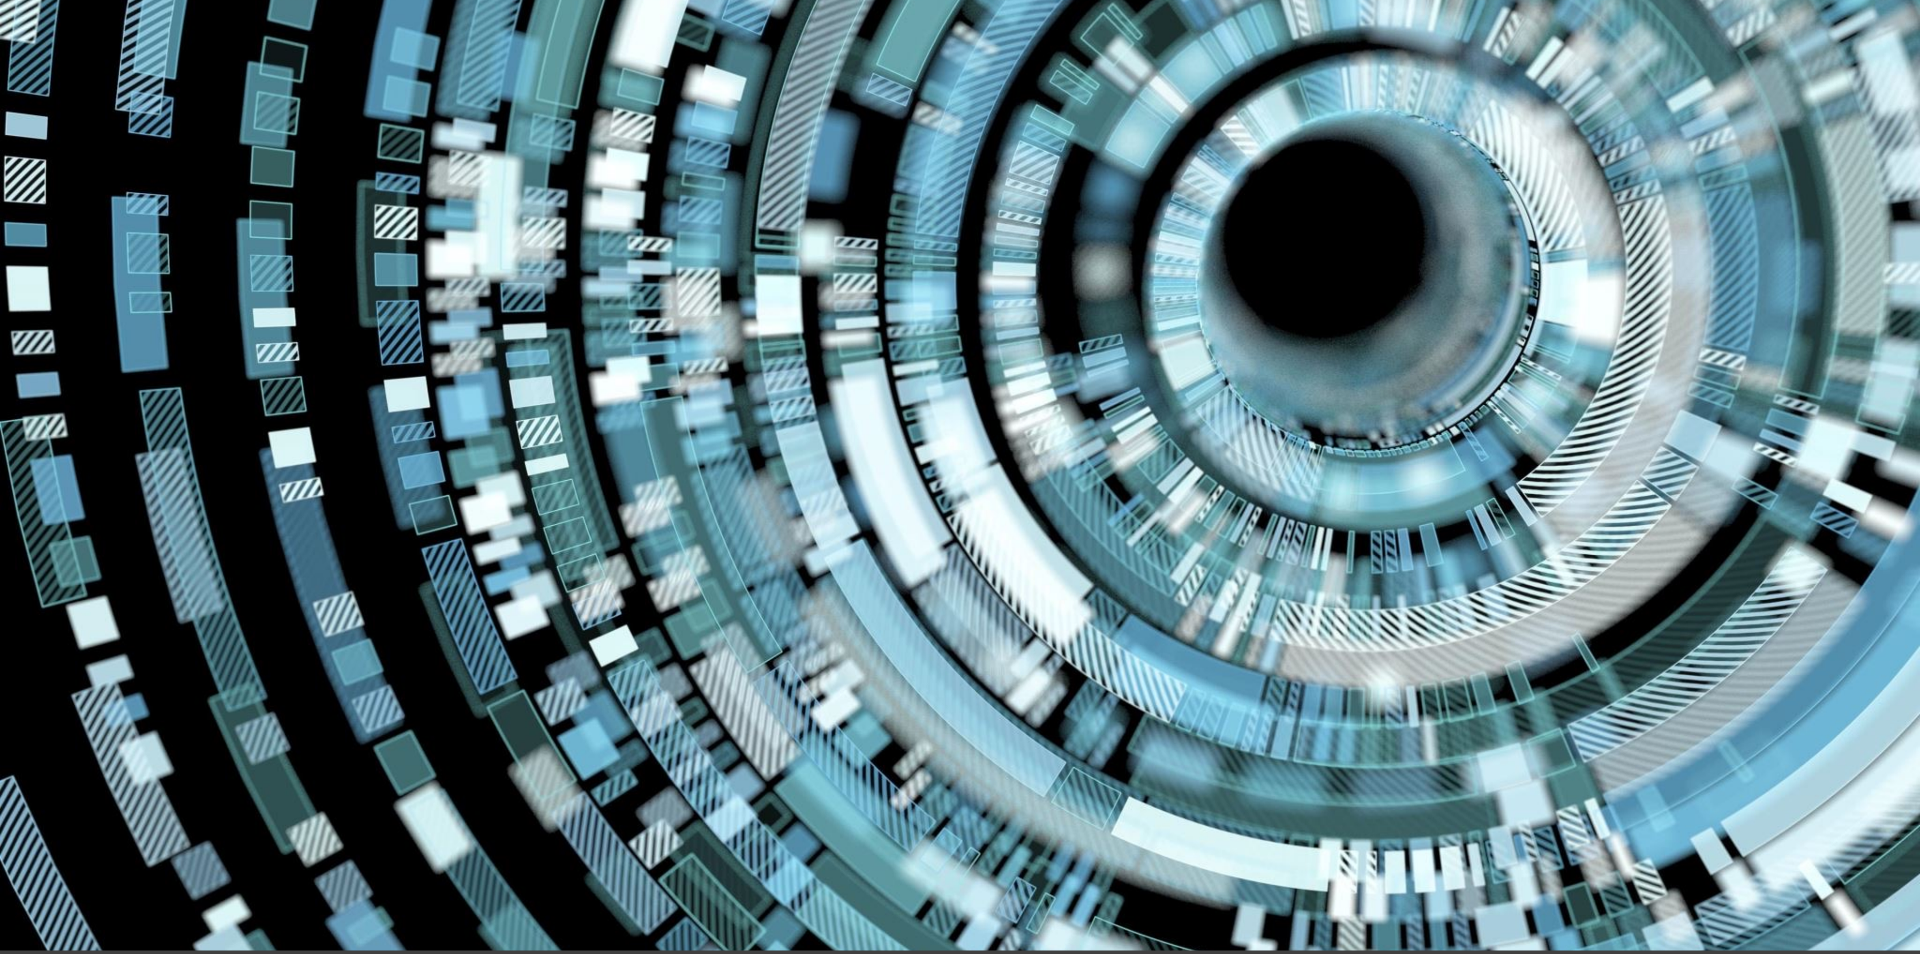

DEEP LEARNING DANS LE DOMAINE MÉDICAL

Master 1 Recherche Biomédicale

E. Reynaud, Labo. EMC, Univ. Lyon 2

Intelligence Artificielle : de l'histoire ancienne…

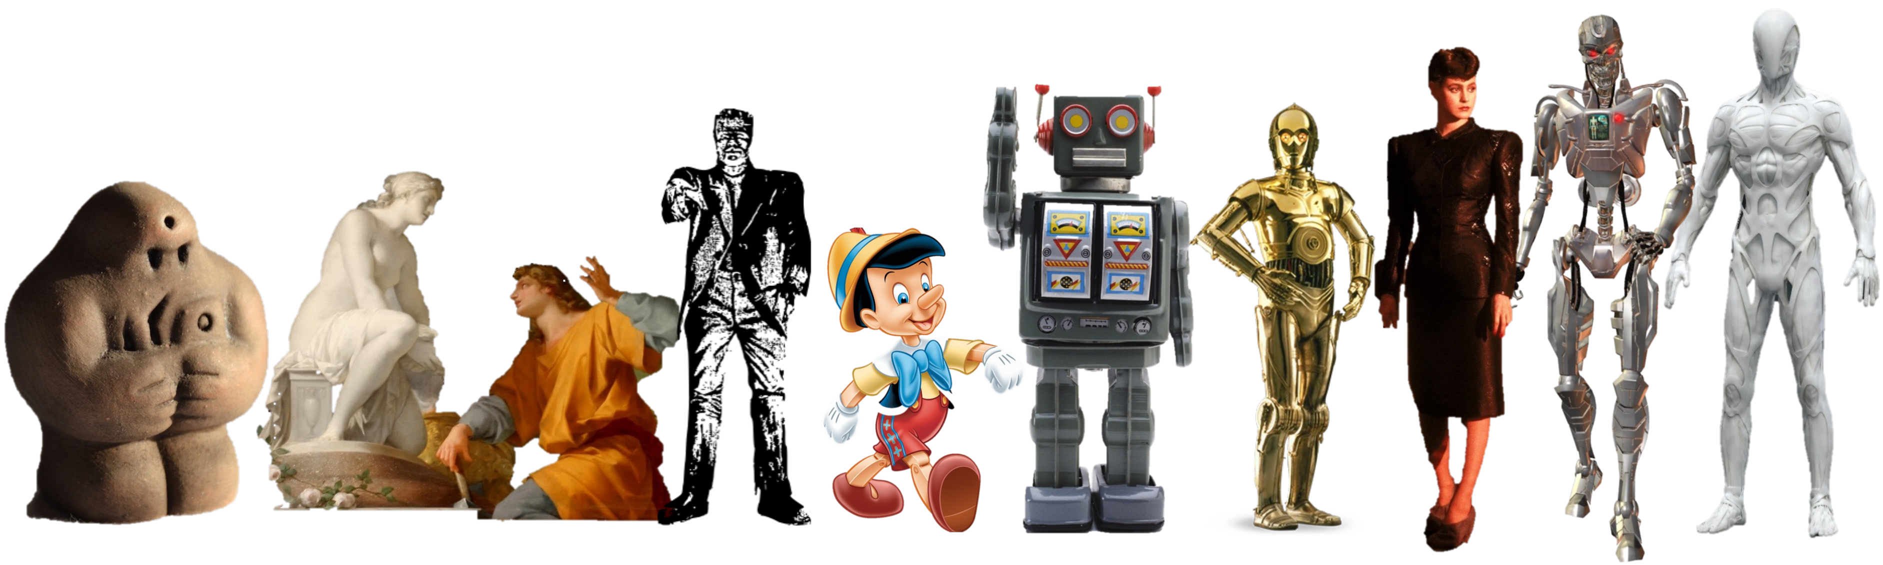

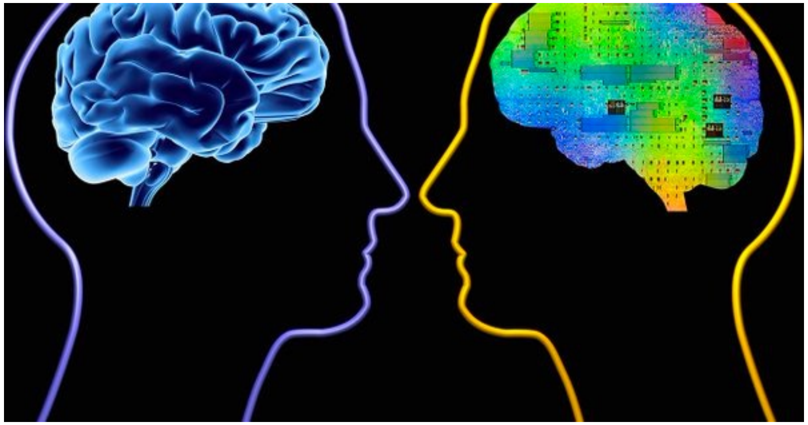

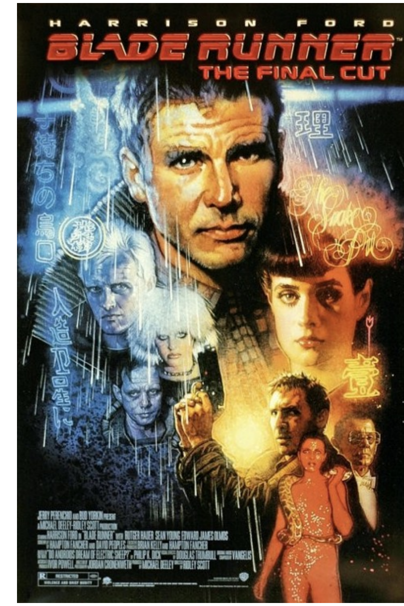
HARRISON FORD
BLADE RUNNER
THE FINAL CUT

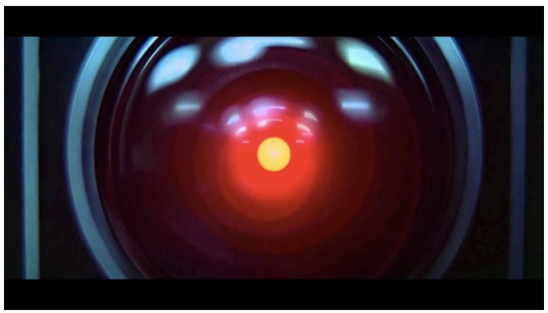

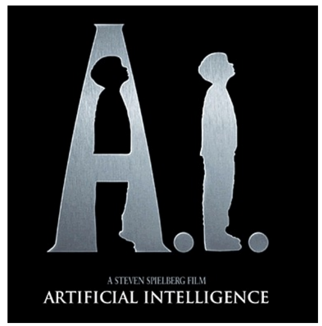
A.I.
A STEVEN SPIELBERG FILM
ARTIFICIAL INTELLIGENCE


EX_MACHINA
LE 3 JUIN

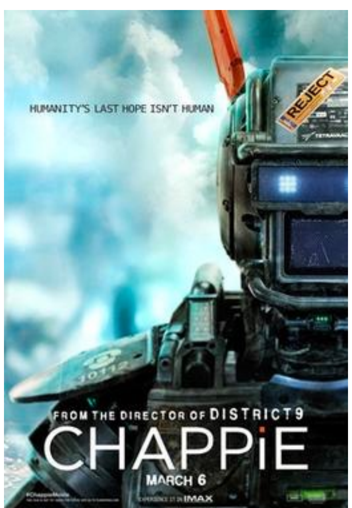
HUMANITY'S LAST HOPE ISN'T HUMAN
REJECT
FROM THE DIRECTOR OF DISTRICT 9
CHAPPiE
MARCH 6

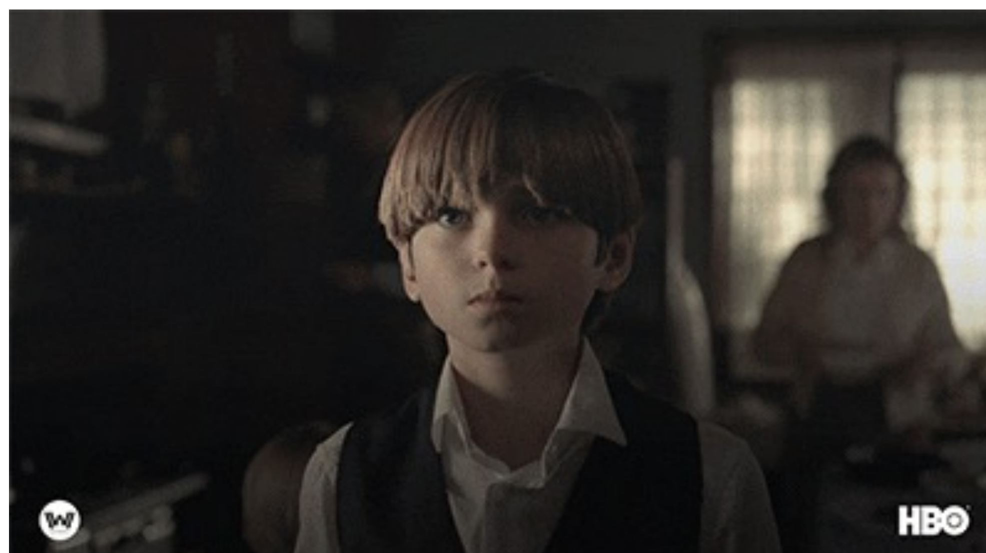

Artificial intelligence

AI wants to build powerful machines
than can think / act intelligently

What is the most intelligent machine ?

Powerful
Portable
Low cost in energy
Adaptative
Can learn
Small
Self-repair
etc

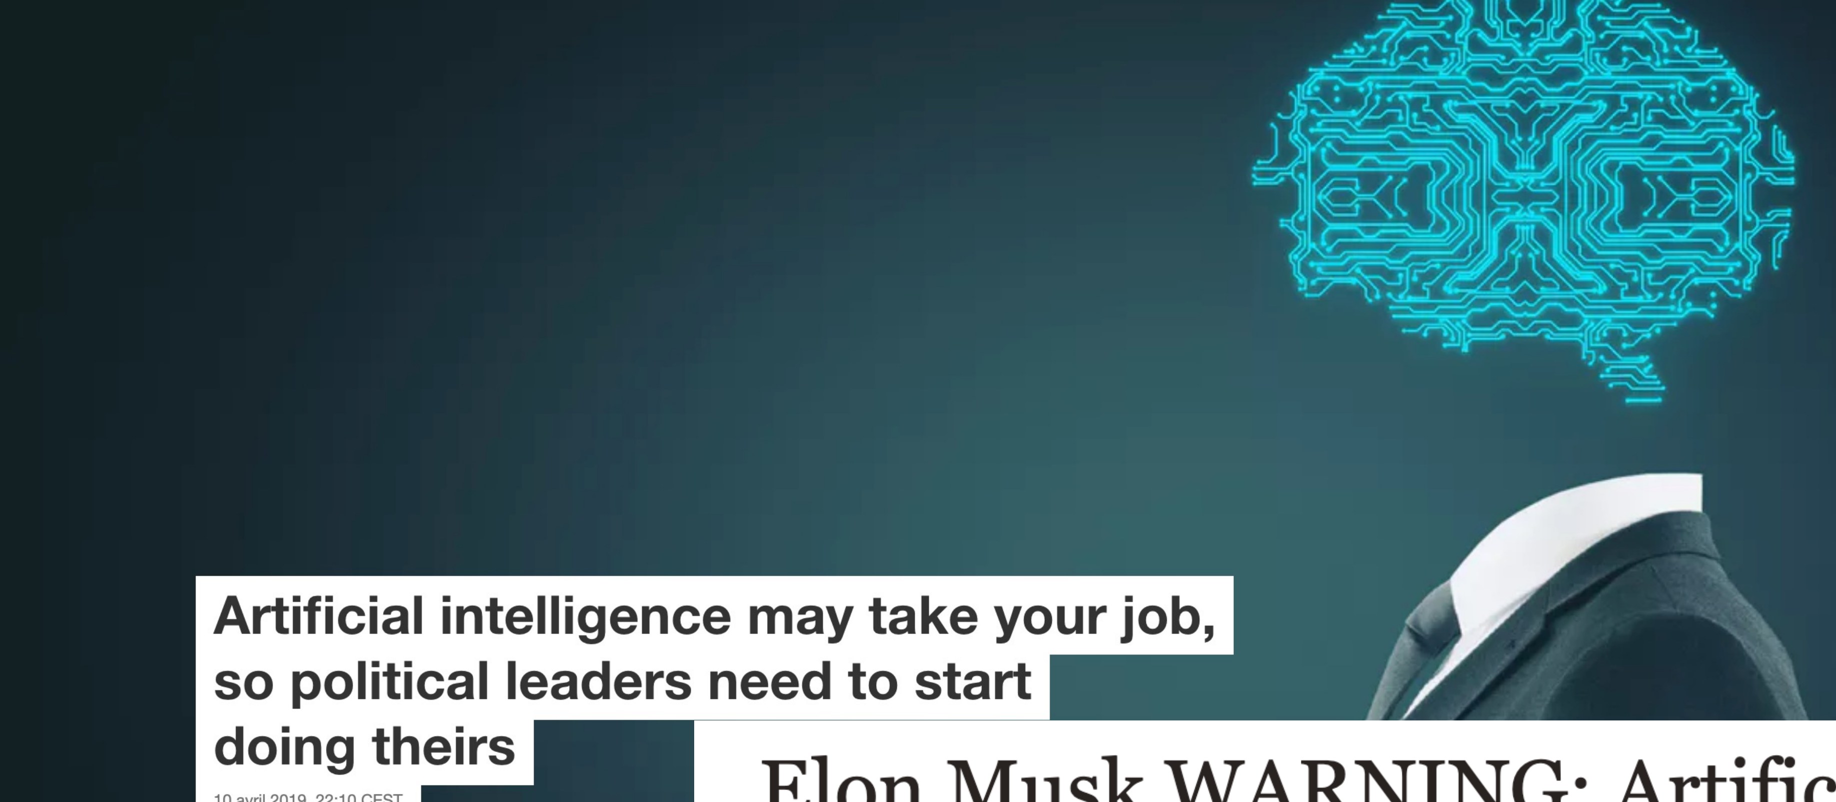

Elon Musk WARNING: Artificial Intelligence could be an 'IMMORTAL DICTATOR'

ial intelligence (AI) could bring with it an
y, that they do not "have to be evil to destroy

Hinton, LeCun, Bengio : la « conspiration » du deep learning

LES REBELLES DE LA SCIENCE 9/10. Au début des années 2000, contre l'opinion de leurs pairs, trois chercheurs en intelligence artificielle ont remis au goût du jour une voie jugée sans avenir : les réseaux de neurones.

Artificial Intelligence

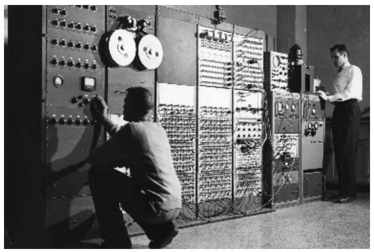

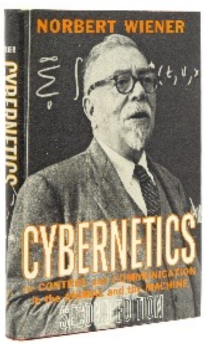

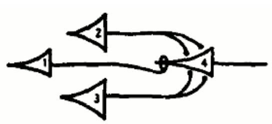

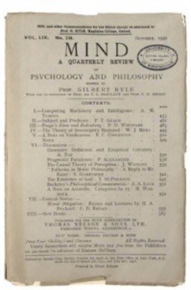

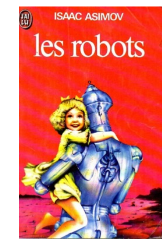

"Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it. » (Dartmouth Proposal, 1956)

Artificial Neural Networks aka ANN

- McCulloch & Pitts, 1943 : MCP neuron model

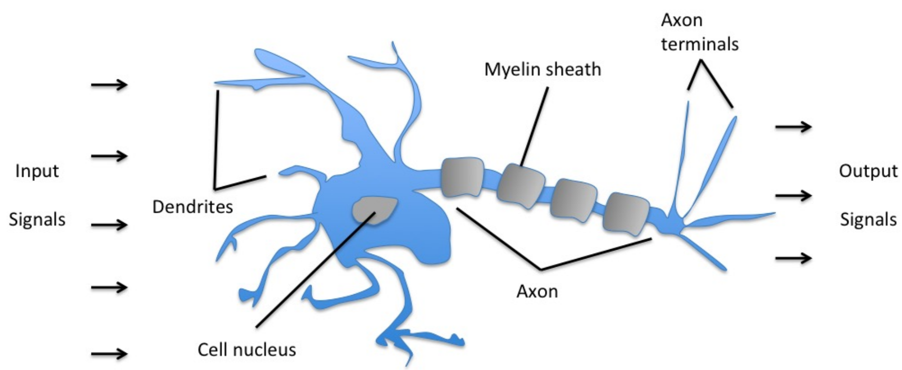

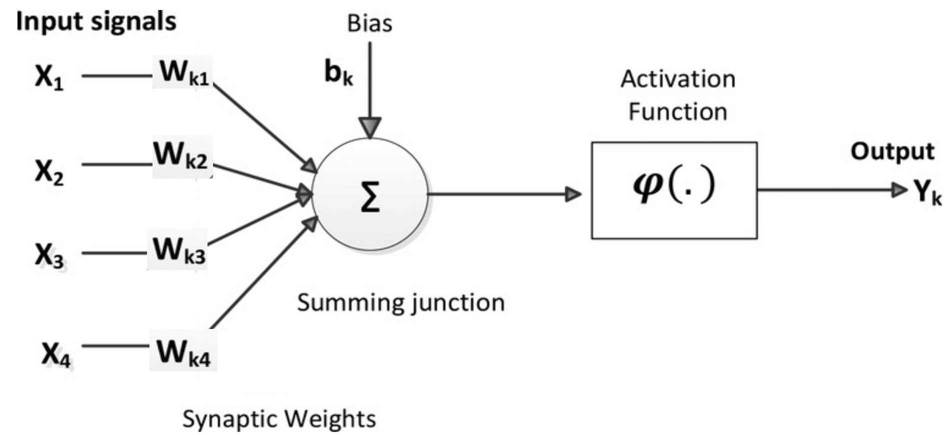

Dendrites Soma Axon

MCP neuron k :

Activation $A = x_1 . w_{k1} + x_2 . w_{k2} + x_3 . w_{k3} + x_4 . W_{k4} - b_k$

Output $Y_k = \varphi(A)$

Transfer function : $\varphi(x) = \begin{cases} 1 \text{ if } x \geq 0 \\ 0 \text{ if } x < 0 \end{cases}$

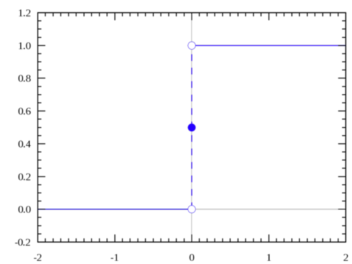

Artificial Neural Networks aka ANN

For neuron k with n pre-synaptic inputs :

$$y_k = \varphi\left(\sum_{i=1}^{n} x_i w_{ki} - \theta_k\right)$$

In code :

```
for i = 1 to n :
        sum = sum + x_i * w_ki
sum = sum - θ
out = activation_function(sum)
```

Activation function

Name	Plot	Equation
Identity	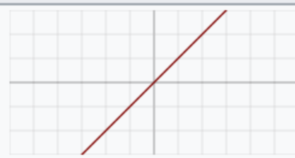	$f(x) = x$
Binary step	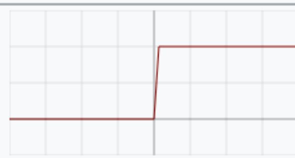	$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$
Logistic (a.k.a. Sigmoid or Soft step)	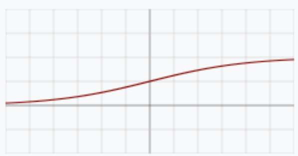	$f(x) = \sigma(x) = \frac{1}{1+e^{-x}}$ [1]
TanH	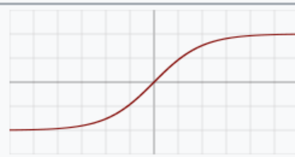	$f(x) = \tanh(x) = \frac{(e^x - e^{-x})}{(e^x + e^{-x})}$
ArcTan	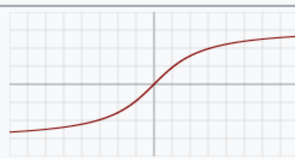	$f(x) = \tan^{-1}(x)$
Rectified linear unit (ReLU)[15]	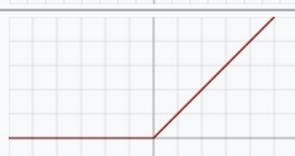	$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}$

Artificial Neural Networks aka ANN

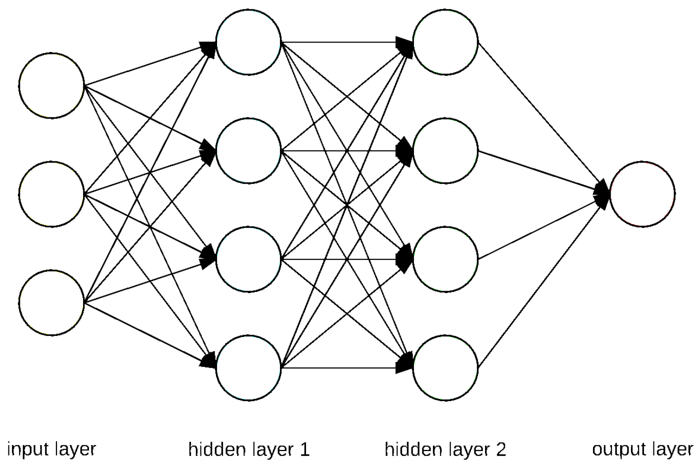

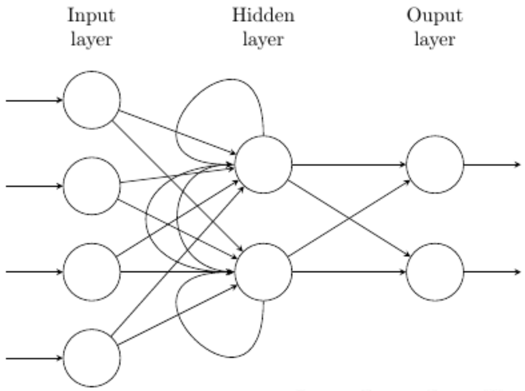

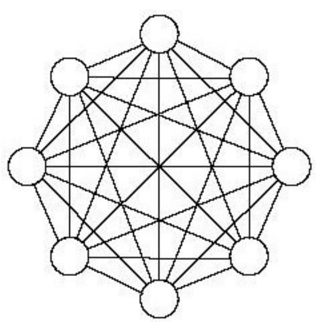

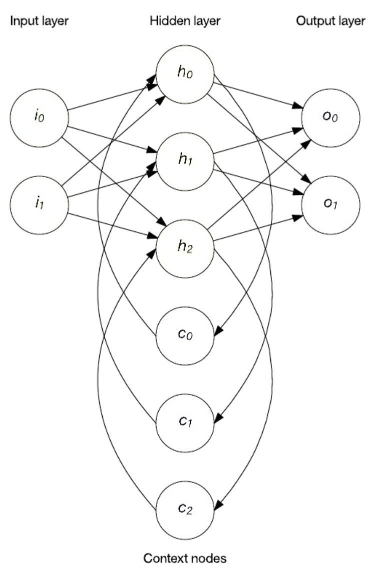

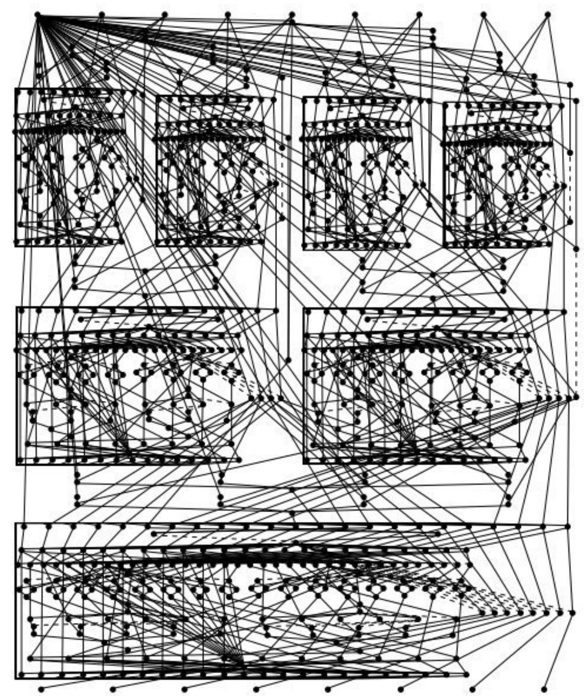

Artificial Neural Networks aka ANN

- Simple units
- Parallel Distributed Processing
- Units in a network : « knowledge » can emerge
- « Knowledge » lies in the synaptic weights

