Deep Learning for Medical Images: An appetizer

Stefano Gualandi
Università di Pavia, Dipartimento di Matematica

email: stefano.gualandi@unipv.it
twitter: @famo2spaghi
blog: http://stegua.github.com
web: http://matematica.unipv.it/gualandi/

Research Group Marketing

UNIVERSITÀ DI PAVIA

Sezione:SI PARLA DI NOI

PROVINCIA PAVESE
Dir. Resp.:Andrea Filippi
Tiratura: 9.006 Diffusione: 11.420 Lettori: 78.000

UNIVERSITÀ

Matematica in ospedale ricercatori pavesi premiati negli Stati Uniti

PAVIA
Un modello per ottimizzare le sale operatorie. Il team The HospITAls del Dipartimento di Matematica Felice Casorati dell'Università di Pavia si è classificato al secondo posto nell'edizione annuale dell'Aimms-Mopta Optimization Modeling Competition. Si tratta di una gara internazionale per gruppi di studenti che prevede la soluzione di problemi reali attraverso modelli e metodi di ottimizzazione, una branca della matematica applicata. Il gruppo composto dagli studenti Ambrogio Maria Bernardelli, Lorenzo Bonasera ed Eleonora Vercesi, del dottorato in Computational Mathematics and Decision Sciences, sotto la supervisione del ricercatore Davide Duma, si è aggiudicato un premio di 600 dollari.

La finale si è svolta alla Lehigh University a Bethlehem, Pennsylvania, all'interno dell'International Conference on Continuous Optimization (Iccopt 2022). La challenge riguardava un problema di pianificazione delle sale operatorie, tema di ricerca tra i più indagati in letteratura scientifica. —

ALESSIO MOLTENI



Outline

1 Introduction

2 Multilayer Neural Networks

3 Convolutional Neural Networks

4 Transfer Learning

5 Generative Models

Outline

1. Introduction (*DISCLAIMER: math perspective!*)
2. Neural Networks (NNs) for Discriminative Models
   1. Multilayer NNs
   2. **Convolutional NNs**
3. Neural Networks for Generative Models
   1. ThisPersonDoesNotExists.com
   2. Data augmentation for brain tumoral images
4. Reserach Perspective

Math and Images

Disclaimer: "Images are just numbers!"

Math and Images

Disclaimer: "Images are just numbers!"



64x64 pixels
MedNIST Challenge

Math and Images

Disclaimer: "Images are just numbers!"



64x64 pixels
MedNIST Challenge



16x16 pixels
Downsampling

Math and Images

Disclaimer: "Images are just numbers!"



64x64 pixels
MedNIST Challenge



16x16 pixels
Downsampling

17	14	13	12	15	14	14	87	82	16	14	14	13	14	16	21
15	12	11	10	104	77	10	123	116	23	101	13	10	12	14	19
13	10	10	11	115	108	13	146	146	63	117	29	8	10	11	17
12	9	90	64	106	140	11	159	145	114	148	21	6	7	10	16
11	7	99	123	85	152	98	184	184	165	161	17	8	8	13	16
11	6	44	135	110	166	152	174	177	168	171	5	8	93	114	15
11	6	5	139	162	153	185	195	189	196	163	5	102	171	139	16
10	5	6	46	161	168	208	218	214	227	177	6	165	173	54	15
9	4	25	22	170	188	212	210	210	230	170	158	207	156	6	11
8	3	241	147	119	217	224	225	225	244	182	215	214	86	6	11
5	1	91	249	228	211	247	253	253	254	234	216	158	5	5	11
5	1	16	186	81	195	255	255	255	255	252	152	5	4	4	11
6	2	1	1	3	234	255	255	255	253	138	4	0	1	3	10
6	2	2	2	5	255	255	255	255	218	27	5	1	2	3	10
6	3	2	2	95	255	251	255	255	115	2	2	1	1	3	10
8	4	4	3	171	255	255	255	255	5	3	3	3	3	4	10

Image as matrix of numbers

Math, Color Images, and Tensors



RGB Image

Math, Color Images, and Tensors



RGB Image



Red channel

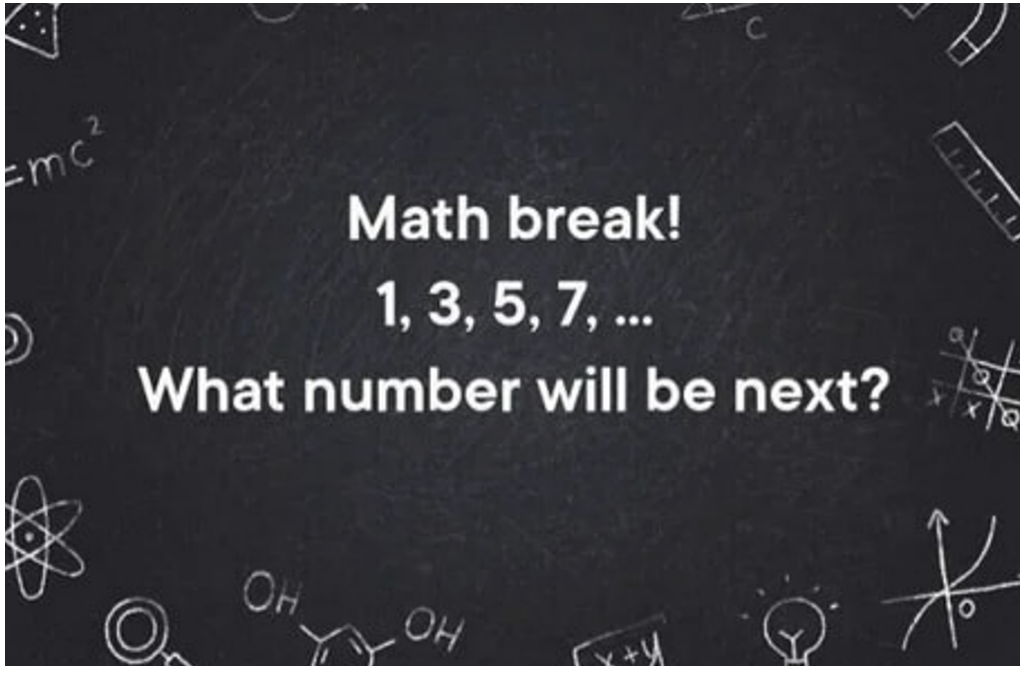


Green channel

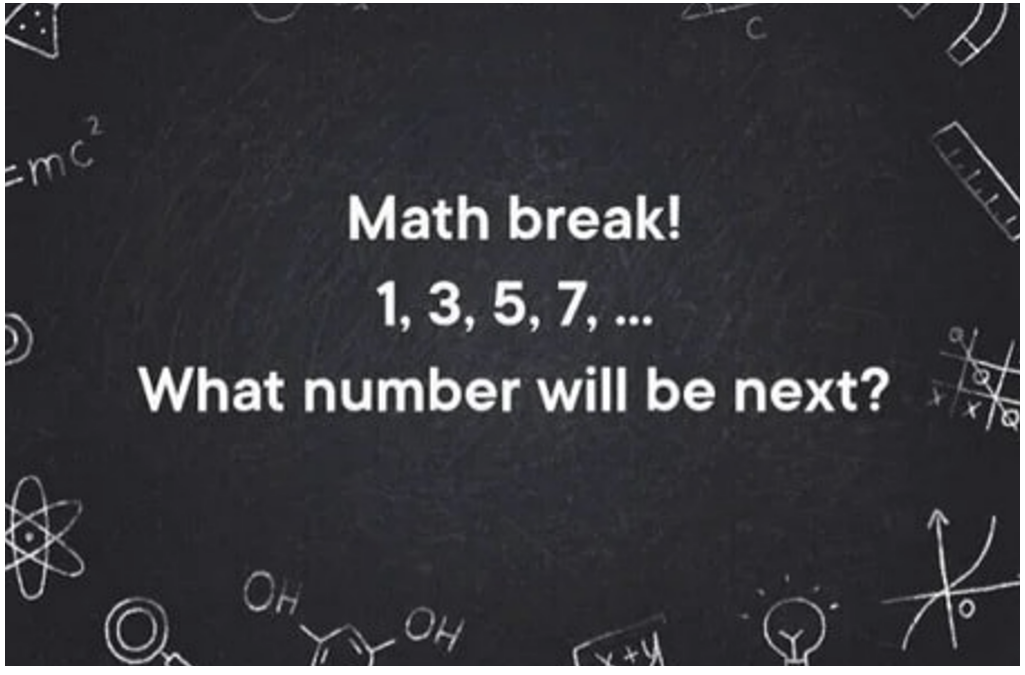


Blu channel

Math and Learning



Math and Learning

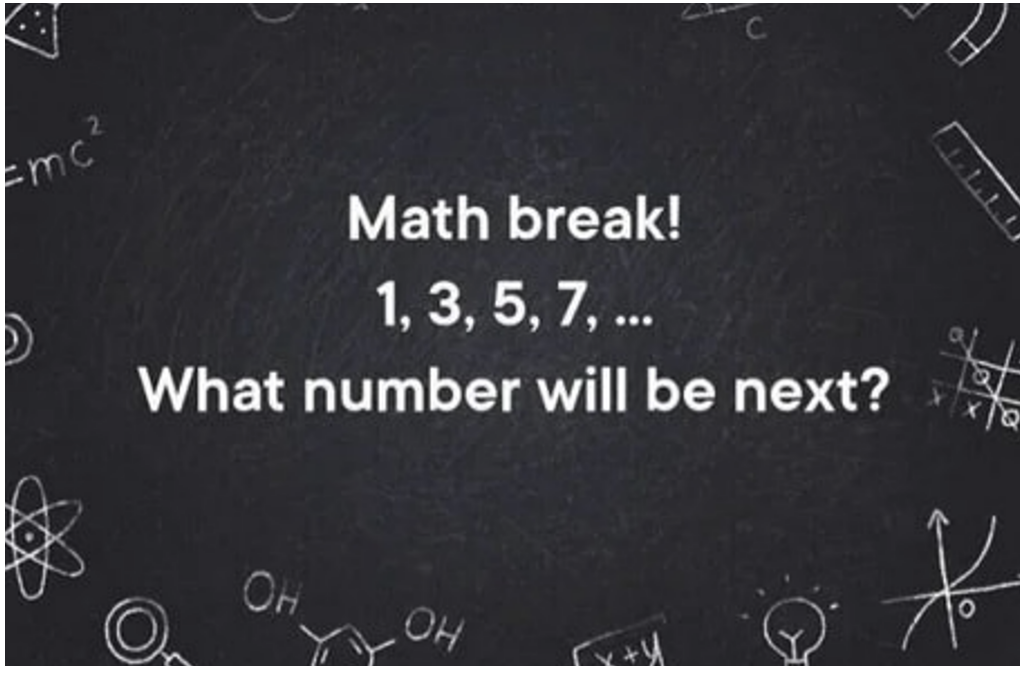


The next number is 217341 if

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

Math and Learning



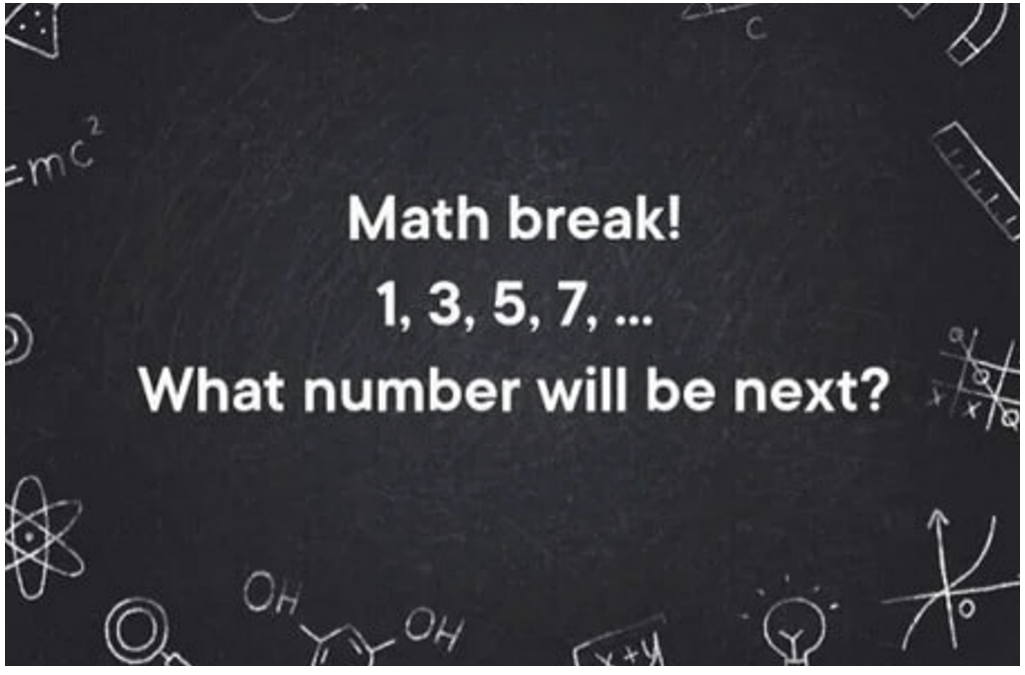
The next number is 217341 if

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$f(1) = 1$,

Math and Learning



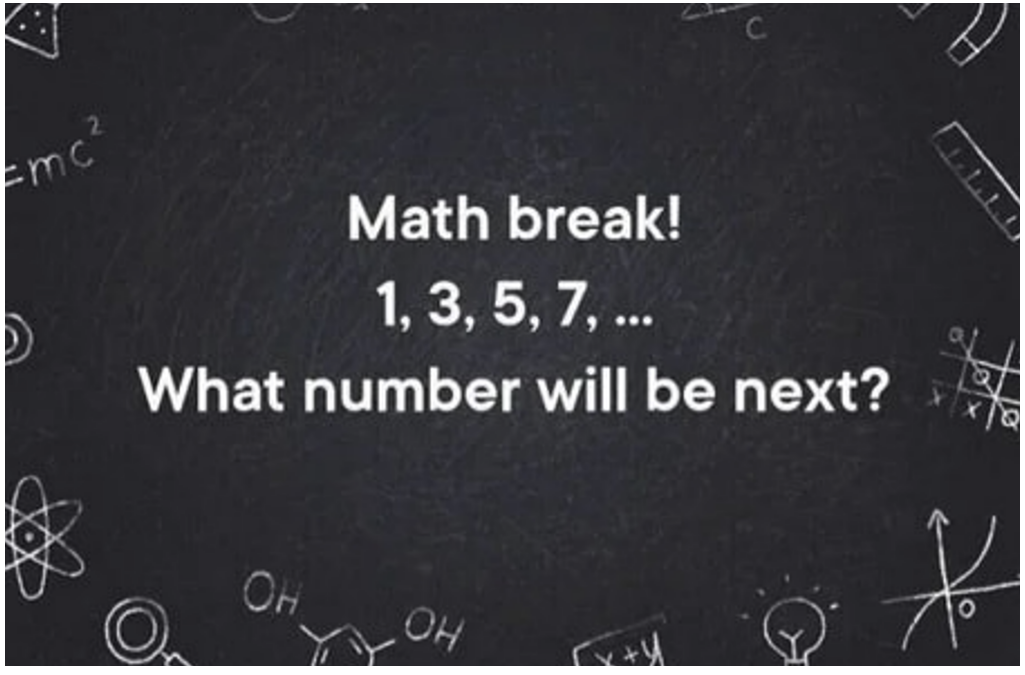
The next number is 217341 if

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3,$$

Math and Learning



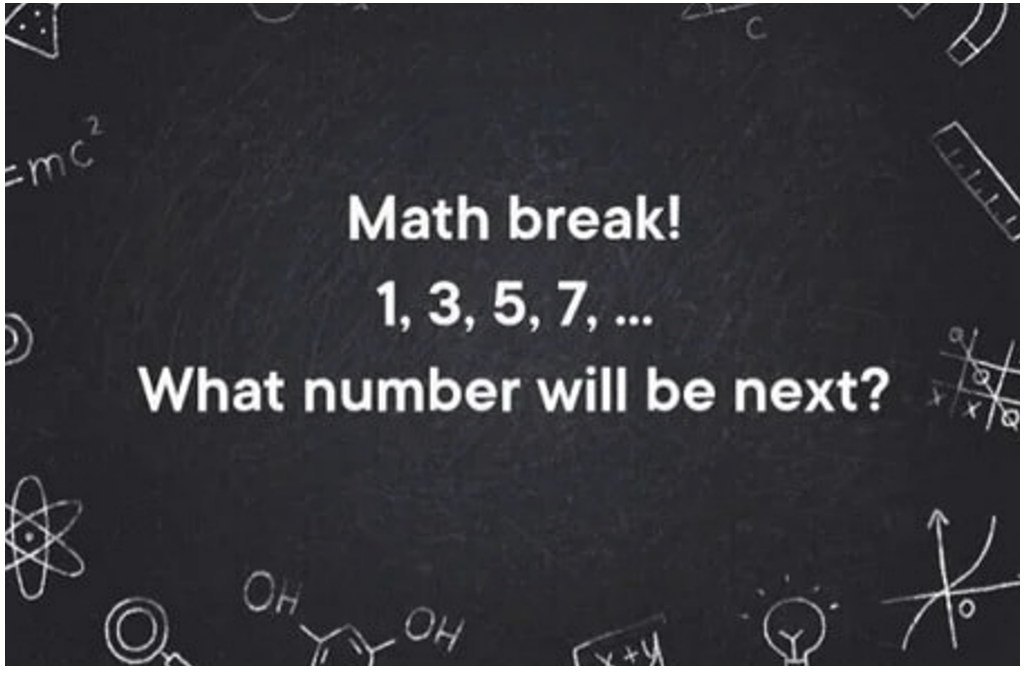
The next number is 217341 if

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3, \quad f(3) = 5,$$

Math and Learning



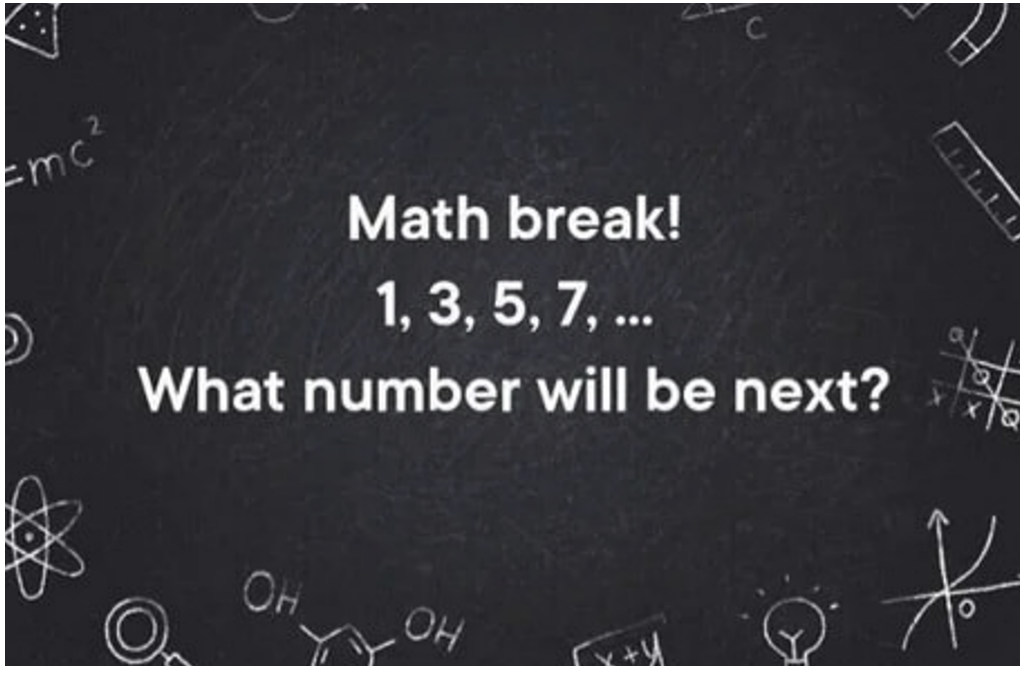
The next number is 217341 if

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3, \quad f(3) = 5, \quad f(4) = 7,$$

Math and Learning



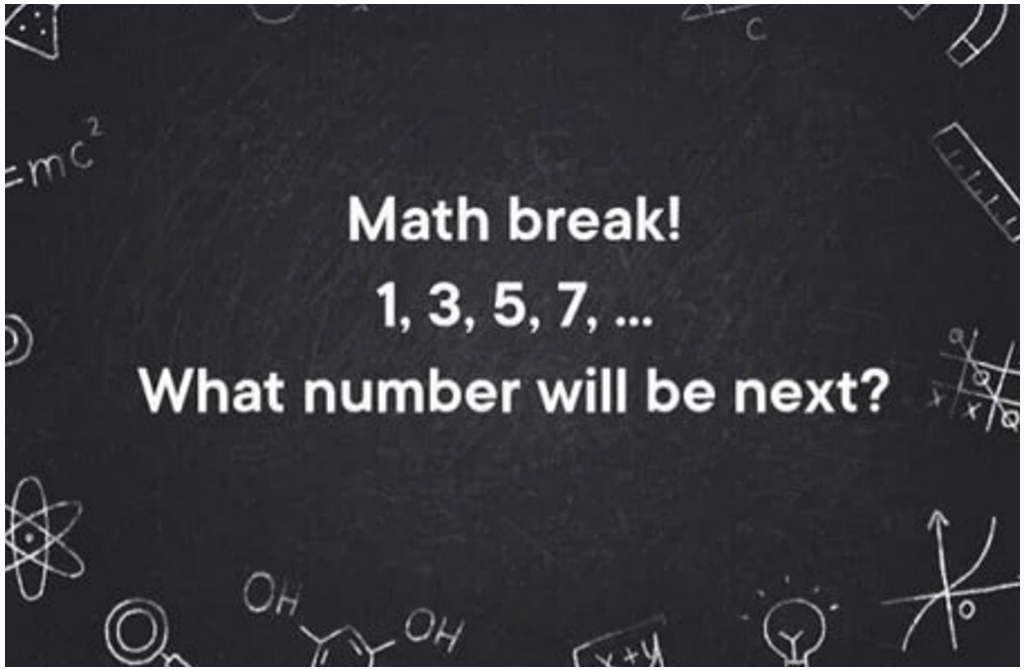
The next number is 217341 if

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3, \quad f(3) = 5, \quad f(4) = 7, \quad f(5) = 217341$$

Math and Learning



The next number is 217341 if

$$f(x) = \frac{18111}{2}x^4 - 90555x^3 + \frac{633885}{2}x^2 - 452773x + 217331$$

since then

$$f(1) = 1, \quad f(2) = 3, \quad f(3) = 5, \quad f(4) = 7, \quad f(5) = 217341$$

REMARK: Learning is finding a (composite) function that "*explains*" the data.

Learning Classification Tasks: RSNA Classroom

In 2017 and 2018, the RSNA Deep Learning Classroom, presented by NVIDIA Deep Learning Institute, used the **MNIST image dataset** to show radiologists how to train a classifier to recognize handwritten digits.

This year [2019], the Classroom will feature a medically originated set of images, **MedNIST Database** [..] (1) chest radiographs, (2) head CTs, (3) chest CTs, (4) abdomen CTs, (5) breast MRIs, and (6) brain MRIs [..] The goal is to build and train an AI model that accurately classifies images into those six categories.

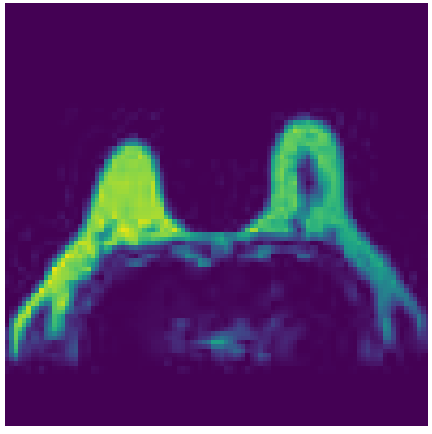
Charles E. Kahn. We All Need A Little Maigc (Editorial). Radiology: Artificial Intelligence. Vol.1, n.4, 2019.

