



CSCI 497P/597P: Computer Vision



Lecture 3: Convolution and its Properties

CSCI 497P/597P: Computer Vision



Lecture 3: Convolution and its Properties

Announcements

- HW1 is out today
 - Covers filtering and convolution
 - Due in 1 week (10pm next Tuesday)

Goals

- Understand the distinction between cross-correlation and [convolution](#).
- Know the properties of cross-correlation and convolution:
 - Linearity and shift-invariance (both)
Associativity and commutativity (convolution only)
- Understand the design of several common image filters:
 - Box blur and Gaussian blur
 - Sharpening
- Understand the limitations of linear filtering


Computing Cross-Correlation

$$g = f \otimes w$$

output image \swarrow f \nwarrow weights, or filter, or kernel w

input image \swarrow

```
for x = 0 to w:  
  for y = 0 to h:  
    for i in -k to k:  
      for j in -k to k:  
        out[x,y] += w[i,j] * in[x+i, y+j]
```




Computing Cross-Correlation

$$g = f \otimes w$$

output image \nearrow f \nwarrow input image

weights, or filter, or kernel \nwarrow w

```
for x = 0 to w:  
  for y = 0 to h:  
    for i in -k to k:  
      for j in -k to k:  
        out[x,y] += w[i,j] * in[x+i, y+j]
```



A bit of practice

In groups: work on Problems #1-4

- Problems are linked from the course webpage on the Schedule table
- Write answers in your Google Doc (pinned in group Discord channels)

$$\begin{pmatrix} 3 & 6 & 3 \\ 4 & 8 & 4 \\ 3 & 6 & 3 \end{pmatrix} \quad \begin{pmatrix} 4 & 8 & 4 \\ 4 & 8 & 4 \\ 4 & 8 & 4 \end{pmatrix} \quad [8]$$

$$\begin{array}{cccc|cccc}
 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
 3 & 4 & 2 & 6 & 3 & 4 & 3 & 0
 \end{array}$$

$$\begin{array}{ccc}
 1 & 2 & 1 \\
 2 & 4 & 2 \\
 1 & 2 & 1
 \end{array}$$

Questions remain

- What happens at the edges?
- What properties does this operator have?
- What can and can't this operator do?

A shift filter

Cross-correlate the image f with the kernel w .

Use "same" output size, with zero-padding for out-of-bounds values.

1	2	1
1	3	1
0	1	3

f

 \otimes

0	0	0
0	0	1
0	0	0

w

 $=$

Cross-correlation vs Convolution

- Cross-correlation: $g = f \otimes w$

$$g(x, y) = \sum_{i=-k}^k \sum_{j=-k}^k w(i, j) f(x + i, y + j)$$

\uparrow \uparrow

- Convolution: $g = f * w$

$$g(x, y) = \sum_{i=-k}^k \sum_{j=-k}^k w(i, j) f(x - i, y - j)$$

\downarrow \downarrow

These are related:

$$f * w = f \otimes \text{flip horiz}(\text{flip vert}(w))$$

Properties

Assume: f is an image; w and v are filters; s, t are scalars.

filter img = shifting, filter shift result

Shift invariance (both) $f(x, y) \otimes w = [f(x - s, y - t) \otimes w](x - s, y - t)$

Linearity (both)

and

$$(f \otimes w) + (f \otimes v) = f \otimes (w + v)$$

$$(f \otimes sw) = s(f \otimes w)$$

Commutativity (conv only)

$$f * w = w * f$$

$$f \otimes w \neq w \otimes f$$

expensive

expensive!


Associativity (conv only)

$$(f * w) * v = f * (w * v)$$

cheap!

What can we do with this?


What can we do with this?



0	0	0
0	1	0
0	0	0


=

What can we do with this?




Identity filter: output = input

What can we do with this?




What can we do with this?



left shift

What can we do with this?




$*$ $\frac{1}{9}$

1	1	1
1	1	1
1	1	1



=

What can we do with this?



mean filter, or
box blur


Blurring: Another example



What motivated the mean filter?

Notice: lattice-like texture

Blurring: Another example





Notice: lattice-like texture

What motivated the mean filter?

