# Principles of neural network design

Francois Belletti, CS294 RISE

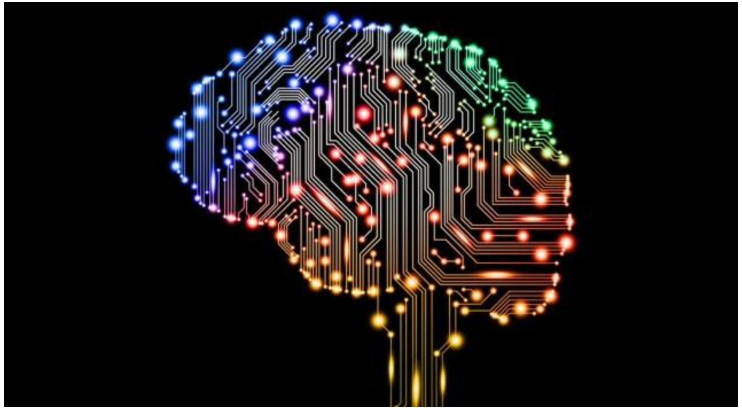

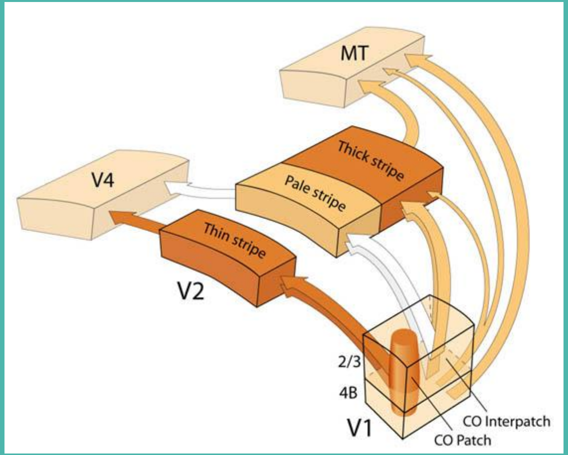

MT
V4
Thick stripe
Pale stripe
Thin stripe
V2
2/3
4B
V1
CO Interpatch
CO Patch


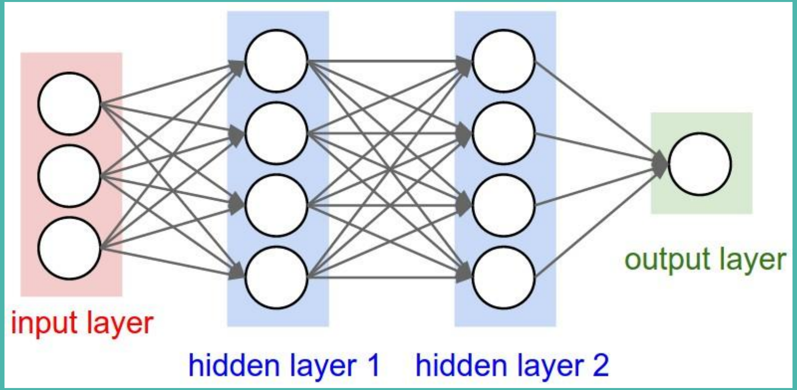

input layer
hidden layer 1
hidden layer 2
output layer

# Human brains as metaphors of statistical models

### *Biological analogies*

The visual cortex of mammals

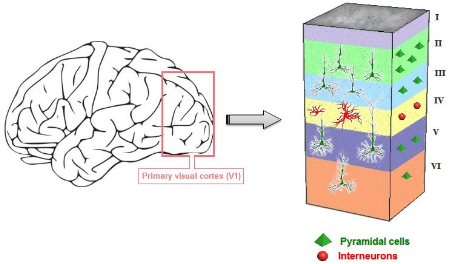


Multiple sensing channels

Memory and attention

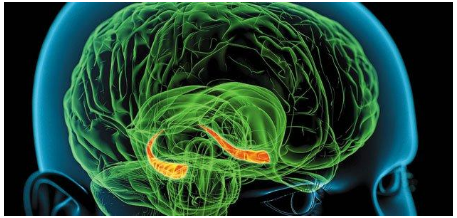

### *Machine learning instantiations*

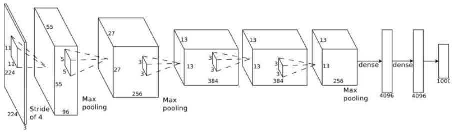


Deep convolutional neural networks

Multimodal neural networks

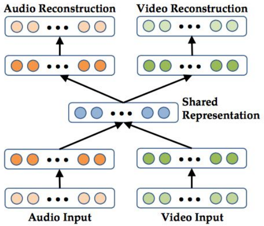


LSTMs and GRUs

