

E1 216 COMPUTER VISION

LECTURE 07: GEOMETRIC TRANSFORMATIONS

Venu Madhav Govindu
Department of Electrical Engineering
Indian Institute of Science, Bengaluru

2024

Geometric Transformations

- Use multiple or single image(s)
- Geometric - pure 3D rotations - mosaics
- Radiometric - high dynamic range imaging
- Focus on **geometric** transformations

Geometric Transformations



coolopticalillusions.com

Geometric Transformations




coolopticalillusions.com

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \mathbf{K} \left[\mathbf{R} \mid \mathbf{T} \right] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

Pinhole Camera

- Effects of rotations and translations are mixed
- **Only rotations ? (Mosaics)**
- Only translations ? (Stereo; considered later)
- Both ? (Multiview Geometry; considered later)

Geometric Transformations



$$\mathbf{p}_1 = \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} = \mathbf{K} \left[\mathbf{I} \mid \mathbf{0} \right] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{K}\mathbf{P}$$


$$\mathbf{p}_2 = \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} = \mathbf{K} \left[\mathbf{R} \mid \mathbf{0} \right] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{K}\mathbf{R}\mathbf{P}$$

$$\mathbf{p}_2 = \mathbf{K}\mathbf{R}\mathbf{K}^{-1}\mathbf{p}_1$$

Pure 3D Camera Rotation

- $\mathbf{P} = [X, Y, Z]^T$
- Pure 3D Rotations is a special case
- \mathbf{p}_1 and \mathbf{p}_2
 - related via camera parameters
 - does not depend on 3D geometry

Geometric Transformations



Rotating Camera

- Centre of projection same for all cameras
- Each image samples from same parametric ray set
- No “parallax” problem
- Depth plays no role
- Excellent for mosaics
- Equivalent to wider FOV camera

Why Mosaic?

Are you getting the whole picture?

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



Why Mosaic?

Are you getting the whole picture?

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




Why Mosaic?

Are you getting the whole picture?

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



Geometric Transformations



$$\mathbf{p}_1 = \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} = \mathbf{K} \left[\mathbf{I} \mid \mathbf{0} \right] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

$$\mathbf{p}_2 = \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} = \mathbf{K} \left[\mathbf{I} \mid \mathbf{T} \right] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

$$\mathbf{p}_1 = \mathbf{K}\mathbf{P} \quad \text{and} \quad \mathbf{p}_2 = \mathbf{K}(\mathbf{P} + \mathbf{T})$$

$$x_2 - x_1 = \frac{fB}{Z}$$

Pure 3D Camera Translation



- $\mathbf{P} = [X, Y, Z]^T$
- \mathbf{p}_1 and \mathbf{p}_2 related via translation and depth
- No simple relationship like pure rotations
- Used to recover 3D depth (stereo)

Geometric Transformations



urixblog.com

Geometric Transformations




Pure 3D Translations

- No single geometric (parametric) transformation
- Non-linear dependence on depth
- Use to estimate depth (stereo)
- Effects of 3D rotation and translation are complementary

We can also take a purely 2D geometric transformation view
Following slides borrowed from Noah Snavely

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

- $g(x) = h(f(x))$



- image warping: change *domain* of image

- $g(x) = f(h(x))$



Parametric (global) warping

- Examples of parametric warps:



translation



rotation



aspect



affine



perspective



cylindrical

Parametric (global) warping

- Examples of parametric warps:



translation



rotation




aspect

Parametric (global) warping



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



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

- Transformation T is a coordinate-changing machine:
$$\mathbf{p}' = T(\mathbf{p})$$
- What does it mean that T is global?
 - Is the same for any point \mathbf{p}
 - can be described by just a few numbers (parameters)
- Let's consider *linear* xforms (can be represented by a 2D matrix):

$$\mathbf{p}' = \mathbf{T}\mathbf{p} \quad \begin{bmatrix} x' \\ y' \end{bmatrix} = \mathbf{T} \begin{bmatrix} x \\ y \end{bmatrix}$$

Common linear transformations

- Uniform scaling by s :




$$\mathbf{S} = \begin{bmatrix} s & 0 \\ 0 & s \end{bmatrix}$$

What is the inverse?