- ANN need to be trained in most cases : supervised learning
- Training = weights adjustments
- Error = difference between the network outputs and the desired outputs for a set of stimuli

The father of all ANN : The Perceptron

1957, Rosenblatt

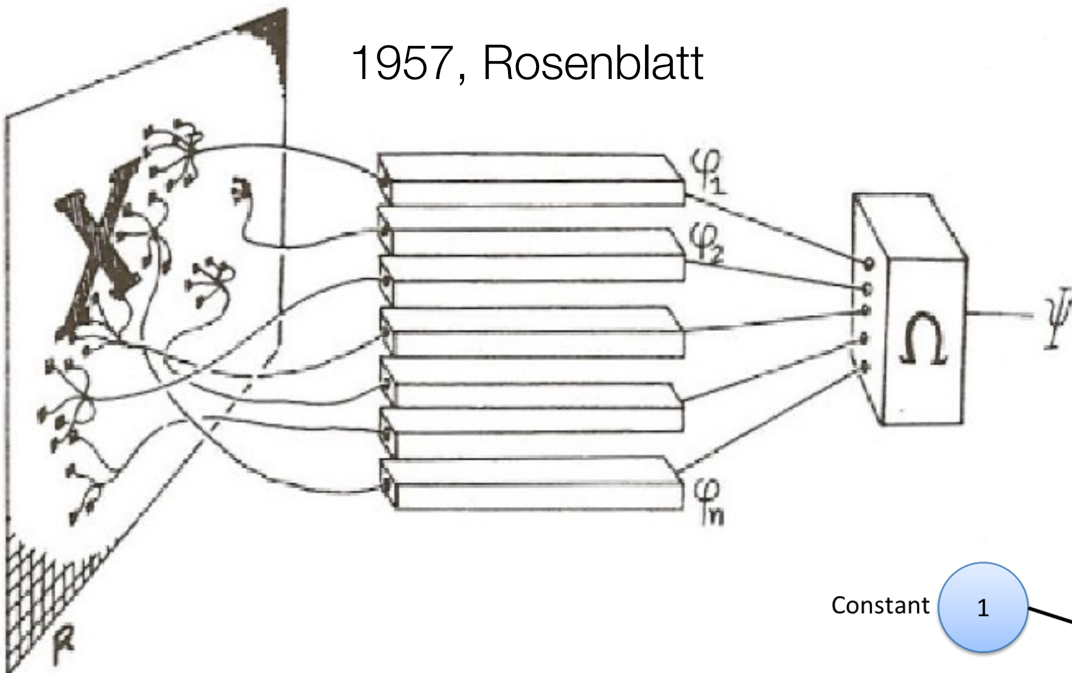

Function to learn :

Letter	d :
X	1
T	0

Learning rule :

$$w_j(t+1) = w_j(t) + \eta\,(d - y)\,x$$

with :

- d : desired output
- y : network output
- η : learning rate

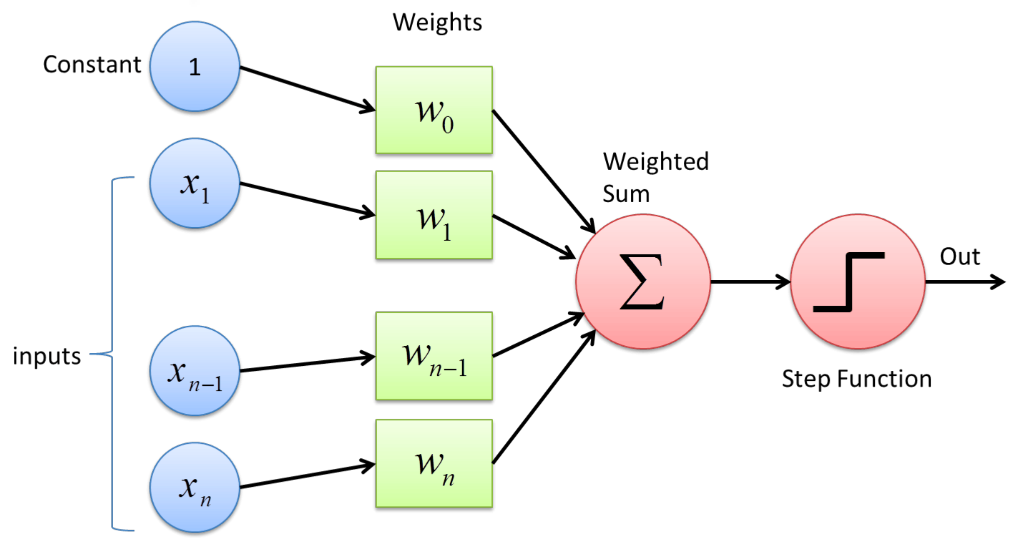

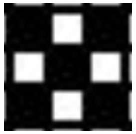

X1 = [1 0 1 0 1 0 1 0 1]

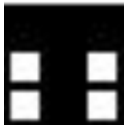

X2 = [1 1 1 0 1 0 0 1 0]

$W_{t=0}$ = [0 0 0 0 0 0 0 0 0]

$\eta = 0.5$

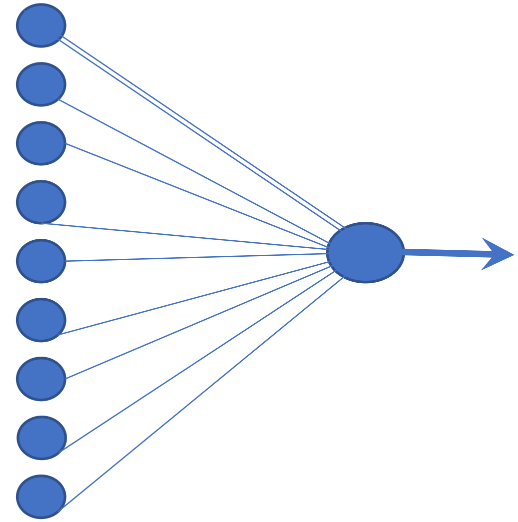

```
X1 = [1 0 1 0 1 0 1 0 1 ];
X2 = [ 1 1 1 0 1 0 0 1 0 ];

W=[0 0 0 0 0 0 0 0 0];
%W = randn(9,1) ;
disp(W);

eta = 0.5;
d1 = 1;
d2 = 0;

i=1;
while i < 10
    iter = sprintf('EPOQUE n. %d',i);
    disp("_______________________________
    disp(iter)
    A=sum(X1.*W);
    if A>=0 y = 1 ;
    else y = 0;
    end
    fprintf('Presentation X1 lettre X.,
    W=W + (eta * (d1-y))*X1;
    disp((W'));

    A=sum(X2.*W);
    if A>=0 y = 1 ;
    else y = 0;
    end
    fprintf('Presnetation X2 lettre T., output reseau: %d\n',y);
    W=W + (eta * (d2-y))*X2;
    disp(W');
i=i+1;
end
```

	d	Σ	y	d-y	eta*(d-y)	nouveaux poids
X	1	0	1	0	0	[0 0 0 0 0 0 0 0 0]
T	0	0	1	-1	-0.5	[-0.5 -0.5 -0.5 0 -0.5 0 0 -0.5 0]
X	1	-1.5	0	1	0.5	[0 -0.5 0 0 0 0 0.5 -0.5 0.5]
T	0	-1	0	0	0	[0 -0.5 0 0 0 0 0.5 -0.5 0.5]
X	1	1	1	0	0	[0 -0.5 0 0 0 0 0.5 -0.5 0.5]

The Perceptron complexified : MLPs

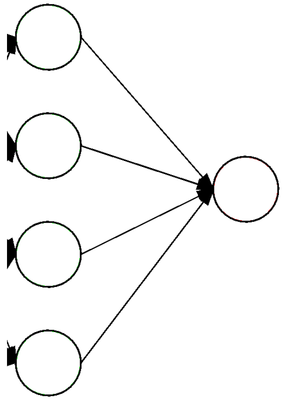

MLP = perceptron

+ hidden layers to complexify decision regions
+ activation function : sigmoïd / tanh
+ backpropagation ☺

Backpropagation simplified

1 – From Layer 1 to Output layer : input is propagated
2 – Network output is computed
3 – Error is computed (from desired -output)
4 – Weights are updated from Output Layer to previous layers,
Error is distributed backwards according to which weights contributed the most to this error, then to the error in the next layers etc until Layer 1

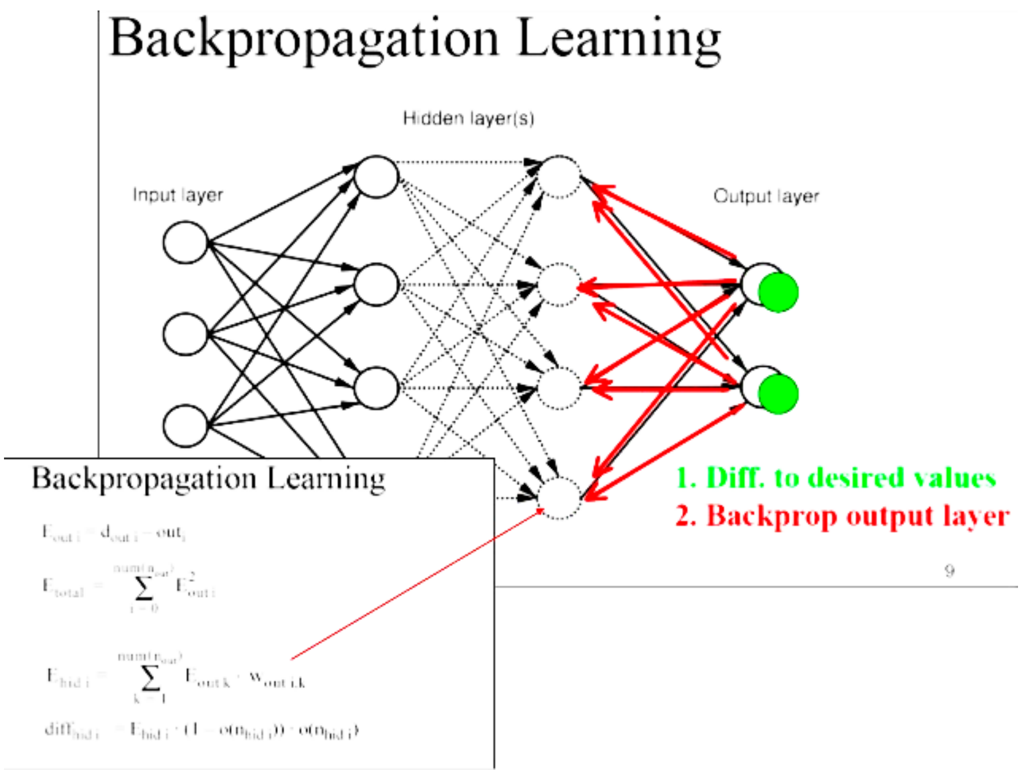

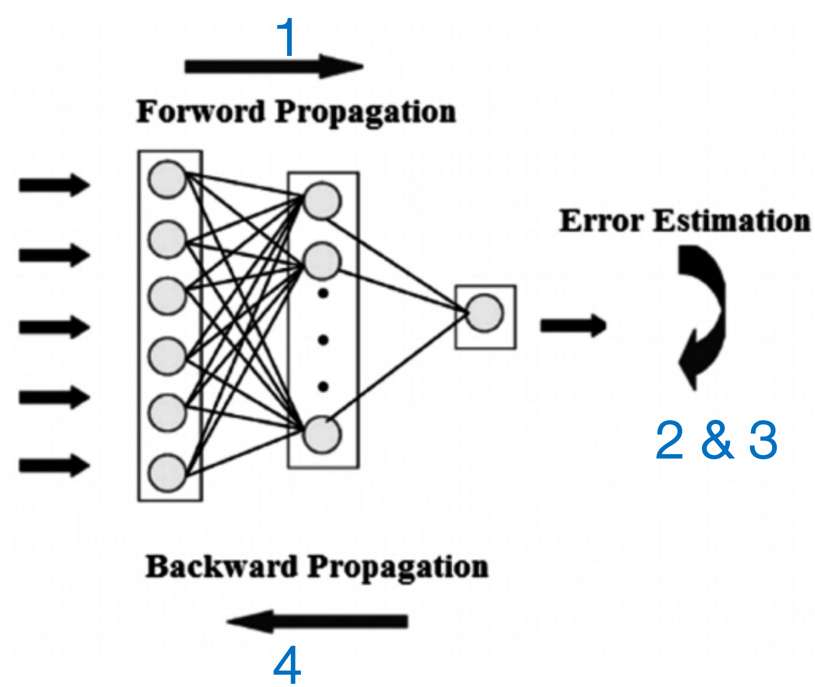

Deep Learning

- « Brain » / « Cognition » inspired
- MLP with many layers + modified architecture
- Hierarchical representation of information through the layers (see visual cortex for example)
- Modified neurons (eg. : ReLU units)
- Less weights, more learning !
- Krizhevsky et al. (2012): won the ImageNet challenge with AlexNet, similar to LeNet (LeCun 1998)

Computer and images

Variables – red

picture | red | green | blue

227x227 uint8

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
1	125	127	126	125	122	123	122	120	119	120	118	118	119	120	122	128	152	164	163	163	163	162
2	125	125	124	124	122	122	122	119	116	116	117	118	118	118	119	130	158	164	165	165	162	162
3	128	123	122	121	121	123	122	118	113	114	115	116	117	117	118	129	159	165	167	165	162	163
4	128	122	120	120	119	120	120	118	115	116	115	114	116	117	116	137	164	170	168	167	162	164
5	129	124	122	121	119	121	120	119	117	118	117	114	113	114	116	142	168	171	170	168	166	167
6	128	125	123	122	119	118	118	118	116	118	117	120	116	115	115	146	172	172	170	171	171	170
7	125	125	123	122	121	118	118	116	114	116	114	115	113	113	119	153	168	175	174	173	173	173
8	125	124	125	122	118	118	116	114	112	113	114	113	111	112	128	157	167	170	173	177	175	177
9	126	122	122	119	117	119	114	114	115	112	115	113	110	111	130	157	165	166	169	176	181	183
10	128	125	122	121	119	118	114	114	115	114	113	111	111	111	134	157	163	166	168	170	176	182
11	129	127	125	121	119	118	117	115	111	112	111	110	112	110	138	155	160	163	165	167	172	176
12	128	123	124	122	120	118	117	116	110	110	110	111	111	113	139	155	158	162	164	165	167	171
13	133	125	126	123	119	118	115	111	107	109	110	110	110	113	139	154	156	160	164	165	168	168
14	134	128	125	121	118	119	114	110	109	107	108	110	109	109	135	152	154	158	162	164	164	167
15	131	127	126	122	118	117	116	113	112	109	106	108	108	107	133	150	152	158	159	162	163	166
16	131	127	127	122	121	119	117	112	112	110	109	108	107	105	127	149	150	155	156	161	163	163
17	134	129	127	123	121	118	115	112	110	110	110	107	108	105	121	144	148	152	155	159	161	163
18	134	130	127	124	121	116	113	111	109	108	106	105	109	108	115	138	146	148	150	149	147	151
19	133	131	128	127	121	114	113	112	107	107	106	107	106	105	111	134	142	144	143	143	145	145
20	134	132	127	126	121	115	110	110	108	108	106	107	105	105	107	128	141	140	139	147	150	151
21	135	133	127	124	118	115	111	109	109	107	105	105	104	105	105	112	135	138	141	147	150	150
22	137	132	126	123	117	115	111	111	109	107	105	103	104	104	105	105	118	133	141	147	146	149
23	135	129	127	124	120	115	108	109	109	107	105	102	104	104	103	106	111	116	133	143	140	141
24	136	132	130	127	119	115	110	107	107	107	105	102	102	104	102	105	111	113	115	132	140	135
25	137	133	129	124	118	116	112	108	106	106	104	102	101	101	102	104	111	112	111	116	136	141

What we see

Image of dimension N*M

What computers see

tensor of N * M * 3 integers between 0 and 255
(for 8-bits RGB images)

Convolutional Neural Networks aka CNN

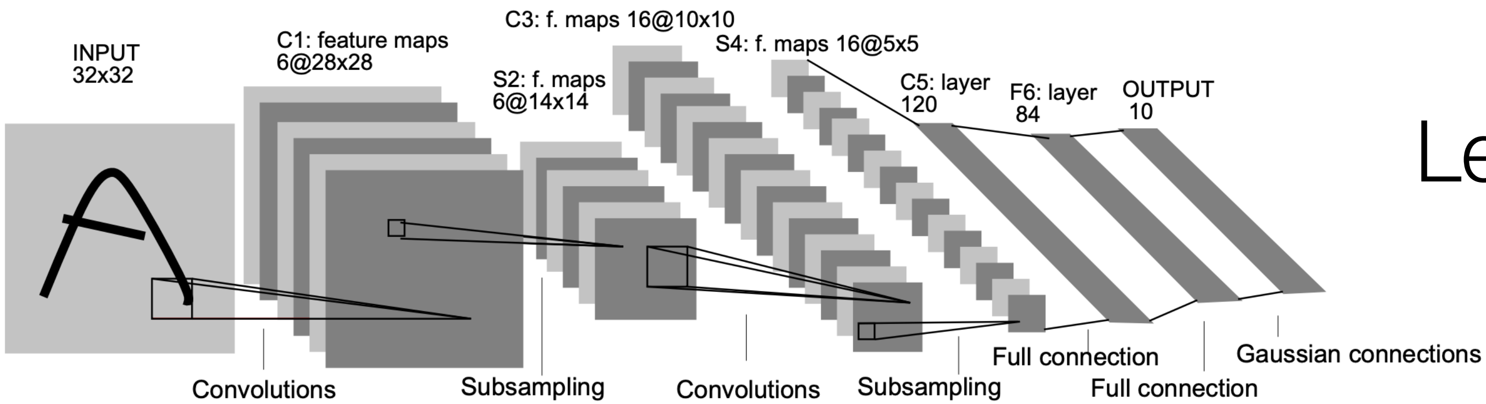

LeNet-5

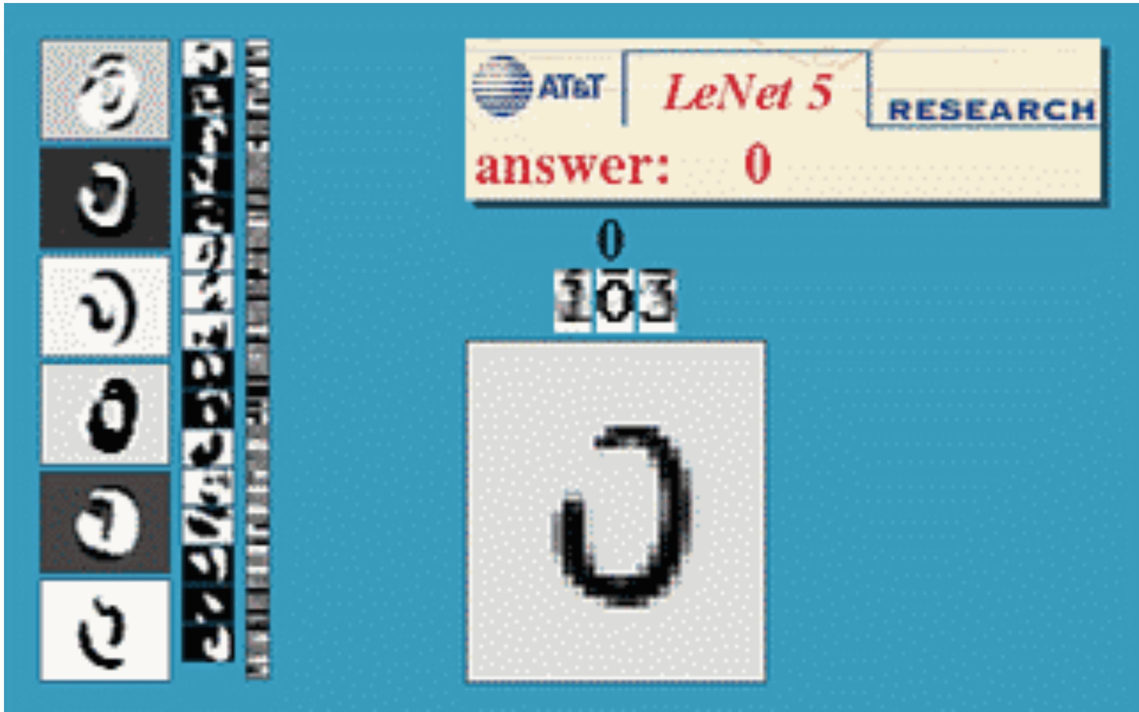

Convolution Layer

« receptive field »

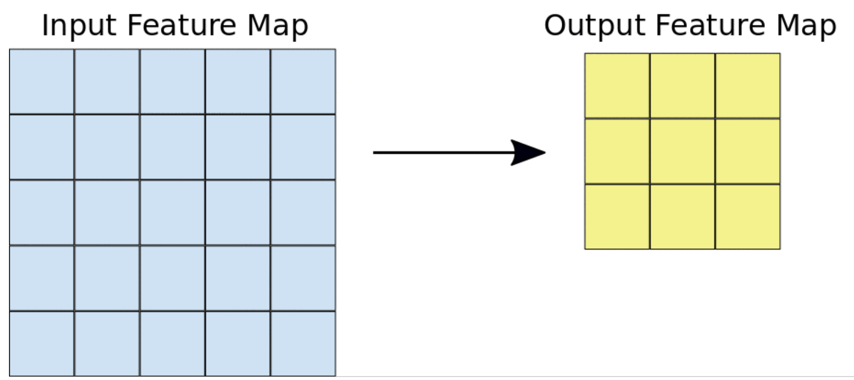

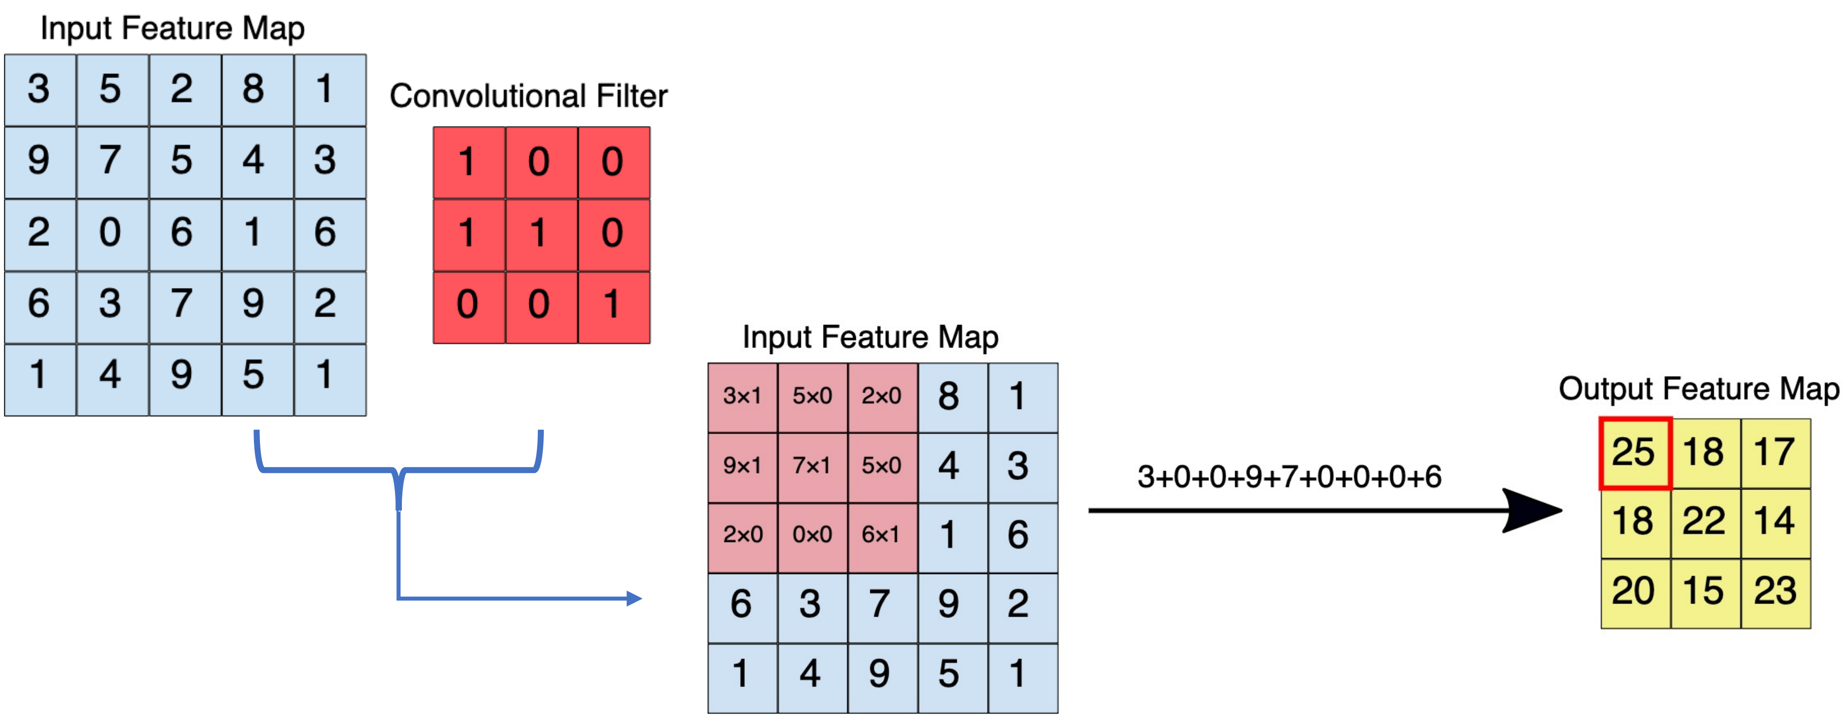

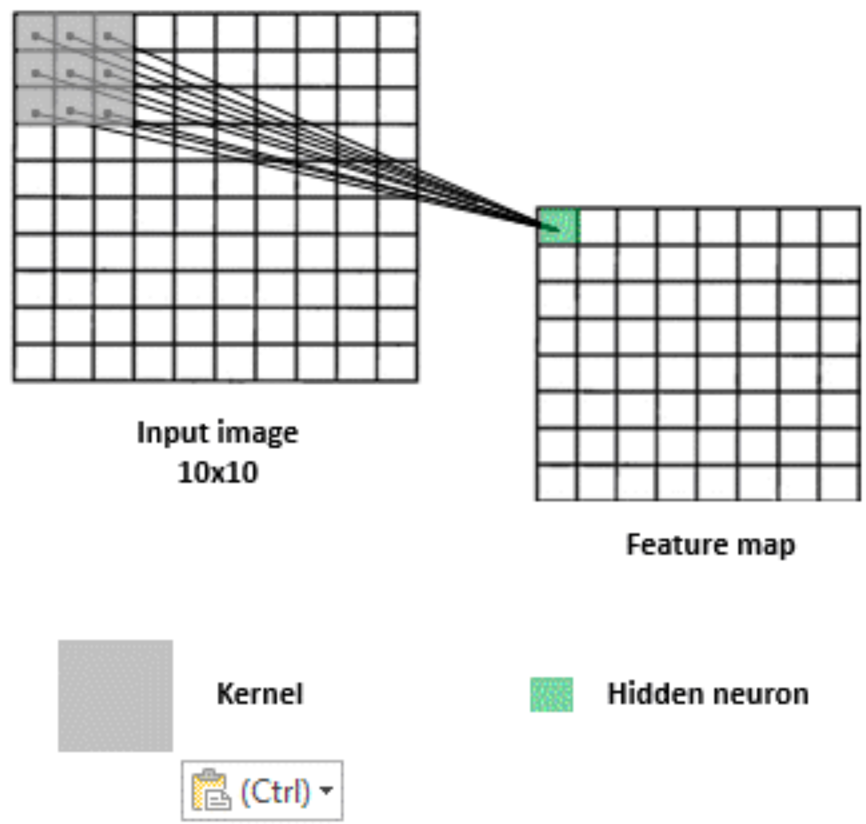
Input image
10x10
Feature map
Kernel
Hidden neuron
(Ctrl)