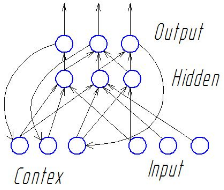

# Neural Networks For Computer Vision

# Neural Networks in Computer Vision

Neural networks for classification of handwritten digits

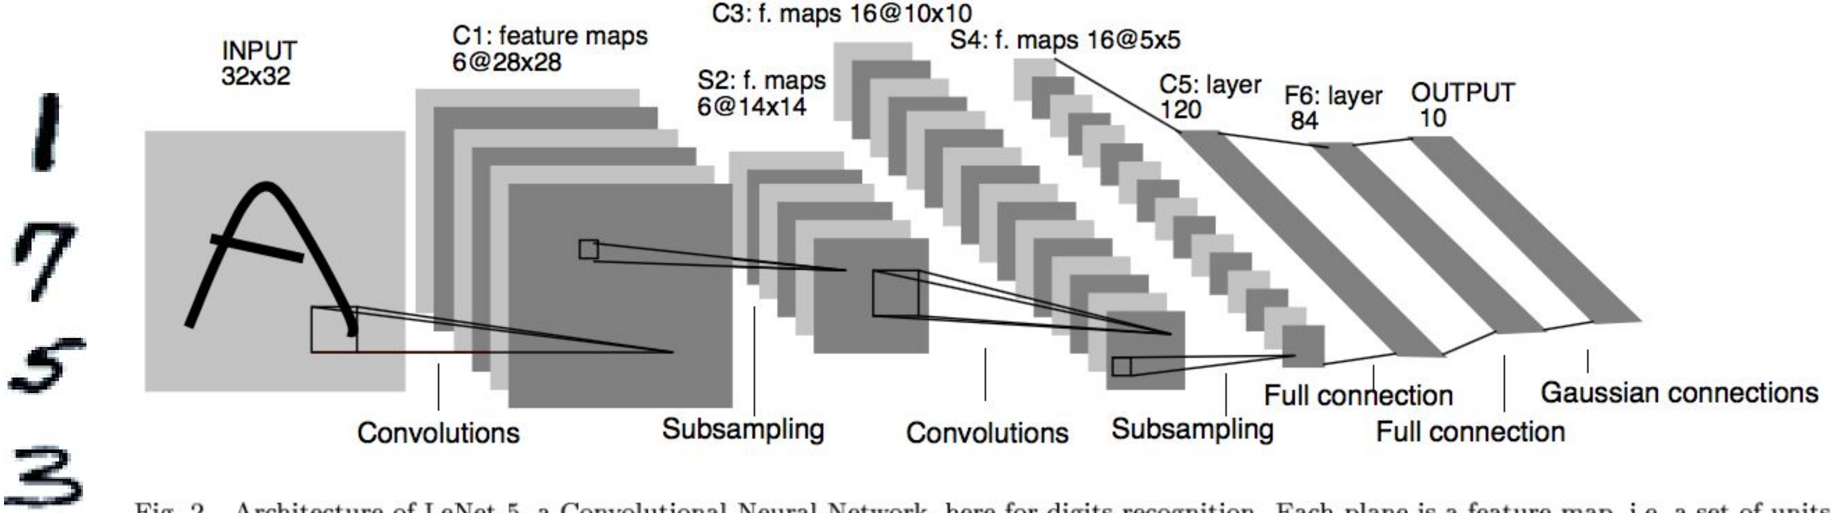


Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

# Learning Mechanism: Correction of Mistakes

Nature used a single tool to get to today's success: mistake

$$W_k = W_{k-1} - \epsilon \frac{\partial E^{p_k}(W)}{\partial W}$$

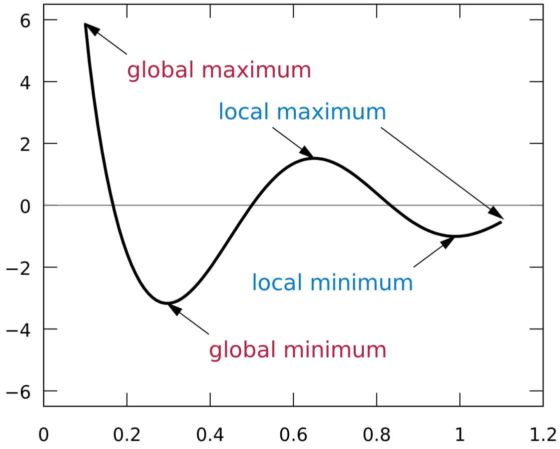


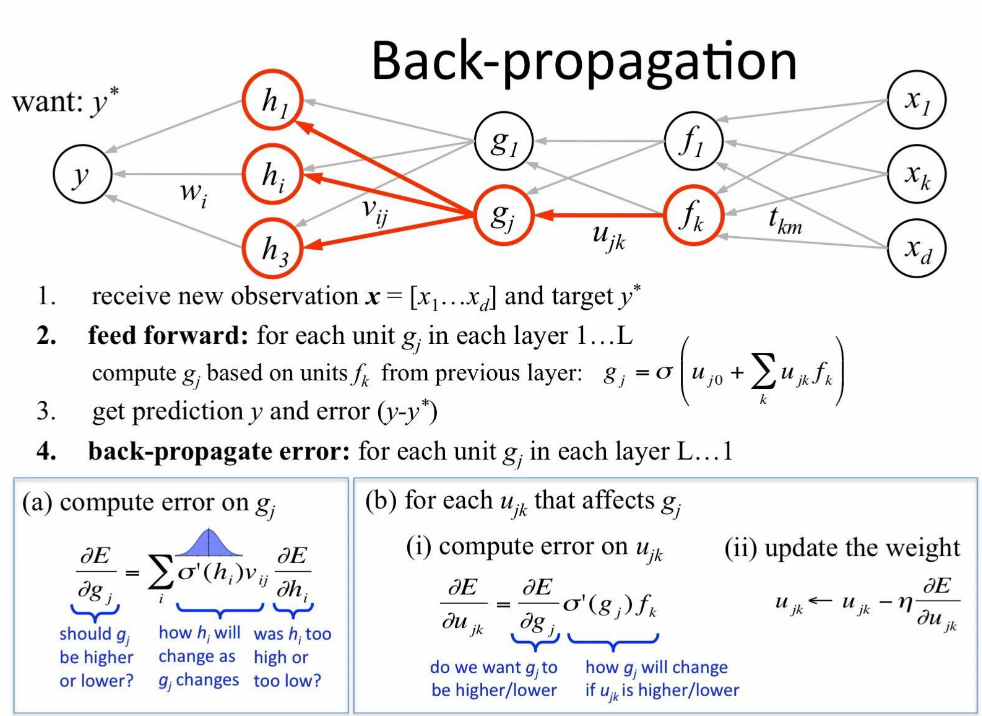

# Modularity Is Back-Prop's Perk for Software Eng.

Back-propagation is a recursive algorithm

This module provides `function()`, commonly accessed as `theano.function`, **the interface for compiling graphs into callable objects.**