Common linear transformations

- Rotation by angle θ (about the origin)



$$\mathbf{R} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

What is the inverse?

For rotations:

$$\mathbf{R}^{-1} = \mathbf{R}^T$$

2x2 Matrices

- What types of transformations can be represented with a 2x2 matrix?

2D mirror about Y axis?

$$\begin{aligned}x' &= -x \\ y' &= y\end{aligned}\quad \mathbf{T} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}$$

2D mirror across line $y = x$?

$$\begin{aligned}x' &= y \\ y' &= x\end{aligned}\quad \mathbf{T} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

2x2 Matrices

- What types of transformations can be represented with a 2x2 matrix?

2D Translation?

$$x' = x + t_x \quad \text{NO!}$$

$$y' = y + t_y$$

Translation is not a linear operation on 2D coordinates

All 2D Linear Transformations

- Linear transformations are combinations of ...

- Scale,
- Rotation,
- Shear, and
- Mirror

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

- Properties of linear transformations:

- Origin maps to origin
- Lines map to lines
- Parallel lines remain parallel
- Ratios are preserved
- Closed under composition


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

Homogeneous coordinates

Trick: add one more coordinate:

$$(x, y) \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

homogeneous image
coordinates



Converting *from* homogeneous coordinates

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow (x/w, y/w)$$

Translation

- Solution: homogeneous coordinates to the rescue

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

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

Affine transformations

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$$



any transformation with
last row $[0 \ 0 \ 1]$ we call an
affine transformation

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

Basic affine transformations

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

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

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

2D *in-plane* rotation

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

Shear

Affine Transformations

- Affine transformations are combinations of ...
 - Linear transformations, and
 - Translations
- $$\begin{bmatrix} x' \\ y' \\ w \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$
- Properties of affine transformations:
 - Origin does not necessarily map to origin
 - Lines map to lines
 - Parallel lines remain parallel
 - Ratios are preserved
 - Closed under composition

Where do we go from here?

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





what happens when we
mess with this row?

affine transformation

Projective Transformations aka Homographies aka Planar Perspective Maps

$$\mathbf{H} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix}$$

Called a *homography*
(or *planar perspective map*)




Homographies



Homographies

- Homographies ...
 - Affine transformations, and
 - Projective warps
- $$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$
- Properties of projective transformations:
 - Origin does not necessarily map to origin
 - Lines map to lines
 - Parallel lines do not necessarily remain parallel
 - Ratios are not preserved
 - Closed under composition

2D image transformations



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

These transformations are a nested set of groups


- Closed under composition and inverse is a member

Homographies




Image Warping

- Given a coordinate xform $(x',y') = T(x,y)$ and a source image $f(x,y)$, how do we compute an xformed image $g(x',y') = f(T(x,y))$?




Forward Warping

- Send each pixel $f(\mathbf{x})$ to its corresponding location $(\mathbf{x}', \mathbf{y}') = T(\mathbf{x}, \mathbf{y})$ in $g(\mathbf{x}', \mathbf{y}')$
- What if pixel lands “between” two pixels?




Forward Warping

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




Inverse Warping

- Get each pixel $g(x',y')$ from its corresponding location $(x,y) = T^{-1}(x',y')$ in $f(x,y)$
- Requires taking the inverse of the transform
- What if pixel comes from “between” two pixels?



Inverse Warping

- Get each pixel $g(\mathbf{x}')$ from its corresponding location $\mathbf{x}' = \mathbf{h}(\mathbf{x})$ in $f(\mathbf{x})$
- What if pixel comes from “between” two pixels?
- Answer: *resample* color value from *interpolated* (*prefiltered*) source image



Interpolation

- Possible interpolation filters:
 - nearest neighbor
 - bilinear
 - bicubic (interpolating)
 - sinc
- Needed to prevent “jaggies” and “texture crawl”
(with prefiltering)



Geometric Transformations




Enlarged FOV; Why do we have a radial shape ?

