

# Convolutional Neural Networks

Yuzhen Ye


School of Informatics and Computing, Indiana University

# Contents


- CNN basics
- CNN for visual recognition (to explain the concept of convolution)
- CNN for bioinformatics

# A Beginner's Guide To Understanding CNN

- Neural Networks with **Convolution layers**



A Full Convolutional Neural Network (LeNet)




Ref 1: <http://cs231n.stanford.edu>

Ref 2: <https://adeshpande3.github.io/adeshpande3.github.io/A-Beginner%27s-Guide-To-Understanding-Convolutional-Neural-Networks/>

# Convolution Layer

32x32x3 image



5x5x3 filter



**Convolve** the filter with the image  
i.e. “slide over the image spatially,  
computing dot products”

# Moving Average In 2D

$f[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$g[x, y]$

	0	10	20	30	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0	30	60	90	90	90	60	30	
	0	30	50	80	80	90	60	30	
	0	30	50	80	80	90	60	30	
	0	20	30	50	50	60	40	20	
	10	20	30	30	30	30	20	10	
	10	10	10	0	0	0	0	0	

# Correlation filtering

Say the averaging window size is  $2k+1 \times 2k+1$ :

$$g(i, j) = \frac{1}{(2k+1)^2} \sum_{u=-k}^k \sum_{v=-k}^k f(i+u, j+v)$$

*Attribute uniform weight to each pixel*      *Loop over all pixels in neighborhood around image pixel  $f[i,j]$*

Now generalize to allow different weights depending on neighboring pixel's relative position:

$$g(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k \underbrace{h(u, v)}_{\text{Non-uniform weights}} f(i+u, j+v)$$

*Non-uniform weights*

# Correlation filtering

$$g(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k h(u, v) f(i + u, j + v)$$


This is called cross-correlation, denoted  $g = h \otimes f$

Filtering an image: replace each pixel with a linear combination of its neighbors.

The filter “kernel” or “mask”  $h[u, v]$  is the prescription for the weights in the linear combination.

# Averaging filter

- What values belong in the kernel  $h$  for the moving average example?



$$g = h \otimes f$$



# Gaussian filter

- What if we want nearest neighboring pixels to have the most influence on the output?


0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$f[x, y]$

$$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} h[u, v]$$

This kernel is an approximation of a 2d Gaussian function:

$$h(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{\sigma^2}}$$



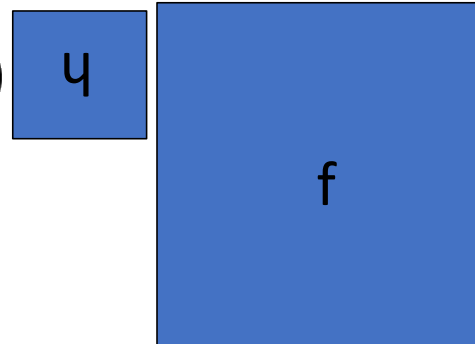
# Convolution

- Convolution is a simple mathematical operation which is fundamental to many common image processing operators.
- Convolution is performed by sliding the kernel over the image, generally starting at the top left corner, so as to move the kernel through all the positions where the kernel fits entirely within the boundaries of the image.
- Convolution:
  - Flip the filter in both dimensions (bottom to top, right to left)
  - Then apply cross-correlation

$$g(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k h(u, v) f(i - u, j - v)$$

$$g = h * f$$

↑  
*Notation for convolution operator*



# Convolution vs. correlation

Convolution

$$g(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k h(u, v) f(i - u, j - v)$$

$$g = h * f$$


Cross-correlation

$$g(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k h(u, v) f(i + u, j + v)$$

$$g = h \otimes f$$

# Derivatives and edges

An edge is a place of rapid change in the image intensity function.



# Derivatives with convolution

For 2D function,  $f(x,y)$ , the partial derivative is:

$$\frac{\partial f(x, y)}{\partial x} = \lim_{\varepsilon \rightarrow 0} \frac{f(x + \varepsilon, y) - f(x, y)}{\varepsilon}$$

For discrete data, we can approximate using finite differences:

$$\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x + 1, y) - f(x, y)}{1}$$





To implement above as convolution, what would be the associated filter?