You've already seen example usage in the basic tutorial... something like this:

```
>>> x = theano.tensor.dscalar()
>>> f = theano.function([x], 2*x)
>>> print f(4)  # prints 8.0
```

input

output

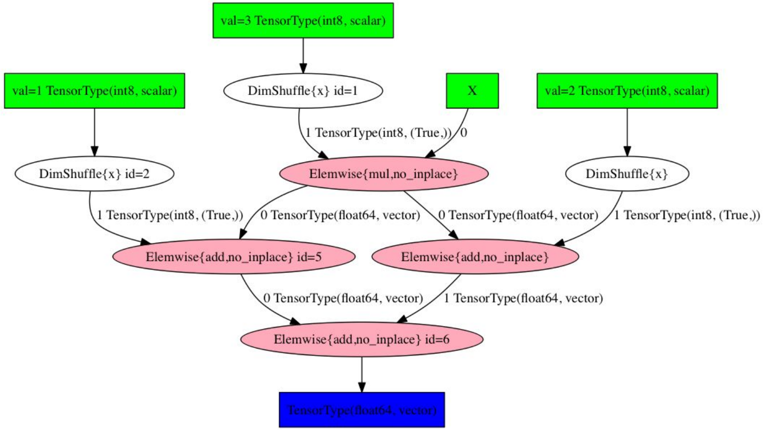