Following slides on impact of geometry of virtual camera plane
Taken from Magnus Oskarsson's slides

Panoramas


For calibrated cameras:




Panoramas



Panoramas




Panoramas



Points are transformed to the first image.




Panoramas



Panoramas


For calibrated cameras:



Distances are not preserved. Points close to the x -axis tend to infinity.




Panoramas



Panoramas

For calibrated cameras:




Cannot transfer all points into the first image.



Panoramas

For calibrated cameras:



Project onto a cylinder instead.



Panoramas

For calibrated cameras:



Distances are roughly preserved. Lines may not appear straight.



Geometric Transformations




Figure 3: A simplistic model showing how Projected Coordinate Systems are created using a sphere. Source:

[Britannica](#).

- Topology of sphere \neq that of 2D plane
- Issue has plagued map making!

<https://medium.com/nightingale/understanding-map-projections-8b23ecbd2a2f>

Geometric Transformations



Figure 4: The Mercator projection exaggerates the size of the countries as you move away from the Equator.
Source: snippet from The True Size Of.

Geometric Transformations




Figure 5: Map of the world in three different projections: (a) is in azimuthal projection that preserves distance from the center point, (b) is in a Mercator projection that preserves shape, and (c) is in cylindrical equal-area projection that preserves area. Source: [Wikipedia](https://en.wikipedia.org/wiki/Map_projection).

<https://medium.com/nightingale/understanding-map-projections-8b23ecbd2a2f>

Choosing the right projection system


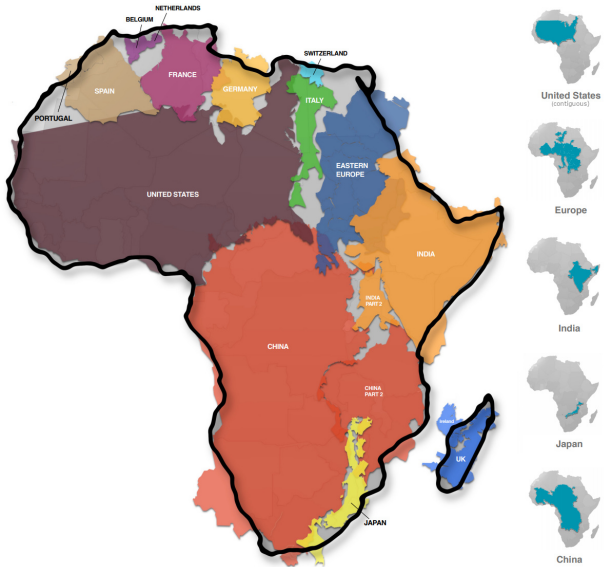


Figure 6: Robinson projection of the world. The projection is a compromise between the area and the shape of the world. Source: [Wikipedia](#).

<https://medium.com/nightingale/understanding-map-projections-8b23ecbd2a2f>

Geometric Transformations



AFRICA IS BIG!

Geometric Transformations



Why is north 'up' ?

Recovering Geometry

- Recall pure 3D rotations
- $\boldsymbol{p}_2 = \boldsymbol{K}\boldsymbol{R}\boldsymbol{K}^{-1}\boldsymbol{p}_1$
- Do we need to know \boldsymbol{K} and \boldsymbol{R} ?
- $\boldsymbol{H} = \boldsymbol{K}\boldsymbol{R}\boldsymbol{K}^{-1}$
- \boldsymbol{H} is 3×3 projective matrix
- \boldsymbol{H} is a homography/collineation/projective transformation
- $\boldsymbol{p}_2 = \boldsymbol{H}\boldsymbol{p}_1$

Recovering Geometry

- Recall pure 3D rotations
- $\boldsymbol{p}_2 = \boldsymbol{K}\boldsymbol{R}\boldsymbol{K}^{-1}\boldsymbol{p}_1$
- Do we need to know \boldsymbol{K} and \boldsymbol{R} ?
- $\boldsymbol{H} = \boldsymbol{K}\boldsymbol{R}\boldsymbol{K}^{-1}$
- \boldsymbol{H} is 3×3 projective matrix
- \boldsymbol{H} is a homography/collineation/projective transformation
- $\boldsymbol{p}_2 = \boldsymbol{H}\boldsymbol{p}_1$