Classification Tasks: RSNA Classroom

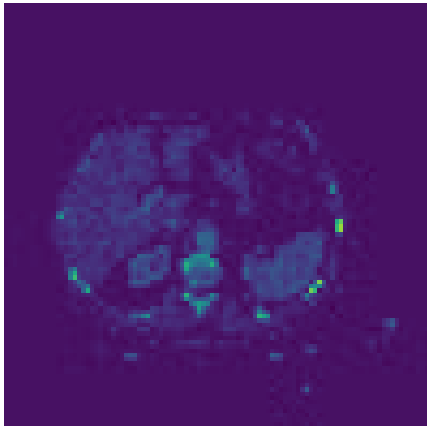
Hand



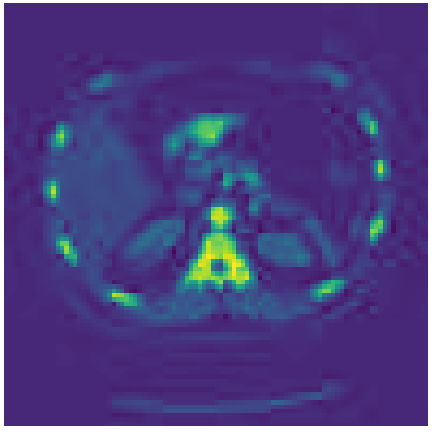
BreastRMI



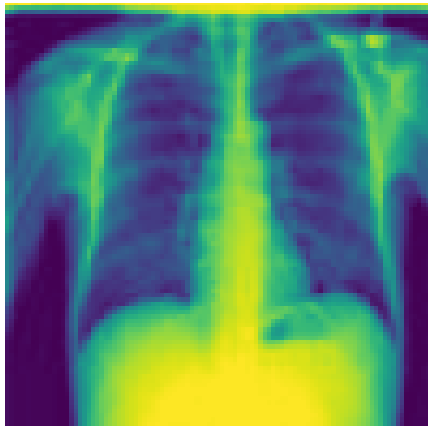
ChestCT



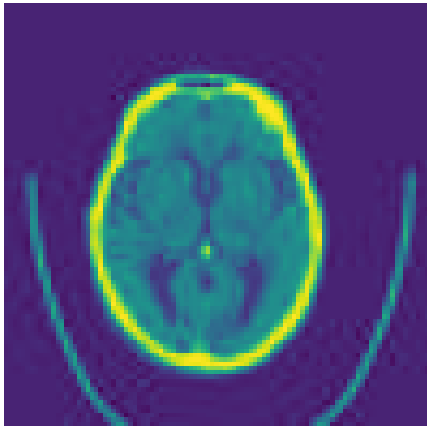
CTabdome



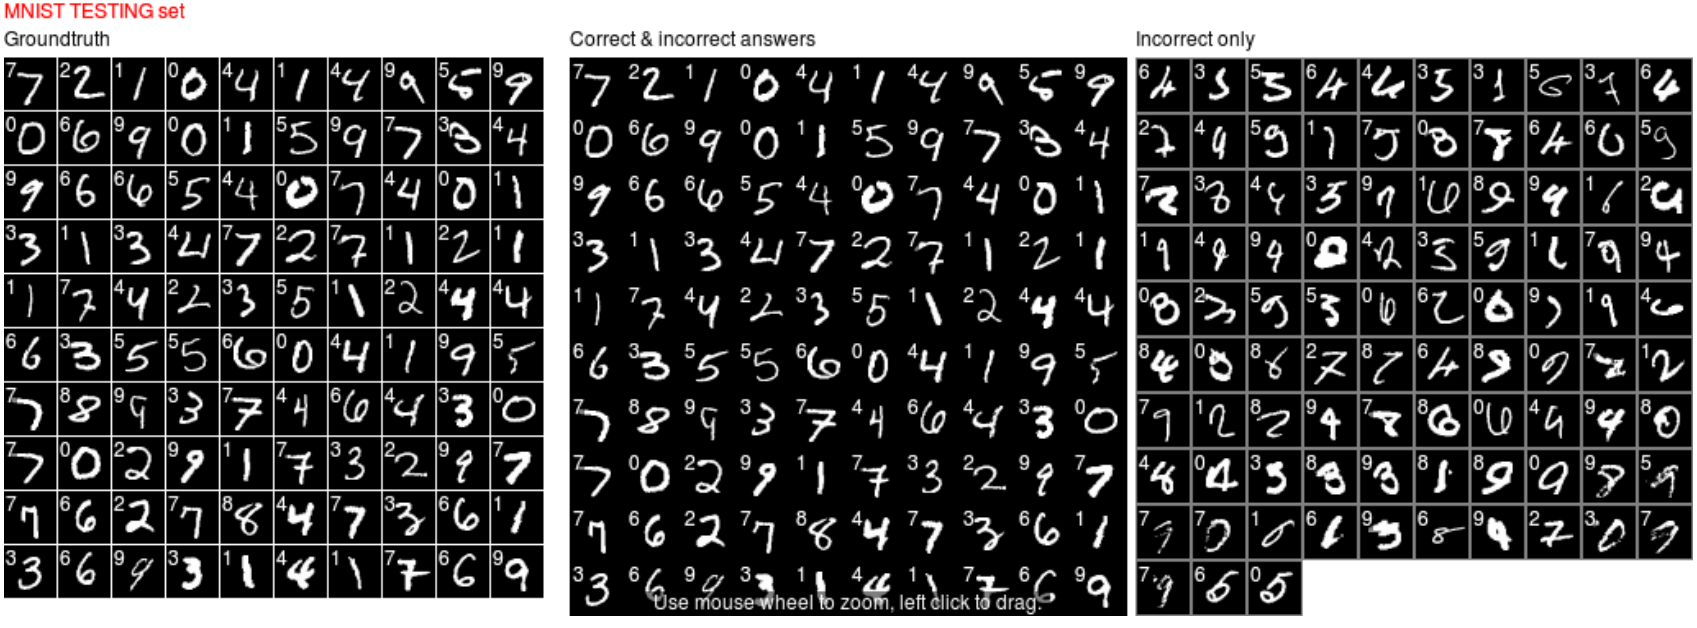
CXR



HeadCT



Most used dataset in the research: Digits



Outline

1 Introduction

2 Multilayer Neural Networks

3 Convolutional Neural Networks

4 Transfer Learning

5 Generative Models

Artificial Neural Networks



https://mashable.com/2016/12/01/facebook-ai-video-lessons

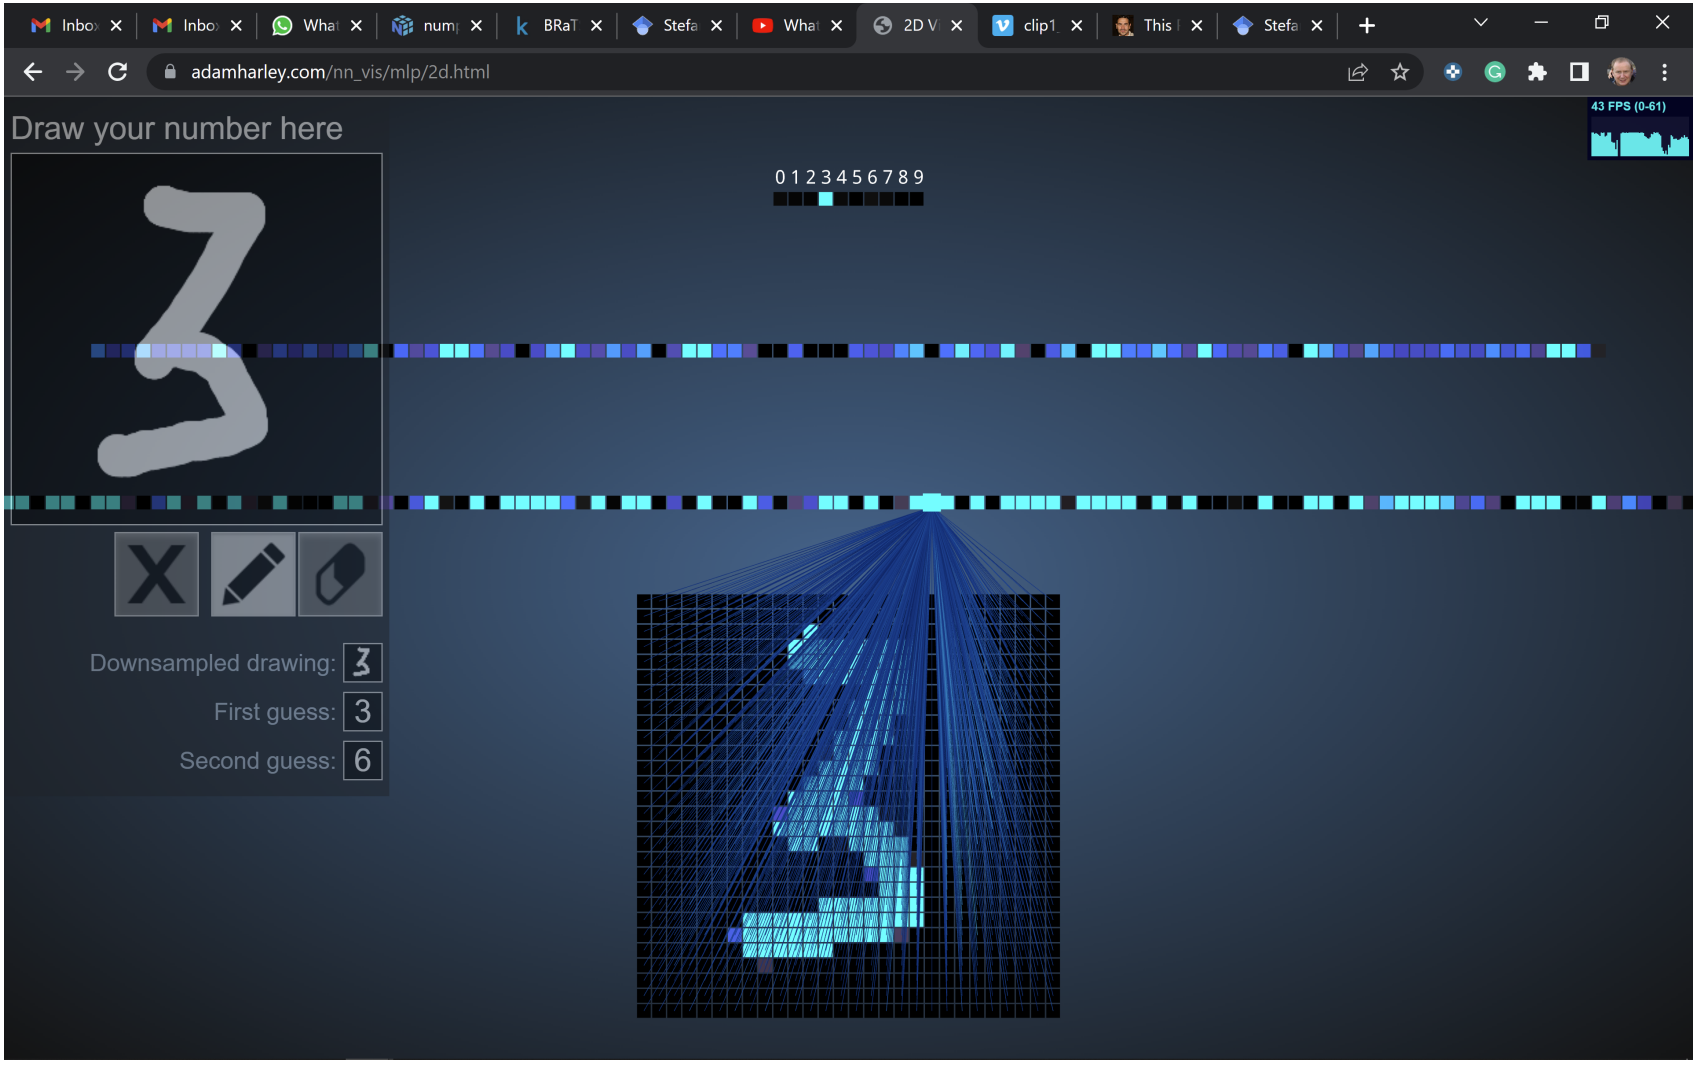
Biological Motivation

1. The human brain is estimated to contain a densely interconnected network of around 10^{11} neurons
2. Each neuron is connected to around others 10^4 neurons. Neurons activity is typically excited or inhibited through connections to other neurons.
3. The switching of neurons is of order of 10^{-3} seconds, much slower than modern computers, which have switching speeds of 10^{-10} seconds
4. Humans take around 10^{-1} seconds to recognize a friend (compare to switching speed...)
5. Human brain performs a highly parallel tasks

Biological Neurons

VIDEO RECORDED

Multilayer Feedforward Neural Networks



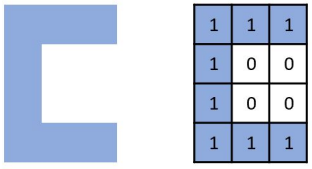
https://adamharley.com/nn_vis/mlp/2d.html

Backpropagation as Automatic Differentiation

DEMO ON LEARNING WEIGHTS: Backpropagation Algorithms

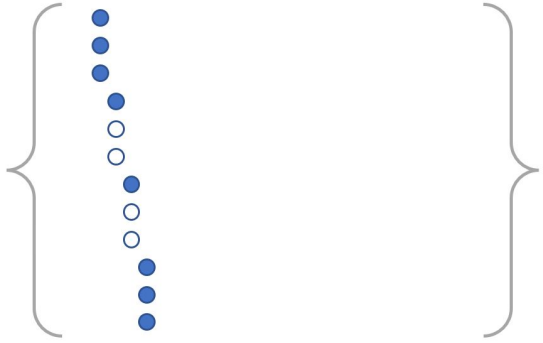
https://www.youtube.com/watch?v=Ilg3gGewQ5U

Basic Math Operations for an Artificial Neuron

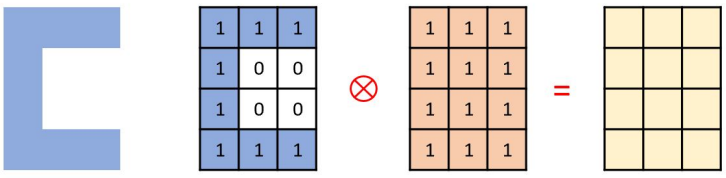


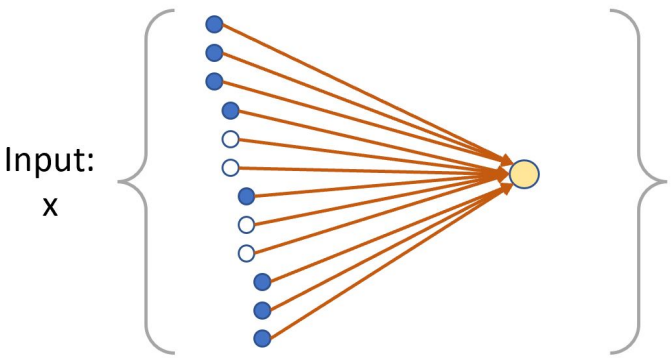
1	1	1
1	0	0
1	0	0
1	1	1

Input:
x

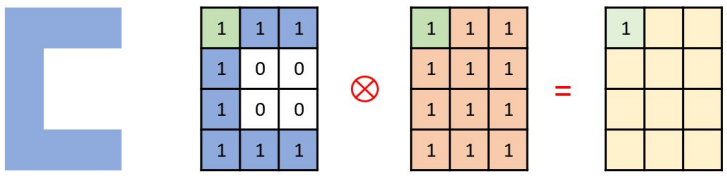


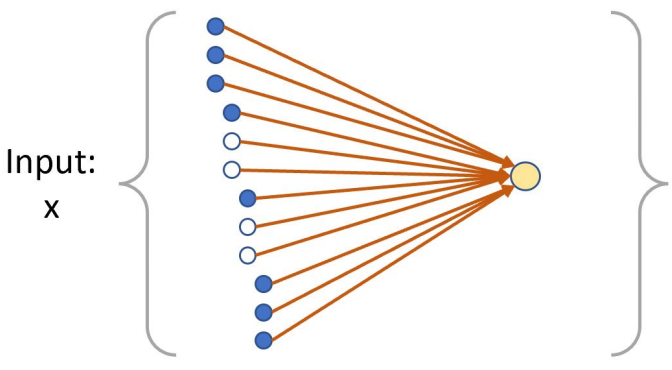
Basic Math Operations for Artificial NNs



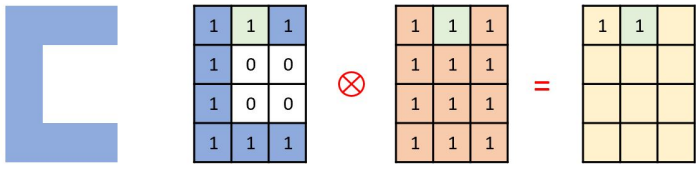


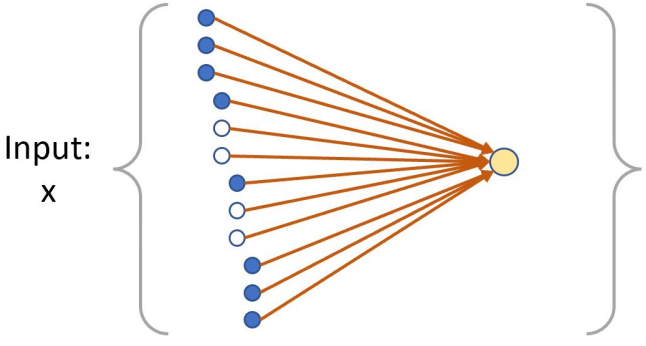
Basic Math Operations for Artificial NNs



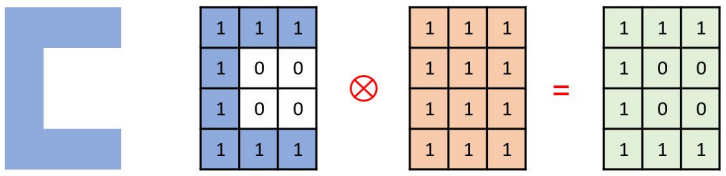


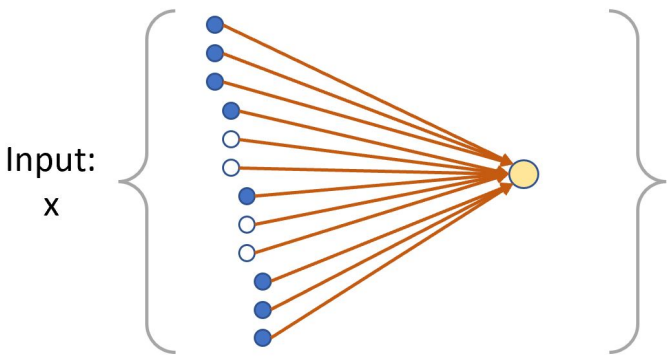
Basic Math Operations for Artificial NNs



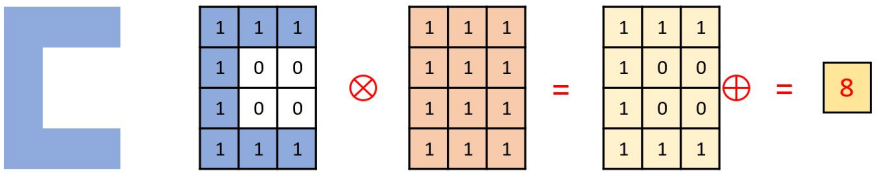


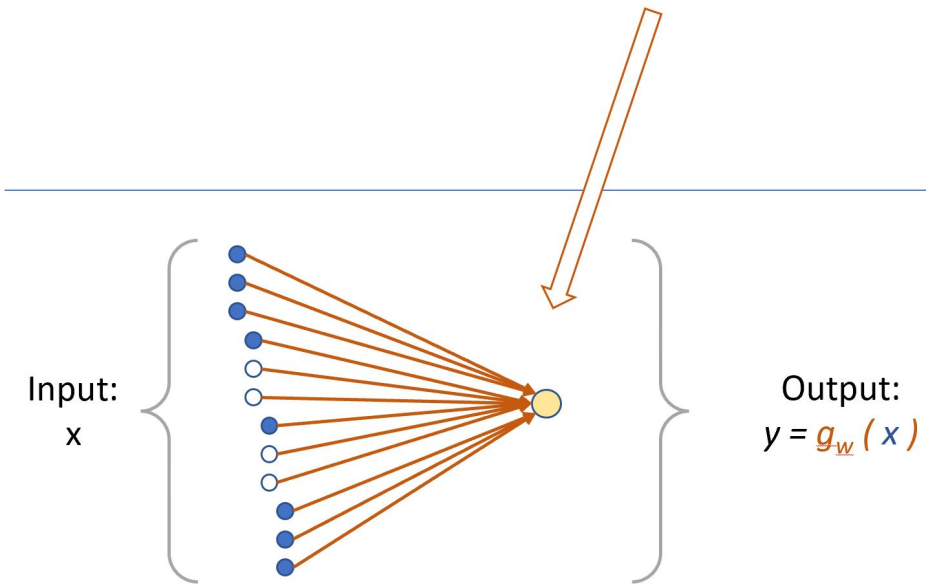
Basic Math Operations for Artificial NNs



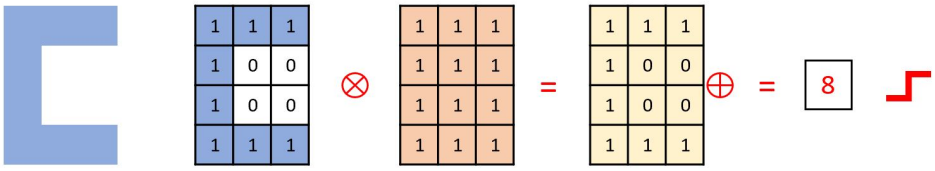


Basic Math Operations for Artificial NNs

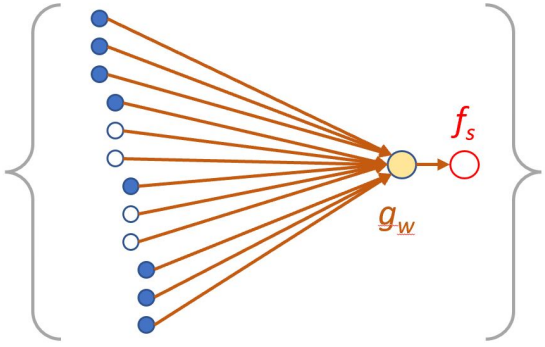




Basic Math Operations for Artificial NNs



If result > 0, then output 1 (= is letter “C”)

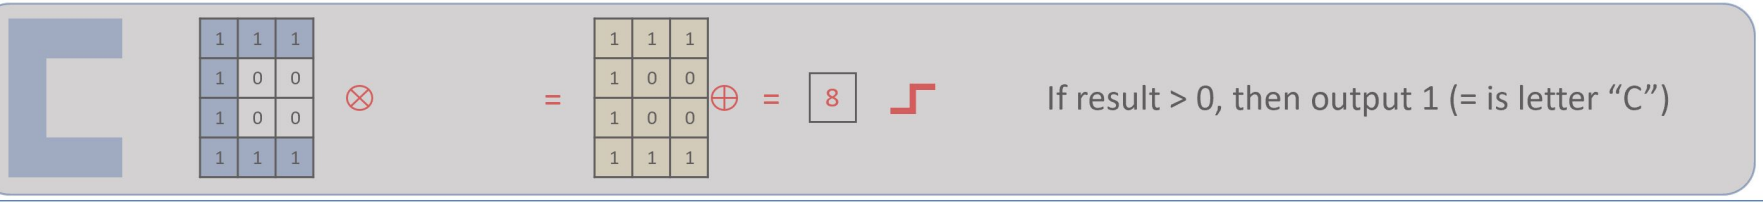


Output:
$y = f_s (g_w (x))$

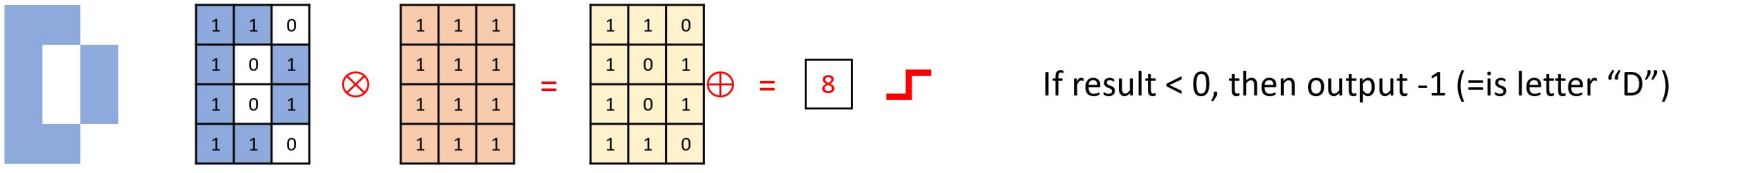
Historically, this is a perceptron (‘60)

Mathematically, we have a function composition (weighted linear combination and linear threshold).

