



University of Cyprus – MSc Artificial Intelligence

# MAI644 – COMPUTER VISION

## Lecture 8: Feature Descriptors and Image Transforms

**Melinos Averkiou**


CYENS Centre of Excellence

University of Cyprus - Department of Computer Science

[m.averkiou@cyens.org.cy](mailto:m.averkiou@cyens.org.cy)



**CYENS**  
CENTRE OF EXCELLENCE



## Last time

- Features
- Self-difference
- Harris corner detection

# Today's Agenda

- Basic feature descriptor and matching
- Histogram of Oriented Gradients
- SIFT
- Image transformations
- Estimate transformations


**[material based on Joseph Redmon's course]**

# Today's Agenda

- Basic feature descriptor and matching
- Histogram of Oriented Gradients
- SIFT
- Image transformations
- Estimate transformations


# Ok, we found corners, now what?

- Need to match image patches to each other
- Need to figure out transform between images




# Ok, we found corners, now what?

- Need to match image patches to each other
  - What is a match? How do we look for matches? Pixels?
- Need to figure out transform between images



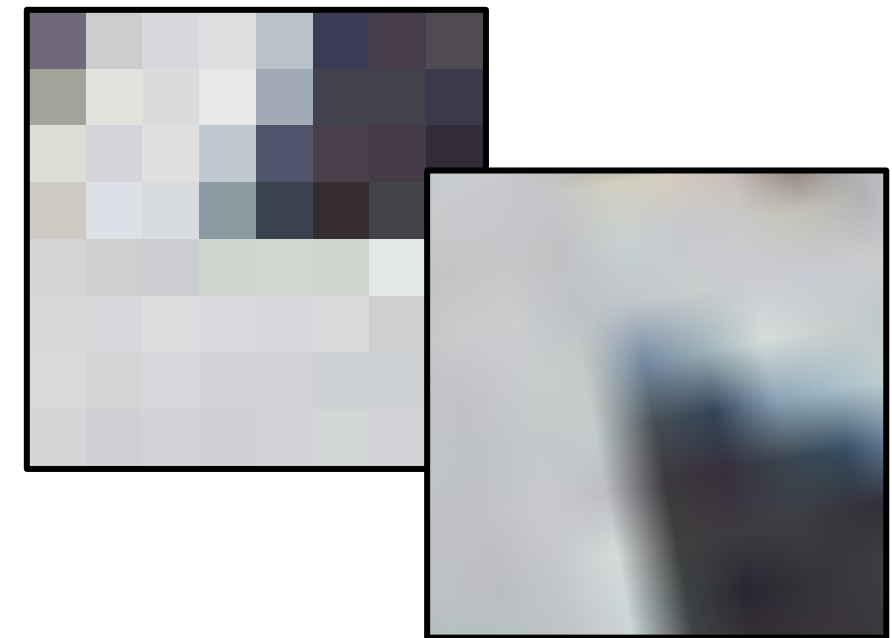
# Matching patches: descriptors!

- We want a way to represent an image patch
- Can be very simple, just pixels!
- Finding matching patch is easy, distance metric:
  - $\sum_{x,y} (I(x,y) - J(x,y))^2$
  - What problems are there with just using pixels?



# Matching patches: descriptors!


- We want a way to represent an image patch
- Can be very simple, just pixels!
- Finding matching patch is easy, distance metric:
  - $\sum_{x,y} (I(x,y) - J(x,y))^2$
  - Not invariant to some image transformations (e.g. rotation, scaling) !





# Matching patches: descriptors!

- We want a way to represent an image patch
- Can be very simple, just pixels!
- Finding matching patch is easy, distance metric:
  - $\sum_{x,y} (I(x,y) - J(x,y))^2$
  - Not invariant to lighting changes !



# Matching patches: descriptors!

- We want feature descriptors invariant to lighting and image transforms !
- Descriptors can be more complex
  - Gradient information
  - How much context?

# Today's Agenda

- Basic feature descriptor and matching
- Histogram of Oriented Gradients
- SIFT
- Image transformations
- Estimate transformations

# Histogram of Oriented Gradients (HOG)

- By Dalal and Triggs 2005
- Better image descriptor
- Not reliant on magnitude, just direction
  - Invariant to some lighting changes
- They used it to train an SVM to recognize people

Binning from <http://aishack.in/tutorials/sift-scale-invariant-feature-transform-features/>

# Histogram of Oriented Gradients (HOG)

Steps to calculate HOG Feature Descriptor

1. Compute gradients
2. Bin gradient directions to create histogram
3. Normalize histograms of gradients



# Histogram of Oriented Gradients (HOG)

## Steps to calculate HOG Feature Descriptor

### 1. Compute gradients

Gaussian smoothing (experimented with various  $\sigma$ ), followed by a derivative filter

- $\sigma = 0$  , i.e., no smoothing gave best results
- 1D filter [-1, 0, 1] gave best results




# Histogram of Oriented Gradients (HOG)

Steps to calculate HOG Feature Descriptor

2. Bin gradient directions to create histogram

Split image into 8x8 'cells' and compute histogram for each cell

- Unsigned gradients, i.e.,  $\theta=0-180$  degrees gave best results
- 9 bins gave best results



# Histogram of Oriented Gradients (HOG)

## Steps to calculate HOG Feature Descriptor

### 3. Normalize histograms of gradients

Gather overlapping 'cells' into 'blocks', concatenate histograms and normalize

- 16x16 blocks of 4 (2x2) cells gave best results
- L2 Normalization gave best results




# Histogram of Oriented Gradients (HOG)

For each training image of 64x128 there are 7x15 blocks, so the overall descriptor is  $7 \times 15 \times 36 = 3780$  dimensions



Training image



HOG descriptor of the image visualized for each 16x16 block




Descriptor weighted by the SVM weights


# This is as good as it gets ?

- Not so fast...
- Harris has some issues:
  - Corner detection is rotation invariant
  - Harris not invariant to scale
- Descriptors are also hard
  - Just looking at pixels is not rotation invariant!
  - HOG also not rotation invariant

