

AGENDA



GENERAL IDEA

The origin of CNNs



CONVOLUTION

How does a CNN work?



CNN ARCHITECTURE


Example VGGnet

GENERAL IDEA

THE ORIGIN OF CNNs


THE ORIGIN OF CNNs

- USED FOR IMAGE RECOGNITION SINCE 1980s
- **INSPIRED BY THE BRAIN'S VISUAL CORTEX**
 - NEURONS IN VISUAL CORTEX HAVE A **SMALL LOCAL RECEPTIVE FIELD**
 - RECEPTIVE FIELDS OF DIFFERENT NEURONS **OVERLAP**
 - TOGETHER THEY TILE THE **WHOLE VISUAL FIELD**
- SOME NEURONS ONLY REACT TO **SPECIFIC SHAPES**
- SOME NEURONS REACT TO MORE **COMPLEX SHAPES FROM LOWER LEVELS**
- POWERFUL ARCHITECTURE OF LOWER AND HIGHER-LEVEL NEURONS TO DETECT COMPLEX PATTERNS




FIELDS OF APPLICATIONS


- IMAGE DETECTION



- VOICE RECOGNITION



- NATURAL LANGUAGE PROCESSING (NLP)



CONVOLUTION

HOW A CNN WORKS


HOW DOES A CNN WORK?



WHAT DOES A COMPUTER SEE?



Input image
3600 x 2400



Resized image
30 x 30


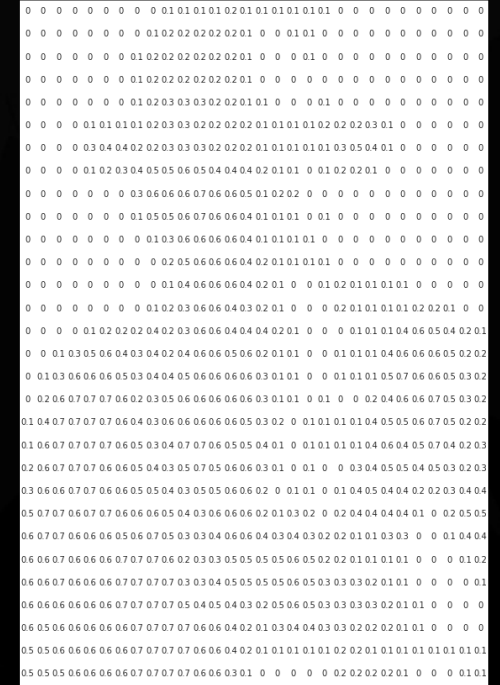


Image as matrix
of numbers [0, 1]



Matrix of numbers
[0, 1]
(900 values)

