

CPSC 340: Machine Learning and Data Mining

Convolutional Neural Networks

Fall 2017

Admin



- **Assignment 5:**
 - Due tonight, 1 late day for Wednesday, 2 for Friday.
- **Final:**
 - Next Tuesday, details and previous exams posted on Piazza.
- **Extra office hours:**
 - 3:00-?:?? Thursday in ICICS 146 (with me).
 - Monday we'll have office hours at 11-12 (1 TA) and 1-2 (2 TAs).
 - Tuesday we'll have office hours from 12-2 (1 TA).

Last Lectures: Deep Learning

- We've been discussing **neural network / deep learning** models:

$$\hat{y}_i = v^T h(W^{(m)} h(W^{(m-1)} h(\dots h(W^{(2)} h(W^{(1)} x_i)) \dots)))$$

- We discussed **unprecedented vision/speech performance**.




Last Lectures: Deep Learning

- Last time we discussed **heuristics to make it work**:
 - Parameter initialization and data transformations.
 - Setting the **step size(s)** in stochastic gradient.
 - Alternative non-linear functions like **ReLU**.
 - Different forms of regularization:
 - L2-regularization, early stopping, dropout.
- These are often **still not enough** to get deep models working.
- Deep computer vision models are all **convolutional neural networks**:
 - The $W^{(m)}$ are **very sparse and have repeated parameters** (“tied weights”).
 - Drastically reduces number of parameters (speeds training, reduces overfitting).

Motivation: Automatic Brain Tumor Segmentation


- Task: segmentation tumors and normal tissue in multi-modal MRI data.



- Applications:
 - Radiation therapy target planning, quantifying treatment responses.
 - Mining growth patterns, image-guided surgery.
- Challenges:
 - Variety of tumor appearances, similarity to normal tissue.
 - “You are never going to solve this problem.”

Naïve Voxel-Level Classifier

- We could treat classifying a voxel as **supervised learning**:




$$x_i = (98, 187, 246)$$

$$y_i = \text{"tumour"}$$

- We can formulate predicting y_i given x_i as supervised learning.
- But it **doesn't work** at all with these features.

Need to Summarize Local Context


- The individual voxel values are almost meaningless:
 - This x_i could lead to different y_i .



- Intensities not standardized.
- Non-trivial overlap in signal for different tissue types.
- “Partial volume” effects at boundaries of tissue types.

Need to Summarize Local Context

- We need to represent the spatial “context” of the voxel.




$$x_i = (\underbrace{\quad\quad\quad}_{\text{neighbouring voxels}}, \underbrace{\quad\quad\quad}_{\text{summary statistics}}, \underbrace{\quad\quad\quad}_{\text{convolutions}})$$

- Include all the values of **neighbouring voxels**?
 - Variation on coupon collection problem: **requires lots of data** to find patterns.
- Measure neighbourhood **summary statistics** (mean, variance, histogram)?
 - Variation on bag of words problem: loses **spatial information** present in voxels.
- Standard approach uses **convolutions** to represent neighbourhood.

