

# CPSC 340: Machine Learning and Data Mining

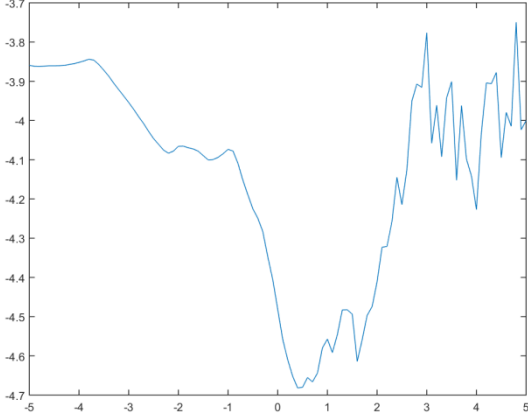
Convolutional Neural Networks

Fall 2018

# Admin


- Mike and I finish CNNs on Wednesday.
- After that, we will cover different topics:
  - Mike will do a demo of [training CNNs with cloud/GPU](#) resources.
  - I am planning to cover [boosting](#) (the other type of ensemble method).
    - The [lecture will probably be 90 minutes](#) (I won't be offended if you leave early, extra time won't be testable).
- Friday's lectures will also be different:
  - Mike will do a [course review](#) in his section.
  - Aline Tabet will give a [guest lecture](#) in this section ("ML Applications in Medicine").
- Final: Thursday December 13<sup>th</sup> at 8:30am in WOOD 2.
  - Similar style of questions to midterm.
  - 2 pages of notes.
- CPSC 532M students: course project due December 19 (details on Piazza).

# Last Time: Convolutions

- Consider our original “signal”:
- For each “time”:
  - Compute dot-product of signal at surrounding times with a “filter”.

$$w = [-0.1416 \quad -0.1781 \quad -0.2746 \quad 0.1640 \quad 0.8607 \quad 0.1640 \quad -0.2746 \quad -0.1781 \quad -0.1416]$$

- This gives a new “signal”:
  - Measures a property of “neighbourhood”.
  - This particular filter shows a local “how spiky” value.



# 1D Convolution

- 1D convolution example:

- Signal:

$$x = [0 \quad 1 \quad 1 \quad 2 \quad 3 \quad 5 \quad 8 \quad 13]$$

Indices 1 to 8 are shown below the array. A blue box highlights the elements from index 2 to 6. The element at index 4 (value 2) is circled in red. A green bracket under indices 2 to 6 is labeled 'n'.

- Filter:

$$w = [0 \quad -1 \quad 2 \quad -1 \quad 0]$$

Indices  $w_{-2}$ ,  $w_{-1}$ ,  $w_0$ ,  $w_1$ ,  $w_2$  are shown below the array. A blue bracket under the entire filter is shown.

- Convolution:

$$z = [ \quad \quad \quad 0 \quad \quad \quad ]$$

Indices 1 to 8 are shown below the array. The element at index 4 (value 0) is circled in red. A green bracket under indices 2 to 6 is labeled 'n'.

Let's compute  $z_4$ :

$$x_{4-2:4+2} = [1 \quad 1 \quad 2 \quad 3 \quad 5]$$

A blue arrow points from the circled '2' in the signal array to this equation. A blue bracket under the array is shown.

Multiply element-wise

$$[0 \quad -1 \quad 4 \quad -3 \quad 0]$$

↓ Add

$$z_4 = 0 - 1 + 4 - 3 + 0 = \underline{0}$$

# 1D Convolution

- 1D convolution example:

- Signal:

$$x = [0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13]$$

Indices 1 through 8 are marked below the signal. A blue box highlights the segment from index 3 to 8. A red box highlights the value 3 at index 5. A green bracket labeled 'n' spans from index 3 to 5.

- Filter:

$$w = [0 \ -1 \ 2 \ -1 \ 0]$$

Indices  $w_{-2}$ ,  $w_{-1}$ ,  $w_0$ ,  $w_1$ ,  $w_2$  are marked below the filter. A blue arrow points from the  $w_0$  position to the value 3 in the signal above.

- Convolution:

$$z = [ \quad 0 \quad -1 \quad \quad \quad ]$$

Indices 1 through 8 are marked below the convolution result. A red box highlights the value -1 at index 5. A green bracket labeled 'n' spans from index 3 to 5.

Let's compute  $z_5$ :

$$[1 \ 2 \ 3 \ 5 \ 8]$$

A blue arrow points from the value 3 in the signal above to the value 3 in this array.

Multiply ↓

$$[0 \ -2 \ 6 \ -5 \ 0]$$

A blue arrow points from the value 3 in the array above to the value 6 in this array.