# Image Classification

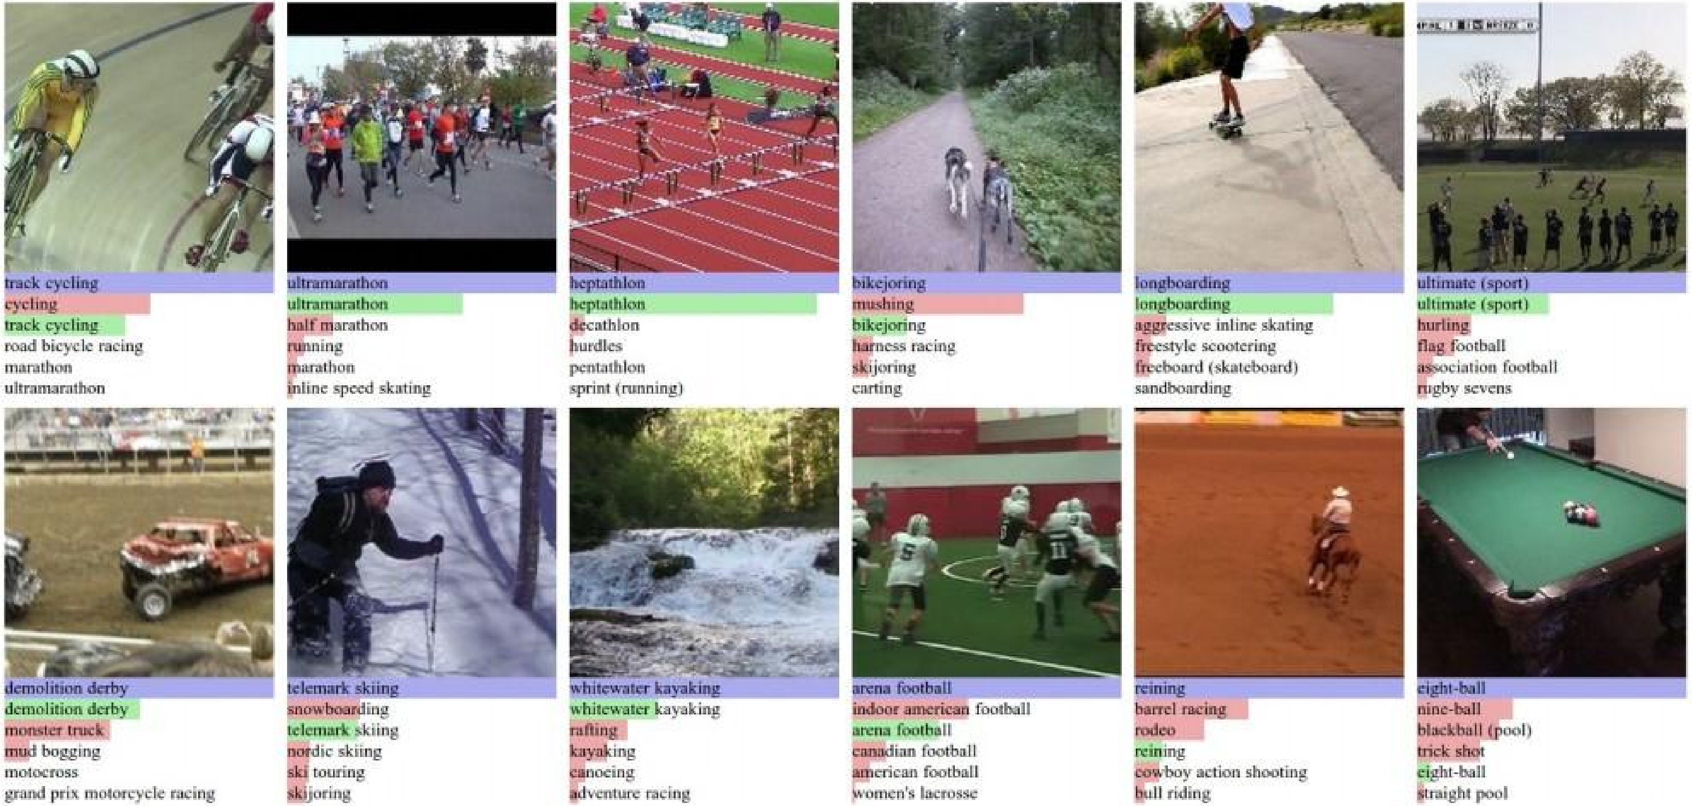

# Successful Architecture In Computer Vision

An example of a wide network: AlexNet

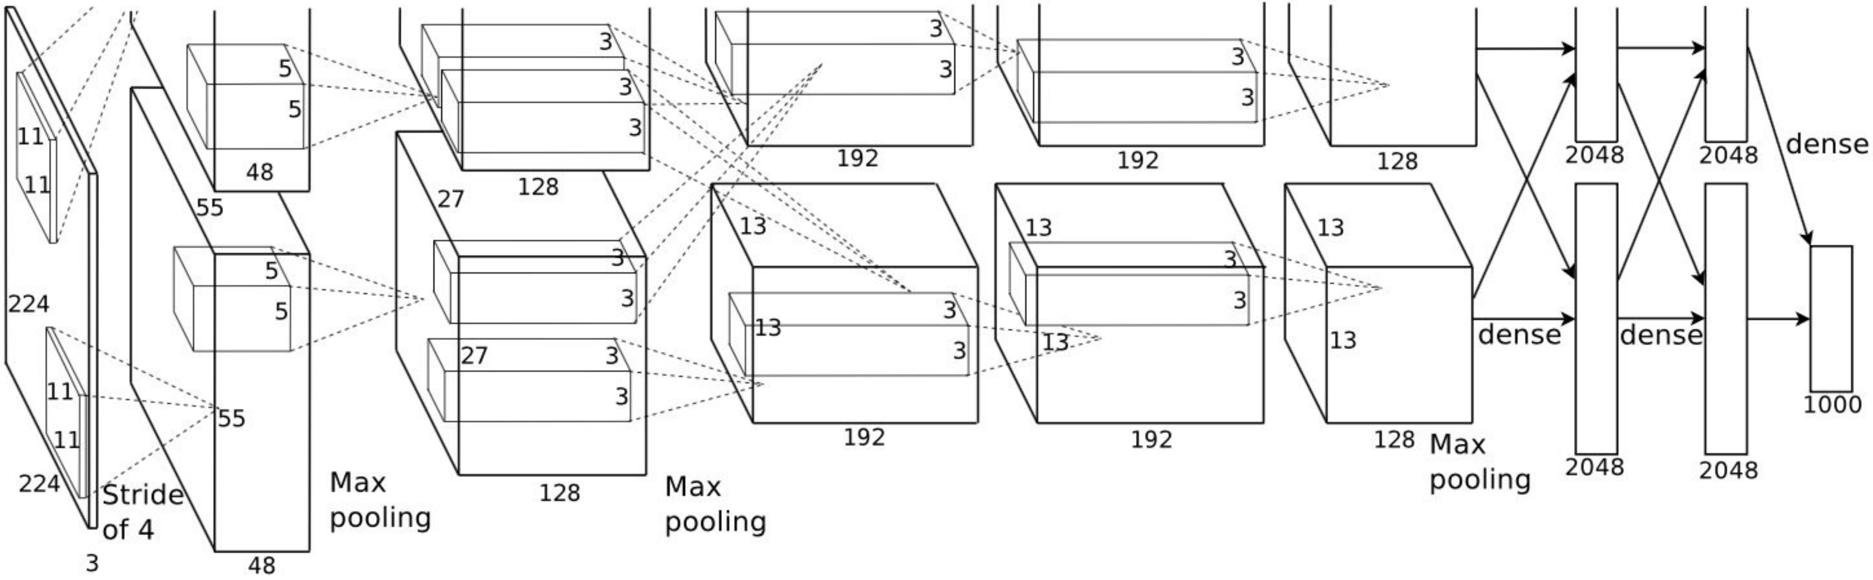

# Understanding What Happens Within A Deep NN

Examining convolution filter banks

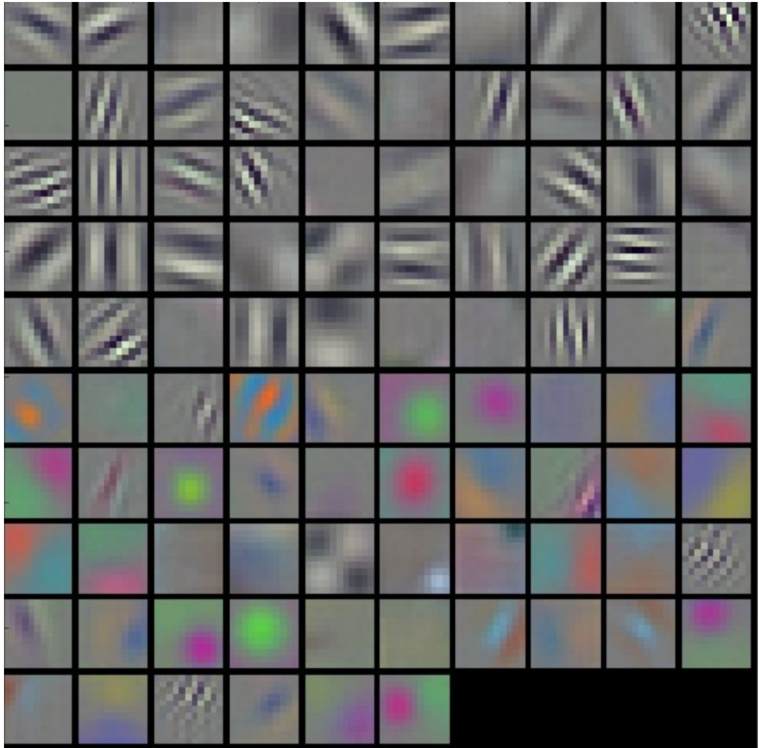