Corner



Still Corner



Not Corner



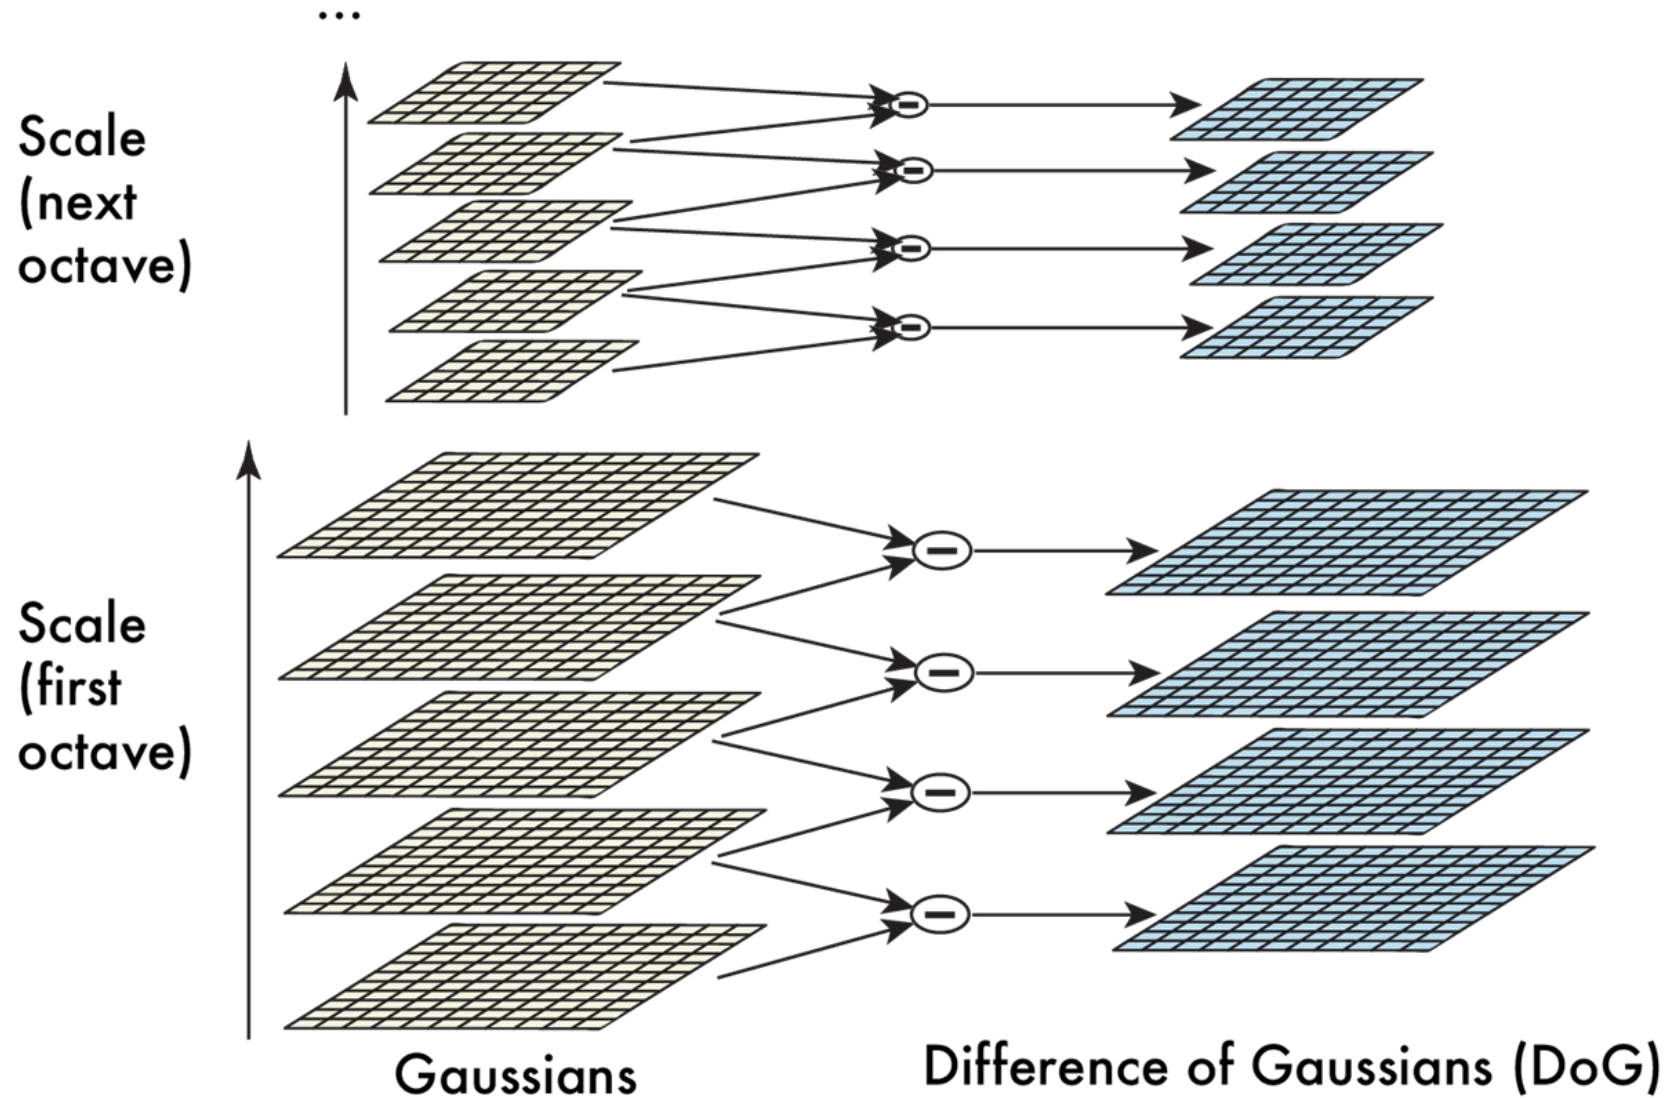
# Today's Agenda

- Basic feature descriptor and matching
- Histogram of Oriented Gradients
- SIFT
- Image transformations
- Estimate transformations

Want features invariant to scaling, rotation, etc.