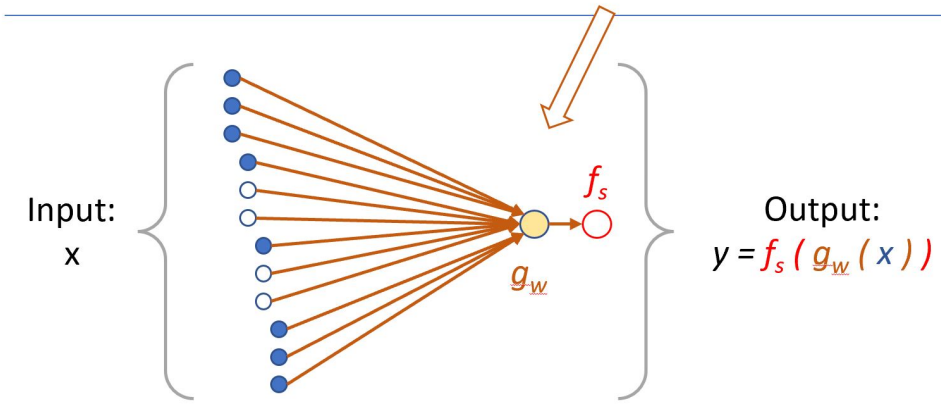
Basic Math Operations for Artificial NNs



If result > 0, then output 1 (= is letter “C”)



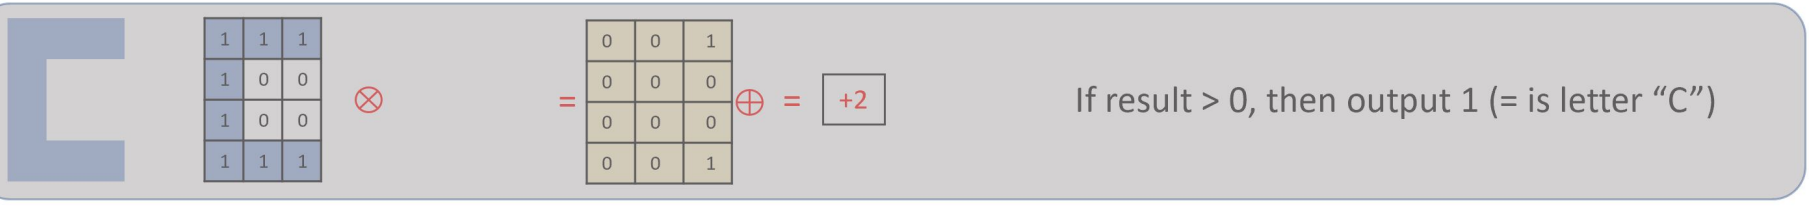
If result < 0, then output -1 (=is letter “D”)

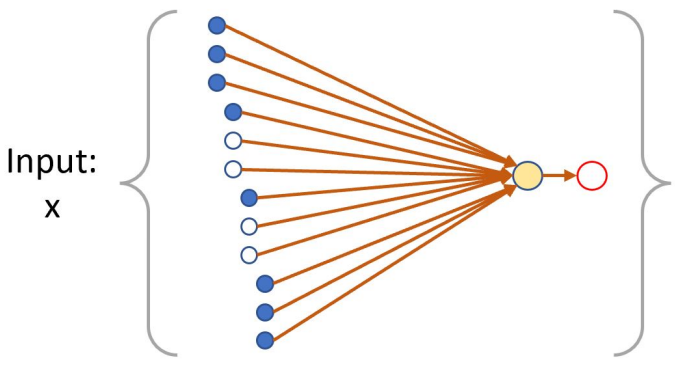


Historically, this is a perceptron (‘60)

Mathematically, we have a function composition (weighted linear combination and linear threshold).

Basic Math Operations for Artificial NNs





Historically, this is a perceptron (‘60)

Basic Math Operations for Artificial NNs



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

⊗

weights w

0	0	1
0	0	-1
0	0	-1
0	0	1

=

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

⊕ = +2

activation function



If result > 0, then output 1 (= is letter “C”)



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

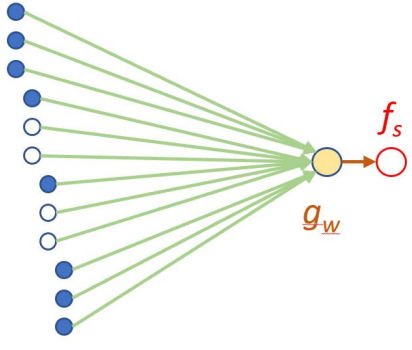
⊗ =

0	0	0
0	0	-1
0	0	-1
0	0	0

⊕ = -2

If result < 0, then output -1 (=is letter “D”)

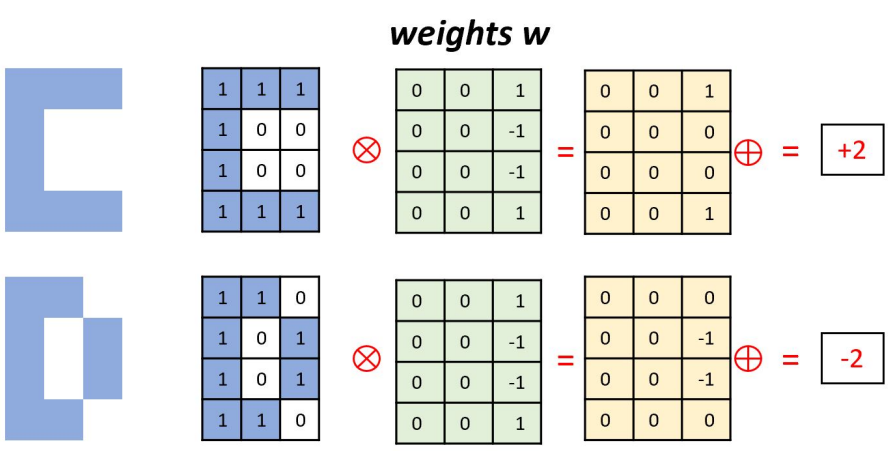
Input:
x

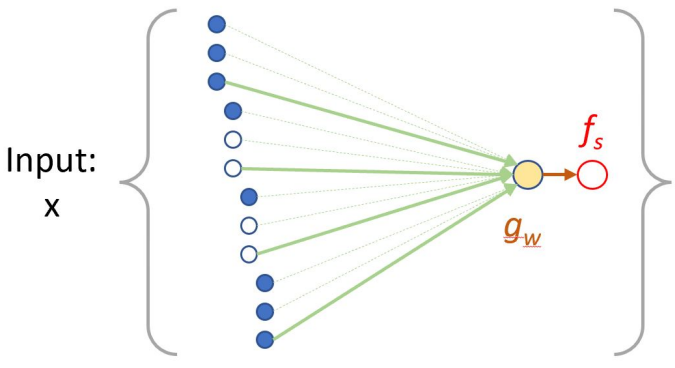


Output:
$y = f_s (g_w (x))$

Historically, this is a perceptron (‘60)

Basic Math Operations for Artificial NNs



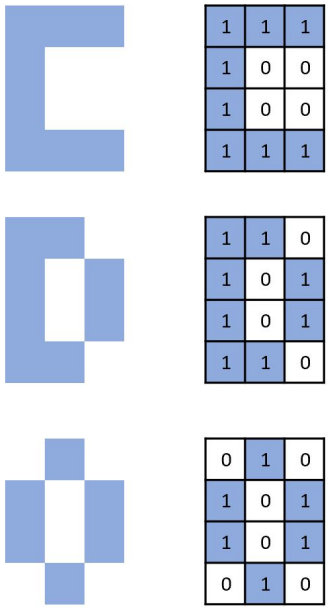


Historically, this is a perceptron (‘60)

Basic Math Operations for Artificial NNs



Multiclass Classification



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

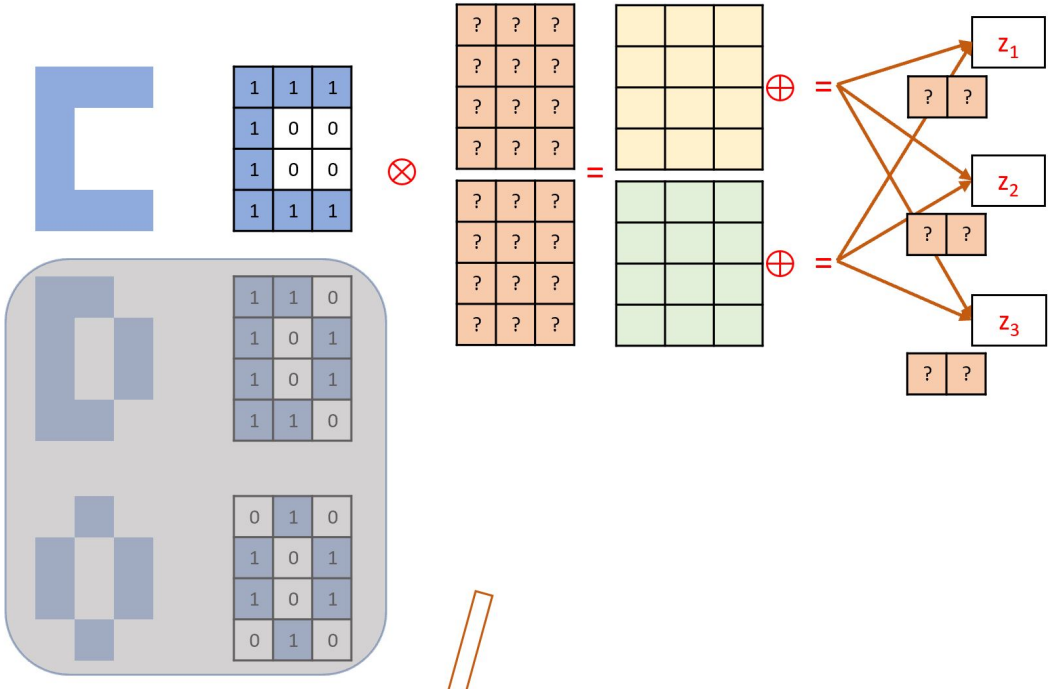
1	1	0
1	0	1
1	0	1
1	1	0

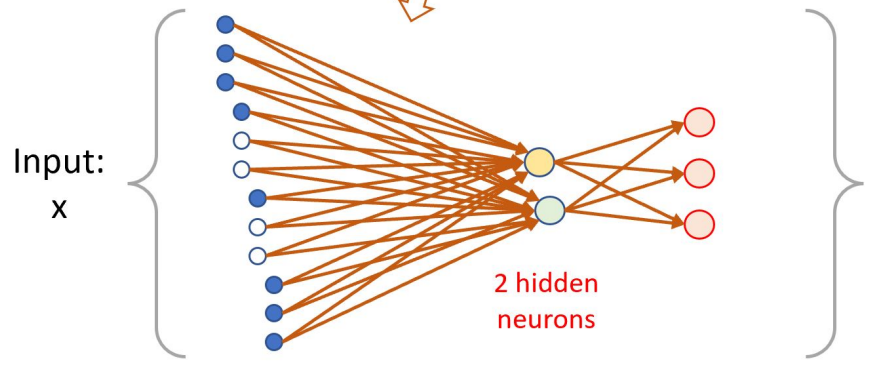
0	1	0
1	0	1
1	0	1
0	1	0

Input:
x

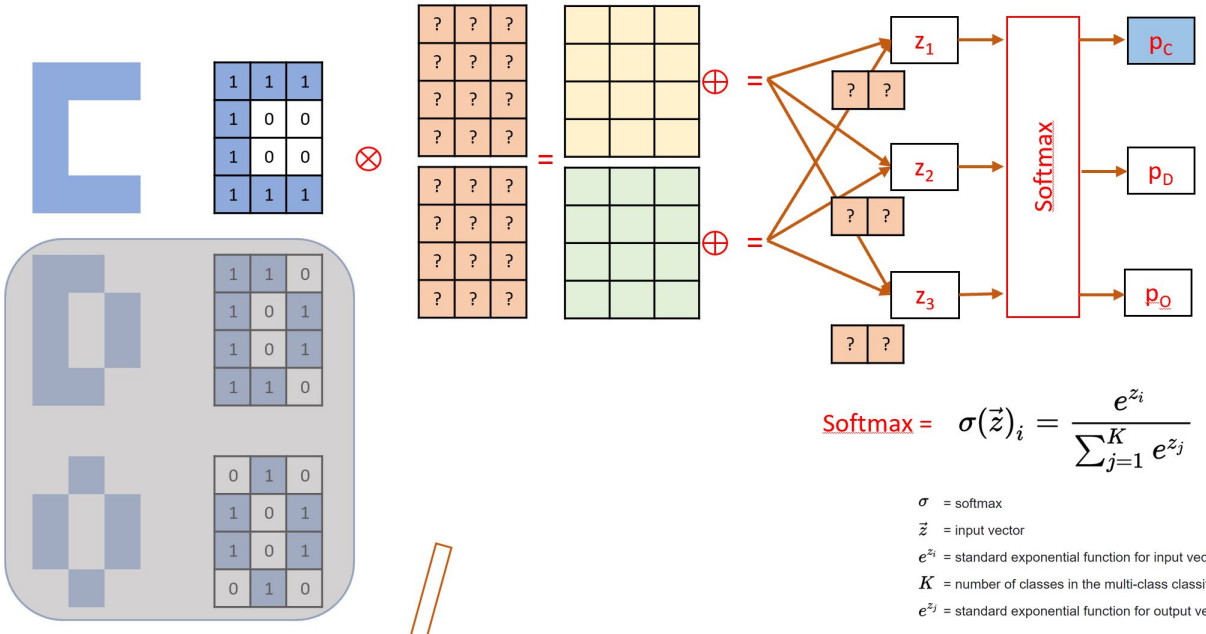
Output:
y = *?*

Multiclass Classification





Multiclass Classification

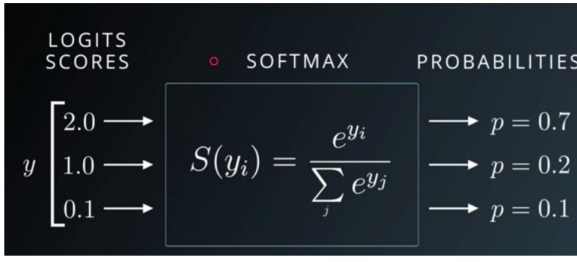


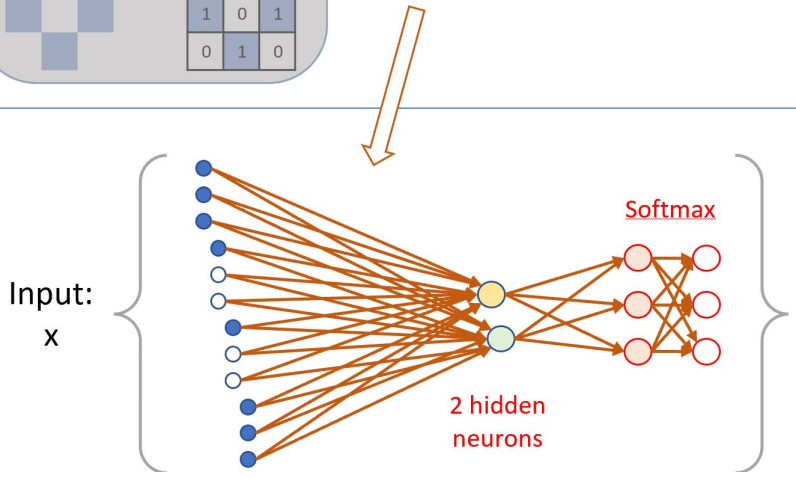
The output is the probability to be a given letter (or class):

$p_C + p_D + p_O = 1.0$

Softmax = $\sigma(\vec{z})_i = \frac{e^{z_i}}{\sum_{j=1}^{K} e^{z_j}}$

σ = softmax
$\vec{z}$ = input vector
e^{z_i} = standard exponential function for input vector
K = number of classes in the multi-class classifier
e^{z_j} = standard exponential function for output vector

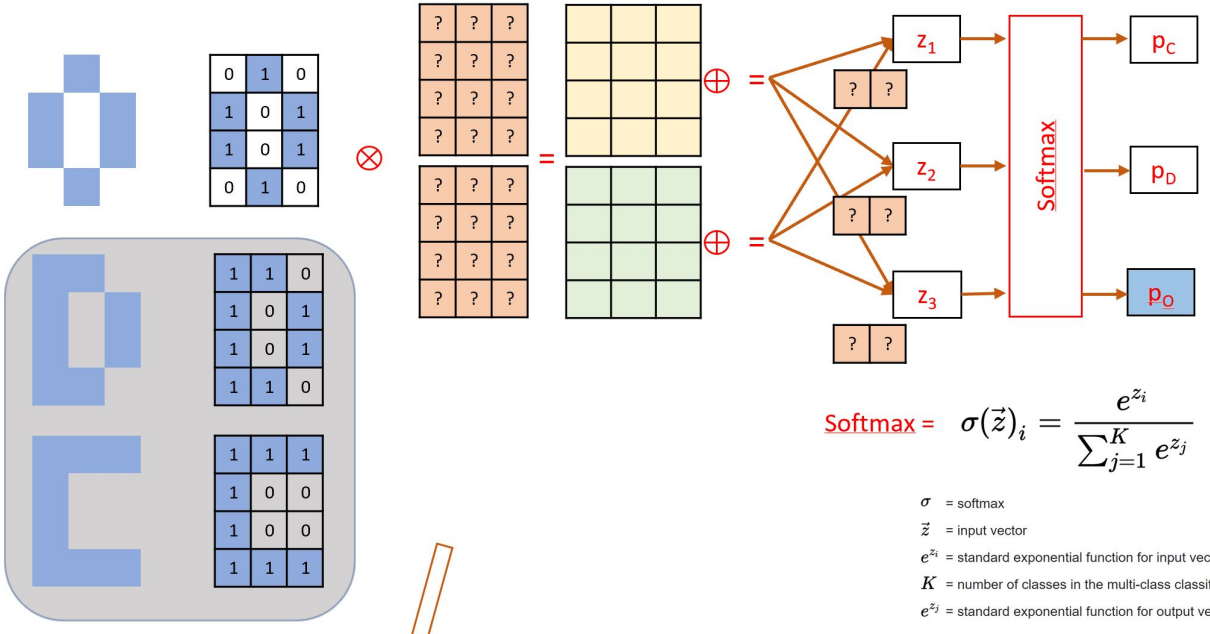




Input:
x

Output:
$y = Softmax(f_s (g_w (x)))$

Multiclass Classification

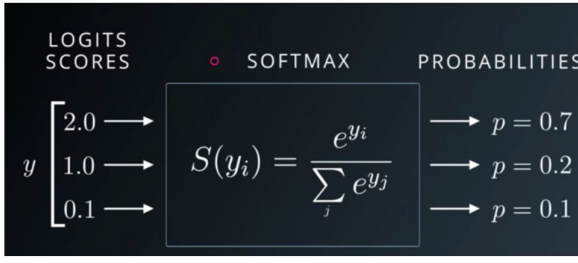


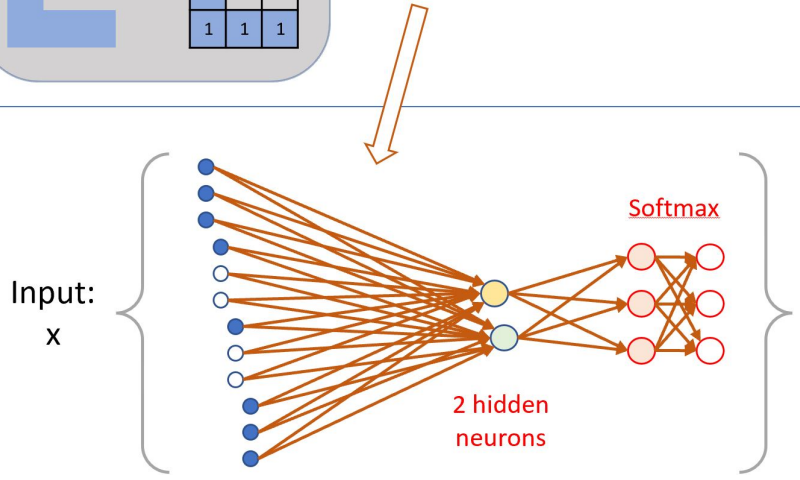
The output is the probability to be a given letter (or class):

$p_C + p_D + p_O = 1.0$

Softmax = $\sigma(\vec{z})_i = \frac{e^{z_i}}{\sum_{j=1}^{K} e^{z_j}}$

σ = softmax
$\vec{z}$ = input vector
e^{z_i} = standard exponential function for input vector
K = number of classes in the multi-class classifier
e^{z_j} = standard exponential function for output vector

