

Descriptors

EECS 442 – David Fouhey


Winter 2023, University of Michigan

https://web.eecs.umich.edu/~fouhey/teaching/EECS442_W23/

Administrivia

- Extra OH, zoom on Tuesday. Post HW topics you think are challenging; will try to cover top-voted questions.


Recap: Motivation




1: find corners+features

Last Time – Gradients

Image gradients – treat image like function of x, y – gives edges, corners, etc.



$$\nabla f = \left[\frac{\partial f}{\partial x}, 0 \right]$$




$$\nabla f = \left[0, \frac{\partial f}{\partial y} \right]$$




$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right]$$

Last Time – Corner Detection


Can localize the location, or any shift →
big intensity change.



“flat” region:
no change in
all directions



“edge”:
no change
along the edge
direction




“corner”:
significant
change in all
directions


Last Time – Corner Detection

Zoom-In at x, y

Window with and w/o Offset




“Window”
At $x+u, y+v$
Here: $u=-2, v=-3$



“Window”
At x, y

Last Time – Corner Detection

Zoom-In at x, y



Error (Sum Sqs) for u, v offset

$$E(u, v) = \sum_{(x, y) \in W} (I[x + u, y + v] - I[x, y])^2$$

$$\left(\begin{array}{c} \text{Red Box} \\ \text{Blue Box} \end{array} \right)^2$$

