Lecture 13: Deep Learning

Historical review of deep learning

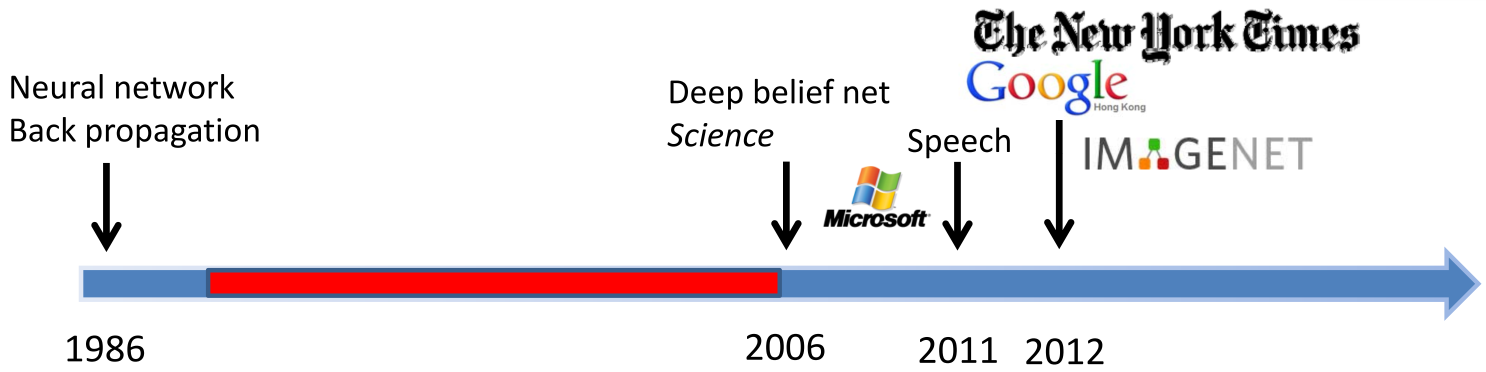

- Google and Baidu announced their deep learning based visual search engines (2013)
  - Google
    - "on our test set we saw **double the average precision** when compared to other approaches we had tried. We acquired the rights to the technology and went full speed ahead adapting it to run at large scale on Google's computers. We took cutting edge research straight out of an academic research lab and launched it, in just a little over six months."
  - Baidu

Historical review of DL (con't)

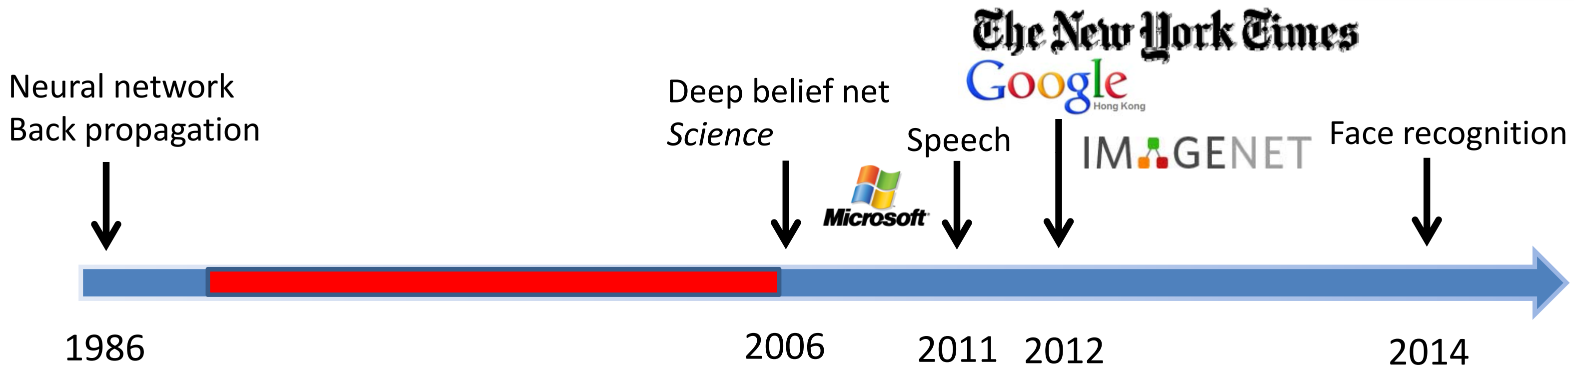

- Deep learning achieves 99.47% face verification accuracy on Labeled Faces in the Wild (LFW), higher than human performance

Y. Sun, X. Wang, and X. Tang. Deep Learning Face Representation by Joint Identification-Verification. NIPS, 2014.

Y. Sun, X. Wang, and X. Tang. Deeply learned face representations are sparse, selective, and robust. CVPR, 2015.

Major types

Deep Boltzmann machine:

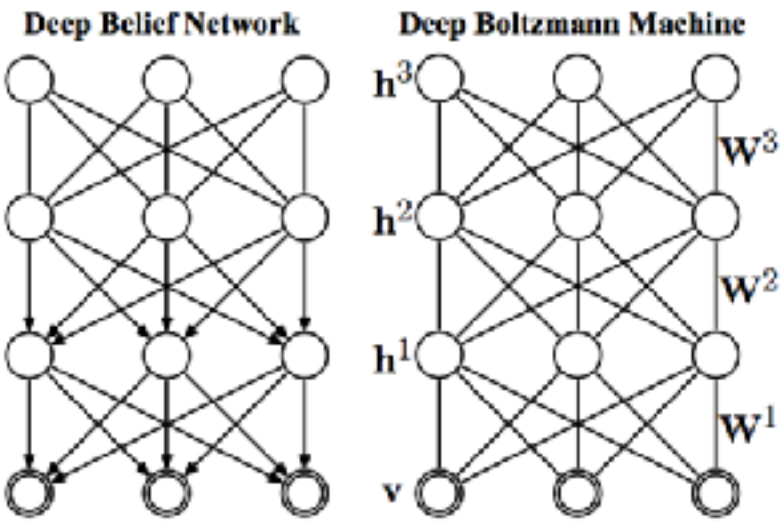

Auto-encoder:

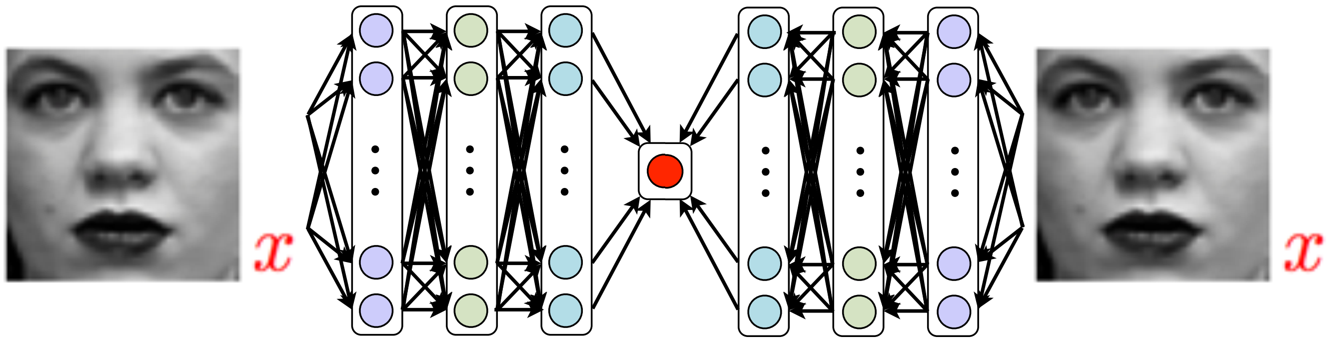

Major types

Convolutional neural networks:

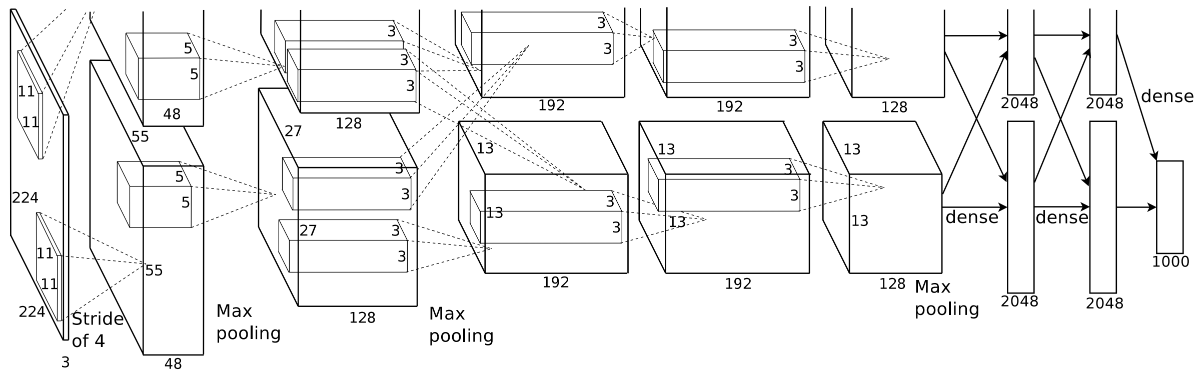

Recurrent neural networks:

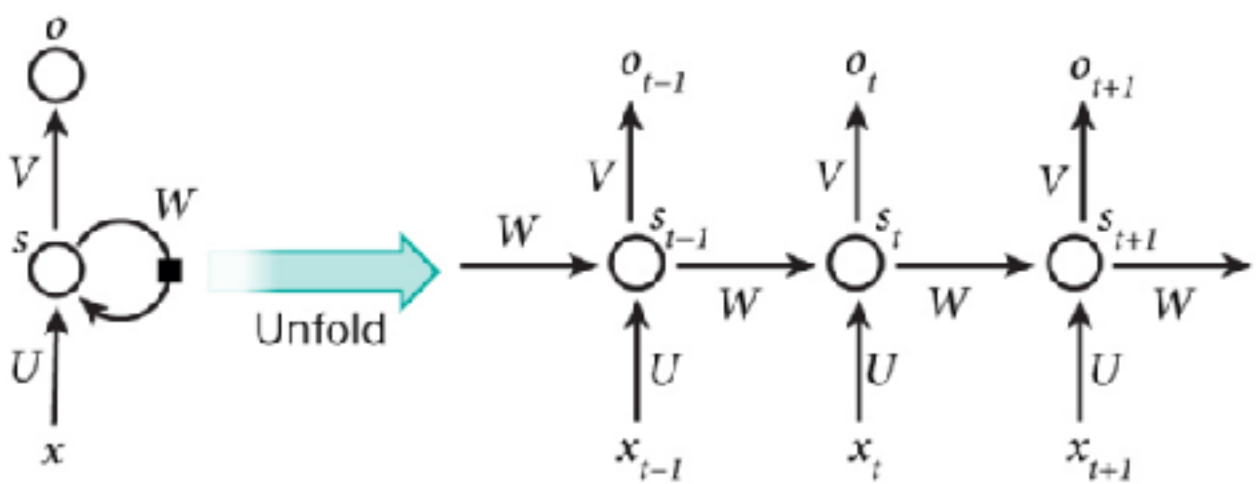

Autoencoder

autoencoder

restricted Boltzmann machine
a type of associative memory network

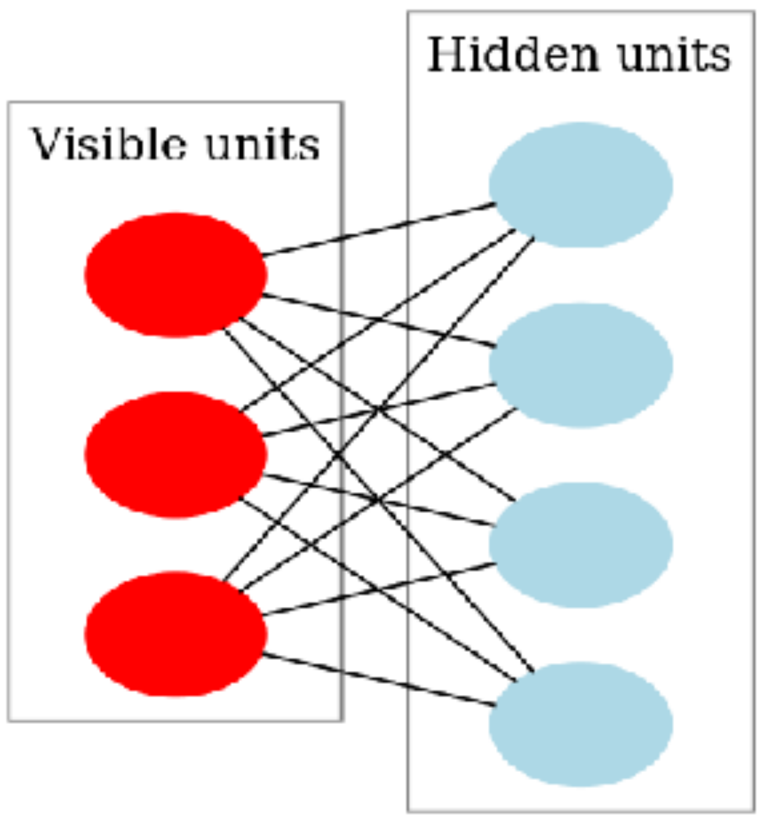

[image from http://en.wikipedia.org/wiki/Restricted_Boltzmann_machine

Autoencoder

autoencoder

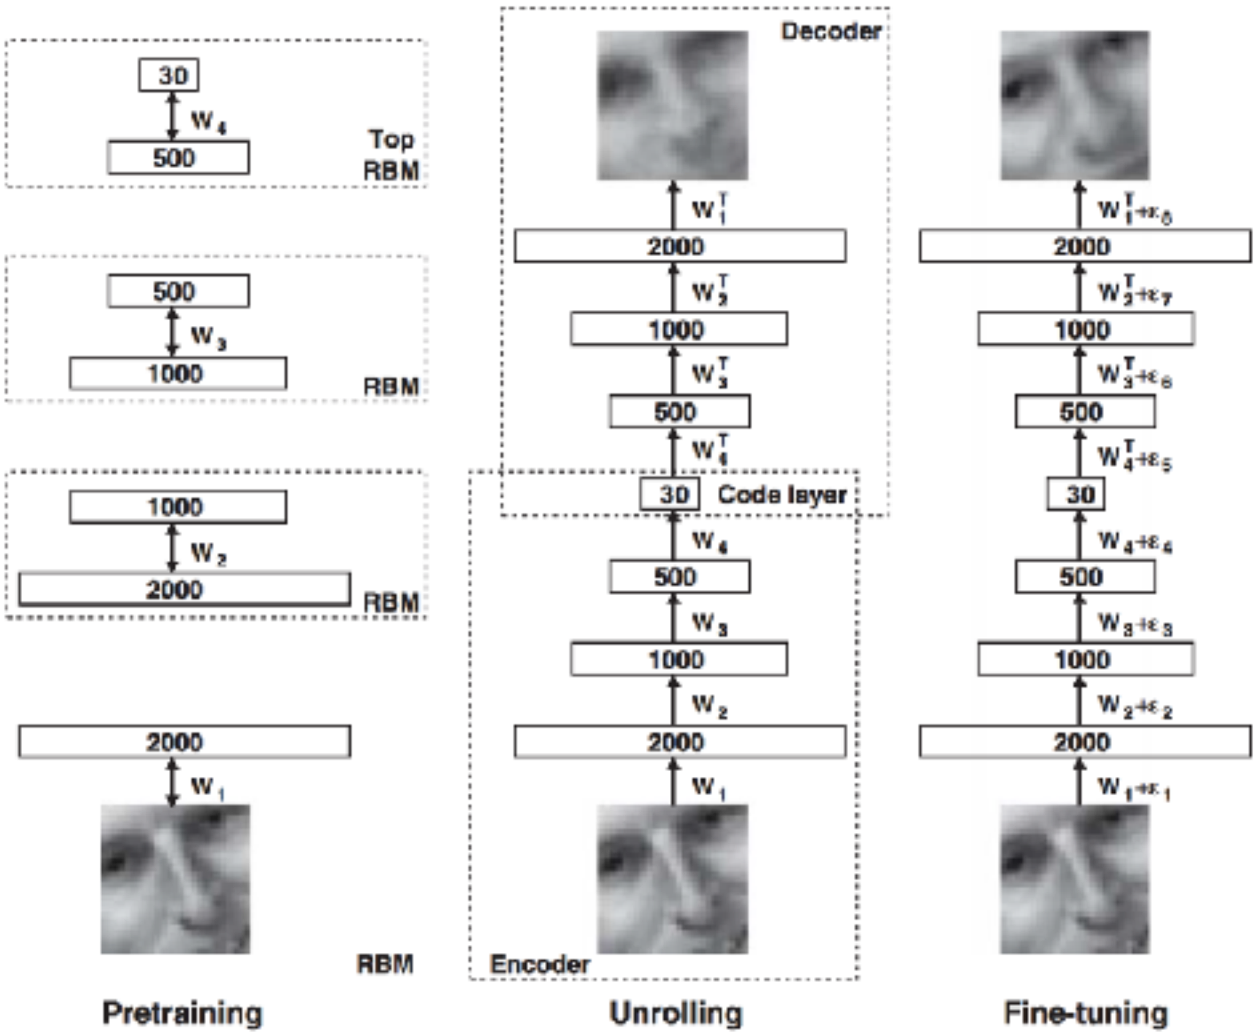

[image from [G. E. Hinton and R. R. Salakhutdinov, Science 2006]]