↓ Add

$$z_5 = -2 + 6 - 5 = -1$$

# 1D Convolution Examples

- Examples:

- “Identity”

$$\hookrightarrow w = [0 \ 1 \ 0]$$

- “Translation”

$$\hookrightarrow w = [0 \ 0 \ 1]$$

$$\text{Let } x = [0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13]$$

$$z = [0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13]$$

$0 \cdot x_0 + 1 \cdot x_1 + 0 \cdot x_2$        $0 \cdot x_1 + 1 \cdot x_2 + 0 \cdot x_3$

$$z = [1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13 \ ?]$$

$0 \cdot x_0 + 0 \cdot x_1 + 1 \cdot x_2$

# 1D Convolution Examples

- Examples:

- “Identity”

$$\hookrightarrow w = [0 \ 1 \ 0]$$

- “Local Average”

$$\hookrightarrow w = [\frac{1}{3} \ \frac{1}{3} \ \frac{1}{3}]$$

Let  $x = [0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13]$

$$z = [0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13]$$

Diagram illustrating the “Identity” convolution. The input vector  $x$  is shown above the output vector  $z$ . Green brackets and arrows indicate that the output  $z$  is identical to the input  $x$ , with the word “average” written below the arrows.

$$z = [? \ \frac{2}{3} \ \frac{1}{3} \ 2 \ \frac{3}{3} \ \frac{5}{3} \ \frac{8}{3} \ ?]$$

Diagram illustrating the “Local Average” convolution. The input vector  $x$  is shown above the output vector  $z$ . Green brackets and arrows indicate that the output  $z$  is the average of the input  $x$  over a window of size 3. The output values are  $2/3$ ,  $1/3$ ,  $2$ ,  $3/3$ ,  $5/3$ , and  $8/3$ , with question marks for the first and last elements.

# Boundary Issue

- What can we do about the “?” at the edges?

If  $x = [0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13]$  and  $w = [\frac{1}{3} \ \frac{1}{3} \ \frac{1}{3}]$  then  $z = [? \ 2\frac{2}{3} \ 1\frac{1}{3} \ 2 \ 3\frac{1}{3} \ 5\frac{1}{3} \ 8\frac{2}{3} \ ?]$

- Can assign values **past the boundaries**:

- “Zero”:  $x = [0 \ 0 \ 0 \ 0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13] \ 0 \ 0 \ 0$

- “Replicate”:  $x = [0 \ 0 \ 0 \ 0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13] \ 13 \ 13 \ 13$

- “Mirror”:  $x = [2 \ 1 \ 1 \ 0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13] \ 8 \ 5 \ 3$

- Or just ignore the “?” values and **return a shorter vector**:




$$z = [2\frac{2}{3} \ 1\frac{1}{3} \ 2 \ 3\frac{1}{3} \ 5\frac{1}{3} \ 8\frac{2}{3}]$$



# 1D Convolution Examples

- Translation convolution shift signal:



$$w = [1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]$$



# 1D Convolution Examples

- **Averaging** convolution computes local mean:



$$w = \left[ \frac{1}{3} \quad \frac{1}{3} \quad \frac{1}{3} \right]$$



# 1D Convolution Examples

- **Averaging** over bigger window gives coarser view of signal:

$$w = \left[ \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \right]$$






# 1D Convolution Examples

- **Gaussian** convolution blurs signal:  $w_i \propto \exp\left(-\frac{i^2}{2\sigma^2}\right)$ 
  - Compared to averaging it's more smooth and maintains peaks better.

$$W = [0.0001 \quad 0.0644 \quad 0.0540 \quad 0.2420 \quad 0.3989 \quad 0.2420 \quad 0.0540 \quad 0.0644 \quad 0.0001]$$




$(\sigma = 1, m = 4)$



# 1D Convolution Examples

- **Sharpen** convolution enhances peaks.
  - An “average” that places **negative weights** on the surrounding pixels.



$$w = [-1 \quad 3 \quad -1]$$



# 1D Convolution Examples


- **Laplacian** convolution approximates second derivative:
  - “Sum to zero” filters “respond” if input vector looks like the filter

$$w = [-1 \quad 2 \quad -1]$$




# Digression: Derivatives and Integrals

- Numerical derivative approximations can be viewed as filters:
  - Centered difference:  $[-1, 0, 1]$  (derivativeCheck in findMin).



- Numerical integration approximations can be viewed as filters:
  - “Simpson’s” rule:  $[1/6, 4/6, 1/6]$  (a bit like Gaussian filter).



- Derivative filters add to 0, integration filters add to 1,
  - For constant function, derivative should be 0 and average = constant.

# 1D Convolution Examples




- **Laplacian of Gaussian** is a smoothed 2<sup>nd</sup>-derivative approximation:

$$w_i = \left(1 - \frac{i^2}{2\sigma^2}\right) \exp\left(-\frac{i^2}{2\sigma^2}\right)$$

(then subtract mean)

$$w = [-0.1416 \quad -0.1781 \quad -0.2746 \quad 0.1640 \quad 0.8607 \quad 0.1640 \quad -0.2746 \quad -0.1781 \quad -0.1416]$$


$$(\sigma^2 = 1, m = 4)$$






# 1D Convolution Examples

- We often use **maximum over several convolutions** as features:
  - Below is maximum of Laplacian of Gaussian at 'i' and its 16 KNNs.
  - We **use different convolutions as our features** (derivatives, integrals, etc.).