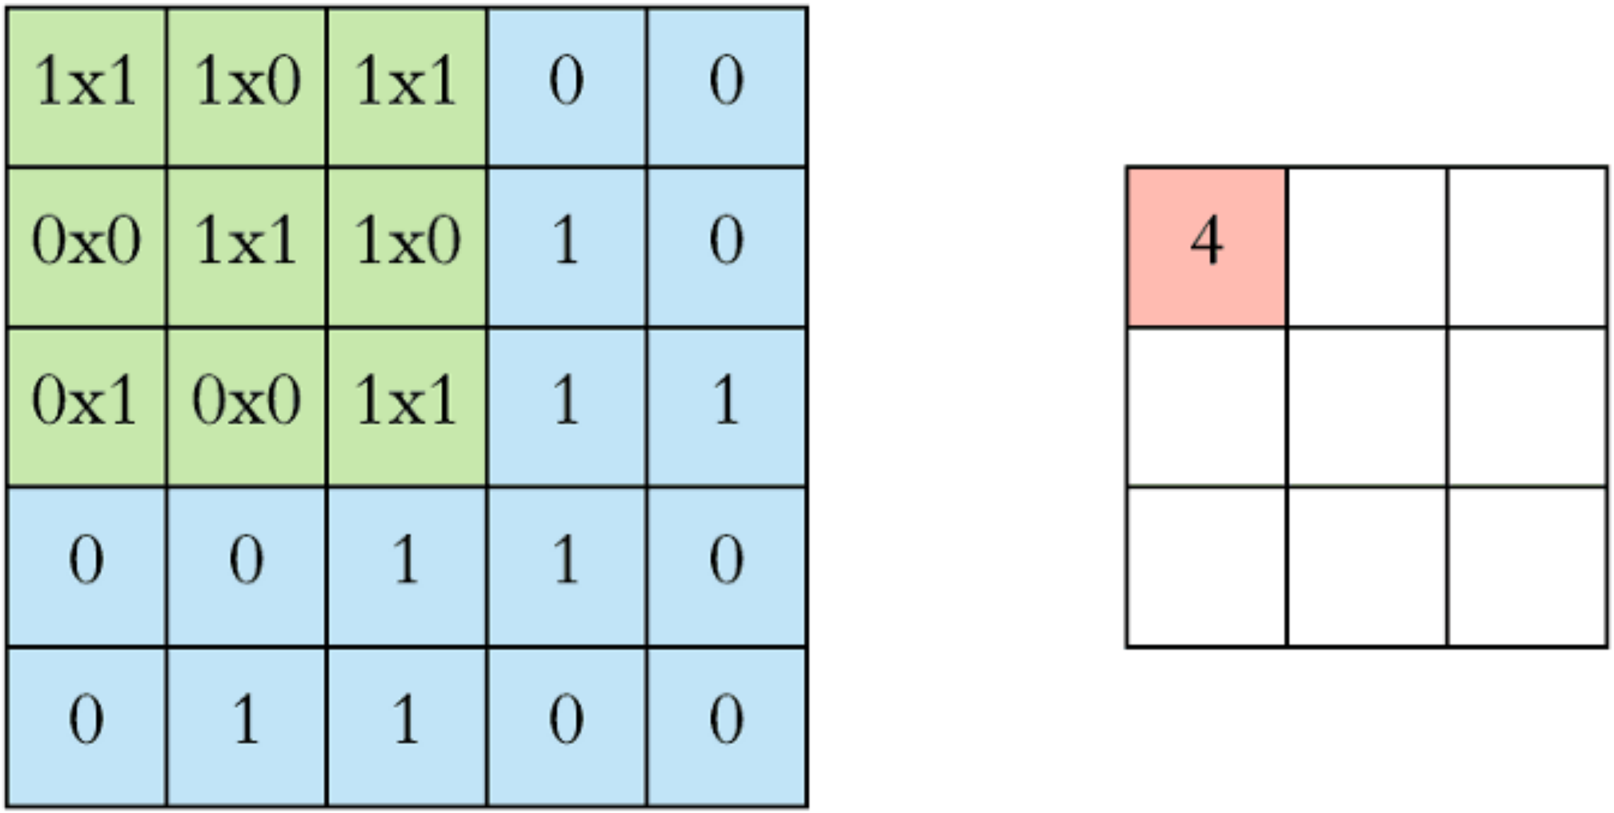
1x1 1x0 1x1 0 0
0x0 1x1 1x0 1 0
0x1 0x0 1x1 1 1
0 0 1 1 0
0 1 1 0 0
4

Convolution Layer : the feature detector

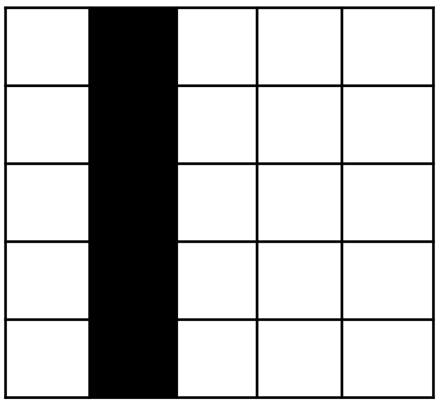

Input

*

0	1	0
0	1	0
0	1	0

→

3	0	0
3	0	0
3	0	0

*

0	0	0
1	1	1
0	0	0

→

1	1	0
1	1	0
1	1	0

filter (weights)

feature map

Pooling Layer

- Downsampling the feature maps by associating with a smaller filter a single value

  (eg : max-pooling with 2*2 filter)

- Saves computation time, reduces dimensions, but preserves critical features

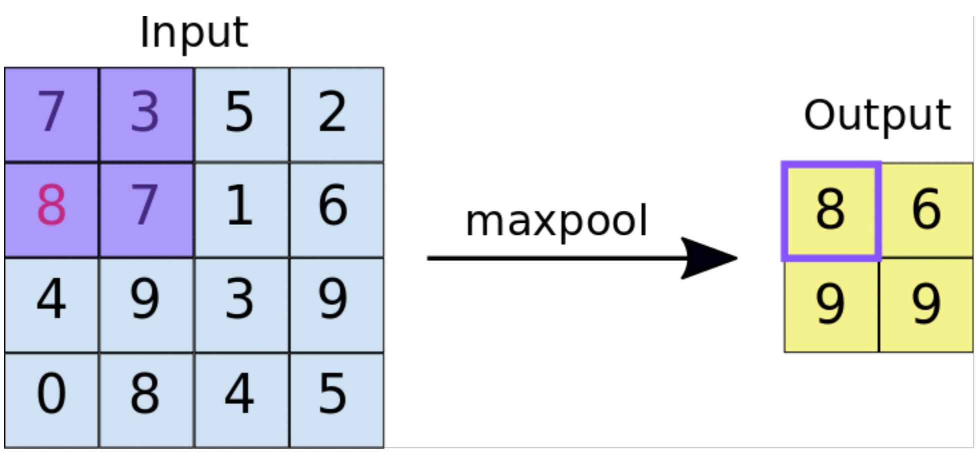

Convolution Layer + Maxpool layer

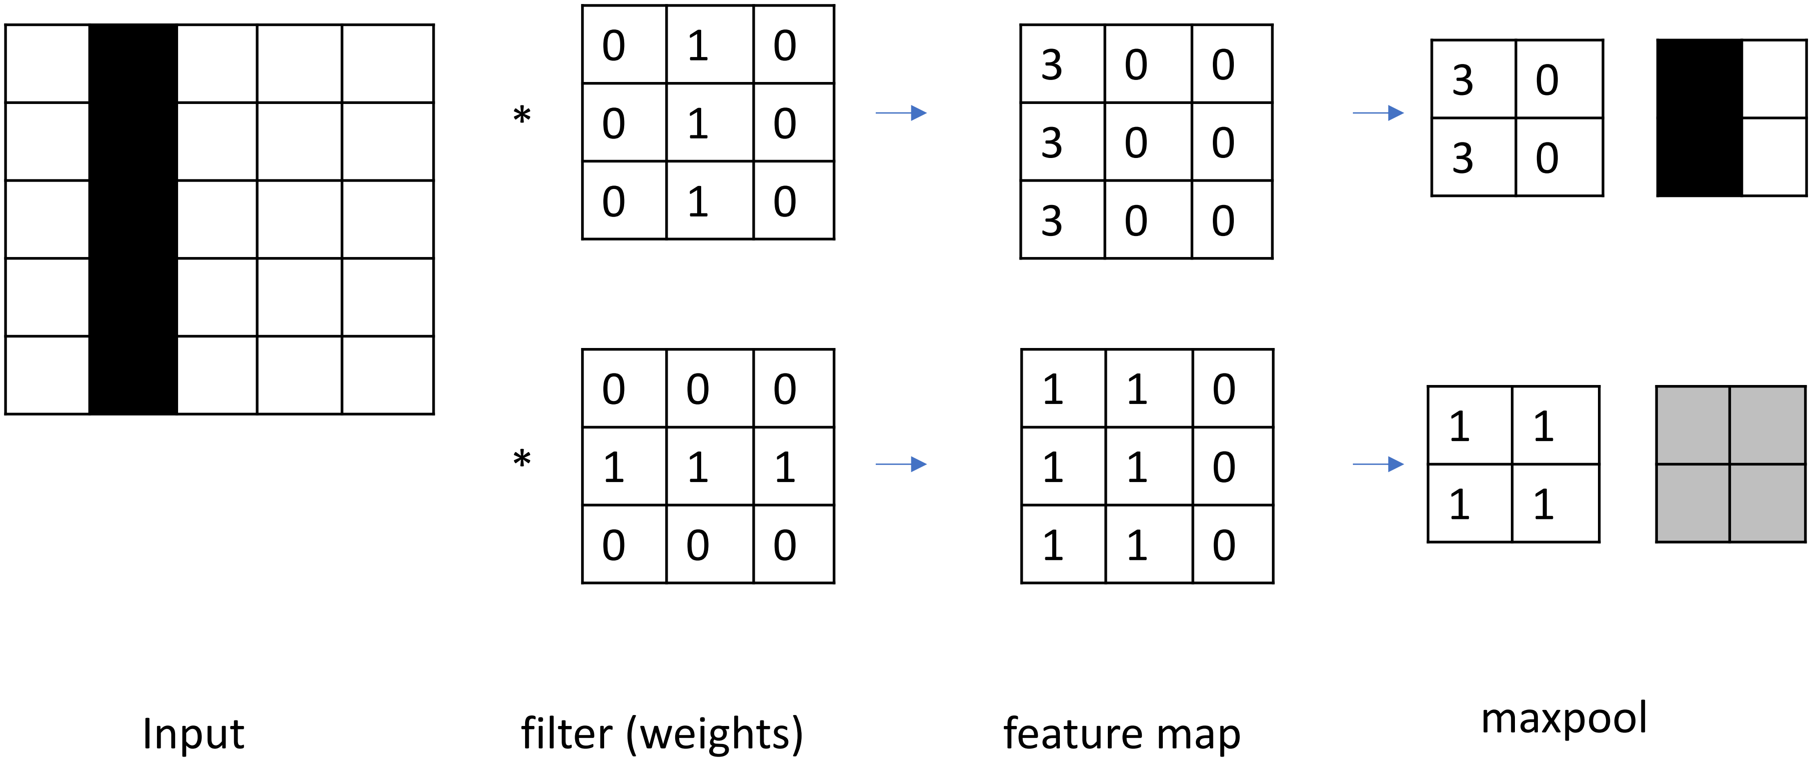

Fully connected layers

- End of CNN : one or more fully connected layers

- Performs classification based on the features extracted by the convolutions.

- Outputs a probability value from 0 to 1 for each of the classification labels the CNN is trying to predict

- One output neuron / category to include

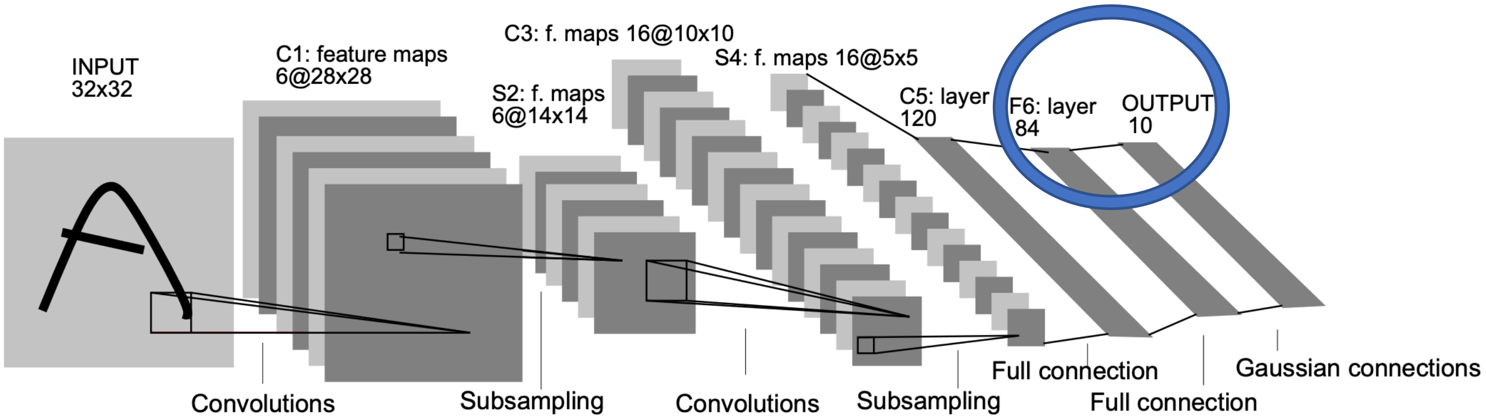

ConvNets for image recognition

CNN builds up more and more abstract representations of the image

Output : probability of each category

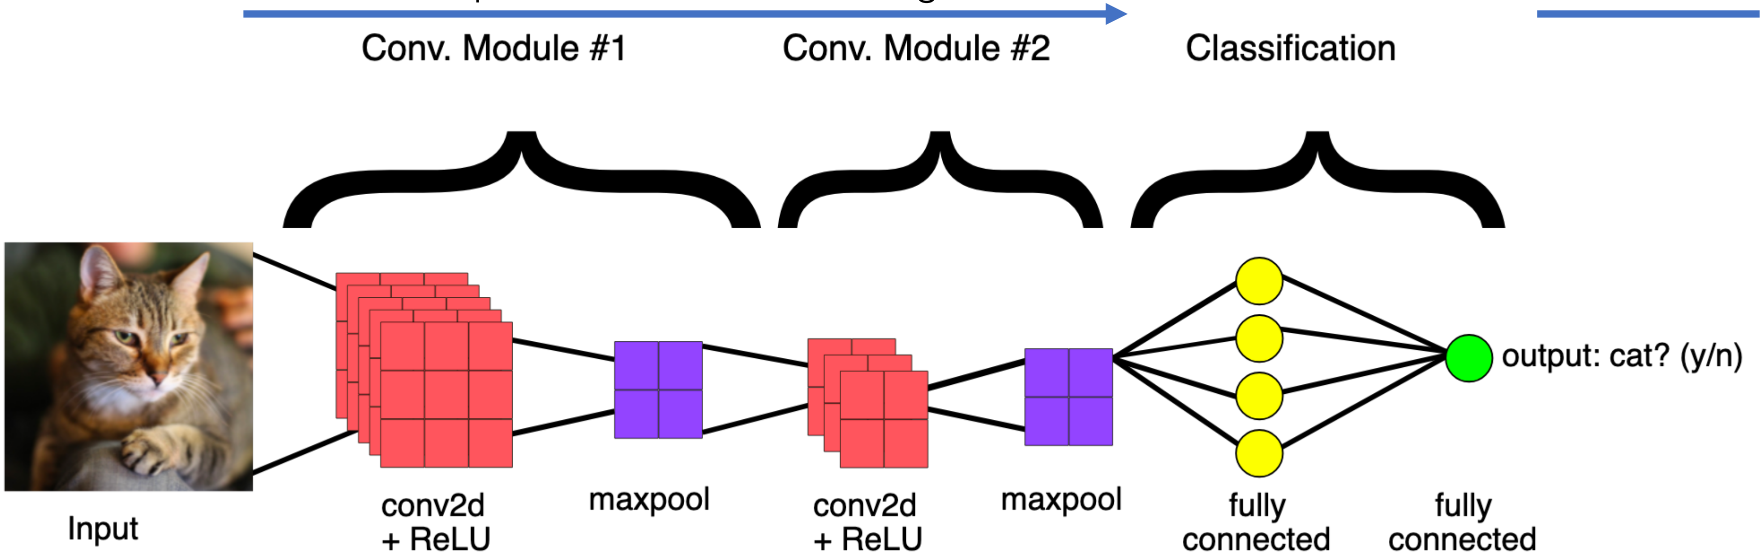

Figure 5.2. The visual world forms a spatial hierarchy of visual modules: hyperlocal edges combine into local objects such as eyes or ears, which combine into high-level concepts such as “cat.”

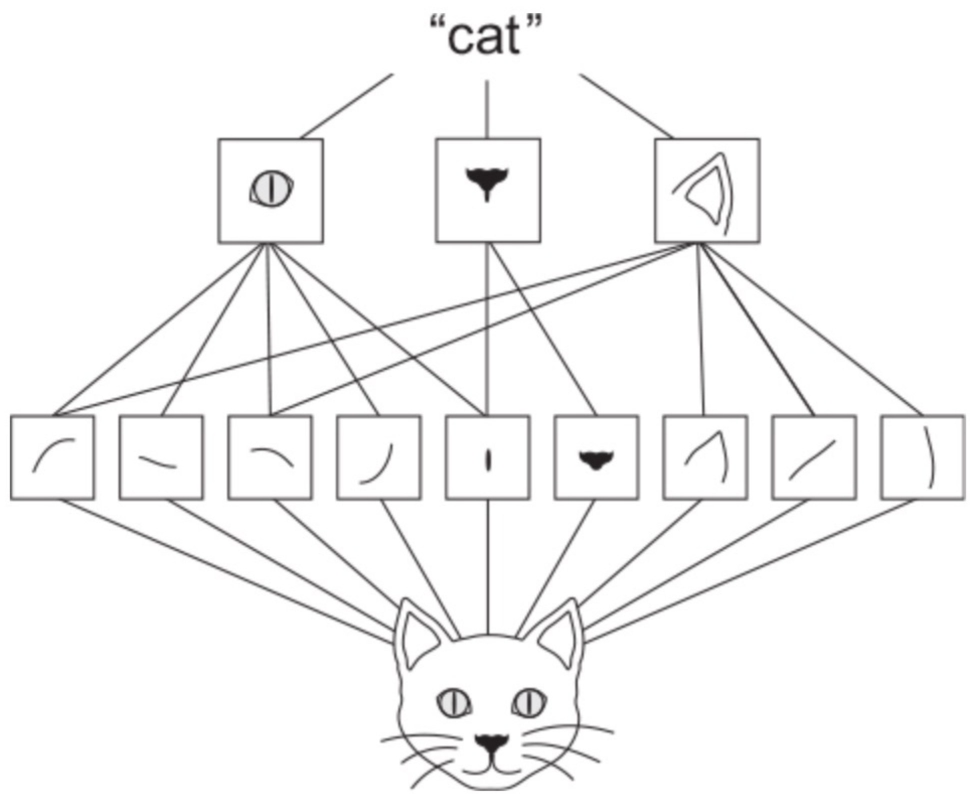

Démo : détection d'objets, YOLOv4

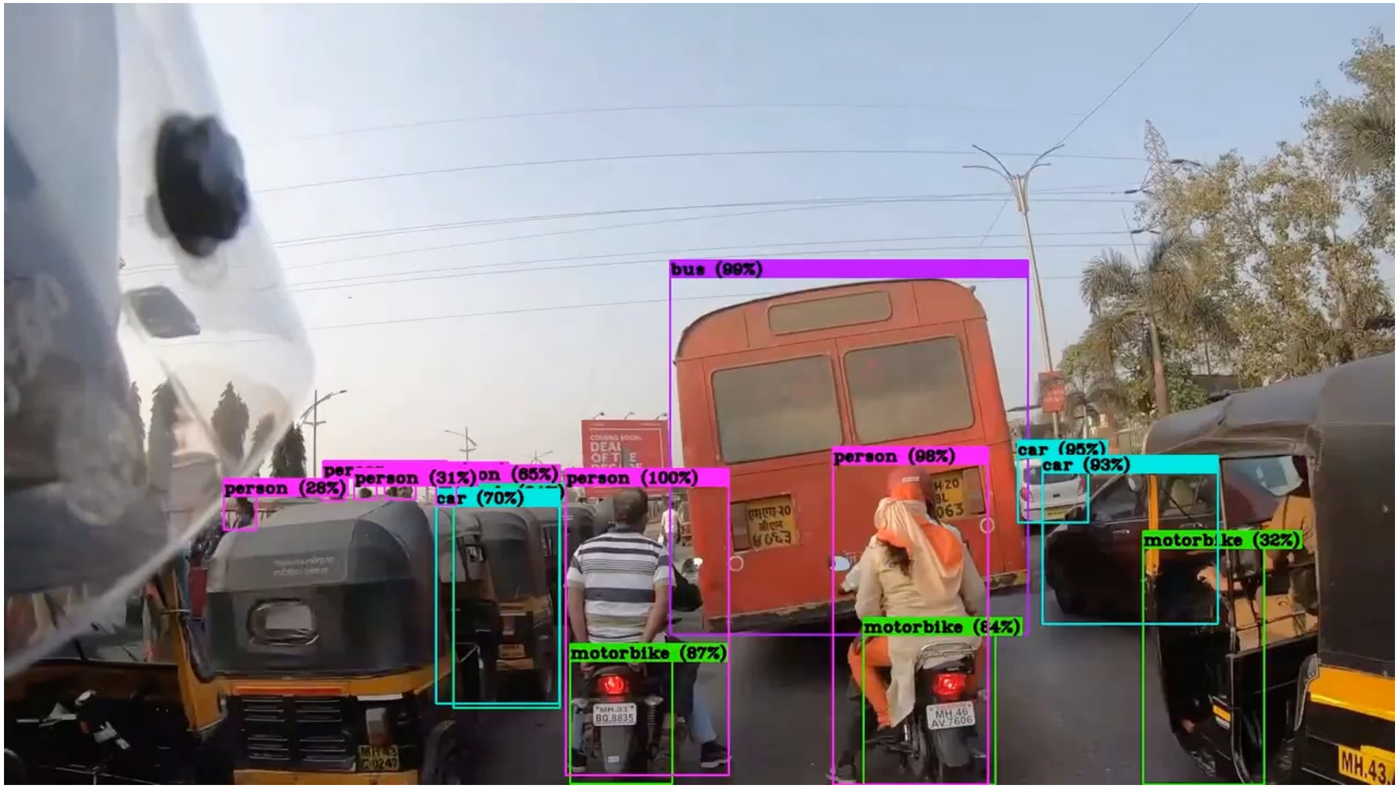

Source : https://www.youtube.com/watch?v=h08N0HX16l8 , GrahamML

Démo : Quickdraw (AI Experiments, Google)
https://quickdraw.withgoogle.com/

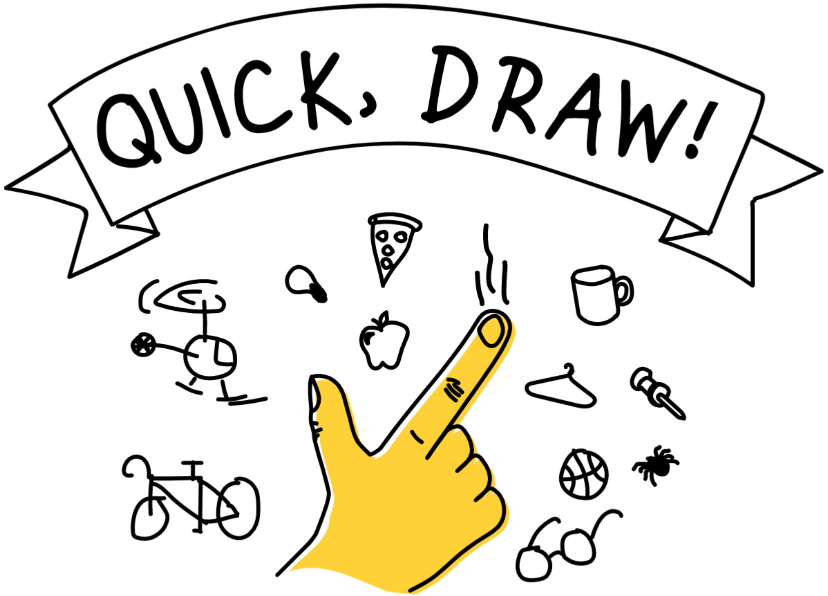

Demo : ChatGPT

LI Please write a dog food advertisement in a magazine read by dogs

"Woof! Are you tired of the same old kibble day after day? Switch things up with our delicious and nutritious dog food! Made with high-quality ingredients, our food will keep your tail wagging and your taste buds satisfied. Plus, our specially-formulated recipes cater to dogs of all sizes and breeds, so every pup can enjoy a hearty meal. Try it today and see the difference it makes in your dog's health and happiness. Order now and receive a free bag of treats with your purchase! Woof woof!"

What is the diameter of an apple divided by the thickness of an apple peel?

The diameter of an apple can vary depending on the variety, but on average it's about 7-10 cm. The thickness of an apple peel is around 1-2 mm.

Dividing the diameter of an apple by the thickness of the peel, we get:
7 cm / 0.1 cm = 70

So, the diameter of an apple divided by the thickness of an apple peel is approximately 70.

Note: The above numbers are estimates and the exact numbers may vary depending on the apple being considered.

So, if Earth was the size of an apple, the oceans would be thinner than the apples peal?

No, if the Earth was the size of an apple, the ocean depths would be much thicker than the apple peel.

To provide some perspective, the diameter of an apple divided by the thickness of the apple peel is around 70, as calculated earlier. However, the diameter of the Earth divided by the deepest ocean depth is approximately 1162.8. This indicates that the ocean depths would be much thicker than the apple peel, if the Earth was scaled down to the size of an apple.

You are actually wrong, it is the other way around. We are talking about fractions, and 1/1162 is actually smaller than 1/70. But you made good improvement with the latest update, good job!

Thank you for your feedback. I apologize for the mistake in my previous answer. Yes, you are correct, if the Earth was the size of an apple, the ocean depths would be much thinner than the apple peel. The fraction 1/1162 is indeed smaller than 1/70. I appreciate the opportunity to correct my mistake. Thank you for your understanding.

SU How do I get my crush to like me?

Getting someone to like you can be difficult, as it depends on many factors such as compatibility and timing. However, there are some things you can do to increase the chances of your crush developing feelings for you:

- Be confident: Confidence is attractive, so try to be comfortable in your own skin and believe in yourself.
- Show interest: Let your crush know that you are interested in them by flirting, making eye contact, and showing genuine interest in their life and interests.
- Be a good listener: Show your crush that you value what they have to say by actively listening to them and asking questions.
- Be kind and respectful: Treat your crush and others with kindness and respect, as people are more likely to be attracted to those who display good character.
- Be yourself: Don't try to be someone you're not in order to impress your crush, as ultimately, they will appreciate the real you.

my creativity and ability to come up with new ideas?"

