

All about convolution

# Last time: Convolution and cross-correlation

- Cross correlation

$$S[f] = w \otimes f$$
$$S[f](m, n) = \sum_{i=-k}^k \sum_{j=-k}^k w(i, j) f(m + i, n + j)$$

- Convolution

$$S[f] = w * f$$
$$S[f](m, n) = \sum_{i=-k}^k \sum_{j=-k}^k w(i, j) f(m - i, n - j)$$

# Cross correlation

0	10	5	7	0
5	11	6	8	3
9	22	4	5	1
2	9	14	6	7
3	10	15	12	9

Local image data



Kernel size =  $2k+1$

$$S[f](m, n) = \sum_{i=-k}^k \sum_{j=-k}^k w(i, j) f(m + i, n + j)$$

# Properties: Linearity

$$(w \otimes f)(m, n) = \sum_{i=-k}^k \sum_{j=-k}^k w(i, j) f(m + i, n + j)$$

$$f' = af + bg$$

$$w \otimes f' = a(w \otimes f) + b(w \otimes g)$$

# Properties: Linearity

$$(w \otimes f)(m, n) = \sum_{i=-k}^k \sum_{j=-k}^k w(i, j) f(m + i, n + j)$$

$$w' = aw + bv$$

$$w' \otimes f = a(w \otimes f) + b(v \otimes f)$$

# Shift **equivariance**

$$f'(m, n) = f(m - m_0, n - n_0)$$

$$(w \otimes f')(m, n) = (w \otimes f)(m - m_0, n - n_0)$$

- Shift, then convolve = convolve, then shift
- Output of convolution does not depend on where the pixel is



$f$



$f'$

# Boundary conditions

$$(w * f)(m, n) = \sum_{i=-k}^k \sum_{j=-k}^k w(i, j) f(m - i, n - j)$$

- What if  $m-i < 0$ ?
- What if  $m-i > \text{image size}$
- Assume  $f$  is defined for  $[-\infty, \infty]$  in both directions, just 0 everywhere else
- Same for  $w$

$$(w * f)(m, n) = \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} w(i, j) f(m - i, n - j)$$

# Boundary conditions

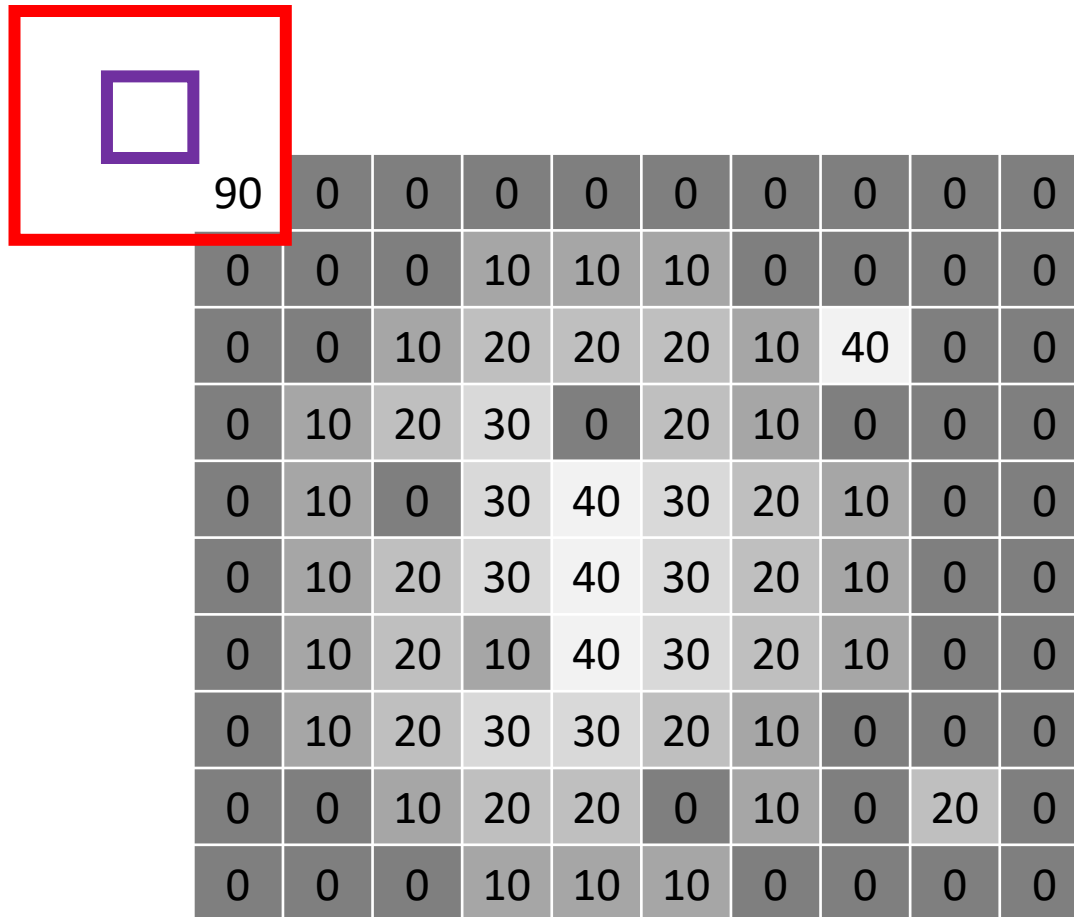
90	0	0	0	0	0	0	0	0	0
0	0	0	10	10	10	0	0	0	0
0	0	10	20	20	20	10	40	0	0
0	10	20	30	0	20	10	0	0	0
0	10	0	30	40	30	20	10	0	0
0	10	20	30	40	30	20	10	0	0
0	10	20	10	40	30	20	10	0	0
0	10	20	30	30	20	10	0	0	0
0	0	10	20	20	0	10	0	20	0
0	0	0	10	10	10	0	0	0	0



