#### **Image Structure** Feature Detection and Matching 4F12: Computer Vision

#### Instructor: Samuel Albanie

Based on course material authored by Roberto Cipolla



#### **Representing Images via Intensities**



Image credit: "pixel geometry", https://commons.wikimedia.org/wiki/File:Pixel\_geometry\_01\_Pengo.jpg

#### Challenge: Nuisance factors in image data



Reference: pixel geometry, "Cafe" by avinashbhat is licensed with CC BY-SA 2.0.

- If a point on an object (in this case, a torta di mele) is visible in view, the intensity I(x, y) is a function of many geometric and photometric variables:
- The position and orientation of the camera
- The geometry of the scene (3D shapes and layout)
- The nature and distribution of light sources
- Reflectance properties of the surfaces: specular-Lambertian, albedo 0 (black) - 1 (white)
- The properties of the camera lens and sensor array

In practice, the point may only be partially visible, or its appearance may also be affected by occlusion.



## Challenge: Data reduction



File:Apple\_iSight\_FireWire\_Camera.jpg

Goals for <u>generic features</u>:

- Allow the image to be discarded, so that all subsequent processing is done on the features themselves. This enables us to <u>reduce the amount of data</u>
- Preserve the useful information in the images (such as the albedo changes and 2D shape of objects in the scene)
- Discard the redundant information in the images (such as the lighting conditions).
- <u>As generic as possible</u> (so the same processing will be useful across a wide range of applications).

OpenAl's powerful 2021 CLIP model Research Trivia (trained on hundreds of GPUs) still only uses up to 448 pixel images



## **Computer vision as hierarchical processing**



Reference: D. Marr, "Vision: A computational investigation into the human representation and processing of visual information" (1982) J. Hawkins, S. Ahmad, and Y. Cui, "A theory of how columns in the neocortex enable learning the structure of the world." Frontiers in neural circuits (2017)



#### Image structure



Let's examine pixel values in three patches in this photo of Claire:

- A featureless region
- An edge
- A corner

124	127	1.38	128	131	312	135	14
1 29	130	132	132	133	135	139	12
138	136	1.37	135	135	136	130	5
1++	142	143	1+1	139	139	92	6
150	152	151	148	138	113	79	8
158	160	160	154	133	113	111	12
163	165	166	158	145	146	147	15
168	171	1.74	173	169	171	172	17
167	171	176	177	176	178	150	15
166	173	179	182	182	184	185	18
166	173	179	183	183	185	189	19

Note that an edge or corner representation imparts a desirable invariance to lighting: the intensity discontinuities are likely to be prominent, whatever the lighting conditions. Computational question: How can we find these structures efficiently?



## **1D Edge Detection**

When developing an edge detection algorithm, it is important to bear in mind the invariable presence of image noise.

Consider this signal I(x) with an obvious edge:



An intuitive approach to edge detection might be to look for maxima and minima in I'(x).



Oh dear: it's hard to spot the edge in this signal!

Our simple strategy was defeated by high-frequency noise (which is amplified by differentiation).

For this reason, all edge detectors start by smoothing the signal to <u>suppress noise</u>.

The most common approach is to use a <u>Gaussian filter</u> (a lowpass filter that suppresses high frequencies).



## **1D Edge Detection (with smoothing)**

1D Edge Detection algorithm

- A broad overview of 1D edge detection is:
- 1. First, convolve the signal I(x) with a Gaussian kernel  $g_{\sigma}(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right)$  to produce smooth s(x).
- 2. Compute s'(x), the derivative of s(x).
- 3. Find maxima and minima of s'(x).
- 4. Use thresholding on the magnitude of the extrema to mark edges.



### **1D Edge Detection: a computational trick**

The smoothing in step 1 was performed by a 1D convolution with a Gaussian:

$$s(x) = I(x) \circledast g_{\sigma}(x) = \int_{-\infty}^{\infty} g_{\sigma}(u)I(x-u)du = \int_{-\infty}^{\infty} g_{\sigma}(x-u)I(u)du$$

The differentiation in step 2 is also performed by a 1D convolution. So it seems that edge detection requires two computationally expensive convolutions.