Sure! Here are some sample behavioral interview questions about your ability to create and execute successful marketing campaigns:

1. Can you describe a time when you had to develop and implement a marketing campaign from start to finish? What was the campaign for, and what were the results?

research when planning a marketing
you use to gather information?
hannels will be most effective for a
campaign where you had to make this

rketing campaign? What metrics do you use,
n future campaigns?
be creative in order to solve a problem or
How did your solution impact the campaign's

CR continue this guitar pattern

```
e||--1-----------------------------1----x------6-------------||
B||--1--------------1--------------1----x------8-------------||
G||--3----3---------3--------------2----x------10*-----------||
D||--3------3--------3----3---------5----x------8-------------||
A||--1--1-----1-0----1------1--------3----x---6----------------||
E||--------------3--3--3-----3-0-1--1----x--------------------||
```

```
e||--1----------------------------
B||--1--------------1-------------
G||--3----3---------3-------------
D||--3------3--------3----3-------
A||--1--1-----1-0----1------1-----
E||--------------3--3--3-----3-0-1
e||--1--4--6--4--1--4--6--4--1--4
B||--1--4--6--4--1--4--6--4--1--4
G||--3--5--7--5--3--5--7--5--3--5
D||--3--5--7--5--3--5--7--5--3--5
A||--1--3--5--3--1--3--5--3--1--3
E||--------------3--3--3-----3-0-1
```

This is an example of how the pa
the desired style or genre of mus

HO Out of these four shoes, which would you say works best with outfit #3: light brown Timberlands, white Air Force Ones, black Air Max, dark red Converse

Of the four shoe options you mentioned, the white Air Force Ones would work best with Outfit #3 (white T-shirt, black leather jacket, and dark blue jeans). The white Air Force Ones are a classic and versatile sneaker that can add a clean and casual look to the outfit.

The black Air Max may be a bit too sporty for the outfit, as the black leather jacket and dark blue jeans have a more edgy and casual vibe. The dark red Converse may clash with the black leather jacket and look out of place with the rest of the outfit. The light brown Timberlands may work, but they could be a bit too bulky and heavy for the outfit, as the outfit already has a lot of texture and weight with the leather jacket and jeans.

Remember, your personal style and preference should ultimately guide your choice of shoes, so if you feel that one of the other options complements your outfit better, then go for it!

Identifier son interlocuteur

Visage Technologies

https://www.visagetechnologies.com/HTML5/latest/Samples/ShowcaseDemo/ShowcaseDemo.html

Deep Learning for AI

- GAFAM :

  DeepMind, AlphaGo, Google Car, Cortana, Siri, Watson etc

- Image recognition, automatic labelling, voice recognition, buying recommandations, automatic translation, art creation, fraud detection, spam filtering, customer profiling, personalization of contents, facial identification, emotion recognition, video creation, chatbots, game playing, natural language processing, etc …

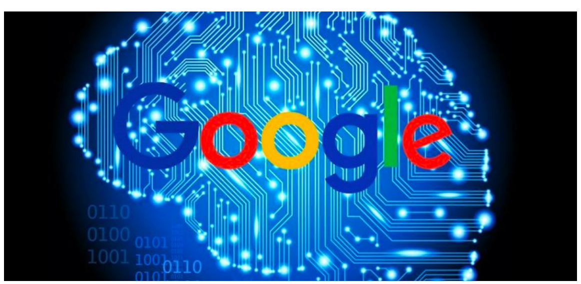

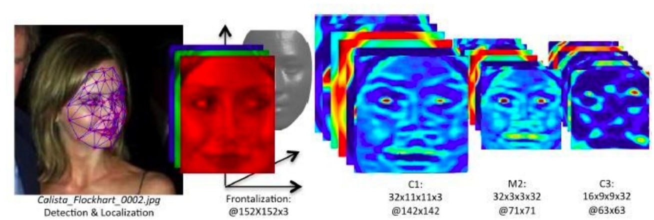

Deep Learning for … medical applications?

- DL is a great tool for image recognition and classification

- All medical diagnostics that are based on medical images and that require image classification / object detection can benefit from DL :

  - Dermatology
  - Radiology
  - Ophtalmology
  - Pathology…

- DL can help physicians to detect relevant areas in images
- DL can help to predict the outcomes
- DL can offer a second opinion
- DL can be « the physician's third eye » (a well-trained eye)

Fig. 1

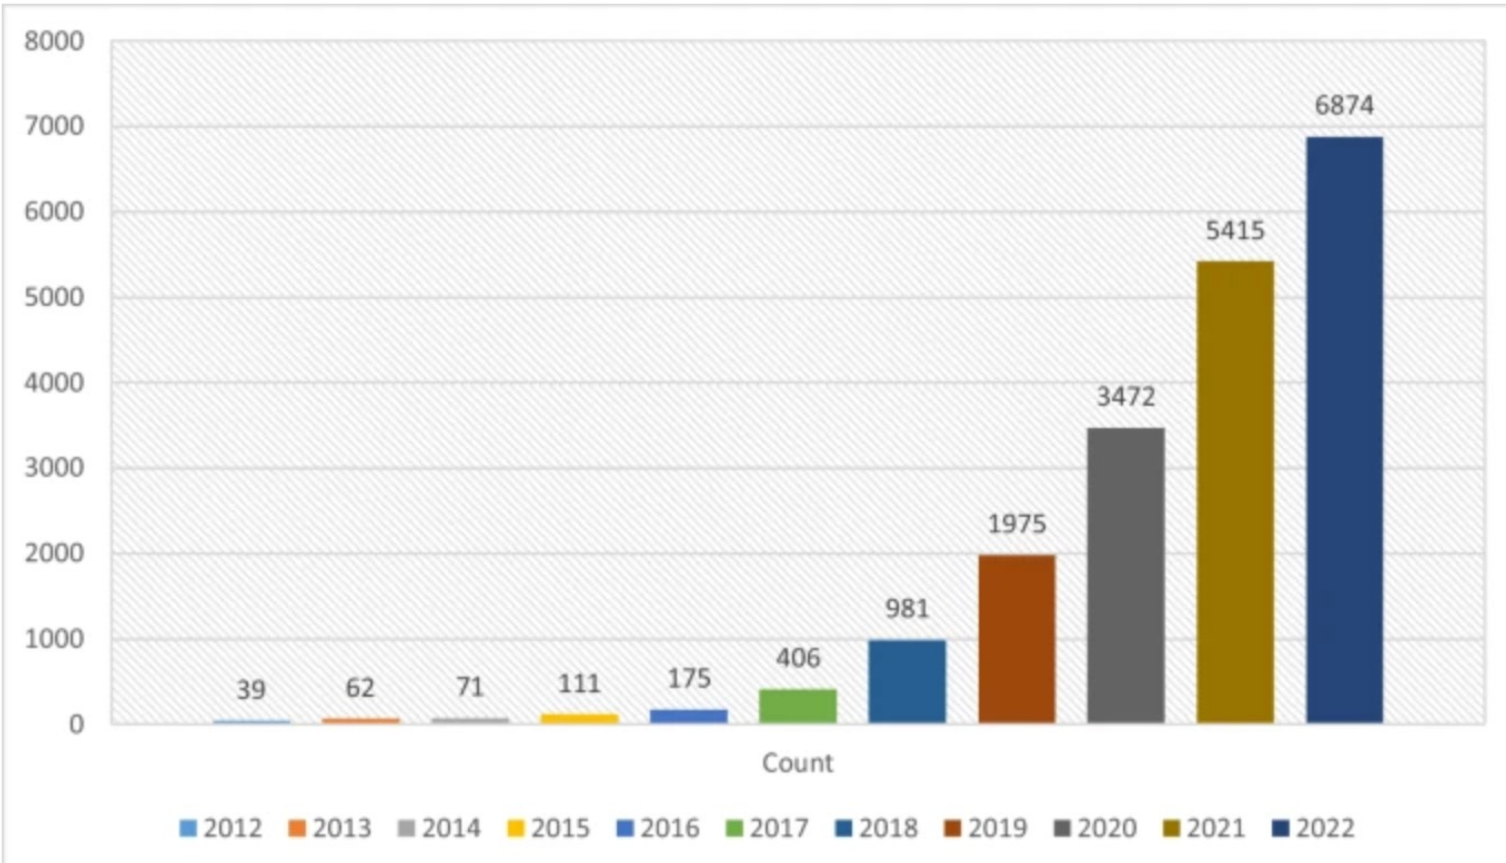

The yearly distribution of DL techniques in HCS in the PubMed database for the last decade

AI in Health : long story, many questions

- Old story, beginning with Expert Systems - GOFAI
- Problem : can we make « experience » explicit ?
- Can we replace humans with machines in the medical domain?
- Can we make patients accept a machine’s opinion ?
- What if there is no agreement ?

Champs d'application du DL en médecine

Clinical Images	Biosignals	Biomedicine	E.H.R.
Oncology	ECG	Genomics	Treatment Plans
Cardiovascular	EEG	Transcriptomics	Disease analisis
Orthopaedics	PCG	Proteomics	Risk analysis
Neurology	PPG	Metabolism	
Pulmonology	EMG		
Other Applications: Gastroenterology Ophthamolody Abdominal	NSS		
	SS		

1 – DL and clinical images

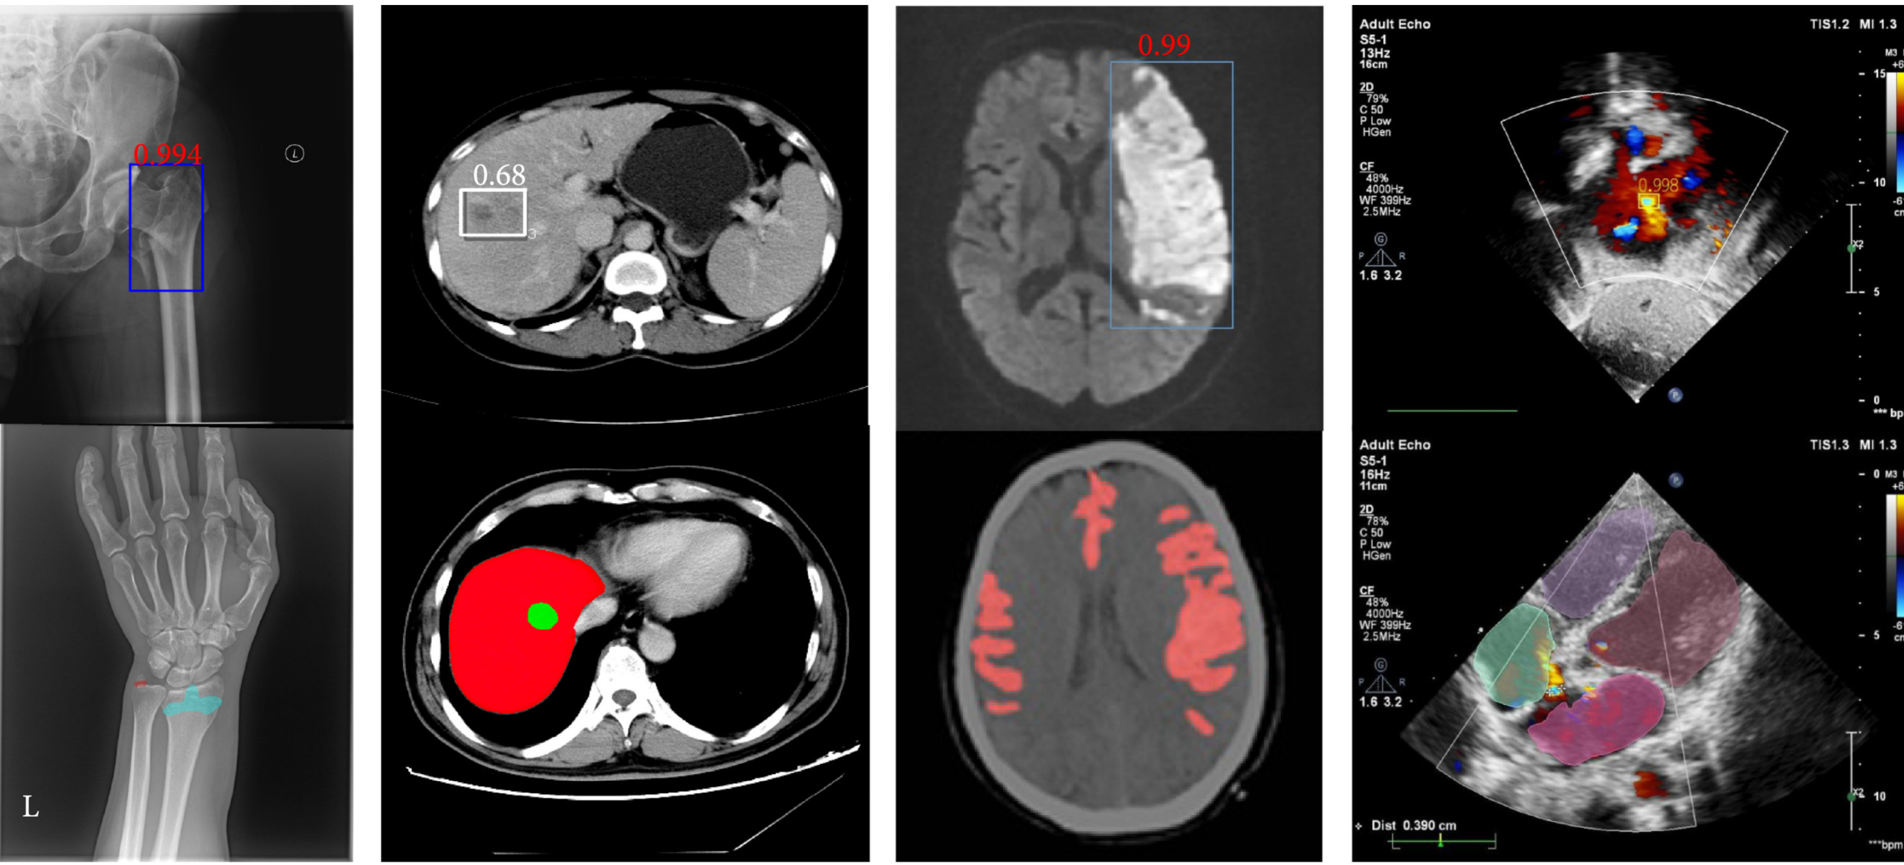

Bone X-ray | Liver CT | Brain MRI | Cardiac ultrasound

Area	Task	Contribution	Model
Oncology	Classification	Benefit from unlabelled data for lung tumour stratification	DBN [115]
		Introduction of a transfer learning approach in rectal cancer prediction	CNN [118]
		Identification of bladder tumour sub-types from histopathological images	ResNet [114]
		Improvement in breast tumour estimation by considering a large set of risk factors	CNN [187]
		Estimation of the cancer grade	CNN [116]
		Estimation of the cancer type	CNN [111,112], ResNet [119]
	Detection	Optimal localization of lung cancer sub-types	CNN [113]
	Segmentation	Analysis of colour contrast and parameter variability issues in pancreatic tumour	U-Net [127]
		Impact of dimension variations on DL model performance in thyroid melanomas	U-Net [128]
		Limitation of the overfitting problem in bone cancer	CNN [188], GAN+U-Net [91]
		Improvement in image accuracy in lung and prostate cancer	U-Net [124,125], GAN [92]
	Registration	Optimized DL model in terms of time complexity and accuracy in lung melanoma estimation	CNN [129,130]
	Dose estimation	Estimation of the right substance doses	GAN [131,132]
Cardiovascular	Classification	Limitation of overfitting	GAN [142], ResNet [140]
		Analysis of the particular characteristics of the heart by using echocardiograms	ResNet [141]
	Detection	Low-cost object detection for malaria	YOLO [143]
	Segmentation	DL model for multi-step integration and registration error reduction in atrial fibrillation analysis	CNN+LSTM [189]
		Accuracy in the analysis of irregular pelvic hematoma images	U-Net [190]
		Improvement in aortic disease analysis with the introduction of new accuracy measures	U-Net [144]
		Introduction of the transfer learning approach in atrium study	U-Net [140]
	Reconstruction	Improvement in the Signal-to-Noise Ratio Multi-data integration	CNN [145]
		Improvement in image quality at high levels in the study of coronary diseases	CNN [146]
Orthopaedics	Classification	Improvement in bone image quality	U-Net [153]
		Analysis of the impact of gender on skeletal muscles	CNN [154]
	Segmentation	Analysis of the impact of the image quality in osteoarthritis	U-Net [149], RCNN [152]
		Introduction of transfer learning and attention mechanism in the study of the knees	VGGNet+U-Net [150]
		Improvement in image accuracy of the cartilage	U-Net [148], HNN [42], U-Net+GAN [151], RCNN
		Combination of the region-based approach with U-Net for bone diseases	RCC+U-Net [191]
	Reconstruction	Application of CNNs to computed tomography for chest digital images	CNN [192]
Neurology	Classification	Automatic estimation of brain diseases risk	AlexNet [158], CNN [193]
	Segmentation	Limitation of overfitting in White Matter analysis	GAN [160]
		Colour quality improvement in orbital analysis	U-Net [159]
	Reconstruction	Introduction of a DAE as a priori model for noise density in magnetic resonance	DAE [162]
		Analysis of perturbation effects	CNN [157]
		Introduction of transfer learning into magnetic resonance	CNN [163]
Pulmonology	Classification	Improvement of accuracy and efficiency in COP diseases	ResNet [167], VGGNet+CNN [168], DBN [169]
		Analysis of interstitial lung diseases	CNN [165]
	Segmentation	Segmentation of lung lob using different types of datasets	U-Net [170]
Gastroenterology	Classification	Estimation of the normal levels of the pancreas	CNN [172,174]
	Detection	Improvement in image accuracy in neoplasia analysis	ResNet [175]
	Segmentation	Analysis of image effects in neoplasia and catheter detection	U-Net [175], RNN [176]
Ophthalmology	Classification	Improvement in image quality	CNN [179], CNN+LSTM [181]
Abdominal	Classification	Improvement in accuracy in abdominal ultrasounds	CNN [185]

DL in medecine : artificial eye

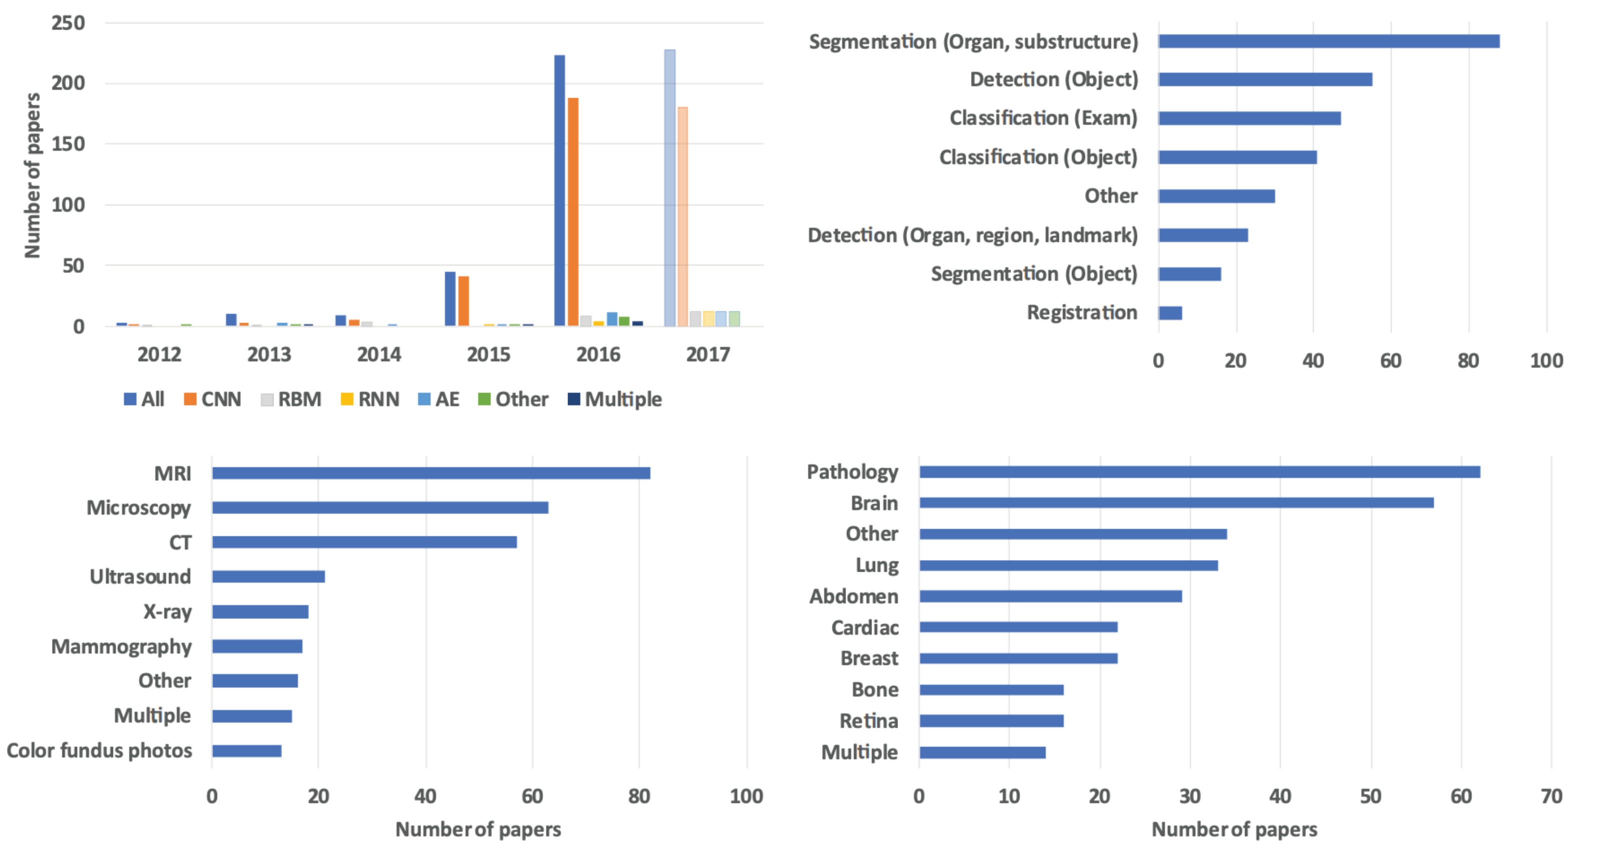

Figure 1: Breakdown of the papers included in this survey in year of publication, task addressed (Section 3), imaging modality, and application area (Section 4). The number of papers for 2017 has been extrapolated from the papers published in January.

With MRI images, a CNN network can do these things :

- Tissue/anatomy/lesion/tumor segmentation → 1
- Image (re)construction/enhancement → 2
- Disorder classification (eg. AD, MCI, Schizophrenia) → 3
- Lesion/tumor detection and classification
- Survival/disease activity/development prediction
- Other

All image-based tasks can benefit from DL

From images, DL can help to :

- Identifying moles vs melanomas
- Diabetic retinopathy
- Cardiovascular risks
- Breast lesion detection in mammograms
- Flagging large artery occlusions
- Predicting survival probabilities
- Combining imaging modalities

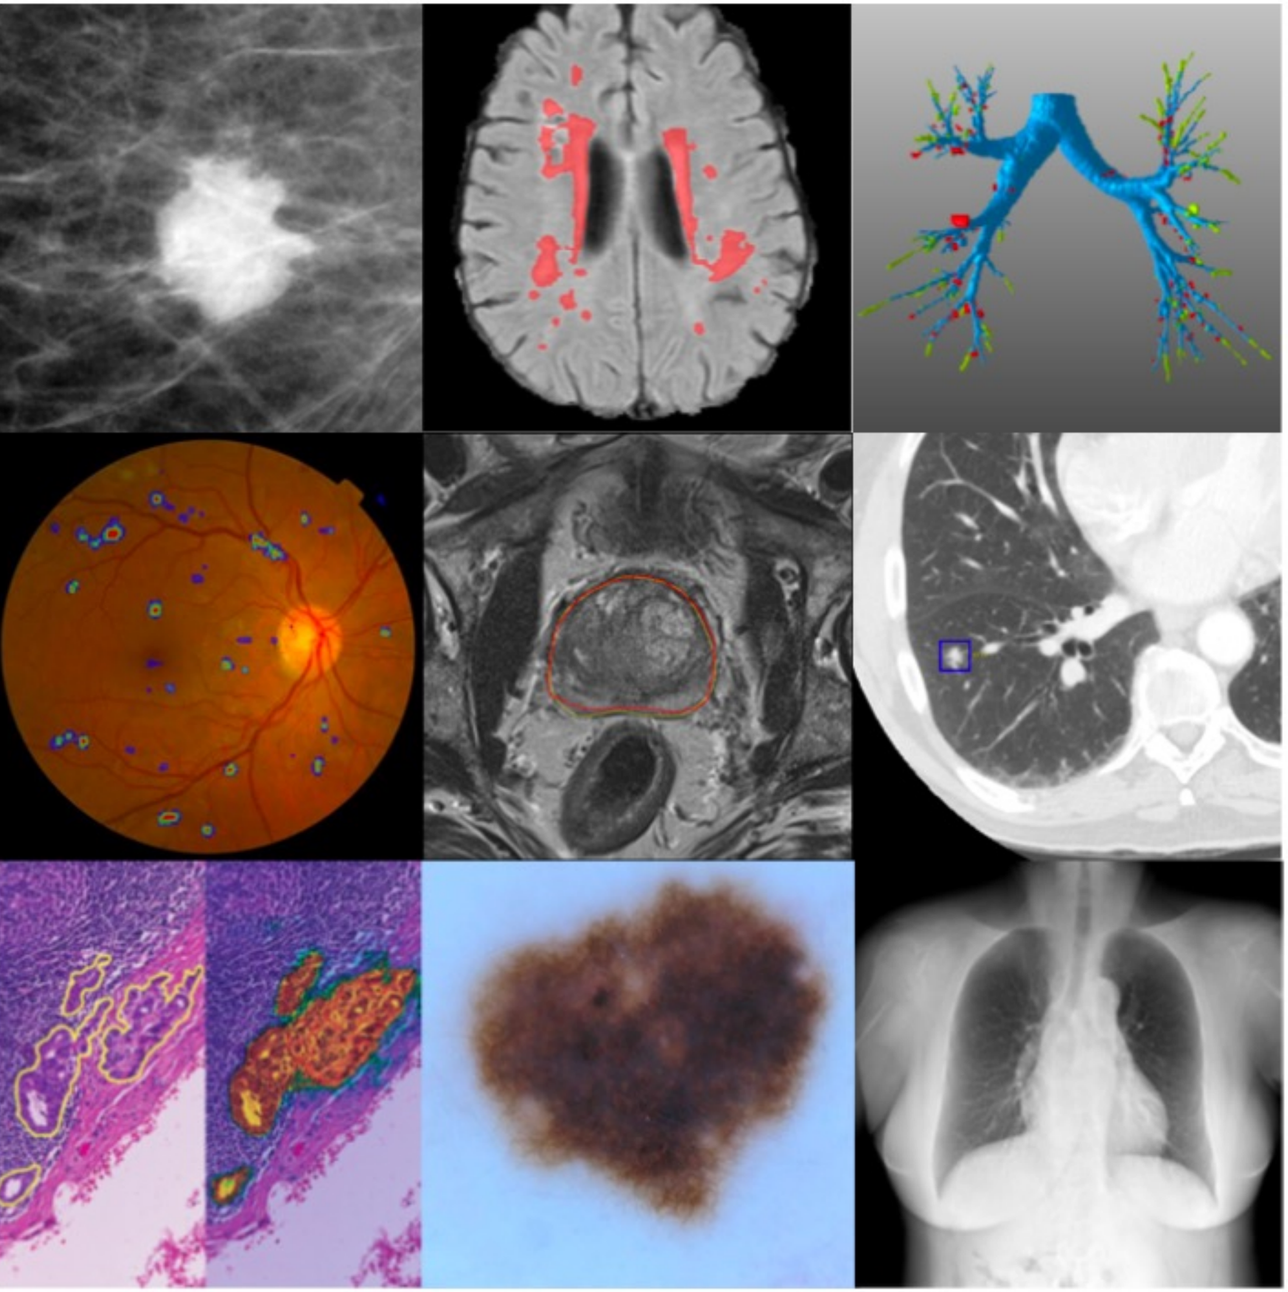

What DL can do :

Mammographic mass classification,
Segmenting brain lesions
Detecting leaks in airway tree,
Classifying diabetic retinopathy
Segmenting prostate,
Classifying nodules
Detecting breast cancer metastases in lymph nodes
Classifying skin lesions,
Suppressing bones in chest X-Rays

How does it works ?