# Boundary conditions

90	0	0	0	0	0	0	0	0	0
0	0	0	10	10	10	0	0	0	0
0	0	10	20	20	20	10	40	0	0
0	10	20	30	0	20	10	0	0	0
0	10	0	30	40	30	20	10	0	0
0	10	20	30	40	30	20	10	0	0
0	10	20	10	40	30	20	10	0	0
0	10	20	30	30	20	10	0	0	0
0	0	10	20	20	0	10	0	20	0
0	0	0	10	10	10	0	0	0	0


# Boundary conditions



The diagram illustrates boundary conditions for a numerical grid. A red square highlights a purple square, which is positioned above the first cell of the first row. The value 90 is written below the purple square. The grid consists of 10 rows and 10 columns of cells, with the top-left cell containing the value 90. The values in the grid are as follows:


90	0	0	0	0	0	0	0	0	0
0	0	0	10	10	10	0	0	0	0
0	0	10	20	20	20	10	40	0	0
0	10	20	30	0	20	10	0	0	0
0	10	0	30	40	30	20	10	0	0
0	10	20	30	40	30	20	10	0	0
0	10	20	10	40	30	20	10	0	0
0	10	20	30	30	20	10	0	0	0
0	0	10	20	20	0	10	0	20	0
0	0	0	10	10	10	0	0	0	0

# Boundary conditions



90	0	0	0	0	0	0	0	0	0
0	0	0	10	10	10	0	0	0	0
0	0	10	20	20	20	10	40	0	0
0	10	20	30	0	20	10	0	0	0
0	10	0	30	40	30	20	10	0	0
0	10	20	30	40	30	20	10	0	0
0	10	20	10	40	30	20	10	0	0
0	10	20	30	30	20	10	0	0	0
0	0	10	20	20	0	10	0	20	0
0	0	0	10	10	10	0	0	0	0

# Boundary conditions




90	0	0	0	0	0	0	0	0	0
0	0	0	10	10	10	0	0	0	0
0	0	10	20	20	20	10	40	0	0
0	10	20	30	0	20	10	0	0	0
0	10	0	30	40	30	20	10	0	0
0	10	20	30	40	30	20	10	0	0
0	10	20	10	40	30	20	10	0	0
0	10	20	30	30	20	10	0	0	0
0	0	10	20	20	0	10	0	20	0
0	0	0	10	10	10	0	0	0	0