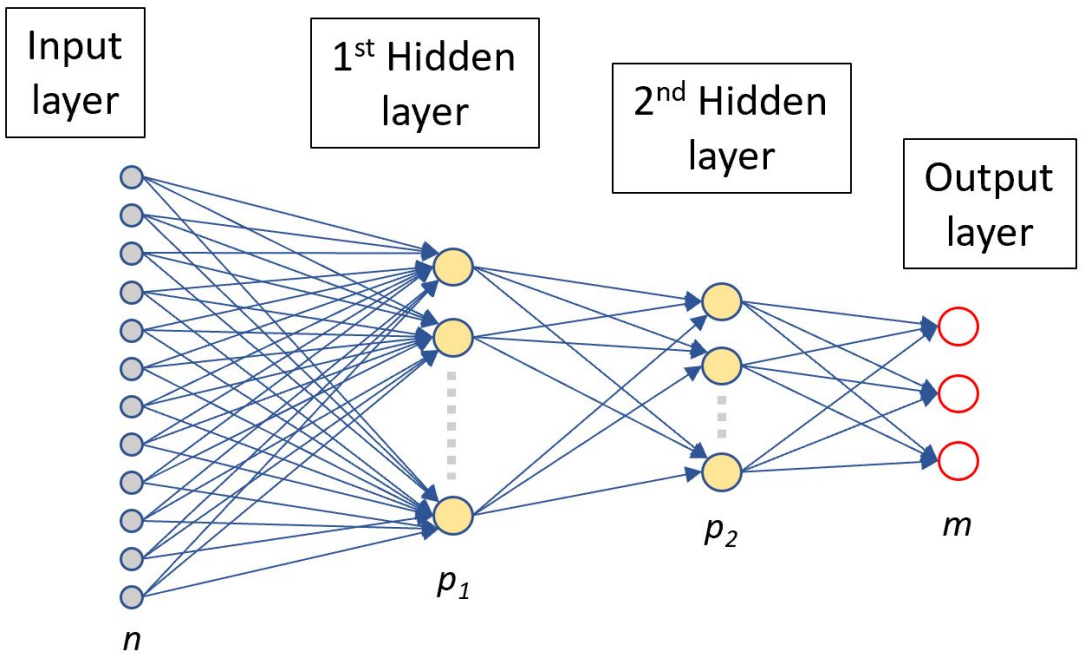




Input:
x

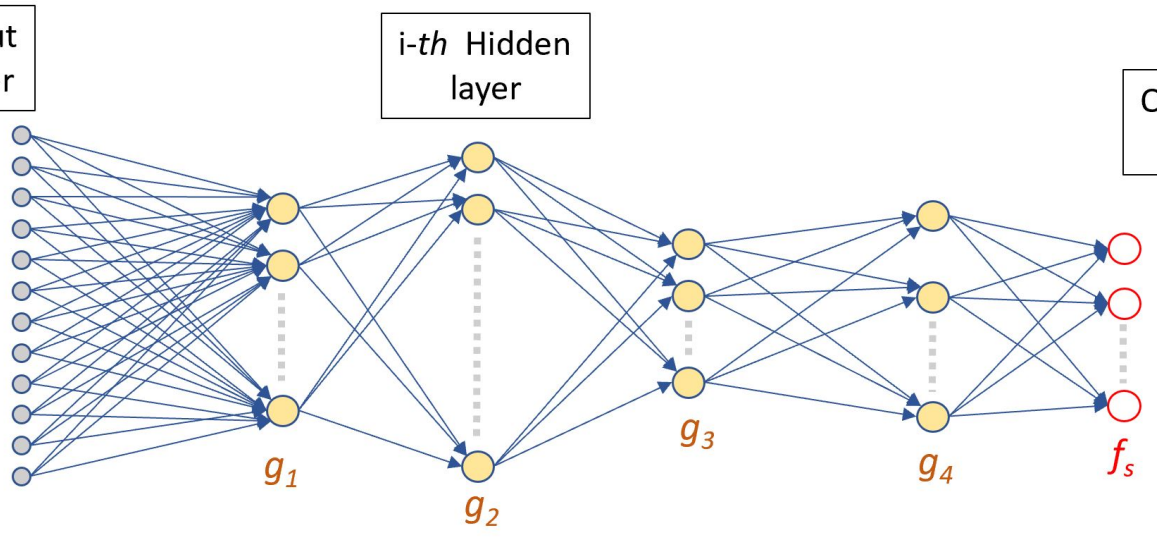
Output:
$y = Softmax(f_s (g_w (x)))$

Neural Network Architecture



We have in total to "learn"
$n\,p_1 + p_1\,p_2 + p_2\,m$
parameters

Multiclass Classification



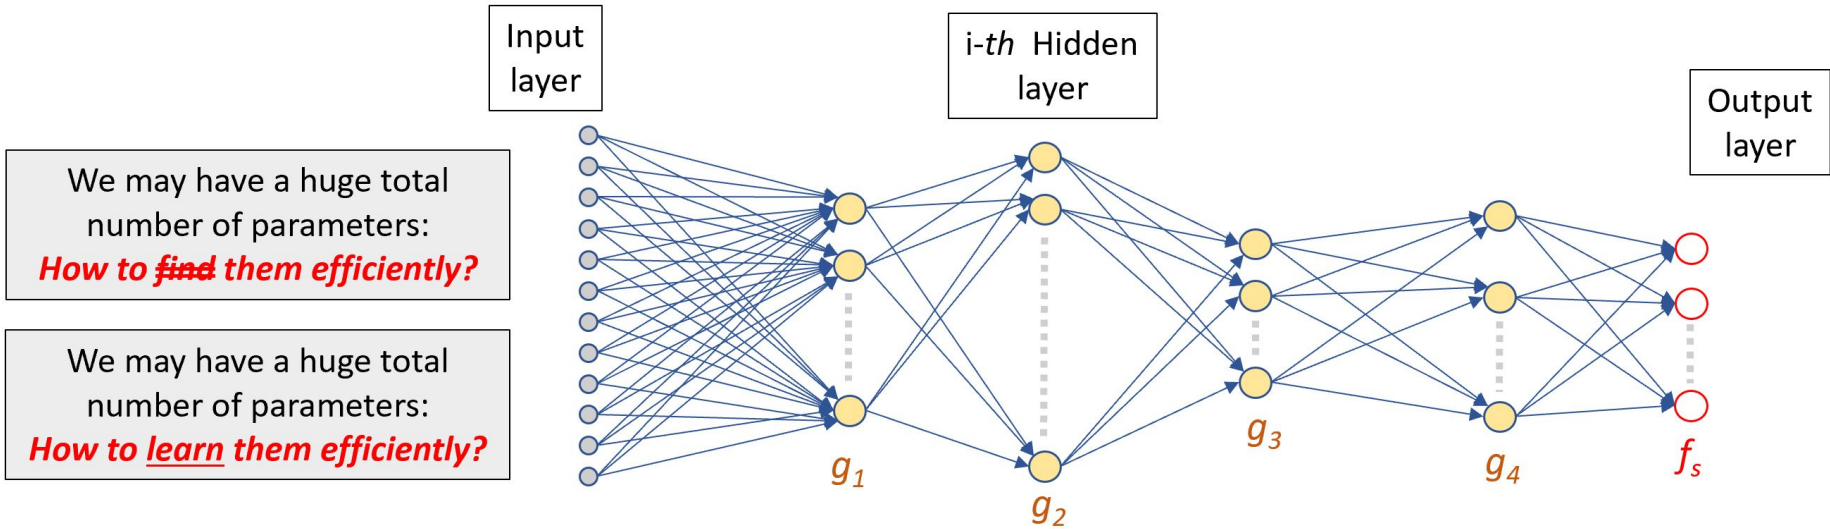
We may have a huge total number of parameters:
How to ~~find~~ them efficiently?

We may have a huge total number of parameters:
How to learn them efficiently?

By computing the first derivative of composite function:

$$\mathbf{y} = f_s (g_4 (g_3 (g_2 (g_1 (\mathbf{x})))))$$

Multiclass Classification



By computing the first derivative of composite function: $\mathbf{y} = f_s(g_4(g_3(g_2(g_1(\mathbf{x})))))$

1. Number of hidden layers
2. Number of neurons at each layer
3. Type of activation function at each layer
4. Type of output function

Question: Do we really want fully connected layers?

Universal Approximation Theorem

Theorem 3

Let $d \in \mathbb{N}$, $K \subset \mathbb{R}^d$ compact, $f : K \to \mathbb{R}$ continuous, $\rho : \mathbb{R} \to \mathbb{R}$ continuous and not a polynomial. Then, for each $\epsilon > 0$, there exists $N \in \mathbb{N}$ and $a_k, b_k \in \mathbb{R}, w_k \in \mathbb{R}^d, 1 \leq k \leq N$, such that:

$$\left\| f - \sum_{k=1}^{N} a_k \rho(\langle w_k, \cdot \rangle - b_k) \right\|_\infty \leq \epsilon.$$

REMARK 1: Each continuous function on a compact domain can be approximated up to an arbitrary accuracy by a shallow neural network.

Outline

1 Introduction

2 Multilayer Neural Networks

3 Convolutional Neural Networks

4 Transfer Learning

5 Generative Models

Convolutional Neural Network Architecture

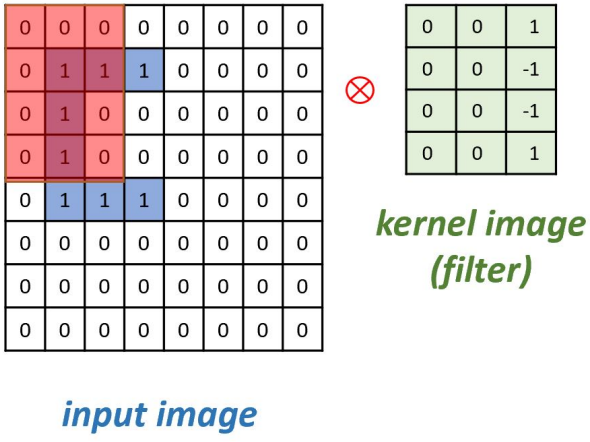
ANIMATED APPLET

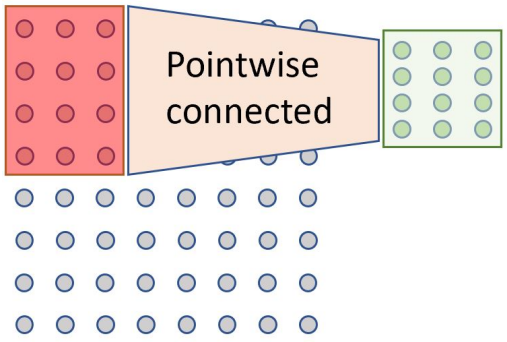
http://www.cs.cmu.edu/~aharley/vis/

Convolution

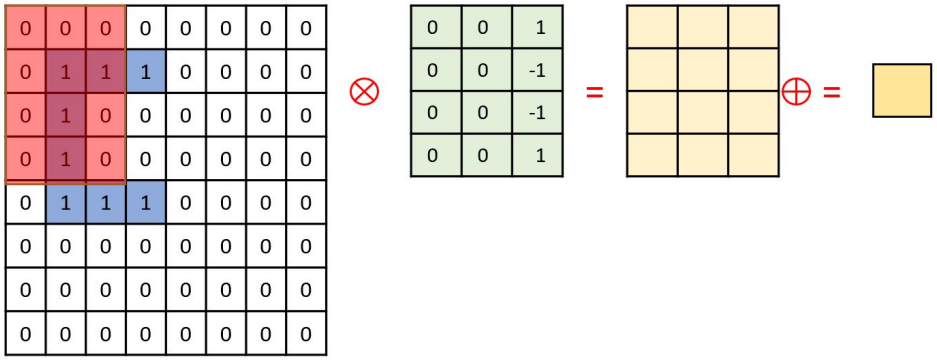
0	0	0	0	0	0	0	0
0	1	1	1	0	0	0	0
0	1	0	0	0	0	0	0
0	1	0	0	0	0	0	0
0	1	1	1	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