Examining activations

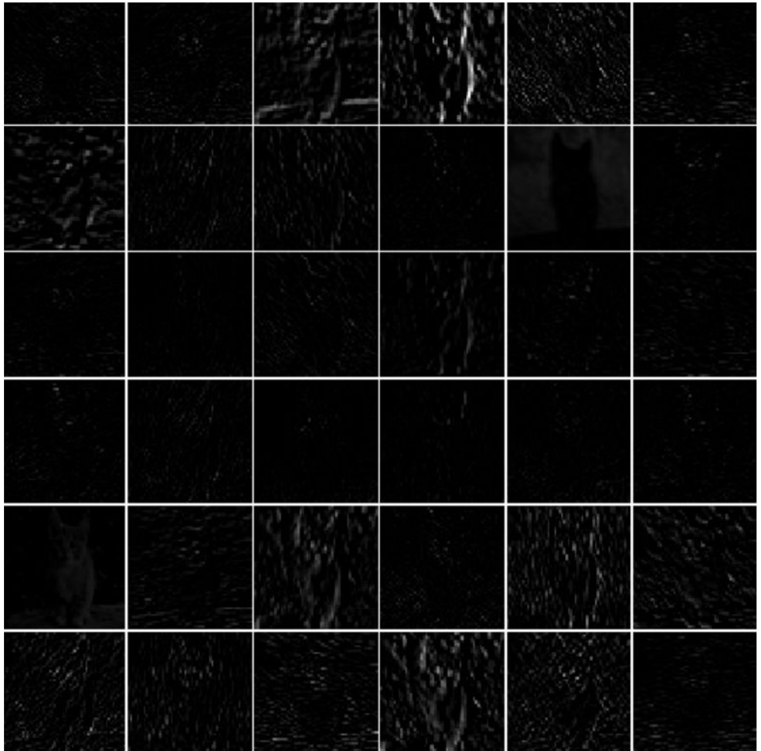

# Determining A Neuron’s Speciality

Images that triggered the highest activations of a neuron:

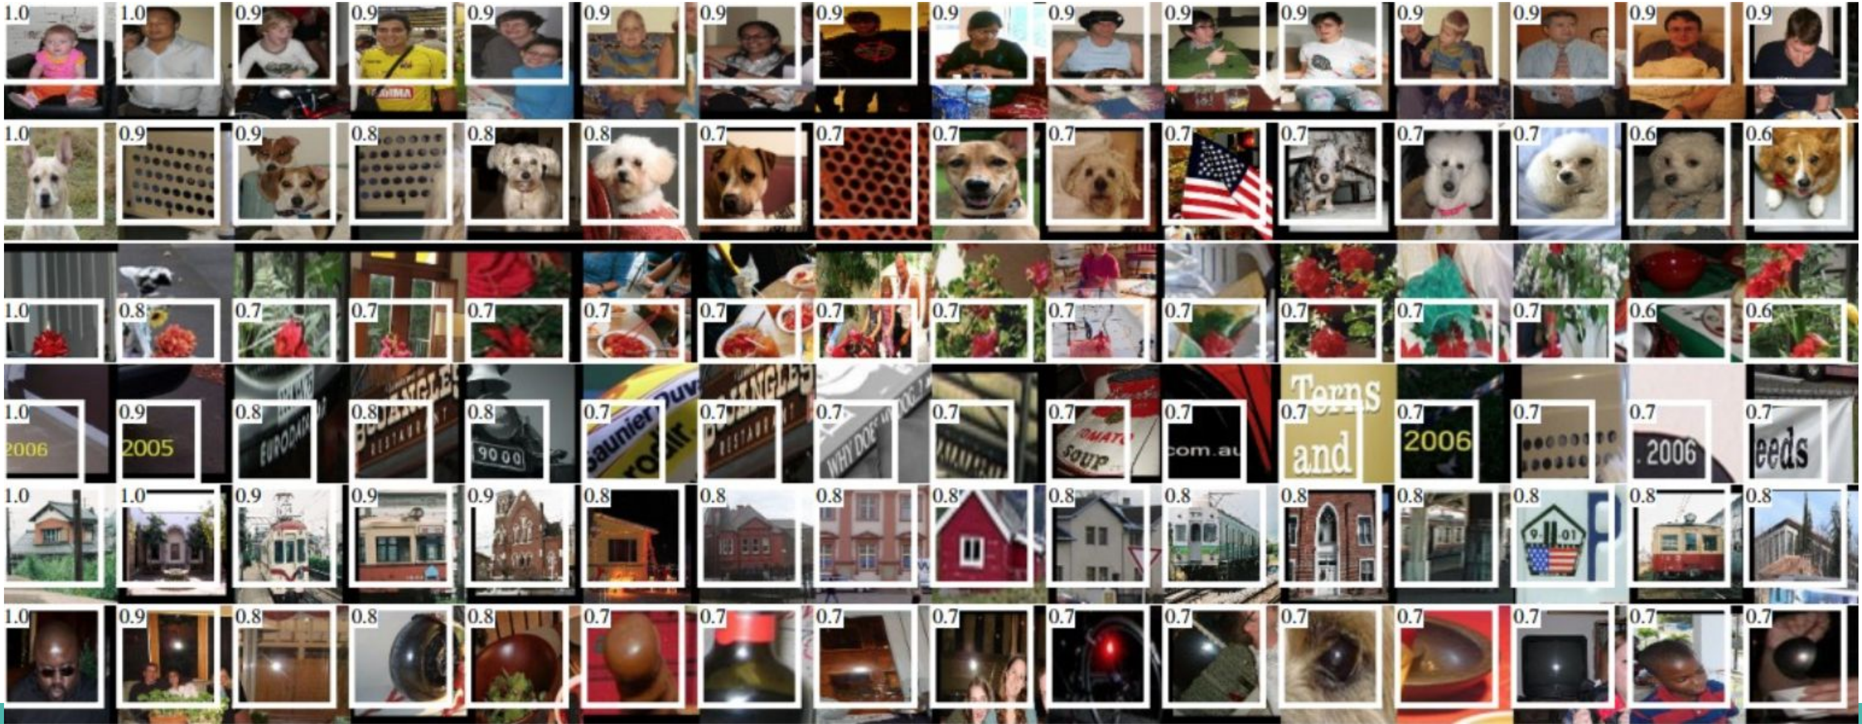

# Another Successful Architecture For CV


Inception:

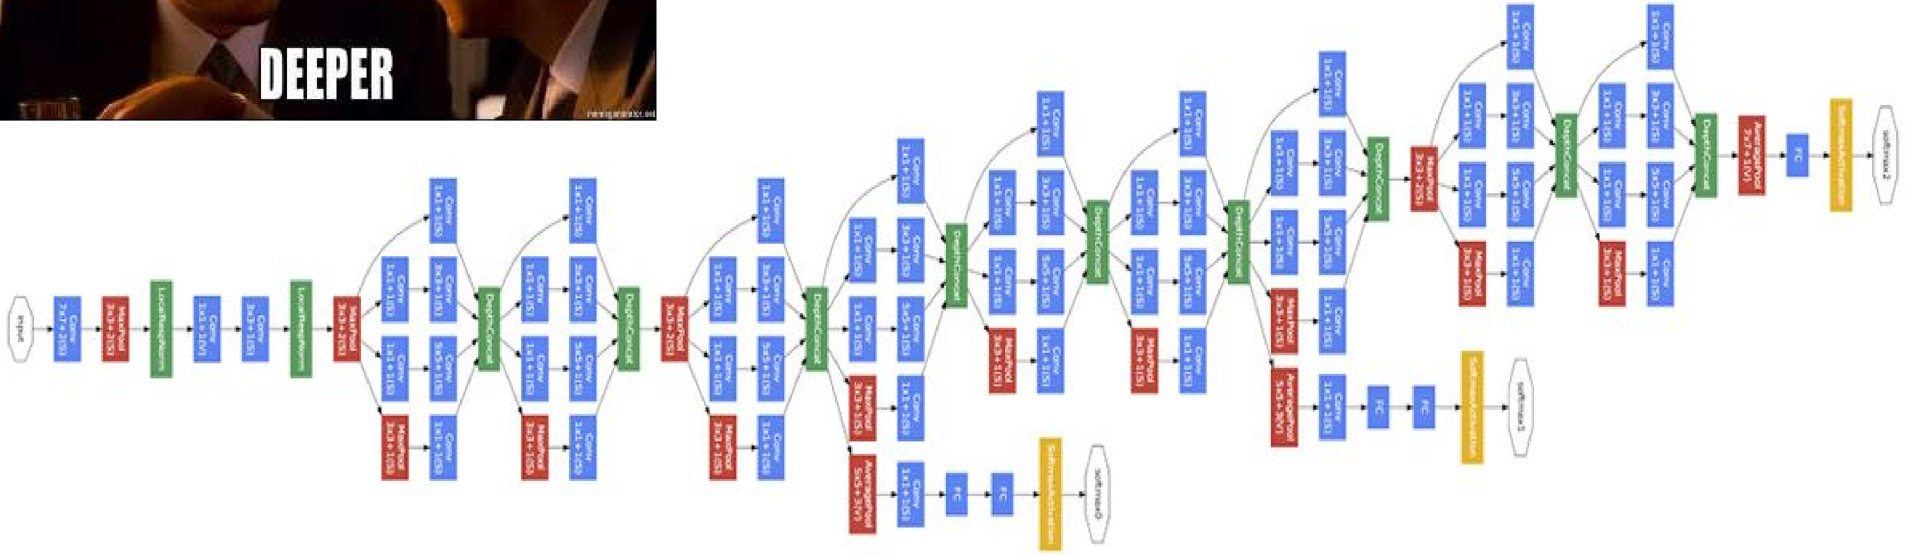

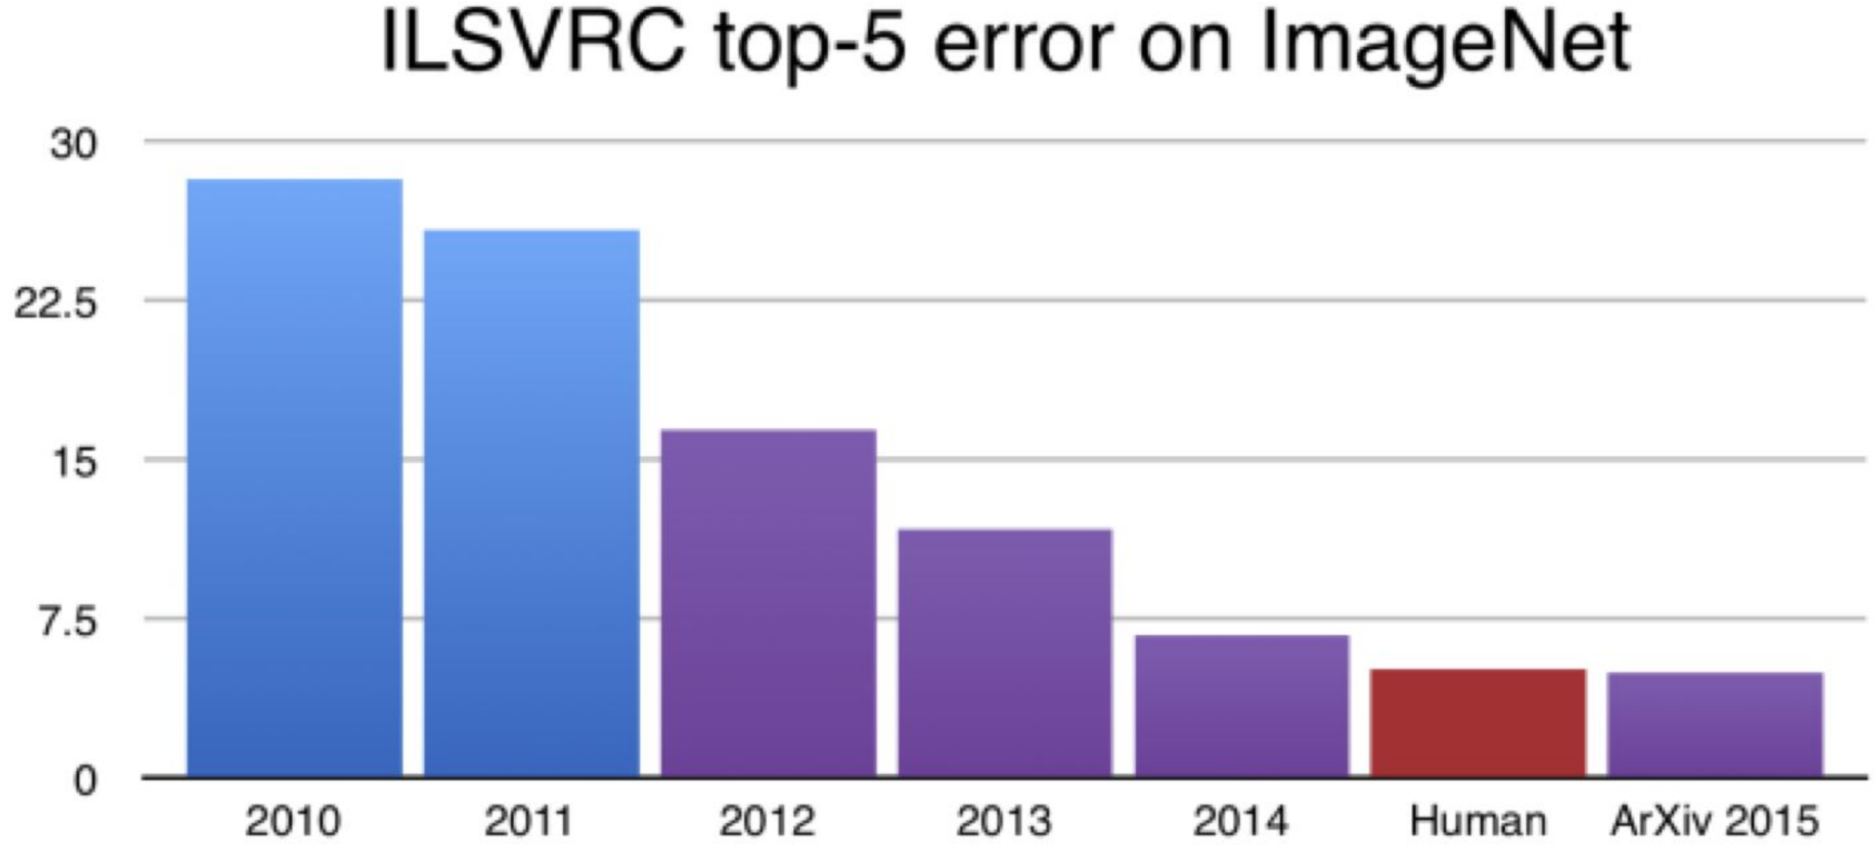
ILSVRC top-5 error on ImageNet
30
22.5
15
7.5
0
2010
2011
2012
2013
2014
Human
ArXiv 2015