In this case,  
these 2 features  
are all we need.




# Images and Higher-Order Convolution

- **2D convolution:**
  - Signal 'x' is the pixel intensities in an 'n' by 'n' image.
  - Filter 'w' is the pixel intensities in a '2m+1' by '2m+1' image.
- The **2D convolution** is given by:


$$z[i_1, i_2] = \sum_{j_1=-m}^m \sum_{j_2=-m}^m w[j_1, j_2] x[i_1 + j_1, i_2 + j_2]$$

- **3D and higher-order convolutions** are defined similarly.

$$z[i_1, i_2, i_3] = \sum_{j_1=-m}^m \sum_{j_2=-m}^m \sum_{j_3=-m}^m w[j_1, j_2, j_3] x[i_1 + j_1, i_2 + j_2, i_3 + j_3]$$




# Image Convolution Examples



Identity convolution:  
(zeroes with a '1' at  $w_{0,0}$ )

$$w * =$$




Compute  $z[i,j]$   
for this  
location

multiply element-wise  
and add up result to get


# Image Convolution Examples

$x$



Identity convolution:  
(zeroes with a '1' at  $w_{0,0}$ )

$w$



$*$

$=$

$z$




Compute  $z[i,j]$  for this location

multiply element-wise and add up result to get

# Image Convolution Examples



Translation Convolution:




Boundary: "zero"



# Image Convolution Examples



Translation Convolution:




Boundary: "replicate"



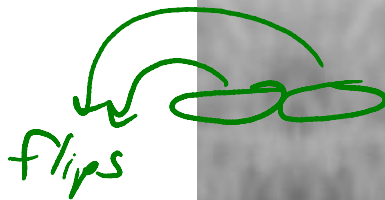
# Image Convolution Examples




Translation Convolution:



Boundary: "mirror"


 flips

Handwritten green arrows pointing from the word "flips" to the mirrored top portion of the image on the right.

# Image Convolution Examples



Translation Convolution:



Boundary: "ignore"








# Image Convolution Examples



Gaussian Convolution:

\*



=

blurs image to represent  
average  
(smoothing)




# Image Convolution Examples



Gaussian Convolution:

\*



=

(smaller variance)

blurs image to represent  
average  
(smoothing)






# Image Convolution Examples



Laplacian of Gaussian

\*



=

(larger variance)

Similar preprocessing may be done in basal ganglia and LGN.



Black/white as sides of edge

# Image Convolution Examples



"Emboss" filter:

$$* \begin{bmatrix} -2 & -1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & 2 \end{bmatrix} =$$


Many Photoshop effects  
are just convolutions.



# Image Convolution Examples




Gabor filter  
(Gaussian multiplied by  
sine or cosine)



\*


=

||




Gaussian

\*



Parallel Sine functions



horizontal  
bright  
to dark


horizontal dark to bright

# Image Convolution Examples



Gabor filter  
(Gaussian multiplied by  
sine or cosine)

\*



=

Different orientations of  
the sine/cosine let us  
detect changes with different  
orientations.






# Image Convolution Examples




Gabor filter  
(Gaussian multiplied by  
sine or cosine)

\*



=

(smaller variance)




# Image Convolution Examples



Gabor filter  
(Gaussian multiplied by  
sine or cosine)

\*



=



(smaller variance)

Vertical orientation


- Can obtain other orientations by  
rotating.

- May be similar to effect of V1 "simple cells."

# Image Convolution Examples



Max absolute value  
between horizontal and  
vertical Gabor:




"Horizontal/vertical edge detector"


# 3D Convolution



Represent  
as RGB



Can apply 3D  
convolutions



# 3D Convolution



Gaussian filter



# 3D Convolution



Gaussian filter  
(higher variance on  
green channel)



# 3D Convolution



Sharpen the blue  
channel.



# 3D Convolution




Gabor filter on  
each channel.





# Filter Banks

- To characterize context, we used to use **filter bank** like “MR8”:
  - 1 Gaussian filter, 1 Laplacian of Gaussian filter.
  - 6 max(Gabor) filters: 3 scales of sine/cosine (maxed over orientations).



- **Convolutional neural networks** are now replacing filter banks.