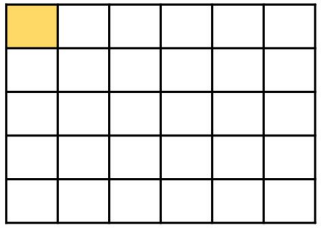
Convolution

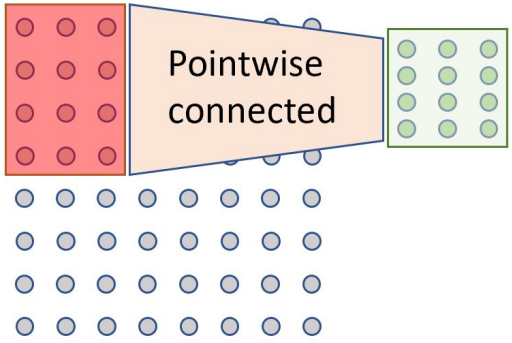




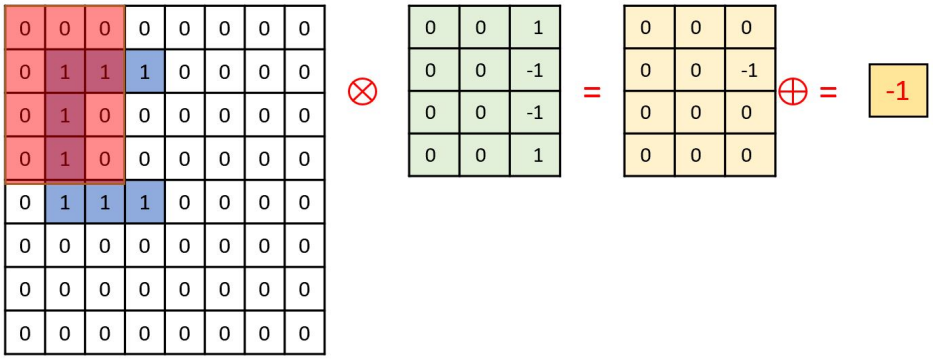
Convolution

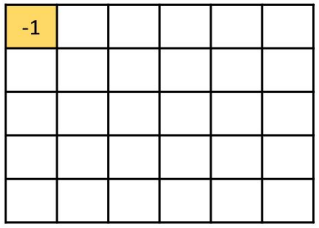


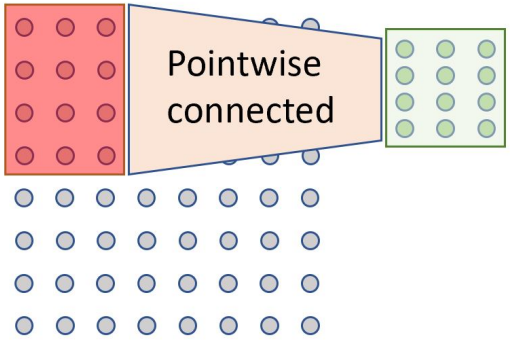




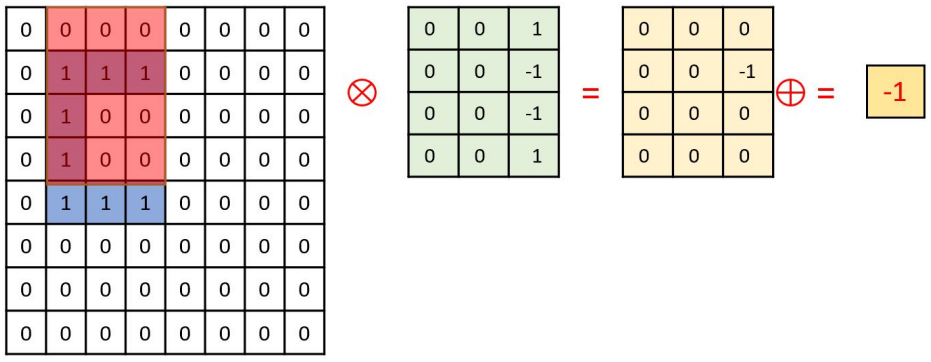
Convolution

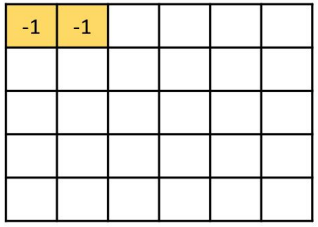


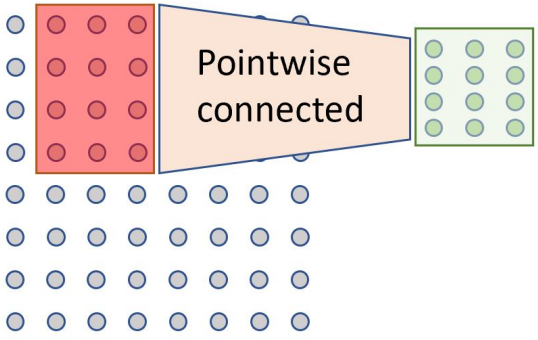




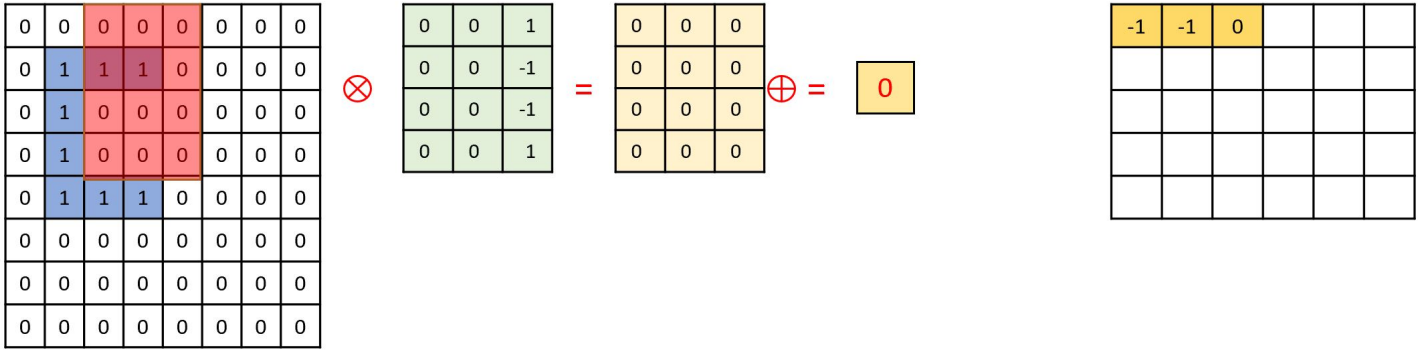
Convolution

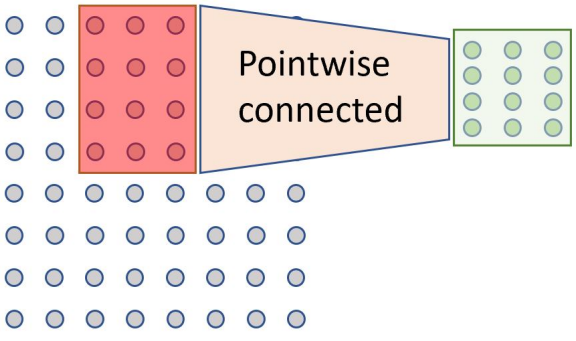




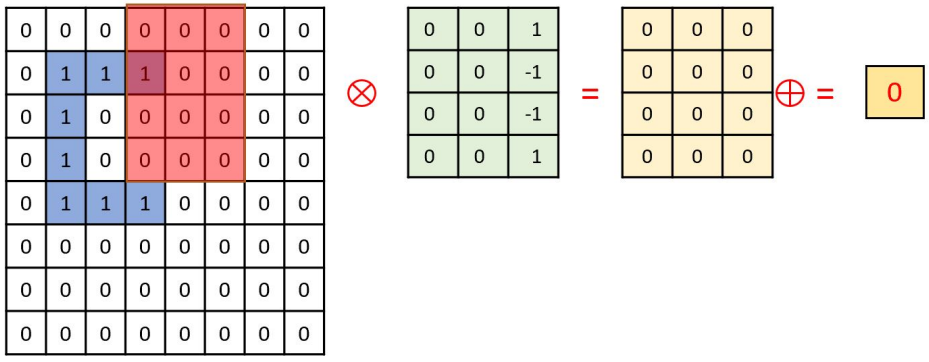


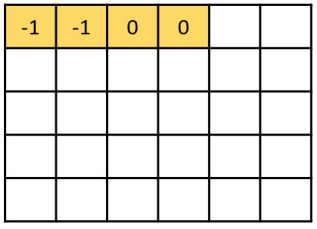
Convolution

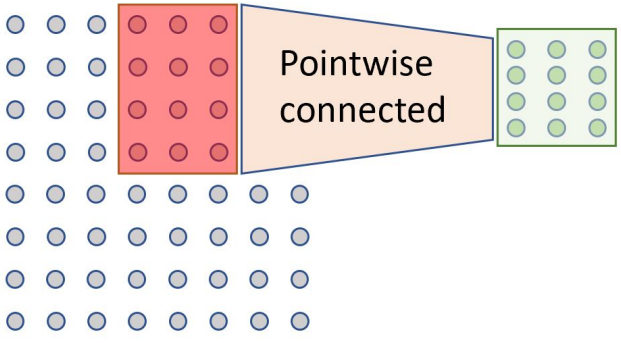




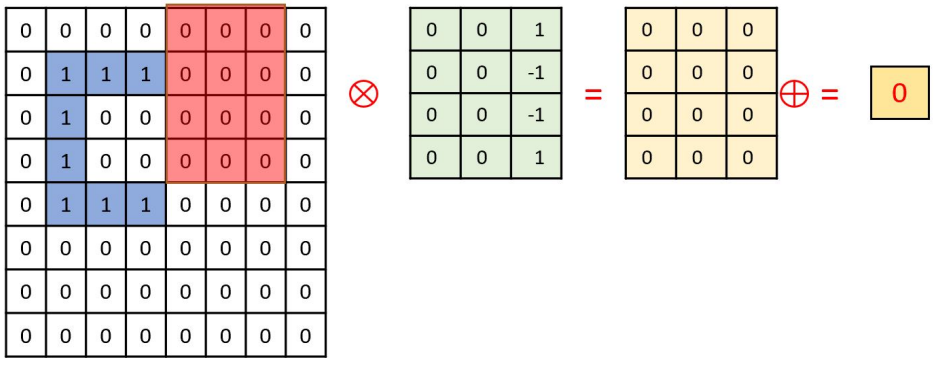
Convolution

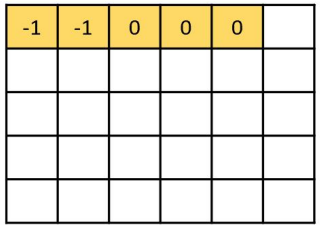


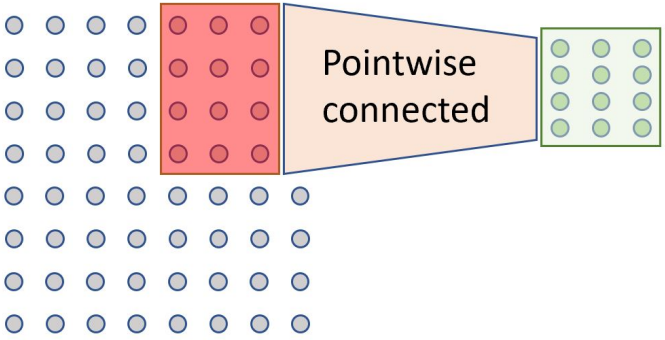




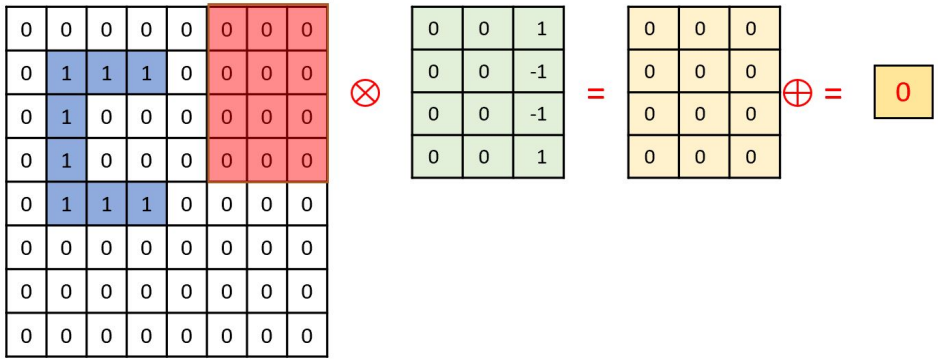
Convolution

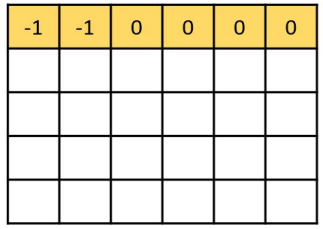


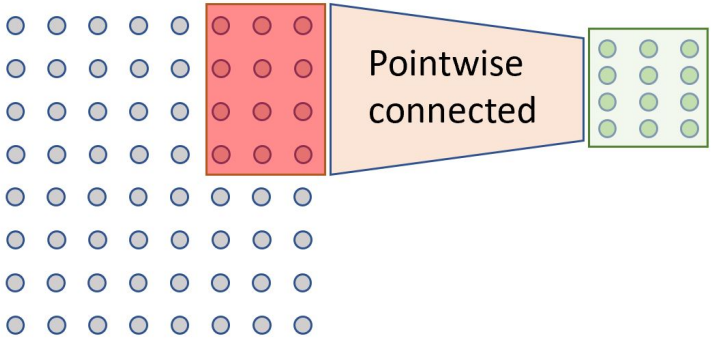




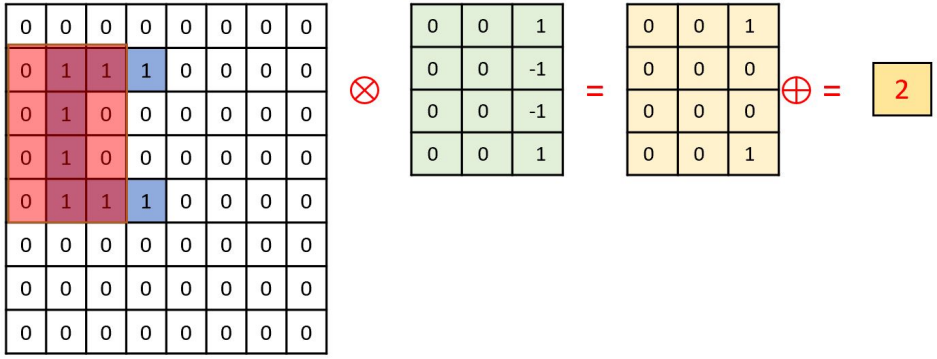
Convolution

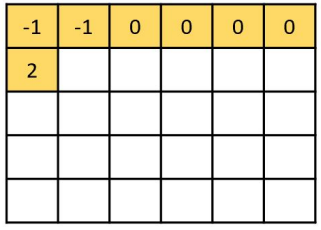


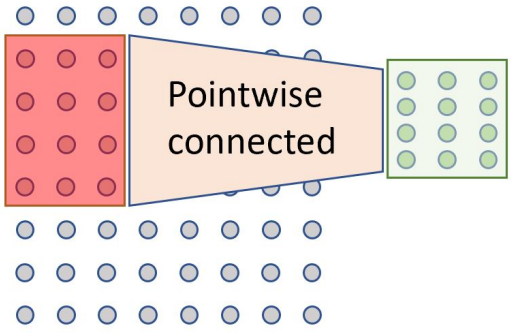




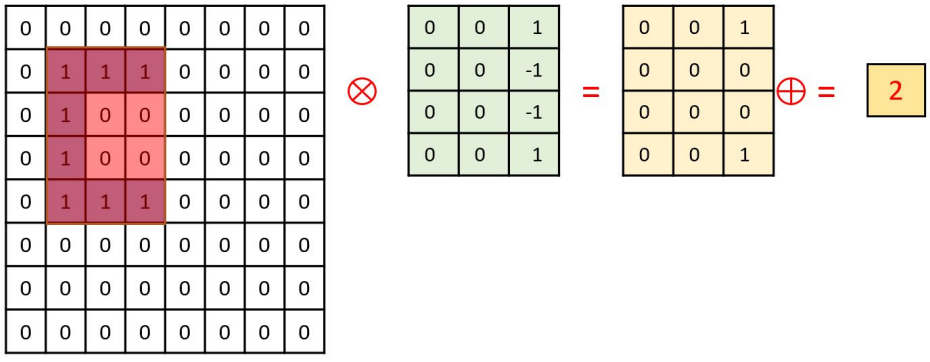
Convolution

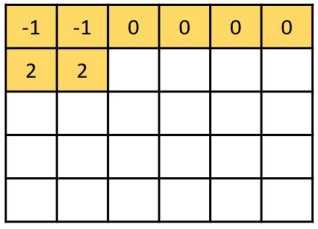


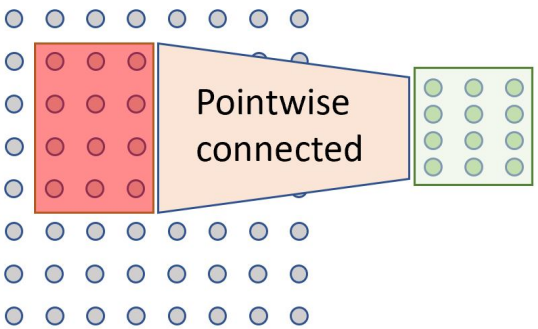




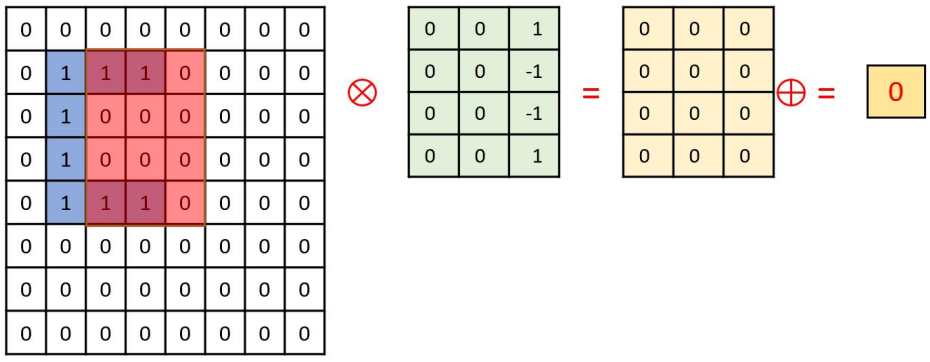
Convolution

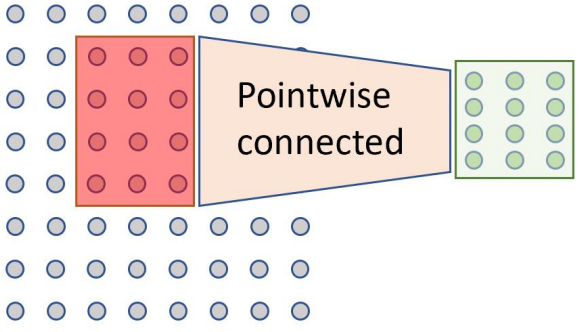




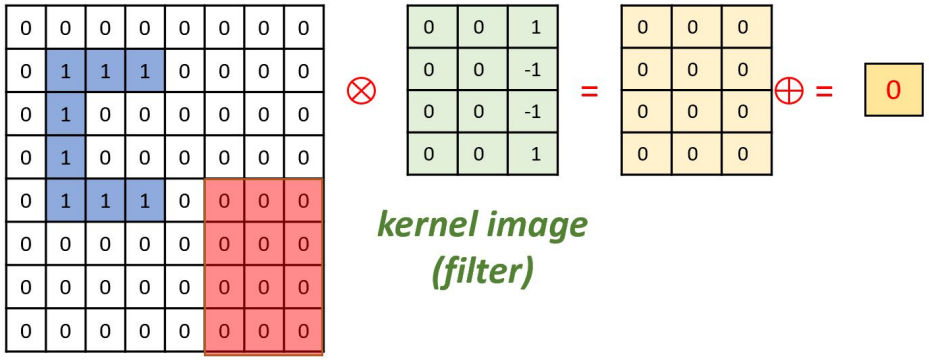


Convolution

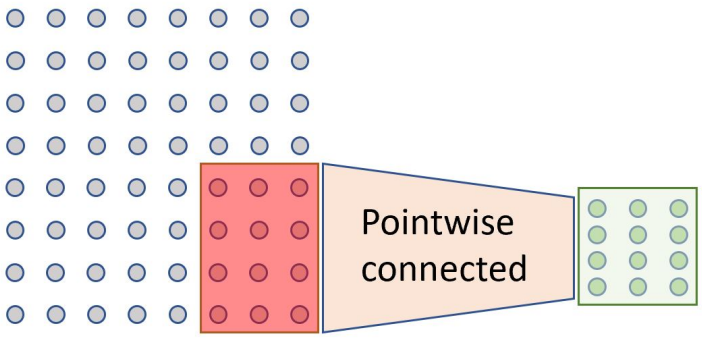




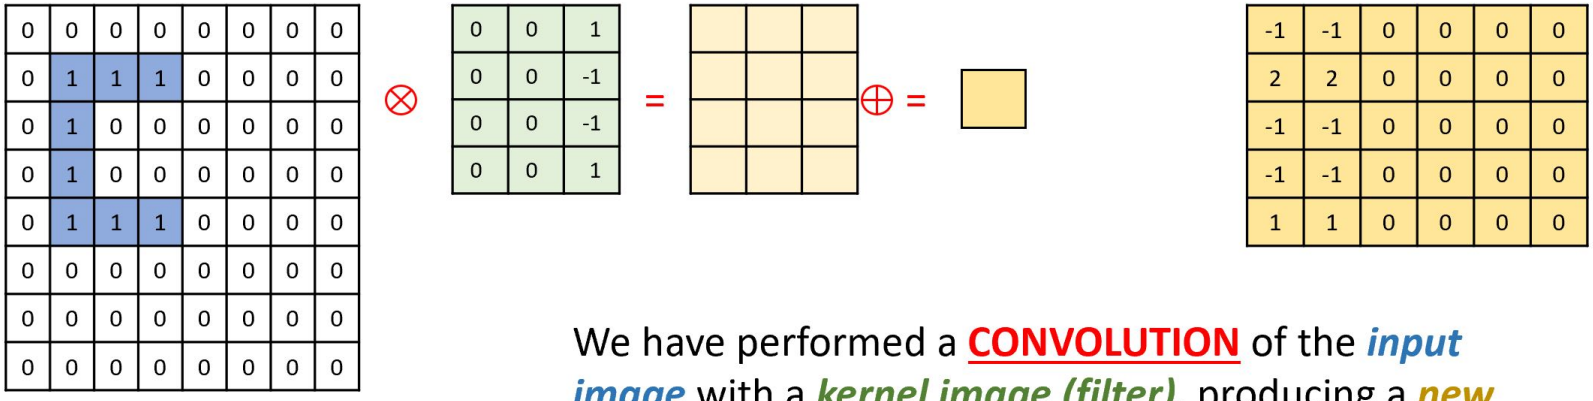
Convolution





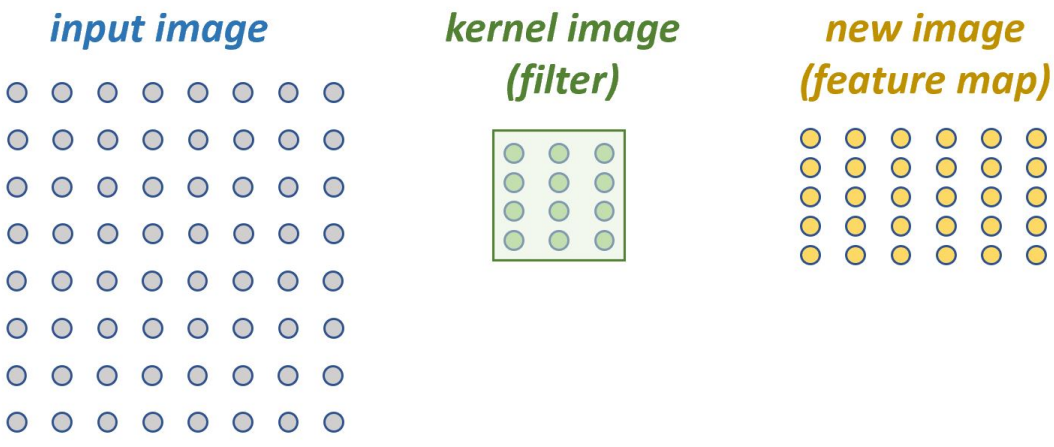


Convolution

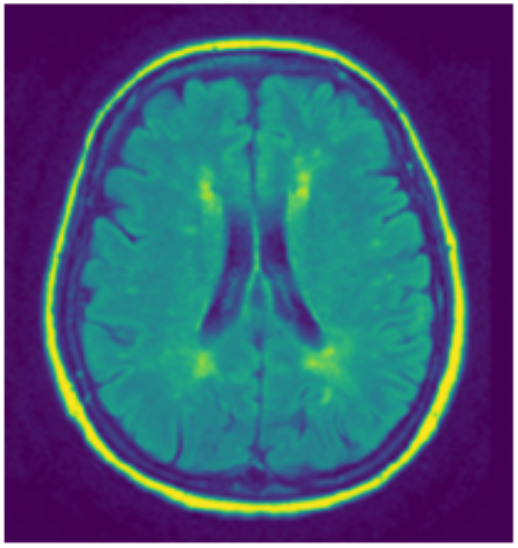


We have performed a **CONVOLUTION** of the *input image* with a *kernel image (filter)*, producing a *new image (feature map)*.

Many important image transformations are convolutions where you *convolve* the image function with a very small "*kernel function*"



