



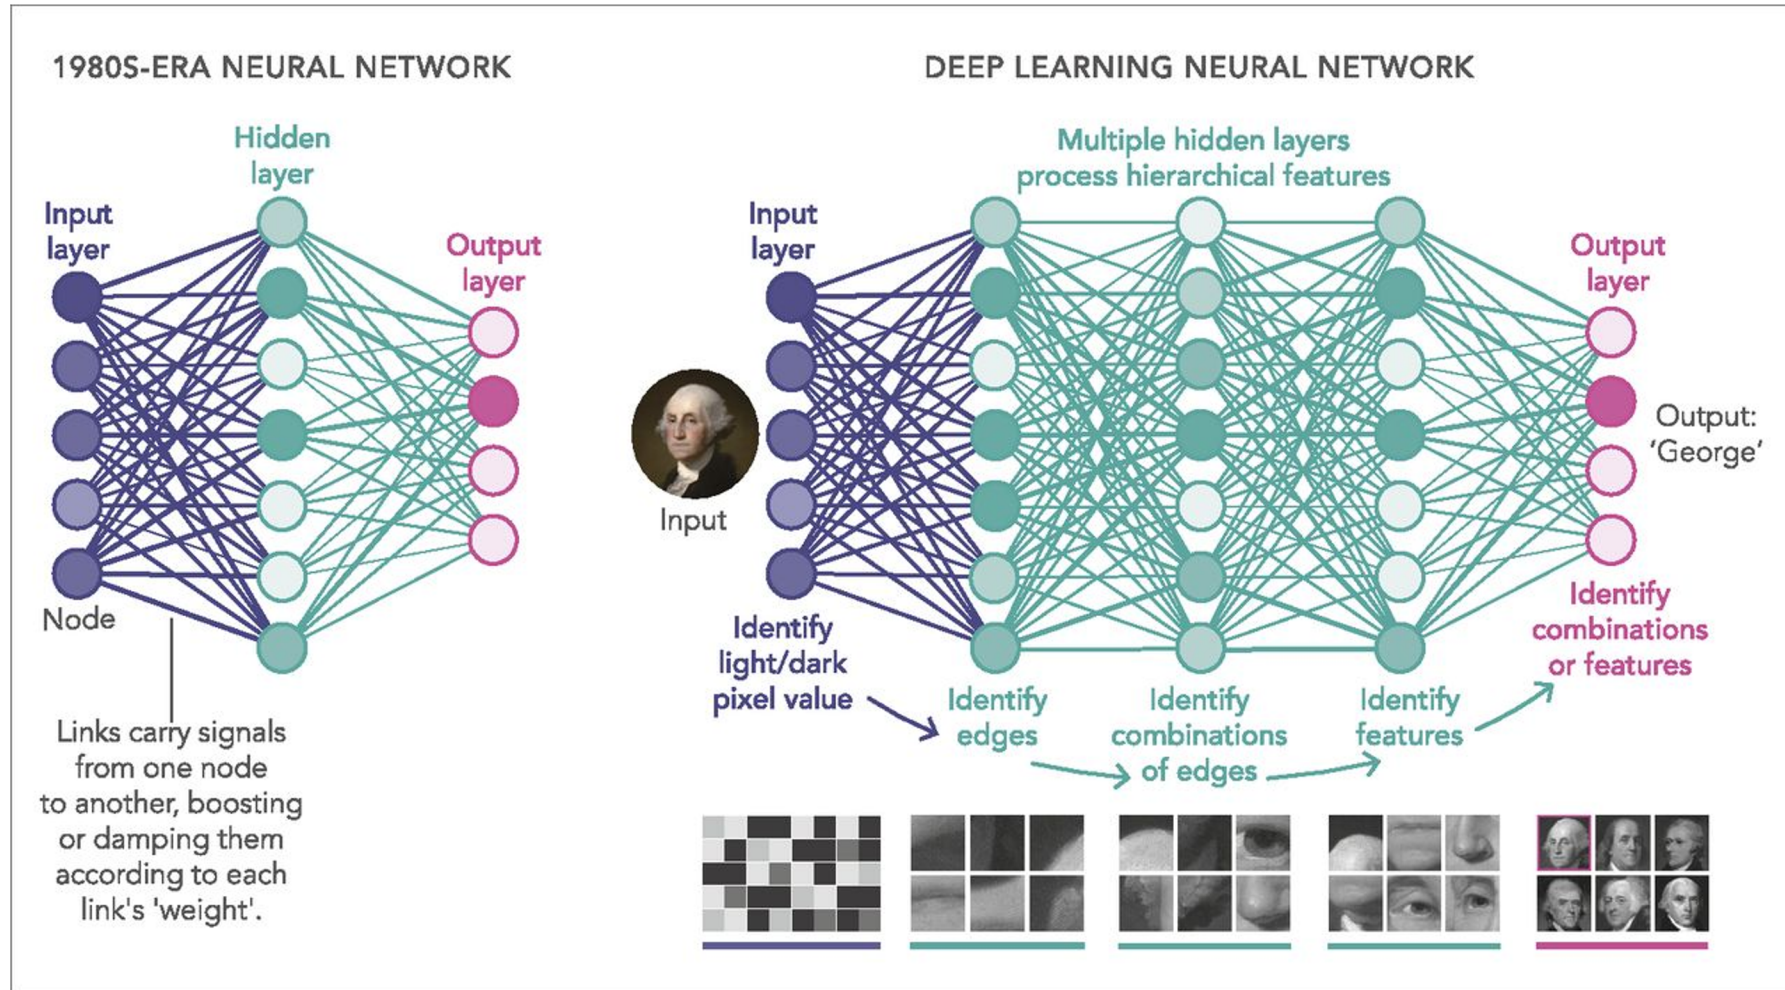
ARTIFICIAL INTELLIGENCE & CYBER SECURITY