Representing Neighbourhoods with Convolutions

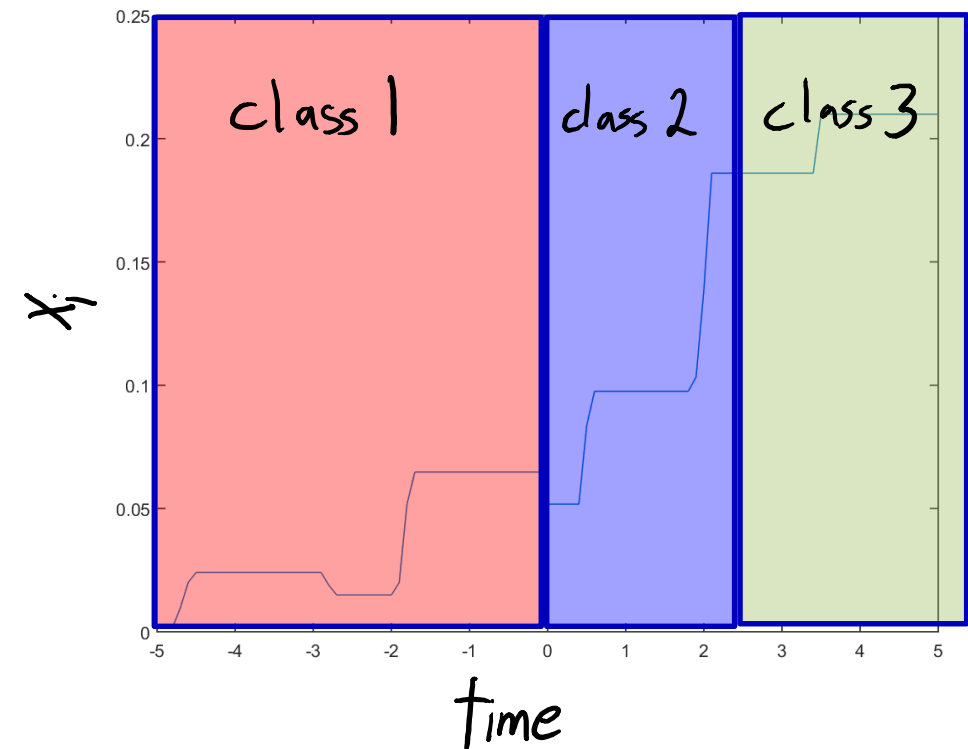

- Consider a 1D dataset:
 - Want to classify each time into y_i in $\{1,2,3\}$.
 - Example: speech data.



- Easy to distinguish class 2 from the other classes (x_i are smaller).
- Harder to distinguish between class 1 and class 3 (similar x_i range).
 - But convolutions can represent that class 3 is in “spiky” region.

Representing Neighbourhoods with Convolutions

- Original features (left) and features from **convolutions** (right):



- Easy to distinguish the 3 classes with these 2 features.

1D Convolution (notation is specific to this lecture)

- 1D convolution input:

- Signal 'x' which is a vector length 'n'.

- Indexed by $i=1,2,\dots,n$.

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

- Filter 'w' which is a vector of length '2m+1':

- Indexed by $i=-m,-m+1,\dots,-2,0,1,2,\dots,m-1,m$

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

$w_{-2} \quad w_{-1} \quad w_0 \quad w_1 \quad w_2$

- Output is a vector of length 'n' with elements:

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

- You can think of this as centering w at z_i and taking a dot product.

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 elements 1 through 5, and a red circle highlights the element 2 at index 4. A green bracket under indices 1-5 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.

- Convolution:

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

Indices 1 to 8 are shown below the array. A red circle highlights the element 0 at index 4. A green bracket under indices 1-5 is labeled 'n'.

Let's compute z_4 :

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

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 shown below the array. A blue box highlights the elements from index 3 to 7 (values 1, 2, 3, 5, 8). A red box highlights the element at index 5 (value 3). 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 shown below the array. A blue arrow points from the w_0 position (value 2) to the start of the blue box in the signal array.

- Convolution:

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

Indices 1 through 8 are shown below the array. A red box highlights the element at index 5 (value -1). 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 blue box in the signal array to this array. A blue arrow labeled "Multiply" points from this array to the next one.

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

A blue arrow labeled "Add" points from this array to the final calculation.

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

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 x is shown above z . Green brackets and arrows labeled "average" point from the first three elements of x (0, 1, 1) to the first three elements of z (0, 1, 1), and from the next three elements of x (2, 3, 5) to the next three elements of z (2, 3, 5). This indicates that the output z is identical to the input x .

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

Diagram illustrating the local average convolution. The input x is shown above z . Green brackets and arrows labeled "average" point from the first three elements of x (0, 1, 1) to the first three elements of z ($2\frac{1}{3}$, $1\frac{1}{3}$, 2), and from the next three elements of x (2, 3, 5) to the next three elements of z ($3\frac{1}{3}$, $5\frac{1}{3}$, $8\frac{2}{3}$). The first and last elements of z are unknown, indicated by question marks.

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.

$$W = [0.0001 \quad 0.0644 \quad 0.0540 \quad 0.2420 \quad 0.3989 \quad 0.2420 \quad 0.0540 \quad 0.0044 \quad 0.0001]$$




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



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.jl).