Autoencoder

autoencoder

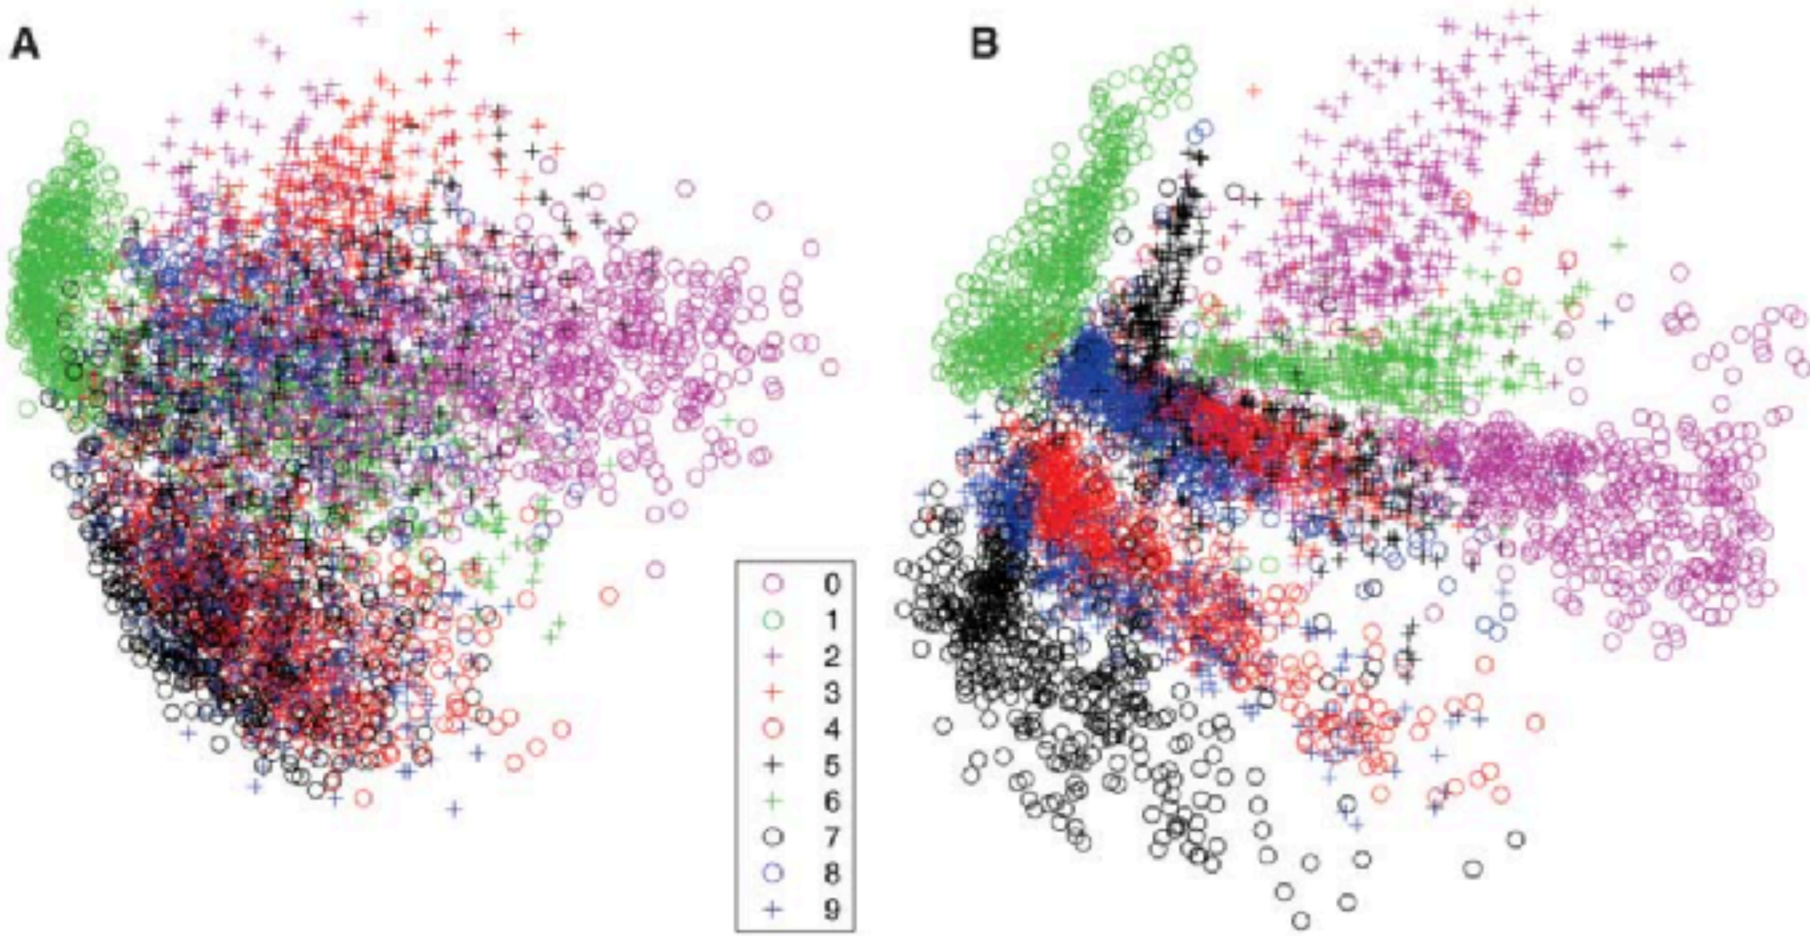

PCA autoencoder

[image from [G. E. Hinton and R. R. Salakhutdinov, Science 2006]]

CNN

Convolutional Neural Networks (CNN/LeNet) for general image feature extraction

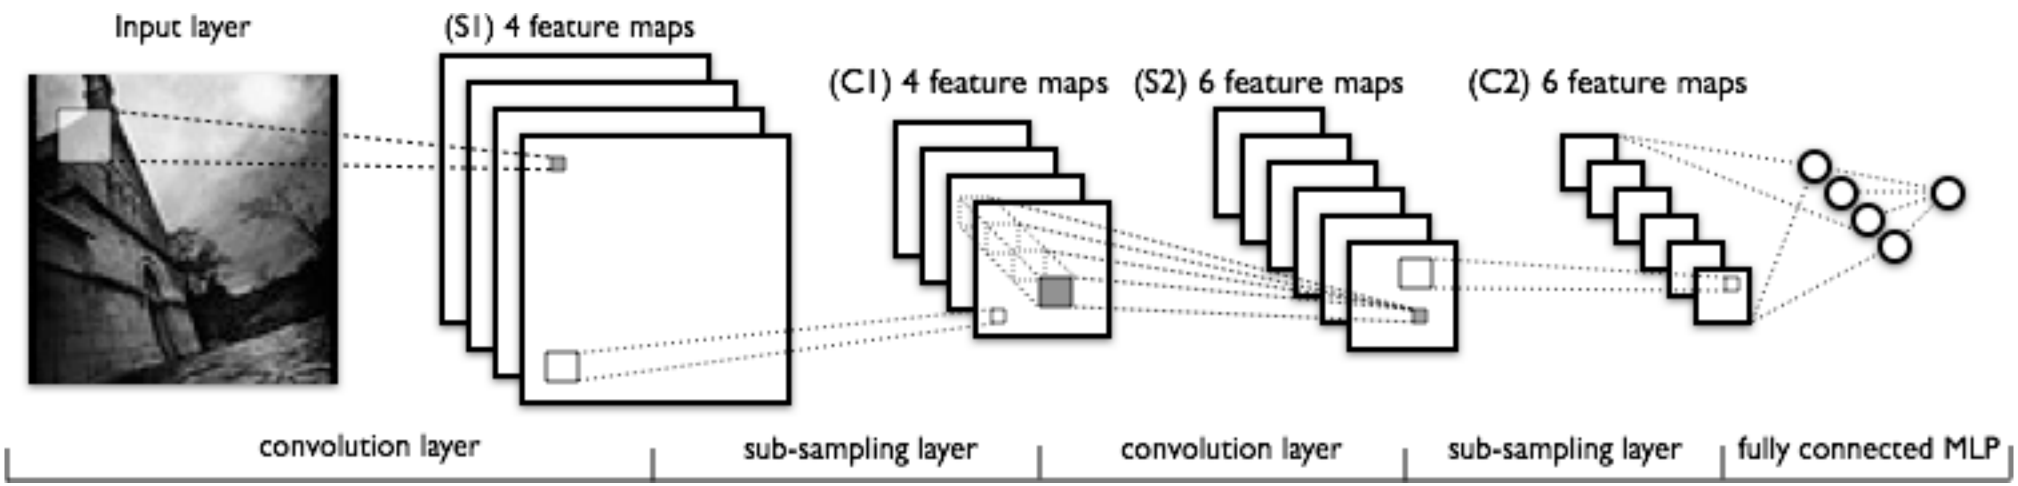

[image from http://deeplearning.net/tutorial/lenet.html]

CNN

Convolution layer

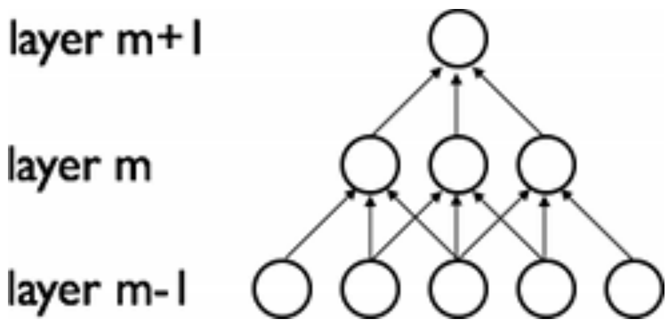

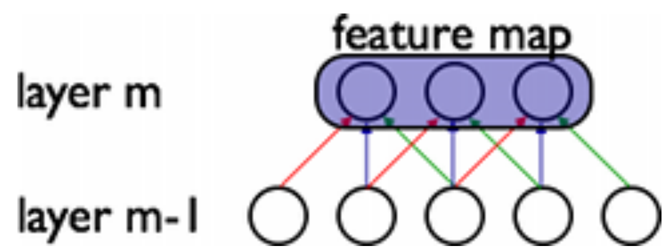

sparse connectivity

shared weights

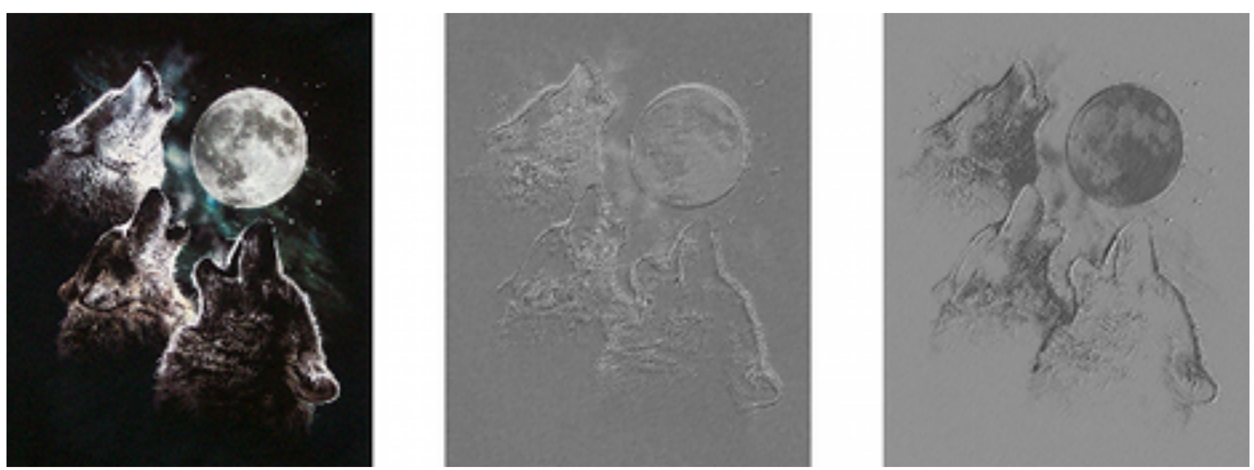

[image from http://deeplearning.net/tutorial/lenet.html]

Subsampling layer

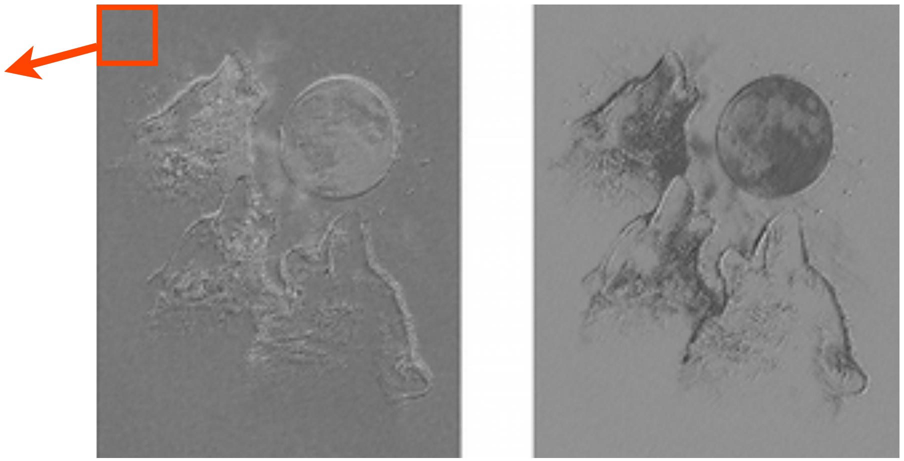

[image from http://deeplearning.net/tutorial/lenet.html]