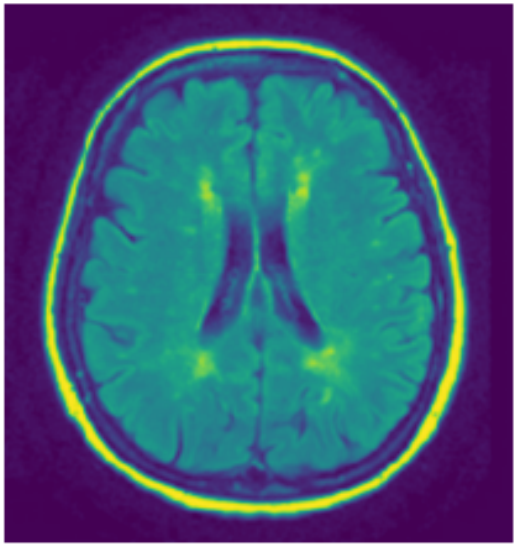
Kernels for Edge Detection



Input image:
Brain MRI
1190x1123 pixels

Two kernels/filters

Kernels for Edge Detection



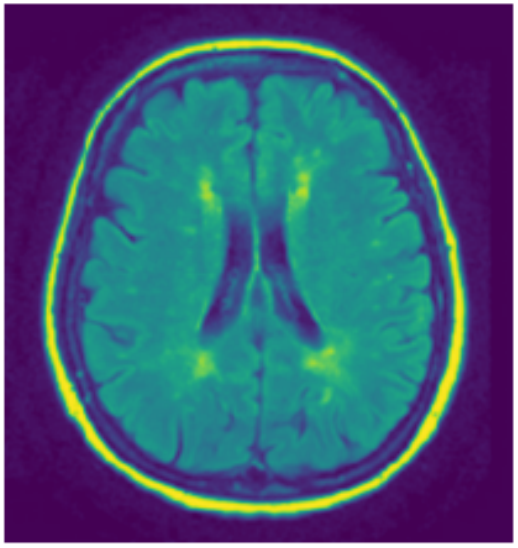
Input image:
Brain MRI
1190x1123 pixels

Two kernels/filters

$K_1 =$

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

Kernels for Edge Detection



Input image:
Brain MRI
1190x1123 pixels

Two kernels/filters

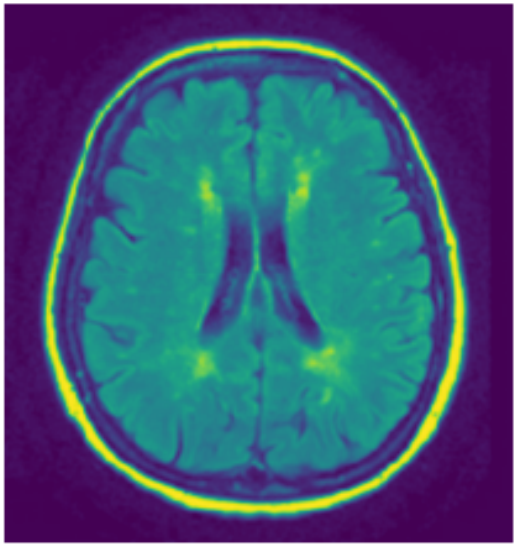
$K_1 =$

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

$K_2 =$

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

Kernels for Edge Detection



Input image:
Brain MRI
1190x1123 pixels

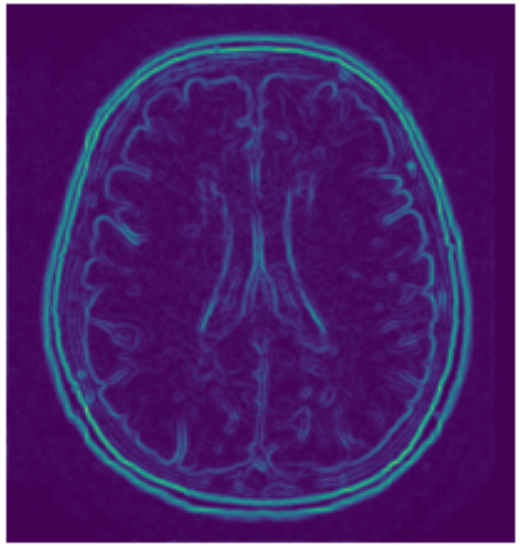
Two kernels/filters

$K_1 =$

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

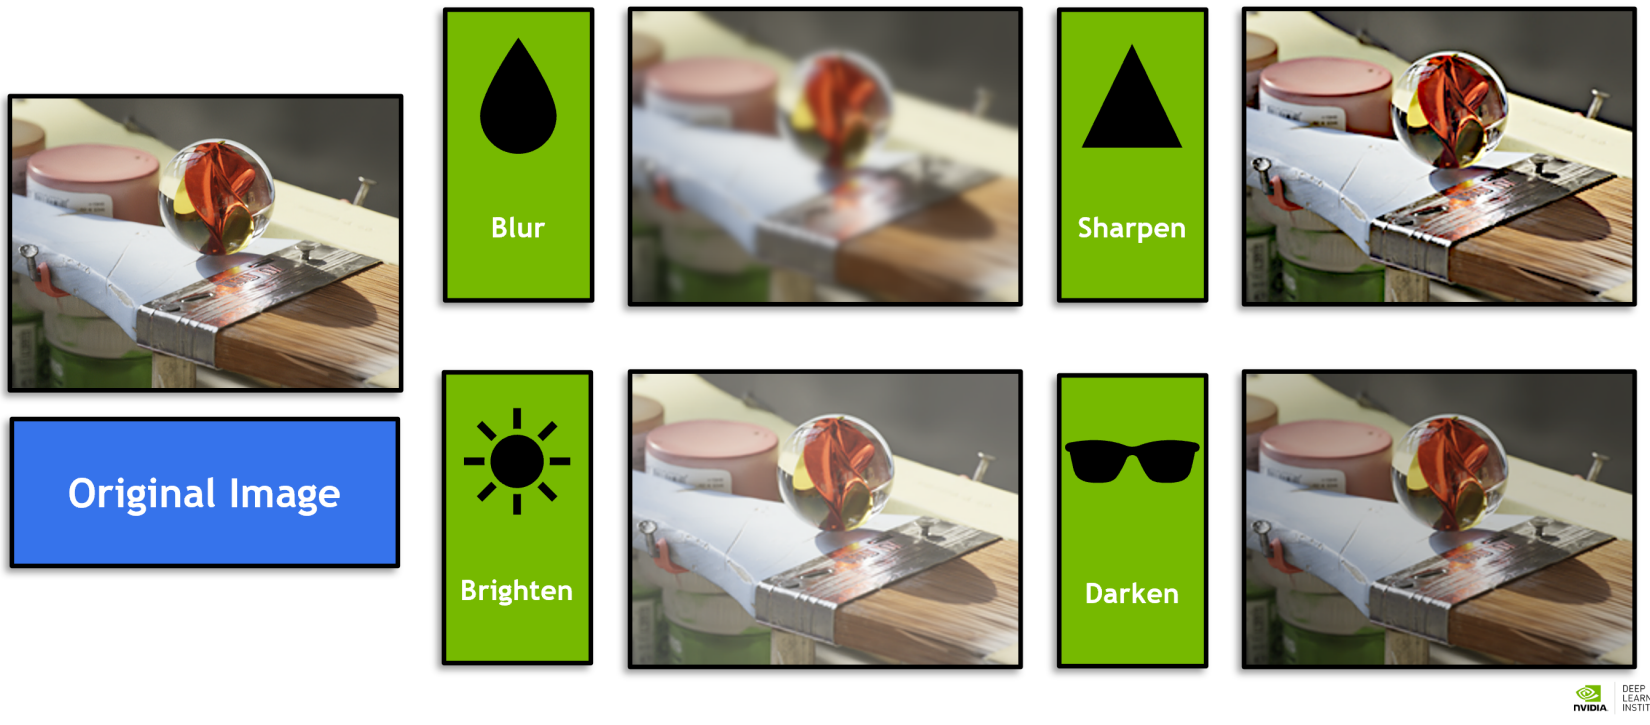
$K_2 =$

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



Feature Map:
Edge detection

Common Kernels



Common Kernels

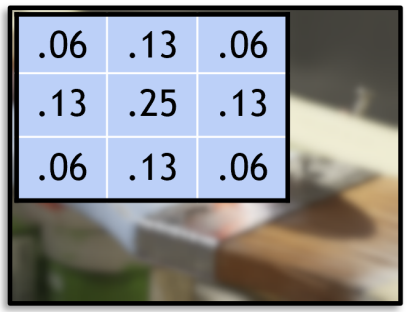
KERNELS AND CONVOLUTION



Original Image

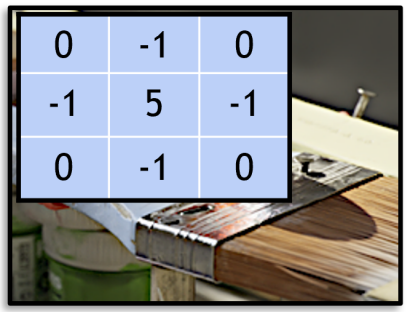
Blur

.06	.13	.06
.13	.25	.13
.06	.13	.06



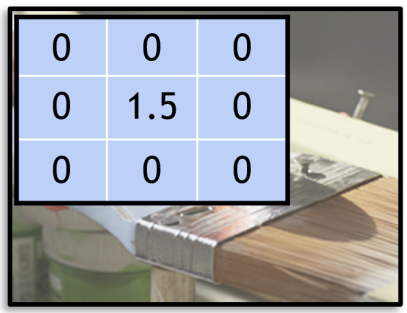
Sharpen

0	-1	0
-1	5	-1
0	-1	0



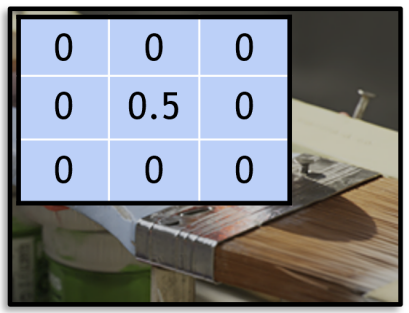
Brighten

0	0	0
0	1.5	0
0	0	0



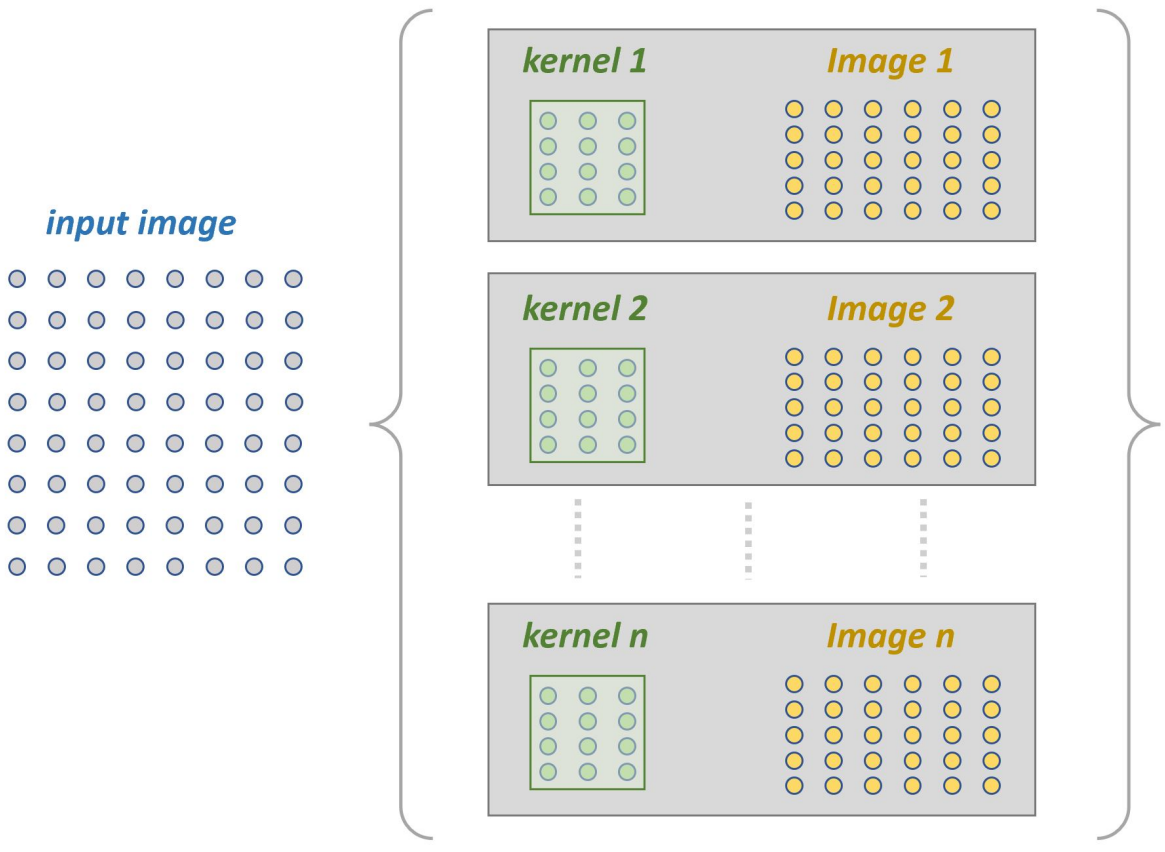
Darken

0	0	0
0	0.5	0
0	0	0



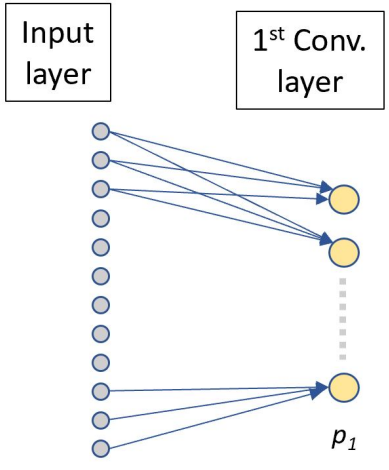
NVIDIA DEEP LEARNING INSTITUTE

Convolution

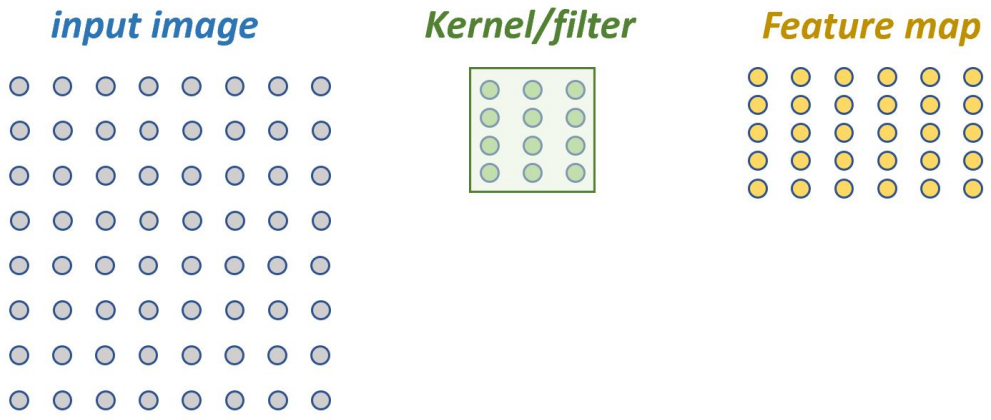


Each ***kernel function/matrix*** may be ***LEARNED*** be since is nothing else than a small matrix of weights (weighted connections of our sparse-ANN)

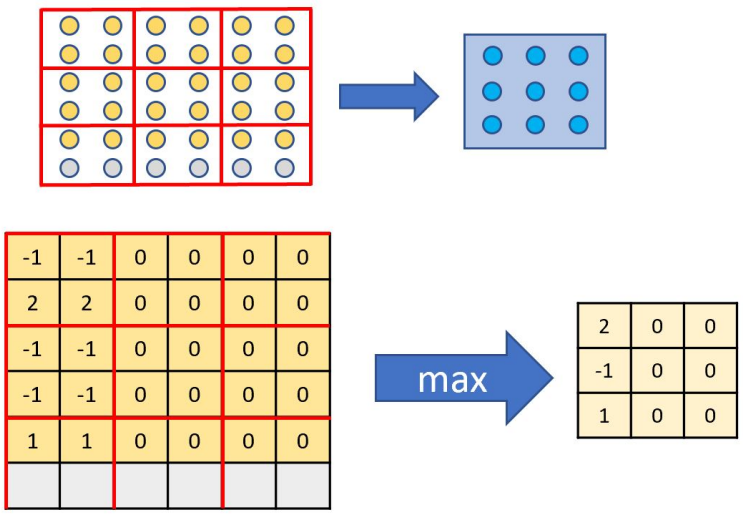
Convolutional Neural Network layer:



Pooling (subsampling)

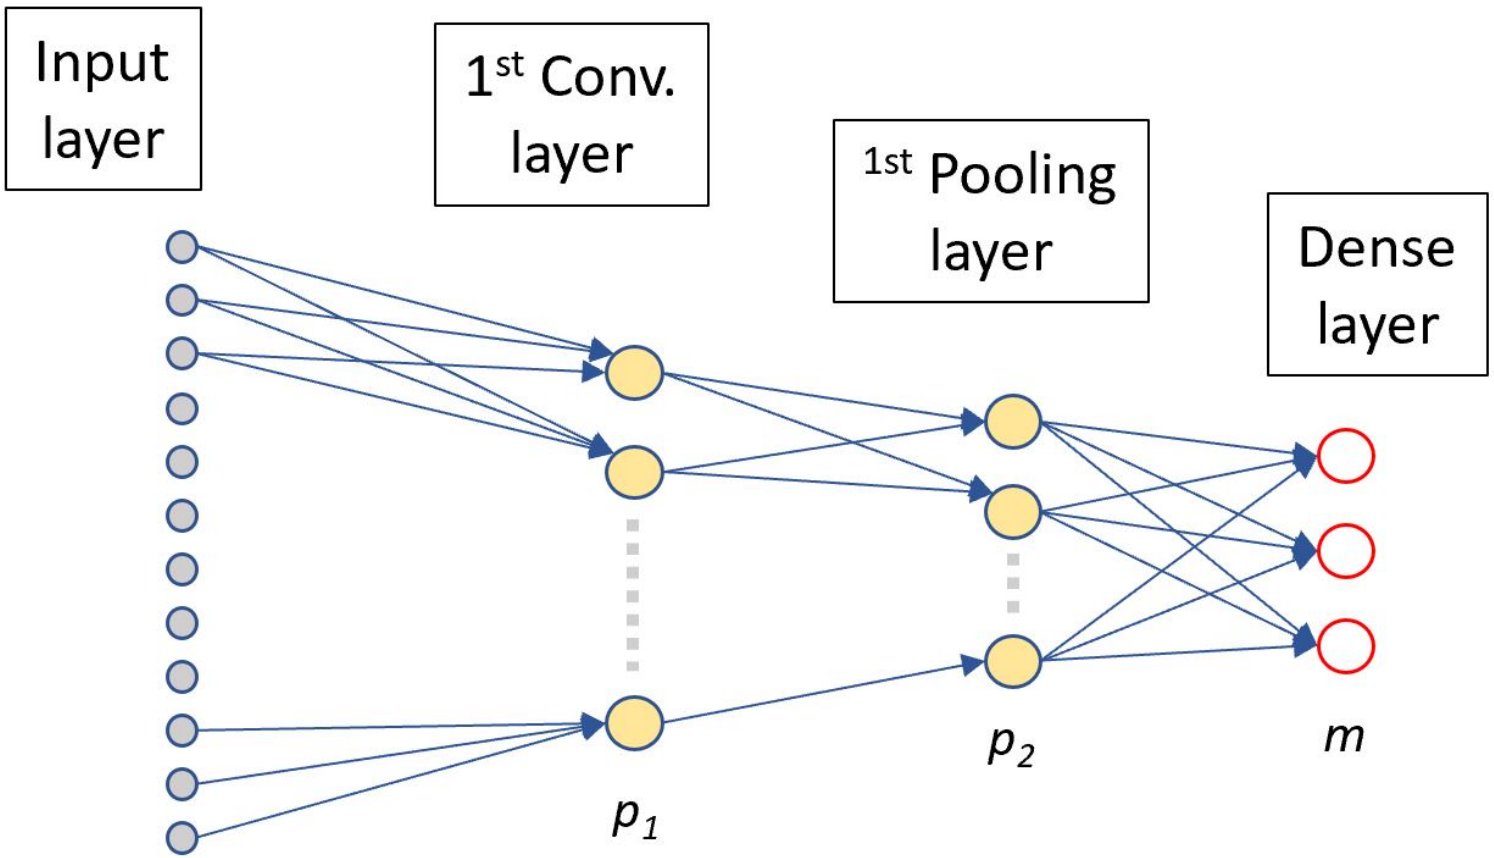


Compress images by ***maximum pooling***
For each submatrix takes the *maximum* value.



Convolutional Neural Network Architecture

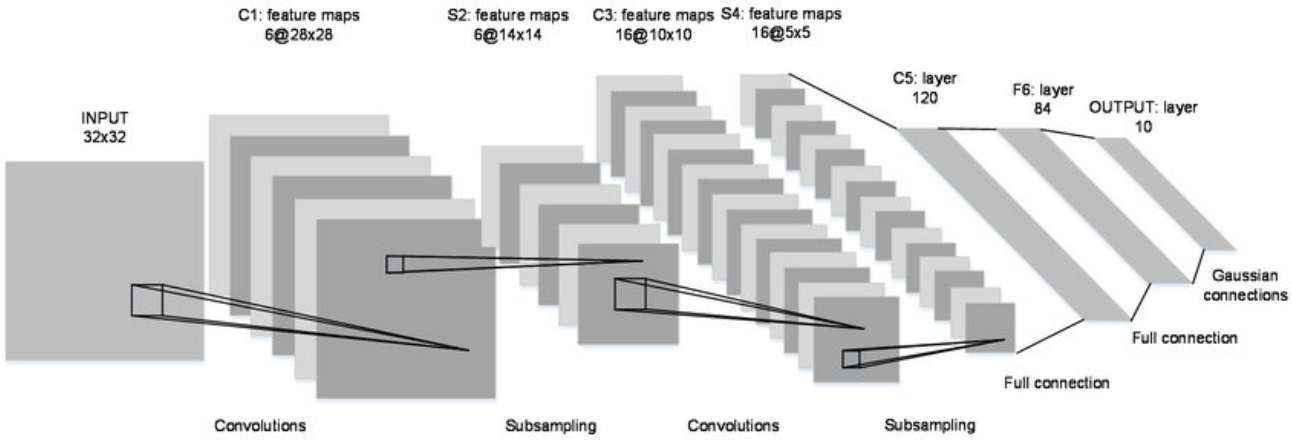
Convolutional Neural Network architecture:



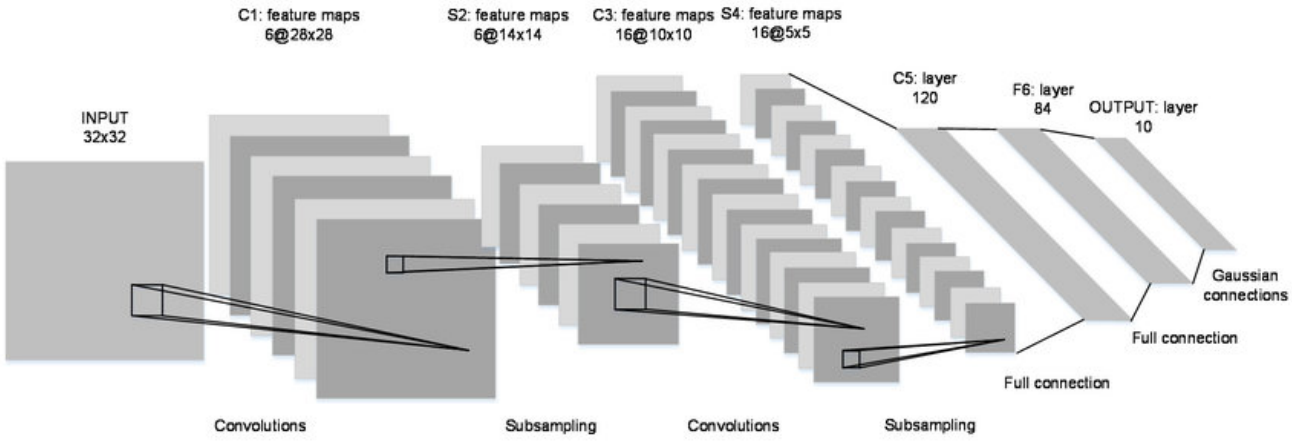
CNN in action

DEMO CLIP (FILE)

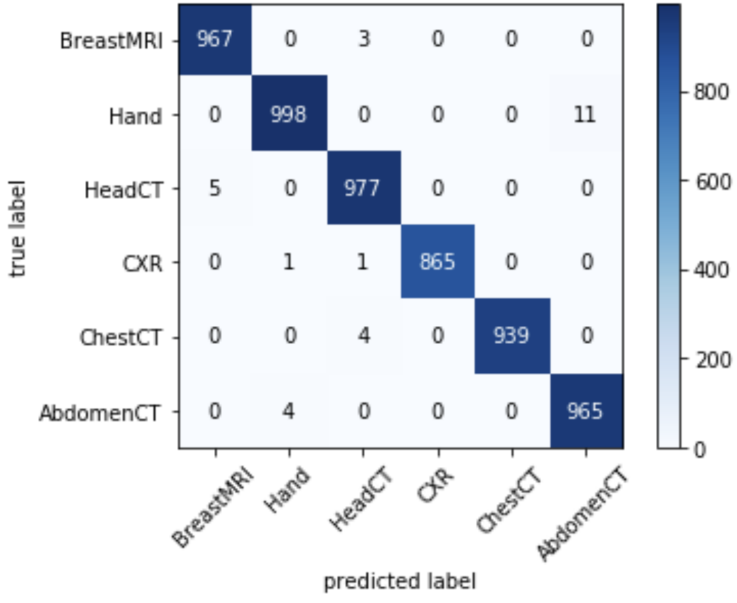
Convolutional Neural Network: MedNIST Classification



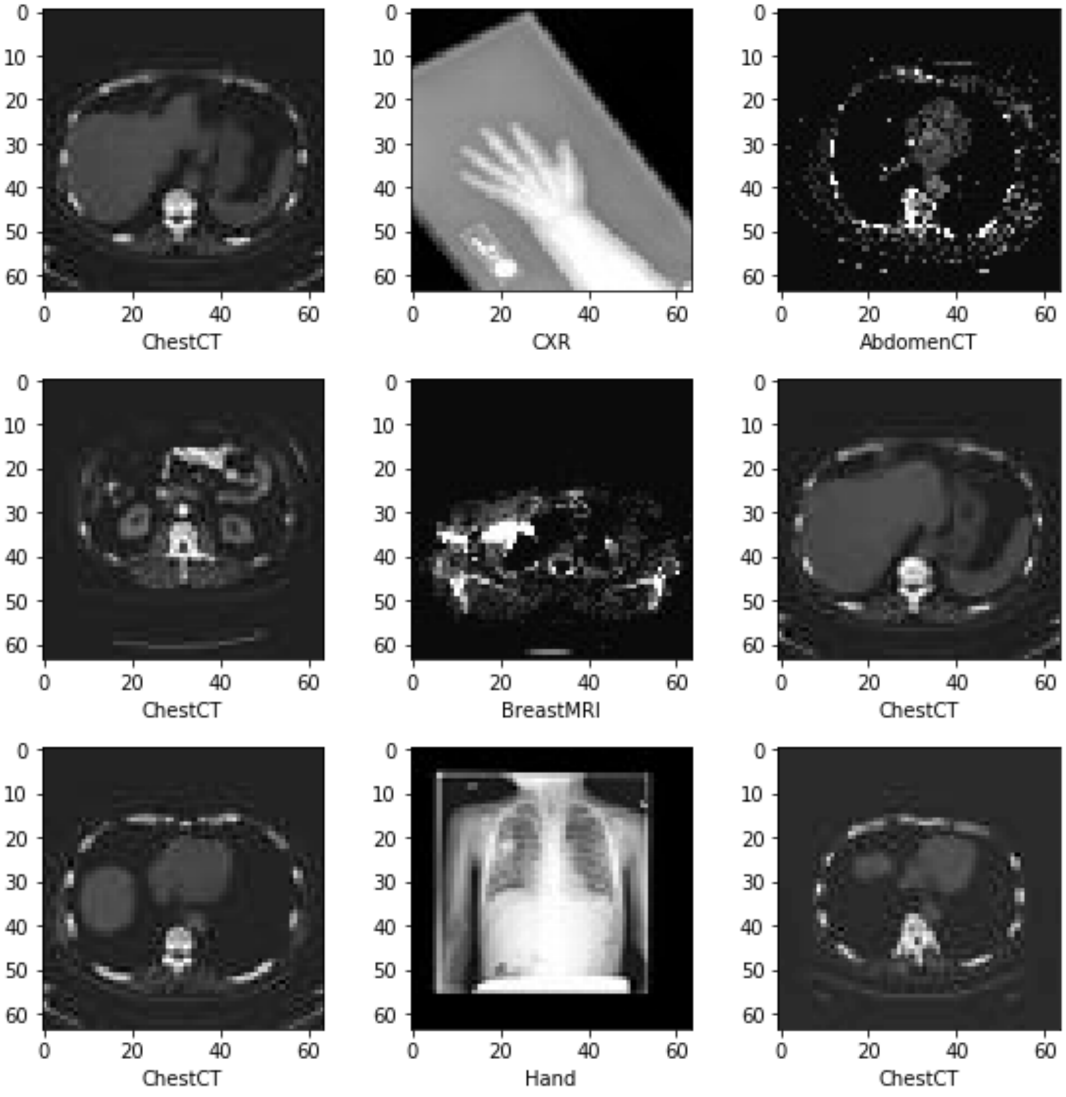
Convolutional Neural Network: MedNIST Classification