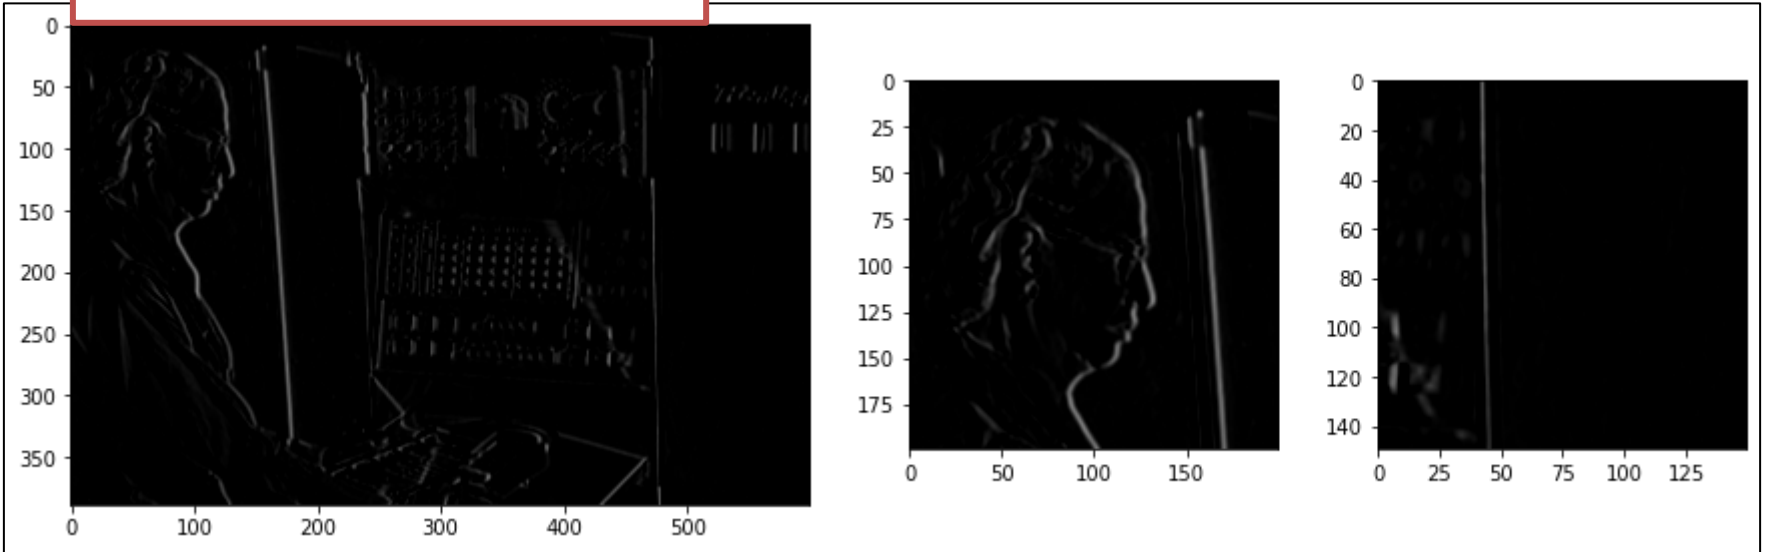
Formalizing Corner Detection

By linearizing image, we can approximate $E(u,v)$ with quadratic function of u and v

$$E(u, v) \approx \sum_{(x,y) \in W} (I_x^2 u^2 + 2I_x I_y uv + I_y^2 v^2)$$
$$= [u, v] \mathbf{M} [u, v]^T$$