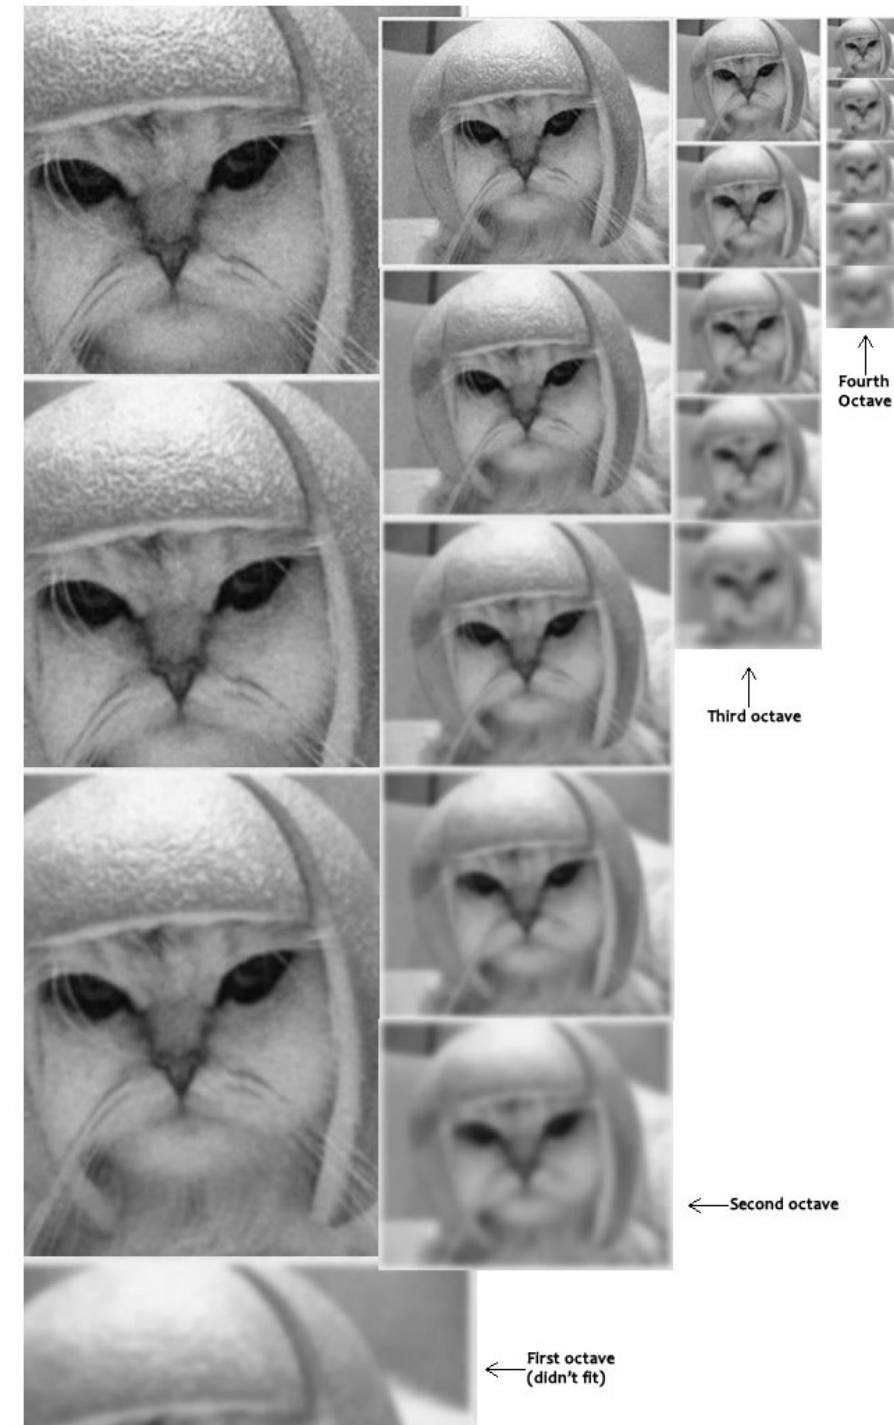
- Scale Invariant Feature Transform (SIFT)
  - Lowe et al. 2004, many images from that paper
- Get scale-invariant response map
- Find keypoints
- Extract rotation-invariant descriptors
  - Normalize based on orientation
  - Normalize based on lighting

# Extract DoG features at multiple scales




## Scale space

	scale →				
octave	0.707107	1.000000	1.414214	2.000000	2.828427
	1.414214	2.000000	2.828427	4.000000	5.656854
	2.828427	4.000000	5.656854	8.000000	11.313708
	5.656854	8.000000	11.313708	16.000000	22.627417



# Find local-maxima in location and scale



# Throw out weak responses and edges

- Estimate gradients
  - Similar to before, look at nearby responses
  - Not whole image, only a few points! Faster!
  - Throw out weak responses
- Find cornery things
  - Same deal, structure matrix, use det and trace information
  - $r$  : ratio of larger to smaller eigenvalue

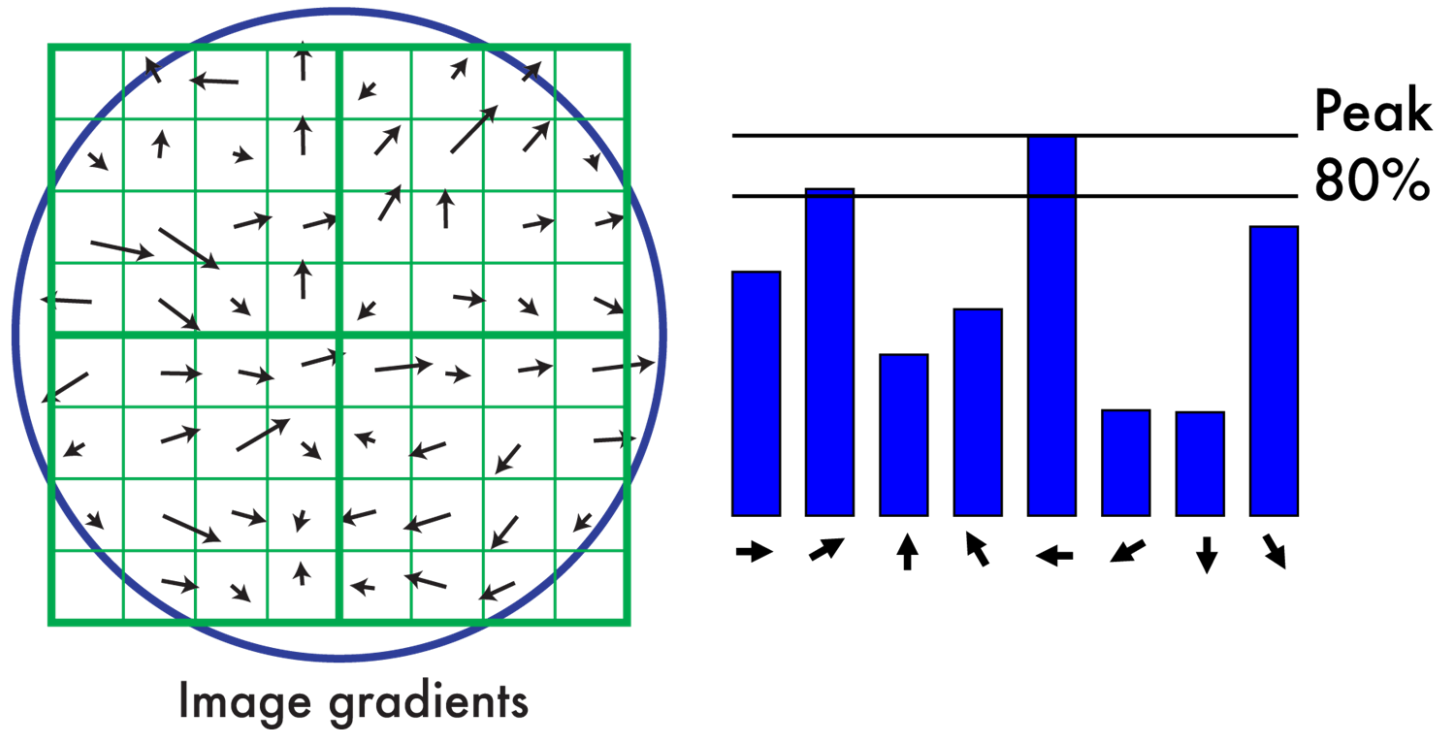
$$\frac{\text{Tr}(\mathbf{H})^2}{\text{Det}(\mathbf{H})} = \frac{(\alpha + \beta)^2}{\alpha\beta} = \frac{(r\beta + \beta)^2}{r\beta^2} = \frac{(r + 1)^2}{r};$$

$$\frac{\text{Tr}(\mathbf{H})^2}{\text{Det}(\mathbf{H})} < \frac{(r + 1)^2}{r}$$




# Find main orientation of patches

- Look at weighted histogram of nearby gradients
  - Any gradient within 80% of peak gets its own descriptor
    - Multiple keypoints per pixel
  - Descriptors are normalized based on main orientation



## Keypoints are normalized gradient histograms

- Divide into subwindows (4x4)
- Bin gradients within subwindow, get histogram
  - Normalize to unit length
  - Clamp at maximum .2
  - Normalize again
  - Helps with lighting changes!