# Boundary conditions in practice

- “Full convolution”: compute if *any* part of kernel intersects with image
  - requires padding
  - Output size =  $m+k-1$
- “Same convolution”: compute if center of kernel is in image
  - requires padding
  - output size =  $m$
- “Valid convolution”: compute only if *all* of kernel is in image
  - no padding
  - output size =  $m-k+1$

# Filters: examples




Original (f)



$\frac{1}{9}$


1	1	1
1	1	1
1	1	1

Kernel (k)



Blur (with a mean filter) (g)

# Filters: examples




Original (f)




0	0	0
0	1	0
0	0	0

Kernel (k)




Identical image (g)

# Sharpening



**before**




**after**




# Sharpening


- What does blurring take away?



Let's add it back:




# Sharpening



# Sharpening

- What does blurring take away?




Let's add it back:




# Sharpening

$$\begin{aligned}f_{sharp} &= f + \alpha(f - f_{blur}) \\&= (1 + \alpha)f - \alpha f_{blur} \\&= (1 + \alpha)(w * f) - \alpha(v * f)\end{aligned}$$



$$= ((1 + \alpha)w - \alpha v) * f$$

# Sharpening filter



Original

$$* \left( \begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 2 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array} - \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array} \right) =$$




**Sharpening filter**  
(accentuates edges)



Another example



# Another example





# More properties of convolution

$$\begin{aligned}(w * f)(m, n) &= \sum_i \sum_j w(i, j) f(m - i, n - j) \\ &= \sum_i \sum_j w(m - i', n - j') f(i, j) \\ &= (f * w)(m, n)\end{aligned}$$

$$\begin{aligned}i' &= m - i \Rightarrow i = m - i' \\ j' &= n - j \Rightarrow j = n - j'\end{aligned}$$

# More properties of convolution

- Convolution is linear
- Convolution is shift-invariant
- Convolution is commutative ( $w * f = f * w$ )
- Convolution is *associative* ( $v * (w * f) = (v * w) * f$ )
- Every linear shift-invariant operation is a convolution

# More convolution filters

- Mean filter


1/25

1	1	1	1	1
1	1	1	1	1
1	1	1	1	1
1	1	1	1	1
1	1	1	1	1


- But nearby pixels are more correlated than far-away pixels
- Weigh nearby pixels more

# Gaussian filter

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



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



# Gaussian filter

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

- Ignore factor in front, instead, normalize filter to sum to 1


0.003	0.013	0.022	0.013	0.003
0.013	0.060	0.098	0.060	0.013
0.022	0.098	0.162	0.098	0.022
0.013	0.060	0.098	0.060	0.013
0.003	0.013	0.022	0.013	0.003

5x5,  $\sigma=1$

# Gaussian filter



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

- Ignore factor in front, instead, normalize filter to sum to 1




5x5,  $\sigma=1$


# Gaussian filter



21x21,  $\sigma=0.5$




21x21,  $\sigma=1$




21x21,  $\sigma=3$




# Difference of Gaussians



21x21,  $\sigma=1$



21x21,  $\sigma=3$





# Time complexity of convolution

- Image is  $w \times h$
- Filter is  $k \times k$
- Every entry takes  $O(k^2)$  operations
- Number of output entries:
  - $(w+k-1)(h+k-1)$  for full
  - $wh$  for same
- Total time complexity:
  - $O(whk^2)$


# Optimization: separable filters

- basic alg. is  $O(r^2)$ : large filters get expensive fast!
- definition:  $w(x,y)$  is *separable* if it can be written as:

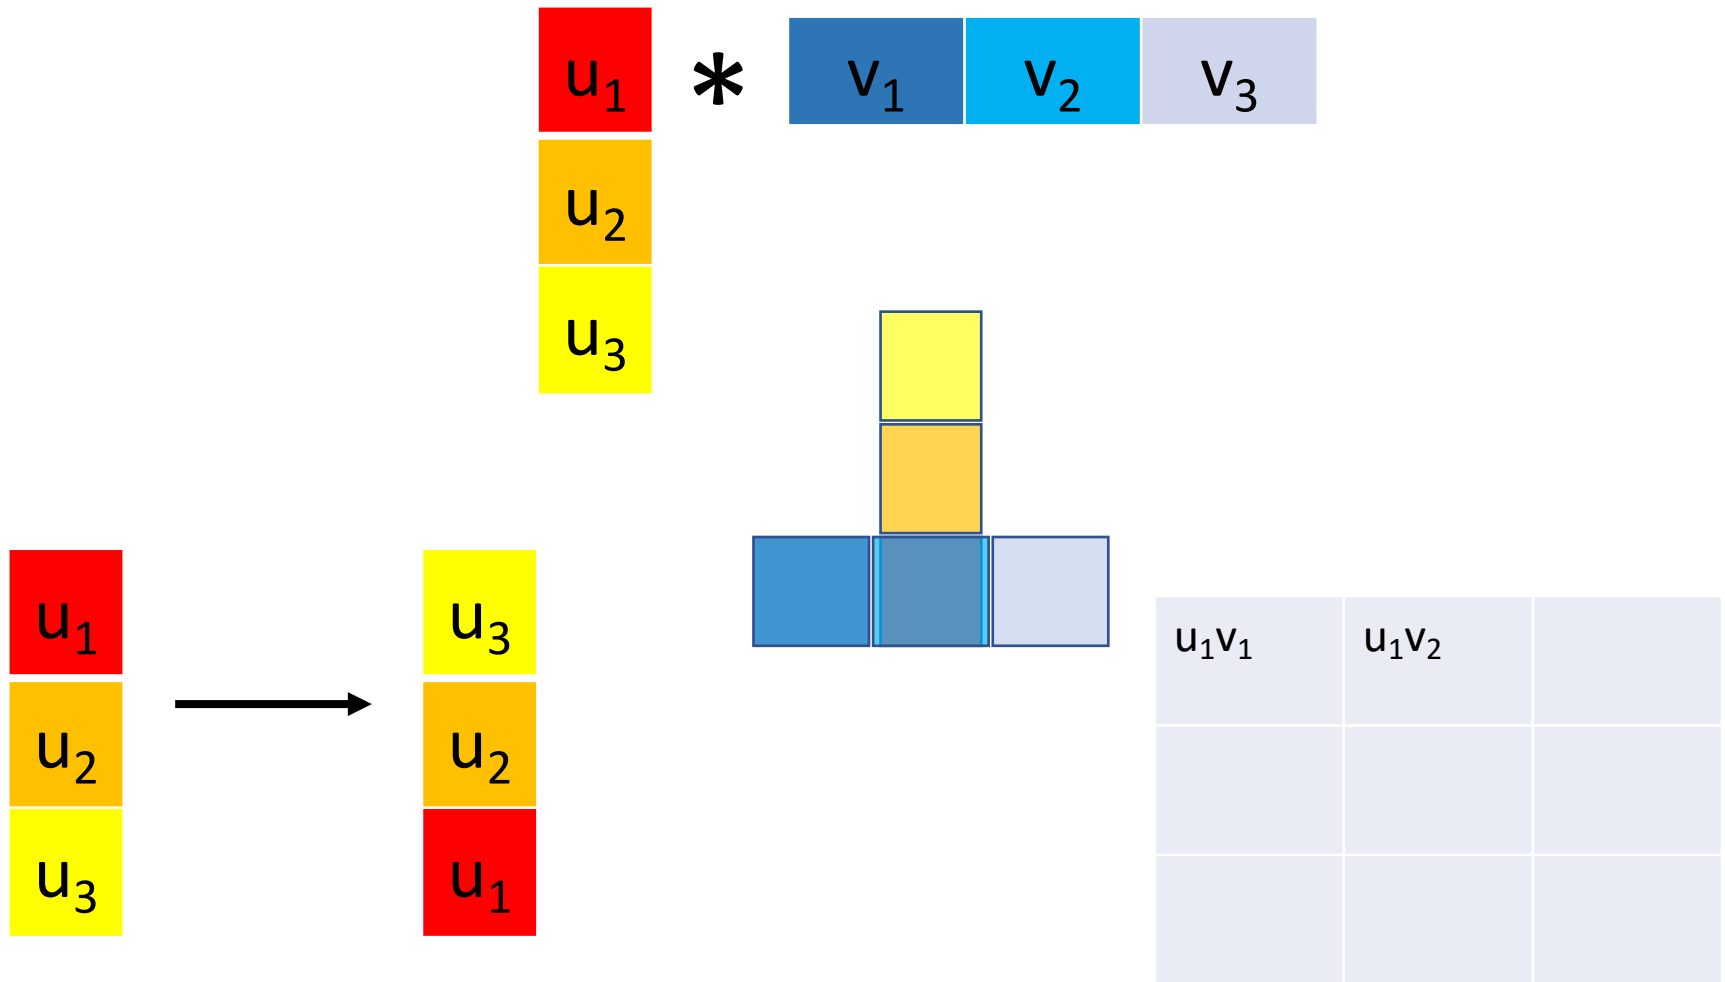
$$w(i, j) = u(i)v(j)$$

- Write  $u$  as a  $k \times 1$  filter, and  $v$  as a  $1 \times k$  filter
- Claim:  $w = u * v$


# Separable filters






# Separable filters



# Separable filters






# Separable filters






$u_1v_1$	$u_1v_2$	$u_1v_3$
$u_2v_1$		

# Separable filters



$u_1v_1$	$u_1v_2$	$u_1v_3$
$u_2v_1$	$u_2v_2$	




# Separable filters



$u_1v_1$	$u_1v_2$	$u_1v_3$
$u_2v_1$	$u_2v_2$	$u_2v_3$