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix}$$

$I_x = x$ derivative

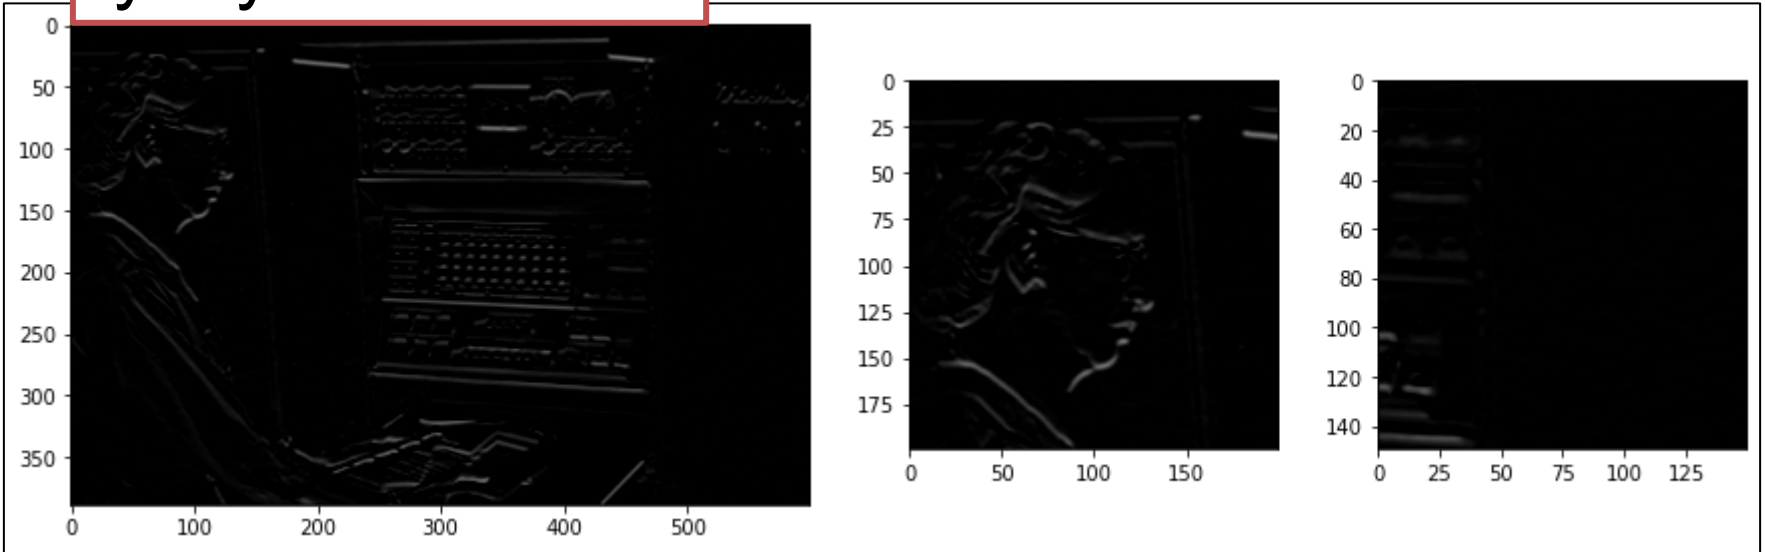


$(x,y) \in W$

$$= [u, v] \mathbf{M} [u, v]^T$$

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix}$$

$I_y = y$ derivative



$(x, y) \in W$

$$= [u, v] \mathbf{M} [u, v]^T$$

$$\mathbf{M} = \begin{bmatrix} \sum_{x, y \in W} I_x^2 & \sum_{x, y \in W} I_x I_y \\ \sum_{x, y \in W} I_x I_y & \sum_{x, y \in W} I_y^2 \end{bmatrix}$$

Sum goes over all the pixels in window W:

$$\sum_{x,y \in W} I_x^2 = \sum_{x,y \in W} (I_x[y, x])^2$$

i.e., sum of squares of x gradients in window

$$= [u, v] \mathbf{M} [u, v]^T$$

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix}$$

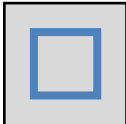


Intuitively what is M?

Pretend gradients are *either* vertical or horizontal

**Obviously
Wrong!**

at a pixel (so $I_x I_y = 0$)

$$M = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix} \approx \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}$$

a,b both small:	flat		$\begin{bmatrix} 0.1 & 0 \\ 0 & 0.1 \end{bmatrix}$
One big, other small:	edge		$\begin{bmatrix} 50 & 0 \\ 0 & 0.1 \end{bmatrix}$ or $\begin{bmatrix} 0.1 & 0 \\ 0 & 50 \end{bmatrix}$
a,b both big:	corner		$\begin{bmatrix} 50 & 0 \\ 0 & 50 \end{bmatrix}$

Intuitively what is M?

Pretend gradients are *either* vertical or horizontal at a pixel (so $I_x I_y = 0$)

$$M = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix} \approx? \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}$$

a,b both small:

flat





Image might be rotated by rotation θ !


One big,
other small:

edge



a,b both big:

corner




Intuitively what is M?

Pretend gradients are *either* vertical or horizontal at a pixel (so $I_x I_y = 0$)

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix} = \mathbf{V}^{-1} \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix} \mathbf{V}$$

a,b both small:


flat



If image rotated by rotation θ / matrix \mathbf{V}

One big,
other small:

edge




\mathbf{M} will look like

$$\mathbf{V}^{-1} \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix} \mathbf{V}$$

a,b both big:

corner



So What Now?

Can calculate \mathbf{M} at pixel, by summing nearby gradients, but need access to a and b .

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix} = \mathbf{V}^{-1} \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix} \mathbf{V}$$

Given \mathbf{M} , can decompose it into eigenvectors \mathbf{V} and eigenvalues λ_1, λ_2 with $\mathbf{M} = \mathbf{V}^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \mathbf{V}$.

Really slow. Why?

So What Now?

Can calculate M at pixel, by summing nearby gradients, but need access to a and b .

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix} = \mathbf{V}^{-1} \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix} \mathbf{V}$$

Instead: compute quantity R from \mathbf{M}

$$R = \det(\mathbf{M}) - \alpha \text{trace}(\mathbf{M})^2 = \lambda_1 \lambda_2 - \alpha (\lambda_1 + \lambda_2)^2$$

Easy fast formula
for 2x2

Fast – sum the diagonal

Empirical value,
usually 0.04-0.06

The tl;dr

TL;DR: Taylor expansion for error $E(u,v)$. All terms in equation are sums of image gradients and in \mathbf{M}