(pause)

# 1D Convolution as Matrix Multiplication

- Each element of a convolution is an **inner product**:

$$z_i = \sum_{j=-m}^m w_j x_{i+j}$$

$$= w^T x_{(i-m:i+m)}$$

$$= \tilde{w}^T x \quad \text{where } \tilde{w} = [0 \ 0 \ 0 \ \underbrace{\quad w \quad}_{\text{positions } i-m \text{ through } i+m} \ 0 \ 0]$$

- So **convolution is a matrix multiplication** (I'm ignoring boundaries):

$$z = \tilde{W}x \quad \text{where } \tilde{W} = \begin{bmatrix} \underbrace{\quad w \quad}_{\text{positions } i-m \text{ through } i+m} & 0 & 0 & 0 \\ 0 & \underbrace{\quad w \quad}_{\text{positions } i-m \text{ through } i+m} & 0 & 0 \\ 0 & 0 & \underbrace{\quad w \quad}_{\text{positions } i-m \text{ through } i+m} & 0 \\ 0 & 0 & 0 & \underbrace{\quad w \quad}_{\text{positions } i-m \text{ through } i+m} \end{bmatrix}$$

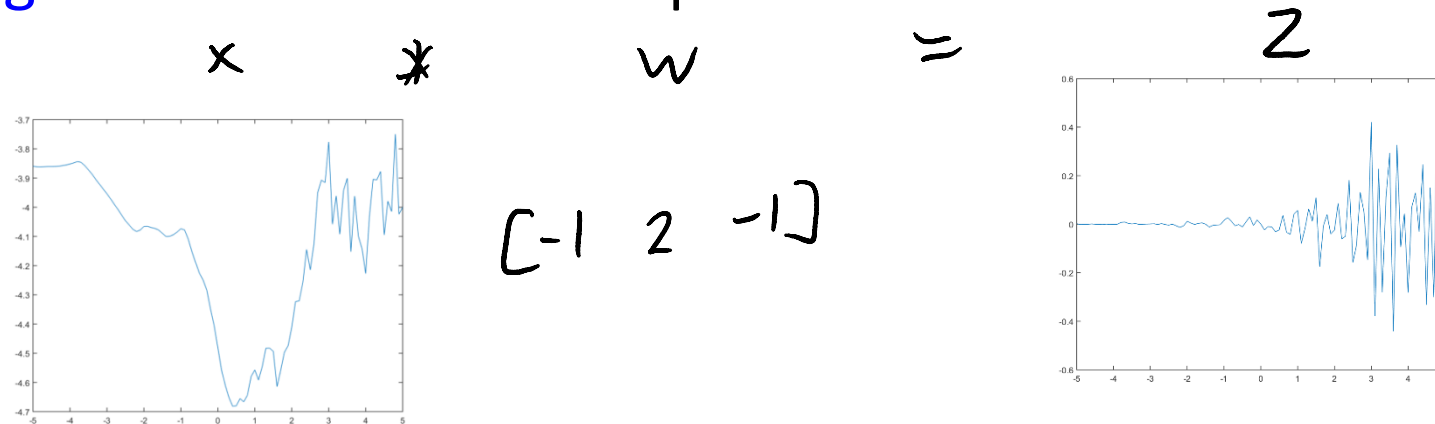
} matrix can be very sparse and only has  $2m+1$  variables.

- The shorter 'w' is, the more sparse the matrix is.

# 1D Convolution as Matrix Multiplication

- 1D convolution:

- Takes signal 'x' and filter 'w' to produce vector 'z':



- Can be written as a matrix multiplication:

$$W_x = \begin{bmatrix} 1 & -2 & 1 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 1 & -2 & 1 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -2 & 1 & 0 & \dots & 0 & 0 & 0 & 0 \\ \vdots & & & \vdots & & & & & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 1 & -2 & 1 \end{bmatrix} x = z$$