Accuracy = 0.9949477351916376



Convolutional Neural Network: MedNIST Classification



Outline

1 Introduction

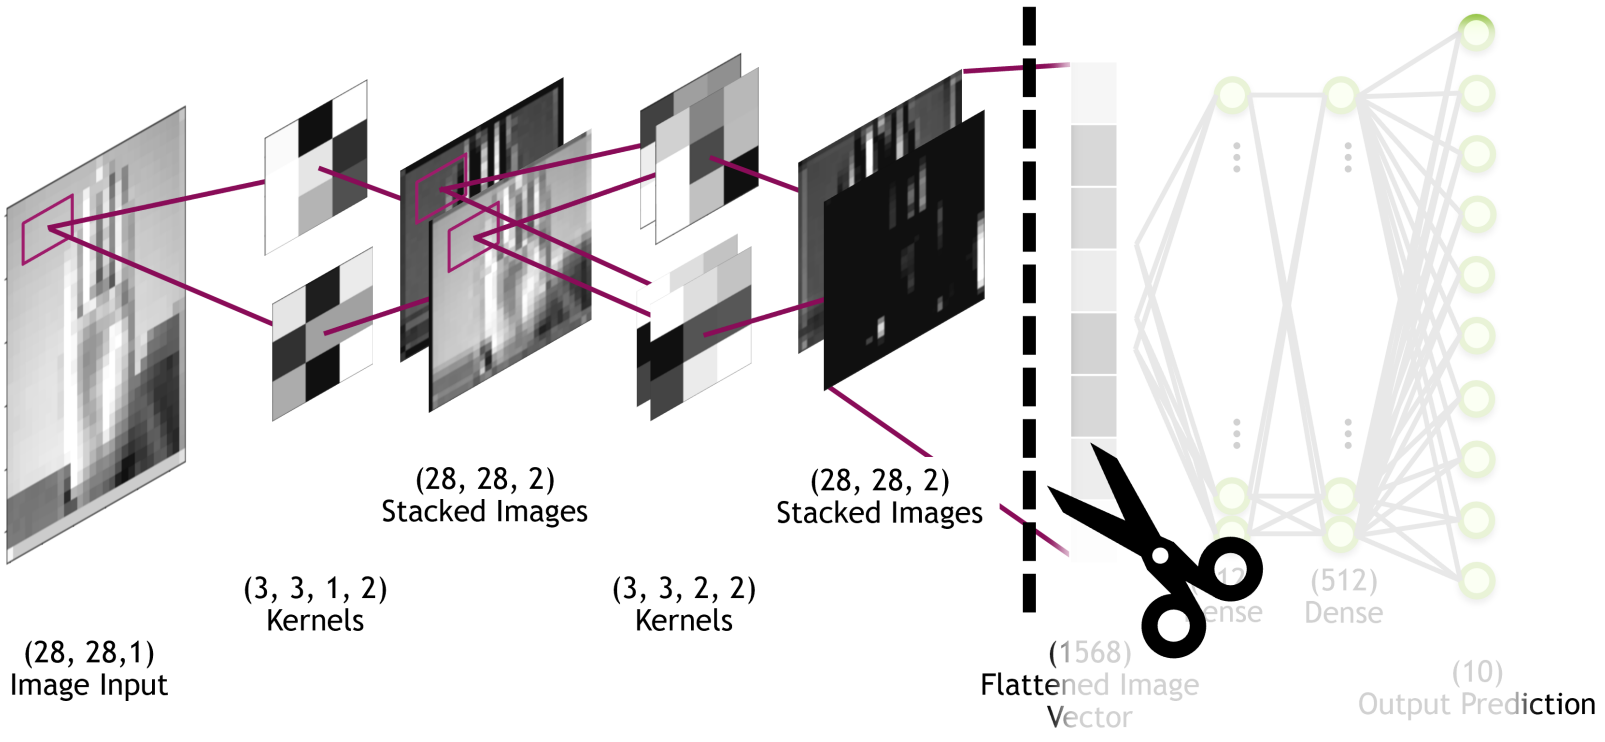
2 Multilayer Neural Networks

3 Convolutional Neural Networks

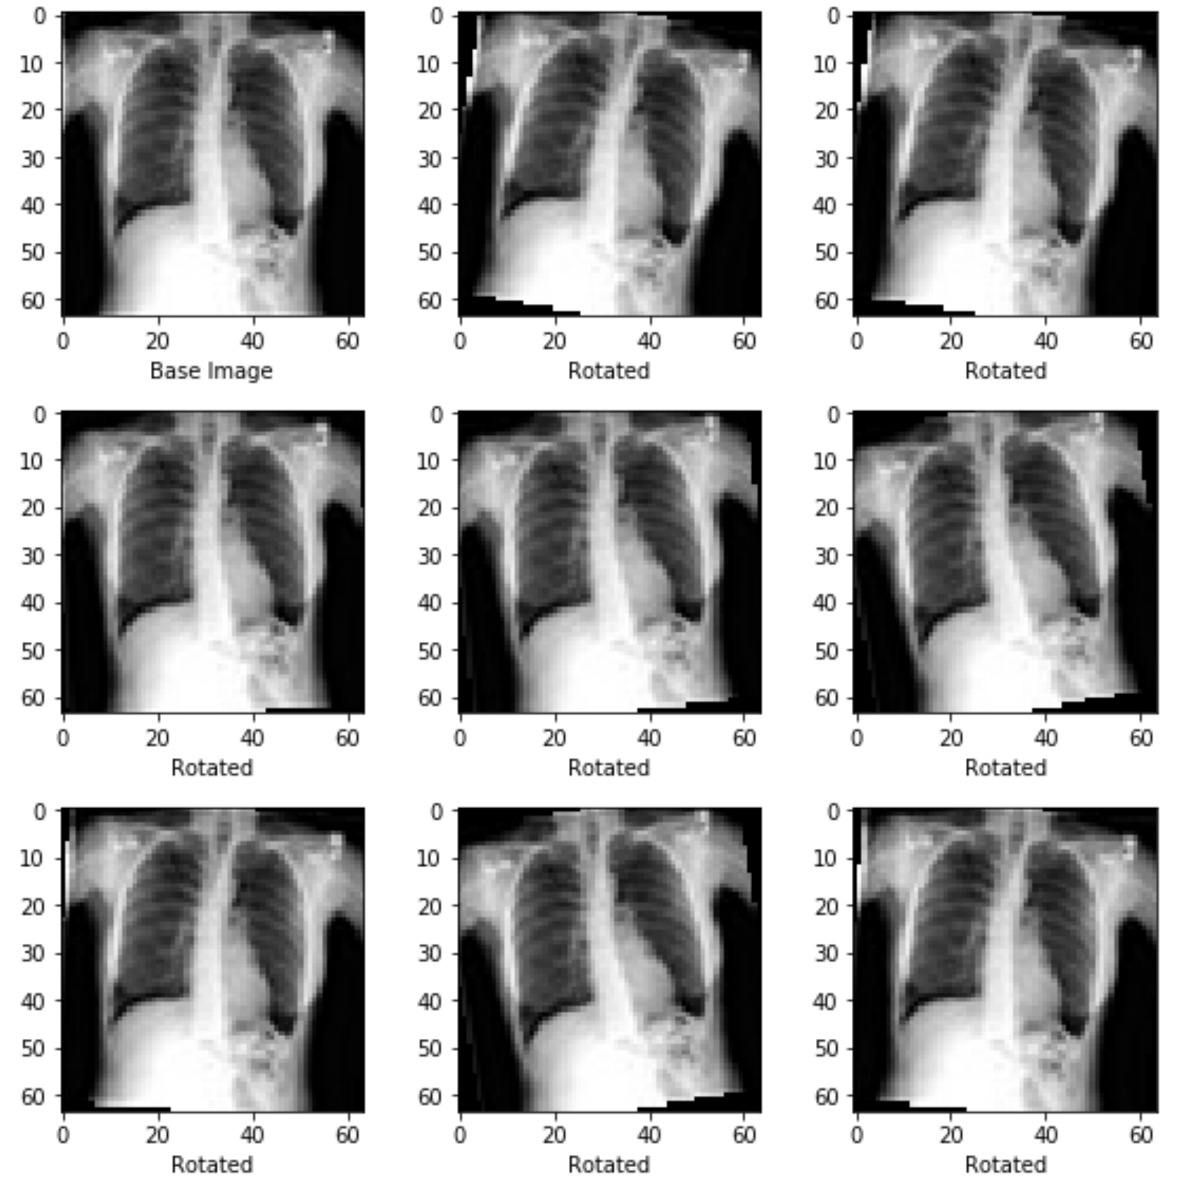
4 Transfer Learning

5 Generative Models

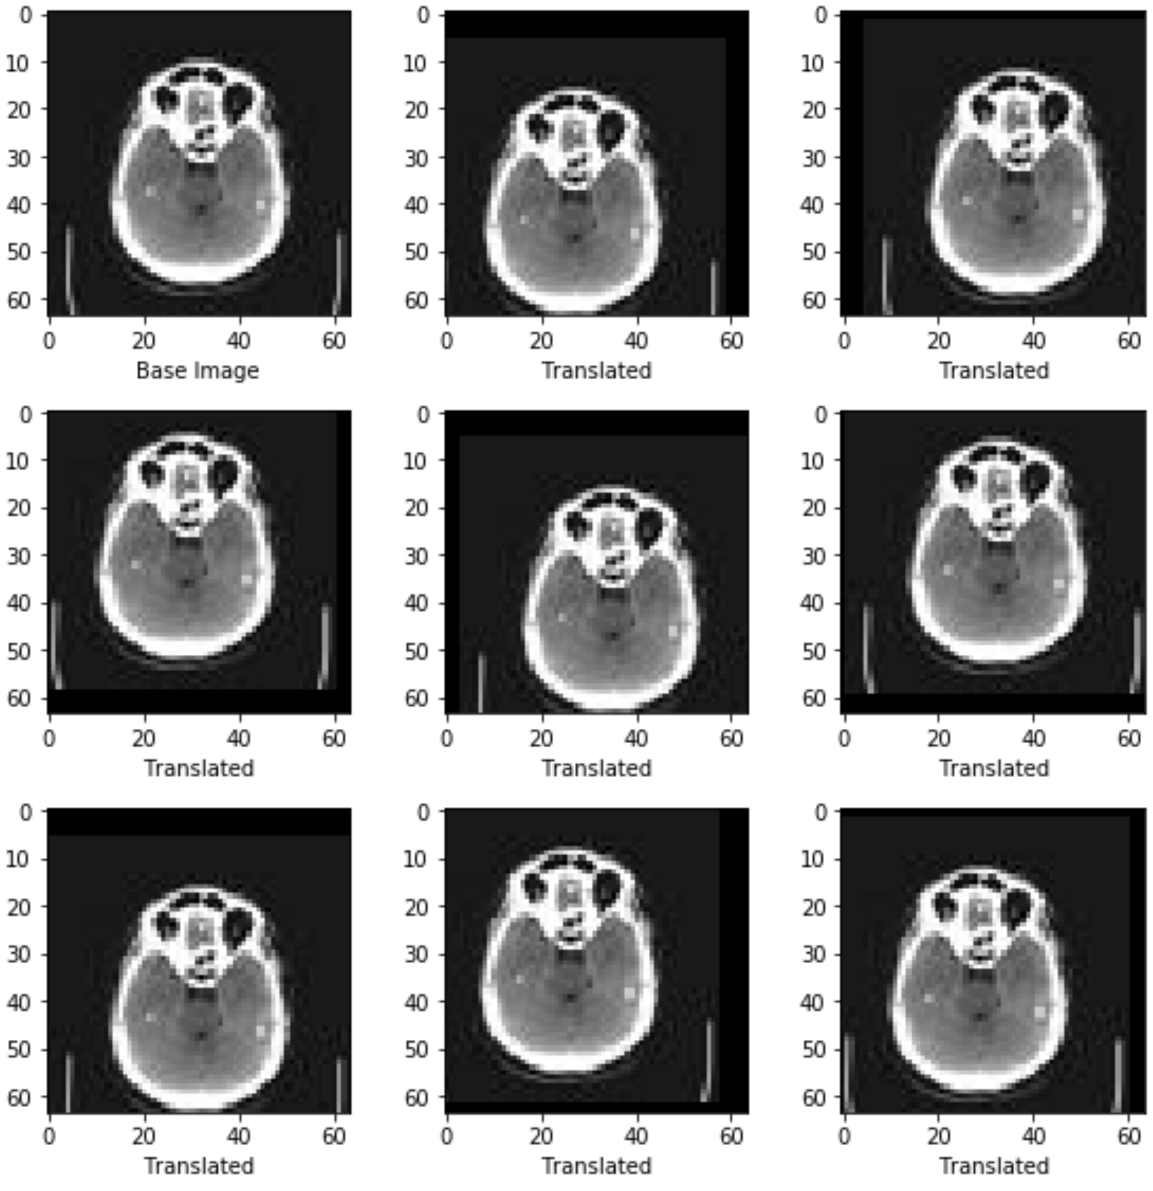
Transfer Learning



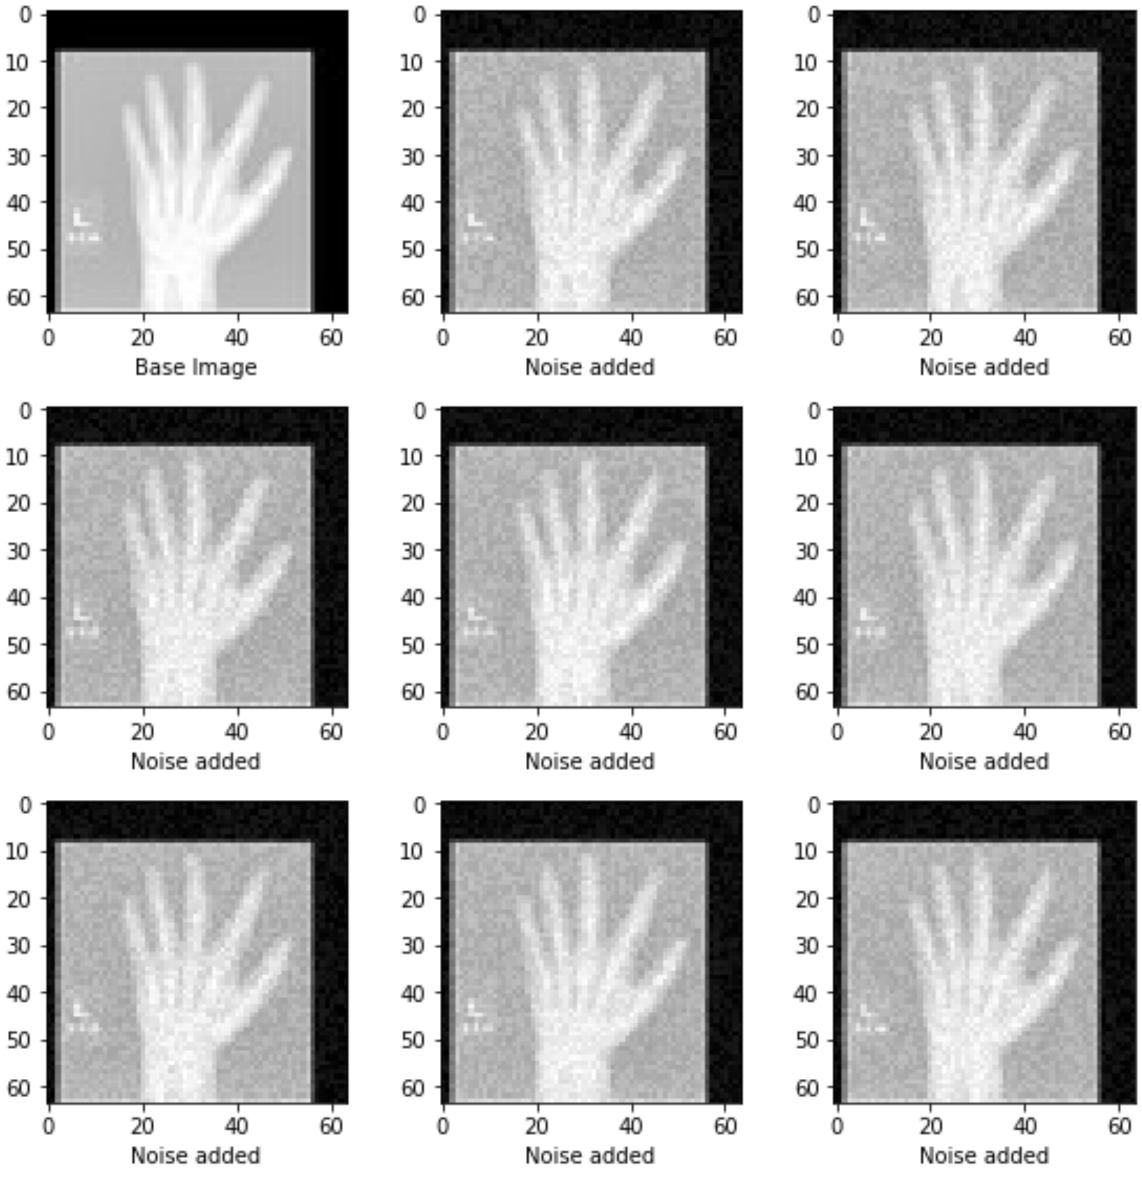
Data Augmentation: Rotations



Data Augmentation: Translation



Data Augmentation: White Noise



Outline

1 Introduction

2 Multilayer Neural Networks

3 Convolutional Neural Networks

4 Transfer Learning

5 Generative Models

Which is a fake picture?





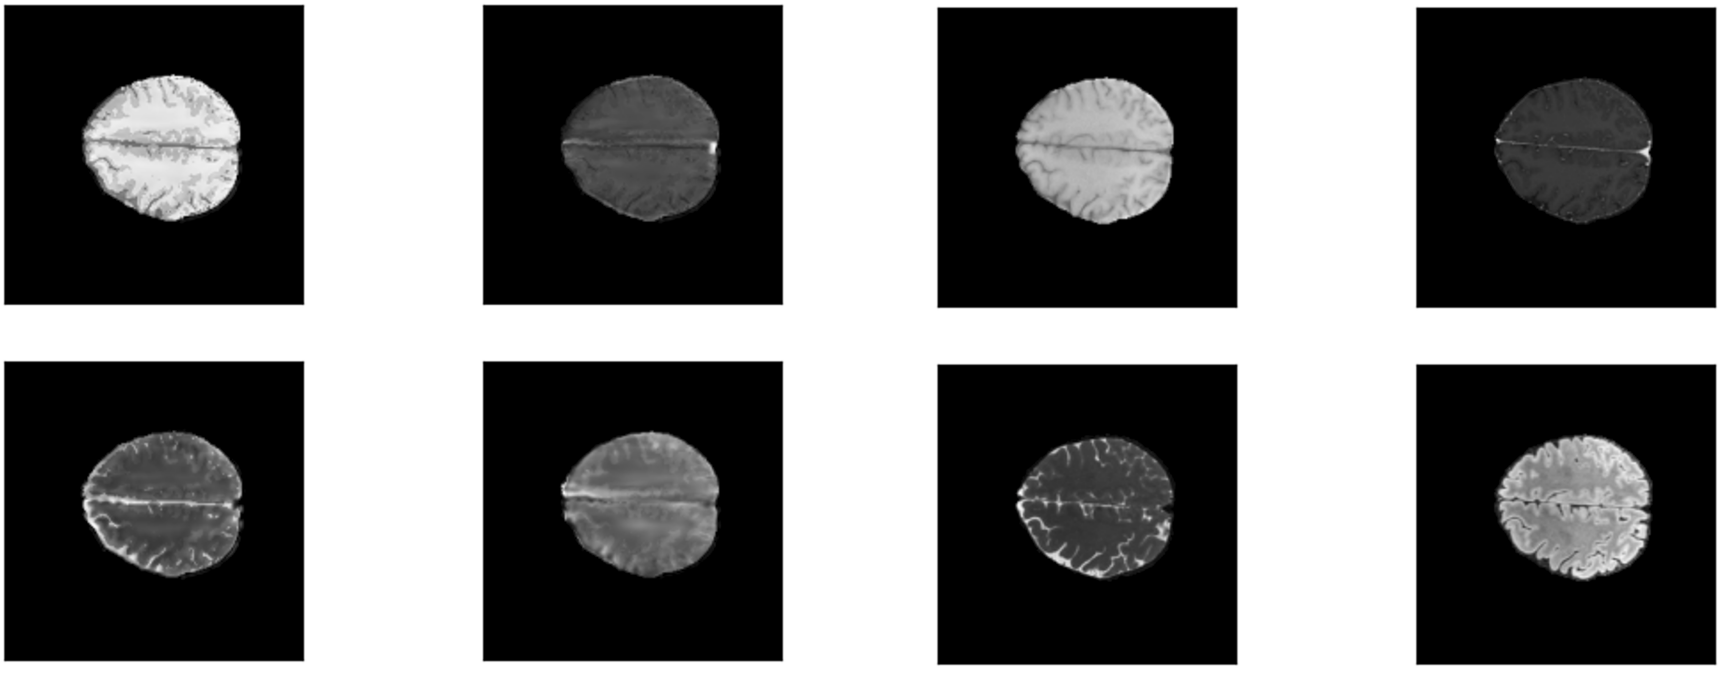
Which is a fake picture?





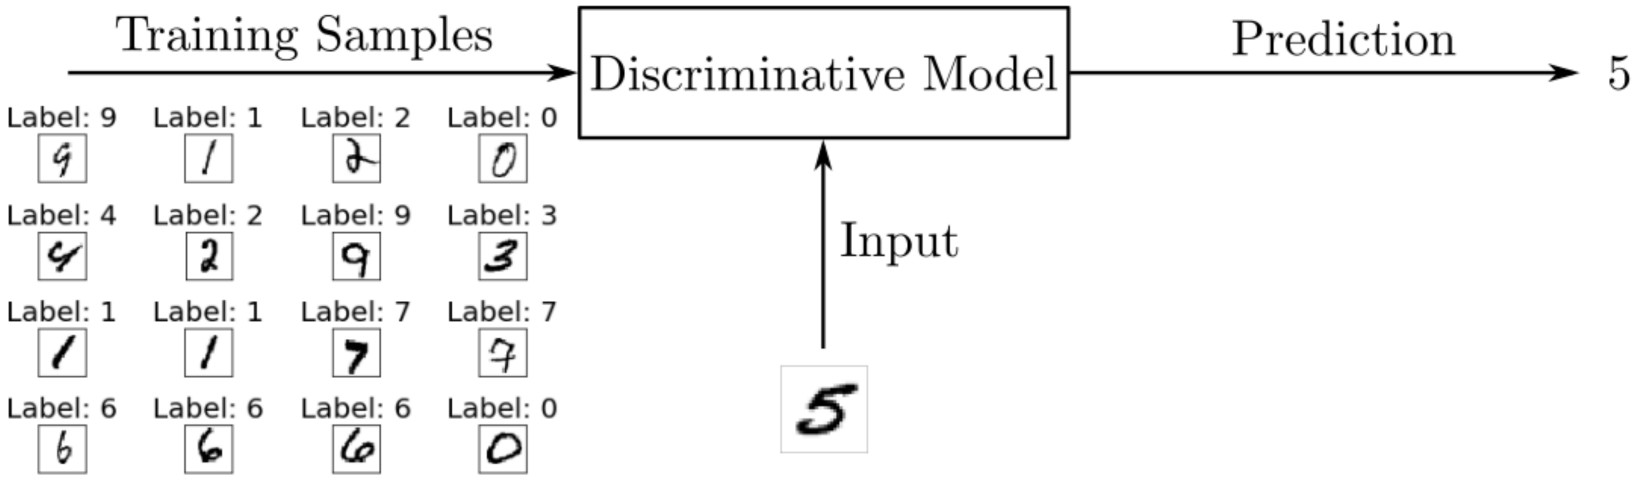
https://thispersondoesnotexist.com/

Which is a fake scan?



Credits to: **Abramo Agosti, Alex Astolfi, Andrea Codegoni, S.G.**,
Math. Dept. at UniPv

Using Deep NNs as Discriminative Models



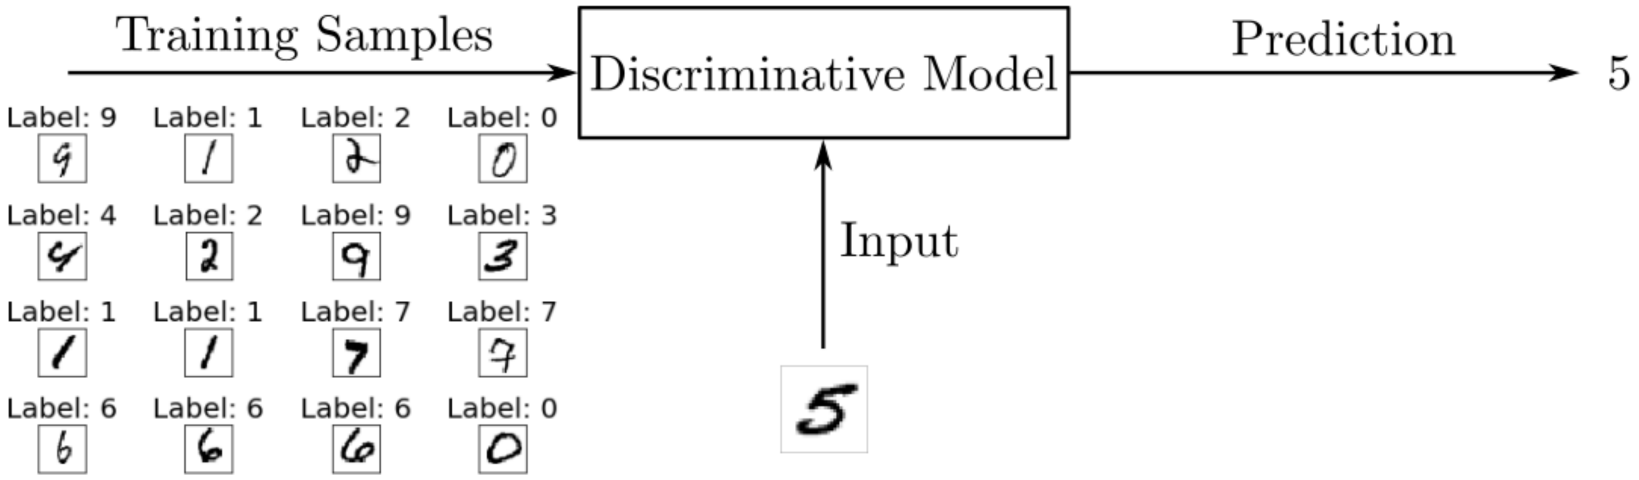
Step 1: **Train-test split of the dataset**.

Step 2: **Choice of architecture**.

Step 3: **Training**.

Step 4: **Testing**.

Using Deep NNs as Discriminative Models



Step 1: **Train-test split of the dataset**.

Step 2: **Choice of architecture**.

