameters!)

$$E_{\text{TLS}} = \sum_i (\mathbf{a}_i \mathbf{x})^2$$

$$= \|\mathbf{A}\mathbf{x}\|^2 \quad \text{(matrix form)}$$

$$\|\mathbf{x}\|^2 = 1 \quad \text{constraint}$$

minimize	$\ \mathbf{A}\mathbf{x}\ ^2$		minimize	$\frac{\ \mathbf{A}\mathbf{x}\ ^2}{\ \mathbf{x}\ ^2}$
subject to	$\ \mathbf{x}\ ^2 = 1$			(Rayleigh quotient)

Solution is the eigenvector corresponding to smallest eigenvalue of

$$\mathbf{A}^\top \mathbf{A}$$

(equivalent)

Solution is the column of \mathbf{V} corresponding to smallest singular value

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$$



$$A^T A = (U \Sigma V^{-1})^T U \Sigma V^{-1}$$

$$A^T A = V \Sigma U^{-1} U \Sigma V^{-1}$$

$$A^T A = V \Sigma \Sigma V^{-1}$$

$$A^T A V = V \Sigma \Sigma$$

$$A^T A V = \Sigma^2 V$$



Homogeneous Linear Least Squares problem

$$A\mathbf{x} = \mathbf{0}$$

$$A = U\Sigma V^T = \sum_{i=1}^9 \sigma_i \mathbf{u}_i \mathbf{v}_i^T$$

- If the homography is *exactly determined*, then $\sigma_9 = 0$, and there exists a homography that fits the points exactly.
- If the homography is *overdetermined*, then $\sigma_9 \geq 0$. Here σ_9 represents a “residual” or goodness of fit.
- We will not handle the case of the homography being *underdetermined*.



Solving for H using DLT

Given $\{x_i, x'_i\}$ solve for H such that $x' = Hx$

1. For each correspondence, create 2x9 matrix A_i
2. Concatenate into single $2n \times 9$ matrix A
3. Compute SVD of $A = U\Sigma V^T$
4. Store singular vector of the smallest singular value $h = v_{\hat{i}}$
5. Reshape to get H



Recap: Two Common Optimization Problems

Problem statement

$$\text{minimize } \|\mathbf{Ax} - \mathbf{b}\|^2$$

least squares solution to $\mathbf{Ax} = \mathbf{b}$

Solution

$$\mathbf{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$$

```
import numpy as np
x, resid, rank, s = np.linalg.lstsq(A, b)
```

Problem statement

$$\text{minimize } \mathbf{x}^T \mathbf{A}^T \mathbf{A} \mathbf{x} \text{ s.t. } \mathbf{x}^T \mathbf{x} = 1$$