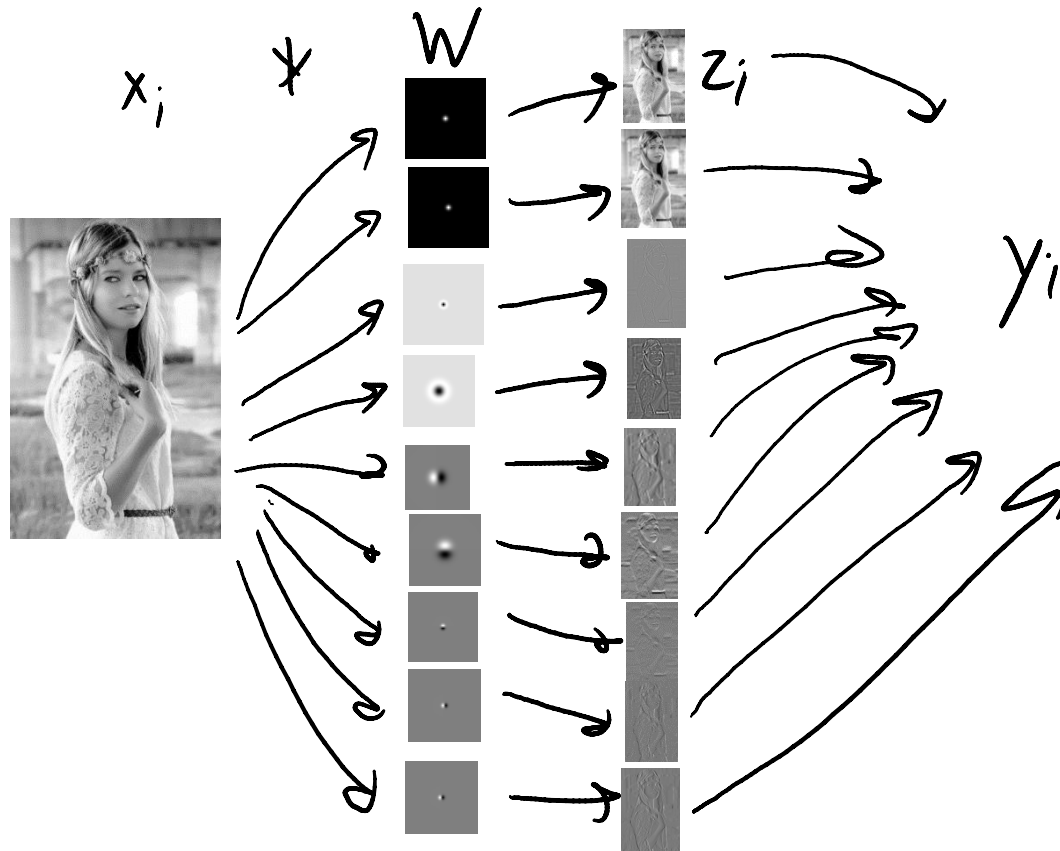
# Motivation for Convolutional Neural Networks

- Consider training neural networks on 256 by 256 images.
  - This is 256 by 256 by 3  $\approx$  200,000 inputs.
- If first layer has  $k=10,000$ , then it has **about 2 billion parameters**.
  - We want to avoid this huge number (due to storage and overfitting).
- Key idea: make  $Wx_i$  act like several convolutions (to make it sparse):
  1. Each row of  $W$  only applies to part of  $x_i$ .
  2. Use the same parameters between rows.
- Forces most weights to be zero, reduces number of parameters.

$$w_1 = [0 \ 0 \ 0 \ \text{---} \ w \ \text{---} \ 0 \ 0 \ 0]$$
$$w_2 = [0 \ \text{---} \ w \ \text{---} \ 0 \ 0 \ 0 \ 0]$$

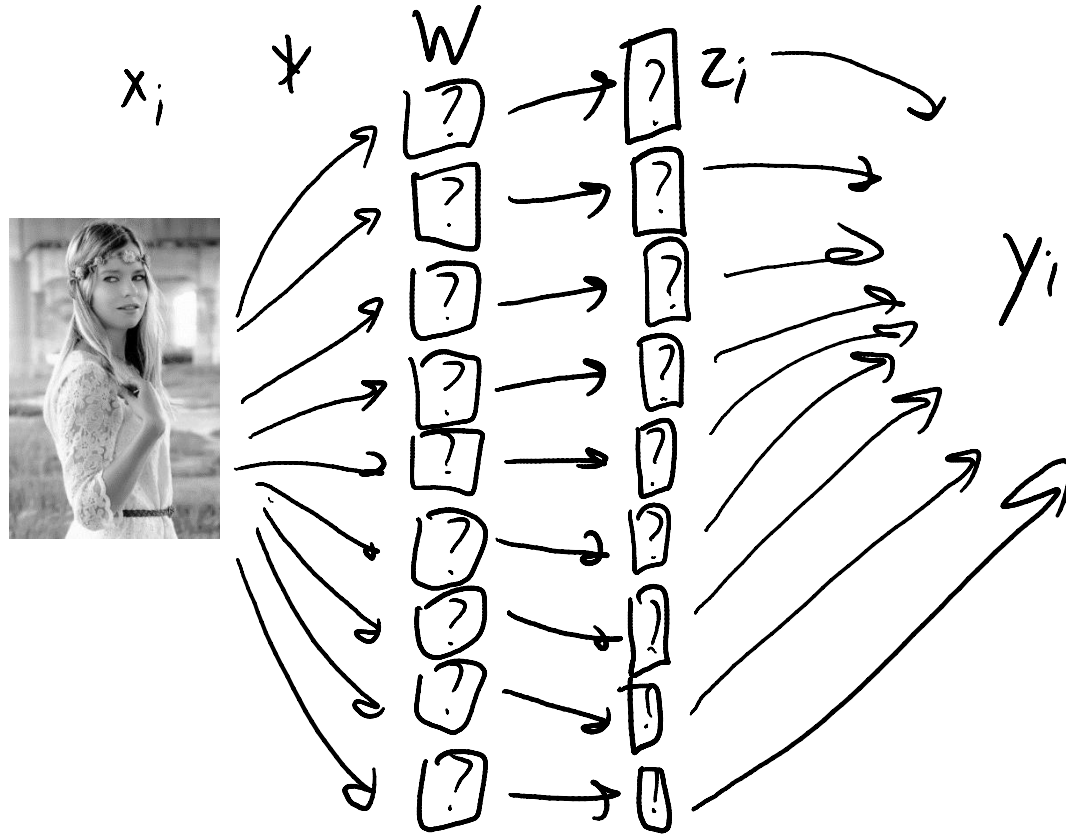
# Motivation for Convolutional Neural Networks

- Classic vision methods uses **fixed convolutions** as features:
  - Usually have **different types/variances/orientations**.
  - Can do subsampling or take **maxes across locations/orientations/scales**.