CNN

Convolutional Neural Networks (CNN/LeNet) for general image feature extraction

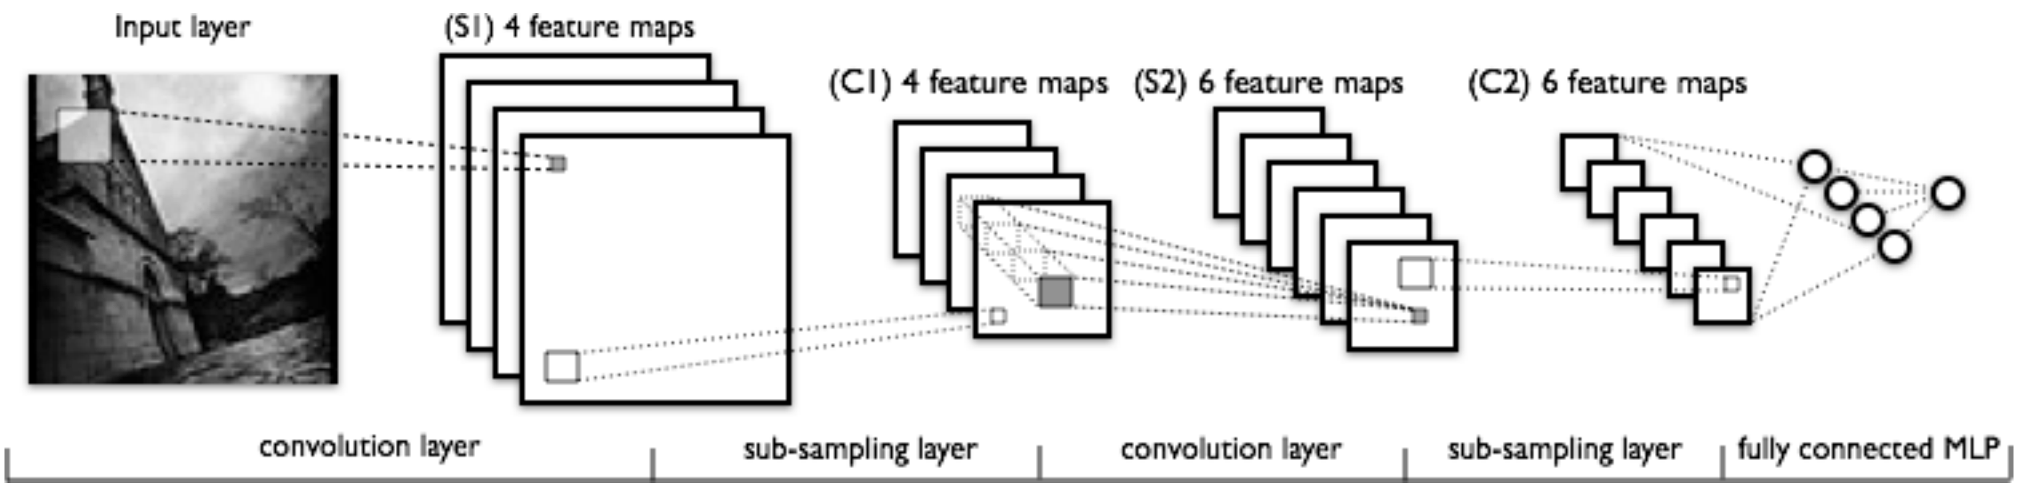

[image from http://deeplearning.net/tutorial/lenet.html]

Activation functions (con't)

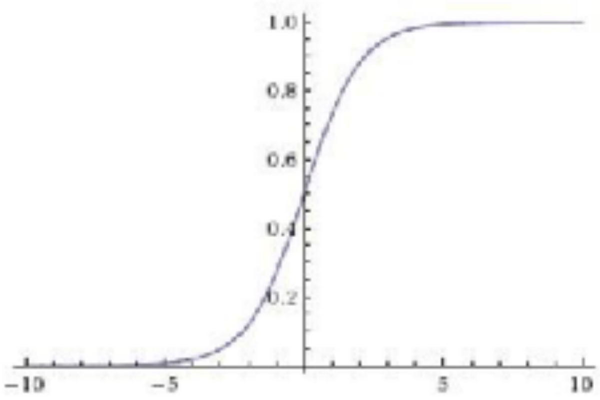

Sigmoid

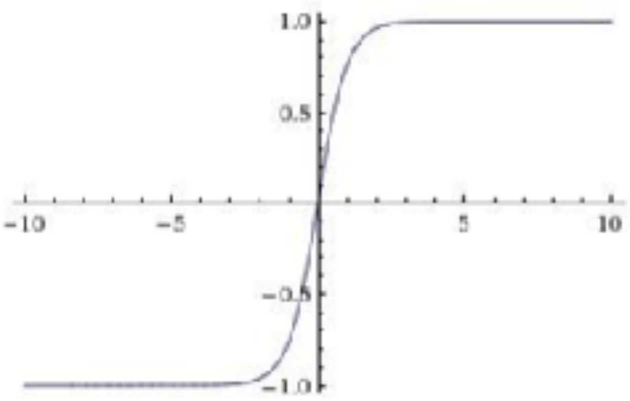

tanh(x)

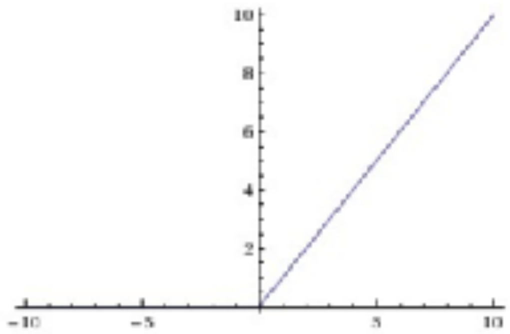

ReLU
(Rectified Linear Unit)

And many more …

CNN

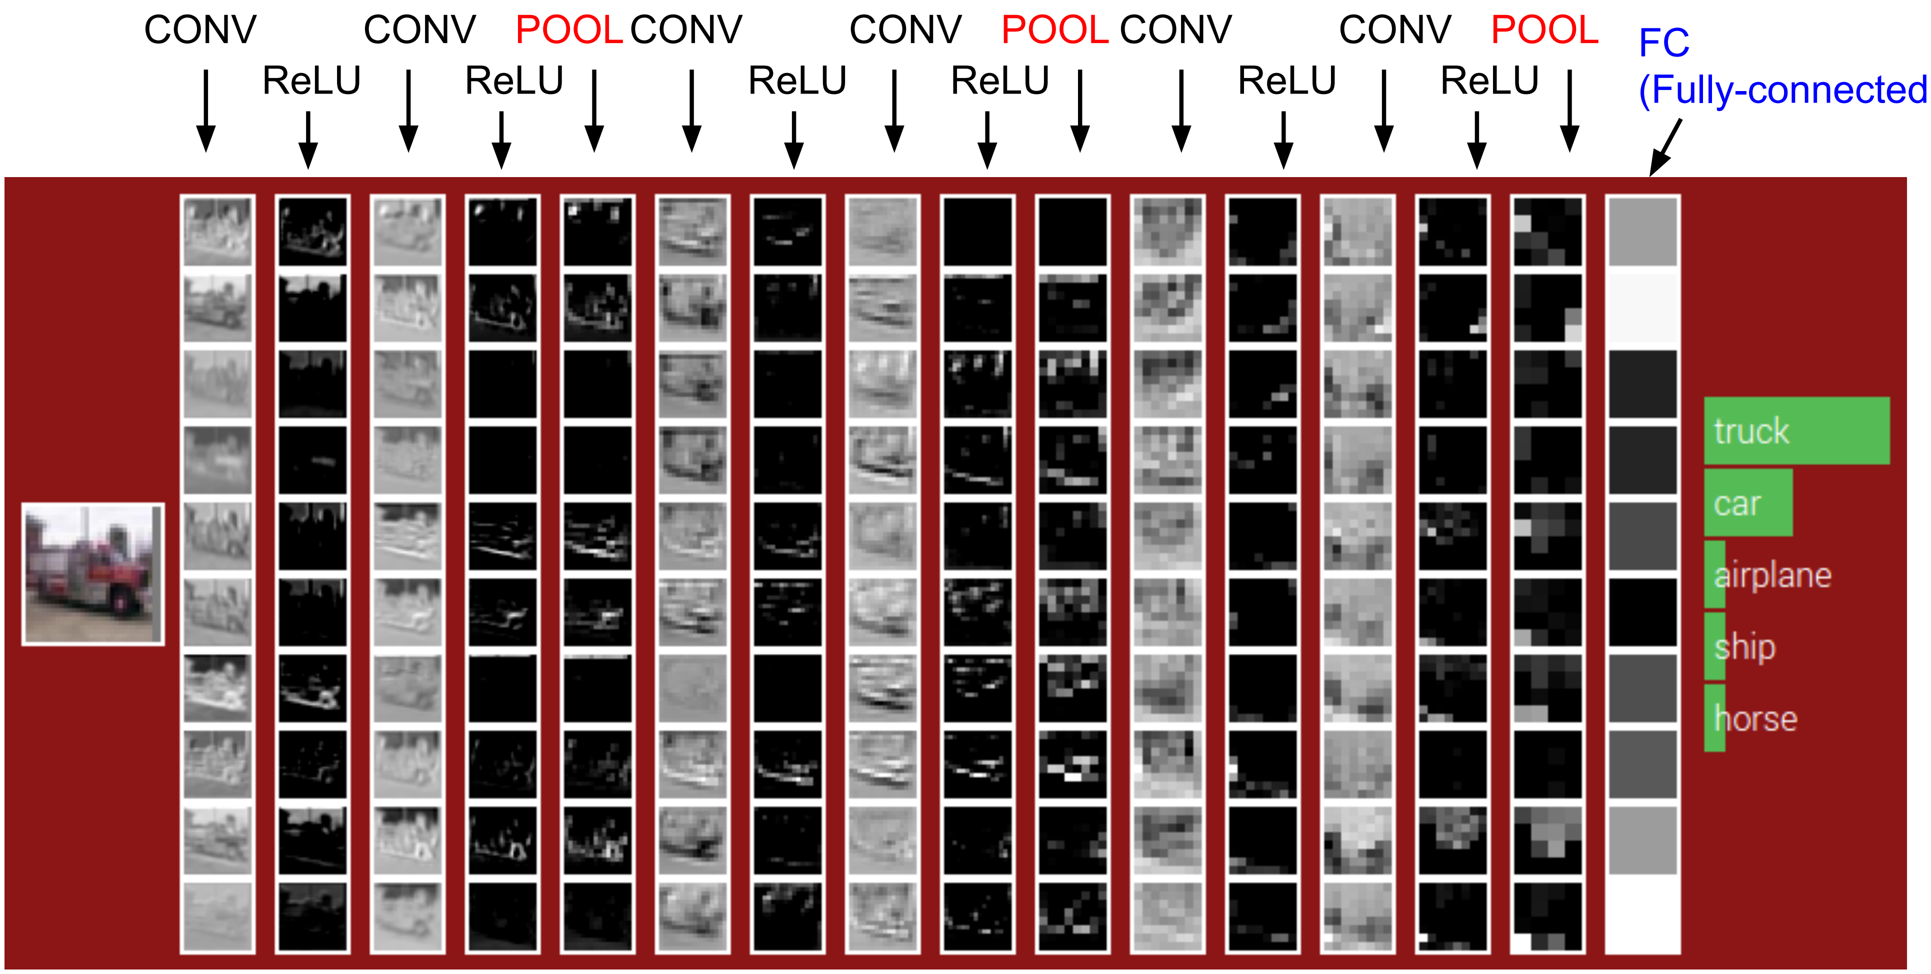
CONV
ReLU
CONV
ReLU
POOL
CONV
ReLU
CONV
ReLU
POOL
CONV
ReLU
CONV
ReLU
POOL
FC
(Fully-connected)
truck
car
airplane
ship
horse

CNN

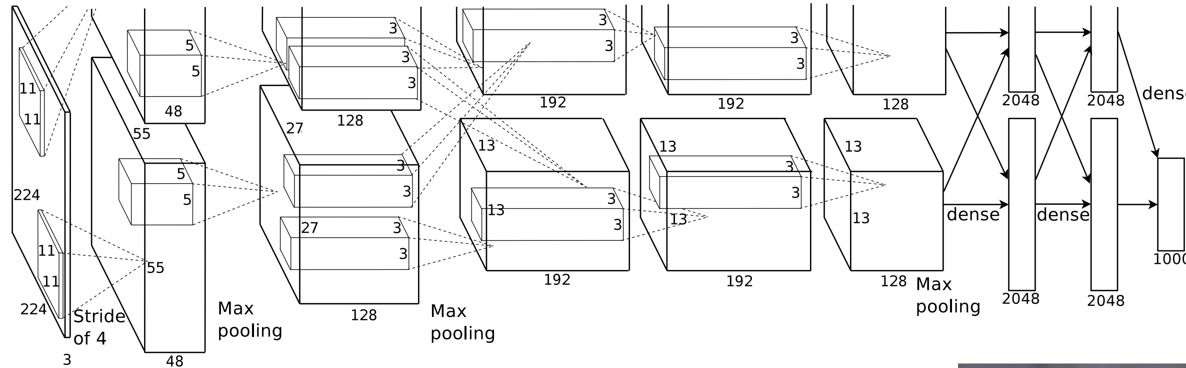

4.94% (DL) vs 5.1% (human)