SIFT is great!

- Find good keypoints, describe them
- Finding objects, recognition, panoramas, etc.



## SIFT is great!

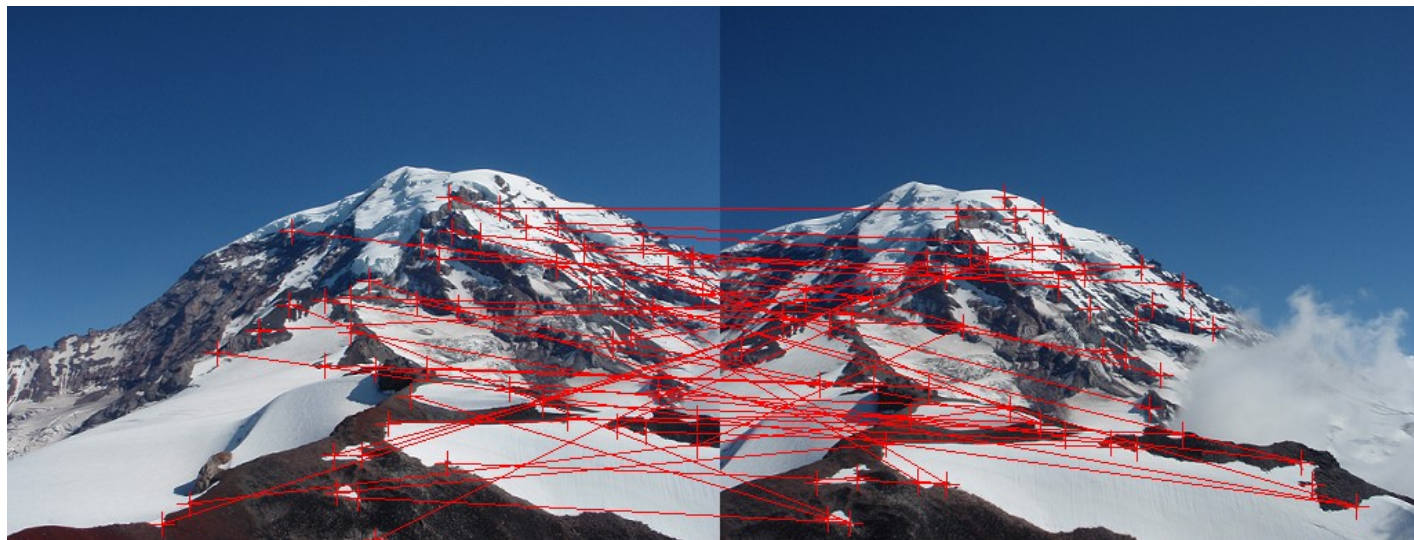


# Today's Agenda

- Basic feature descriptor and matching
- Histogram of Oriented Gradients
- SIFT
- Image transformations
- Estimate transformations


# Matching patches: descriptors!

- Already have our patches that are likely “unique”-ish
- Loop over good patches in one image
  - Find best match in other image
- Do something with them?




# Ok, we found corners, now what?

- Need to match image patches to each other
- Need to figure out transform between images



# Ok, we found corners, now what?

- Need to match image patches to each other
- Need to figure out transform between images
  - How can we transform images?
  - How do we solve for this transformation given matches?





# How can we transform images?

- Need to warp one image into the other
- Many different image transforms
  - Nested hierarchy of transformations

# How can we transform images?

- $\mathbf{x}$  is a point in our image where:
  - $\mathbf{x} = (x, y)$  or in matrix terms

$$\mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$



# Say we want new coordinate system

- Map points from one image into another
- Often we can use matrix operations
- Given a point  $x$ , map to new point  $x'$  using  $M$

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



# Scaling is just a matrix operation


- Map points from one image into another
- Often we can use matrix operations
- Given a point  $x$ , map to new point  $x'$  using  $M$

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

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

# Translation is harder...

- $x' = M x$ 
  - Want to move  $x'$  by  $dx$  and  $y'$  by  $dy$
  - How do we pick  $M$ ?
  - Can only add up multiples of  $x$  or  $y$
  - No easy way to add a constant!



# Translation: add another row

- $\bar{x}$  is  $x$  but with an added 1
- *Augmented vector*

$$\bar{x} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

# Translation: add another row

- $\bar{x}$  is  $x$  but with an added 1
- *Augmented vector*
- Now translation is easy

$$\bar{x} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$x' = [ I \quad t ] \bar{x}$$

# Reminder, I = Identity

Common to just use I as a generic, whatever size identity fits here.

$$\mathbf{I} = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ & & & \dots & \\ 0 & 0 & 0 & \dots & 1 \end{bmatrix} \quad \mathbf{I}_{2 \times 2} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \mathbf{I}_{3 \times 3} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$



# Translation: add another row

- $\bar{\mathbf{x}}$  is  $\mathbf{x}$  but with an added 1
- *Augmented vector*
- Now translation is easy

$$\bar{\mathbf{x}} = \begin{bmatrix} \mathbf{x} \\ y \\ 1 \end{bmatrix}$$

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

$$\mathbf{x}' = [\mathbf{I} \ \mathbf{t}] \bar{\mathbf{x}}$$

# Translation: add another row

- $\bar{x}$  is  $x$  but with an added 1
- *Augmented vector*
- Now translation is easy
- $x' = 1*x + 0*y + dx*1$
- $y' = 0*x + 1*y + dy*1$

$$\bar{x} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

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

$$x' = [I \ t] \bar{x}$$

# Translation: add another row

- $\bar{x}$  is  $x$  but with an added 1
- *Augmented vector*
- Now translation is easy
- $x' = 1*x + 0*y + dx*1$
- $y' = 0*x + 1*y + dy*1$


$$\bar{x} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

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

$$x' = [I \ t] \bar{x}$$

# Euclidean: rotation + translation

- Want to translate and rotate at same time
- Still just matrix operation





# Euclidean: rotation + translation

- Want to translate and rotate at same time
- Still just matrix operation

$$\mathbf{x}' = [ \mathbf{R} \quad \mathbf{t} ] \bar{\mathbf{x}}$$



# Euclidean: rotation + translation

- Want to translate and rotate at same time
- Still just matrix operation
- $\mathbf{R}$  is rotation matrix,  $\mathbf{t}$  is translation

$$\mathbf{x}' = [ \mathbf{R} \quad \mathbf{t} ] \bar{\mathbf{x}}$$

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

# Euclidean: rotation + translation


- Want to translate and rotate at same time
- Still just matrix operation
- $\mathbf{R}$  is rotation matrix,  $\mathbf{t}$  is translation