# Motivation for Convolutional Neural Networks


- Convolutional neural networks learn the features:
  - Learning 'W' and 'v' automatically chooses types/variances/orientations.
  - Don't pick from fixed convolutions, but learn the elements of the filters.





# Convolutional Neural Networks


- **Convolutional Neural Networks** classically have 3 layer “types”:
  - **Fully connected layer**: usual neural network layer with unrestricted  $W$ .



# Convolutional Neural Networks

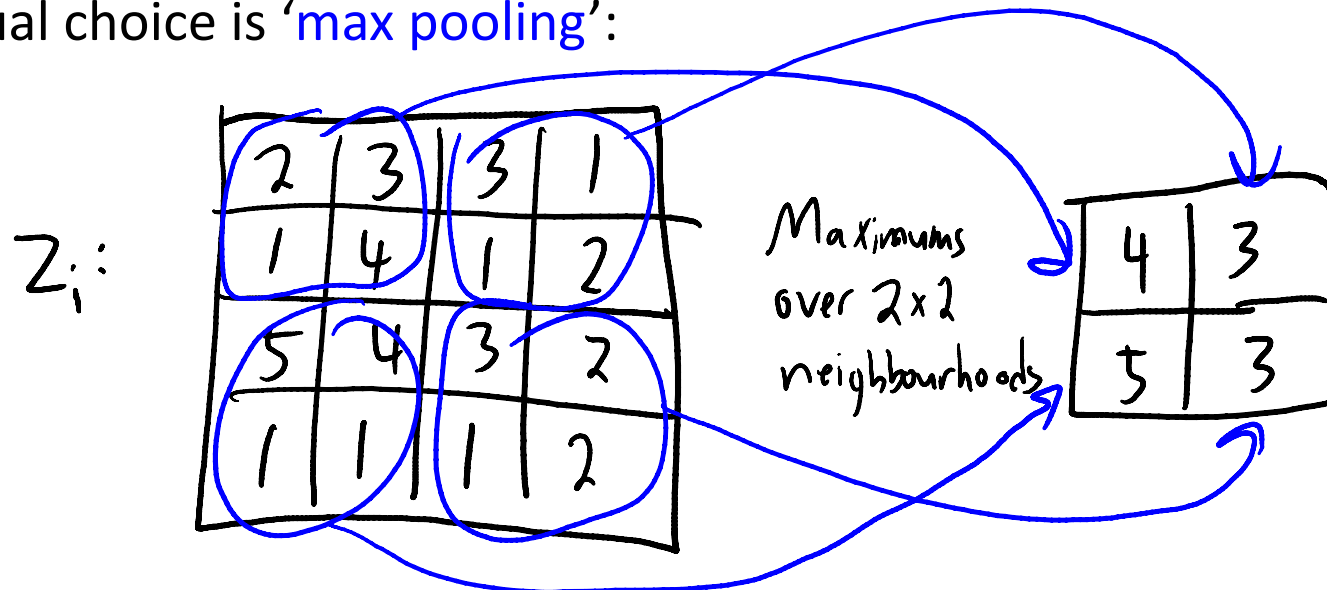
- **Convolutional Neural Networks** classically have 3 layer “types”:
  - **Fully connected layer**: usual neural network layer with unrestricted  $W$ .
  - **Convolutional layer**: restrict  $W$  to act like several convolutions.