# Partial derivatives of an image



$$\frac{\partial f(x, y)}{\partial x}$$

$$\frac{\partial f(x, y)}{\partial y}$$



Which shows changes with respect to x?

# Filters as feature (edge) detectors


Prewitt:  $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$

Sobel:  $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$

Roberts:  $M_x = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$



sobel filter



Slide credit: Kristen Grauman




<http://homepages.inf.ed.ac.uk/rbf/HIPR2/index.htm>

# Image gradient

The gradient of an image:

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

The gradient points in the direction of most rapid change in intensity



The **gradient direction** (orientation of edge normal) is given by:

$$\theta = \tan^{-1} \left( \frac{\partial f / \partial y}{\partial f / \partial x} \right)$$

The **edge strength** is given by the gradient magnitude

$$\|\nabla f\| = \sqrt{\left( \frac{\partial f}{\partial x} \right)^2 + \left( \frac{\partial f}{\partial y} \right)^2}$$







# Effects of noise

Consider a single row or column of the image

- Plotting intensity as a function of position gives a signal

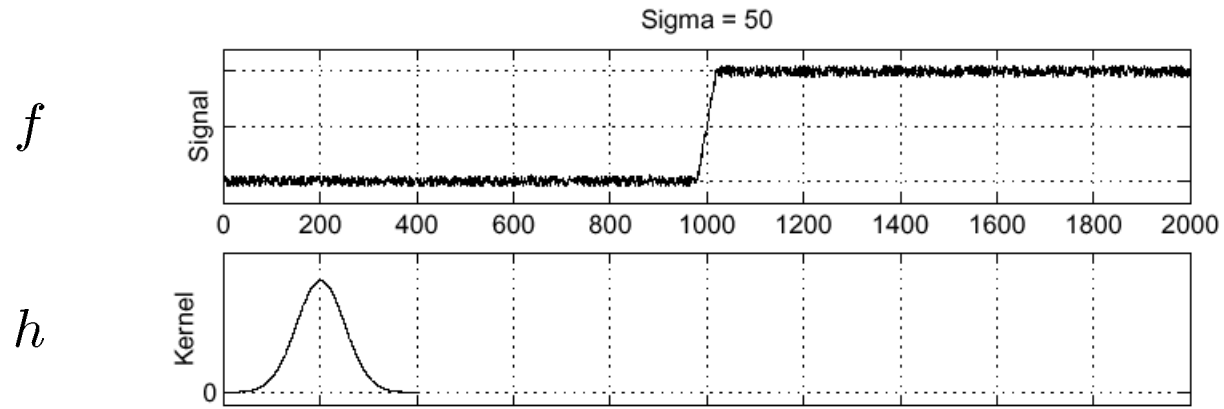


Where is the edge?

# Effects of noise

- Difference filters respond strongly to noise
  - Image noise results in pixels that look very different from their neighbors
  - Generally, the larger the noise the stronger the response
- What can we do about it?

# Solution: smooth first




Where is the edge?  
Slide credit: Kristen Grauman

Look for peaks in  $\frac{\partial}{\partial x}(h \star f)$

# Derivative theorem of convolution

$$\frac{\partial}{\partial x}(h \star f) = \left(\frac{\partial}{\partial x}h\right) \star f$$


Differentiation property of convolution.




# Laplacian of Gaussian

Consider  $\frac{\partial^2}{\partial x^2}(h \star f)$


$f$



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



$(\frac{\partial^2}{\partial x^2}h) \star f$




Where is the edge?