- 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 2nd-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)

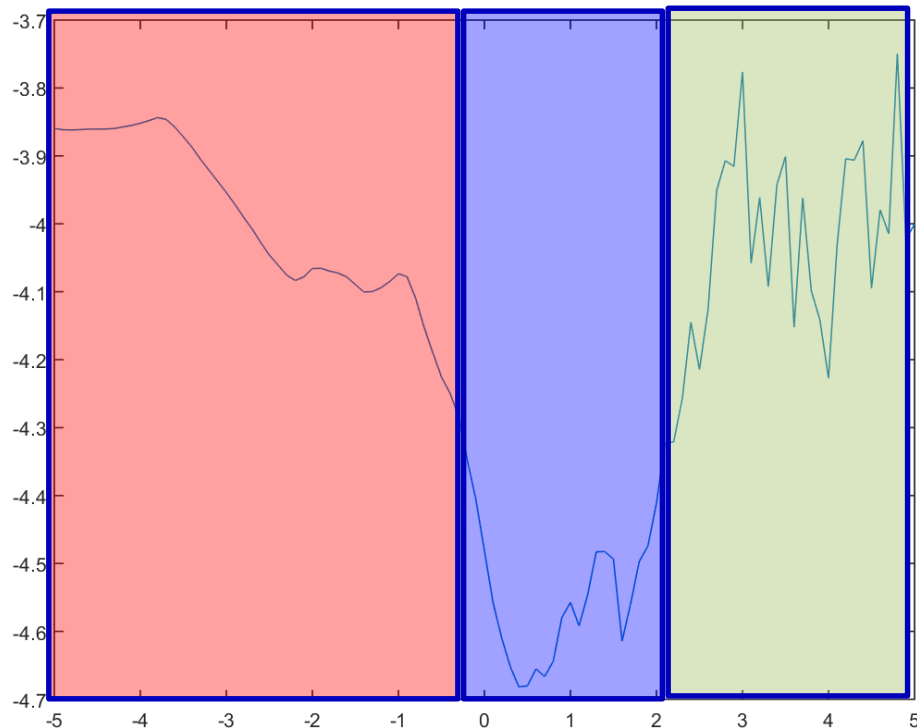
$$w = [-0.1416 \quad -0.1781 \quad -0.2746 \quad 0.1640 \quad 0.8607 \quad 0.1640 \quad -0.2746 \quad -0.1781 \quad -0.1416]$$

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




1D Convolution Examples

- We often use **maximum over several convolutions** as features:
 - We could take maximum of Laplacian of Gaussian over x_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

$z[i,j]$


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

Image Convolution Examples



Translation Convolution:



Boundary: "ignore"



Image Convolution Examples



Average convolution:

$$* \frac{1}{51} \left[\begin{array}{c} | | | | | \\ | | | | | \\ | | | | | \\ | | | | | \\ | | | | | \end{array} \right] =$$




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



"How much does it look like a black dot surrounded by white?"




"signed" image
(gray is 0)

Image Convolution Examples



Laplacian of Gaussian

*



=

(larger variance)

Similar preprocessing may be done in basal ganglia and LGN.



Image Convolution Examples



"Emboss" filter:

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

Many Photoshop effects
are just convolutions.

<http://setosa.io/ev/image-kernels>




Image Convolution Examples




Gabor filter
(Gaussian multiplied by
sine or cosine)

*




=

||



Gaussian

*



Parallel Sine functions





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.

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.

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 convolutions (to make it smaller):


1. Each row of W only applies to part of x_i .
2. Use the same parameters between rows.

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

- Forces most weights to be zero, reduces number of parameters.

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 results of 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 results of 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 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.

Number of parameters?


- Example with 1 conv/pool layer and 2 fully connected layers:
 - you start with a 28x28x3 RGB image
 - 32 filters each of size 5x5x3
 - 2x2 max pooling
 - fully connected layer with 128 hidden units
 - fully connected layer going to 10 output units for 10-class classification
- How many parameters does this model have?
 - the first convolutional layer has $5 \times 5 \times 3 \times 32$ (+32 bias).
 - this results in images of size 24x24 (this depends on how you handle convolutions at boundaries).
 - After 2x2 max pooling they are 12x12.
 - When we flatten this representation, we get 12x12x32 activations. This gives us $12 \times 12 \times 32 \times 128$ (+128 bias).
 - Finally we have a dense layer with 128×10 (+10 bias) parameters.
 - The grand total is $5 \times 5 \times 32 \times 3 + 12 \times 12 \times 32 \times 128 + 128 \times 10 + 32 + 128 + 10 = 2400 + 589824 + 1280 + 170 = 593674$.
- Most of the parameters come from the dense layer in this case (non-sparse).
- This kind of calculation is tedious but it's a good way to understand the details.

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”.




(Look different because we're showing middle slice and alignment is in 3D.)

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.