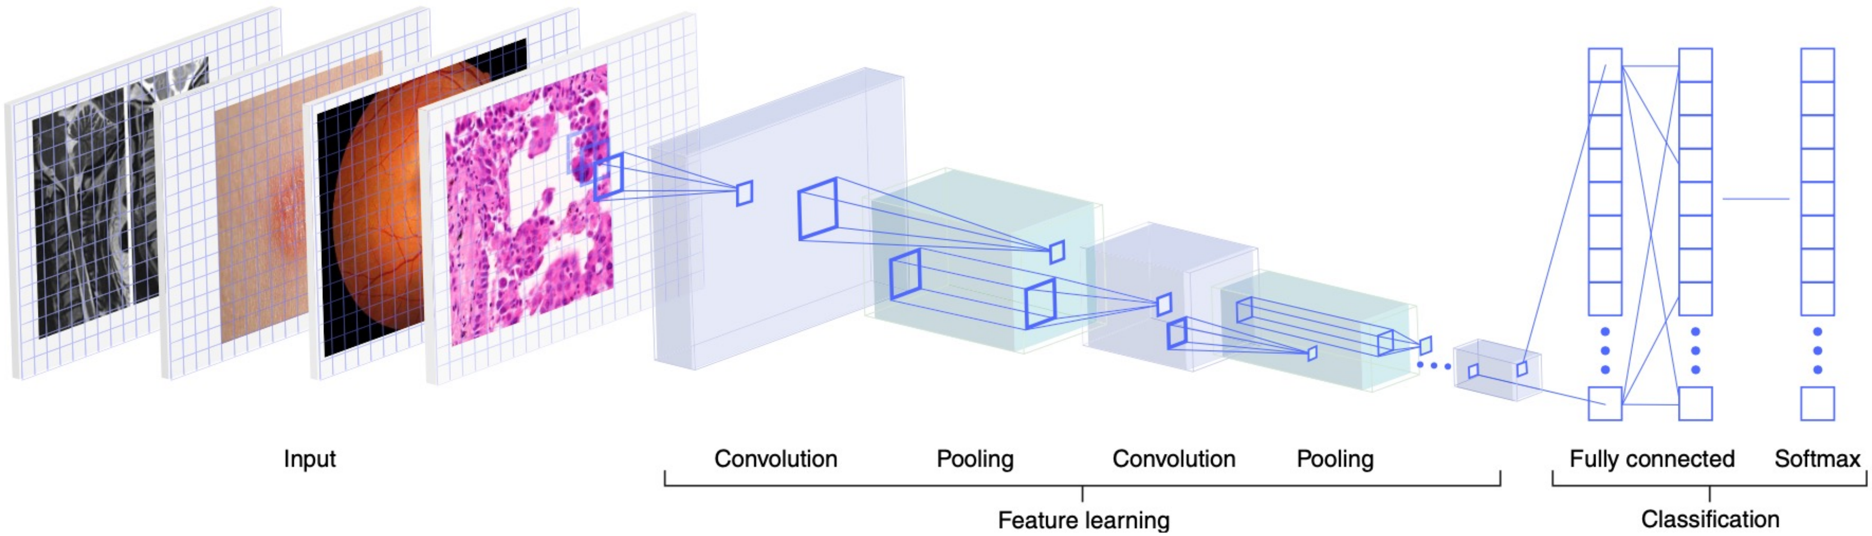

1.1- DL for segmenting medical images

- Segmenting from patches of images + context
- Segmenting = classifying « in » vs « out »
- Class balance is skewed severely towards « out » class in a training setting.
- Usually : « out » samples are easy to discriminate, preventing the deep learning method to focus on the challenging samples
- Segmenting organs, lesions, surgical objects (stents)

Segmenting the pancreas

- Ex. : pancreas segmentation with CNN from CT scans
- Pancreas has a very high anatomical variability : challenging

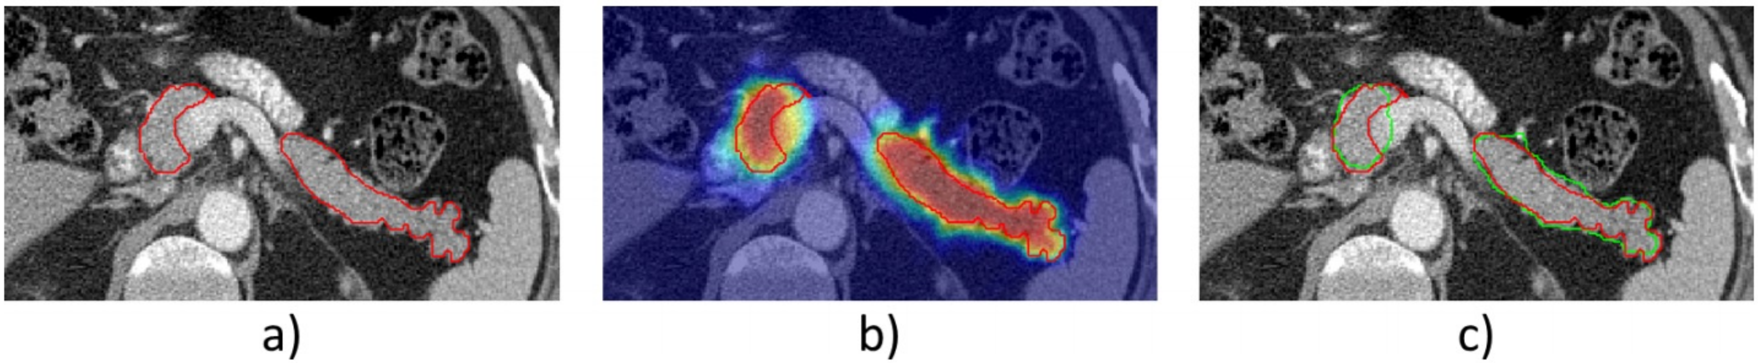

Figure 6: Example of pancreas segmentation using the proposed R_2−ConvNet approach in testing. a) The manual ground truth annotation (in red outline); b) the $G(P_2(x))$ probability map; c) the final segmentation (in green outline) at $p_2 = 0.6$ (DSC=82.7%).

Segmenting the brain

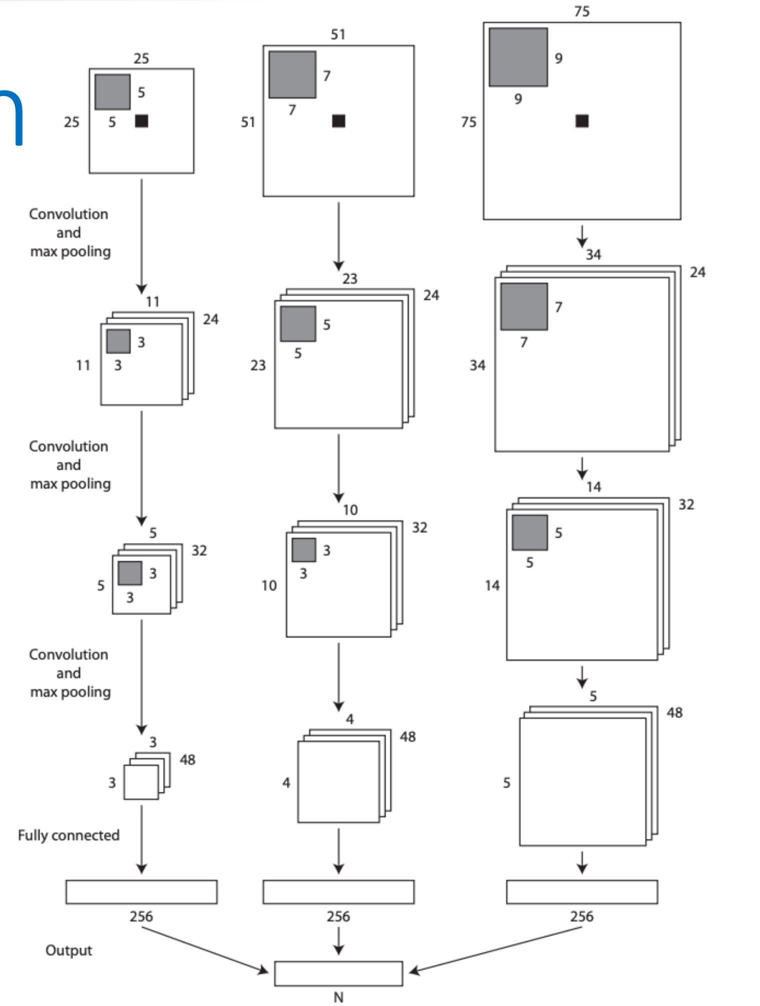

Fig. 1: Schematic overview of the convolutional neural network. The number of output classes, N, was set to 9 (8 tissue classes and background) for the neonatal images, to 8 (7 tissue classes and background) for the ageing adult images, and to 7 (6 tissue classes and background) for the young adult images. After the third convolution layer, max-pooling is only

Inside the CNNs…

TABLE I: Acquisition parameters for the images used in this paper.

	Cor. 30 wks	**Cor. 40 wks**	**Ax. 40 wks**	**Ageing adults**	**Young adults**
Age	30 weeks PMA	40 weeks PMA	40 weeks PMA	70 years	23 years
Acquisition protocol	Coronal T_2-weighted	Coronal T_2-weighted	Axial T_2-weighted	Axial T_1-weighted	Sagittal T_1-weighted
Number of images	10	5	7	20	15
Reconstruction matrix	384 × 384 × 50	512 × 512 × 110	512 × 512 × 50	240 × 240 × 48	256 × 256 × (261–334)
Reconstructed voxel sizes [mm^3]	0.34 × 0.34 × 2.0	0.35 × 0.35 × 1.2	0.35 × 0.35 × 2.0	0.96 × 0.96 × 3.0	1.0 × 1.0 × 1.0

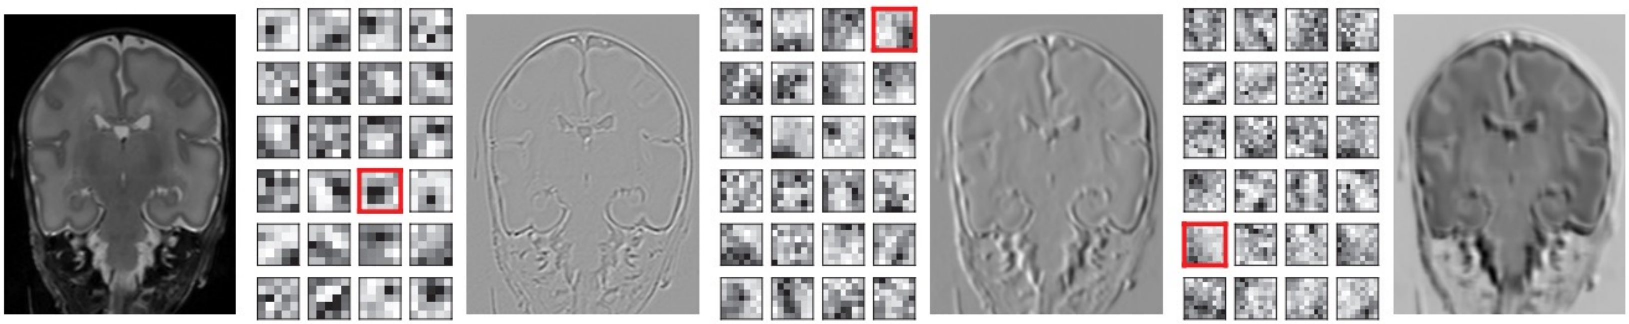

Fig. 2: Trained convolution kernels in the first layer after 10 epochs using the 5 training images acquired at 30 weeks PMA, and the kernels indicated in red applied to a test image. From left to right: the T_2-weighted test image, the kernels of 5 × 5 voxels, the image convolved with the indicated 5 × 5 kernel, the kernels of 7 × 7 voxels, the image convolved with the indicated 7 × 7 kernel, the kernels of 9 × 9 voxels, and the image convolved with the indicated 9 × 9 kernel.

30 - 40 weeks PMA aging young

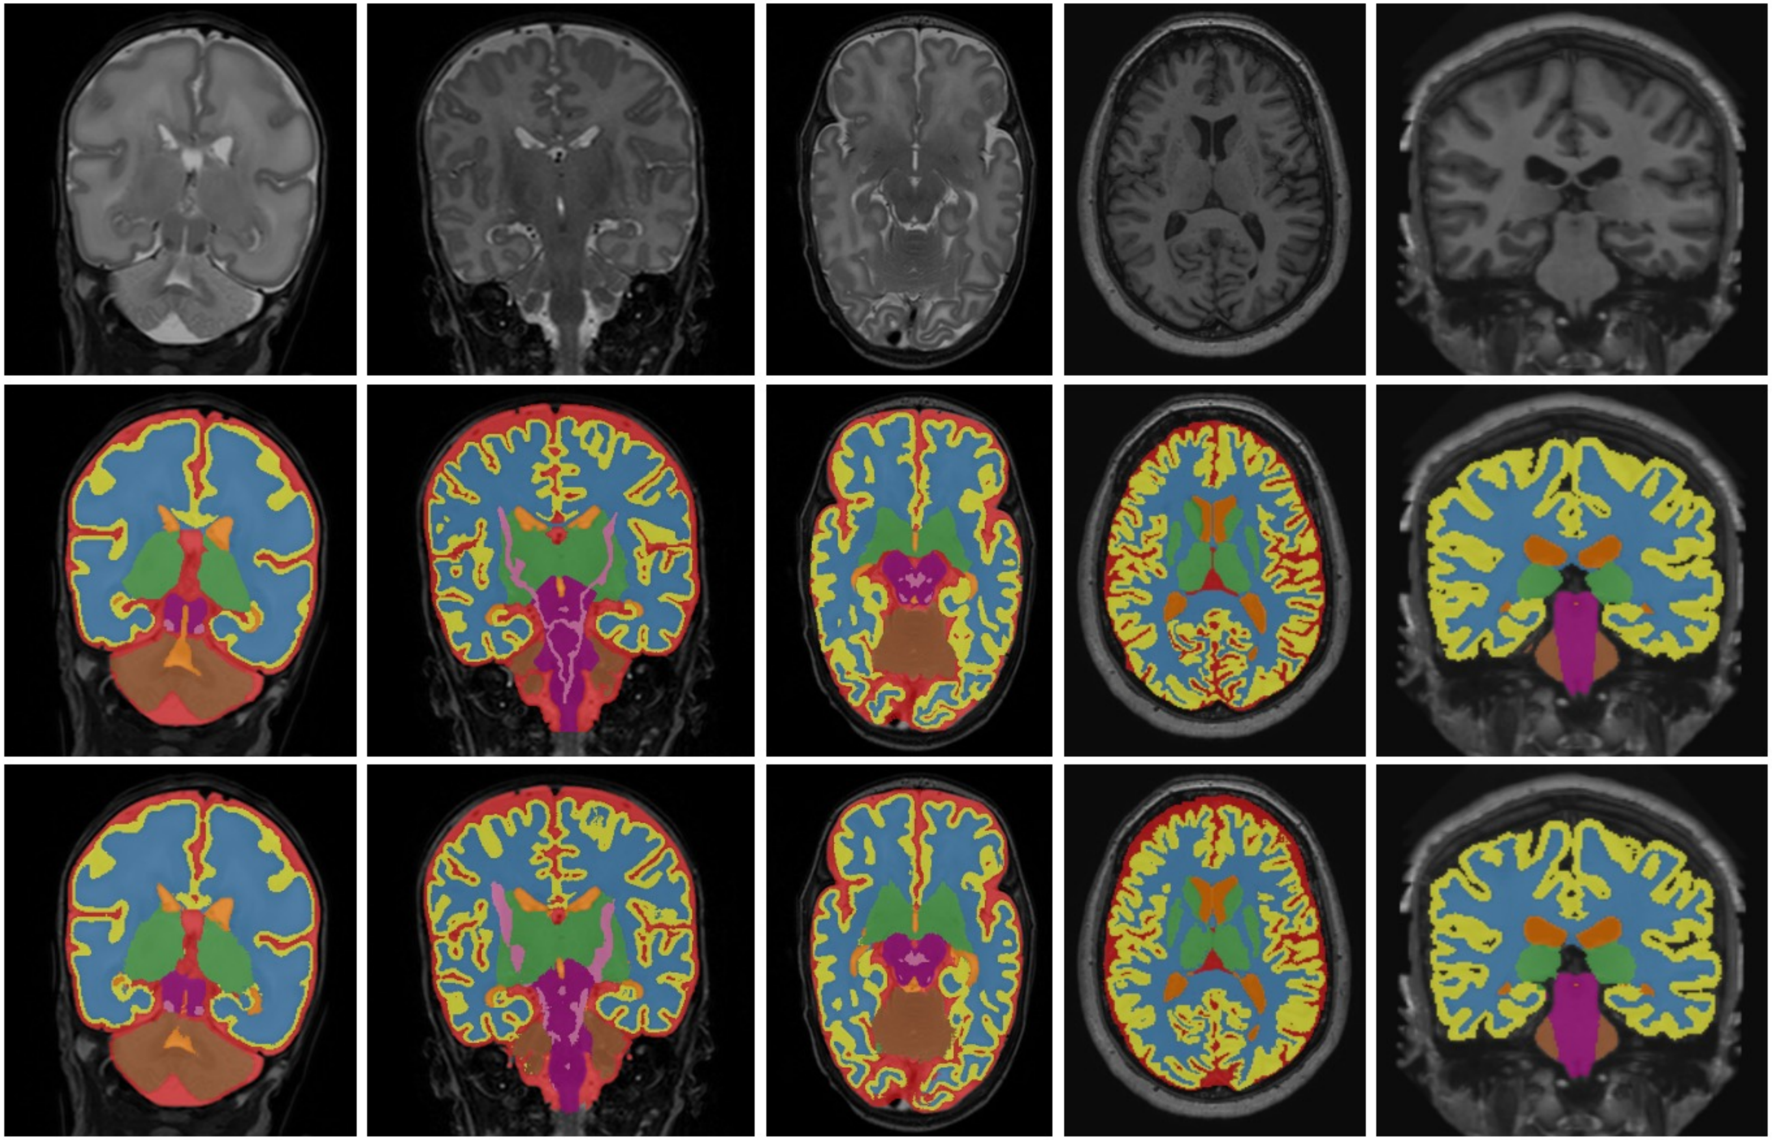

Fig. 4: Segmentation results for CB (brown), mWM (pink), BGT (green), vCSF (orange), (u)WM (blue), BS (purple), cGM (yellow), eCSF (red) in coronal images acquired at 30 weeks PMA (first column), coronal images acquired at 40 weeks PMA (second column), axial images acquired at 40 weeks PMA (third colum), axial images of ageing adults (fourth colum), and

Inside a sulcus of a 30 weeks PMA brain

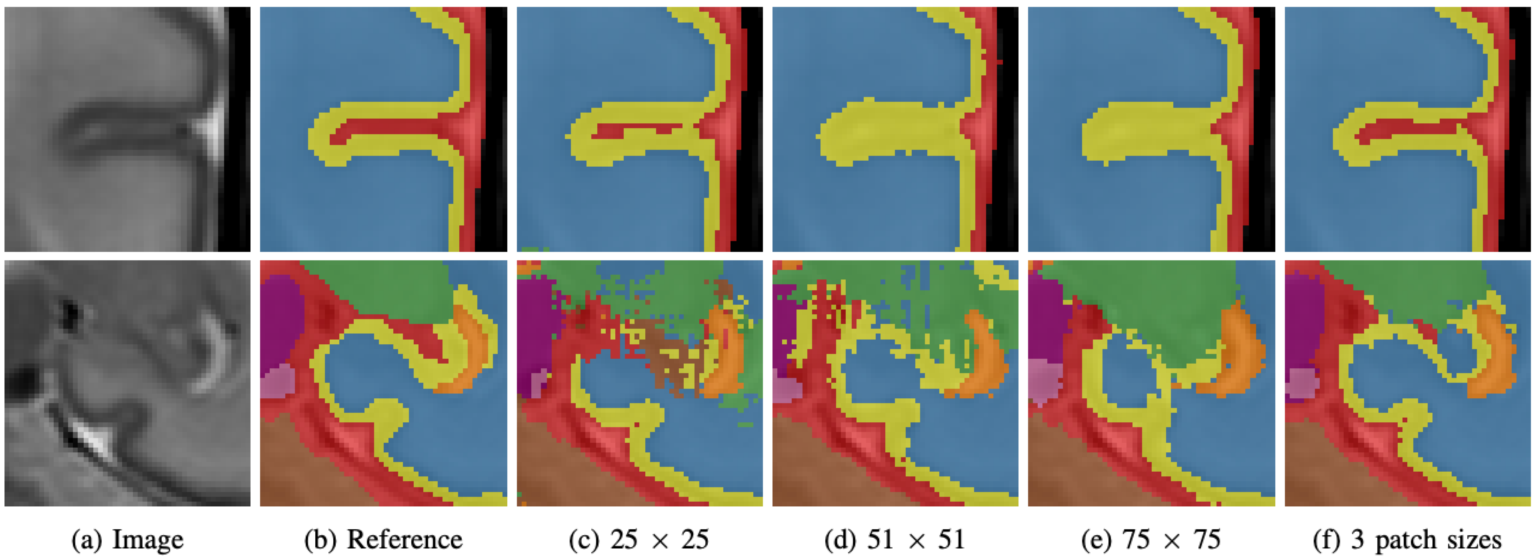

Fig. 5: Segmentation results in a T_2-weighted image (left) acquired at 30 weeks PMA for the lateral sulcus (top) and the hippocampus (bottom) using (from left to right), manual segmentation, only a patch of 25 × 25 voxels, only a patch of 51 × 51 voxels, only a patch of 75 × 75 voxels, and these 3 patch sizes combined. The tissues are labelled as follows: CB in brown, mWM in pink, BGT in green, vCSF in orange, uWM in blue, BS in purple, cGM in yellow, and eCSF in red.

Left ventricle (heart) segmentation for cardiovascular diseases

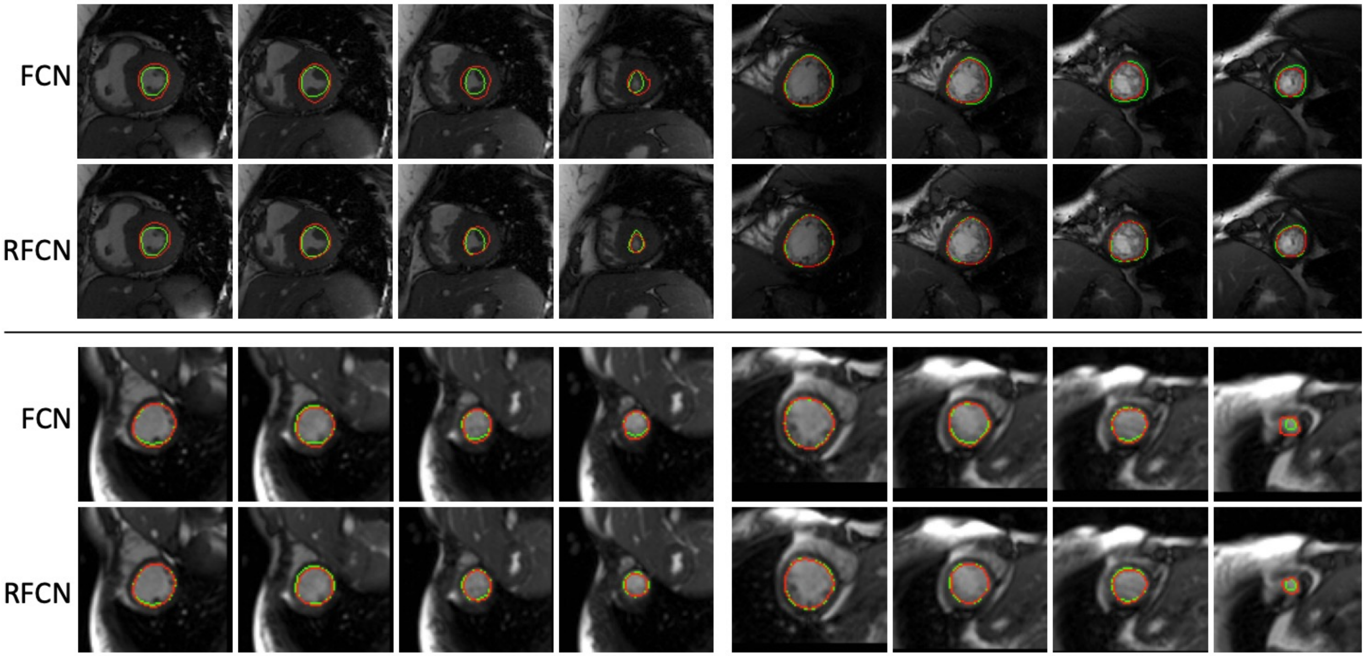

Fig. 3. Example of segmented left ventricle using RFCN and FCN architectures from MICCAI dataset (top two rows) [21] and PRETERM dataset (bottom two rows). Green contours represent the ground truth and red contours are the predicted contours. RFCN is often able to better delineate the left-ventricle contours with weaker boundaries compared to FCN.

1.2- Image registration : example of fMRI

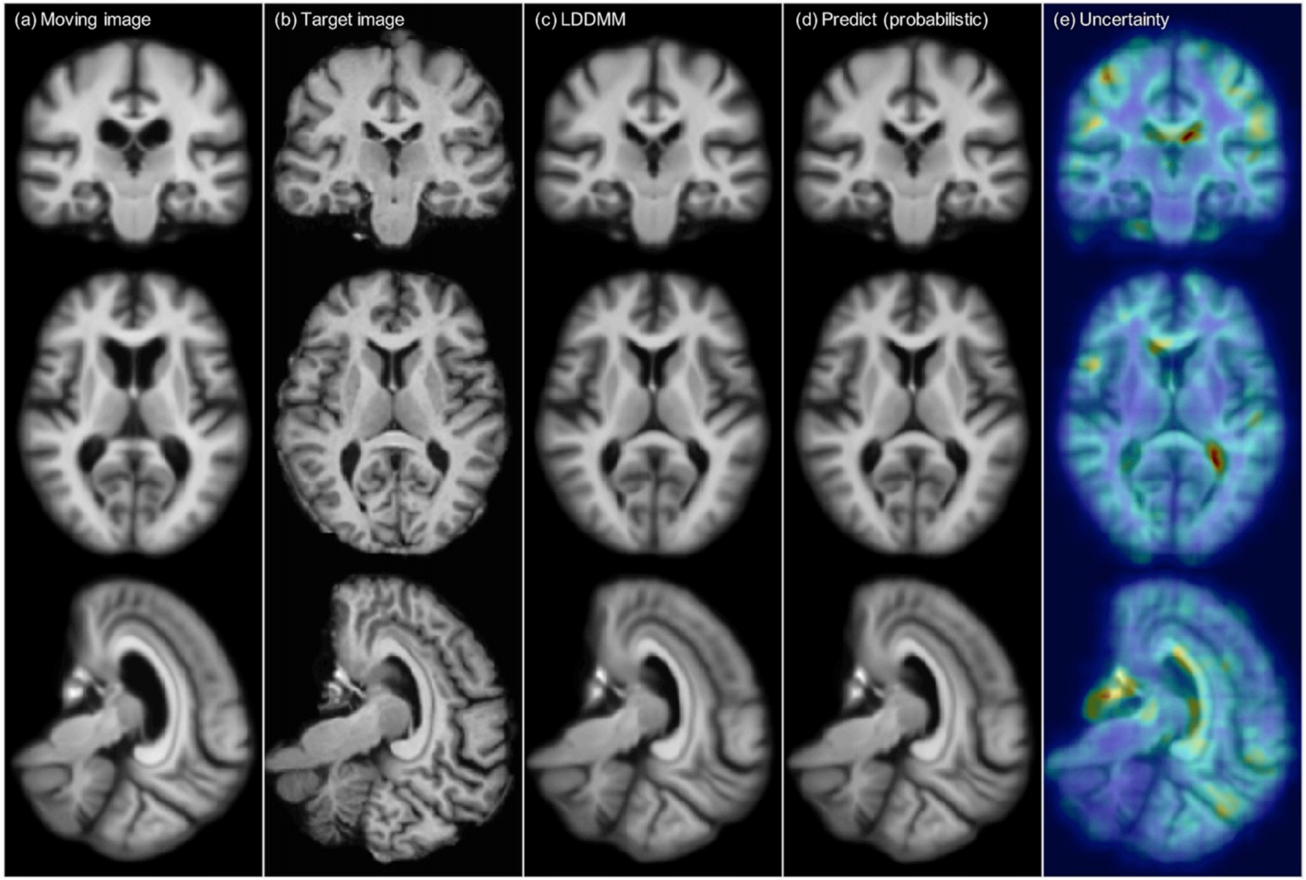

1.3 - DL for classifying medical images

- Predict the category of an image
- Classify images (group similar images)
- Ex. : cancer, melanoma, AD, depression, schizophrenia etc
- From images (anat and functional) + other data