Recovering Geometry

- Recall pure 3D rotations
- $\mathbf{p}_2 = \mathbf{K}\mathbf{R}\mathbf{K}^{-1}\mathbf{p}_1$
- Do we need to know \mathbf{K} and \mathbf{R} ?
- $\mathbf{H} = \mathbf{K}\mathbf{R}\mathbf{K}^{-1}$
- \mathbf{H} is 3×3 projective matrix
- \mathbf{H} is a homography/collineation/projective transformation
- $\mathbf{p}_2 = \mathbf{H}\mathbf{p}_1$

Homography relationship

How can we use this relationship $\mathbf{p}_2 = \mathbf{H}\mathbf{p}_1$

- Radiometric: $I_1(\mathbf{p}) = I_2(\mathbf{H}\mathbf{p})$
- Is this always true?
- Geometric: $\mathbf{p}_2 = \mathbf{H}\mathbf{p}_1$
- Need correspondences $\mathbf{p}_1 \leftrightarrow \mathbf{p}_2$

Geometric Transformations

$$\begin{array}{ll} \mathbf{H} & = \arg \min_{\mathbf{H}} \|I_1(\mathbf{p}) - I_2(\mathbf{H}\mathbf{p})\|^2 \\ \text{Update step} & \mathbf{H} \leftarrow \mathbf{H} + \delta\mathbf{H} \\ \text{Use} & I_2((\mathbf{H} + \delta\mathbf{H})\mathbf{p}) \approx I_2(\mathbf{H}\mathbf{p}) + \mathbf{J}^T \delta\mathbf{H} \\ \text{Minimise} & \|\mathbf{J}^T \delta\mathbf{H} - (I_1(\mathbf{p}) - I_2(\mathbf{H}\mathbf{p}))\|^2 \end{array}$$

Estimating Homographies

- Solution: Least square fit of intensities
- Is it a linear problem?
- Warp, Update, Warp, till convergence
- Use all pixels in overlapping area
- Robust loss $\rho(\cdot)$ for each pixel
- Multiscale approaches used. Why?
- Many issues in estimation

$$\begin{aligned} \mathbf{p}_2 &= \mathbf{H}\mathbf{p}_1 \\ \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} &= \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} \end{aligned}$$

Geometric Estimation

- Correspondences $\mathbf{p}_1 \leftrightarrow \mathbf{p}_2$ (SIFT etc.)
- $\mathbf{p}_2 = \mathbf{H}\mathbf{p}_1$ is a projective relationship
- Non-linear relationship?

Geometric Transformations

$$\begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Implies

$$x_2 = \frac{h_{11}x_1 + h_{12}y_1 + h_{13}}{h_{31}x_1 + h_{32}y_1 + h_{33}} \quad (1)$$

$$y_2 = \frac{h_{21}x_1 + h_{22}y_1 + h_{23}}{h_{31}x_1 + h_{32}y_1 + h_{33}} \quad (2)$$

Can solve using non-linear least squares on equations

Geometric Transformations

$$\begin{aligned}x_2 &= \frac{h_{11}x_1 + h_{12}y_1 + h_{13}}{h_{31}x_1 + h_{32}y_1 + h_{33}} \\y_2 &= \frac{h_{21}x_1 + h_{22}y_1 + h_{23}}{h_{31}x_1 + h_{32}y_1 + h_{33}}\end{aligned}$$

Linear in entries of \mathbf{H} , carry-over will result in

$$\begin{aligned}x_2(h_{31}x_1 + h_{32}y_1 + h_{33}) - (h_{11}x_1 + h_{12}y_1 + h_{13}) &= 0 \\y_2(h_{31}x_1 + h_{32}y_1 + h_{33}) - (h_{21}x_1 + h_{22}y_1 + h_{23}) &= 0\end{aligned}$$