non - trivial lsq solution to $\mathbf{Ax} = 0$

Solution

$$[\mathbf{v}, \lambda] = \text{eig}(\mathbf{A}^T \mathbf{A})$$

$$\lambda_1 < \lambda_{2..n} : \mathbf{x} = \mathbf{v}_1$$



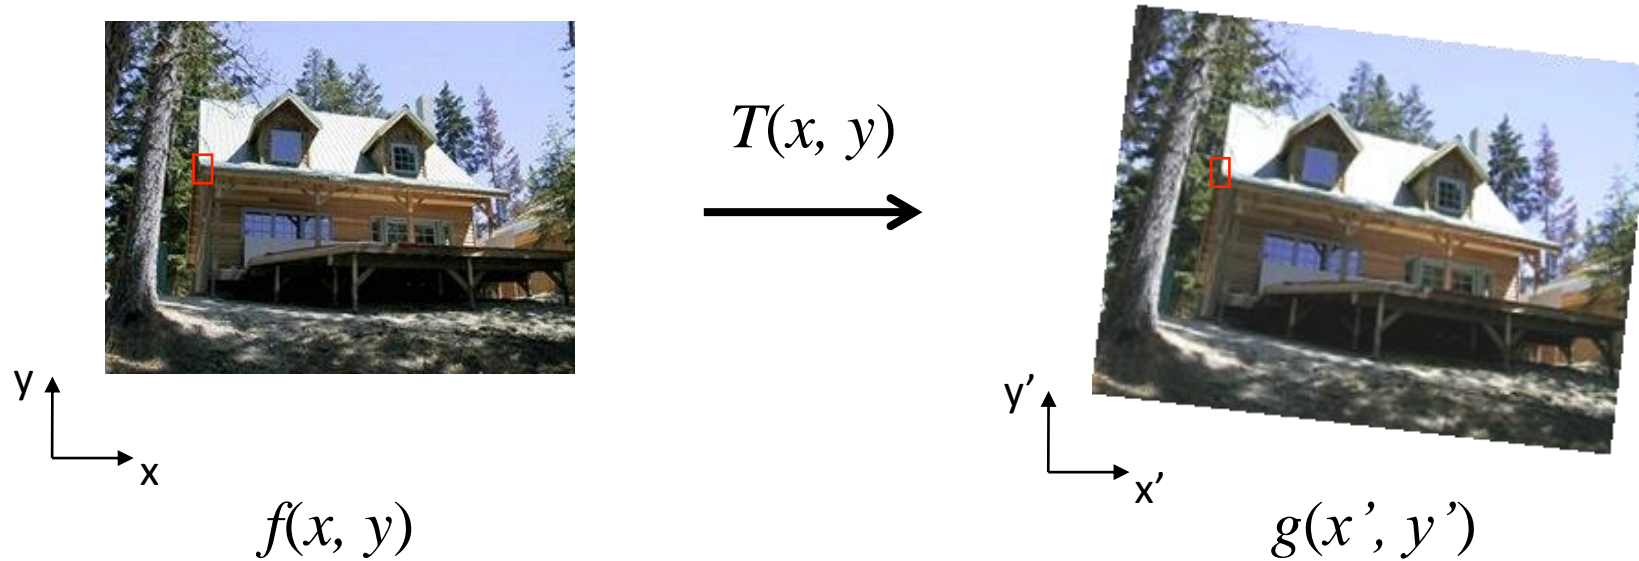
Outline

- Linear algebra
- Image transformations
- 2D transformations.
- Projective geometry 101.
- Transformations in projective geometry.
- Classification of 2D transformations.
- Determining unknown 2D transformations.
- **Determining unknown image warps.**

Determining unknown image warps

Suppose we have two images.


- How do we compute the transform that takes one to the other?



Forward warping

Suppose we have two images.

- How do we compute the transform that takes one to the other?




1. Form enough pixel-to-pixel correspondences between two images ← later lecture
2. Solve for linear transform parameters as before
3. Send intensities $f(x, y)$ in first image to their corresponding location in the second image

Forward warping

Suppose we have two images.

- How do we compute the transform that takes one to the other?



what is the problem with this?


1. Form enough pixel-to-pixel correspondences between two images
2. Solve for linear transform parameters as before
3. Send intensities $f(x, y)$ in first image to their corresponding location in the second image



Forward warping

Pixels may end up between two points


- How do we determine the intensity of each point?



Forward warping

Pixels may end up between two points

- How do we determine the intensity of each point?
- ✓ We distribute color among neighboring pixels (x',y') (“splatting”)




- What if a pixel (x',y') receives intensity from more than one pixels (x,y) ?

Forward warping

Pixels may end up between two points

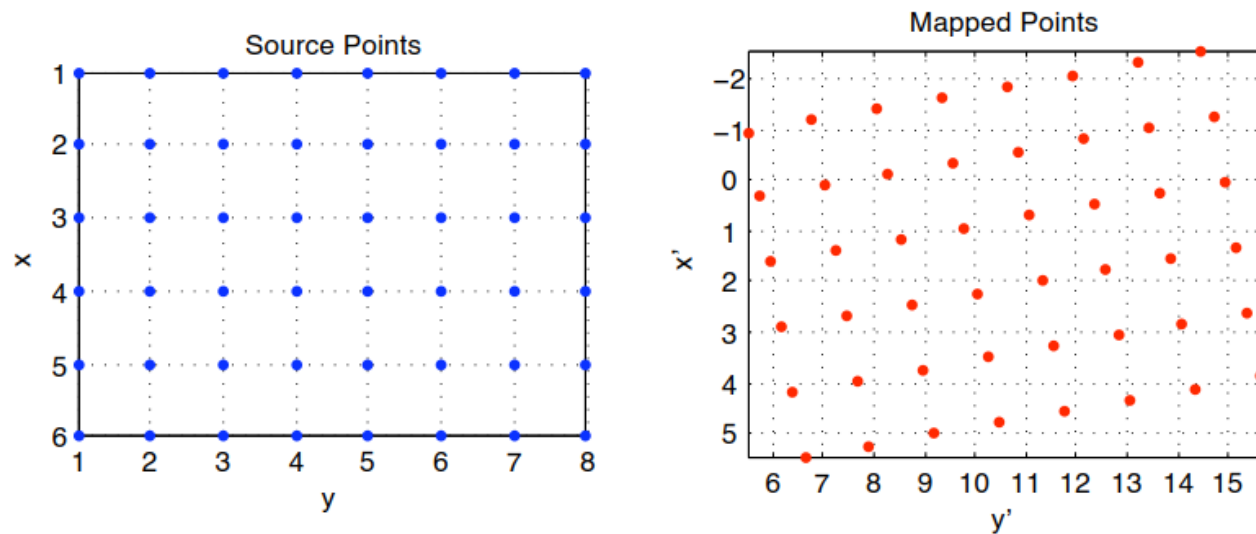
- How do we determine the intensity of each point?
- ✓ We distribute color among neighboring pixels (x',y') (“splatting”)



- What if a pixel (x',y') receives intensity from more than one pixels (x,y) ?
- ✓ We average their intensity contributions.

Forward mapping example

- Rotation Scale and Translation Mapping




The mapped points do not have integer coordinates!

Inverse warping

Suppose we have two images.