Source:
<https://www.youtube.com/watch?v=iaSUYvmCekI>

TASK IN COMPUTER VISION



Input image

```
0 0 0 0 0 0 0 0 0 0 0 1 0 1 0 1 0 2 0 1 0 1 0 1 0 1 0 1 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 1 0 2 0 2 0 2 0 2 0 1 0 0 0 1 0 1 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 1 0 2 0 2 0 2 0 2 0 2 0 1 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 1 0 2 0 2 0 2 0 2 0 2 0 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 1 0 2 0 3 0 3 0 3 0 2 0 2 0 1 0 1 0 1 0 1 0 2 0 2 0 3 0 1 0 0 0 0 0 0
0 0 0 0 0 1 0 1 0 1 0 2 0 3 0 3 0 2 0 2 0 2 0 1 0 1 0 1 0 1 0 2 0 2 0 3 0 1 0 0 0 0 0 0
0 0 0 0 0 3 0 4 0 4 0 2 0 2 0 3 0 3 0 3 0 2 0 2 0 1 0 1 0 1 0 1 0 1 0 3 0 5 0 4 0 1 0 0 0 0 0 0
0 0 0 0 0 1 0 2 0 3 0 4 0 5 0 5 0 6 0 5 0 4 0 4 0 4 0 2 0 1 0 1 0 0 1 0 2 0 2 0 1 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 3 0 6 0 6 0 6 0 7 0 6 0 6 0 5 0 1 0 2 0 2 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 1 0 5 0 5 0 6 0 7 0 6 0 6 0 4 0 1 0 1 0 1 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 1 0 3 0 6 0 6 0 6 0 6 0 4 0 1 0 1 0 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 2 0 5 0 6 0 6 0 6 0 4 0 2 0 1 0 1 0 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 1 0 4 0 6 0 6 0 6 0 4 0 2 0 1 0 0 0 1 0 2 0 1 0 1 0 1 0 1 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 1 0 2 0 3 0 6 0 6 0 4 0 3 0 2 0 1 0 0 0 0 2 0 1 0 1 0 1 0 1 0 2 0 2 0 1 0 0
0 0 0 0 0 1 0 2 0 2 0 2 0 4 0 2 0 3 0 6 0 6 0 4 0 4 0 4 0 2 0 1 0 0 0 0 1 0 1 0 1 0 4 0 6 0 5 0 4 0 2 0 1
0 0 0 1 0 3 0 5 0 6 0 4 0 3 0 4 0 2 0 4 0 6 0 6 0 5 0 6 0 2 0 1 0 1 0 0 0 1 0 1 0 1 0 4 0 6 0 6 0 5 0 2 0 2
0 0 1 0 3 0 6 0 6 0 5 0 3 0 4 0 4 0 5 0 6 0 6 0 6 0 3 0 1 0 1 0 0 0 1 0 1 0 1 0 5 0 7 0 6 0 6 0 5 0 3 0 2
0 0 2 0 6 0 7 0 7 0 7 0 6 0 2 0 3 0 5 0 6 0 6 0 6 0 6 0 3 0 1 0 1 0 0 1 0 0 0 2 0 4 0 6 0 6 0 7 0 5 0 3 0 2
0 1 0 4 0 7 0 7 0 7 0 6 0 4 0 3 0 6 0 6 0 6 0 6 0 6 0 5 0 3 0 2 0 0 1 0 1 0 1 0 1 0 4 0 5 0 5 0 6 0 7 0 5 0 2 0 2
0 1 0 6 0 7 0 7 0 7 0 6 0 5 0 3 0 4 0 7 0 7 0 6 0 5 0 5 0 4 0 1 0 0 1 0 1 0 1 0 1 0 4 0 6 0 4 0 5 0 7 0 4 0 2 0 3
0 2 0 6 0 7 0 7 0 6 0 6 0 5 0 4 0 3 0 5 0 7 0 5 0 6 0 6 0 3 0 1 0 0 1 0 0 0 3 0 4 0 5 0 5 0 4 0 5 0 3 0 2 0 3
0 3 0 6 0 6 0 7 0 7 0 6 0 6 0 5 0 5 0 4 0 3 0 5 0 5 0 6 0 6 0 2 0 0 1 0 1 0 0 1 0 4 0 5 0 4 0 4 0 2 0 2 0 3 0 4 0 4
0 5 0 7 0 7 0 6 0 7 0 7 0 6 0 6 0 6 0 5 0 4 0 3 0 6 0 6 0 6 0 2 0 1 0 3 0 2 0 0 2 0 4 0 4 0 4 0 4 0 1 0 0 2 0 5 0 5
0 6 0 7 0 7 0 6 0 6 0 6 0 5 0 6 0 7 0 5 0 3 0 3 0 4 0 6 0 6 0 4 0 3 0 4 0 3 0 2 0 2 0 1 0 1 0 3 0 3 0 0 0 1 0 4 0 4
0 6 0 6 0 7 0 6 0 6 0 6 0 7 0 7 0 7 0 6 0 2 0 3 0 3 0 5 0 5 0 5 0 6 0 5 0 2 0 2 0 1 0 1 0 1 0 1 0 0 0 0 0 1 0 2
0 6 0 6 0 7 0 6 0 6 0 6 0 7 0 7 0 7 0 3 0 3 0 4 0 5 0 5 0 5 0 6 0 5 0 3 0 3 0 3 0 2 0 1 0 1 0 0 0 0 0 0 0 1
0 6 0 6 0 6 0 6 0 6 0 7 0 7 0 7 0 5 0 4 0 5 0 4 0 3 0 2 0 5 0 6 0 5 0 3 0 3 0 3 0 2 0 1 0 1 0 0 0 0 0
0 6 0 5 0 6 0 6 0 6 0 6 0 7 0 7 0 7 0 6 0 6 0 4 0 2 0 1 0 3 0 4 0 4 0 3 0 3 0 2 0 2 0 2 0 1 0 1 0 0 0 0 0
0 5 0 5 0 6 0 6 0 6 0 6 0 7 0 7 0 7 0 6 0 6 0 4 0 2 0 1 0 1 0 1 0 1 0 2 0 2 0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1
0 5 0 5 0 5 0 6 0 6 0 6 0 7 0 7 0 7 0 6 0 6 0 3 0 1 0 0 0 0 0 0 2 0 2 0 2 0 2 0 1 0 0 0 0 0 1 0 1
```