However, the derivative theorem of convolution comes to the rescue!

$$s'(x) = \frac{d}{dx} [g_{\sigma}(x) \circledast I(x)] = g'_{\sigma}(x) \circledast I(x)$$

So we can compute s'(x) with just a single convolution - a <u>major saving!</u>

du



## **1D Edge Detection (faster)**

Faster 1D Edge Detection algorithm

The fast variant of the edge detection algorithm becomes

- 1. Convolve the signal I(x) with a derivative of Gaussian Kernel  $g'_{\sigma}(x)$  to produce s'(x) directly.
- 2. Find maxima and minima of s'(x).
- 3. Use thresholding on the magnitude of the extrema to mark edges.

Figure credits: S. Seitz



## **1D Edge Detection: zero-crossings**

Fastest 1D Edge Detection algorithm

Finding maxima and minima of s'(x) is the same as looking for <u>zero-crossings</u> of s''(x)!

In many implementations of edge detection algorithms, the signal is convolved with the Laplacian of a Gaussian ("LoG" kernel),  $g''_{\sigma}(x)$ , by applying the derivative theorem of convolution a second time:

 $s''(x) = g_{\sigma}''(x) \circledast I(x)$ 

The zero crossings of s''(x) mark possible edges.

Figure credits: S. Seitz



## 1D Edge Detection: scale



#### Consider this signal:



Does the signal have one "positive" edge or a number of "positive" and "negative" edges?

It's up to you to choose!

Using a small  $\sigma$  brings out all the edges.



#### Fine detail edges

#### Coarse detail edges As $\sigma$ increases, the signal is smoothed more and more, and only the central edge survives. signal I(x) 1200 1000 1400 1600 1800 800 Derivative of Gaussian kernel $g'_{\sigma}$ $\sigma = 50$ Kernel 0 800 1000 1200 1400 1600 1800 2000 200 400 600 **Derivative** s'(x)

1000

800

1200

1400

1600

1800

ð,

0

200

400

600



## 1D Edge Detection: multi-scale

#### The link between smoothing and scale

The amount of smoothing controls the scale at which we analyse the image. There is no right or wrong size for the Gaussian kernel: it all depends on the scale we're interested in.

Modest smoothing (a Gaussian kernel with small  $\sigma$ ) brings out edges at a fine scale. More smoothing (larger  $\sigma$ ) identifies edges at larger scales, suppressing the finer detail.



<u>Note</u>: Fine scale edge detection is particularly sensitive to noise (less of an issue when analysing images at coarse scales).

# **2D Edge Detection (step 1: smoothing)**

Step 1: smoothing

*K* 

The 1D edge detection scheme can be extended to work in two dimensions.

First we smooth the image I(x, y) by convolving with a 2D Gaussian  $G_{\sigma}(x, y)$ :

$$G_{\sigma}(x,y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$

$$S(x, y) = G_{\sigma}(x, y) \circledast I(x, y)$$
$$= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G_{\sigma}(u, v)I(x - u, y - v)dudv$$



#### Effects of Gaussian smoothing





Original image

 $\sigma = 3$  pixels

 $\sigma = 4$  pixels



# 2D Edge Detection (step 2: gradients)

Step 2: gradients

The second step is to compute the gradient of the smoothed image S(x, y) at every pixel:

$$7S = \nabla (G_{\sigma} \circledast I)$$
$$= \left[ \frac{\partial (G_{\sigma} \circledast I)}{\partial x} \right]_{\partial (G_{\sigma} \circledast I)} = \left[ \frac{\partial G_{\sigma}}{\partial x} \right]_{\partial G_{\sigma}}$$

$$\frac{\partial (G_{\sigma} \circledast I)}{\partial y} = \frac{\partial G_{\sigma}}{\partial y} \circledast I$$

\*

#### Example: fruity gradients



Original image



Edge strength  $|\nabla S|$ 

#### 2D Edge Detection (step 3: NMS & step 4: thresholding)

Step 3: Non-Maxima Suppression

The third stage of the edge detection algorithm is Non-Maxima Suppression (NMS).