Slide credit: Steve Seitz


Zero-crossings of bottom graph

# 2D edge detection filters



**Gaussian**


$$h_{\sigma}(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$



**derivative of Gaussian**

$$\frac{\partial}{\partial x} h_{\sigma}(u, v)$$

**Laplacian of Gaussian**



$$\nabla^2 h_{\sigma}(u, v)$$

0	-1	0
-1	4	-1
0	-1	0

-1	-1	-1
-1	8	-1
-1	-1	-1

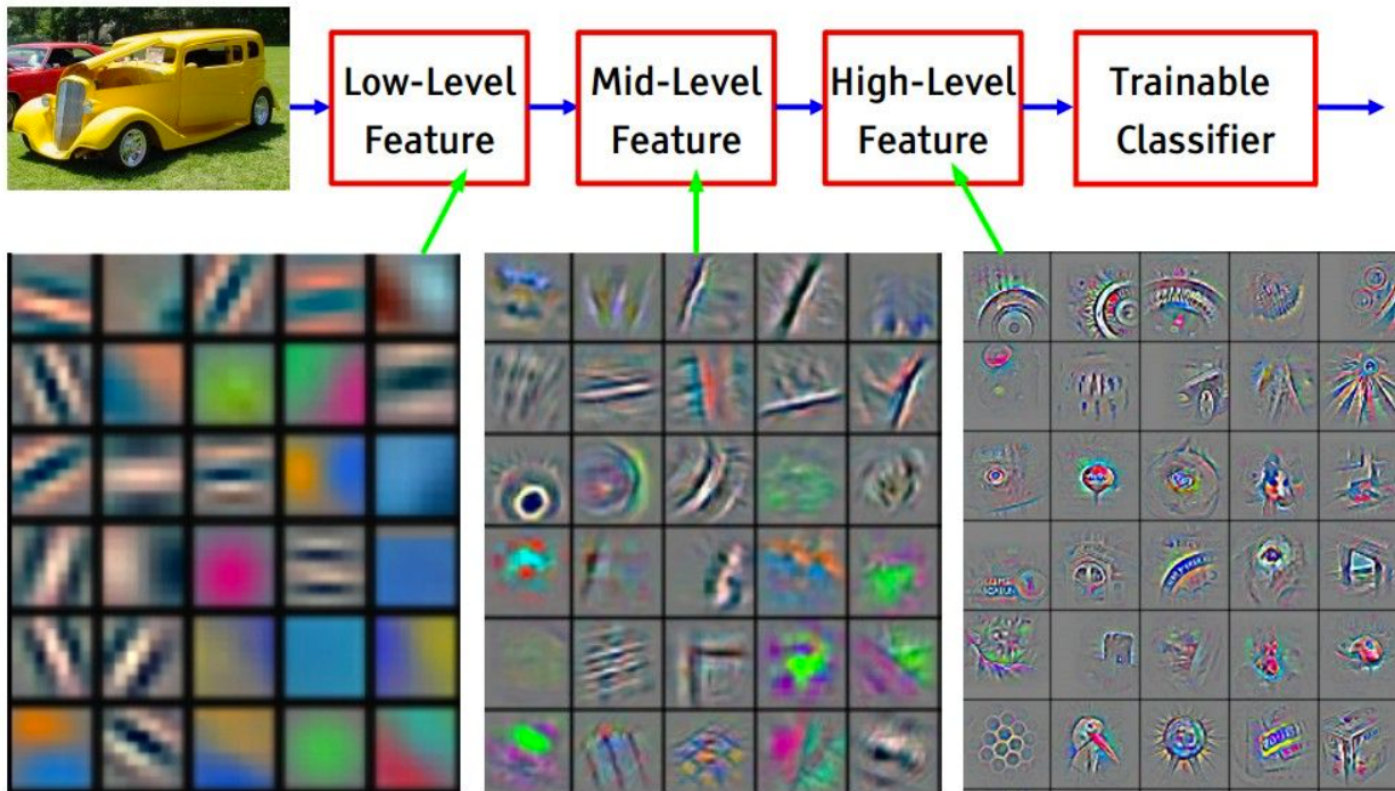
- $\nabla^2$  is the Laplacian operator:

$$\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

Two commonly used discrete approximations to the Laplacian filter

# Preview


[From recent Yann LeCun slides]



Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

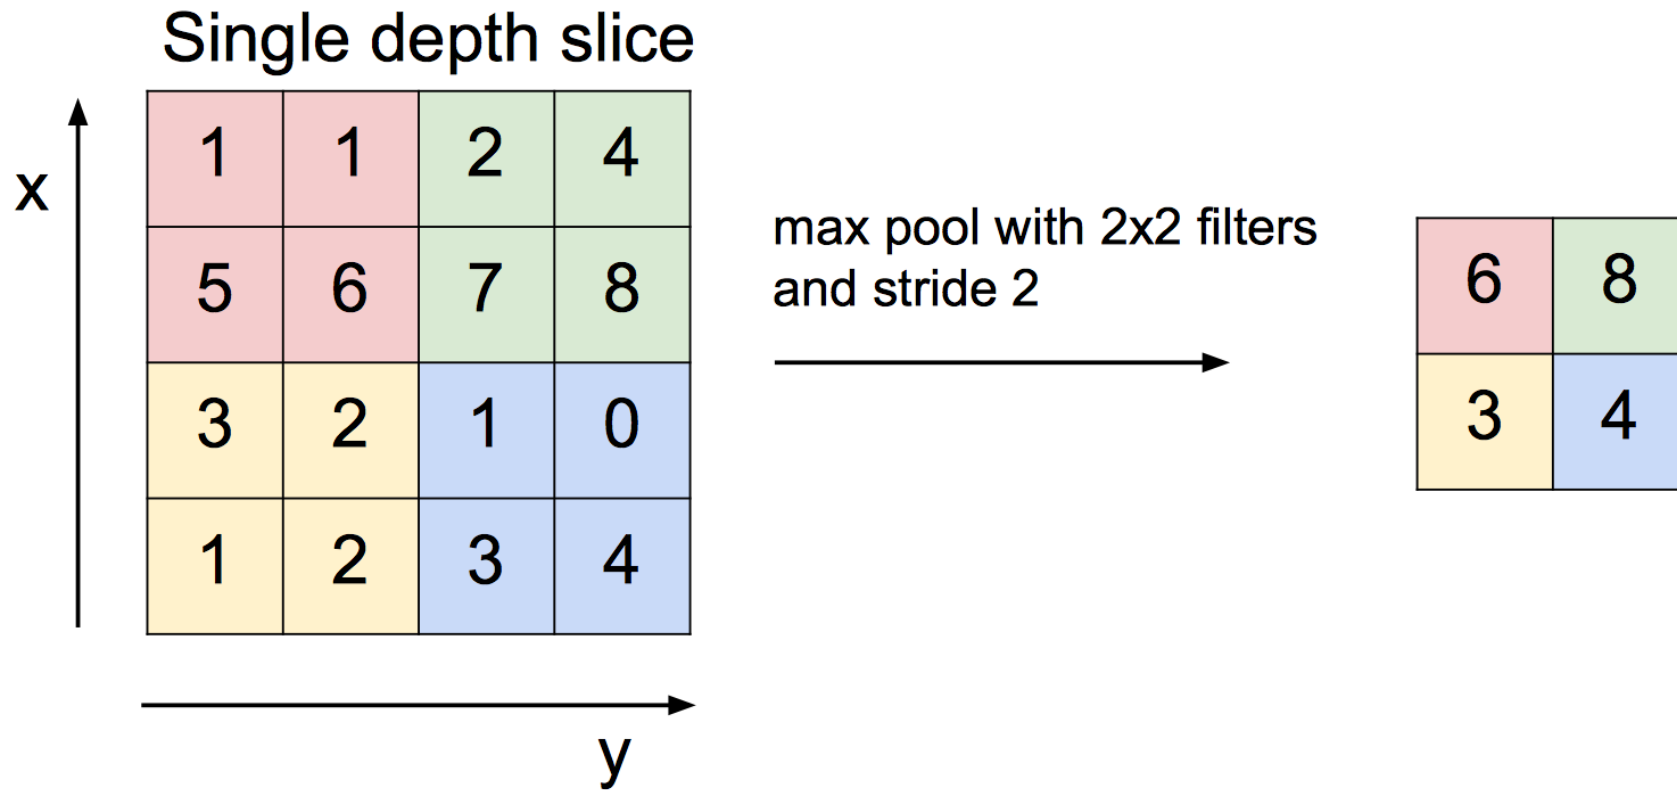
# Pooling layer

- makes the representations smaller and more manageable
- operates over each activation map independently:

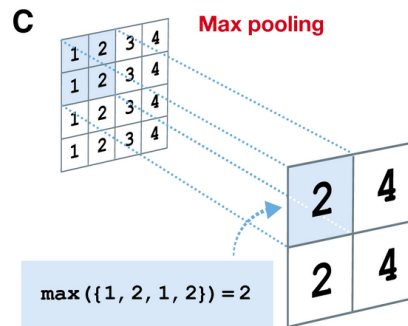

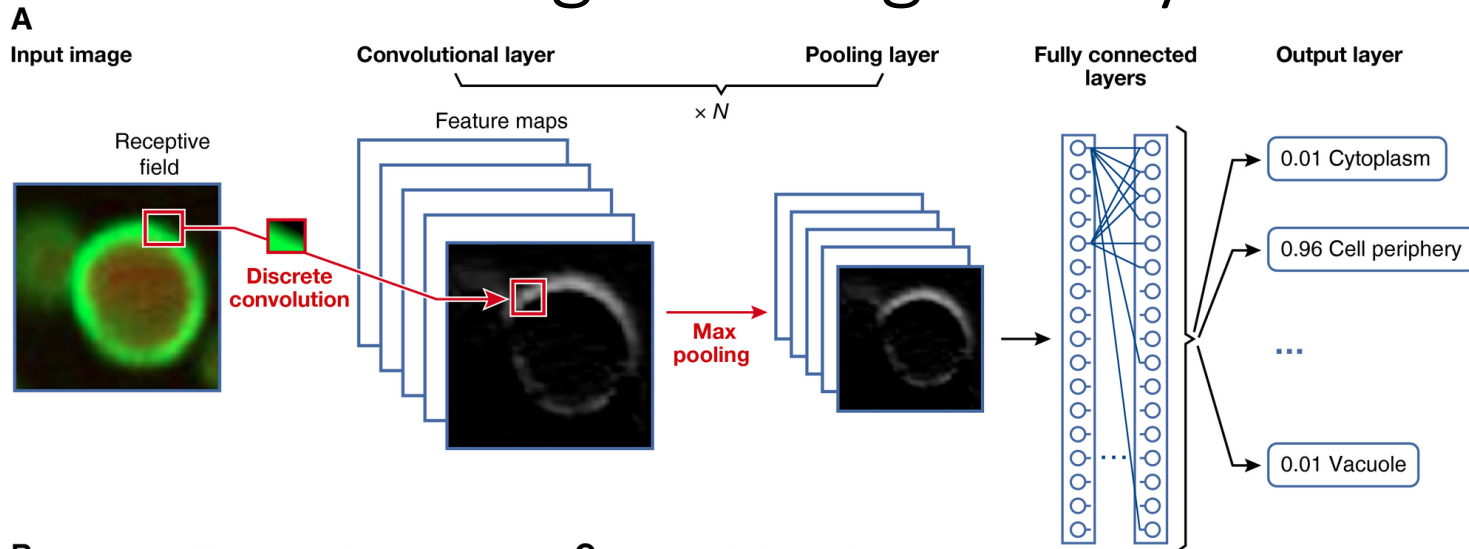




# MAX POOLING




# CNN for biological image analysis





<http://onlinelibrary.wiley.com/doi/10.15252/msb.20156651/full#msb156651-fig-0002>


# CNN for predicting molecular traits



Input data: one-dimensional genomic sequences with one channel per nucleotide



[Visual recognition: 2D-image with three color channels]



Molecular Systems Biology

Volume 12, Issue 7, 29 JUL 2016 DOI: 10.15252/msb.20156651

<http://onlinelibrary.wiley.com/doi/10.15252/msb.20156651/full#msb156651-fig-0002>