NEURAL NETWORKS

What comes next? A brief intro to DL & CV

Estefanía Talavera Martínez
e.talaveramartinez@utwente.nl

DEEP LEARNING



DEEP LEARNING (IMPACT)

The Deep Learning Revolution (by NVIDIA)




[Video](#)
[\[Click here\]](#)

ARTIFICIAL VS CONVOLUTIONA NEURAL NETWORKS

Question : what is the difference between Convolutional Neural Networks (CNNs) and Artificial Neural Networks (NNs)?

ANN



ARTIFICIAL VS CONVOLUTIONA NEURAL NETWORKS

Question : what is the difference between Convolutional Neural Networks (CNNs) and Artificial Neural Networks (NNs)?


PROC. OF THE IEEE, NOVEMBER 1998

Gradient-Based Learning Applied to Document Recognition

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner

all trained ('end-to-end'). representation on different levels of detail/abstraction


CNN



ARTIFICIAL VS CONVOLUTIONAL NEURAL NETWORKS


Question : what is the difference between Convolutional Neural Networks (CNNs) and Artificial Neural Networks (NNs)?

ANN



Spatial information is lost!

CNN

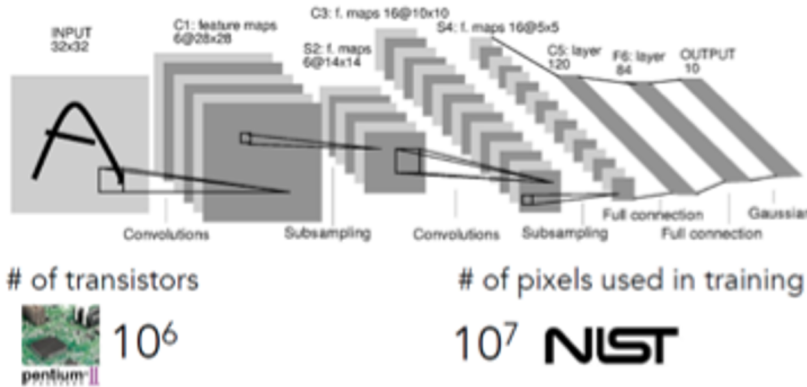


Can learn spatial correlations!