Edge elements, or edgels, are placed at locations where  $|\nabla S|$  is greater than local values of  $|\nabla S|$ in the directions  $\pm \nabla S$ . This aims to ensure that all edgels are located at ridge-points of the surface  $|\nabla S|$ .



Edge strength  $|\nabla S|$  after NMS

Step 4: Thresholding

In the fourth and final step, the edgels are thresholded, so that only those with  $|\nabla S|$  above a certain value are retained.



Edge strength  $|\nabla S|$  after NMS and thresholding

## 2D Edge Detection (variations)

Canny Edge Detection

The edge detection algorithm we have been describing is due to Canny (1986).

The output is a list of edgel positions, each with a strength  $|\nabla S|$  and an orientation  $\nabla S/|\nabla S|$ .

The Canny detector is a <u>directional</u> edge finder (both the gradient magnitude and <u>direction</u> are computed)

**References:** 

Canny, J. A computational approach to edge detection. TPAMI 1986 D. Marr and E. Hildreth, "Theory of edge detection." Proc. of the Royal Society of London. Series B. Biological Sciences, 1980

Marr-Hildreth Edge Detection

An alternative approach to edge detection was developed by Marr and Hildreth (1980).

Unlike the directional Canny edge detector, the Marr-Hildreth operator is isotropic.

It finds zero-crossings of  $\nabla^2 G_{\sigma} \circledast I$ , where  $\nabla^2 G_{\sigma}$  is the Laplacian of  $G_{\sigma}$  (recall the Laplacian operator  $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ ).  $\nabla^2 G_{\sigma}$ 

## **Edge Detection: Implementation details**



Truncation: how much is acceptable?

For acceptable accuracy, kernels are generally truncated so that the discarded samples are less than 1/1000 of the peak value.

	1.0	1.5	3	6
L	7	11	23	45

Another computational trick!

The 2D convolutions would appear to be computationally expensive.

However, they can be decomposed into two 1D convolutions:

 $G_{\sigma}(x, y) \circledast I(x, y) = g_{\sigma}(x) \circledast [g_{\sigma}(y) \circledast I(x, y)]$ 

The computational saving is:

 $\frac{(2n+1)^2}{2(2n+1)}$ 



### **Edge Detection: Implementation details**

Differentiation via convolution

Differentiation of the smoothed image is also implemented with a discrete convolution.

By considering the Taylor-series expansion of S(x, y) one can show that a simple finite-difference approximation to the first-order spatial derivative of S(x, y)with respect to x is given by:

$$\frac{\partial S}{\partial x} = \frac{S(x+1,y) - S(x-1,y)}{2}$$

VE

Implementing first order derivatives

We can calculate the finite-difference approximation to  $\partial S/\partial x$  by convolving the rows of smoothed image samples, S(x, y), with the 3-element kernel:

1/2 0 -1/2

Recall that when convolving, we flip the kernel and sum the element-wise products under each kernel position:



**Note:** this is often called "valid" convolution (the kernel is not allowed to run off the edges of S(x, y)).

# End of Image Structures Lecture 1

#### Appendix

- GoPro 10 raw data rate "5.3K at 60 fps" = colour depth × width × height × fps = stored is 100 Mbits/sec)
- iSight camera (24-bit, 480p @ 30 fps): = 24 × 640 × 480 × 30 = 221 Mbit/s.

Much of the content of these slides is based on material by Roberto Cipolla. For other cases, I have tried to credit the figures where possible, but finding the original sources can be challenging (useful figures naturally propagate across many slide decks!) Please let me know if you spot a missing reference.

Back-of-the-envelope calculations for <u>uncompressed</u> data rates (slide 4)

 $8 \times 5312 \times 2988 \times 60 = 7.6$  Gbits/second (in practice, though, the maximum bitrate that can be

#### Slide content credits

Useful reference: https://en.wikipedia.org/wiki/Uncompressed\_video