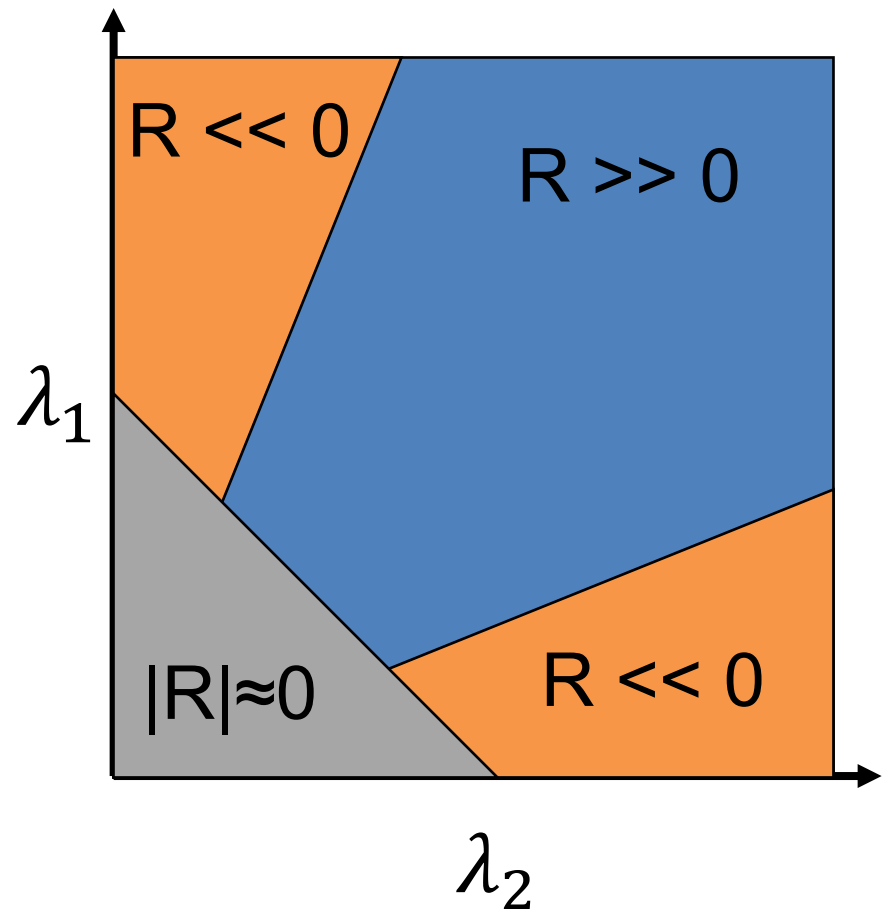
<p><u>Should know</u></p> <p>Can compute at each pixel</p> $\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix} =$ <p>$I_x = I_x$ at point (x,y), $I_y = I_y$ at point (x,y)</p>	<p><u>Optional</u></p> <p>Directions</p> <p>Amounts</p> $\mathbf{V}^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \mathbf{V}$
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------

Putting It Together

$$R = \det(\mathbf{M}) - \alpha \text{trace}(\mathbf{M})^2 = \lambda_1 \lambda_2 - \alpha (\lambda_1 + \lambda_2)^2$$

det, trace are *fast*

flat		$\lambda_1, \lambda_2 \approx 0$
edge		$\lambda_1 \gg \lambda_2 \gg 0$ $\lambda_2 \gg \lambda_1 \gg 0$
corner		$\lambda_1 \approx \lambda_2 \gg 0$



In Practice

1. Compute partial derivatives I_x , I_y per pixel
2. Compute \mathbf{M} at each pixel, using Gaussian weighting w

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} w(x,y) I_x^2 & \sum_{x,y \in W} w(x,y) I_x I_y \\ \sum_{x,y \in W} w(x,y) I_x I_y & \sum_{x,y \in W} w(x,y) I_y^2 \end{bmatrix}$$