ARTIFICIAL VS CONVOLUTIONA NEURAL NETWORKS

Convolutional Neural Networks' Evolution


1998
LeCun et al.
LeNet



2012
Krizhevsky et al.
AlexNet




2015
He et al.
ResNet-152



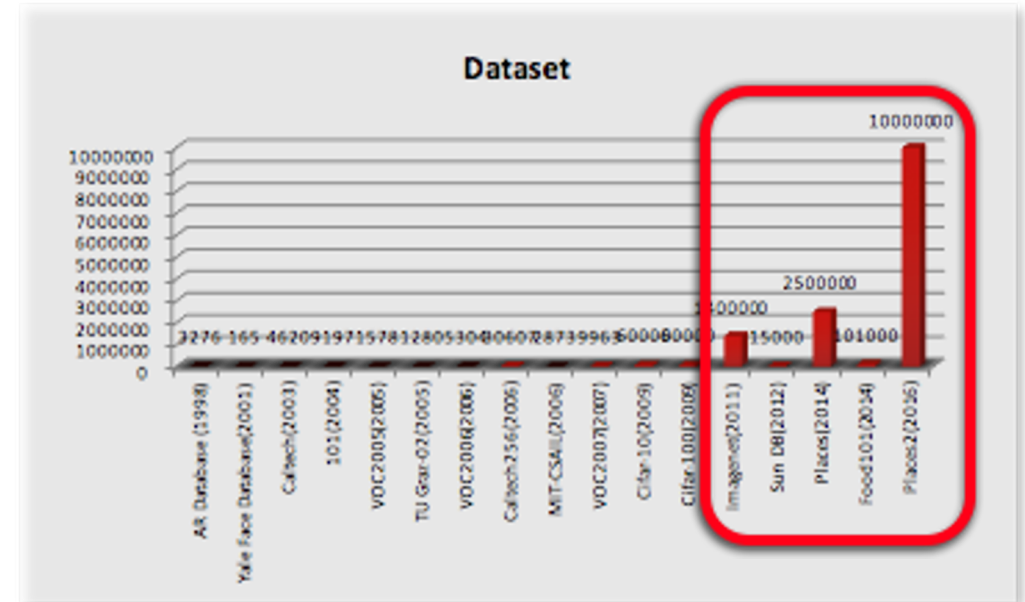
DEEP LEARNING IMPACT

Number of objects/Dataset



ImageNet &
Deep Learning

Number of images/Dataset



APPLICATIONS OF DEEP LEARNING

Input: x



Is this an urban or rural area?

Output: $y = \{-1, +1\}$

Binary Classification

Which city is this?

Output: $y = \{0, 1, 2, \dots, C\}$

Multi-class Classification

APPLICATIONS OF DEEP LEARNING


Input: x



Where is the object in the image?

Output: $y = \{ x, y, w, h \}$ coordinates

Object Detection



What is the semantic class of each pixel?

Output: $y = \{ 0, 1, 2, \dots, C \}$ per pixel

Semantic Segmentation

APPLICATIONS OF DEEP LEARNING

Classification



Retrieval



[Krizhevsky 2012]

Detection



[Faster R-CNN: Ren, He, Girshick, Sun 2015]

Segmentation

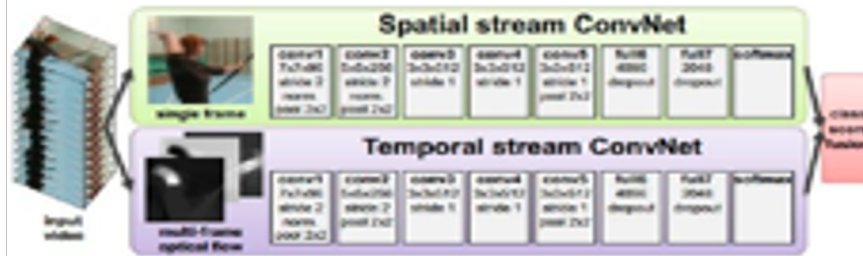



[Farabet et al., 2012]

APPLICATIONS OF DEEP LEARNING



self-driving cars




[Goodfellow 2014]

APPLICATIONS OF DEEP LEARNING



[Toshev, Szegedy 2014]



[Ciresan et al. 2013]



[Sermanet et al. 2011]
[Ciresan et al.]

APPLICATIONS OF DEEP LEARNING



A group of young people playing a game of frisbee.



Two hockey players are fighting over the puck.



A little girl in a pink hat is blowing bubbles.



A refrigerator filled with lots of food and drinks.



A herd of elephants walking across a dry grass field.



A close up of a cat laying on a couch.



A red motorcycle parked on the side of the road.



A yellow school bus parked in a parking lot.



reddit.com/r/deepdream

[Vinyals et al., 2015]

COMPUTER VISION & DEEP LEARNING

'How we're teaching computers to understand pictures'



[Video](#)
[\[Click here\]](#)

Fei-Fei Li · Computer scientist
As Director of Stanford's Artificial Intelligence Lab and Vision Lab