Matrix of numbers

Classification

- Harrison Ford [0.8]
- Sean Connery [0.1]
- Roger Moore [0.05]
- Tom Cruise [0.05]

Prediction

WHY NOT SIMPLY USE A DEEP NETWORK WITH FULLY CONNECTED LAYERS?

Indie resized to 100 x 100



The picture has 10.000 pixels

With a **1.000 neuron** input layer ...


and a **fully connected** 1st hidden layer, ...

this first operation amounts to a total of **10 million** connections (weights, parameters).

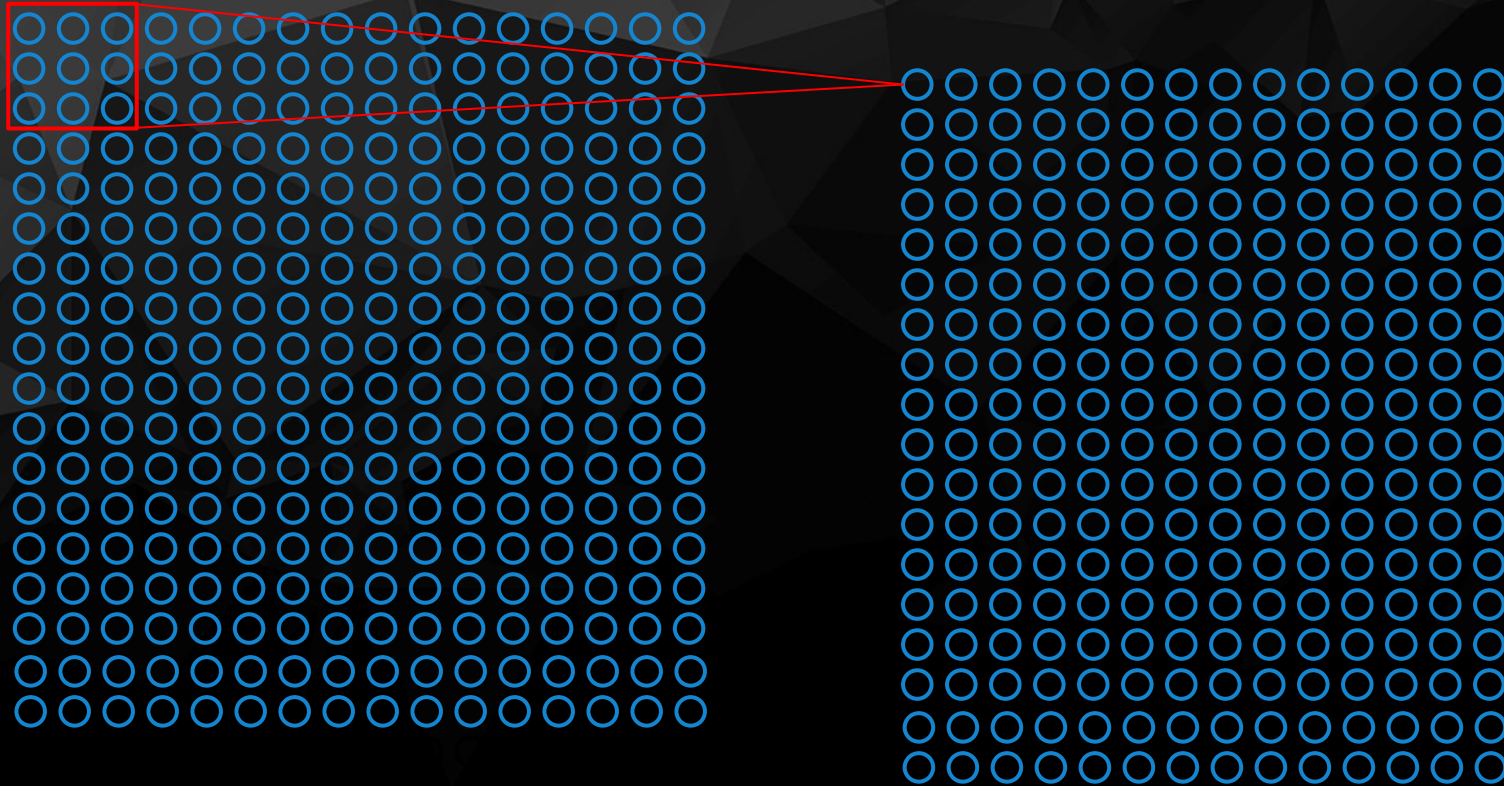
And that is just the **1st** layer!