# Separable filters



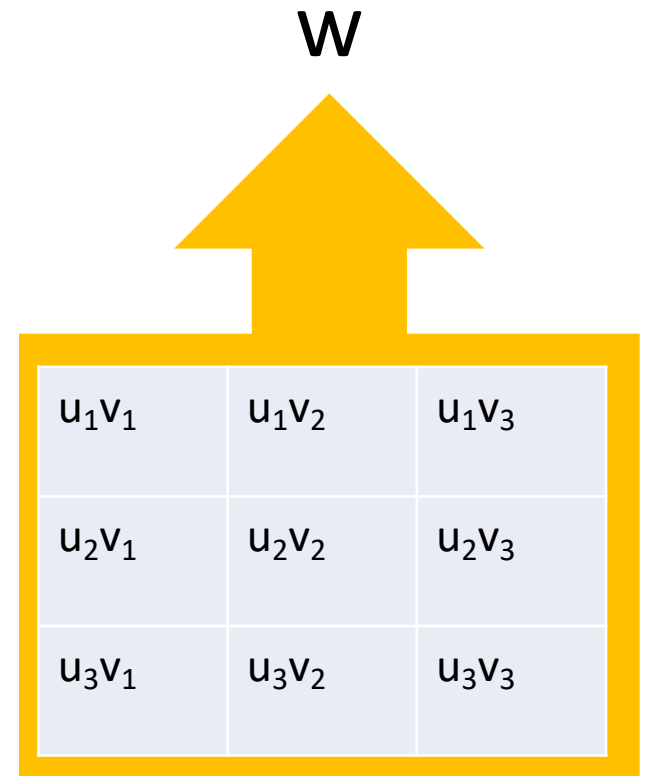



$u_1v_1$	$u_1v_2$	$u_1v_3$
$u_2v_1$	$u_2v_2$	$u_2v_3$
$u_3v_1$		

# Separable filters



$u_1v_1$	$u_1v_2$	$u_1v_3$
$u_2v_1$	$u_2v_2$	$u_2v_3$
$u_3v_1$	$u_3v_2$	

# Separable filters



# Separable filters

$$\begin{aligned}w * f &= (u * v) * f \\ &= u * (v * f)\end{aligned}$$

- Time complexity of original :  $O(whk^2)$
- Time complexity of separable version :  $O(whk)$

Images have structure at various scales