1D example

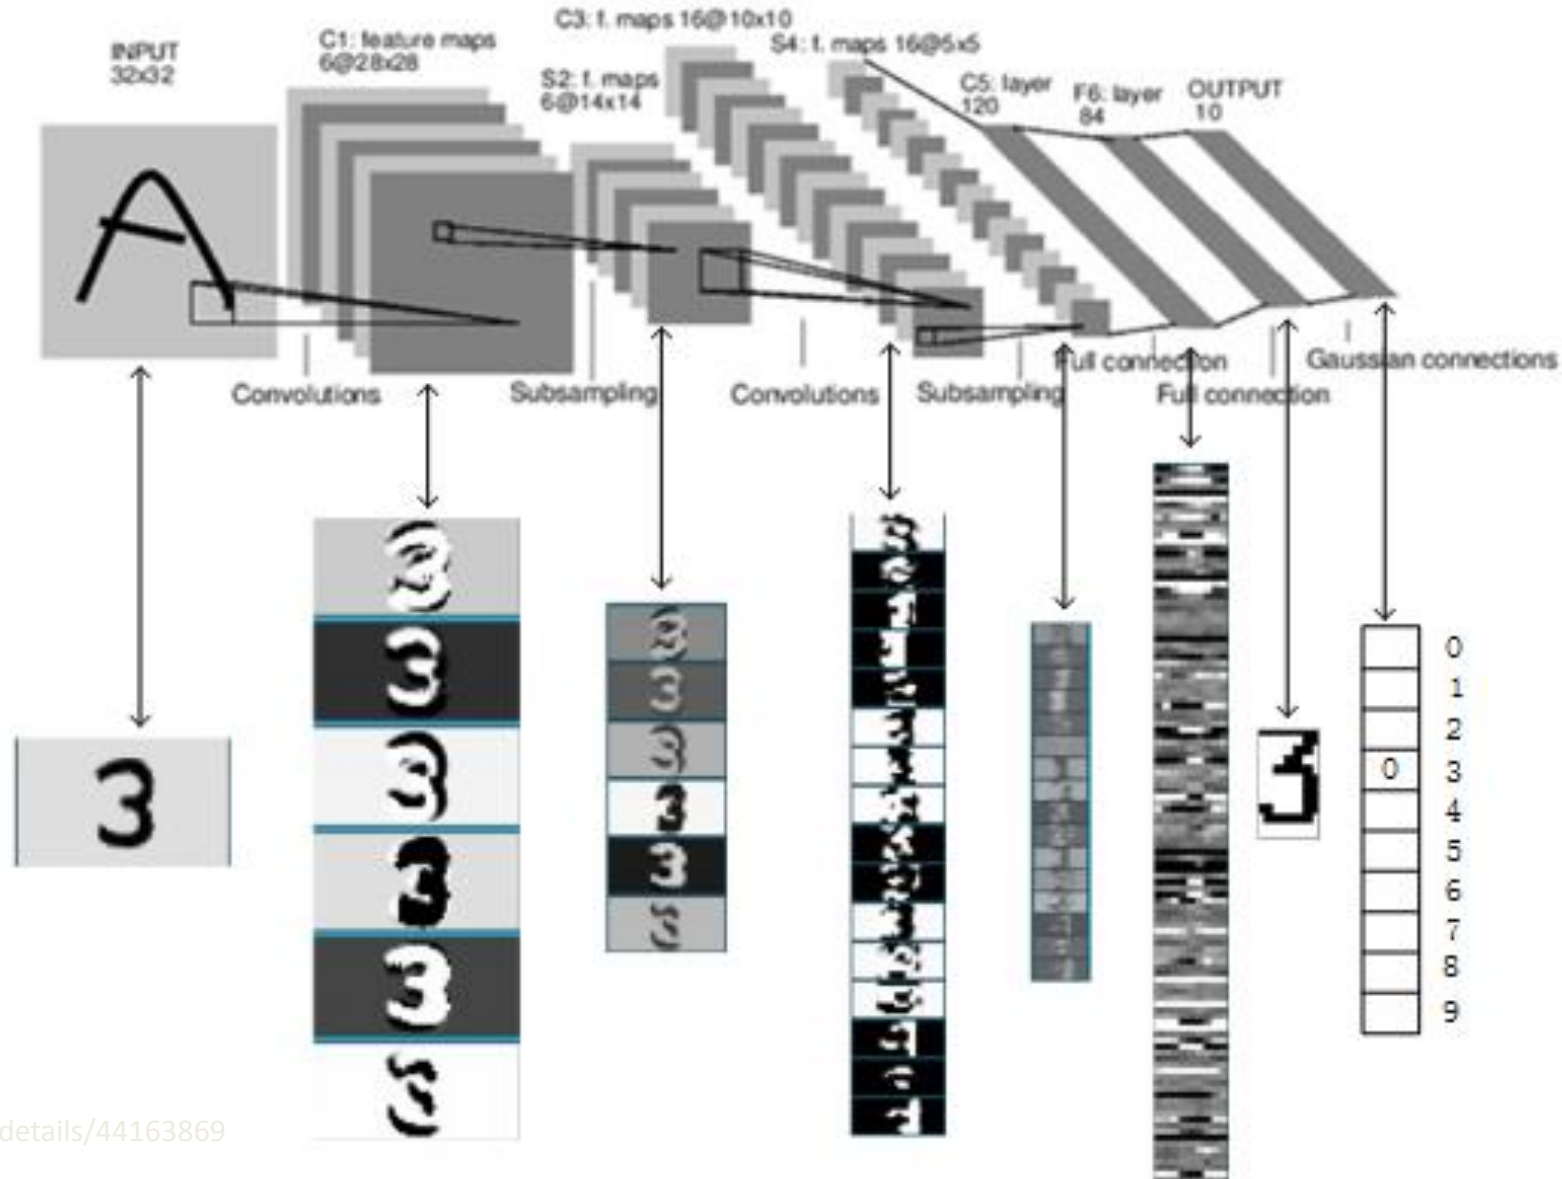


# Convolutional Neural Networks

- **Convolutional Neural Networks** classically have 3 layer “types”:
  - **Fully connected layer**: usual neural network layer with unrestricted  $W$ .
  - **Convolutional layer**: restrict  $W$  to act like several convolutions.
  - **Pooling layer**: combine results of convolutions.
    - Can add invariances or just make the number of parameters smaller.
    - Usual choice is ‘**max pooling**’:



# LeNet for Optical Character Recognition



# Summary


- **Convolutions** are flexible class of signal/image transformations.
  - Can approximate directional derivatives and integrals at different scales.
- **Max(convolutions)** can yield features that make classification easy.
- **Convolutional neural networks:**
  - Restrict  $W^{(m)}$  matrices to represent sets of convolutions.
  - Often combined with max (pooling).
- Next time: modern convolutional neural networks and applications.
  - Image segmentation, depth estimation, image colorization, artistic style.

# FFT implementation of convolution

- Convolutions can be implemented using fast Fourier transform:
  - Take FFT of image and filter, multiply elementwise, and take inverse FFT.
- It has faster asymptotic running time but there are some catches:
  - You need to be using periodic boundary conditions for the convolution.
  - Constants matter: it may not be faster in practice.
    - Especially compared to using GPUs to do the convolution in hardware.
  - The gains are largest for larger filters (compared to the image size).

# Image Coordinates


- Should we use the image coordinates?
  - E.g., the pixel is at location (124, 78) in the image.



- Considerations:
  - Is the interpretation different in different areas of the image?
  - Are you using a linear model?
    - Would “distance to center” be more logical?
  - Do you have enough data to learn about all areas of the image?

# Alignment-Based Features

- The position in the image is important in brain tumour application.
  - But we didn't have much data, so **coordinates didn't make sense**.
- We aligned the images with a “template image”.






# Alignment-Based Features

- The position in the image is important in brain tumour application.
  - But we didn't have much data, so **coordinates didn't make sense**.
- We aligned the images with a "template image".
  - Allowed "alignment-based" features:

Probability of gray matter at this pixel among tons of people aligned with template.

Probability of being brain pixel.




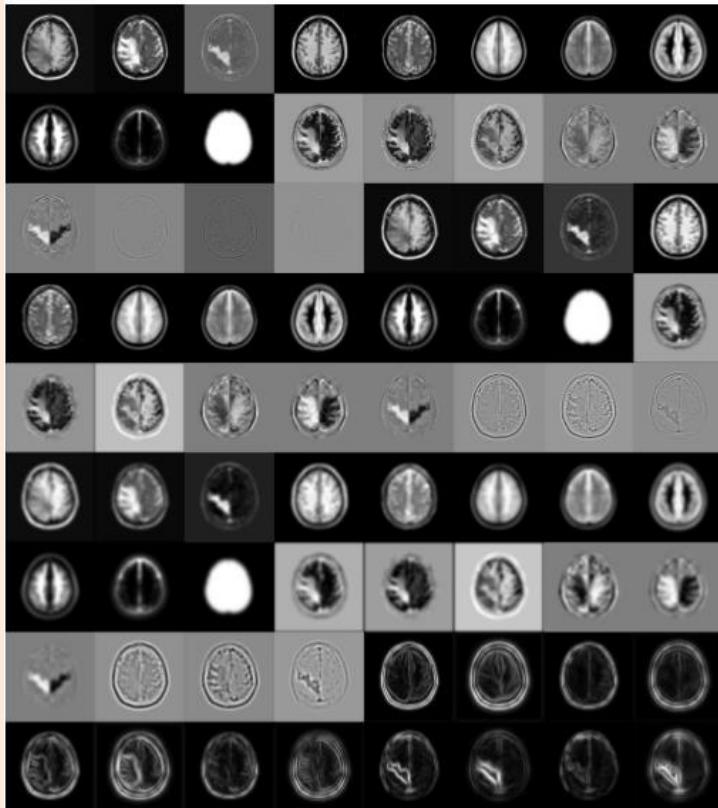
Original pixel values

Actual pixel value of template image at this location.

Left-right symmetry difference.

# Motivation: Automatic Brain Tumor Segmentation

- Final features for brain tumour segmentation:
  - MR8 filter bank applied to original T1, T2, and T1 “contrast” – T1 “original”.
  - Gaussian convolution with 3 variances of alignment-based features.



# SIFT Features

- Scale-invariant feature transform (SIFT):
  - Features used for object detection (“is particular object in the image”?)
  - Designed to detect unique visual features of objects at multiple scales.
  - Proven useful for a variety of object detection tasks.