Skin cancer

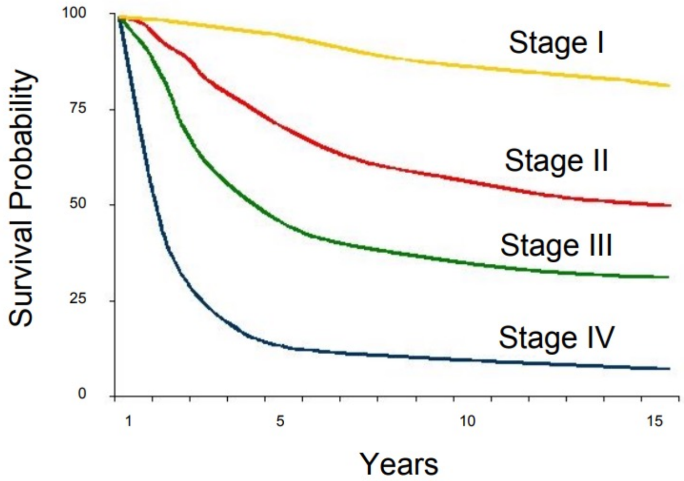
100
75
50
25
0
Survival Probability
Stage I
Stage II
Stage III
Stage IV
1
5
10
15
Years

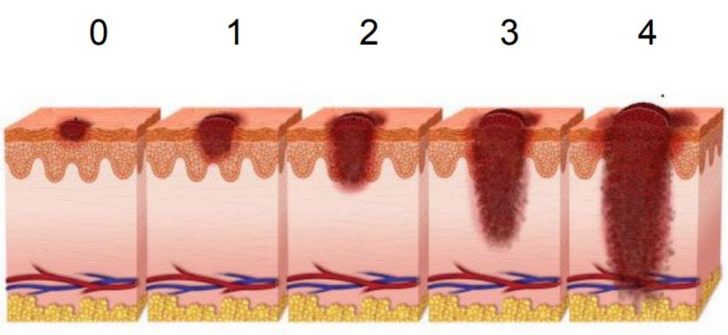
0
1
2
3
4

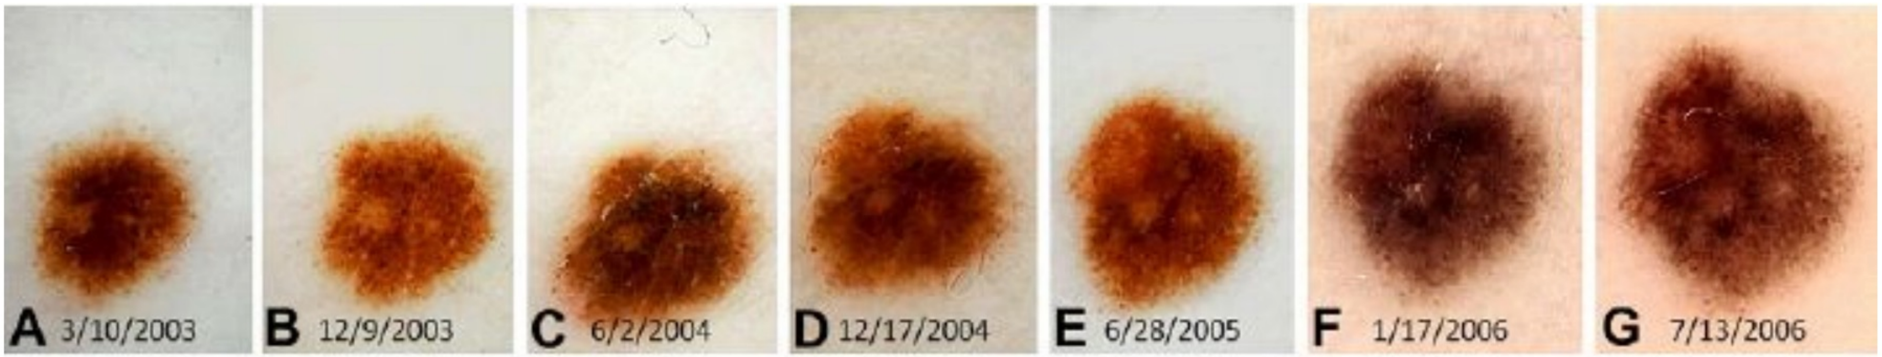
A 3/10/2003
B 12/9/2003
C 6/2/2004
D 12/17/2004
E 6/28/2005
F 1/17/2006
G 7/13/2006

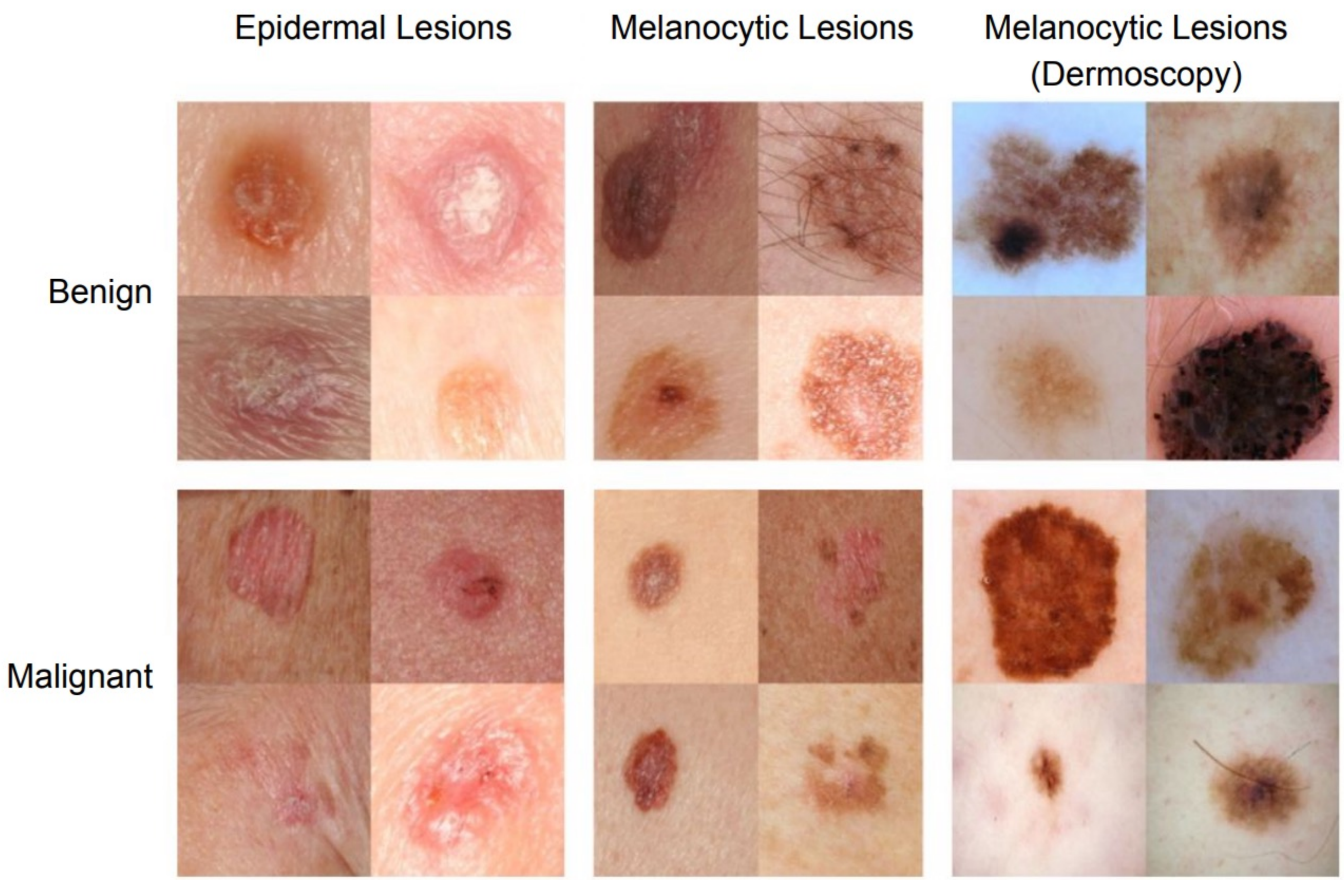
Epidermal Lesions
Melanocytic Lesions
Melanocytic Lesions
(Dermoscopy)
Benign
Malignant

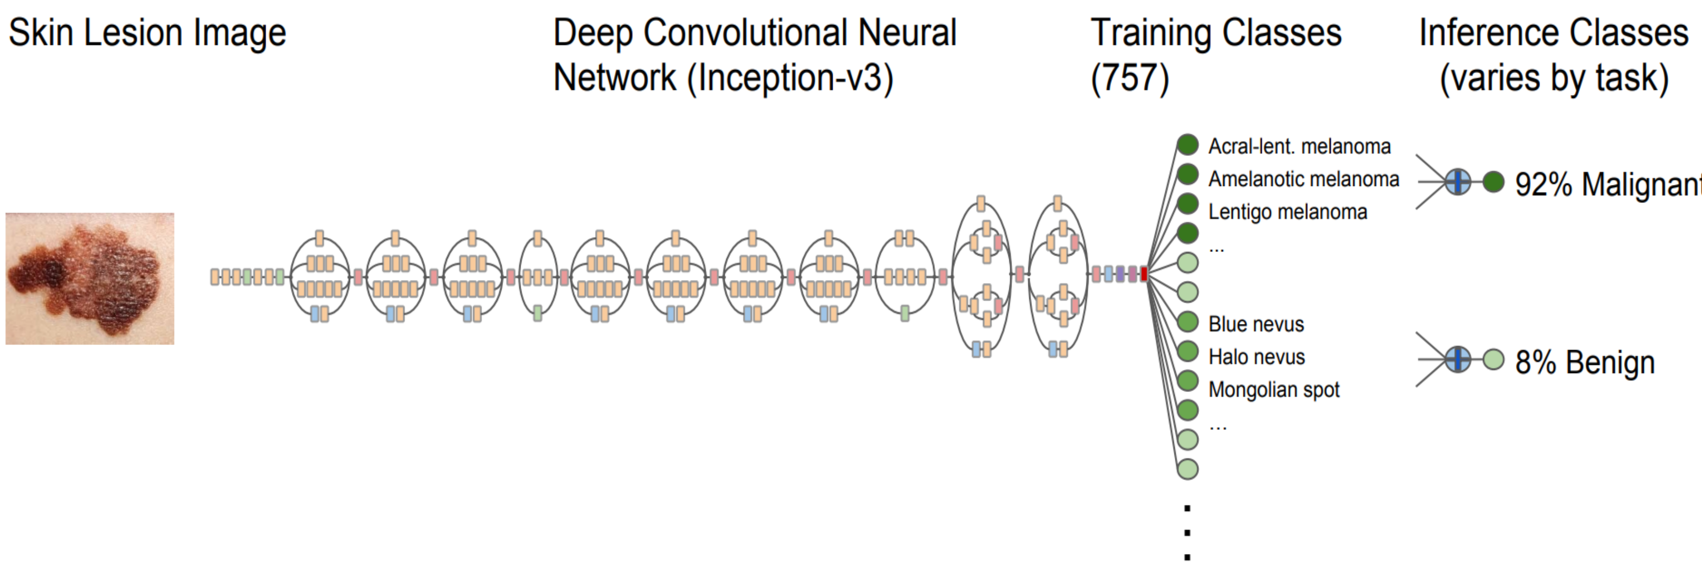
Skin Lesion Image
Deep Convolutional Neural Network (Inception-v3)
Training Classes (757)
Inference Classes (varies by task)
Acral-lent. melanoma
Amelanotic melanoma
Lentigo melanoma
...
Blue nevus
Halo nevus
Mongolian spot
...
92% Malignant
8% Benign

Transfer learning

- Application of a process suited for one specific task to a different problem
    - DL model trained to recognize every day color images, such as animals
    - Same model used to classify radiographs.

- All images share similar features such as edges and blobs
    - The model has learnt to « see »
    - Better than random initialization
    - Re-learning on small medical databases

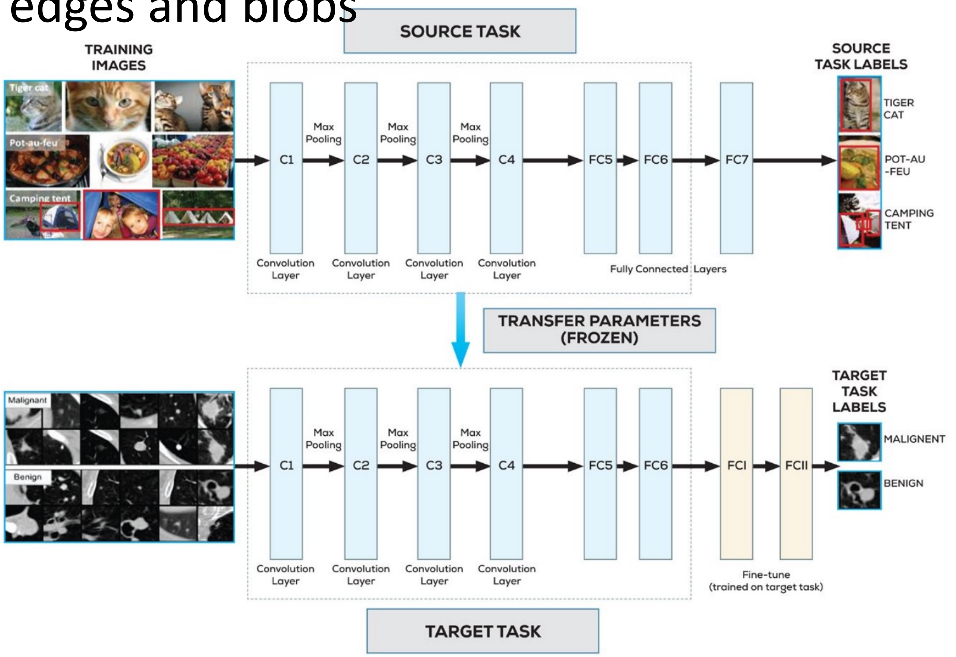

Figure 3. Schematic representation of convolutional neural network (CNN) architecture and the concept of "transfer learning." Because images from multiple sources have common salient features (borders, shapes, etc), the core of a CNN trained for 1 task (*Top row*: eg, cat from pot-au-feu from camping tent) can be "transferred" (ie, used without modification) for a second task (*Bottom row*: benign from malignant tumors in computed tomography images). Because only the distal layers (fully connected [FC] I and FCII) remain to be trained, much less training data are required for the second task. Max indicates maximum.

Skin Cancer Classification

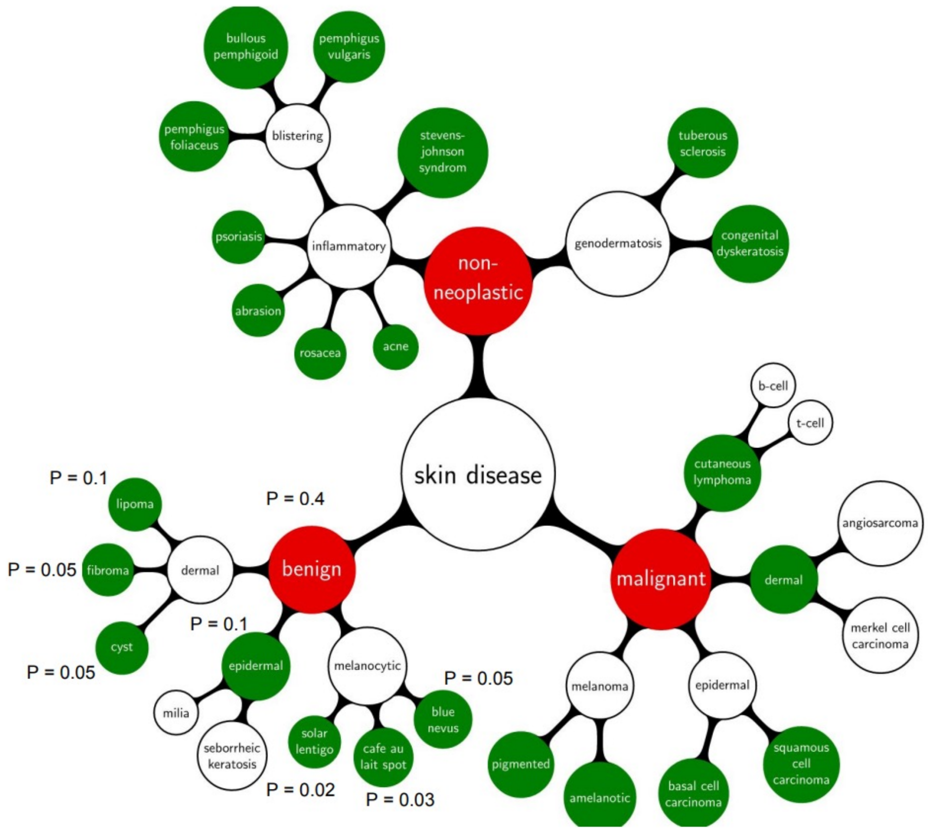

$$P[u] = \sum_{v \in C(u)} P[v]$$

Dermatologist-level Classification of Skin Cancer with Deep Neural Networks

Andre Esteva*, Brett Kuprel*, Rob Novoa, Justin Ko, Susan Swetter, Helen Blau, Sebastian Thrun
Nature, 2017
(Equal contribution authors*)

Classifier	Three-way accuracy
Dermatologist 1	65.6%
Dermatologist 2	66.0%
CNN	69.5%
CNN - PA	**72.0%**

Classifier	Nine-way accuracy
Dermatologist 1	53.3%
Dermatologist 2	55.0%
CNN	48.9%
CNN - PA	**55.3%**

Disease classes: three-way classification

0. Benign single lesions
1. Malignant single lesions
2. Non-neoplastic lesions

Disease classes: nine-way classification

0. Cutaneous lymphoma and lymphoid infiltrates
1. Benign dermal tumors, cysts, sinuses
2. Malignant dermal tumor
3. Benign epidermal tumors, hamartomas, milia, and growths
4. Malignant and premalignant epidermal tumors
5. Genodermatoses and supernumerary growths
6. Inflammatory conditions
7. Benign melanocytic lesions
8. Malignant Melanoma

Signal Detection Theory

- Accuracy (ACC) = (TP + TN)/(TP + TN + FP + FN)
- Sensitivity (SEN) = TP/(TP + FN)
- Specificity (SPEC) = TN/(TN + FP)

	H_1: signal present	H_0: signal absent
Detection	True Positive	False Positive *type I error*
Null result	False Negative *type II error*	True Negative

ROC Curves : sensitivity / specificity

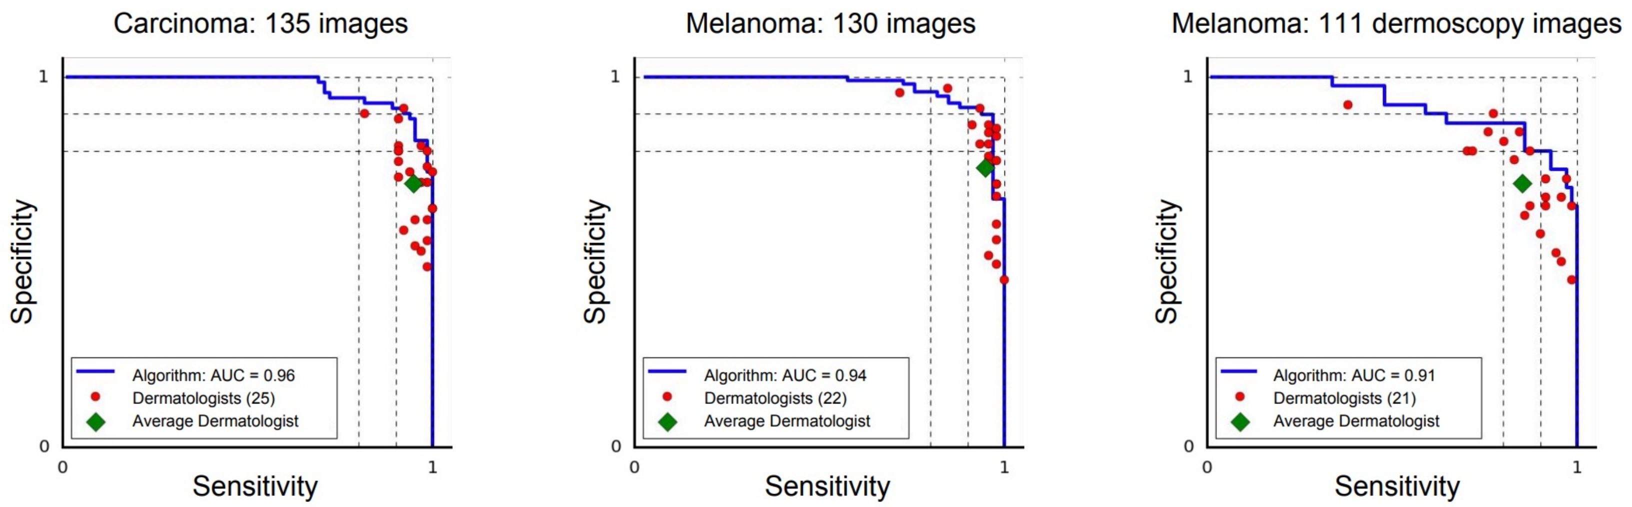

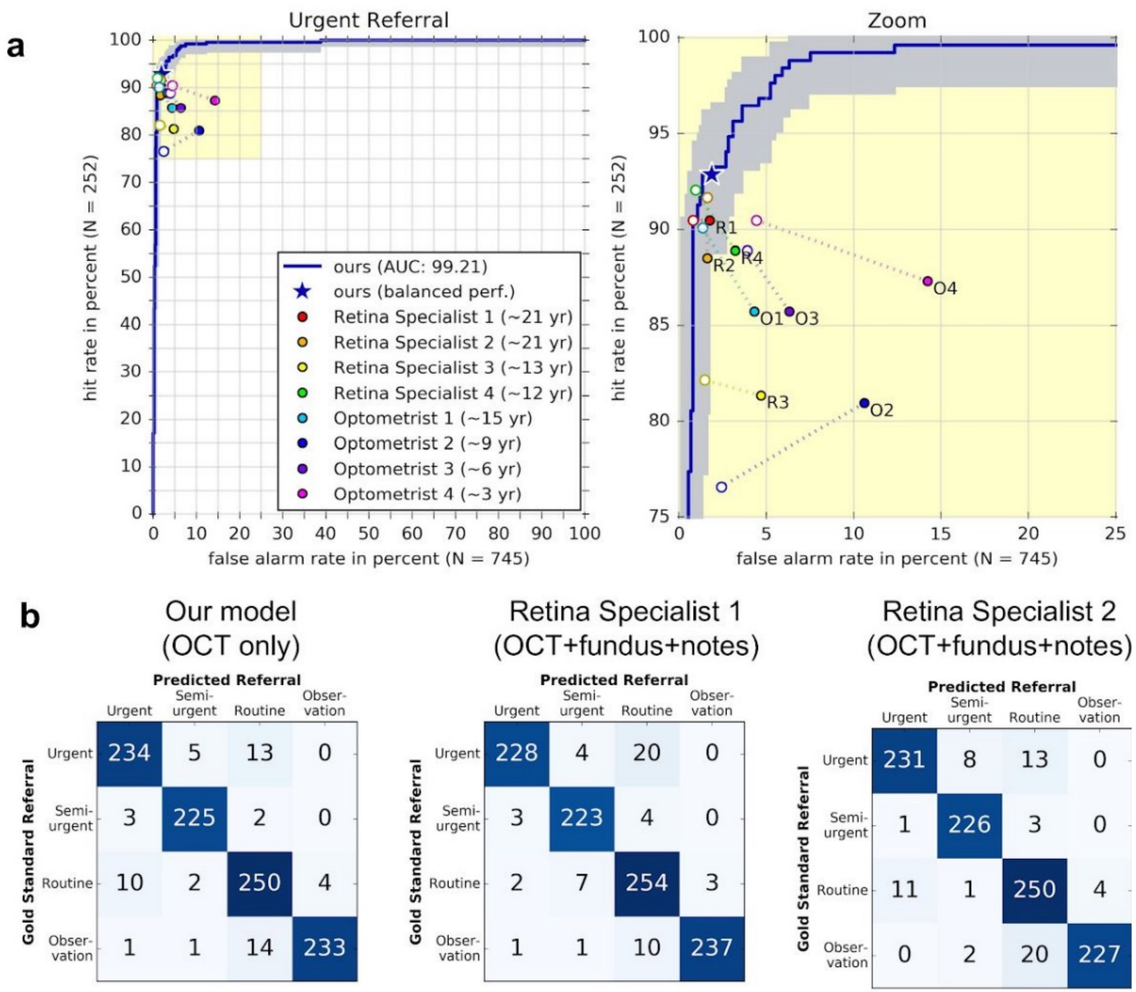
a
Urgent Referral
Zoom
hit rate in percent (N = 252)
false alarm rate in percent (N = 745)
ours (AUC: 99.21)
ours (balanced perf.)
Retina Specialist 1 (~21 yr)
Retina Specialist 2 (~21 yr)
Retina Specialist 3 (~13 yr)
Retina Specialist 4 (~12 yr)
Optometrist 1 (~15 yr)
Optometrist 2 (~9 yr)
Optometrist 3 (~6 yr)
Optometrist 4 (~3 yr)
R1
R2
R3
R4
O1
O2
O3
O4
b
Our model
(OCT only)
Retina Specialist 1
(OCT+fundus+notes)
Retina Specialist 2
(OCT+fundus+notes)
Predicted Referral
Gold Standard Referral
Urgent
Semi-urgent
Routine
Observation
234 5 13 0
3 225 2 0
10 2 250 4
1 1 14 233
228 4 20 0
3 223 4 0
2 7 254 3
1 1 10 237
231 8 13 0
1 226 3 0
11 1 250 4
0 2 20 227

Classifying laserendomicroscopy images of the oral cavity for cancer (Oral Squamous Cell Carcinoma (OSCC)