Leads to

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & -x_1x_2 & -y_1x_2 & -x_2 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -x_1y_2 & -y_1y_2 & -y_2 \end{bmatrix} \begin{bmatrix} h_{11} \\ h_{12} \\ h_{13} \\ \vdots \\ h_{33} \end{bmatrix} = \mathbf{0}$$

Geometric Transformations

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & -x_1x_2 & -y_1x_2 & -x_2 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -x_1y_2 & -y_1y_2 & -y_2 \end{bmatrix} \begin{bmatrix} h_{11} \\ h_{12} \\ h_{13} \\ \vdots \\ h_{33} \end{bmatrix} = \mathbf{0}$$

Linear Method

- 2 eqns per correspondence
- Unknowns in \mathbf{H} ?
- Collect all equations into $\mathbf{A}\mathbf{h} = \mathbf{0}$ problem
- Solution ?
- Two important considerations
 - Robustness (RANSAC or IRLS?)
 - Conditioning (Scale of data)

Geometric Transformations

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & -x_1x_2 & -y_1x_2 & -x_2 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -x_1y_2 & -y_1y_2 & -y_2 \end{bmatrix} \begin{bmatrix} h_{11} \\ h_{12} \\ h_{13} \\ \vdots \\ h_{33} \end{bmatrix} = \mathbf{0}$$

Normalisation

- Recall that all correspondences are noisy
- (x, y) co-ordinates of order of 1000
- Quadratic terms in \mathbf{A}
- Errors in observed \mathbf{A} are not uniform in dimensions
- Leads to very poor conditioning of $\mathbf{A}\mathbf{h} = \mathbf{0}$
- Remedy
 - Scale co-ordinates (x, y) to have magnitude around 1
 - Solve
 - Put back original scale

* Find features using SIFT (vl-feat package) or SURF (inbuilt in MATLAB)

<https://www.vlfeat.org/overview/sift.html>

* Feature matching using vl-feat or MATLAB inbuilt functions.

NOTE: PLEASE BE MINDFUL OF THE CONVENTIONS USED FOR IMAGE COORDINATE SYSTEM,
CAMERA " "

* Raw (putative) matches: outliers included.
Need robustness while estimating homography.

$P' = \begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} \leftrightarrow \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = P$ are matched features in the 2 images.

$$P' = H P$$
$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{pmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}$$

$$x' = \frac{h_{11}x + h_{12}y + h_{13}}{h_{31}x + h_{32}y + h_{33}} \rightarrow (a)$$

$$y' = \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + h_{32}y + h_{33}} \rightarrow (b)$$

Method-1:

$$\begin{aligned}
 (a): (h_{31}x + h_{32}y + h_{33})x' &= h_{11}x + h_{12}y + h_{13} \\
 \Rightarrow h_{31}xx' + h_{32}yx' + h_{33}x' - h_{11}x - h_{12}y - h_{13} &= 0 \\
 (-x)h_{11} + (-y)h_{12} + (-1)h_{13} + (0)h_{21} + (0)h_{22} + (0)h_{23} \\
 + (xx')h_{31} + (yx')h_{32} + (x')h_{33} &= 0 \\
 \hookrightarrow (1)
 \end{aligned}$$

$$\begin{aligned}
 (b): (0)h_{11} + (0)h_{12} + (0)h_{13} + (-x)h_{21} + (-y)h_{22} + (-1)h_{23} \\
 + (xy')h_{31} + (yy')h_{32} + (y')h_{33} &= 0 \\
 \hookrightarrow (2)
 \end{aligned}$$

h_{11}, \dots, h_{33} are the unknowns.

$$\underbrace{\begin{bmatrix} -x & -y & -1 & 0 & 0 & 0 & xx' & yx' & x' \\ 0 & 0 & 0 & -x & -y & -1 & xy' & yy' & y' \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{bmatrix}}_A \begin{bmatrix} h_{11} \\ h_{12} \\ h_{13} \\ h_{21} \\ \vdots \\ h_{33} \end{bmatrix} = 0$$

$Ah = 0 \rightarrow (1)$

Soln: h : least eigen vector of $(A^T A)$

(or)

h : the right singular vector of 'A' corresponding to the least singular value.