Geoffrey E. Hinton
University of Toronto

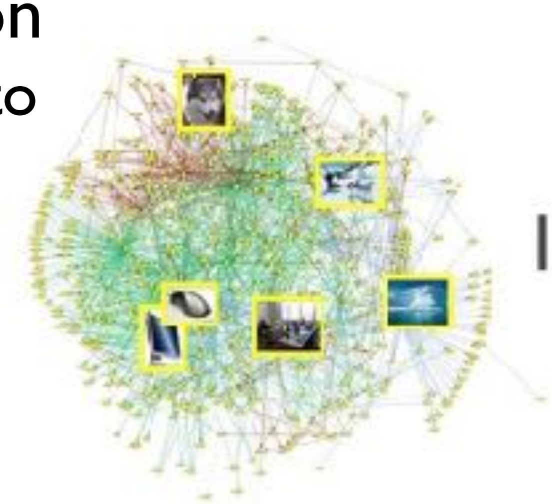

IMAGENET

Fei-Fei Li
Stanford University

CNN toolbox

- Tensorflow (Google)
- PyTorch/Torch (Facebook & NYU)
- Caffe (UC Berkeley)
- …

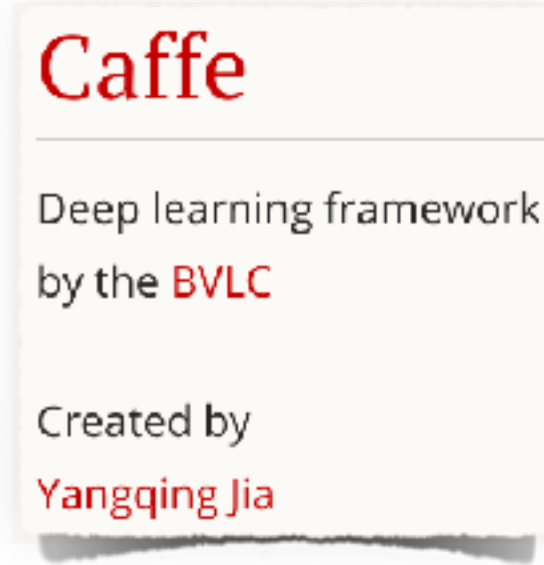

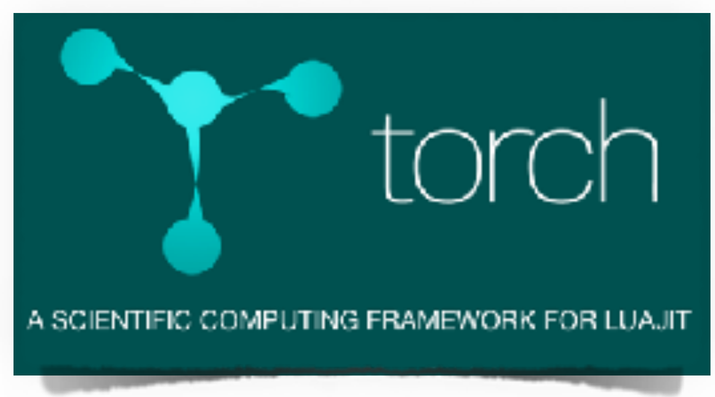

deeplearning.net

NVIDIA-GPUs

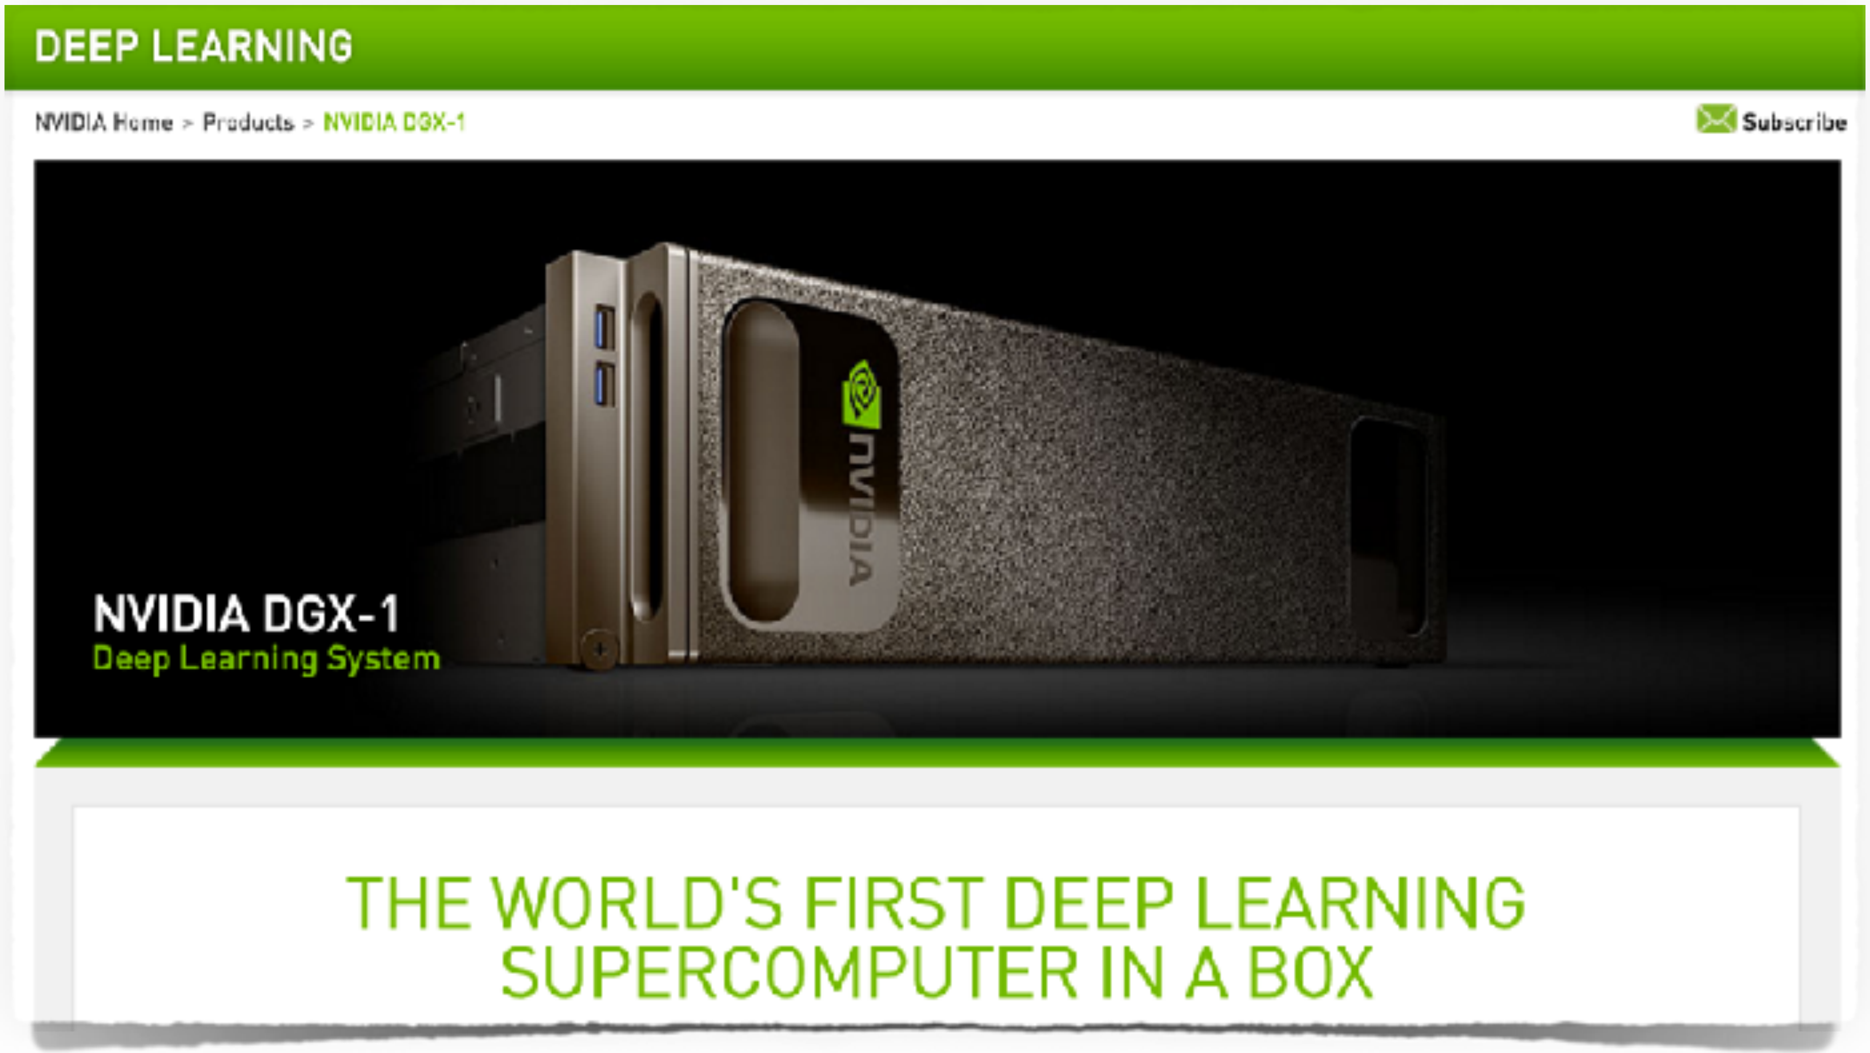
DEEP LEARNING
NVIDIA Home > Products > NVIDIA DGX-1
Subscribe
NVIDIA
NVIDIA DGX-1
Deep Learning System
THE WORLD'S FIRST DEEP LEARNING
SUPERCOMPUTER IN A BOX