C.Harris and M.Stephens. [“A Combined Corner and Edge Detector.”](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

In Practice

1. Compute partial derivatives I_x , I_y per pixel
2. Compute \mathbf{M} at each pixel, using Gaussian weighting w
3. Compute response function R


$$\begin{aligned} R &= \det(\mathbf{M}) - \alpha \operatorname{trace}(\mathbf{M})^2 \\ &= \lambda_1 \lambda_2 - \alpha (\lambda_1 + \lambda_2)^2 \end{aligned}$$

C.Harris and M.Stephens. [“A Combined Corner and Edge Detector.”](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

Computing R



Computing R




In Practice

1. Compute partial derivatives I_x , I_y per pixel
2. Compute \mathbf{M} at each pixel, using Gaussian weighting w
3. Compute response function R
4. Threshold R

C.Harris and M.Stephens. ["A Combined Corner and Edge Detector."](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

Thresholded R




In Practice

1. Compute partial derivatives I_x , I_y per pixel
2. Compute \mathbf{M} at each pixel, using Gaussian weighting w
3. Compute response function R
4. Threshold R
5. Take only local maxima (called non-maxima suppression)

C.Harris and M.Stephens. ["A Combined Corner and Edge Detector."](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

Thresholded



Final Results



Desirable Properties

If our detectors are repeatable, they should be:

- **Invariant** to some things: image is transformed and corners remain the same
- **Covariant/equivariant** with some things: image is transformed and corners transform with it.

Recall Motivating Problem

Images may be different in lighting and geometry




Affine Intensity Change

$$I_{new} = aI_{old} + b$$


M only depends on derivatives, so b is irrelevant

But a scales derivatives and there's a threshold



Partially invariant to affine intensity changes


Image Translation



All done with convolution. Convolution is translation equivariant.

Equivariant with translation

Image Scaling




One pixel can become many pixels and vice-versa.

Not equivariant with scaling

How do we fix this?

Recap: Motivation



1: find corners+features

2: match based on local image data



How?

Today

- Fixing scaling by making detectors in both location **and scale**
- Enabling matching between features by **describing regions**

Key Idea: Scale Space

Left to right: each image is half-sized
Upsampled with big pixels below







Note: I'm also slightly blurring to prevent aliasing (<https://en.wikipedia.org/wiki/Aliasing>)

Key Idea: Scale Space

Left to right: each image is half-sized

If I apply a $K \times K$ filter, how much of the original image does it see in each image?




Note: I'm also slightly blurring to prevent aliasing (<https://en.wikipedia.org/wiki/Aliasing>)

Solution to Scales

Try them all!

Harris Detection




Harris Detection



Harris Detection



Harris Detection




Blob Detection

Another detector (has some nice properties)




$$* \text{ [blob filter] } =$$



Find maxima *and minima* of blob filter response in
scale *and space*


Gaussian Derivatives

Gaussian




1st Deriv

$$\frac{\partial}{\partial y} g$$




2nd Deriv


$$\frac{\partial^2}{\partial^2 y} g$$




$$\frac{\partial}{\partial x} g$$



$$\frac{\partial^2}{\partial^2 x} g$$



Laplacian of Gaussian (LoG)




Slight detail: for technical reasons, you need to scale the Laplacian of Gaussian if you want to compare across sigmas.

$$\nabla_{norm}^2 = \sigma^2 \left(\frac{\partial^2}{\partial x^2} g + \frac{\partial^2}{\partial y^2} g \right)$$


Edge Detection with LoG

f




Edge

$\frac{\partial^2}{\partial^2 x} g$




Laplacian
Of Gaussian
(LoG)

$f * \frac{\partial^2}{\partial^2 x} g$





Edge =
Zero-crossing

Edge Detection with LoG




Edges




Edges *
LoG =
Zero-crossings

Edge Detection with LoG




Edges



f

What happens if we make input 1 unit wide?


$$f * \frac{\partial^2}{\partial x^2} g$$





Edges *

LoG =
Zero-crossings

Edge Detection with LoG




Edge



Edge *
LoG =
Zero-crossing

Scale Selection


Given binary circle and Laplacian filter of scale σ , we can compute the response as a function of the scale.