For **large images**, a deep neural network breaks down.

AND we do not capture the spatial information of the pixels



THE CONVOLUTION LAYER – USING SPATIAL STRUCTURE



Input image
1 neuron for 1 pixel

Hidden layer
1 neuron for all pixels
in **receptive field** (filter)




Idea: connect smaller sections of the input image to respective neuron in the hidden layer.

Receptive field (FILTER) marks a specific area in the input image.


Use a **sliding window** to define the connections.

The GOAL: How to **weight** the **FILTER** to detect particular features in the image?

CONVOLUTION



Input image
1 neuron for 1 pixel



Hidden layer
1 neuron for all pixels
in **receptive field** (filter)

- Apply a **set of weights** – a **filter** – to extract local features
- Use **multiple filters** to extract different features
- **Spatially share parameters** of each filter



ELEMENT-WISE MULTIPLY AND ADD THE OUTPUTS

1	3
5	2

Part of input image



1	2
2	1

Filter

= 19

$(1 * 1) + (3 * 2) + (5 * 2) + (2 * 1) = 19$



APPLICATION OF A FILTER

Convolutional layer: Connection between neurons and only those pixels within their **receptive field**.

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

Image



1	0	1
0	1	0
1	0	1

Filter



4	3	4
2	4	3
2	3	4

Feature map

$$(1 * 1) + (1 * 0) + (1 * 1) + (0 * 0) + (1 * 1) + (1 * 0) + (0 * 1) + (0 * 0) + (1 * 1) = 4$$

$$(1 * 1) + (1 * 0) + (0 * 1) + (1 * 0) + (1 * 1) + (1 * 0) + (0 * 1) + (1 * 0) + (1 * 1) = 3$$

APPLYING DIFFERENT FILTERS


Horizontal edge detection

-1	-2	-1
0	0	0
1	2	1




Vertical edge detection

-1	0	1
-2	0	2
-1	0	1



Mixed edge detection

0	-2	0
-2	1	2
0	2	0



→ Filters

PADDING

Adding additional space to **preserve the same height and width of previous layer.**

Zero padding

0	0	0	0	0	0	0
0	1	1	1	0	0	0
0	0	1	1	1	0	0
0	0	0	1	1	1	0
0	0	0	1	1	0	0
0	0	1	1	0	0	0
0	0	0	0	0	0	0

Image



1	0	1
0	1	0
1	0	1

Filter



2	2	3	1	1
1	4	3	4	1
1	2	4	3	3
1	2	3	4	1
0	2	2	1	1

Feature map

Also, different kinds of paddings possible (i.e., One padding).
Highlight pixels at the edges of image.