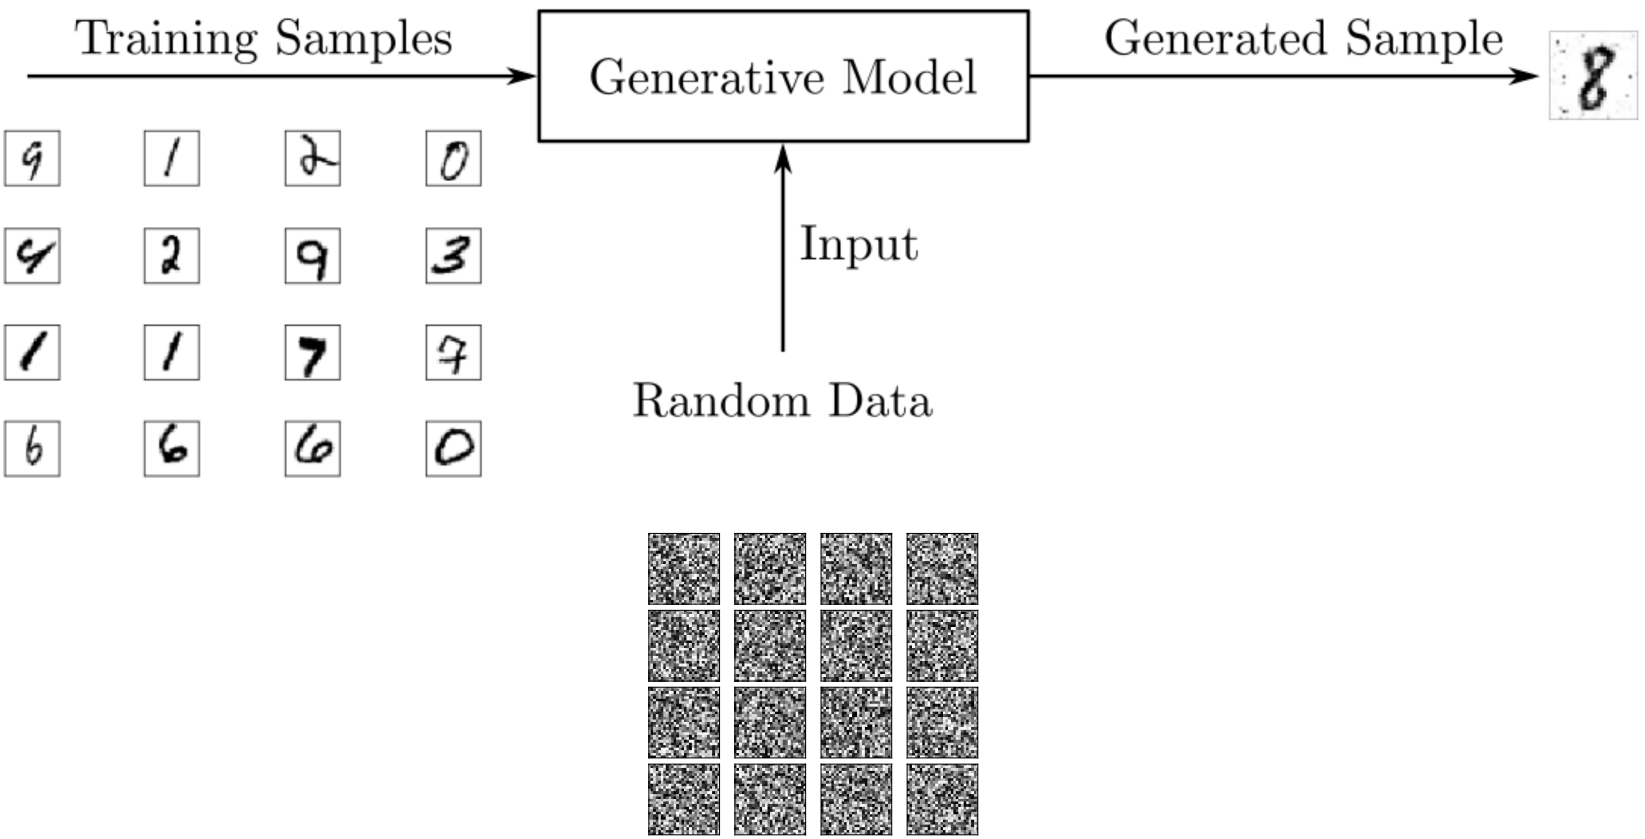
Step 3: **Training**.

Step 4: **Testing**.

This is an example of **Supervised Learning**

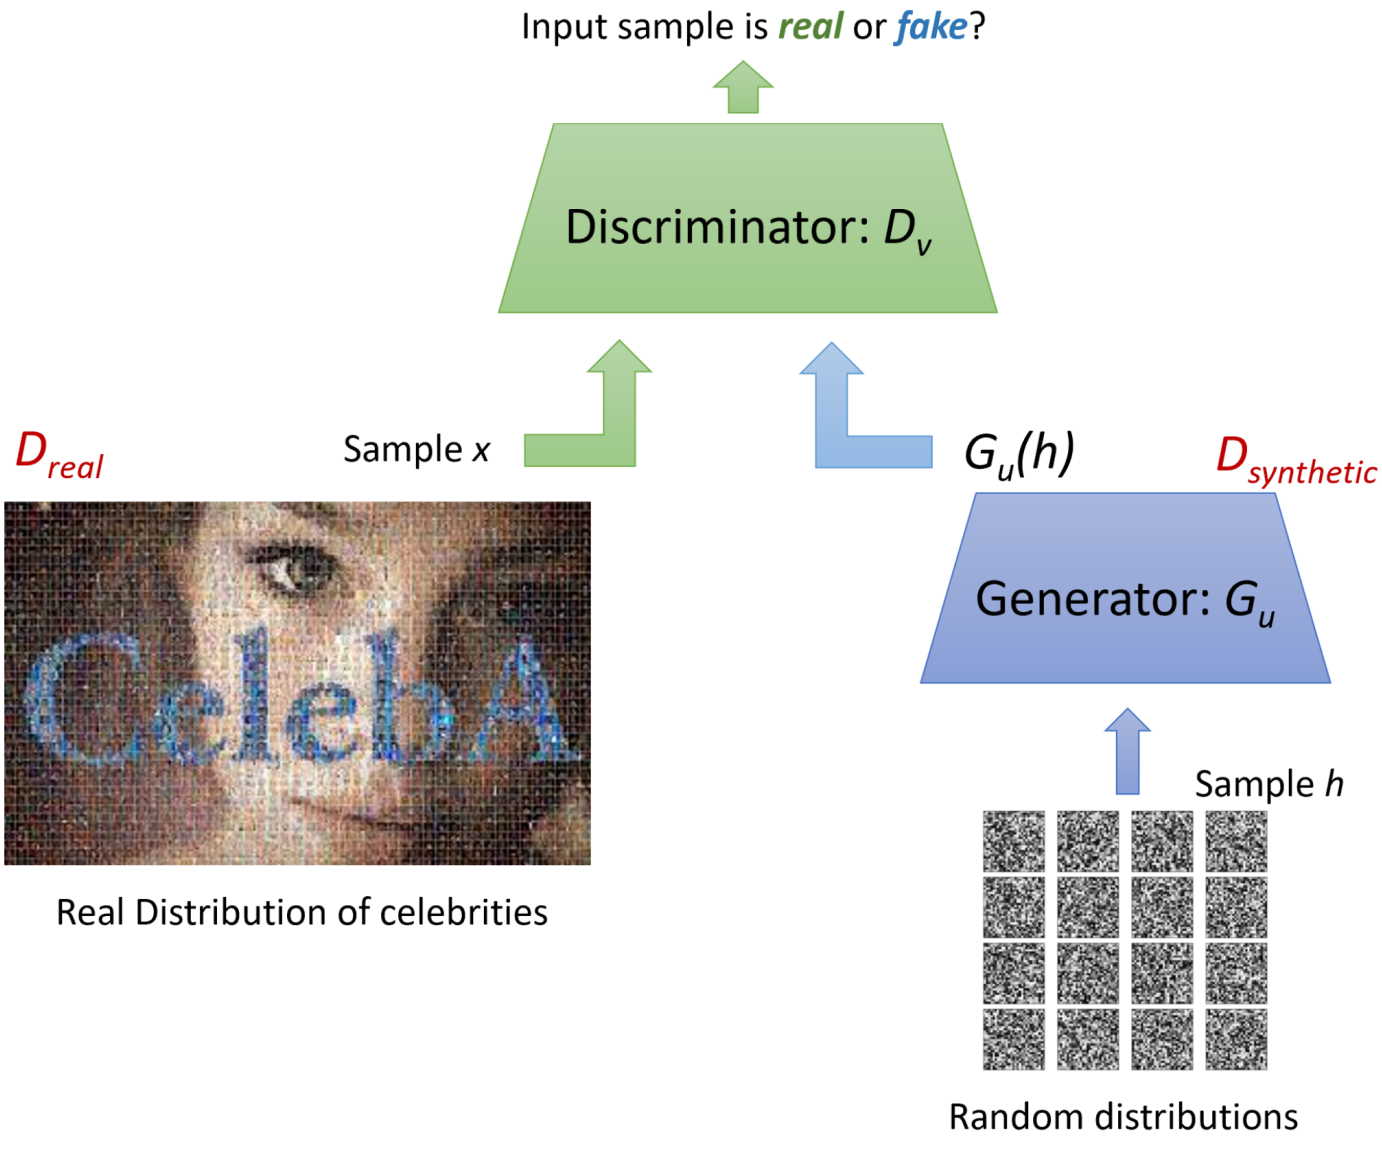
Using Deep NNs as Generative Models

This is an example of **Unsupervised Learning**

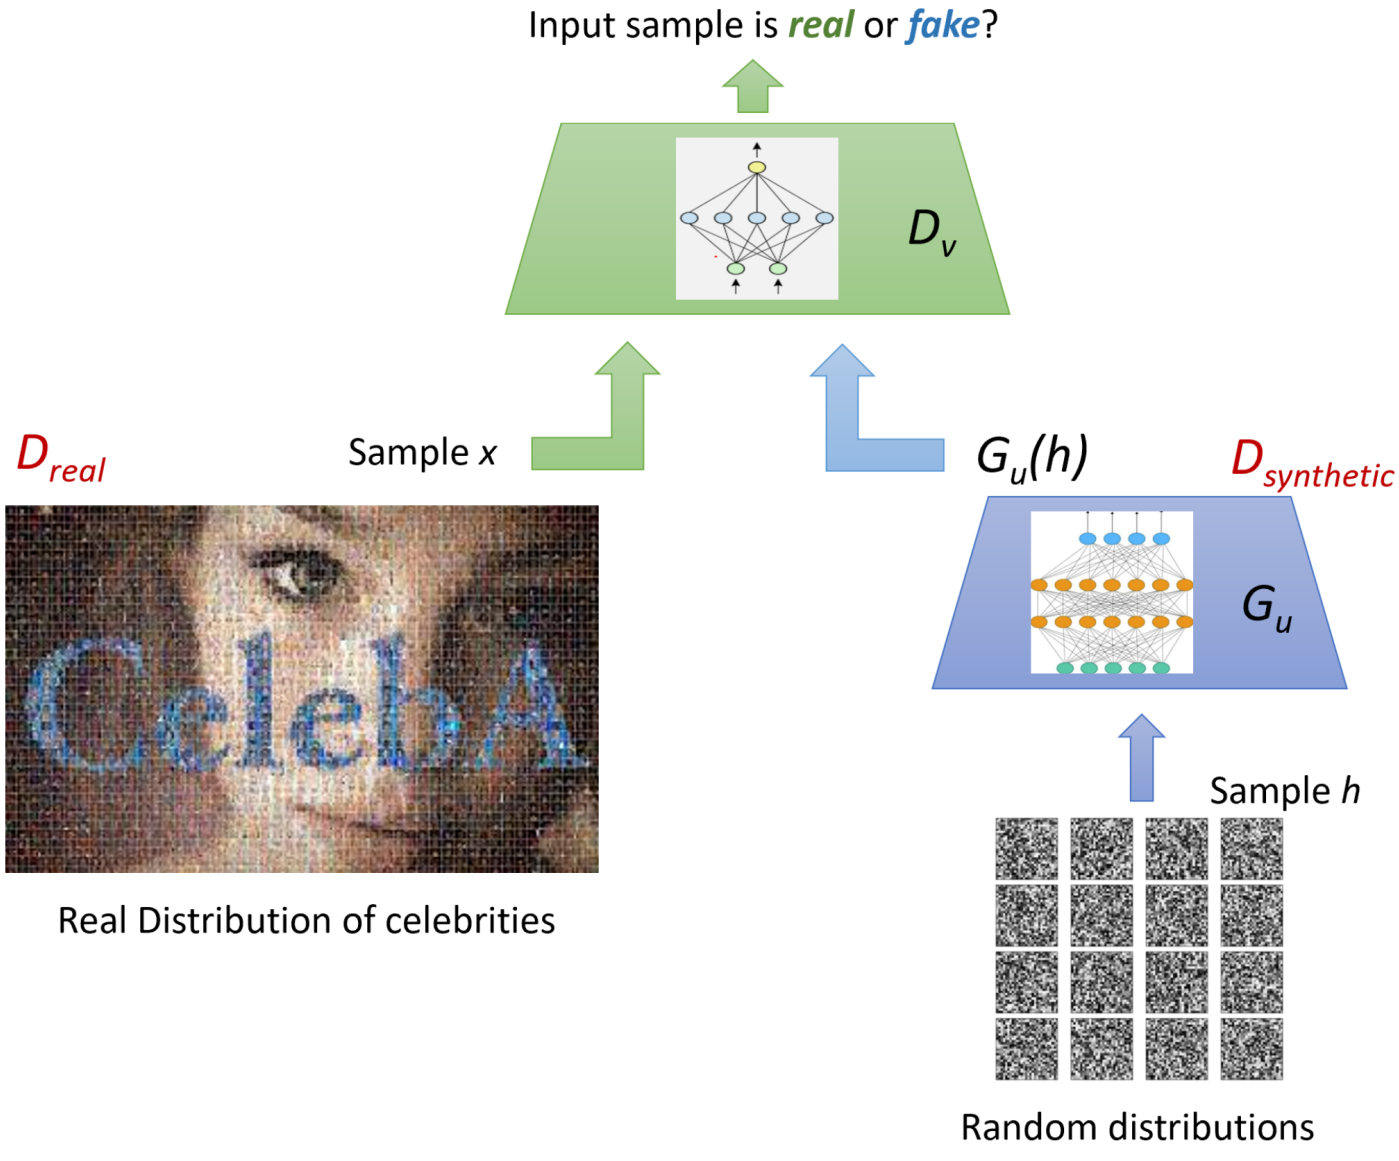


Picture from: `https://realpython.com/generative-adversarial-networks/`

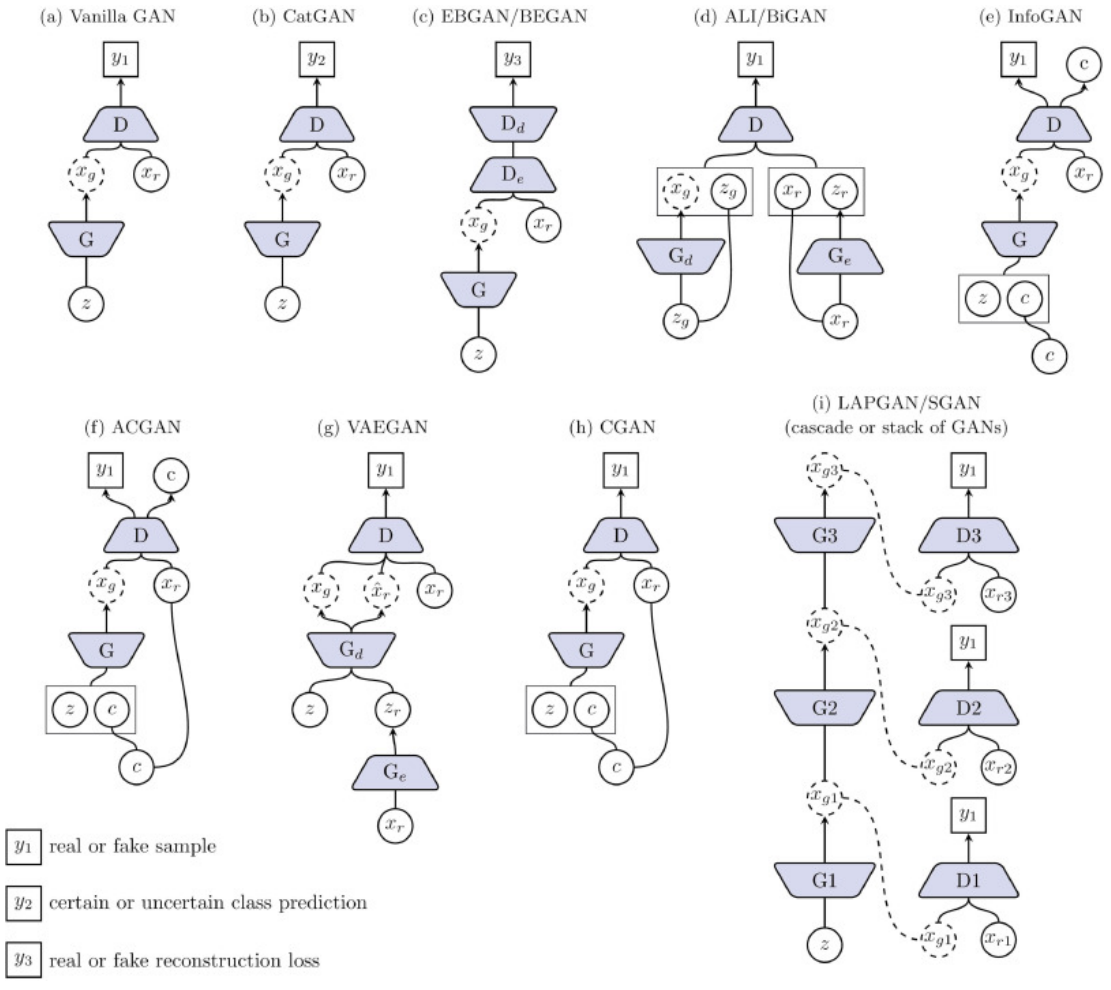
The GAN game


Input sample is real or fake?
Discriminator: D_v
D_{real}
Sample x
$G_u(h)$
$D_{synthetic}$
Generator: G_u
Sample h
Real Distribution of celebrities
Random distributions

The GAN game

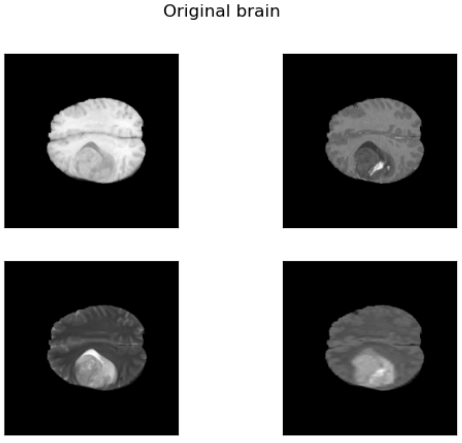

Input sample is real or fake?
D_v
D_{real}
Sample x
$G_u(h)$
$D_{synthetic}$
G_u
Sample h
Real Distribution of celebrities
Random distributions

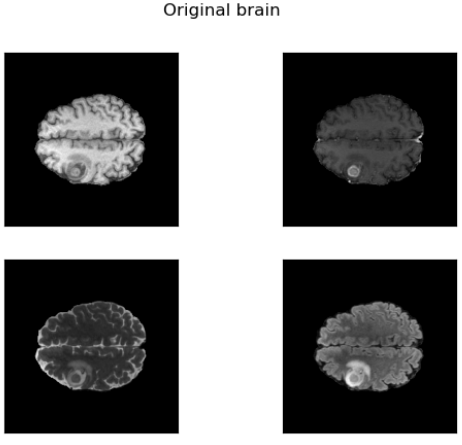
GAN Zoo



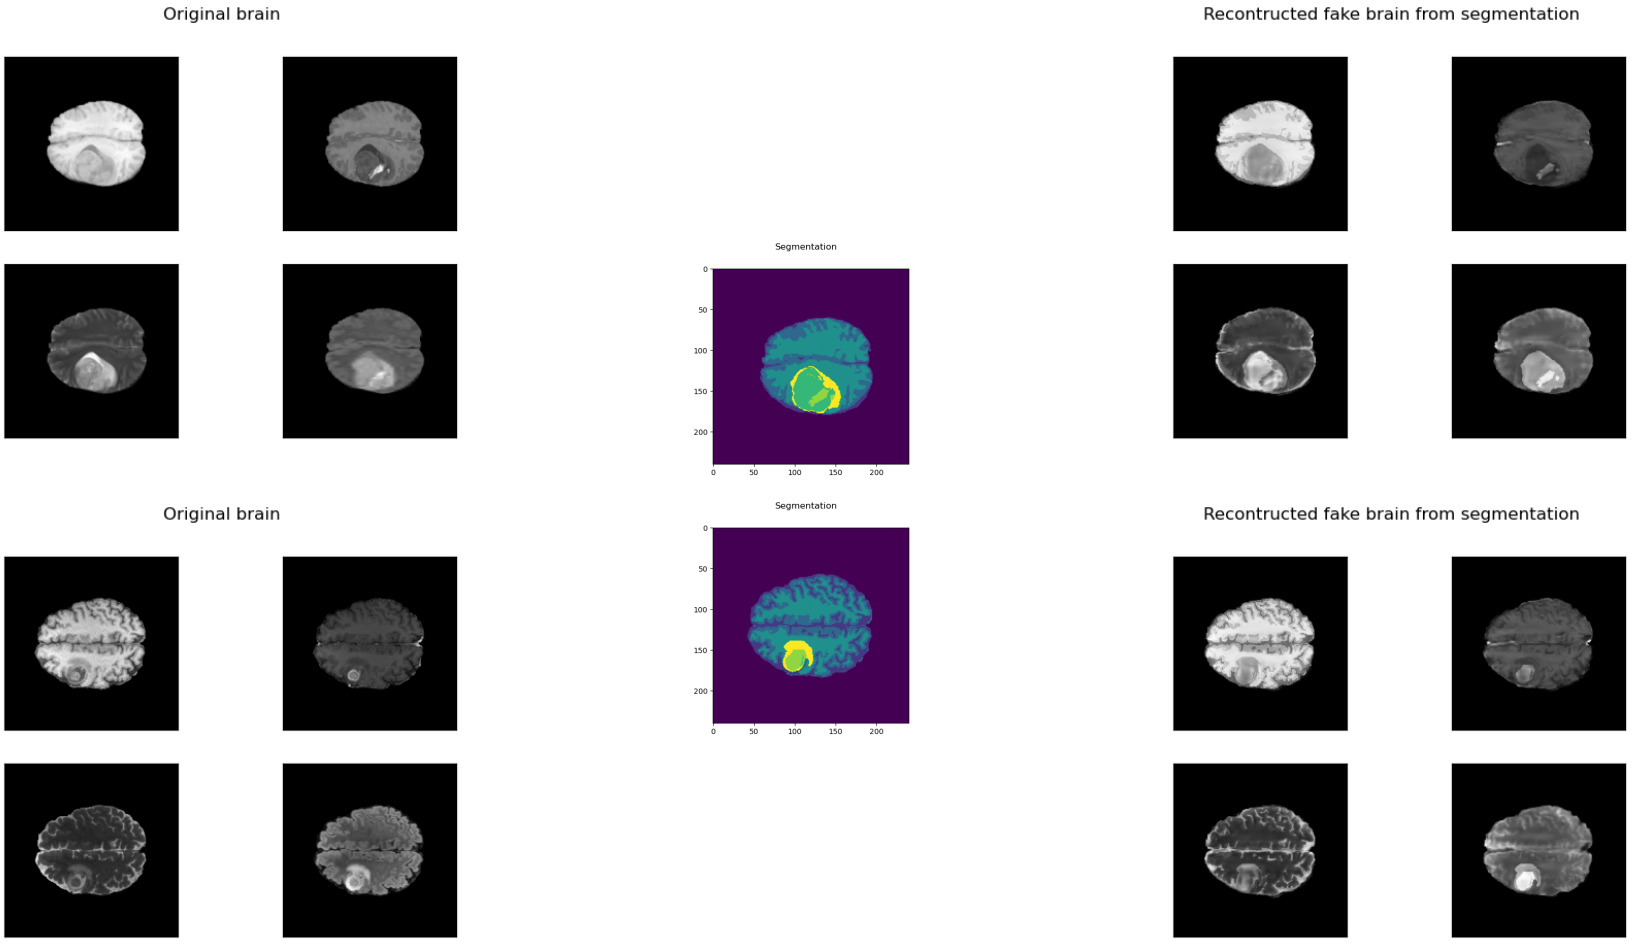
Yi, X., Walia, E. and Babyn, P., 2019. *Generative adversarial network in medical imaging: A review*. Medical image analysis, 58, p.101552.

Ongoing research





Ongoing research



Credits to: **A. Agosti, A. Astolfi, A. Codegoni, S.G.**

QUESTIONS?

Link to videos

- https://mashable.com/2016/12/01/facebook-ai-video-lessons
- https://adamharley.com/nn_vis
- https://www.youtube.com/watch?v=Ilg3gGewQ5U
- https://thispersondoesnotexist.com/
- https://vimeo.com/154085950