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

$$\mathbf{x}' = [\mathbf{R} \ \mathbf{t}] \bar{\mathbf{x}}$$

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


# Similarity: scale, rotate, translate





# Similarity: scale, rotate, translate

$$\mathbf{x}' = [ \mathbf{sR} \quad \mathbf{t} ] \bar{\mathbf{x}}$$




Similarity: scale, rotate, translate


$$\mathbf{x}' = [s\mathbf{R} \quad \mathbf{t}] \bar{\mathbf{x}}$$

$$\mathbf{x}' = \begin{bmatrix} a & -b & dx \\ b & a & dy \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

## Affine: scale, rotate, translate, shear



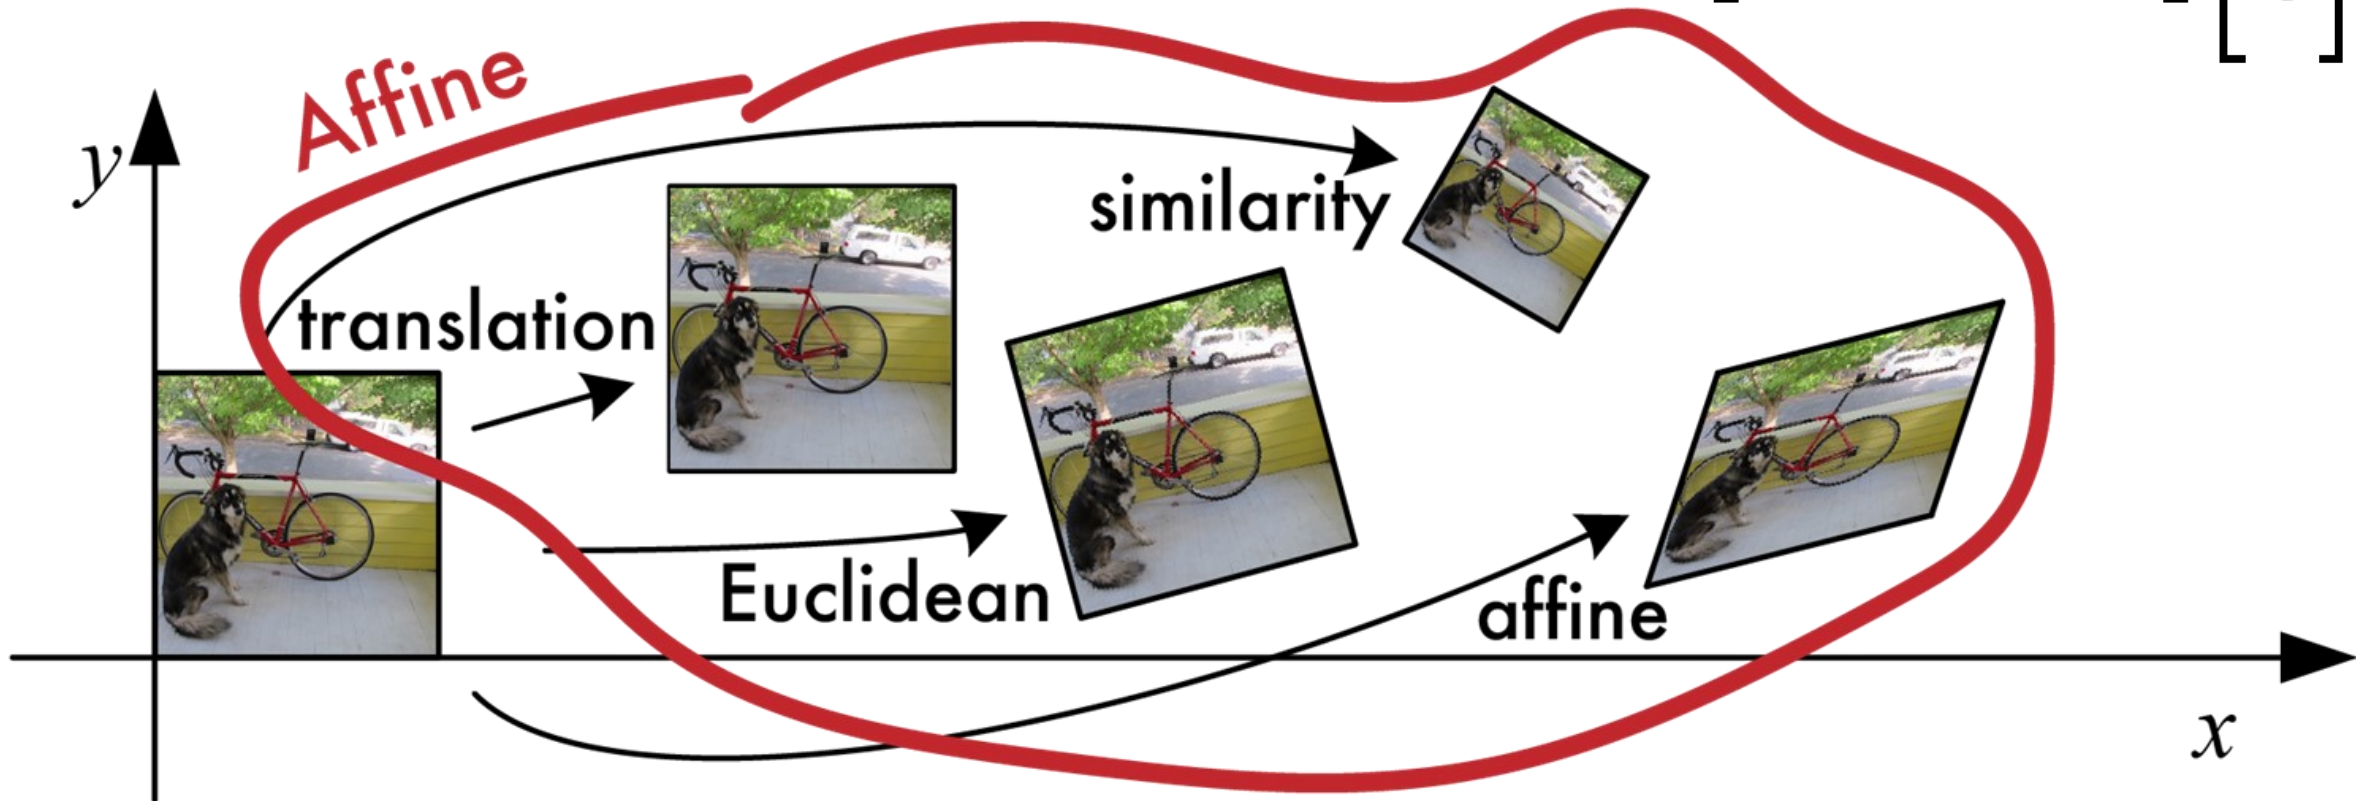
## Affine: scale, rotate, translate, shear



# Affine: scale, rotate, translate, shear

General case of 2x3 matrix

$$\mathbf{x}' = \begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$





# Combinations are still affine

Say you want to translate, then rotate, then translate back, then scale.

$$\mathbf{x}' = \mathbf{S} \mathbf{t} \mathbf{R} \mathbf{t} \bar{\mathbf{x}} = \mathbf{M} \bar{\mathbf{x}},$$

$$\text{If } \mathbf{M} = (\mathbf{S} \mathbf{t} \mathbf{R} \mathbf{t})$$

$\mathbf{M}$  is still affine transformation

Wait, but these are all 2x3, how to we multiply them together?