Some Applications

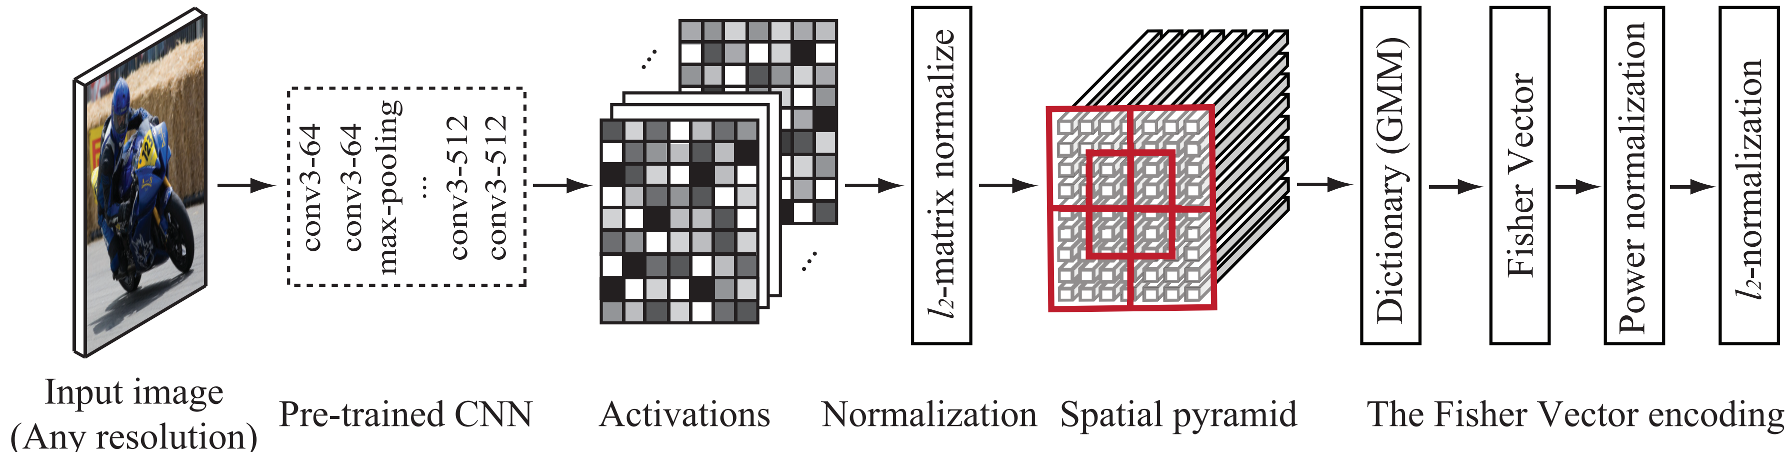

Pre-trained model as feature extractor

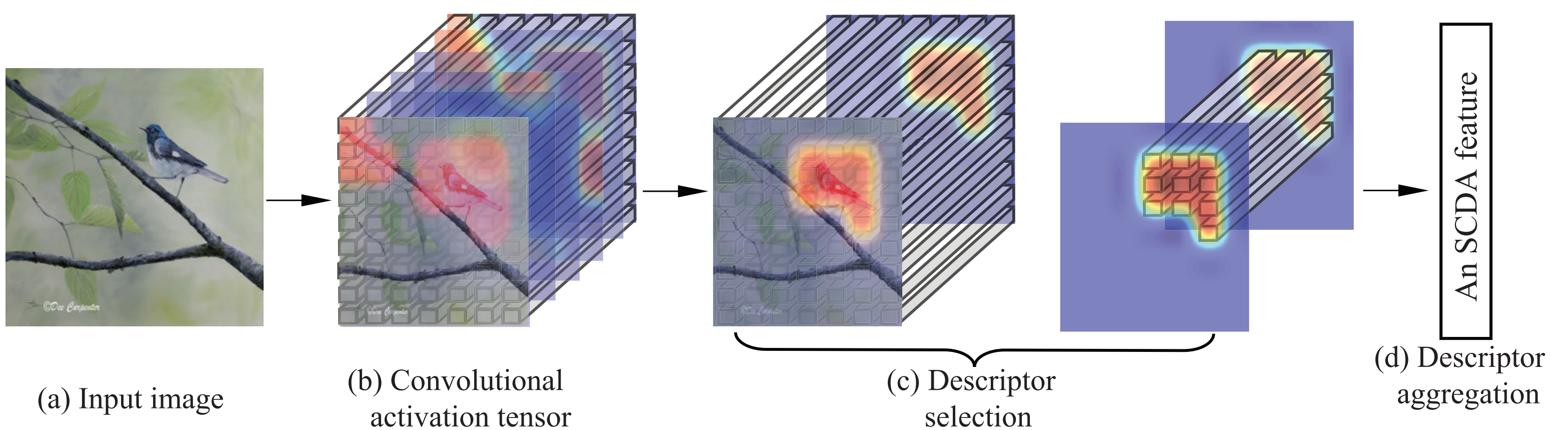

Fine-grained image retrieval

Some Applications

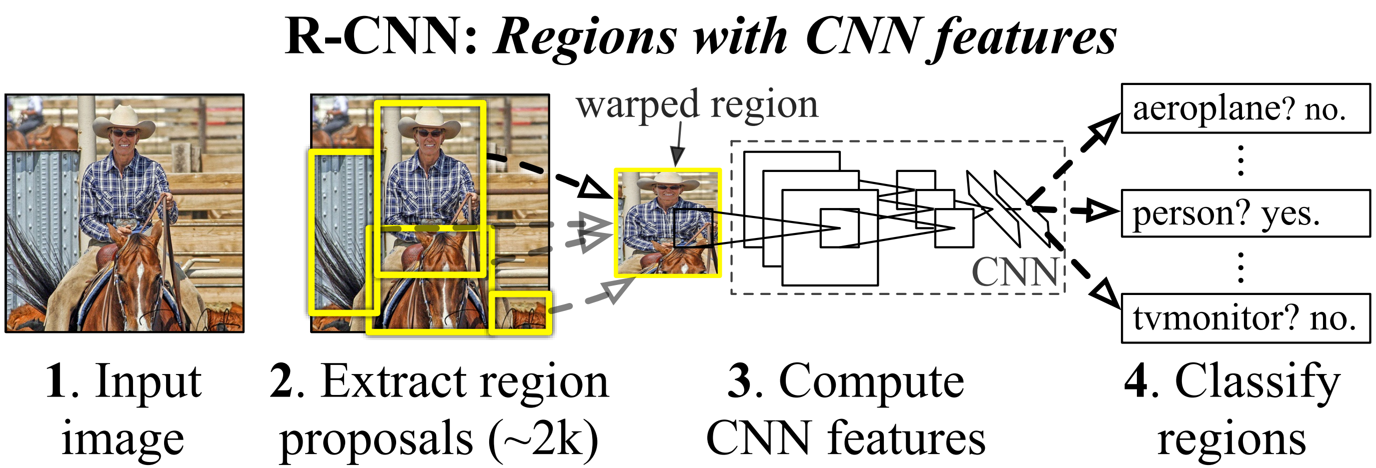

Object detection

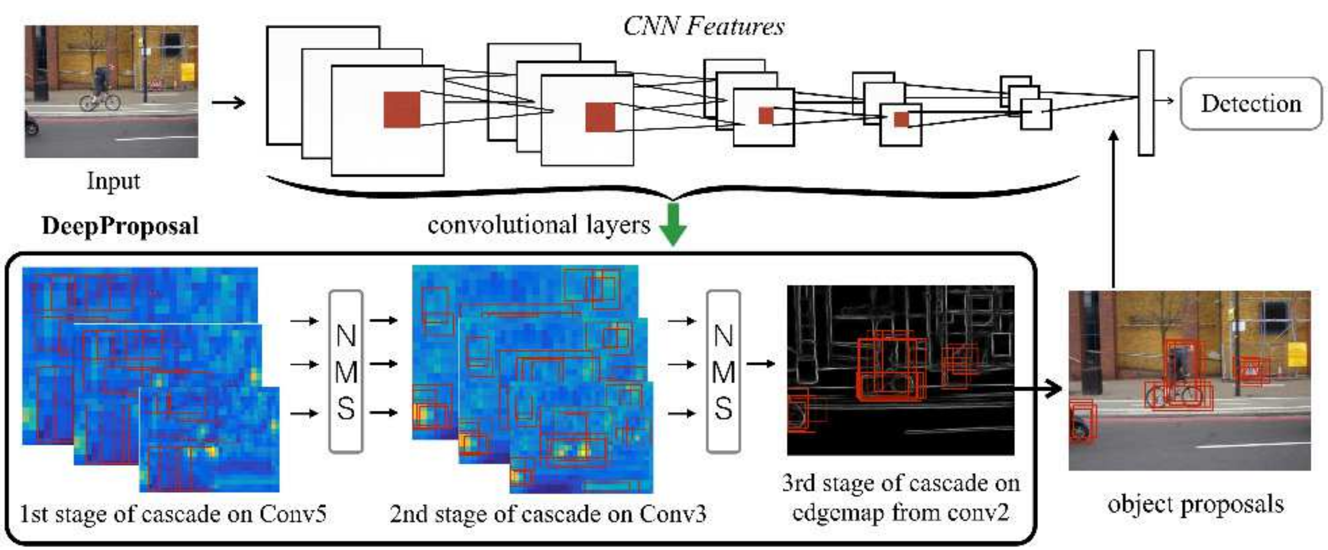

DeepProposal

Some Applications

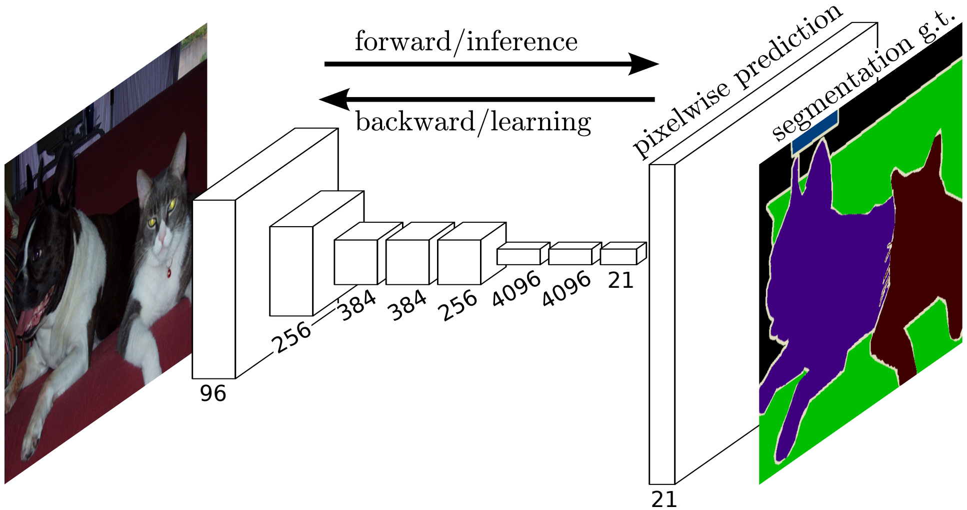

Semantic segmentation

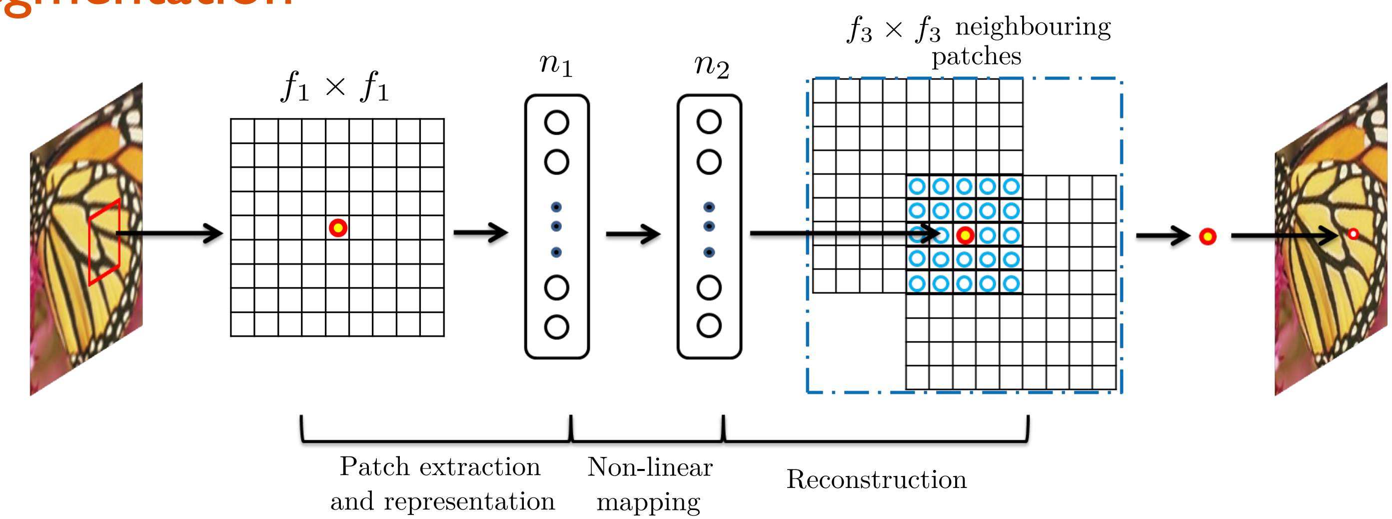

Super-resolution

Some Applications

Object segmentation

Fine-grained classification

Some Applications

Object segn

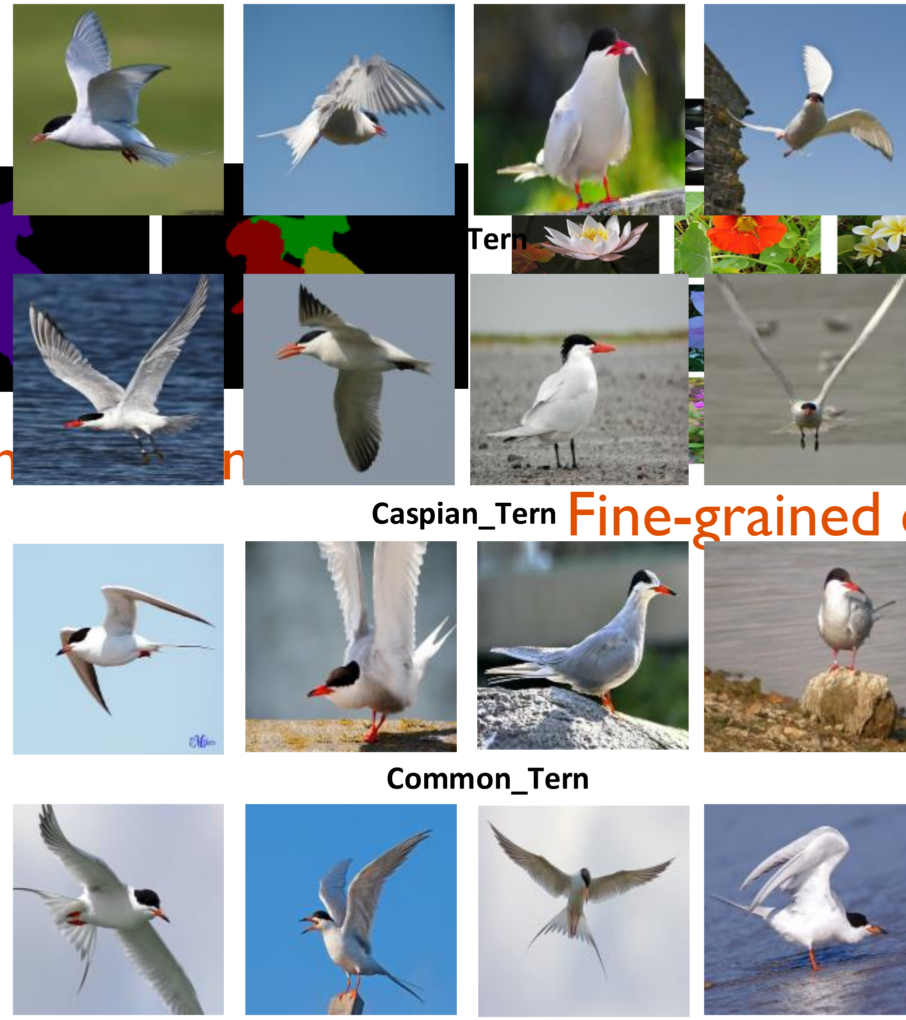

Fine-grained classification

Some Applications

Object segmentation

Fine-grained classification

Face recognition

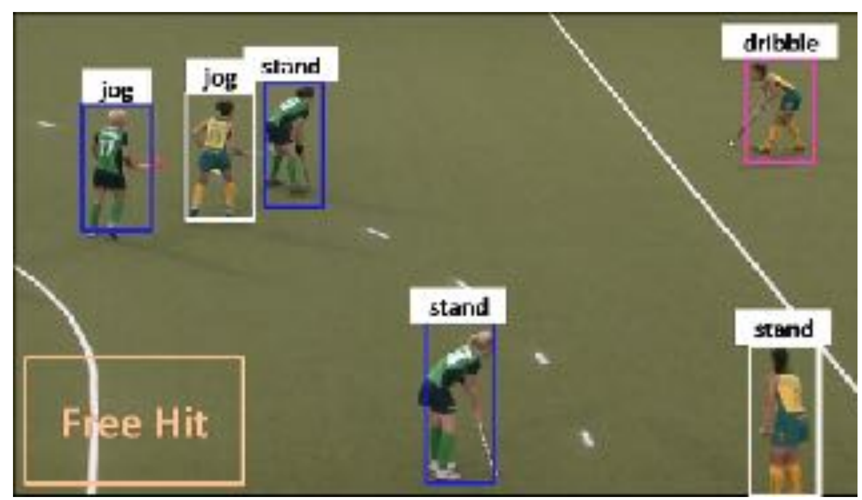

Action recognition

Some Applications

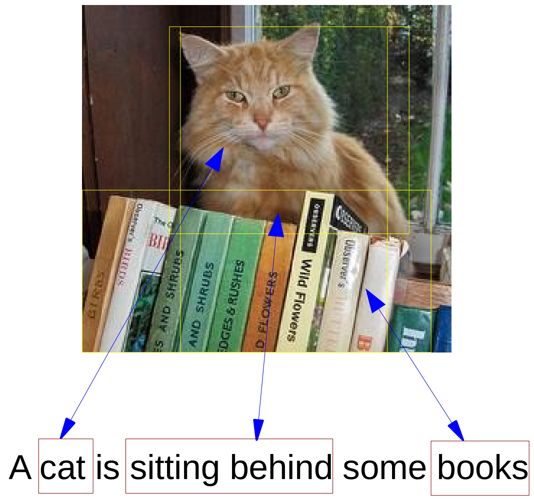

Image caption

Automatic driving

Cultural event recognition

Some Applications

Multimodal Linguistic Regularities

[Kiros et al., TACL 2015]

Some Applications

Multimodal Linguistic Regularities

- day + night =

- flying + sailing =

- bowl + box =

- box + bowl =

Nearest images

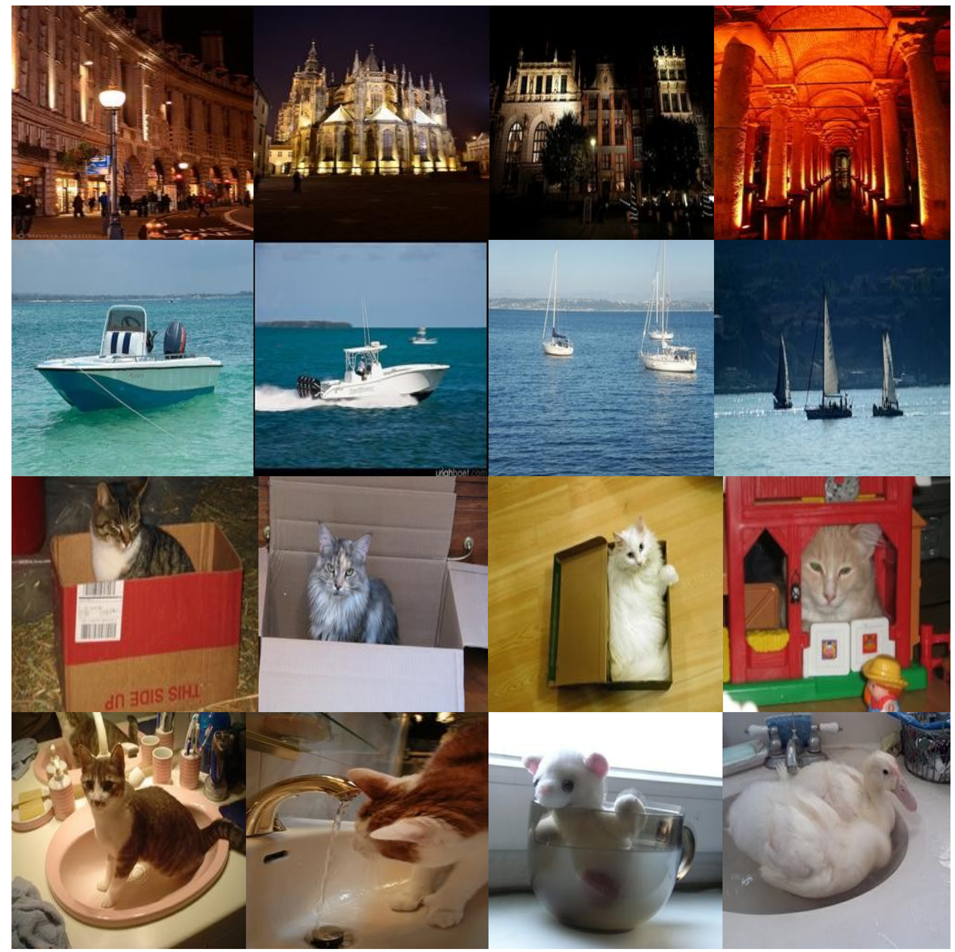

[Kiros et al., TACL 2015]

Transformation for text

One-hot encoding

"I love it." $\longrightarrow$ $\mathbf{x} = [\,0\;0\;1\;0\;0 \,|\, 0\;0\;0\;0\;1 \,|\, 0\;0\;0\;1\;0\,]^\top$

V={"don't", "hate", "I", "it", "love"}

Transformation for text

Seq-CNN for text

$$\text{“}\textit{I love it.}\text{”} \longrightarrow \mathbf{r}_0(\mathbf{x}) = \left[\begin{array}{c} 0 \\ 0 \\ \mathbf{1} \\ 0 \\ 0 \\ \text{—} \\ 0 \\ 0 \\ 0 \\ 0 \\ \mathbf{1} \end{array}\right] \begin{array}{c} \text{don't} \\ \text{hate} \\ \textbf{I} \\ \text{it} \\ \text{love} \\ \\ \text{don't} \\ \text{hate} \\ \text{I} \\ \text{it} \\ \textbf{love} \end{array} \qquad \mathbf{r}_1(\mathbf{x}) = \left[\begin{array}{c} 0 \\ 0 \\ 0 \\ 0 \\ \mathbf{1} \\ \text{—} \\ 0 \\ 0 \\ 0 \\ \mathbf{1} \\ 0 \end{array}\right] \begin{array}{c} \text{don't} \\ \text{hate} \\ \text{I} \\ \text{it} \\ \textbf{love} \\ \\ \text{don't} \\ \text{hate} \\ \text{I} \\ \textbf{it} \\ \text{love} \end{array}$$

V={“don’t”, “hate”, “I”, “it”, “love”}

Transformation for text

bow-CNN for text

“I love it.”

$$\mathbf{r}_0(\mathbf{x}) = \begin{array}{c} \left[\begin{array}{c} 0 \\ 0 \\ \mathbf{1} \\ 0 \\ \mathbf{1} \end{array}\right] \end{array} \begin{array}{c} \text{don't} \\ \text{hate} \\ \textbf{I} \\ \text{it} \\ \textbf{love} \end{array} \qquad \mathbf{r}_1(\mathbf{x}) = \begin{array}{c} \left[\begin{array}{c} 0 \\ 0 \\ 0 \\ \mathbf{1} \\ \mathbf{1} \end{array}\right] \end{array} \begin{array}{c} \text{don't} \\ \text{hate} \\ \text{I} \\ \textbf{it} \\ \textbf{love} \end{array}$$