Normalisation:

Consider $\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 400 \\ 200 \end{pmatrix}$ $\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} 800 \\ 400 \end{pmatrix}$

$$\begin{bmatrix} -4 \times 10^2 & -2 \times 10^2 & -1 & 0 & 0 & 0 & 3.2 \times 10^5 & 1.6 \times 10^5 & 8 \times 10^2 \\ 0 & 0 & 0 & -4 \times 10^2 & -2 \times 10^2 & -1 & 1.6 \times 10^5 & 0.8 \times 10^5 & 4 \times 10^2 \end{bmatrix}$$

Note that the entries are ranging from $0 - 10^5$ in A , then it will range from $0 - 10^{10}$ in $A^T A$.

Lack of homogeneity in the coordinates
 \Rightarrow poor conditioning of A (or) $A^T A$.

\rightarrow We want to reduce the range of the entries in A ($\Leftrightarrow A^T A$) to improve the conditioning (i.e., to reduce the $\kappa(A^T A)$).
image coordinates

→ Apply some transformations to the image coordinates in each image.

(i) translation

(ii) scaling.

(i) Translation: Origin of the new coordinate system should be at the centroid of the image points.

(ii) After translation, the coordinates are scaled s.t mean distance from the origin to a point equals $\sqrt{2}$.

(i) $\{x_i\}, \{y_i\}, i = 1 \dots N$

$$\bar{x} = \frac{1}{N} \sum x_i, \quad \bar{y} = \frac{1}{N} \sum y_i$$

we want, $x'_i = x_i - \bar{x}$

$$\frac{1}{N} \sum x'_i = \frac{1}{N} \sum x_i - \bar{x} = 0$$

$$\begin{bmatrix} x'_i \\ y'_i \\ 1 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & -\bar{x} \\ 0 & 1 & -\bar{y} \\ 0 & 0 & 1 \end{bmatrix}}_{T_{11}} \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix}$$

(ii) Mean distance from the origin to a point equals $\sqrt{2}$.

$$C' = \frac{1}{N} \sum d(x_i', y_i')$$
$$= \frac{1}{N} \sum \sqrt{x_i'^2 + y_i'^2}$$

Now if, $x_i'' = \sqrt{2} x_i' / C'$; $y_i'' = \sqrt{2} y_i' / C'$

Then,

$$C'' = \frac{1}{N} \sum d(x_i'', y_i'')$$
$$= \frac{1}{N} \sum \frac{\sqrt{x_i'^2 + y_i'^2}}{C'} = \frac{\sqrt{2} C'}{C'} = \sqrt{2}.$$

$$\therefore \begin{bmatrix} x_i'' \\ y_i'' \\ 1 \end{bmatrix} = \underbrace{\begin{bmatrix} \sqrt{2}/C' & 0 & 0 \\ 0 & \sqrt{2}/C' & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{T_2'} \begin{bmatrix} x_i' \\ y_i' \\ 1 \end{bmatrix}$$

$$\therefore \begin{bmatrix} x_i'' \\ y_i'' \\ 1 \end{bmatrix} = T_2' T_1' \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix}$$

$$T_1 = T_2^{-1} T_1$$

After transformation, we have

$$\text{in image } \textcircled{1} \leftarrow P_1' = T_1 P_1 ; \quad P_2' = T_2 P_2 \rightarrow \text{in image } \textcircled{2}$$

$$P_2 = H P_1 \Rightarrow T_2^{-1} P_2' = H T_1^{-1} P_1'$$

$$\Rightarrow P_2' = \underbrace{T_2 H T_1^{-1}}_{H'} P_1'$$

Find the homography H' using method-1; then compute

$$H' = T_2 H T_1^{-1}$$

$$T_2^{-1} H' T_1 = H$$

* Using RANSAC for robustness:

Minimum no. of points required: 4 for obtaining one homography.

Do 'N' trials, 'p' inlier probability.