- How do we compute the transform that takes one to the other?



what is the problem with this?


1. Form enough pixel-to-pixel correspondences between two images
2. Solve for linear transform parameters as before, then compute its inverse
3. Get intensities $g(x', y')$ in in the second image from point $(x, y) = T^{-1}(x', y')$ in first image



Inverse warping

Pixel may come from between two points


- How do we determine its intensity?



Inverse warping

Pixel may come from between two points

- How do we determine its intensity?
- ✓ Use interpolation




Inverse warping

Pixel may come from between two points



- How do we determine its intensity?

✓ Use interpolation

- Nearest Neighbors
- Bilinear
- Cubic
- Lanczos



$T^{-1}(x,y)$





The Warping Recipe

- The fundamental steps of a warping algorithm are the following:
 - 1 Computation of the bounding box of the warped image (**forward** mapping).
 - 2 **Backward** mapping of lattice points that sample the bounding box of the warped image (avoid “holes”)
 - 3 Validation of the backward mapped points (must belong to the **domain** of the source image)
 - 4 Intensity transfer via **resampling**


Nearest Neighbor Interpolation

The question

Let x and y be the **integer** coordinates of the lattice. What is the value of f at $[p \ q]^T$?

Nearest Neighbour Answer

$$\hat{f}(p, q) = f(\text{round}(p), \text{round}(q))$$



Problem with NN interpolation



Point Sampled: Aliasing!




Correctly Bandlimited

Bilinear Interpolation

The question

Let x and y be the **integer** coordinates of the lattice. What is the value of f at $[p \ q]^T$?

- $F_{0,0} \stackrel{\text{def}}{=} f(x, y)$
- $F_{1,0} \stackrel{\text{def}}{=} f(x + 1, y)$
- $F_{0,1} \stackrel{\text{def}}{=} f(x, y + 1)$
- $F_{1,1} \stackrel{\text{def}}{=} f(x + 1, y + 1)$
- $\Delta x \stackrel{\text{def}}{=} p - x$ and $\Delta y \stackrel{\text{def}}{=} q - y$



Bilinear Interpolation

The question

Let x and y be the **integer** coordinates of the lattice. What is the value of f at $[p \ q]^T$?


- Linear interpolation in the x direction:

$$f_y(\Delta x) = (1 - \Delta x)F_{0,0} + \Delta x F_{1,0}$$

$$f_{y+1}(\Delta x) = (1 - \Delta x)F_{0,1} + \Delta x F_{1,1}$$

- Linear interpolation in the y direction:

$$\hat{f}(p, q) = (1 - \Delta y)f_y + \Delta y f_{y+1}$$



Bilinear Interpolation


The question

Let x and y be the **integer** coordinates of the lattice. What is the value of f at $[p \ q]^T$?

Bilinear Interpolation Answer

Note that $\hat{f}(p, q)$ "passes through" the samples.


$$\hat{f}(p, q) = (1 - \Delta y)(1 - \Delta x)F_{0,0} + (1 - \Delta y)\Delta x F_{1,0} + \Delta y(1 - \Delta x)F_{0,1} + \Delta y\Delta x F_{1,1}$$



Forward vs inverse warping

Suppose we have two images.

- How do we compute the transform that takes one to the other?




Pros and cons of each?

Forward vs inverse warping

Suppose we have two images.


- How do we compute the transform that takes one to the other?




- Inverse warping eliminates holes in target image
- Forward warping does not require existence of inverse transform

Warping with different transformations


translation



affine



pProjective (homography)




View warping

original view



synthetic top view



synthetic side view



What are these black areas near the boundaries?

Virtual camera rotations



original view

synthetic
rotations



Image rectification

two
original
images



rectified and stitched

Street art



Understanding geometric patterns


What is the pattern on the floor?



magnified view of floor

Understanding geometric patterns

What is the pattern on the floor?




Homography



magnified view of floor



rectified view




reconstruction from
rectified view

Understanding geometric patterns

Very popular in renaissance drawings (when perspective was discovered)



rectified view
of floor



reconstruction

A weird drawing

Holbein, "The Ambassadors"



A weird drawing


Holbein, "The Ambassadors"



What's this???

A weird drawing

Holbein, "The Ambassadors"



rectified view

skull under anamorphic perspective

A weird drawing

Holbein, "The Ambassadors"



DIY: use a polished spoon to see the skull

Panoramas from image stitching

1. Capture multiple images from different viewpoints.



2. Stitch them together into a virtual wide-angle image.





References

Basic reading:

- Szeliski textbook, Section 3.6.

Additional reading:

- Hartley and Zisserman, “Multiple View Geometry in Computer Vision,” Cambridge University Press 2004.
a comprehensive treatment of all aspects of projective geometry relating to computer vision, and also a very useful reference for the second part of the class.
- Richter-Gebert, “Perspectives on projective geometry,” Springer 2011.
a beautiful, thorough, and very accessible mathematics textbook on projective geometry (available online for free from CMU’s library).



Questions?