

# Mathematical techniques in data science

## Lecture 12: Convolutional neural networks

# Reminders

- Office hours:
  - MW 3pm-4pm, Ewing Hall 312
  - By appointments
- Homework 1: due Monday EOD
- Sign up for group projects before class meeting next week

# Feed-forward neural networks

- Structure:
  - Graphical representation
  - Activation functions
  - Loss functions
- Training:
  - Stochastic gradient descent
  - Back-propagation



- High level API for deep learning
- More flexible to define network architecture than sklearn


## Define network architecture (1)

- Define a network as a Sequential object
- Add layers to it one-by-one

```
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Dense

model = Sequential()
model.add(Dense(50, activation='relu'))
model.add(Dense(50, activation='relu'))
model.add(Dense(10, activation='softmax'))
```

## Define network architecture (2)



# One-hot encoding

id	color
1	red
2	blue
3	green
4	blue



id	color_red	color_blue	color_green
1	1	0	0
2	0	1	0
3	0	0	1
4	0	1	0

Labels in Keras are usually encoded as one-hot vectors

## Demo: train an MLP using Keras




## Some intros to computer vision

# Computer vision


A field that enables computers and systems to derive meaningful information from digital images, videos and other visual inputs

- Image classification/object recognition
- Object detection
- Image segmentation
- Image generation
- Image style transfer


# Image classification




# Object detection



# Image segmentation



# Image generation



# Image style transfer




+



=



# Grayscale image representation




12

0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	.6	.8	0	0	0	0	0	0	0
0	0	0	0	0	0	0	.7	.8	0	0	0	0	0	0	0
0	0	0	0	0	0	0	.7	.8	0	0	0	0	0	0	0
0	0	0	0	0	0	0	.5	.8	.4	0	0	0	0	0	0
0	0	0	0	0	0	0	0	.8	.4	0	0	0	0	0	0
0	0	0	0	0	0	0	0	.8	.4	0	0	0	0	0	0
0	0	0	0	0	0	0	0	.8	.7	0	0	0	0	0	0
0	0	0	0	0	0	0	0	.8	.8	0	0	0	0	0	0
0	0	0	0	0	0	0	0	.9	.8	.1	0	0	0	0	0
0	0	0	0	0	0	0	0	.3	.8	.1	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0




# Color image representation



- Use RGB color mode
- Represent a color by 3 values: R (Red) G (Green) B (Blue)
- There are other color modes

# Image representation




- An image is an  $H \times W \times C$  matrix:  $H$  (height),  $W$  (width),  $C$  (depth or number of channels)
- Grayscale image:  $C = 1$
- RGB image:  $C = 3$

# Convolutional neural networks

## Issues with MLP

- Flat vectors lose spatial information
- Sensitive to location of the object
- Cannot capture small regions within an image
- Redundant parameters


# Convolutional neural networks



- two main parts: feature extraction and learning
- formed using 3 types of layers:
  - Convolution
  - Pooling
  - Dense layers


## Convolutional layer

# Convolutional layer



- Do not flatten the input image
- Apply a filter (kernel) to each local region of the image
- Slide the filter through all spatial locations to get the output

# Applying a kernel




$$\begin{aligned} \text{Output}[0, 0] &= (9*0) + (4*2) + (1*1) + \\ &\quad (1*4) + (1*1) + (1*0) + \\ &\quad (1*1) + (2*0) + (1*1) \\ &= 16 \end{aligned}$$


(Adapted from IBM)




## Examples of kernels

*Edge detection*


$$* \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix} =$$

Kernel 



The diagram shows the convolution of an input image (a squirrel's head) with a 3x3 kernel. The kernel is a 3x3 matrix with values [-1, -1, -1; -1, 8, -1; -1, -1, -1]. The result is an edge detection image where only the boundaries of the squirrel's head are visible in black on a dark background. An arrow labeled "Kernel" points to the kernel matrix.

*Sharpen*


$$* \begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix} =$$



The diagram shows the convolution of the same input image (a squirrel's head) with a different 3x3 kernel. The kernel is a 3x3 matrix with values [0, -1, 0; -1, 5, -1; 0, -1, 0]. The result is a sharpened version of the original image, where the edges are more pronounced and the overall image appears crisper.

# Feature extraction


- Historically, convolutional filters have been used to extract image features
- CNNs automate that process by considering the entries of the filters as model parameters

# Applying a kernel

7 x 7 Input Volume



5 x 5 Output Volume



## Applying a kernel


`https://poloclub.github.io/cnn-explainer/assets/figures/convlayer\_detailedview\_demo.gif`

# Stride

- Number of pixels to shift the filter
- Can be different for each dimension

strides = (2, 2)

7 x 7 Input Volume



3 x 3 Output Volume




(Source: Adit Deshpande)


# Padding

2 possible settings: Valid or Same

7 x 7 Input Volume




5 x 5 Output Volume



(Valid)

# Padding

Same padding: add 0 pixels to boundary of input image to get similar output shape



(Same)

## Applying a kernel

`https://vucdinh.github.io/Presentation/figures/cnn/  
filter-run.gif`