Characteristic Scale

Characteristic scale of a blob is the scale that produces the maximum response

Image



Abs. Response



Scale-space blob detector

1. Convolve image with scale-normalized Laplacian at several scales

Scale-space blob detector: Example




Scale-space blob detector: Example



sigma = 11.9912

Scale-space blob detector

1. Convolve image with scale-normalized Laplacian at several scales
2. Find maxima of squared Laplacian response in scale-space



Finding Maxima

Point i,j is maxima (minima if you flip sign) in image I if it's bigger than all neighbors

```
for y=range(i-1,i+1+1):
```

```
    for x in range(j-1,j+1+1):
```

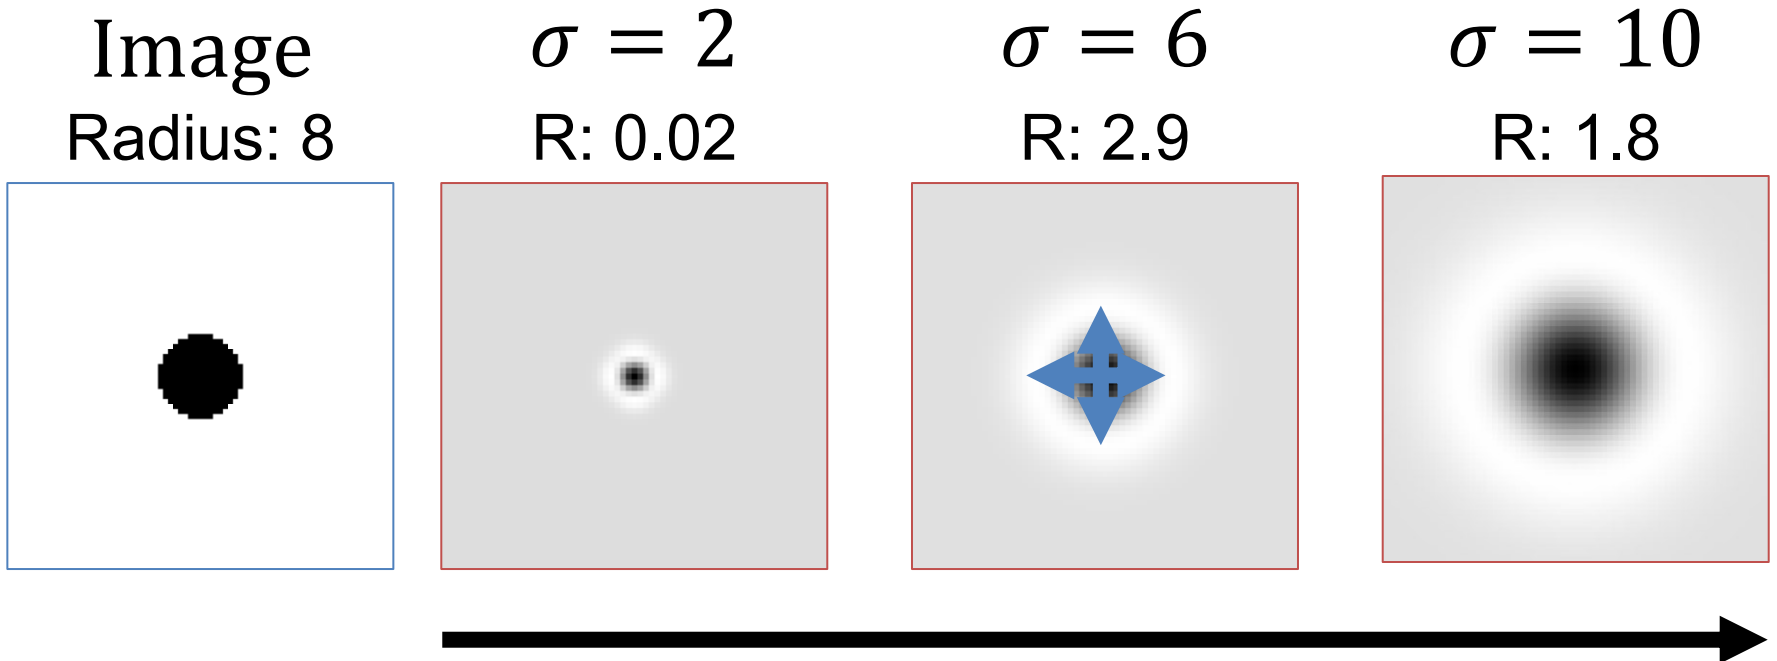
```
        if y == i and x == j: continue
```

```
        #below has to be true
```

```
        I[y,x] < I[i,j]
```