# Added row to transforms

$$\bar{\mathbf{x}}' = \begin{bmatrix} 1 & 0 & dx \\ 0 & 1 & dy \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

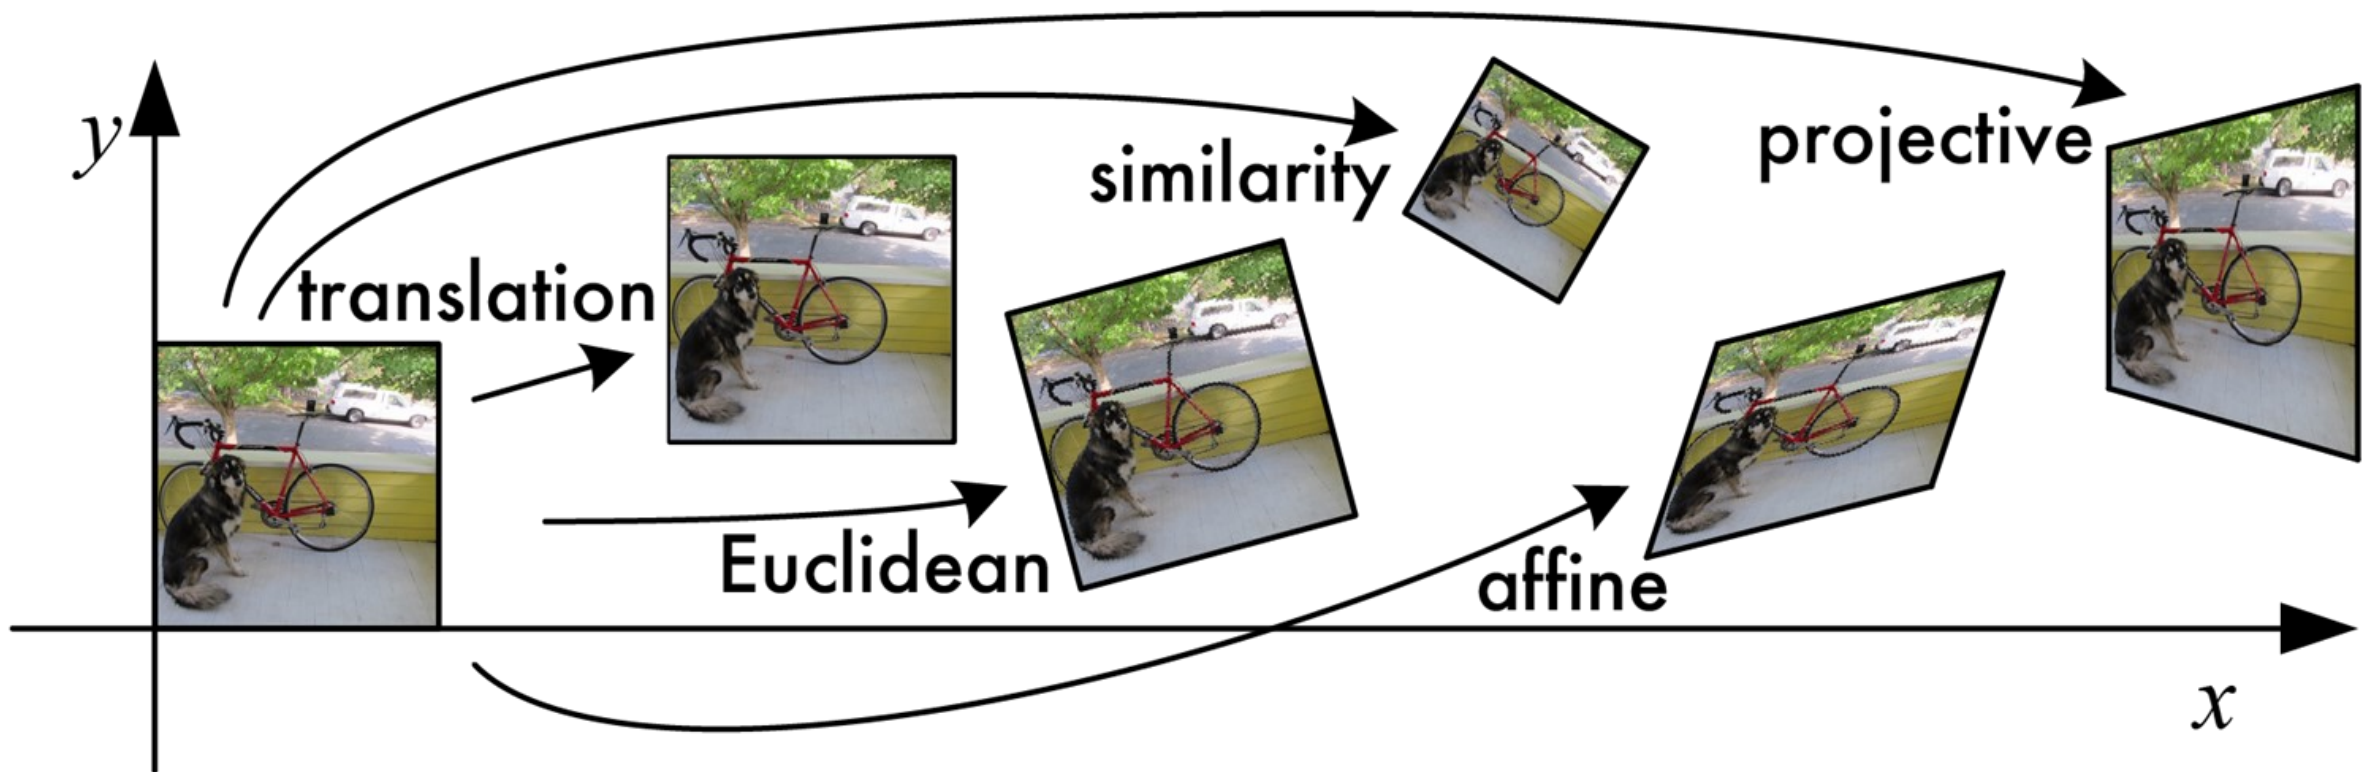
$$\bar{\mathbf{x}}' = \begin{bmatrix} \cos\theta & -\sin\theta & dx \\ \sin\theta & \cos\theta & dy \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\bar{\mathbf{x}}' = \begin{bmatrix} a & -b & dx \\ b & a & dy \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\bar{\mathbf{x}}' = \begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

## Projective transform

- Also known as homography
- Wait but affine was any 2x3 matrix...