V={“don’t”, “hate”, “I”, “it”, “love”}

More for text

Shallow-CNN

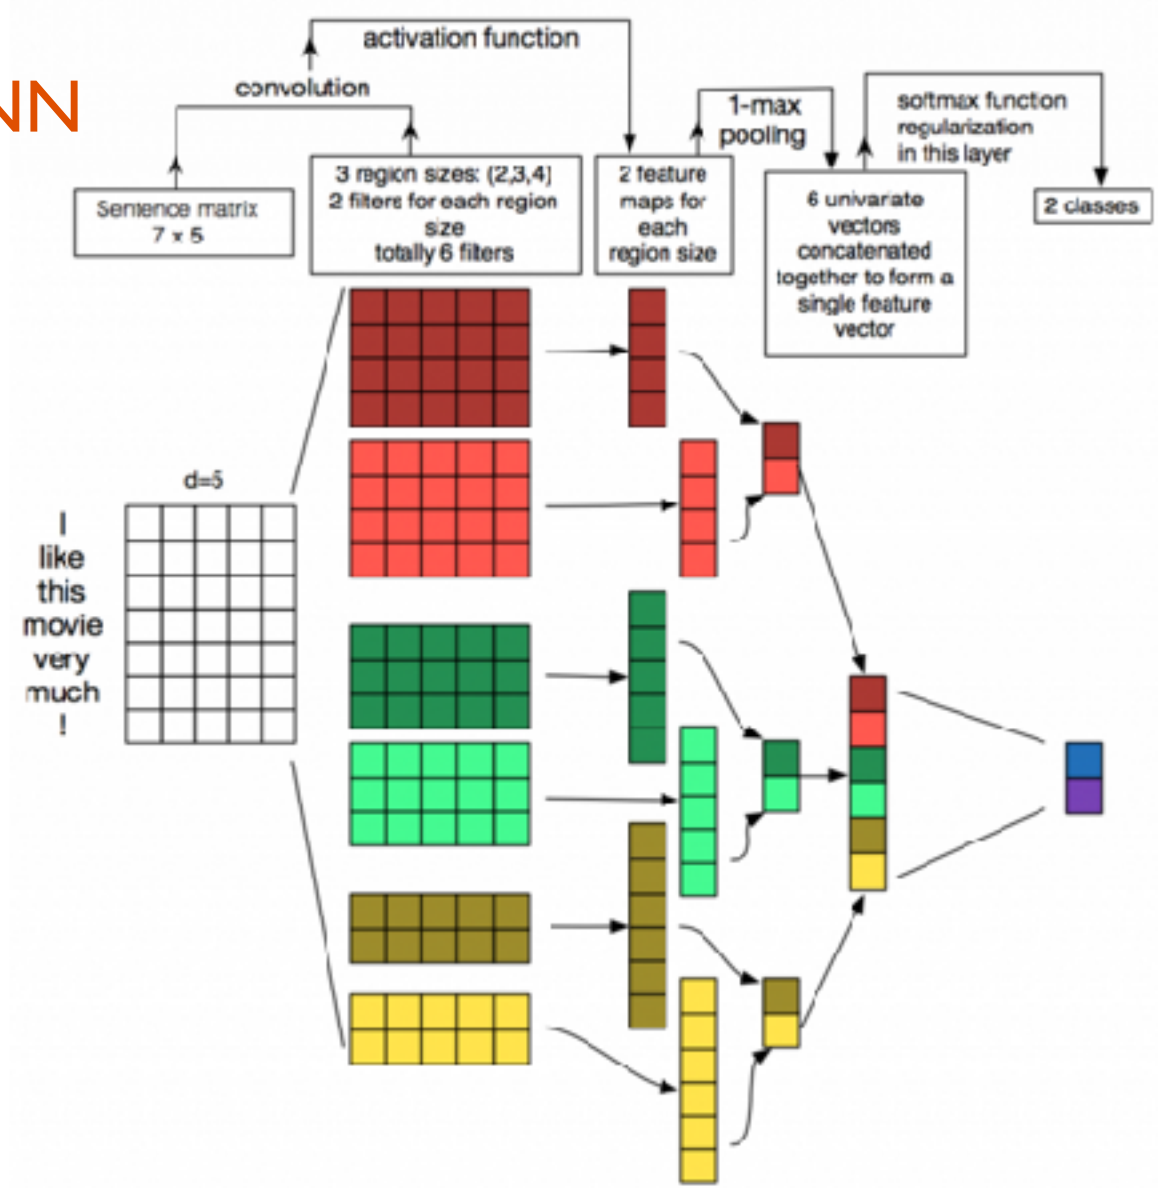

More for text

Deep-CNN

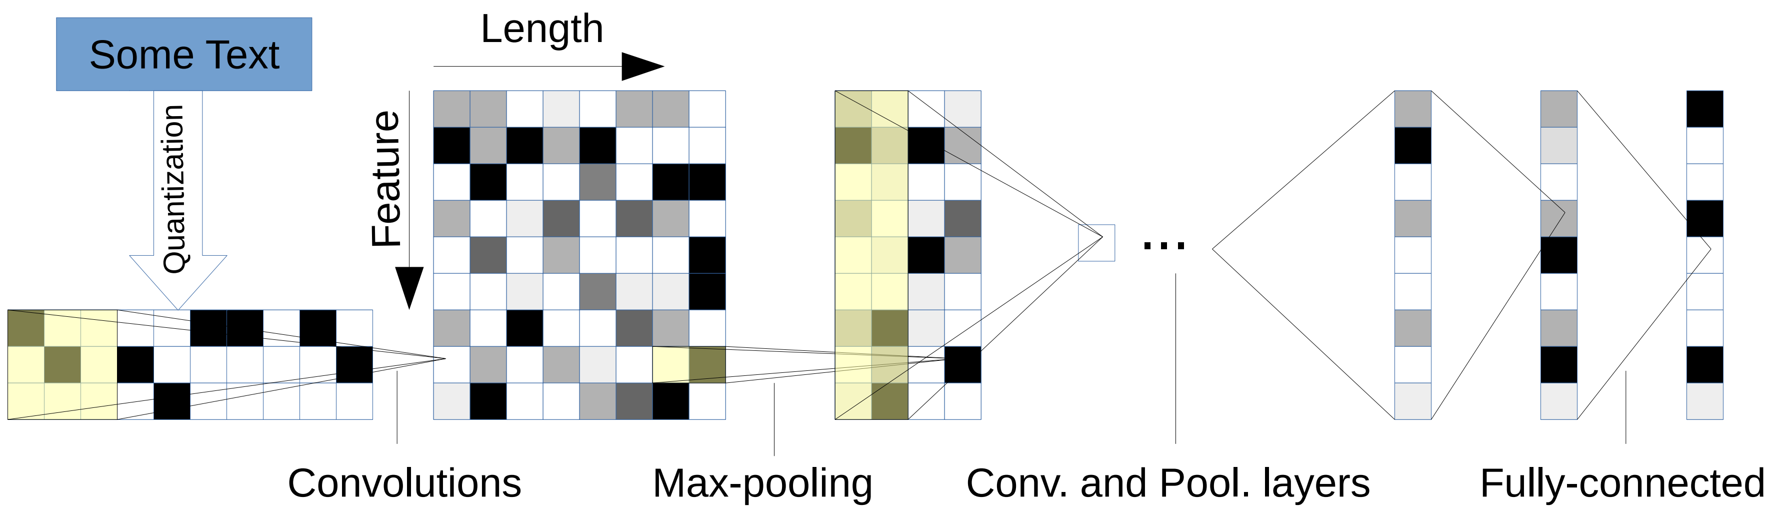

Figure courtesy of [Xiang Zhang et. al, NIPS' 15]

RL in continuous state space

Deep Reinforcement Learning

function approximation by deep neural networks

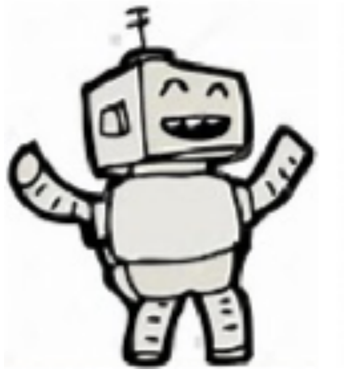

Convolutional neural networks

a powerful neural network architecture for image analysis

- differentiable
- require a lot of samples to train

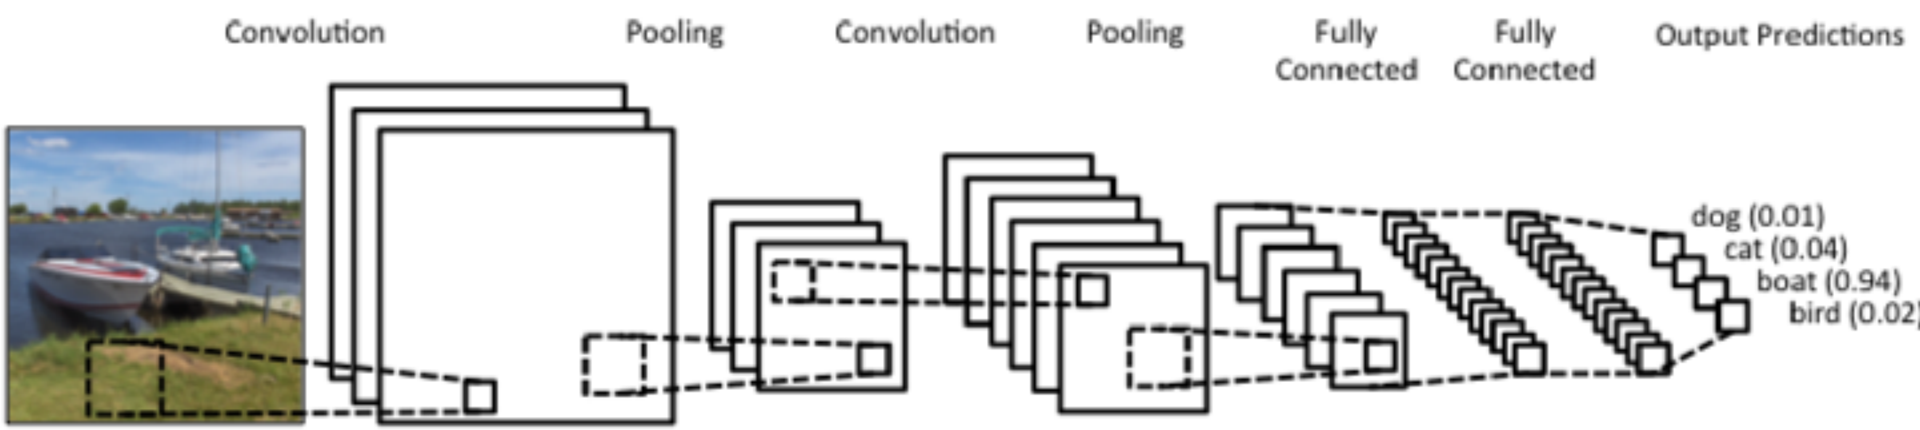

Deep Q-Network

DQN

- using ϵ-greedy policy
- store 1million recent history (s,a,r,s') in replay memory D
- sample a mini-batch (32) from D
- calculate Q-learning target $\tilde{Q}$
- update CNN by minimizing the Bellman error (delayed update)