Scale Space

Blue lines are image-space neighbors (should be just one pixel over but that's impossible to draw)



Scale Space

Red lines are the scale-space neighbors



Finding Maxima

Suppose $I[:, :, k]$ is image at scale k . Point i, j, k is maxima (minima if you flip sign) in image I if:

```
for y=range(i-1,i+1+1):
```

```
    for x in range(j-1,j+1+1):
```

```
        for c in range(k-1,k+1+1):
```


```
            if y == i and x == j and c == k:
```

```
                continue
```

```
            #below has to be true
```

```
            I[y,x,c] < I[i,j,k]
```

Scale-space blob detector: Example



Efficient implementation


- Approximating the Laplacian with a difference of Gaussians:

$$L = \sigma^2 \left(G_{xx}(x, y, \sigma) + G_{yy}(x, y, \sigma) \right)$$


(Laplacian)

$$DoG = G(x, y, k\sigma) - G(x, y, \sigma)$$

(Difference of Gaussians)



Efficient implementation



David G. Lowe. ["Distinctive image features from scale-invariant keypoints."](#) *IJCV* 60 (2), pp. 91-110, 2004.

Slide credit: S. Lazebnik

Problem 1 Solved

- How do we deal with scales: try them all
- **Why is this efficient?**

Vast majority of effort is in the first and second scales

$$1 + \frac{1}{4} + \frac{1}{16} + \frac{1}{64} + \frac{1}{4^i} \dots = \frac{4}{3}$$

Problem 2 – Describing Features

Image – 40

1/2 size, rot. 45°
Lightened+40

Full
Image



100x100 crop
at Glasses




Problem 2 – Describing Features

Once we've found a corner/blobs, we can't just use the image nearby. What about:

1. Scale?
2. Rotation?
3. Additive light?


Handling Scale

Given characteristic scale (maximum Laplacian response), we can just rescale image




Handling Rotation

Given window, can compute “dominant orientation”
and then rotate image




Scale and Rotation


SIFT features at characteristic scales and dominant orientations



Scale and Rotation



SIFT Descriptors



1. Compute gradients
2. Build histogram (2x2 here, 4x4 in practice)

Gradients ignore global illumination changes

SIFT Descriptors

- In principle: build a histogram of the gradients
- In reality: quite complicated
 - Gaussian weighting: smooth response
 - Normalization: reduces illumination effects
 - Clamping
 - Tons of more stuff

Properties of SIFT

- Can handle: up to ~ 60 degree out-of-plane rotation, changes of illumination
- Fast, efficient, code available (but was patented)