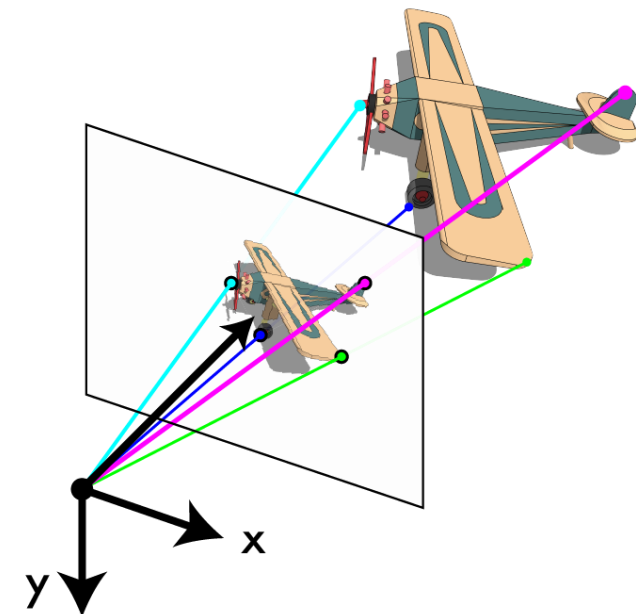

# Need some new coordinates!

- Homogeneous coordinate system
- Each point in 2d is actually a vector in 3d
- Equivalent up to scaling factor
- Have to normalize to get back to 2d

$$\tilde{\mathbf{x}} = \begin{bmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{w} \end{bmatrix} \quad \bar{\mathbf{x}} = \tilde{\mathbf{x}} / \tilde{w}$$

# Why does this make sense?

- Remember our pinhole camera model
- Every point in 3d projects onto our viewing plane through our aperture
- Points along a vector are indistinguishable



# Projective transform

- Also known as homography
- Wait but affine was any 2x3 matrix...
- Homography is general 3x3 matrix
- Multiplication by scalar is equivalent

$$\tilde{\mathbf{x}}' = \tilde{\mathbf{H}} \tilde{\mathbf{x}}$$

# Projective transform

- Also known as homography
- Wait but affine was any 2x3 matrix...
- Homography is general 3x3 matrix
- Multiplication by scalar: equivalent projection

-  $3 * H \sim H$

$$\tilde{\mathbf{x}}' = \begin{bmatrix} h_{00} & h_{01} & h_{02} \\ h_{10} & h_{11} & h_{12} \\ h_{20} & h_{21} & h_{22} \end{bmatrix} \begin{bmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{w} \end{bmatrix}$$

$$\tilde{\mathbf{x}}' = \tilde{\mathbf{H}} \tilde{\mathbf{x}}$$

# Using homography to project point

- Multiply  $\tilde{\mathbf{x}}$  by  $\tilde{\mathbf{H}}$  to get  $\tilde{\mathbf{x}}'$
- Convert to  $\bar{\mathbf{x}}$  by dividing by  $\tilde{w}'$

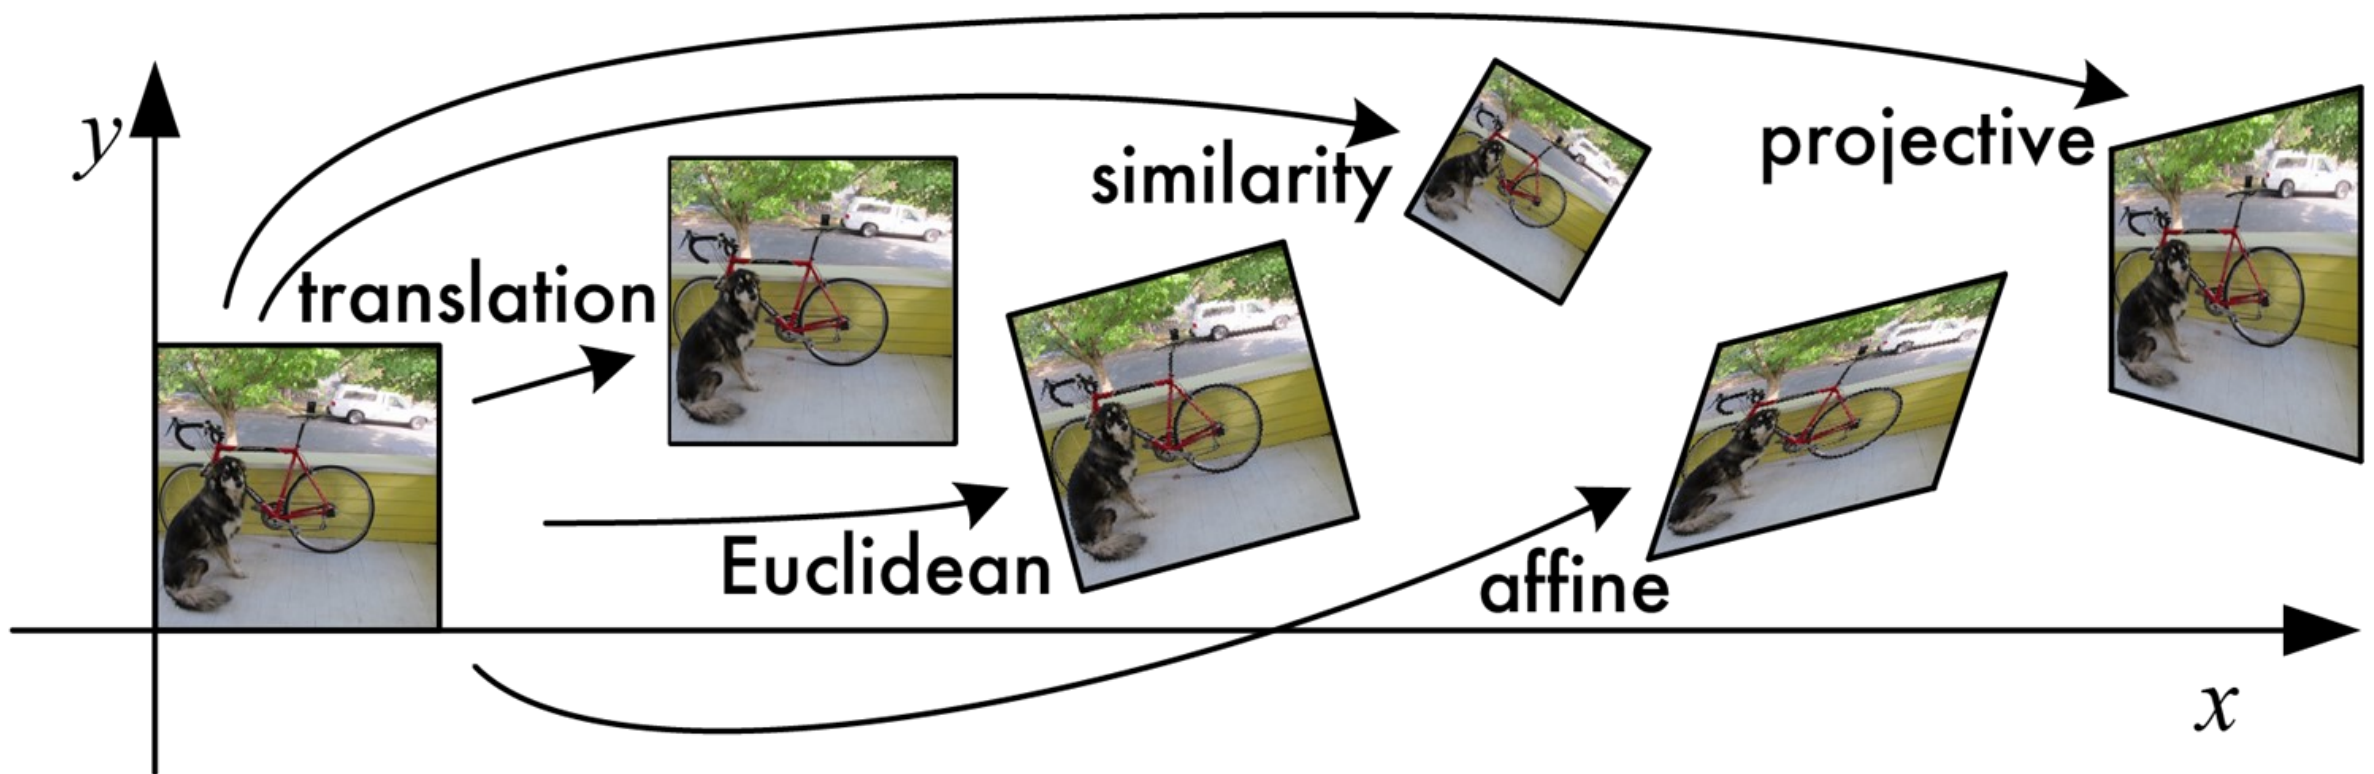
$$\tilde{\mathbf{x}}' = \tilde{\mathbf{H}} \tilde{\mathbf{x}}$$