Idea: the closer the pixel, the more likely it is to be similar

Gaussian Blur

- Idea: weight closer pixels more heavily using a Gaussian kernel:





This is a bivariate (2D) Gaussian function:

$$G_{\sigma} = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$

Gaussian Blur

- Idea: weight closer pixels more heavily using a /
Gaussian kernel:



$$\frac{1}{16} \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array}$$

3x3 approximation


This is a bivariate (2D) Gaussian function:

$$G_{\sigma} = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$



Gaussian Filters





$\sigma = 1$ pixel



$\sigma = 5$ pixels





$\sigma = 10$ pixels




$\sigma = 30$ pixels

Mean vs. Gaussian




Composing Filters

- Recall associativity:



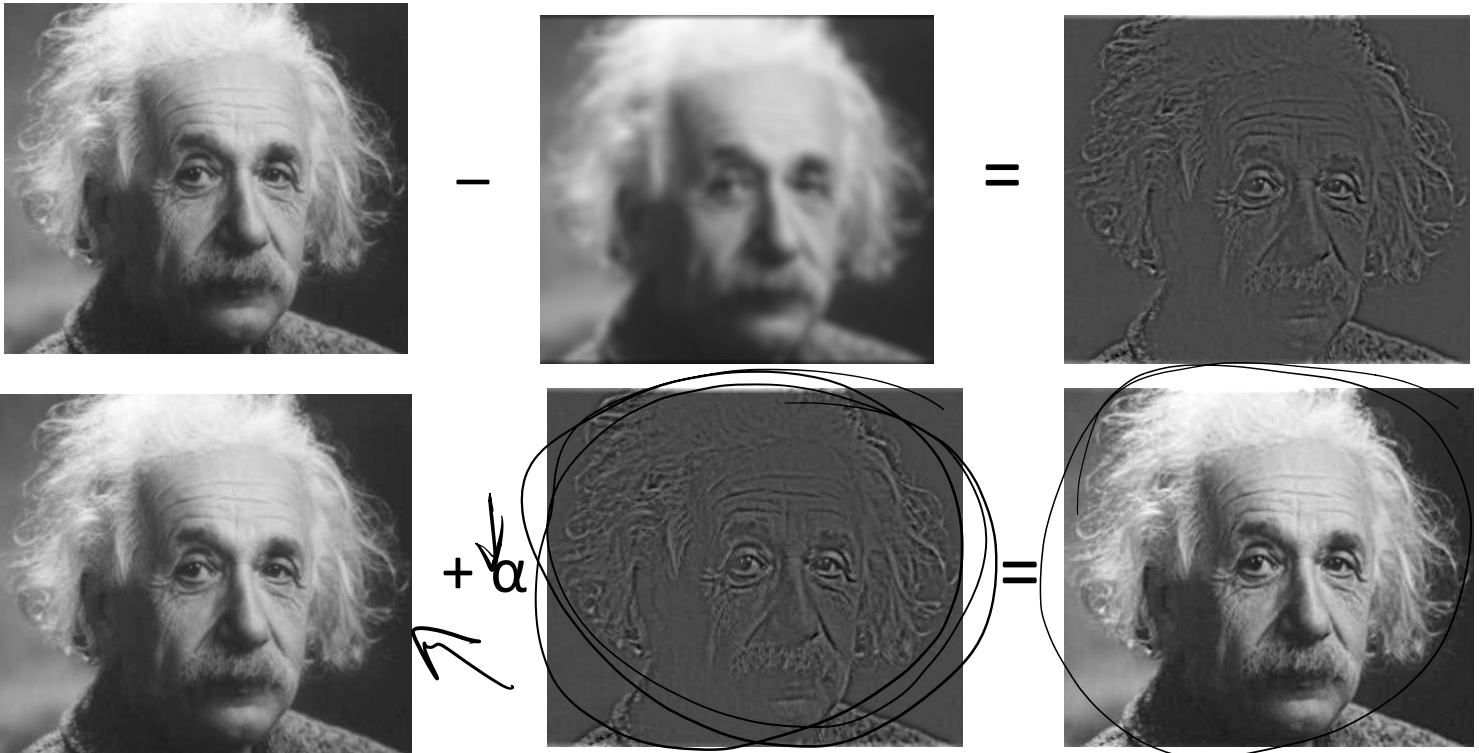
Sharpening!?

- What gets removed when we blur?



Sharpening!?



- What gets removed when we blur?




Sharpening

Before




After



Sharpening: once more with mathing



$$+ (\text{ } - \text{ }) =$$



image

$$f + f - (f * B)$$

blur filter

introduce identity filter

$$2f - f * B$$


$$(f * 2I) - (f * B)$$

Use Linearity





$$f * (2I - B)$$

Sharpening Filter


Sharpening: once more with mathing







$$+ (\quad - \quad) =$$




Sharpening: once more with mathing




$$+ (\quad - \quad) =$$



—



≈



0	-1	0
-1	5	-1
0	-1	0

**Possible 3x3
Approximation**



Effects of Sharpening

Original

Sharpened




Limitations:

What *can't* convolution do?

Problem #5 - discuss in groups.

- Maximum filter?



- Threshold?



- y partial derivative?