Feature Descriptors

Think of feature as some non-linear filter that maps pixels to 128D feature

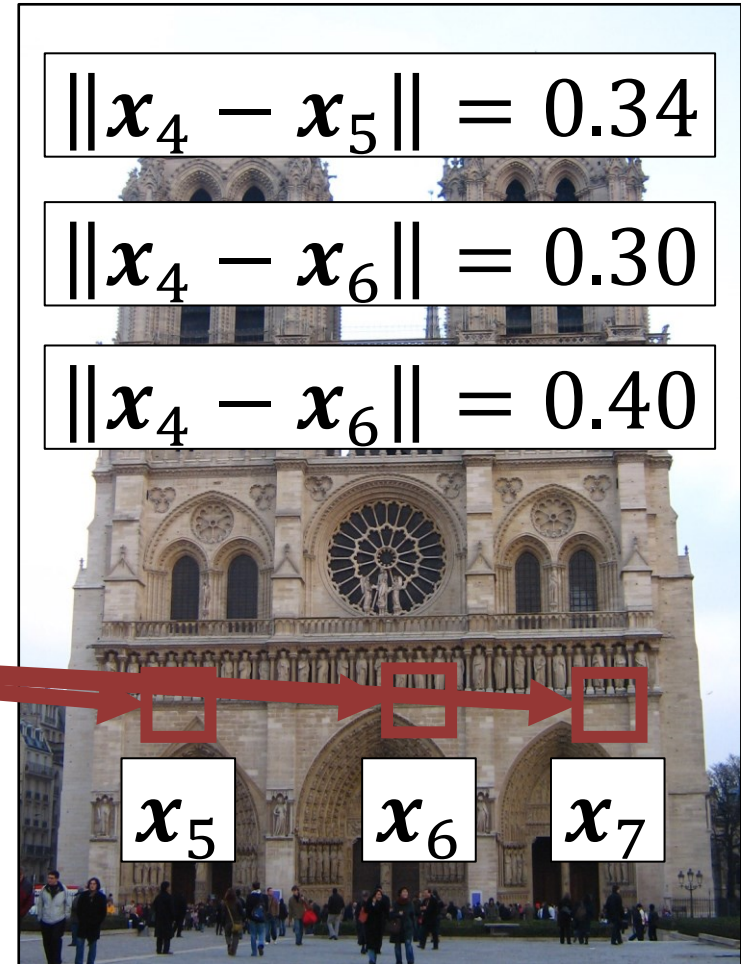



128D
vector x

Instance Matching



Instance Matching




2nd Nearest Neighbor Trick

- Given a feature \mathbf{x}_q , nearest neighbor to \mathbf{x} is a good match, but distances can't be thresholded.
- Instead, find nearest neighbor (\mathbf{x}_{1NN}) and second nearest neighbor (\mathbf{x}_{2NN}). This ratio is a good test for matches:

$$r = \frac{\|\mathbf{x}_q - \mathbf{x}_{1NN}\|}{\|\mathbf{x}_q - \mathbf{x}_{2NN}\|}$$

So Far; What's Next?



1: find corners+features

2: match based on local image data

3: next time: compute offsets from matches


Extra Reading for the Curious

Aside: This Trick is Common

Given a 50×16 person detector, how do I detect:
(a) 250×80 (b) 150×48 (c) 100×32 (d) 25×8 people?



Sample people from image




Aside: This Trick is Common

Detecting all the people
The red box is a fixed size




Sample people from image



Aside: This Trick is Common

Detecting all the people
The red box is a fixed size


Sample people from image



Aside: This Trick is Common

Detecting all the people
The red box is a fixed size

Sample people from image




Affine adaptation

Consider the second moment matrix of the window containing the blob:

$$M = \sum_{x,y} w(x,y) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$$


Recall:

$$\begin{bmatrix} u & v \end{bmatrix} M \begin{bmatrix} u \\ v \end{bmatrix} = \text{const}$$



This ellipse visualizes the “characteristic shape” of the window

Affine adaptation example



Scale-invariant regions (blobs)

Affine adaptation example



Affine-adapted blobs

2nd Nearest Neighbor Trick