$$\begin{bmatrix} \tilde{x}' \\ \tilde{y}' \\ \tilde{w}' \end{bmatrix} = \begin{bmatrix} h_{00} & h_{01} & h_{02} \\ h_{10} & h_{11} & h_{12} \\ h_{20} & h_{21} & h_{22} \end{bmatrix} \begin{bmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{w} \end{bmatrix}$$

$$\bar{\mathbf{x}} = \tilde{\mathbf{x}} / \tilde{w}$$

## Lots to choose from

- What do each of them do?
- Which is right for panorama stitching?





# Today's Agenda

- Basic descriptor and matching
- Image transformations
- Estimate transformations



# How hard are they to recover?

$$\bar{\mathbf{x}}' = \begin{bmatrix} 1 & 0 & dx \\ 0 & 1 & dy \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad \bar{\mathbf{x}}' = \begin{bmatrix} \cos\theta & -\sin\theta & dx \\ \sin\theta & \cos\theta & dy \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\bar{\mathbf{x}}' = \begin{bmatrix} a & -b & dx \\ b & a & dy \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\bar{\mathbf{x}}' = \begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\tilde{\mathbf{x}}' = \begin{bmatrix} h_{00} & h_{01} & h_{02} \\ h_{10} & h_{11} & h_{12} \\ h_{20} & h_{21} & h_{22} \end{bmatrix} \begin{bmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{w} \end{bmatrix}$$






## Lots to choose from

Transformation	Matrix	# DoF	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	



# Say we want affine transformation

- Have our matched points
- Want to estimate  $\mathbf{A}$  that maps from  $\mathbf{x}$  to  $\mathbf{x}'$
- $\mathbf{Ax} = \mathbf{x}'$



# Say we want affine transformation

- Have our matched points
- Want to estimate  $\mathbf{A}$  that maps from  $\mathbf{x}$  to  $\mathbf{x}'$
- $\mathbf{Ax} = \mathbf{x}'$
- How many degrees of freedom?

# Say we want affine transformation

- Have our matched points
- Want to estimate  $\mathbf{A}$  that maps from  $\mathbf{x}$  to  $\mathbf{x}'$
- $\mathbf{Ax} = \mathbf{x}'$
- How many degrees of freedom?
  - 6
- How many knowns do we get with one match?

$$\mathbf{x}' = \begin{bmatrix} \mathbf{a}_{00} & \mathbf{a}_{01} & \mathbf{a}_{02} \\ \mathbf{a}_{10} & \mathbf{a}_{11} & \mathbf{a}_{12} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \\ 1 \end{bmatrix}$$

# Say we want affine transformation

- Have our matched points
- Want to estimate  $\mathbf{A}$  that maps from  $\mathbf{x}$  to  $\mathbf{x}'$
- $\mathbf{Ax} = \mathbf{x}'$
- How many degrees of freedom?
  - 6
- How many knowns do we get with one match?
  - 2
  - $n_x = a_{00} * m_x + a_{01} * m_y + a_{02} * 1$
  - $n_y = a_{10} * m_x + a_{11} * m_y + a_{12} * 1$

$$\mathbf{x}' = \begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

## Say we want affine transformation

- How many knowns do we get with one match?
  - $n_x = a_{00} * m_x + a_{01} * m_y + a_{02} * 1$
  - $n_y = a_{10} * m_x + a_{11} * m_y + a_{12} * 1$
  - Solve linear system of equations  $\mathbf{M} \mathbf{a} = \mathbf{b}$ 
    - $M^{-1} M \mathbf{a} = M^{-1} \mathbf{b} \Rightarrow \mathbf{a} = M^{-1} \mathbf{b}$
    - But  $M^{-1}$  does not exist in general - Why?
- Still works if overdetermined
  - Why???
  - Pseudoinverse – least squares solution
  - $M^T M \mathbf{a} = M^T \mathbf{b}$
  - $(M^T M)^{-1} (M^T M) \mathbf{a} = (M^T M)^{-1} M^T \mathbf{b}$
  - $\Rightarrow \mathbf{a} = (M^T M)^{-1} M^T \mathbf{b}$

$$\begin{matrix} & \mathbf{M} & & \mathbf{a} & = & \mathbf{b} \\ \left[ \begin{array}{cccccc} m_{x1} & m_{y1} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & m_{x1} & m_{y1} & 1 \\ m_{x2} & m_{y2} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & m_{x2} & m_{y2} & 1 \\ m_{x3} & m_{y3} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & m_{x3} & m_{y3} & 1 \\ & \dots & & & & \\ & \dots & & & & \end{array} \right] & \left[ \begin{array}{c} a_{00} \\ a_{01} \\ a_{02} \\ a_{10} \\ a_{11} \\ a_{12} \end{array} \right] & = & \left[ \begin{array}{c} n_{x1} \\ n_{y1} \\ n_{x2} \\ n_{y2} \\ n_{x3} \\ n_{y3} \\ \dots \\ \dots \end{array} \right]
 \end{matrix}$$

**MAI4CAREU**

Master programmes in Artificial  
Intelligence 4 Careers in Europe



**CYENS**  
CENTRE OF EXCELLENCE



# Thank you.