$$\sum\left(r+\gamma \max_{a'} \tilde{Q}(s', a') - Q_w(s, a)\right)^2$$

DQN on Atari

learn to play from pixels

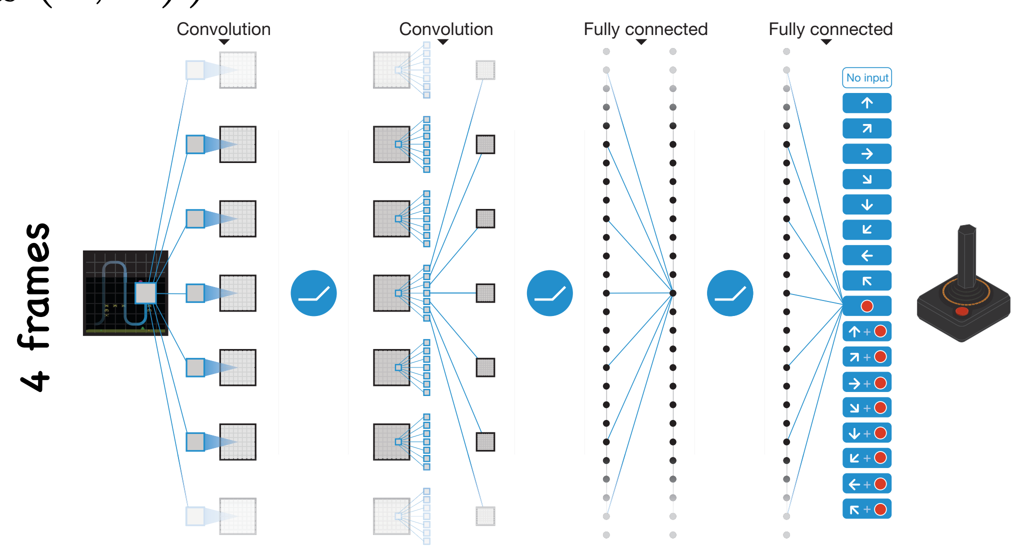

Deep Q-Network

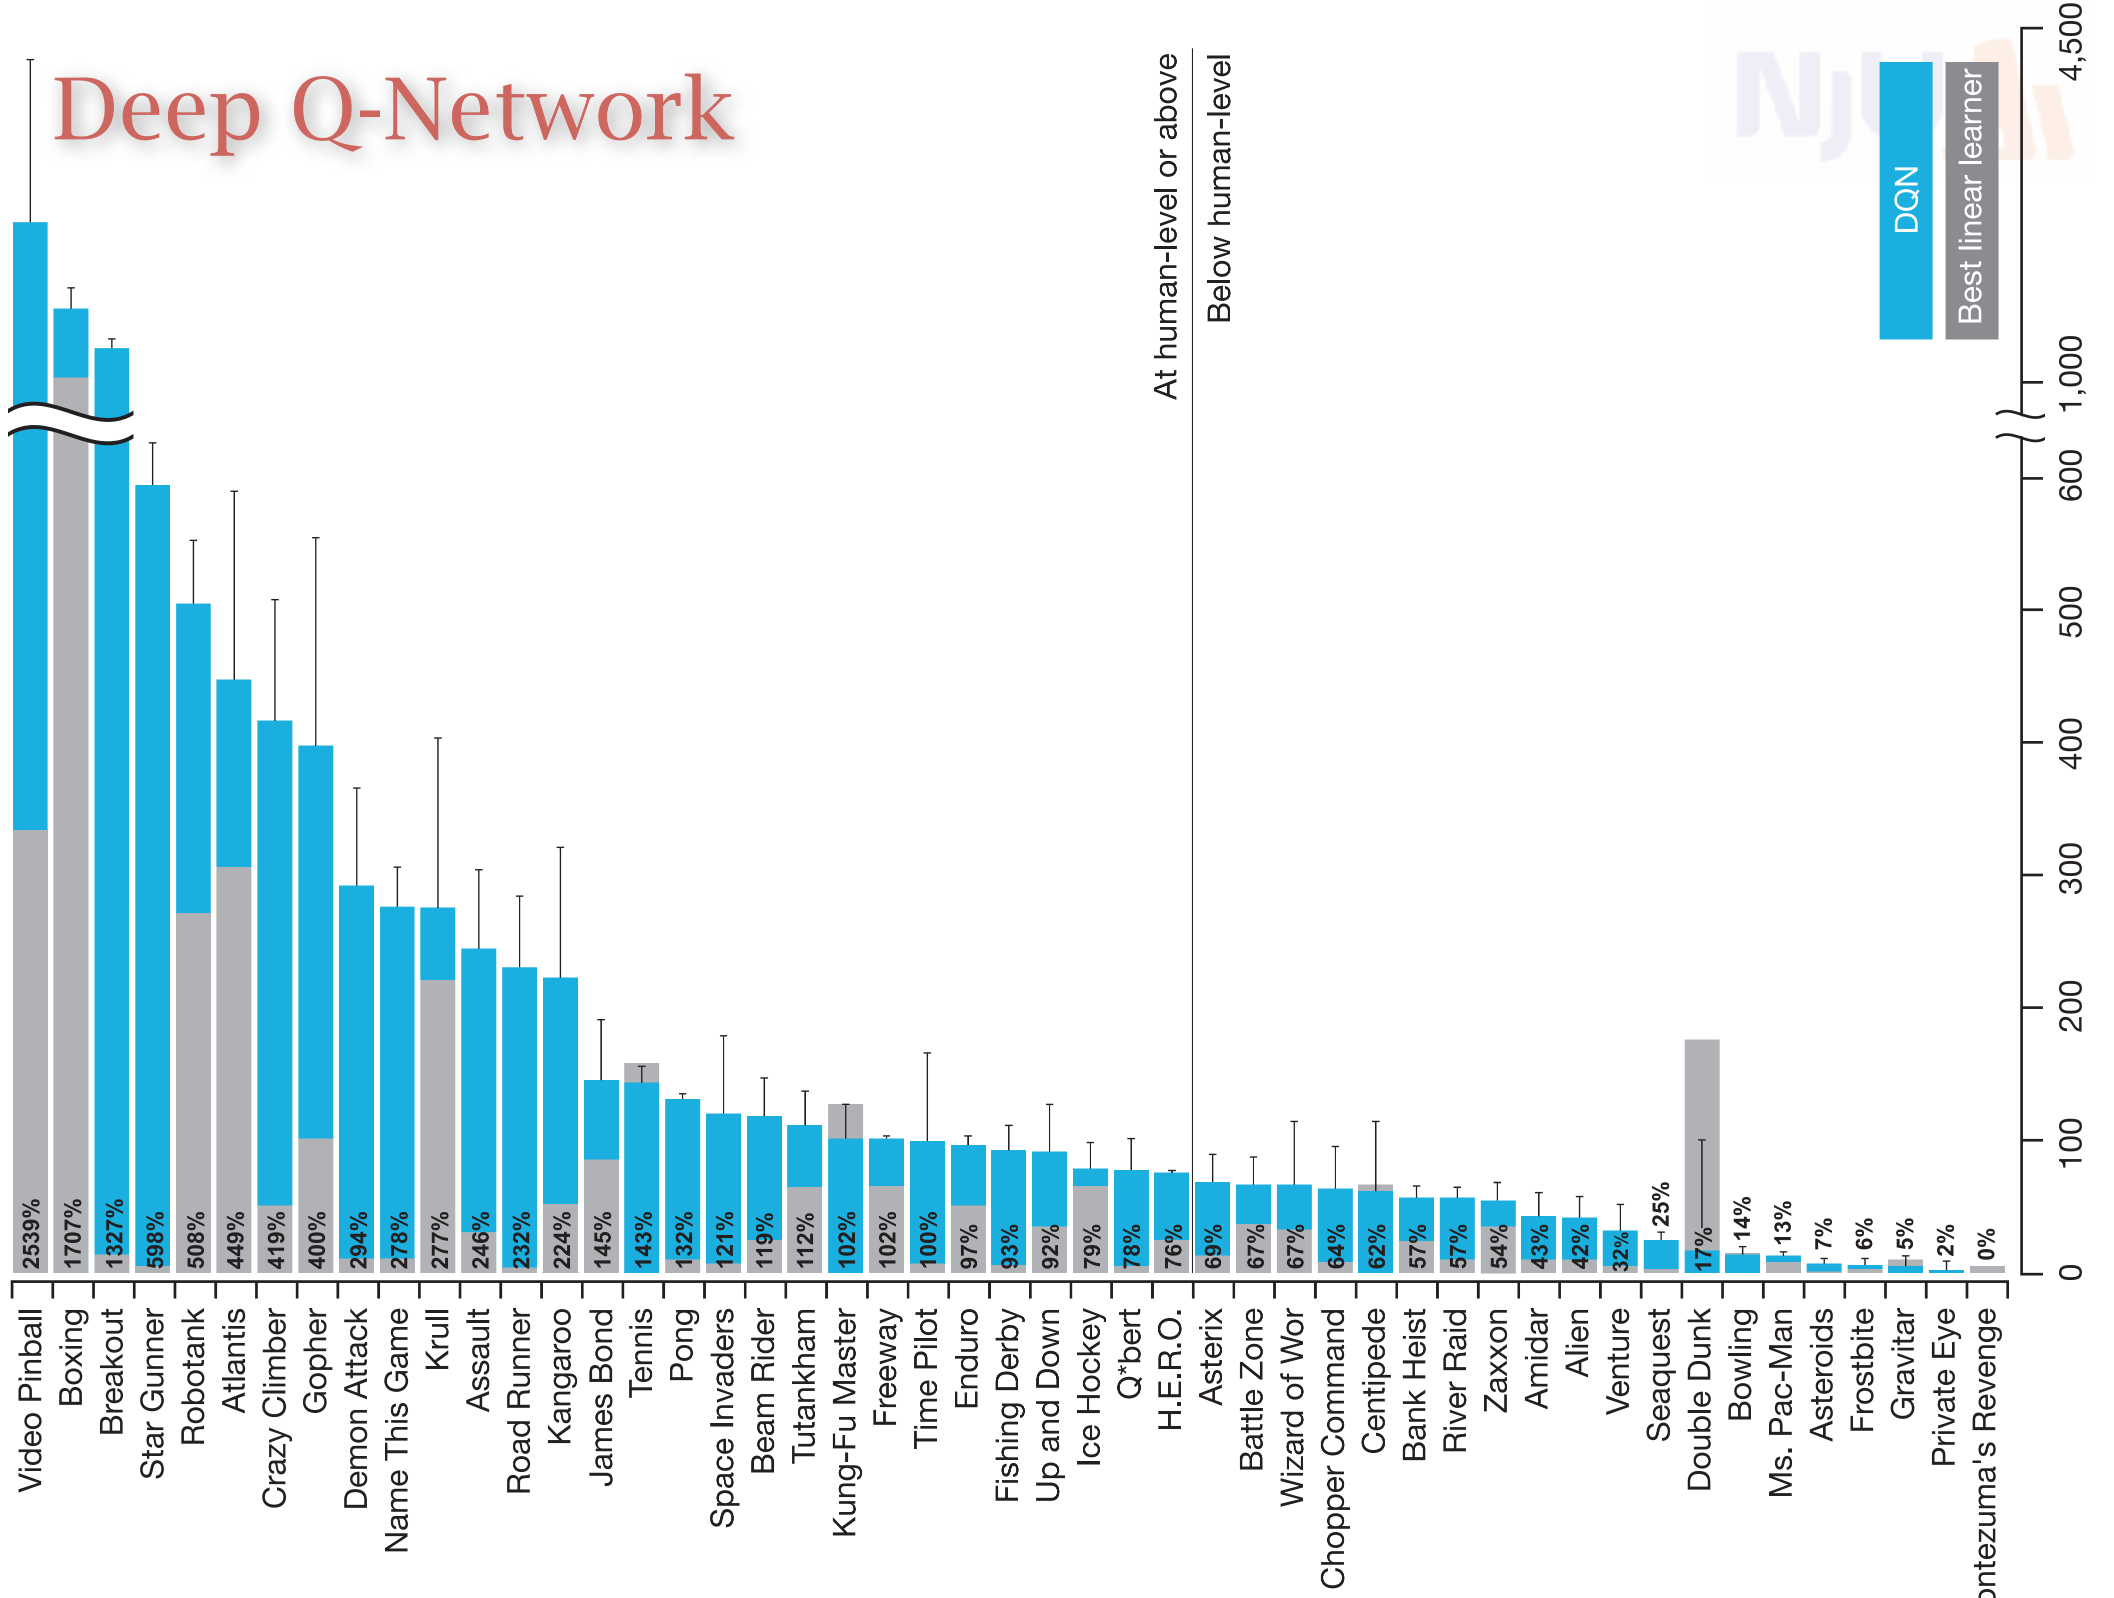

effectiveness

Game	With replay, with target Q	With replay, without target Q	Without replay, with target Q	Without replay, without target Q
Breakout	316.8	240.7	10.2	3.2
Enduro	1006.3	831.4	141.9	29.1
River Raid	7446.6	4102.8	2867.7	1453.0
Seaquest	2894.4	822.6	1003.0	275.8
Space Invaders	1088.9	826.3	373.2	302.0

AlphaGo

A combination of tree search, deep neural networks and reinforcement learning

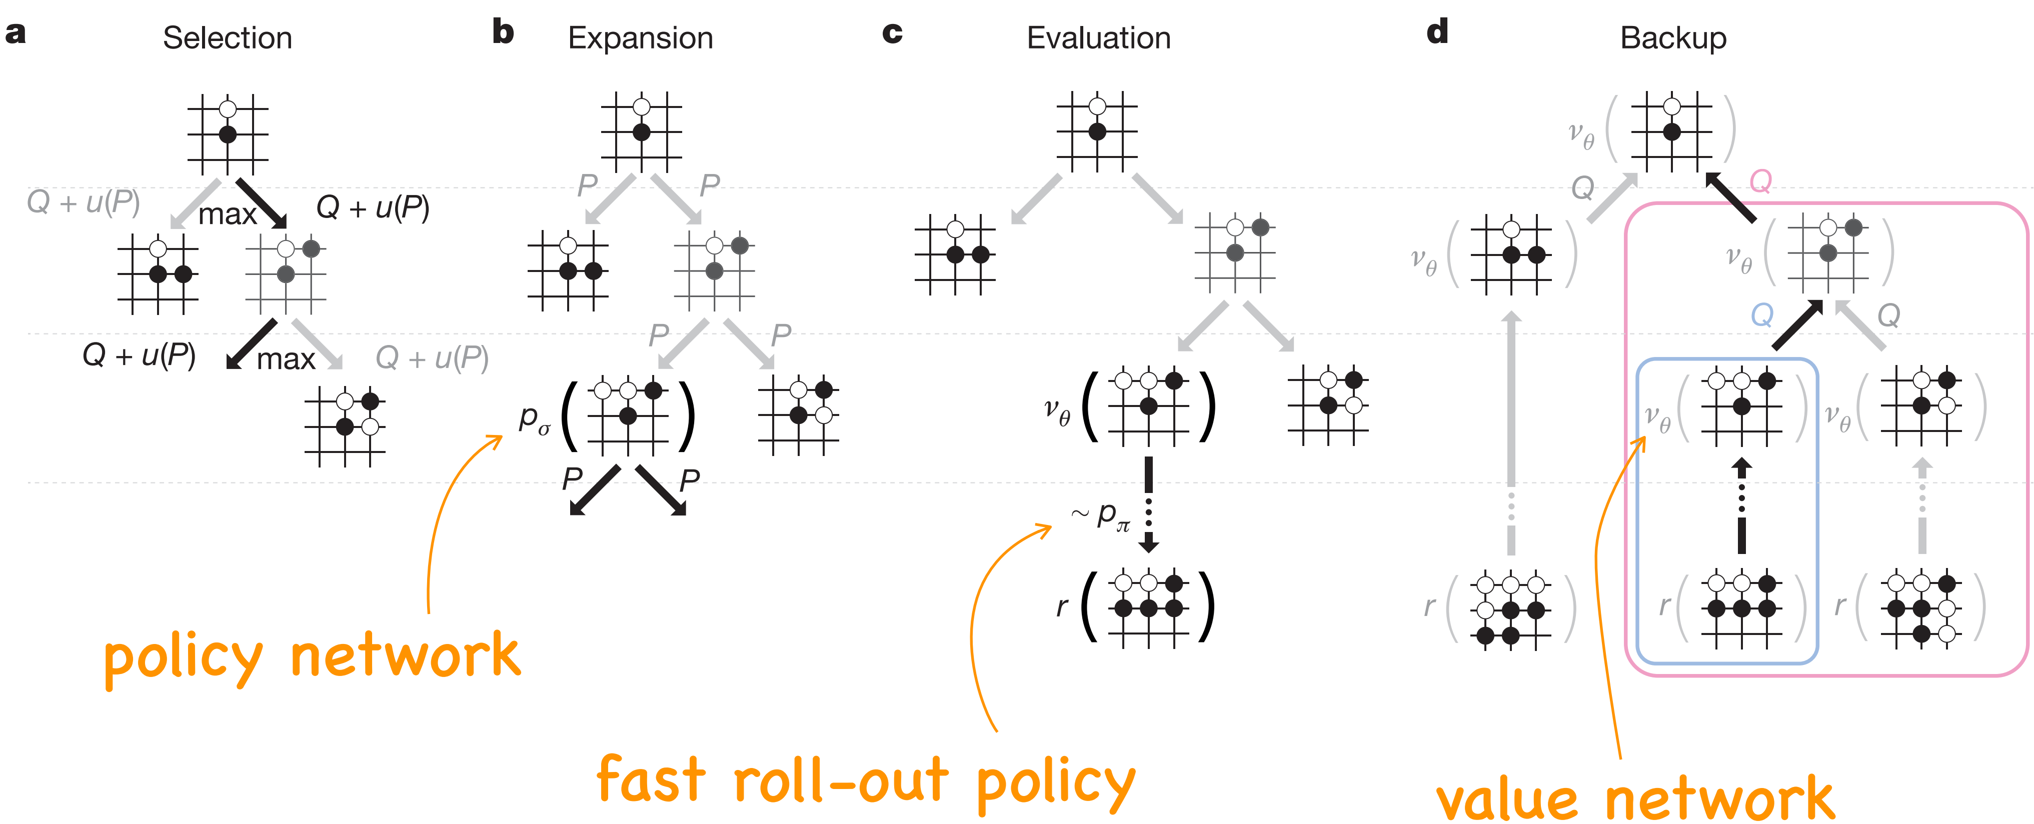

fast roll-out policy:

supervised learning from human v.s. human data

Feature	# of patterns	Description
Response	1	Whether move matches one or more response pattern features
Save atari	1	Move saves stone(s) from capture
Neighbour	8	Move is 8-connected to previous move
Nakade	8192	Move matches a *nakade* pattern at captured stone
Response pattern	32207	Move matches 12-point diamond pattern near previous move
Non-response pattern	69338	Move matches 3×3 pattern around move
Self-atari	1	Move allows stones to be captured
Last move distance	34	Manhattan distance to previous two moves
Non-response pattern	32207	Move matches 12-point diamond pattern centred around move

AlphaGo

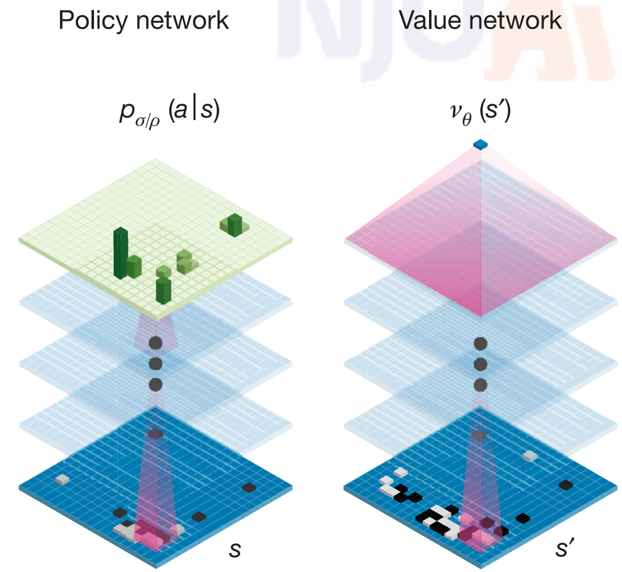

policy network: a CNN output π(s,a)

value network: a CNN output V(s)

Feature	# of planes	Description
Stone colour	3	Player stone / opponent stone / empty
Ones	1	A constant plane filled with 1
Turns since	8	How many turns since a move was played
Liberties	8	Number of liberties (empty adjacent points)
Capture size	8	How many opponent stones would be captured
Self-atari size	8	How many of own stones would be captured
Liberties after move	8	Number of liberties after this move is played
Ladder capture	1	Whether a move at this point is a successful ladder capture
Ladder escape	1	Whether a move at this point is a successful ladder escape
Sensibleness	1	Whether a move is legal and does not fill its own eyes
Zeros	1	A constant plane filled with 0
Player color	1	Whether current player is black

policy network: initialization

supervised learning from human v.s. human data

Architecture			Evaluation				
Filters	Symmetries	Features	Test accuracy %	Train accuracy %	Raw net wins %	*AlphaGo* wins %	Forward time (ms)
128	1	48	54.6	57.0	36	53	2.8
192	1	48	55.4	58.0	50	50	4.8
256	1	48	55.9	59.1	67	55	7.1
256	2	48	56.5	59.8	67	38	13.9
256	4	48	56.9	60.2	69	14	27.6
256	8	48	57.0	60.4	69	5	55.3
192	1	4	47.6	51.4	25	15	4.8
192	1	12	54.7	57.1	30	34	4.8
192	1	20	54.7	57.2	38	40	4.8
192	8	4	49.2	53.2	24	2	36.8
192	8	12	55.7	58.3	32	3	36.8
192	8	20	55.8	58.4	42	3	36.8

AlphaGo

policy network: further improvement

reinforcement learning

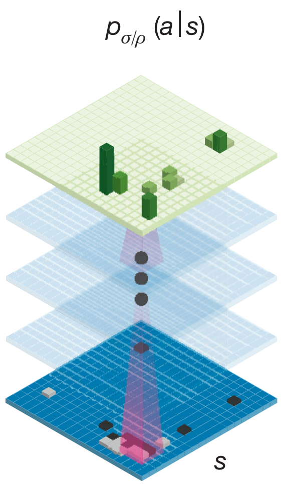

a

s,r

environment

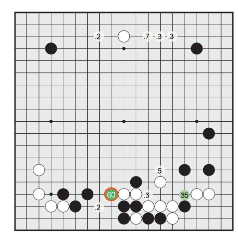

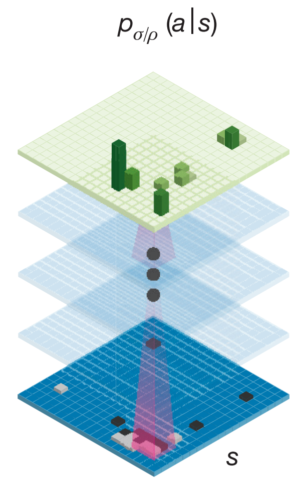

an old policy

a randomly picked policy in previous iterations

a.k.a. self-play

reward:

+1 -- win at terminate state

-1 -- loss at terminate state

AlphaGo

value network: supervised learning from RL data

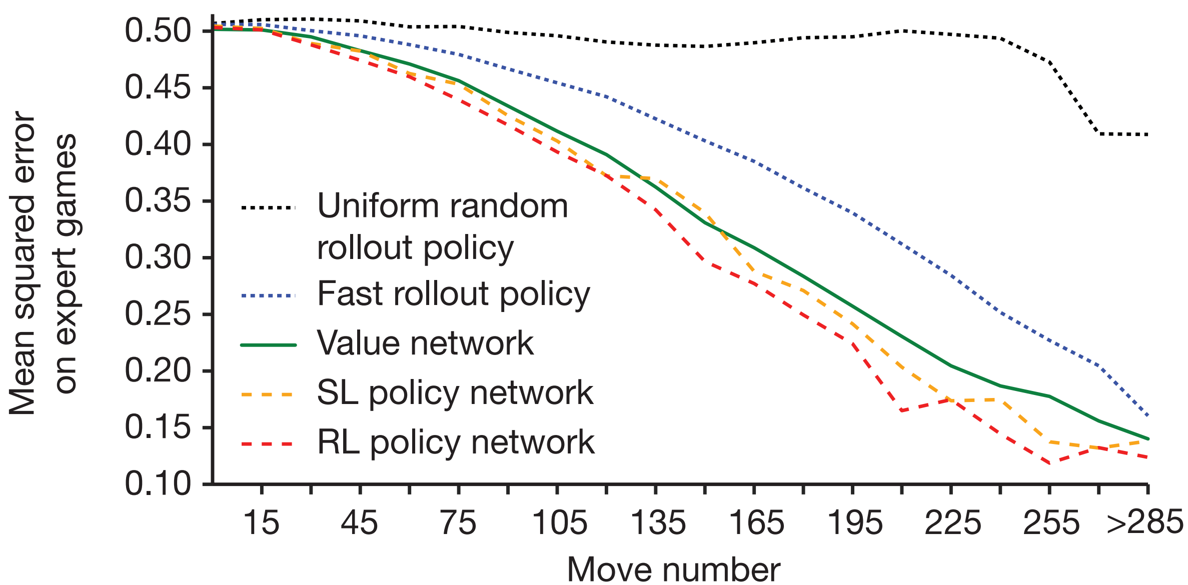

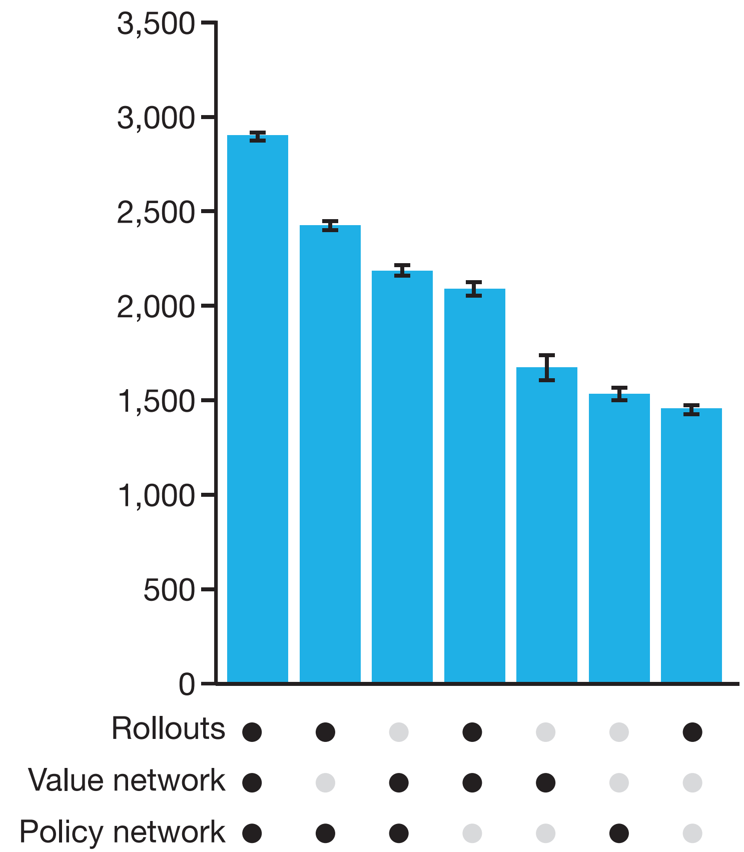


SIMPLE TEST MAP 64X64

What is intelligence?

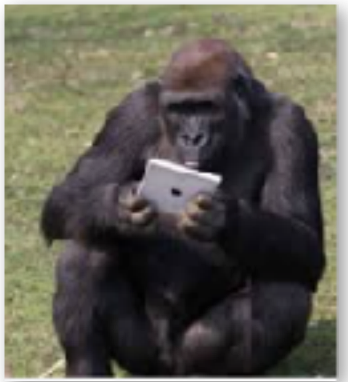

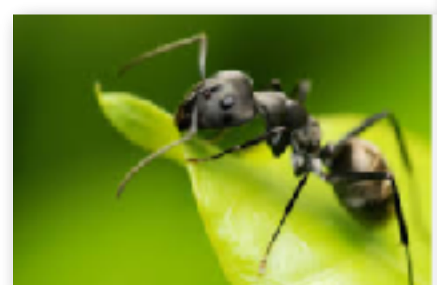

What is intelligence?

The uncertain about intelligence is a fundamental problem of AI

What is AI?

AI is a system that

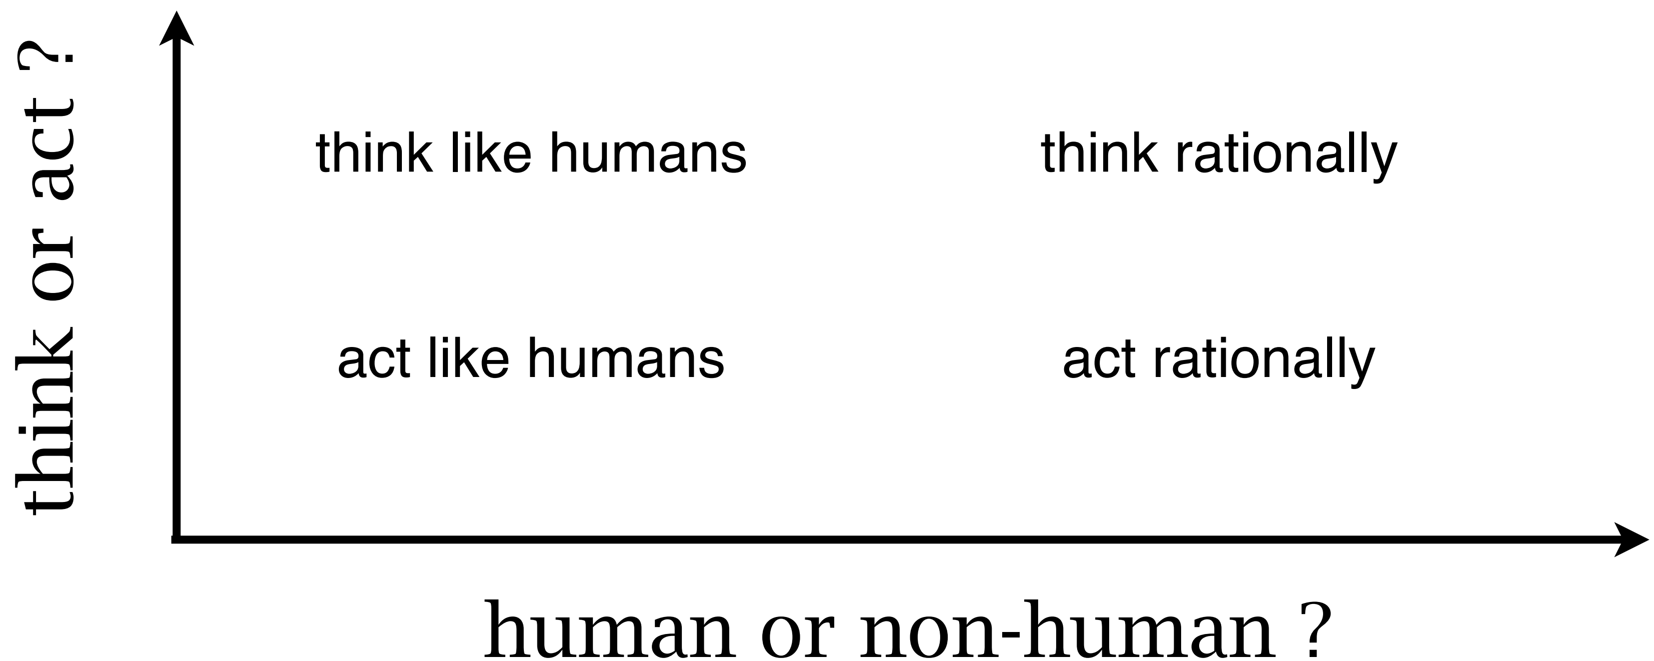

Thinking humanly: Cognitive Science

1960s "cognitive revolution": information-processing psychology replaced prevailing orthodoxy of behaviorism

Requires scientific theories of internal activities of the brain
- What level of abstraction? "Knowledge" or "circuits"?
- How to validate? Requires
  1) Predicting and testing behavior of human subjects (top-down)
  or 2) Direct identification from neurological data (bottom-up)

Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Both share with AI the following characteristic:
the available theories do not explain (or engender) anything resembling human-level general intelligence

Hence, all three fields share one principal direction!

Acting humanly: The Turing test

Turing (1950) "Computing machinery and intelligence":

◇ "Can machines think?" ⟶ "Can machines behave intelligently?"

◇ Operational test for intelligent behavior: the Imitation Game

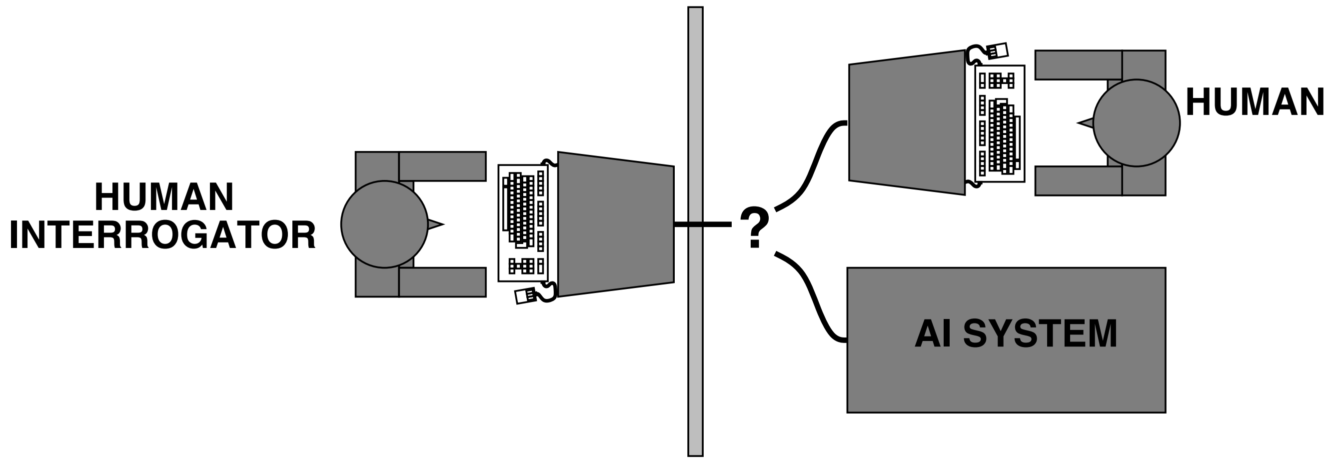

◇ Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes

◇ Anticipated all major arguments against AI in following 50 years

◇ Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problem: Turing test is not **reproducible**, **constructive**, or amenable to **mathematical analysis**

Thinking rationally: Laws of Thought

Normative (or prescriptive) rather than descriptive

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of logic:
    notation and **rules of derivation** for thoughts;
may or may not have proceeded to the idea of mechanization

Direct line through mathematics and philosophy to modern AI

Problems:
1) Not all intelligent behavior is mediated by logical deliberation
2) What is the purpose of thinking? What thoughts **should** I have
    out of all the thoughts (logical or otherwise) that I **could** have?

Acting rationally

Rational behavior: doing the right thing

The right thing: that which is expected to maximize goal achievement, given the available information

Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

Aristotle (Nicomachean Ethics):

> **Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good**

What is AI?

AI is a system that

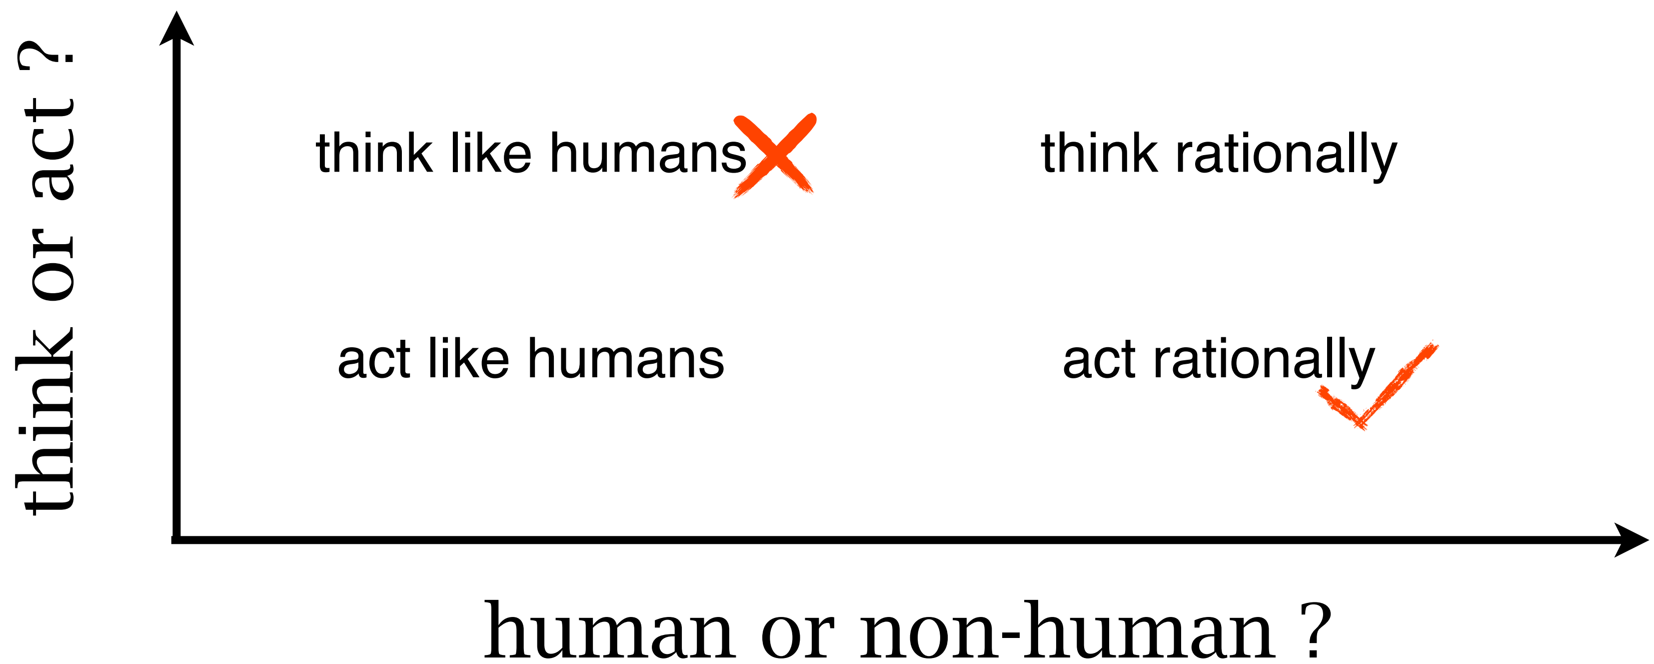

AI IS BLOOMING

HOPE YOU ENJOY

THANK YOU ALL!