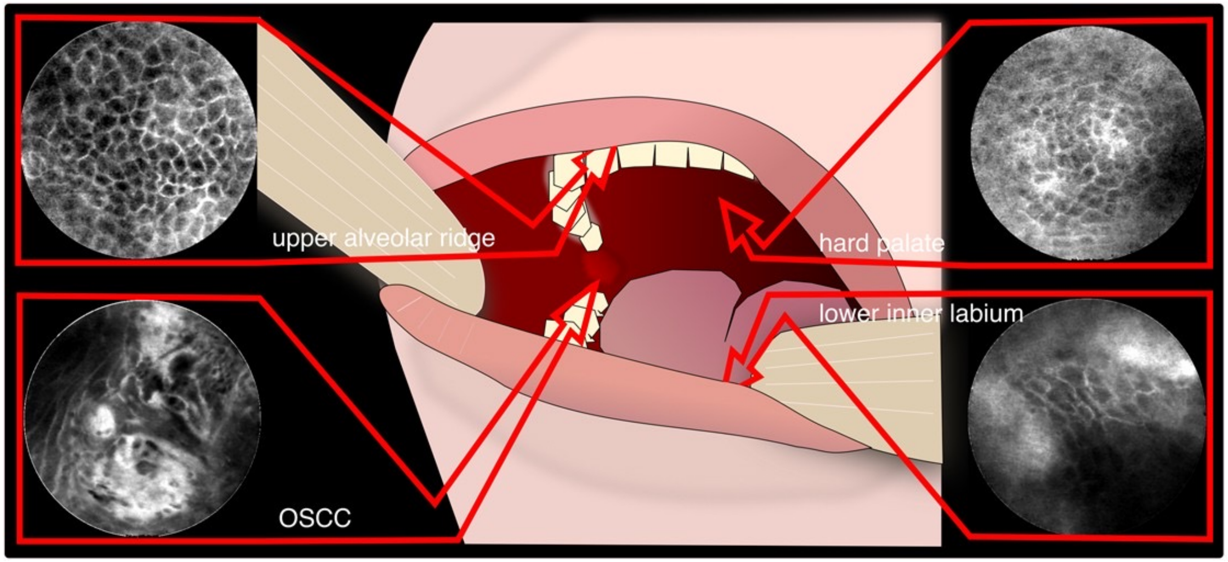

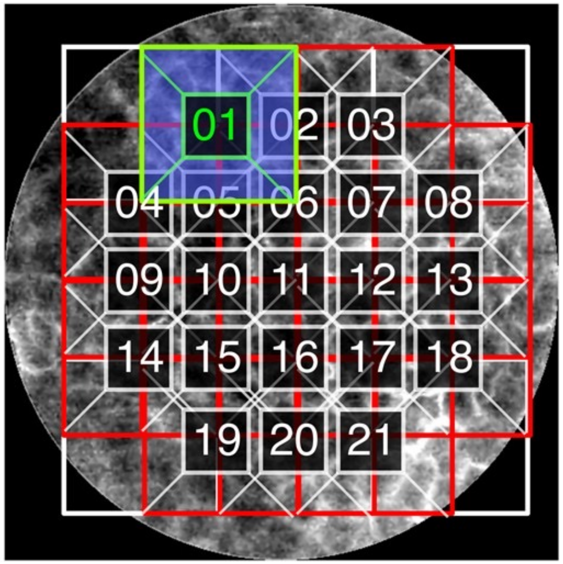

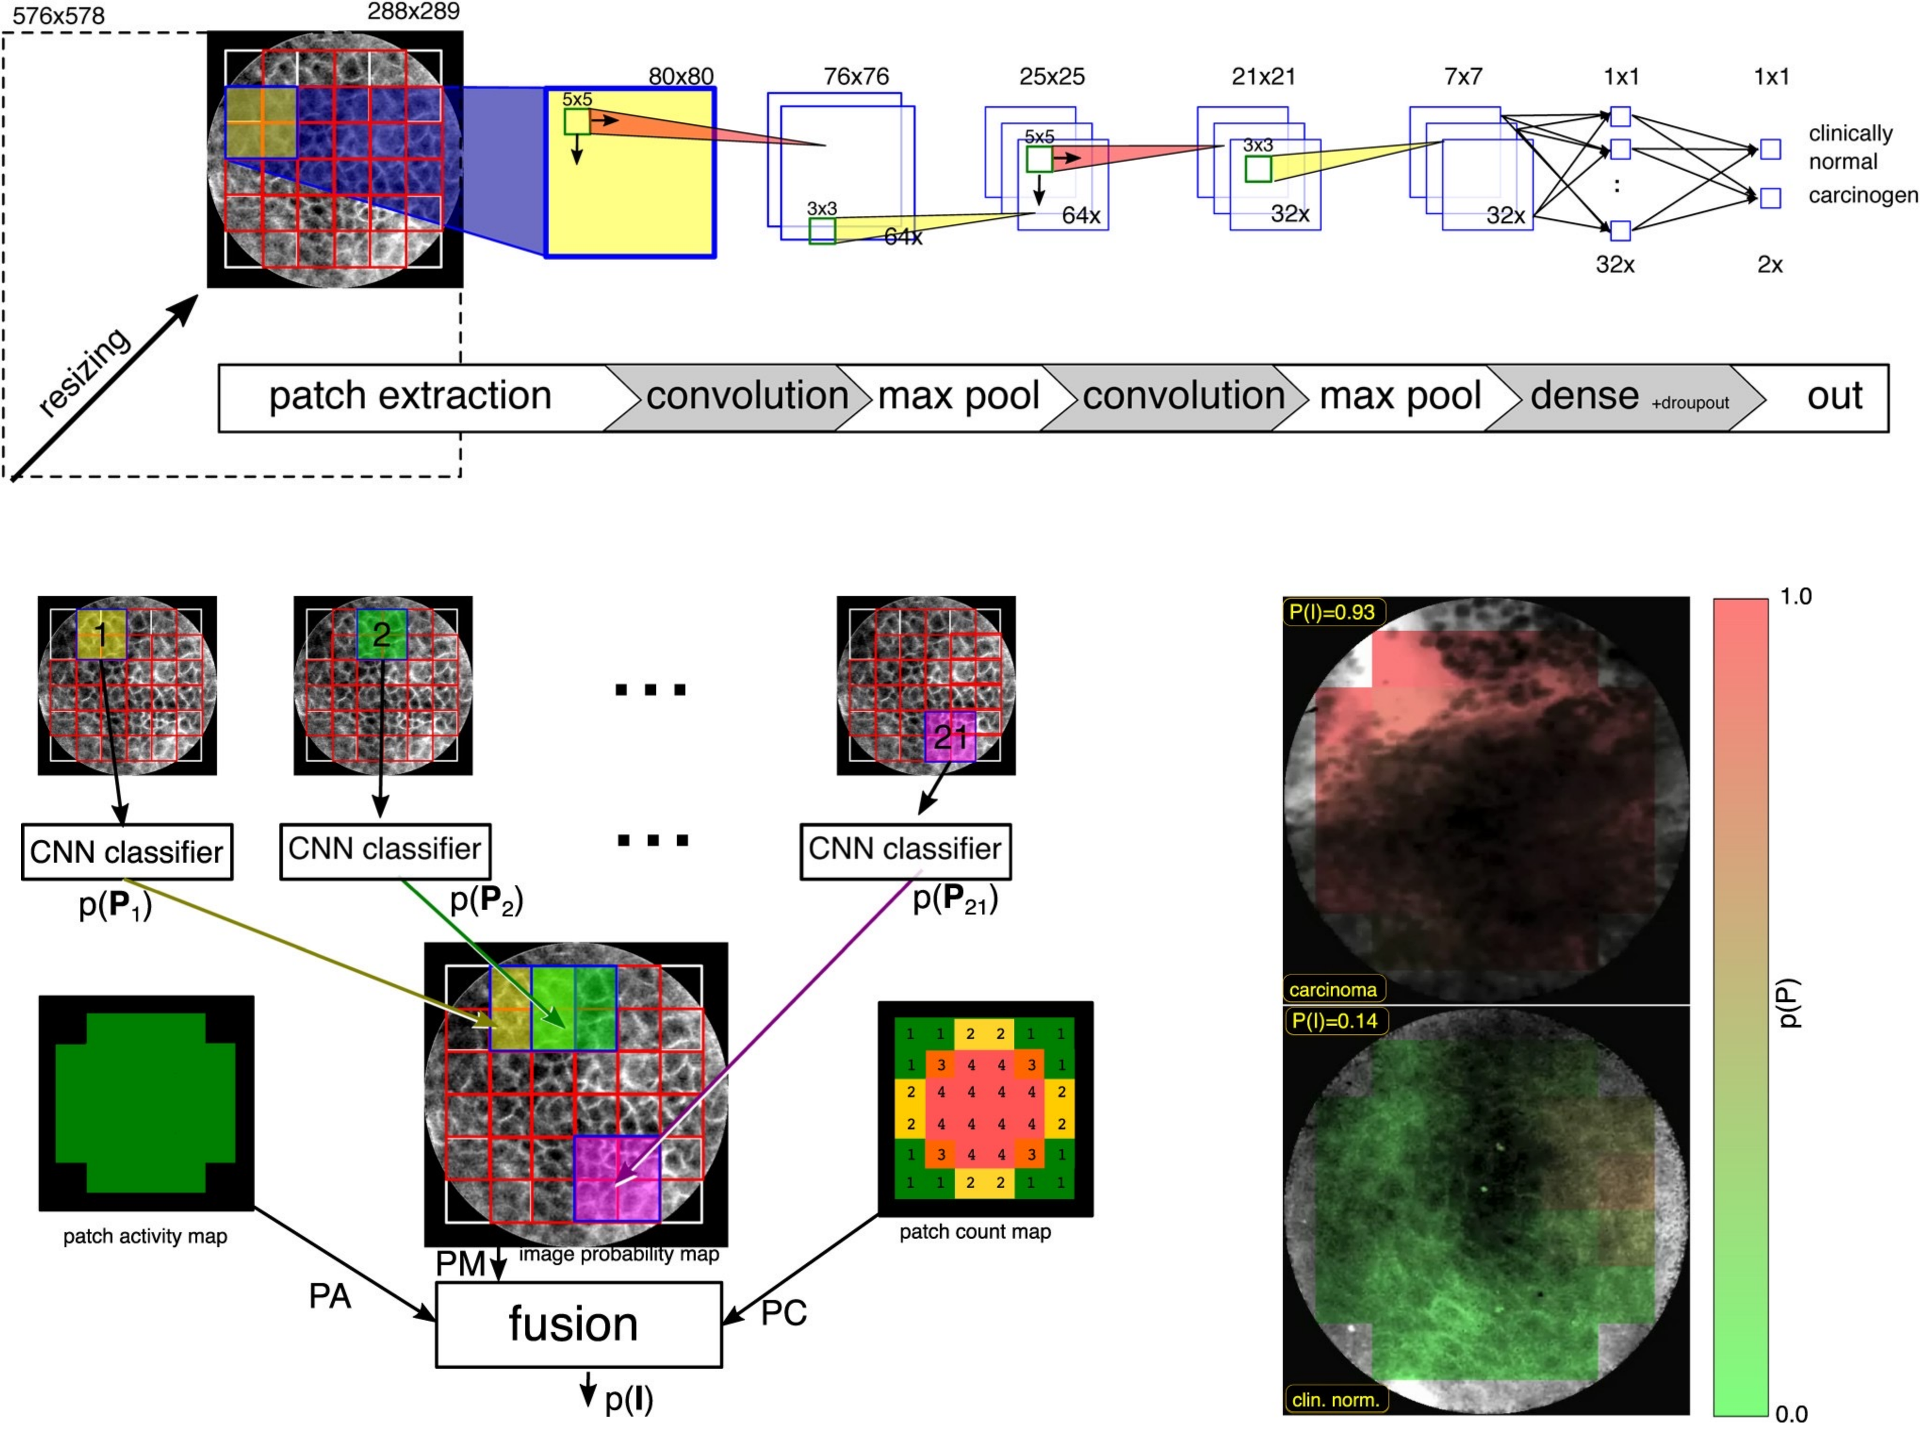

576x578
288x289
resizing
80x80
5x5
76x76
3x3
64x
25x25
5x5
64x
21x21
3x3
32x
7x7
32x
1x1
32x
1x1
2x
clinically normal
carcinogen
patch extraction
convolution
max pool
convolution
max pool
dense +droupout
out
1
2
21
CNN classifier
CNN classifier
CNN classifier
p(P1)
p(P2)
p(P21)
patch activity map
image probability map
patch count map
PM
PA
PC
fusion
p(I)
P(I)=0.93
carcinoma
P(I)=0.14
clin. norm.
1.0
p(P)
0.0

- Patches

images

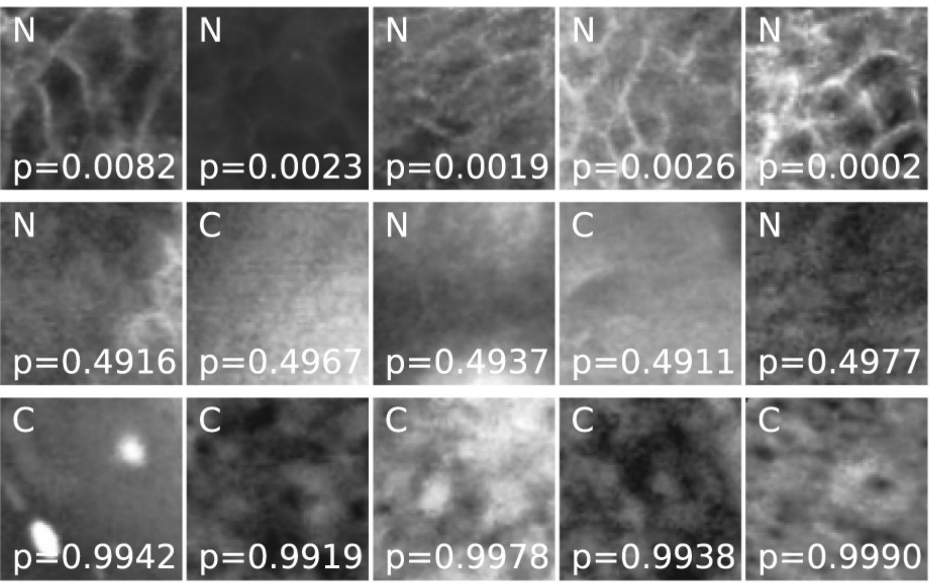

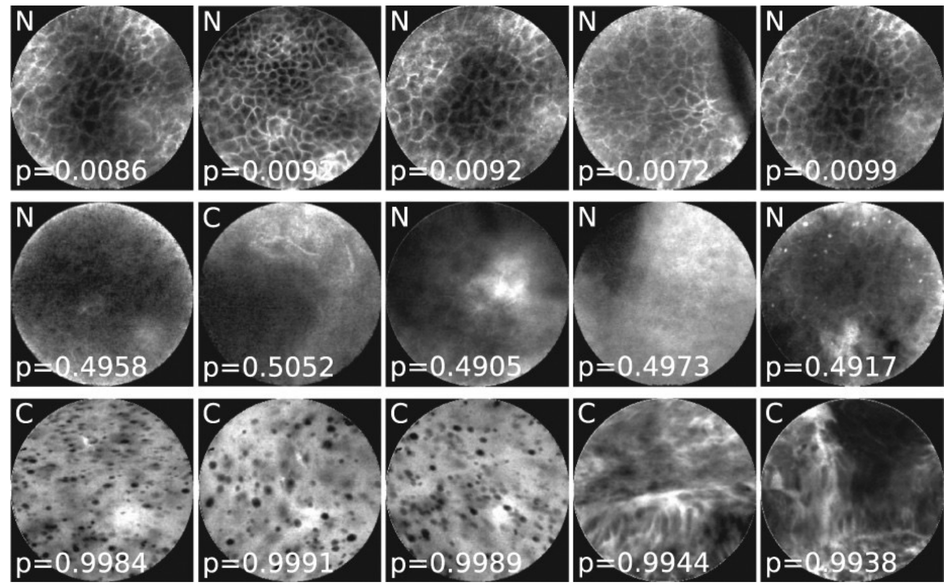

- N = normal, C = cancer

Classifying Mammogram Exams Containing Unregistered Multi-view Images and Segmentation Maps of Lesions

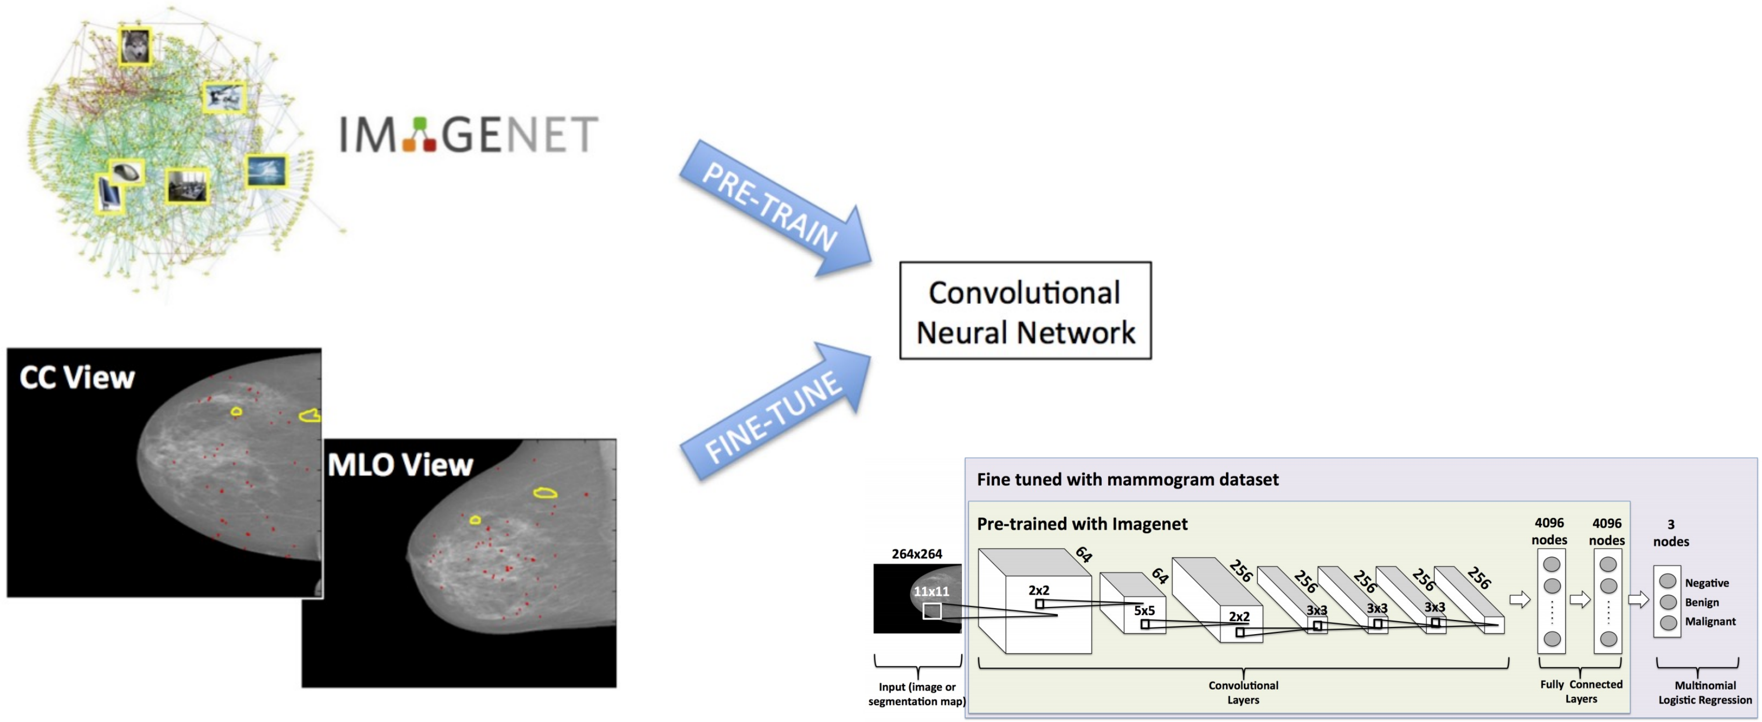

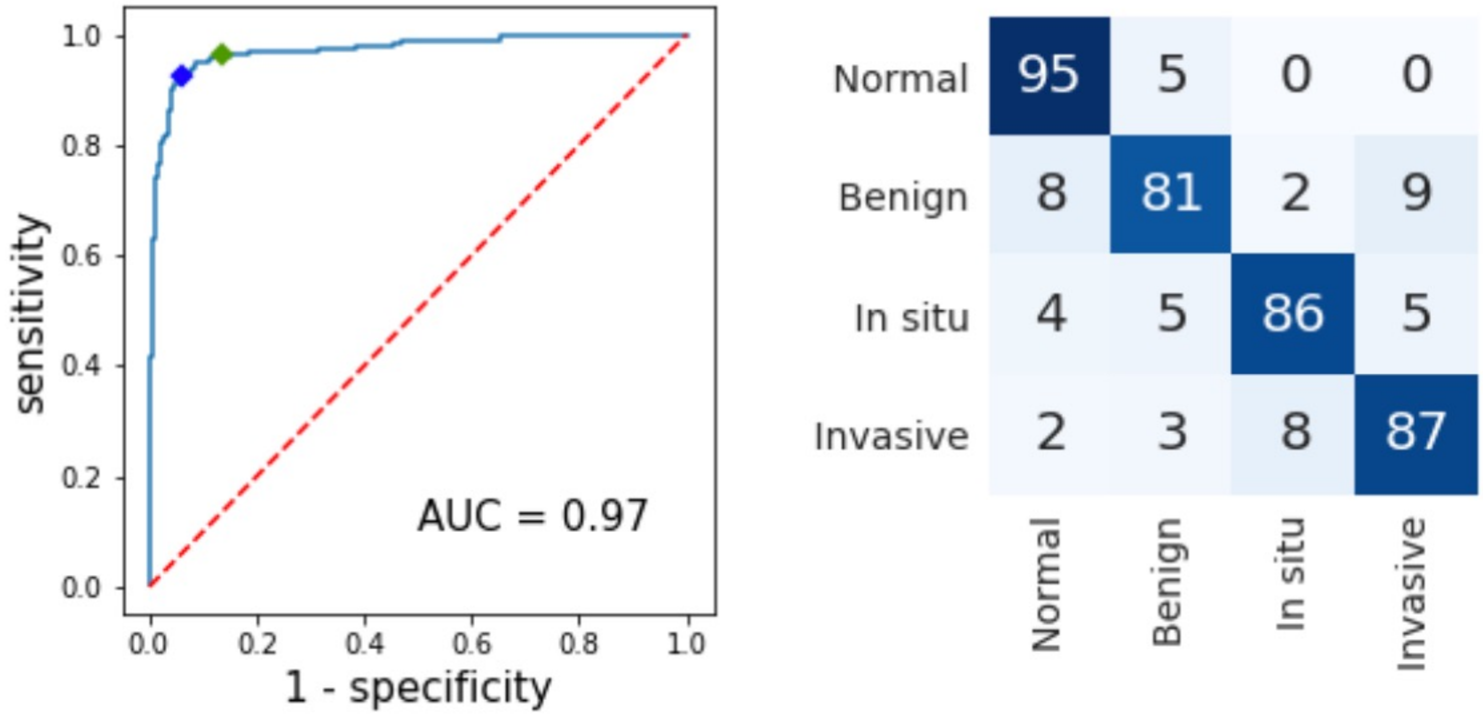

Fig. 4: a) Non-carcinoma vs. carcinoma classification, ROC. High sensitivity setpoint=0.33 (green): 96.5% sensitivity and 88.0% specificity to detect carcinomas. Setpoint=0.50 (blue): 93.0% sensitivity and 94.5% specificity b) Confusion matrix, without normalization. Vertical axis - ground truth, horizontal - predictions.

AD / MCI classification

- When and which MCI patients will develop AD ?
- Identify different progression stages of AD patients based on MRI and PET scans.
- MRI scans of 2146 subjects (803 for training and 1343 for validation) to predict MCI subjects' progression to AD dementia
- Databases :
  - ADNI (http://adni.loni.usc.edu) : MRI scans of 1711 subjects
  - AIBL (www.aibl.csiro.au) : 435 subjects.

A Robust Deep Model for Improved Classification of AD/MCI Patients

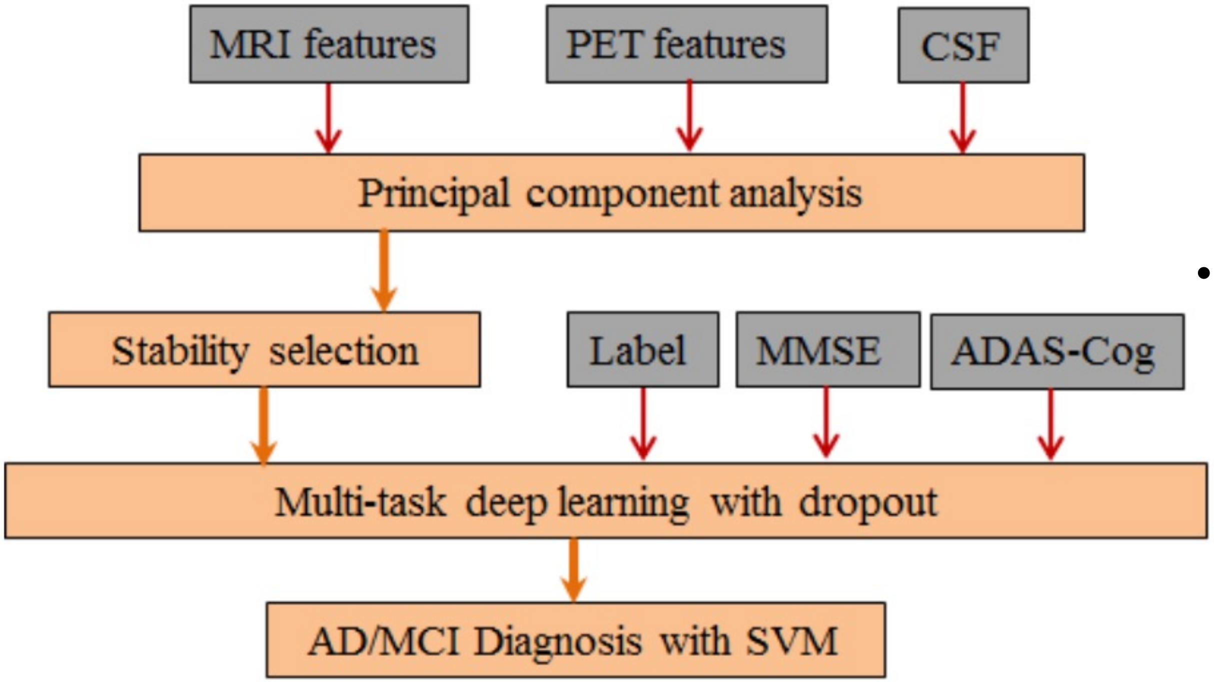

Fig. 1. Diagram of the proposed multi-task deep learning framework.

LEARNING :

- ADNI data set:
  - 51 AD patients,
  - 99 MCI patients
  - (43 MCI patients who conve to AD and 56 MCI patients v did not progress to AD
- 52 healthy normal controls

Minimum Mental State Examination

Alzheimer's Disease Assessment Scale-Cognitive subscale

Results : accuracy of classification

Tasks		Proposed
AD vs HC		**91.4**(1.8)
MCI vs HC		**77.4**(1.7)
AD vs MCI		**70.1**(2.3)
MCI.C vs MCI.NC		57.4(3.6)
Average		**74.1**

AD patients vs Healthy Control subjects (AD vs HC),
MCI patients vs HC (MCI vs HC),
AD patients vs MCI patients (AD vs MCI)
MCI-converted vs MCI-non converted (MCI.C vs MCI.NC).