USING A LARGER STRIDE

The shift from one receptive field to the next one is called **stride**.

Stride = 2

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

Image



1	0	1
0	1	0
1	0	1

Filter



4	4
2	4

Feature map


Connect larger input to smaller layer.

Reduction of the model's computational complexity.

STACKING MULTIPLE FEATURE MAPS

Applying different filters ...

- Convolutional layers with multiple feature maps
- Representation in 3D

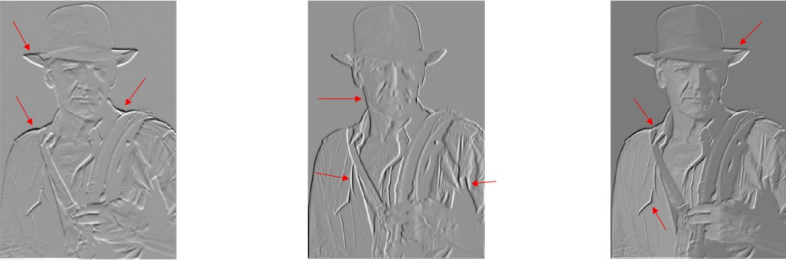
CONVOLUTION **Applying different filters** 

Horizontal edge detection Vertical edge detection Mixed edge detection


-1	-2	-1
0	0	0
1	2	1

-1	0	1
-2	0	2
-1	0	1

0	-2	0
-2	1	2
0	2	0



17



... results in different feature maps.

-1	-2	-1
0	0	0
1	2	1


-1	0	1
-2	0	2
-1	0	1

0	-2	0
-2	1	2
0	2	0

...

→ **Filters**

HOW DOES A CNN LEARN THE FILTERS?



-1	-2	-1
0	0	0
1	2	1

-1	0	1
-2	0	2
-1	0	1

0	-2	0
-2	1	2
0	2	0


...

How are all these filters determined or learned?

- NUMBER AND SIZE OF FILTERS IN EACH CONVOLUTION LAYER ARE SET **BY DESIGN**
- FILTERS INITIALIZED AT RANDOM AND THEN LEARNED (FWD PASS, BACKWARD PROP)
- CONVOLUTIONAL LAYER LEARNS MOST USEFUL FILTERS **AUTOMATICALLY** DURING TRAINING FOR **ITS TASK**
- LAYERS AFTER THIS WILL LEARN TO **COMBINE** THEM TO MORE **COMPLEX PATTERNS**

CAN YOU SPOT THE DIFFERENCE?

Image subsampled by
POOLING



Full sized image



Pooling

POOLING LAYERS

Subsample (i.e., **shrink**) the input image to reduce the computational load (memory usage, number of parameters)

Stride = 2
Filter = 3x3

1	1	2	0	4
0	1	7	1	0
0	8	1	1	1
9	0	1	3	5
0	4	1	0	0

Image

Max pooling

8	7
9	5

Feature map

Set filter size, stride and padding as for convolution layer.

No weights attached.

The layer aggregates input with an aggregation function (i.e., max or mean).