Take randomly 4 feature matches out of N_0 possible matches in each trial.

Compute homography.

$[P]_x$: cross product form of a vector.

$$P' = HP; \quad \underbrace{[P']_X HP} = 0$$

$$\begin{cases} |a_i^T h| < \epsilon \\ |b_i^T h| < \epsilon \end{cases}$$

(or)

$$|[P']_X HP| < \epsilon$$

a_i^T, b_i^T : rows corresponding to one match in A' in ①

ϵ - tolerance

classify other points as inliers/outliers according to the above relations

choose that H_{\max} which has more inliers

→ then recompute H using all those inliers.

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

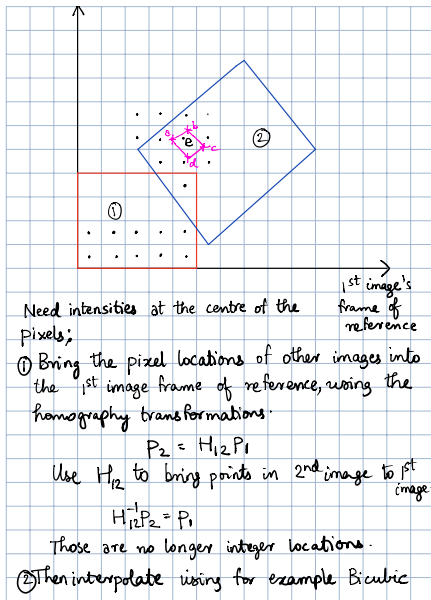
$2N \times 9$

N : max number of inliers for H_{\max}

→ The trials are meant to give us all the inliers, once we get them, we want to fit homography using all of them.


Interpolation:

→ Want to bring all the images into one frame of reference (say the first image's frame of reference).



interpolation (using the 4 nearest neighbours, in the above figure, intensity at 'e' is computed using those at 'a', 'b', 'c', 'd' → which are fractional locations obtained after step ①.)

BLENDING



Once all the images are in the same frame of reference, intensities in the overlap regions are computed using blending.

① Feathering

α varies linearly from 1 to 0 in the overlap region (as discussed in class)

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

② Using Image Pyramids (Not discussed in class, one could look it up.)

Note: Interested people can also use M-estimators for robustness which wasn't discussed in the class.

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & -x_1x_2 & -y_1x_2 & -x_2 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -x_1y_2 & -y_1y_2 & -y_2 \end{bmatrix} \begin{bmatrix} h_{11} \\ h_{12} \\ h_{13} \\ \vdots \\ h_{33} \end{bmatrix} = \mathbf{0}$$

Projective Scaling?

- Are all RHS zeros the same?
- What happens if we set $h_{33} = 1$?
- Yields an $A\mathbf{h} = \mathbf{b}$ problem

Geometric Transformations

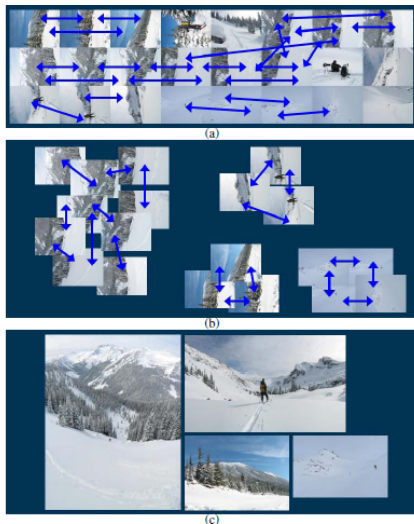





Figure 9.11 Recognizing panoramas (Brown, Szeliski, and Winder 2005), figures courtesy of Matthew Brown: (a) input images with pairwise matches; (b) images grouped into connected components (panoramas); (c) individual panoramas registered and blended into stitched composites.

Recognising Panoramas

Geometric Transformations



Consistency

- Pairwise alignment causes drift
- Use all relationships (“loop closures”)
- $H_{ij} = H_j H_i^{-1}$
- Significantly reduces inconsistencies
- Well-developed method of motion averaging