AD vs Normal from 3D scans

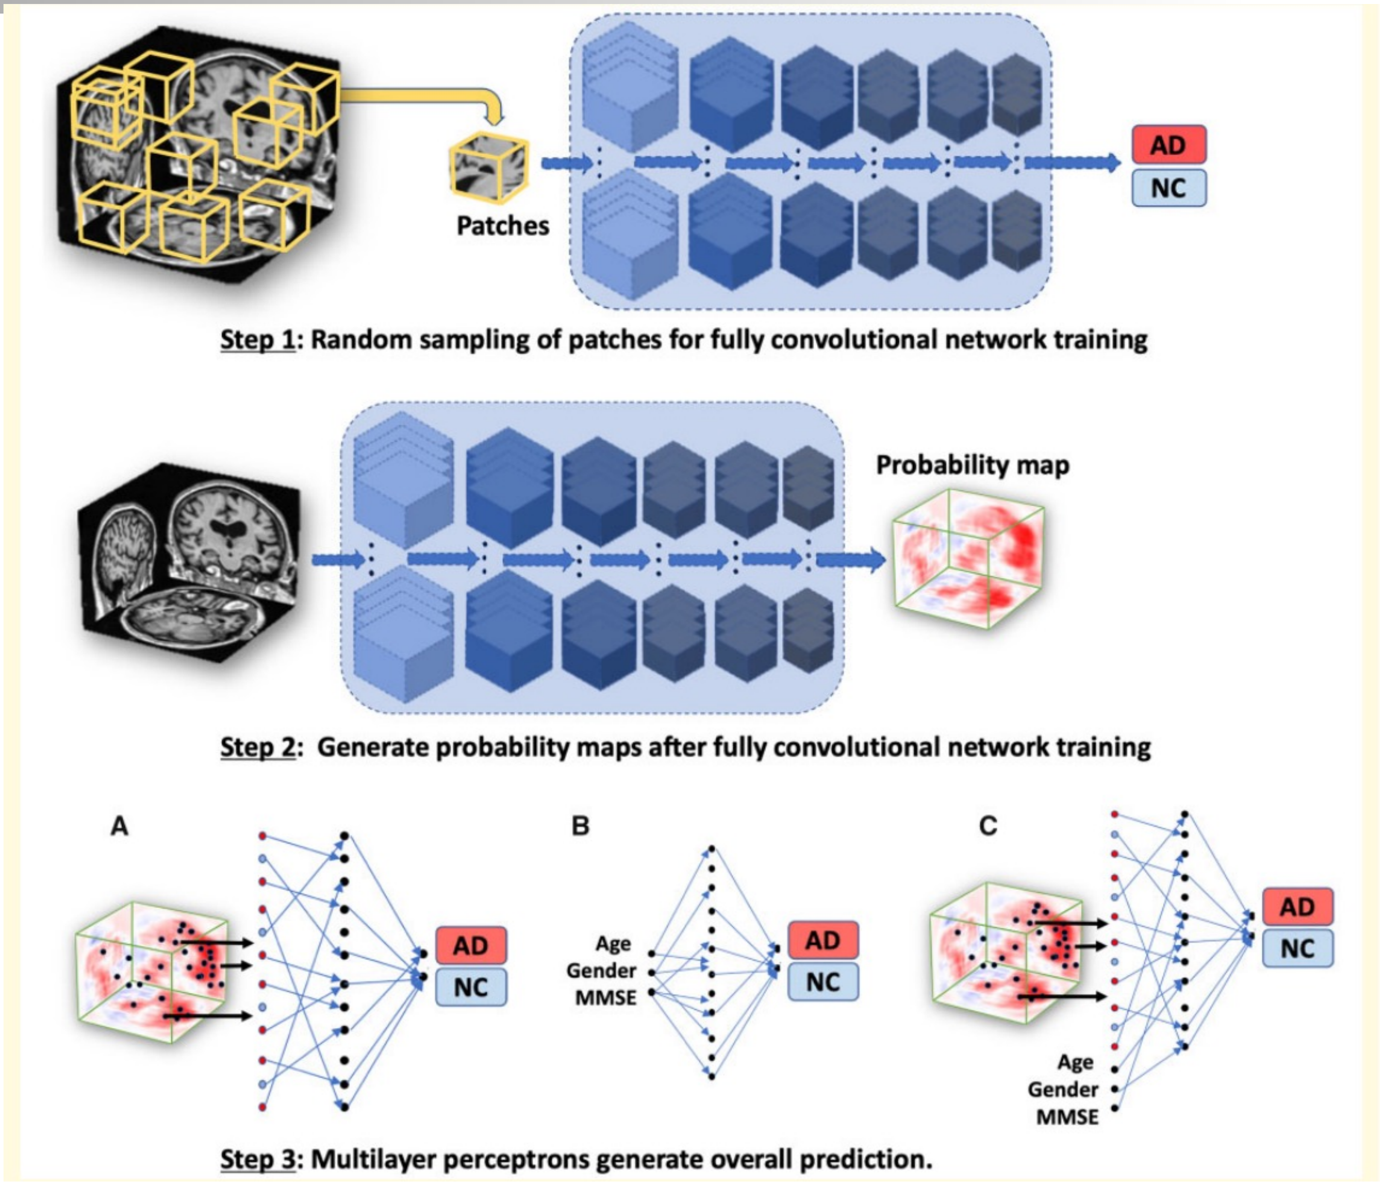

Detection of pneumonia from chest x-rays

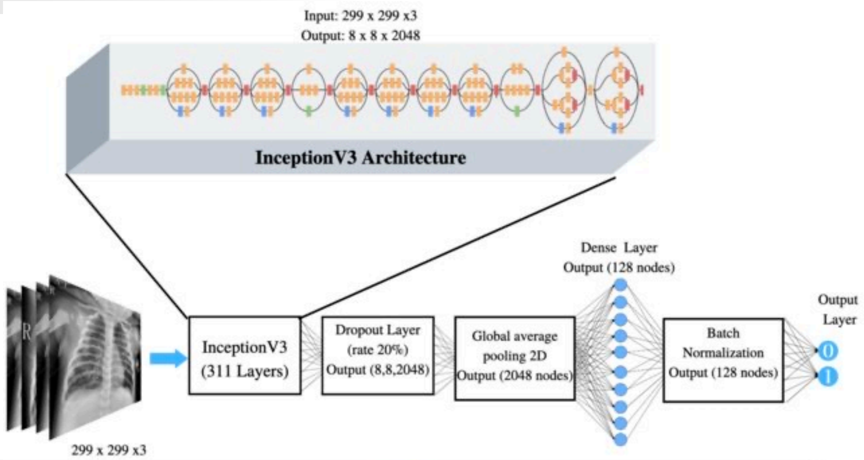

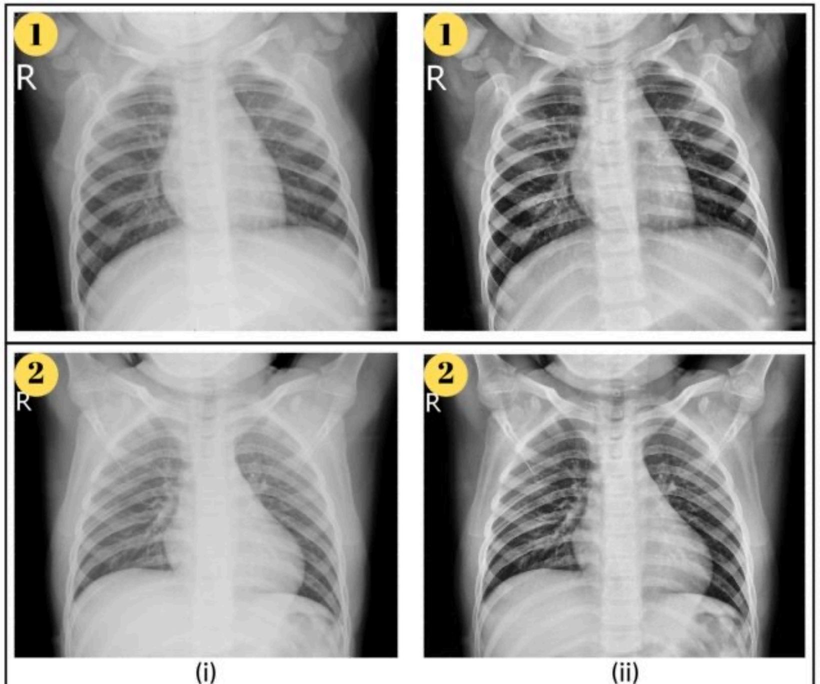

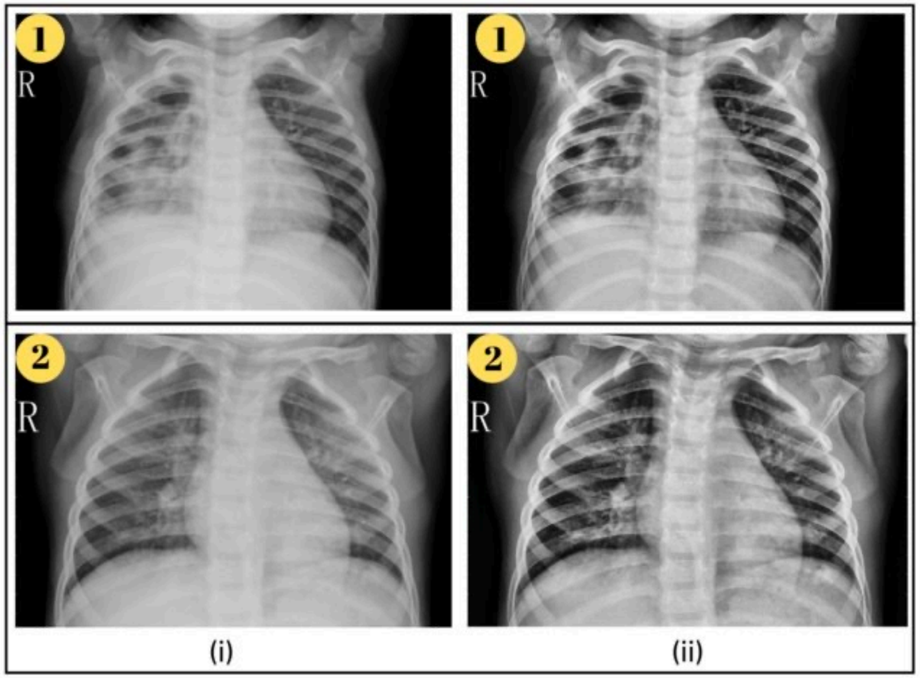

2 – DL and biosignals

- Brain Computer Interface

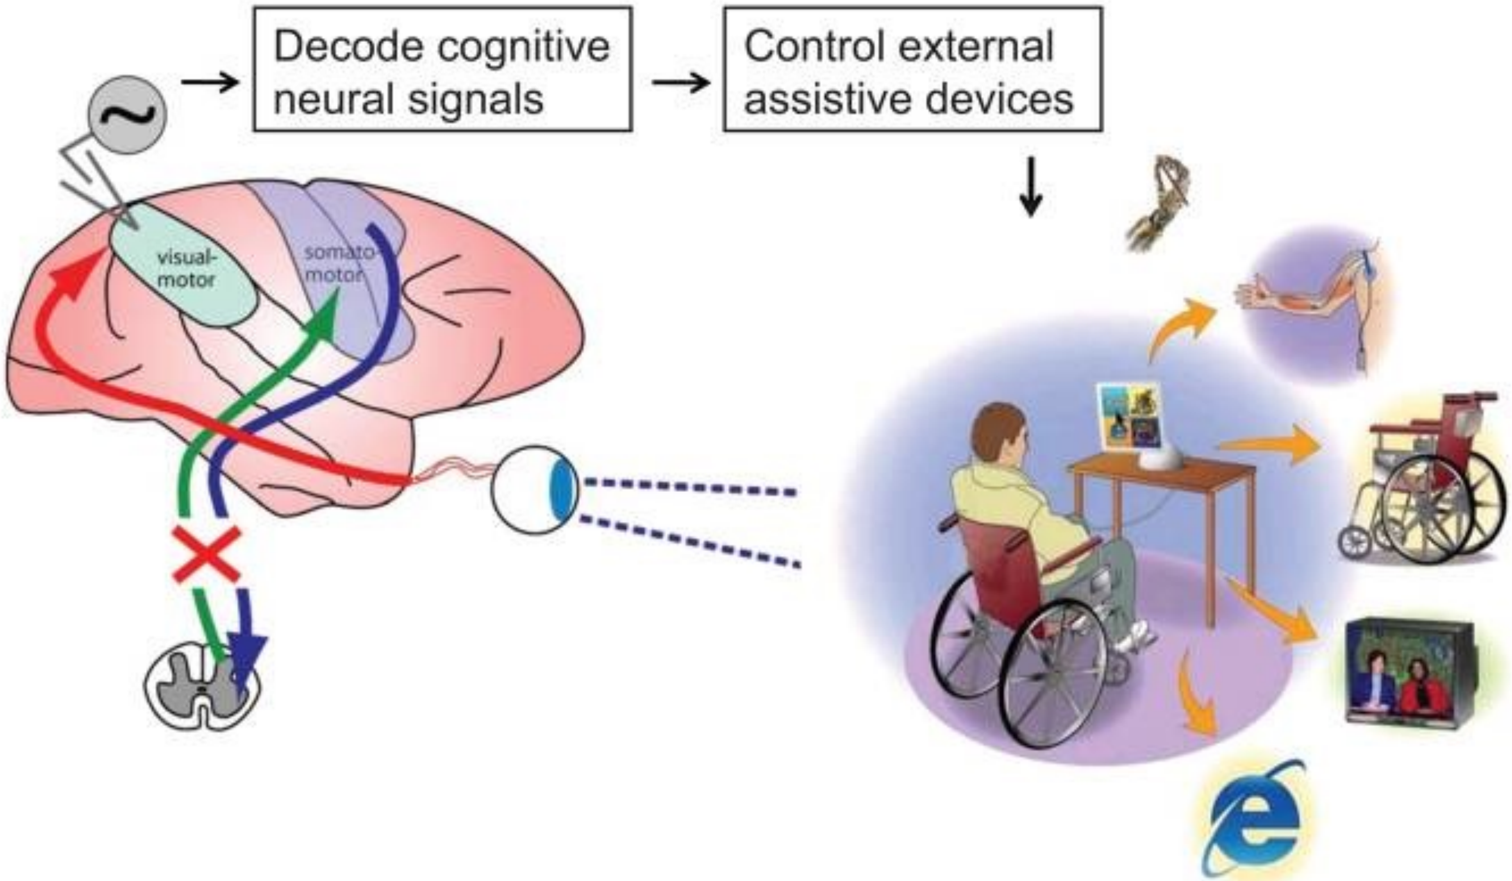

Decode cognitive neural signals
Control external assistive devices
visual-motor
somato-motor

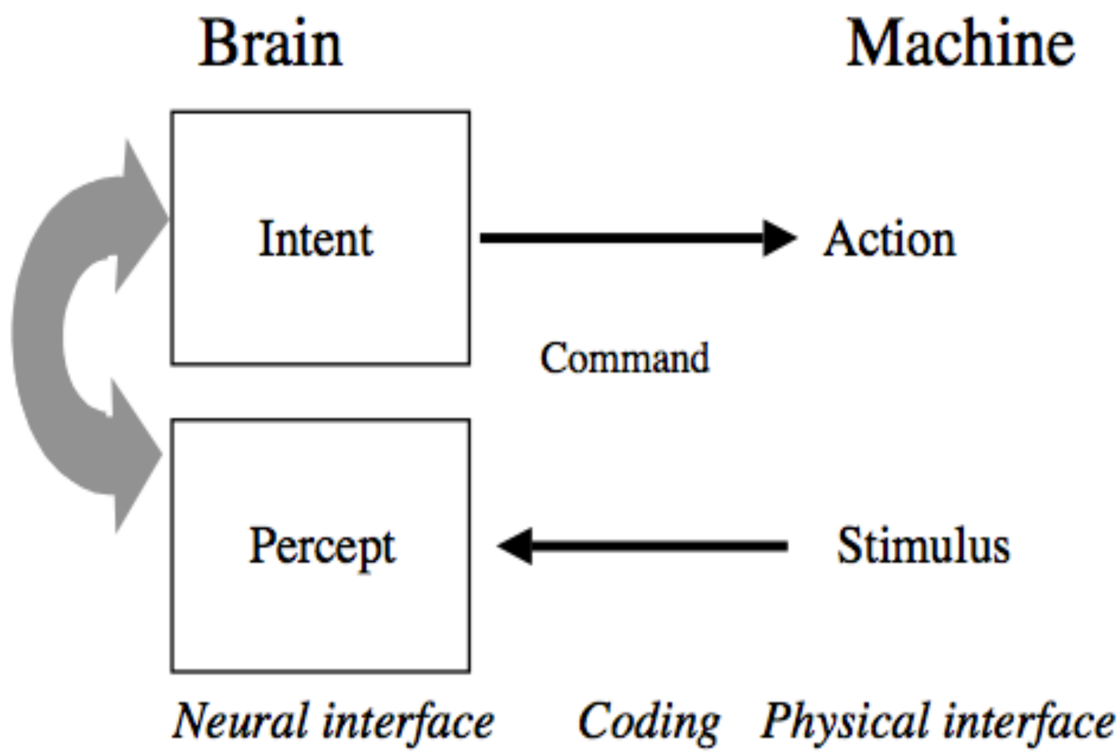
Brain
Machine
Intent
Action
Command
Percept
Stimulus
Neural interface
Coding
Physical interface

BrainGate : Kathy Hutchinson

Table 2

For each medical area this table lists the DL tasks, the obtained contributions and the main used DL models.

Area	Task	Contribution	Model
ECG	Arrhythmia	Classification and detection of arrhythmia levels	RNN [198], LSTM [200], CNN [218], CNN+RNN [21
	Glucose	Evaluation of low glucose levels	CNN [204]
	Ventricular	Data relative to different races	CNN [201]
	Atrial	Choice of low-dimensional datasets	CNN [199]
		Use of ECG with different lengths	CNN [220]
	Foetal	Restriction of data only one ECG channel	AE [221]
	Apnoea	Classification of obstructive sleep apnoea and hypopnoea	CNN [203]
EEG	Epilepsy	Multimodal approach	CNN+LSTM [222]
		Increase in the SNR	CNN [223]
	Emotions	Improvement in performance in heterogeneous samples	DL [222].
PCG	Heart abnormalities	Optimal identification of heart errors	CNN [208]
PPG	Blood analysis	Use of signals correlations	LSTM [211], CNN [212,213]
EMG	Muscle conditions	Improvements in signals accuracy	CNN [214], AE [215]
NNS, SS	Biosignals analysis	Accurate analysis of biosignals	CNN [216,217]

3 – EHR and DL

- CBR , Content-Based Retrieval : technique for knowledge discovery in massive databases (« Big Data »)

- Generating reports from images

- Finding similar cases in previous records

- COVID + connected watches

DL and Electronic health records

- EHR are growing
- 10 million patients over a decade
- A single hospitalisation = 150000 pieces of data

- Then : understand questions in natural language such as : what is this patient’s problem list ?

- Help to transcript patient visits (automatic speech recognition)

Table 4
This table lists the medical tasks, the contributions provided and the DL models most frequently used in EHR.

Area	Task	Contribution	Model
EHR	Disease prediction	Integration of different medical data in cancer analysis	CNN [263]
		High accuracy in dyslipidemia prediction	LSTM [261]
		Improvement in the imbalance problem in heart failure	CNN [264]
		Impact of external databases in sleep staging evaluation	CNN [265]
		Mortality risk estimation by using patients' historical information	LSTM [262]
		Optimal heart failure prediction	LSTM [51,259]
		Integration of structured and unstructured data for the prediction of acute kidney injury subtypes	LSTM [260]
	Risk analysis	Organization of the data on the basis of semantic spheres	GCNN [266]
		Feature extraction by reducing the data dimensions	AE [268]
	Treatment plans	Construction of a treatment plans by using small population sets	LSTM [267]
		Estimation of a treatment plan by exploiting data correlations	CNN [6]

4 – Other applications

Deep learning: new computational modelling techniques for genomics

- As a data-driven science, genomics can use machine learning to capture dependencies in data and derive novel hypotheses.

- The ability to extract new insights from the exponentially growing volume of data requires more expressive machine learning models.

- DL : used for, for example, predicting the impact of genetic variation on gene regulatory mechanisms such as DNA accessibility and splicing.

Table 3
For each medical area this table lists the medical tasks, the contributions obtained and the DL model most frequently used.

Area	Task	Contribution	Model
Genomics	DNA structure	Prediction of DNA missing values from dependences	RNN [230]
		Identification of DNA regions by exploiting the spatial configuration	GCNN [230]
		Simplification of DNA expressions by reducing noisiness	AE+RNN [232], AE [231]
	Disease prediction	High accuracy in the identification of sub-kinds of tumour; Personalized Treatments	Deep Triage [44]
		Improvement in cancer prediction in very sparse molecules	SAE [231]
		Parallel extraction of features from pure DNA expressions	CNN+RNN [234]
		Data integration	DAE [233], AE [232]
Transcriptomics	RNA structure	Reduction of data dimensionality and sparsity	AE [235,237]
		Exploitation of the spatial configuration of RNA molecules	CNN [235]
		Accurate classification of the RNA components	CNN [249], RNN [236]
	Disease prediction	Classification of tumour types	CNN [238]
		RNA variation analysis	CNN [239]
		Heterogeneous data integration	SAE [240]
	Drug discovery	New drug–target interaction identification	DL [250]
Proteomics	Protein structure	Molecular region identification	CNN [251], DeepGSH [252]
		Protein identification	CNN+LSTM [253], AE [254]
	Drug discovery	drug–target interactions	GAN [242]
		Scoring function construction	CNN [241]
Metabolomics	Diseases prediction	Improvement of prediction models	CNN [244]
	Drug discovery	Determination of optimal targets; decrease in drug toxicity	AE [246,255]
		Optimal molecular interactions	CNN [247]

Table 4
This table lists the medical tasks, the contributions provided and the DL models most frequently used in EHR.

Area	Task	Contribution	Model
EHR	Disease prediction	Integration of different medical data in cancer analysis	CNN [263]
		High accuracy in dyslipidemia prediction	LSTM [261]
		Improvement in the imbalance problem in heart failure	CNN [264]
		Impact of external databases in sleep staging evaluation	CNN [265]
		Mortality risk estimation by using patients' historical information	LSTM [262]
		Optimal heart failure prediction	LSTM [51,259]
		Integration of structured and unstructured data for the prediction of acute kidney injury subtypes	LSTM [260]
	Risk analysis	Organization of the data on the basis of semantic spheres	GCNN [266]
		Feature extraction by reducing the data dimensions	AE [268]
	Treatment plans	Construction of a treatment plans by using small population sets	LSTM [267]
		Estimation of a treatment plan by exploiting data correlations	CNN [6]

Some crucial points

- quality and size of the dataset :
  - unavailability of dataset
  - Annotations take time and agreement between experts
  - rare diseases are underrepresented in the data sets.
- Privacy and Legal Issue
  - share the medical data ?
  - Anonymisation, but stil…
- Data Interoperability and Data Standards
  - nature of data differ from hardware to hardware
  - combine several dif-ferent datasets for better algorithms learning and accuracy. I
  - Health data should be standardized and shared between providers
- Black Box and Deep Learning

Neurocomputing – July 2020

Deep Learning for Medical Image Analysis

Heart sounds classification using a novel 1-D convolutional neural network with extremely low parameter consumption
Lung adenocarcinoma diagnosis in one stage
Robust brain extraction tool for CT head images
Multi-label transfer learning for the early diagnosis of breast cancer
Automated hepatobiliary toxicity prediction after liver stereotactic body radiation therapy with deep learning-based portal vein segmentation
Analysis of tuberculosis severity levels from CT pulmonary images based on enhanced residual deep learning architecture
A fully convolutional network feature descriptor: Application to left ventricle motion estimation based on graph matching in short-axis MRI
A framework for hierarchical division of retinal vascular networks
Super-resolution reconstruction of single anisotropic 3D MR images using residual convolutional neural network
CcNet: A cross-connected convolutional network for segmenting retinal vessels using multi-scale features
Deep learning for variational multimodality tumor segmentation in PET/CT
AdaResU-Net: Multiobjective adaptive convolutional neural network for medical image segmentation
Fine-tuning Pre-trained Convolutional Neural Networks for Gastric Precancerous Disease Classification on Magnification Narrow-band Imaging Images
Brain tumor segmentation with deep convolutional symmetric neural network
Computer aided Alzheimer's disease diagnosis by an unsupervised deep learning technology
Bin loss for hard exudates segmentation in fundus images
Non-contact heart rate detection by combining empirical mode decomposition and permutation entropy under non-cooperative face shake
Dynamic MRI reconstruction exploiting blind compressed sensing combined transform learning regularization
Deep learning for ultrasound image caption generation based on object detection
Convolutional neural network based diagnosis of bone pathologies of proximal humerus

DL on MRI images

Table 1: Overview of papers using deep learning techniques for brain image analysis. All works use MRI unless otherwise mentioned.

Reference	Method	Application; remarks
Disorder classification (AD, MCI, Schizophrenia)		
Brosch and Tam (2013)	DBN	AD/HC classification; Deep belief networks with convolutional RBMs for manifold learning
Plis et al. (2014)	DBN	Deep belief networks evaluated on brain network estimation, Schizophrenia and Huntington's disease classification
Suk and Shen (2013)	SAE	AD/MCI classification; Stacked auto encoders with supervised fine tuning
Suk et al. (2014)	RBM	AD/MCI/HC classification; Deep Boltzmann Machines on MRI and PET modalities
Payan and Montana (2015)	CNN	AD/MCI/HC classification; 3D CNN pre-trained with sparse auto-encoders
Suk et al. (2015)	SAE	AD/MCI/HC classification; SAE for latent feature extraction on a large set of hand-crafted features from MRI and PET
Hosseini-Asl et al. (2016)	CNN	AD/MCI/HC classification; 3D CNN pre-trained with a 3D convolutional auto-encoder on fMRI data
Kim et al. (2016b)	ANN	Schizophrenia/NH classification on fMRI; Neural network showing advantage of pre-training with SAEs, and L1 sparsification
Ortiz et al. (2016)	DBN	AD/MCI/HC classification; An ensemble of Deep belief networks, with their votes fused using an SVM classifier
Pinaya et al. (2016)	DBN	Schizophrenia/NH classification; DBN pre-training followed by supervised fine-tuning
Sarraf and Tofighi (2016)	CNN	AD/HC classification; Adapted Lenet-5 architecture on fMRI data
Suk et al. (2016)	SAE	MCI/HC classification of fMRI data; Stacked auto-encoders for feature extraction, HMM as a generative model on top
Suk and Shen (2016)	CNN	AD/MCI/HC classification; CNN on sparse representations created by regression models
Shi et al. (2017)	ANN	AD/MCI/HC classification; Multi-modal stacked deep polynomial networks with an SVM classifier on top using MRI and PET
Tissue/anatomy/lesion/tumor segmentation		
Guo et al. (2014)	SAE	Hippocampus segmentation; SAE for representation learning used for target/atlas patch similarity measurement
de Brebisson and Montana (2015)	CNN	Anatomical segmentation; fusing multi-scale 2D patches with a 3D patch using a CNN
Choi and Jin (2016)	CNN	Striatum segmentation; Two-stage (global/local) approximations with 3D CNNs
Stollenga et al. (2015)	RNN	Tissue segmentation; PyraMiD-LSTM, best brain segmentation results on MRBrainS13 (and competitive results on EM-ISBI12)
Zhang et al. (2015)	CNN	Tissue segmentation; multi-modal 2D CNN
Andermatt et al. (2016)	RNN	Tissue segmentation; two convolutional gated recurrent units in different directions for each dimension
Bao and Chung (2016)	CNN	Anatomical segmentation; Multi-scale late fusion CNN with random walker as a novel label consistency method
Birenbaum and Greenspan (2016)	CNN	Lesion segmentation; Multi-view (2.5D) CNN concatenating features from previous time step for a longitudinal analysis
Brosch et al. (2016)	CNN	Lesion segmentation; Convolutional encoder-decoder network with shortcut connections and convolutional RBM pretraining
Chen et al. (2016a)	CNN	Tissue segmentation; 3D res-net combining features from different layers
Ghafoorian et al. (2016b)	CNN	Lesion segmentation; CNN trained on non-uniformly sampled patch to integrate a larger context with a foviation effect
Ghafoorian et al. (2016a)	CNN	Lesion segmentation; multi-scale CNN with late fusion that integrates anatomical location information into network
Havaei et al. (2016b)	CNN	Tumor segmentation; CNN handling missing modalities with abstraction layer that transforms feature maps to their statistics
Havaei et al. (2016a)	CNN	Tumor segmentation; two-path way CNN with different receptive fields
Kamnitsas et al. (2017)	CNN	Tumor segmentation; 3D multi-scale fully convolutional network with CRF for label consistency
Kleesiek et al. (2016)	CNN	Brain extraction; 3D fully convolutional CNN on multi-modal input
Mansoor et al. (2016)	SAE	Visual pathway segmentation; Learning appearance features from SAE for steering the shape model for segmentation
Milletari et al. (2016a)	CNN	Anatomical segmentation on MRI and US; Hough-voting to acquire mapping from CNN features to full patch segmentations
Moeskops et al. (2016a)	CNN	Tissue segmentation; CNN trained on multiple patch sizes

DL on brain images

Brain		
Brain extraction	[269]	A 3D CNN for skull stripping
Functional connectomes	[270]	Transfer learning approach to enhance deep neural network classification of brain functional connectomes
	[271]	Multisite diagnostic classification of schizophrenia using discriminant deep learning with functional connectivi MRI
Structural connectomes	[272]	A convolutional neural network-based approach (https://github.com/MIC-DKFZ/TractSeg) that directly segments tracts in the field of fiber orientation distribution function (fODF) peaks without using tractography, image registration or parcellation. Tested on 105 subjects from the Human Connectome Project
Brain age	[273]	Chronological age prediction from raw brain T1-MRI data, also testing the heritability of brain-predicted age using a sample of 62 monozygotic and dizygotic twins
Alzheimer's disease	[274]	Landmark-based deep multi-instance learning evaluated on 1526 subjects from three public datasets (ADNI-1, ADNI-2, MIRIAD)
	[275]	Identify different stages of AD
	[276]	Multimodal and multiscale deep neural networks for the early diagnosis of AD using structural MR and FDG-PET images
Vascular lesions	[277]	Evaluation of a deep learning approach for the segmentation of brain tissues and white matter hyperintensities presumed vascular origin in MRI
Identification of MRI contrast	[278]	Using deep learning algorithms to automatically identify the brain MRI contrast, with implications for managi large databases
Meningioma	[279]	Fully automated detection and segmentation of meningiomas using deep learning on routine multiparametric MRI
Glioma	[280]	Glioblastoma segmentation using heterogeneous MRI data from clinical routine
	[281]	Deep learning for segmentation of brain tumors and impact of cross-institutional training and testing
	[282]	Automatic semantic segmentation of brain gliomas from MRI using a deep cascaded neural network
	[283]	AdaptAhead optimization algorithm for learning deep CNN applied to MRI segmentation of glioblastomas (BRATS)
Multiple sclerosis	[284]	Deep learning of joint myelin and T1w MRI features in normal-appearing brain tissue to distinguish between multiple sclerosis patients and healthy controls

Challenges in the DL-medical domain

Name	Summary	Link
Grand-Challenges	Grand challenges in biomedical image analysis. Hosts and lists a large number of competitions	https://grand-challenge.org/
RSNA Pneumonia Detection Challenge	Automatically locate lung opacities on chest radiographs	https://www.kaggle.com/c/rsna-pneumonia-detection-challenge
HVSMR 2016	Segment the blood pool and myocardium from a 3D cardio-vascular magnetic resonance image	http://segchd.csail.mit.edu/
ISLES 2018	Ischemic Stroke Lesion Segmentation 2018. The goal is to segment stroke lesions based on acute CT perfusion data.	http://www.isles-challenge.org/
BraTS 2018	Multimodal Brain Tumor Segmentation. The goal is to segment brain tumors in multimodal MRI scans.	http://www.med.upenn.edu/sbia/brats2018.html
CAMELYON17	The goal is to develop algorithms for automated detection and classification of breast cancer metastases in whole-slide images of histological lymph node sections.	https://camelyon17.grand-challenge.org/Home
ISIC 2018	Skin Lesion Analysis Towards Melanoma Detection	https://challenge2018.isic-archive.com/
Kaggle's 2018 Data Science Bowl	Spot Nuclei. Speed Cures.	https://www.kaggle.com/c/data-science-bowl-2018
Kaggle's 2017 Data Science Bowl	Turning Machine Intelligence Against Lung Cancer	https://www.kaggle.com/c/data-science-bowl-2017
Kaggle's 2016 Data Science Bowl	Transforming How We Diagnose Heart Disease	https://www.kaggle.com/c/second-annual-data-science-bowl
MURA	Determine whether a bone X-ray is normal or abnormal	https://stanfordmlgroup.github.io/competitions/mura/