ACTIVATION BY NON-LINEARITY

- APPLY AFTER EVERY CONVOLUTION OPERATION
- RECTIFIED LINEAR UNIT (RELU)
- $f(x) = \text{MAX}(0, x)$
- PIXEL-BY-PIXEL OPERATION THAT REPLACES ALL NEGATIVE VALUES BY **ZERO**




Vertical edges
activated

CNN ARCHITECTURE

EXAMPLE VGGNET


FROM BASIC TO DETAILED FEATURES



Low level features
(1st Conv Layers)

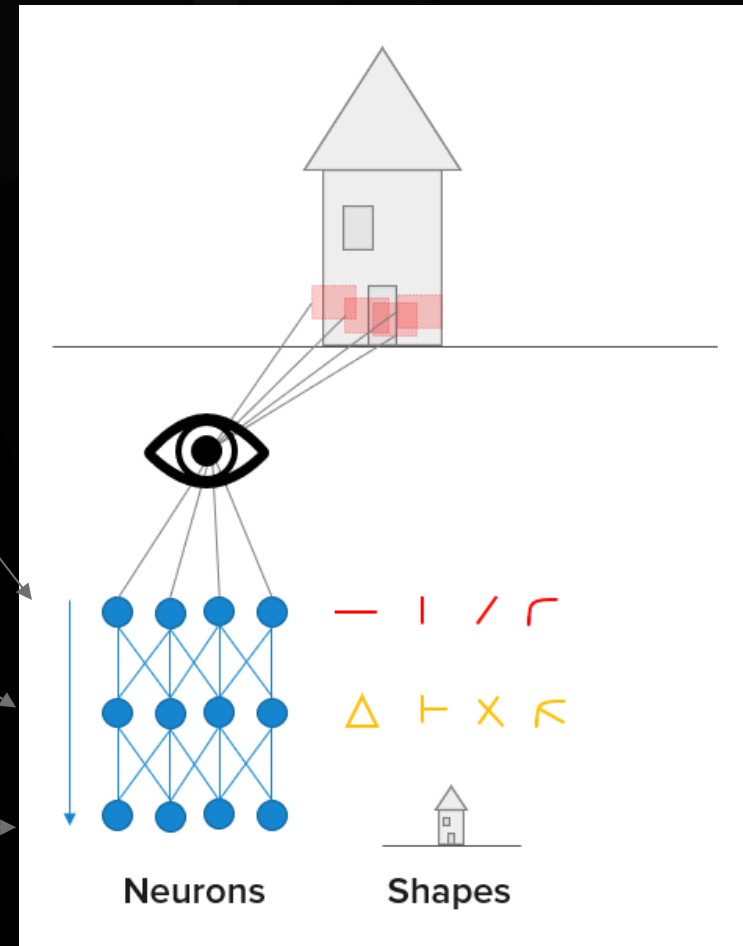


Mid level features
(...)



High level features
(Last Conv Layers)

Remember the visual cortex?




VGGNET

- INVENTED BY SIMONYAN AND ZISSERMAN FROM VISUAL GEOMETRY GROUP (VGG) AT UNIVERSITY OF OXFORD IN 2014 ^[1]
- LARGE SCALE VISUAL RECOGNITION
- FIXED FILTER SIZE OF 3×3 AND THE STRIDE OF 1
- DIFFERENT VERSIONS (VGG16, VGG19, ETC.)
- **WHY?** REDUCE THE # OF PARAMETERS IN THE CONV LAYERS AND IMPROVE ON TRAINING TIME.

[1] K. S. a. A. Zisserman, "Very deep convolutional networks for large-scale image recognition", in International Conference on Learning Representations (ICLR), San Diego, 2015.


VGGNET 16

Each CNN layer learns filters of increasing complexity.




- **Basic feature** detection (edges, corners, etc.)
- **Parts of objects** (for faces i.e., eyes, noses, etc.)
- **Higher representations** (recognize full objects, in different shapes and positions)

A CLOSER LOOK ON FINAL PREDICTION



1	2	1
1	3	5
1	0	0

Flattening



Softmax
1 x 1 x 1.000

Pooling
Layer
7 x 7 x 512

Flattening
1 x 1 x
25.088

Hidden
Layer
1 x 1 x 4.096

Hidden
Layer
1 x 1 x 4.096

Hidden
Layer
1 x 1 x 1.000