# Recurrent Architectures

# Learning To Leverage Context

Memory in Recurrent Architectures: LSTM (Long Short Term Memory Network)

Input x, output y, context c (memory)

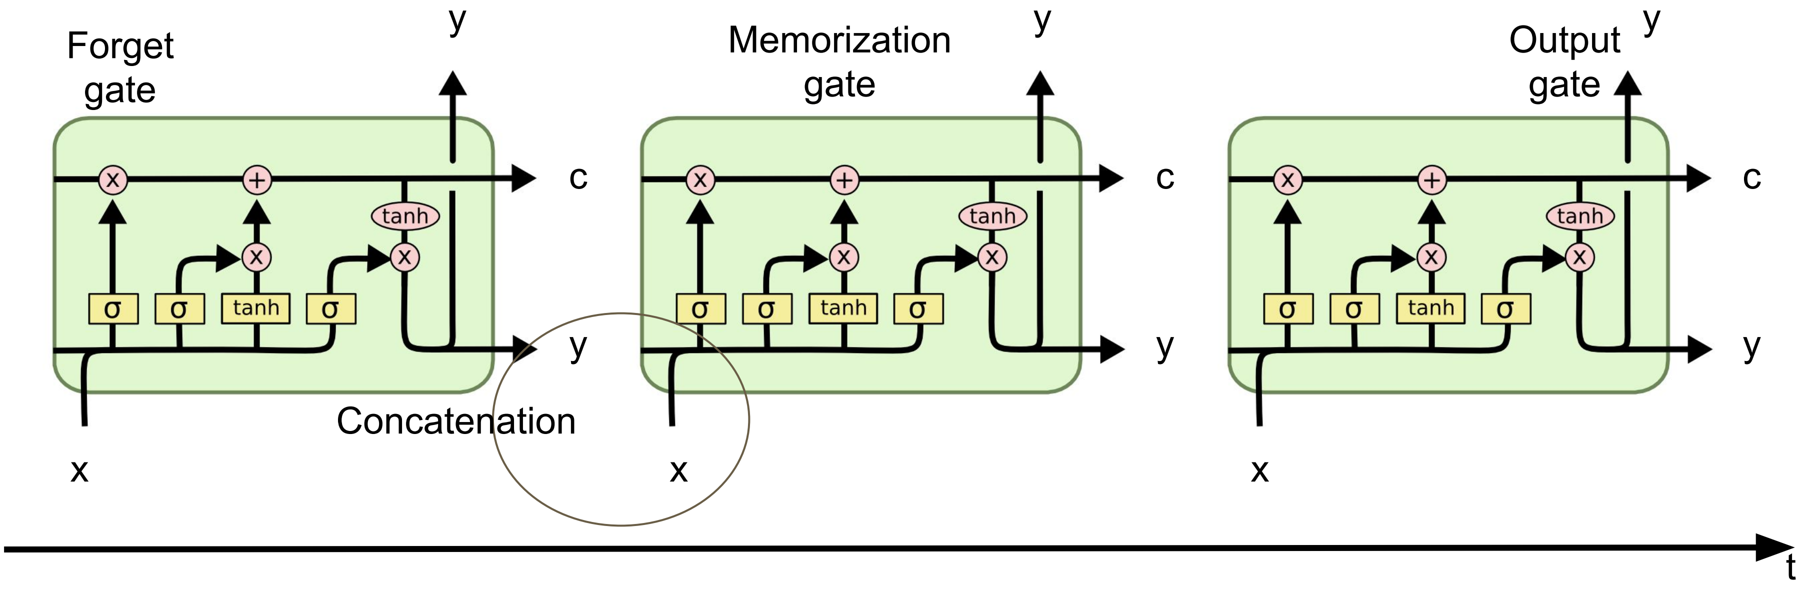

# Other recurrent architectures

Gated recurrent units:

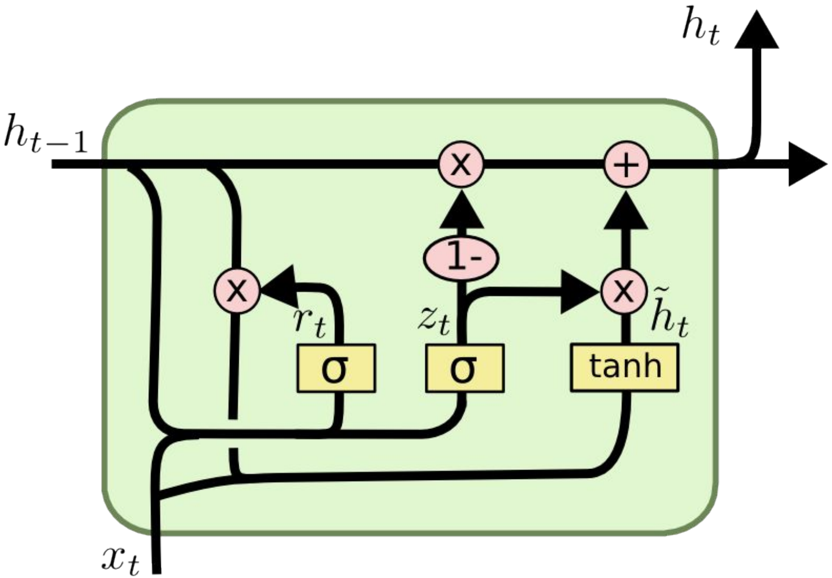


$$z_t = \sigma\left(W_z \cdot [h_{t-1}, x_t]\right)$$

$$r_t = \sigma\left(W_r \cdot [h_{t-1}, x_t]\right)$$

$$\tilde{h}_t = \tanh\left(W \cdot [r_t * h_{t-1}, x_t]\right)$$

$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t$$

# Why Is Context Important?

In language, most grammars are not context free

End-to-end translation, Alex Graves

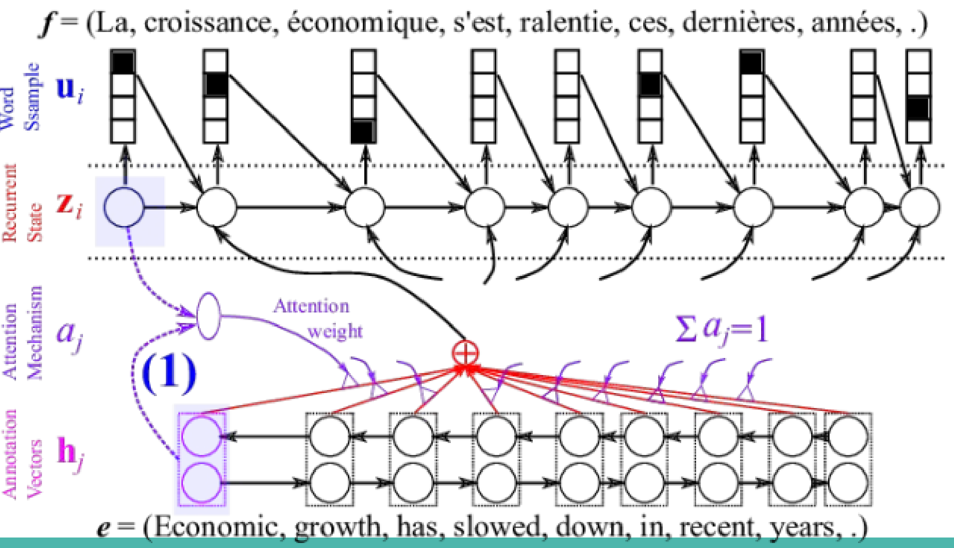


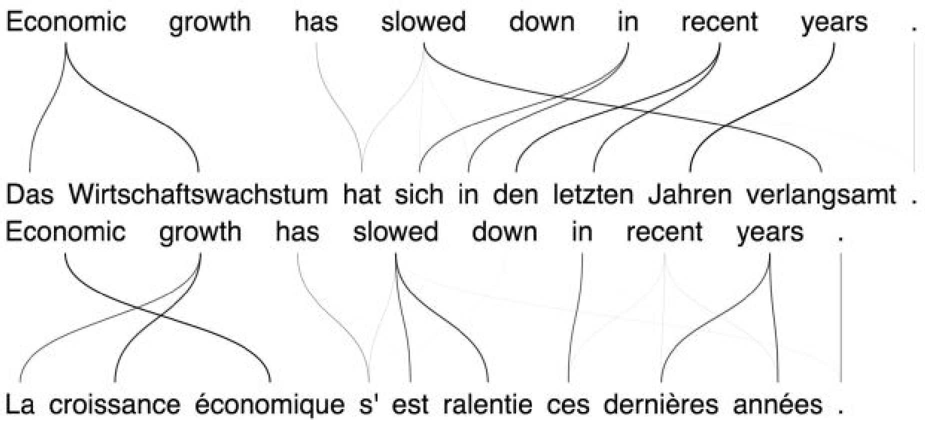

# Context Is Also Important In Control

Remembering what just happened is important for decision making


# Memory is necessary for localization

Latest experiment in asynchronous deep RL: LSTMS for maze running

Memory comes at a cost: a lot of RAM or VRAM is necessary

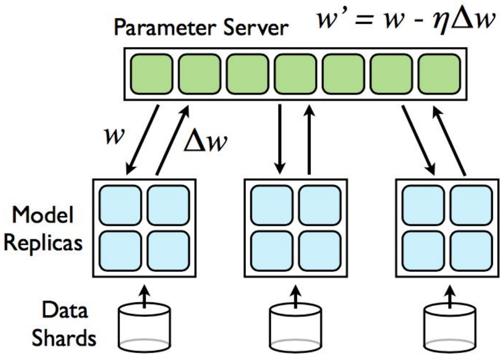

# Conclusion: the distributed brain

# Interaction is crucial in enabling AI

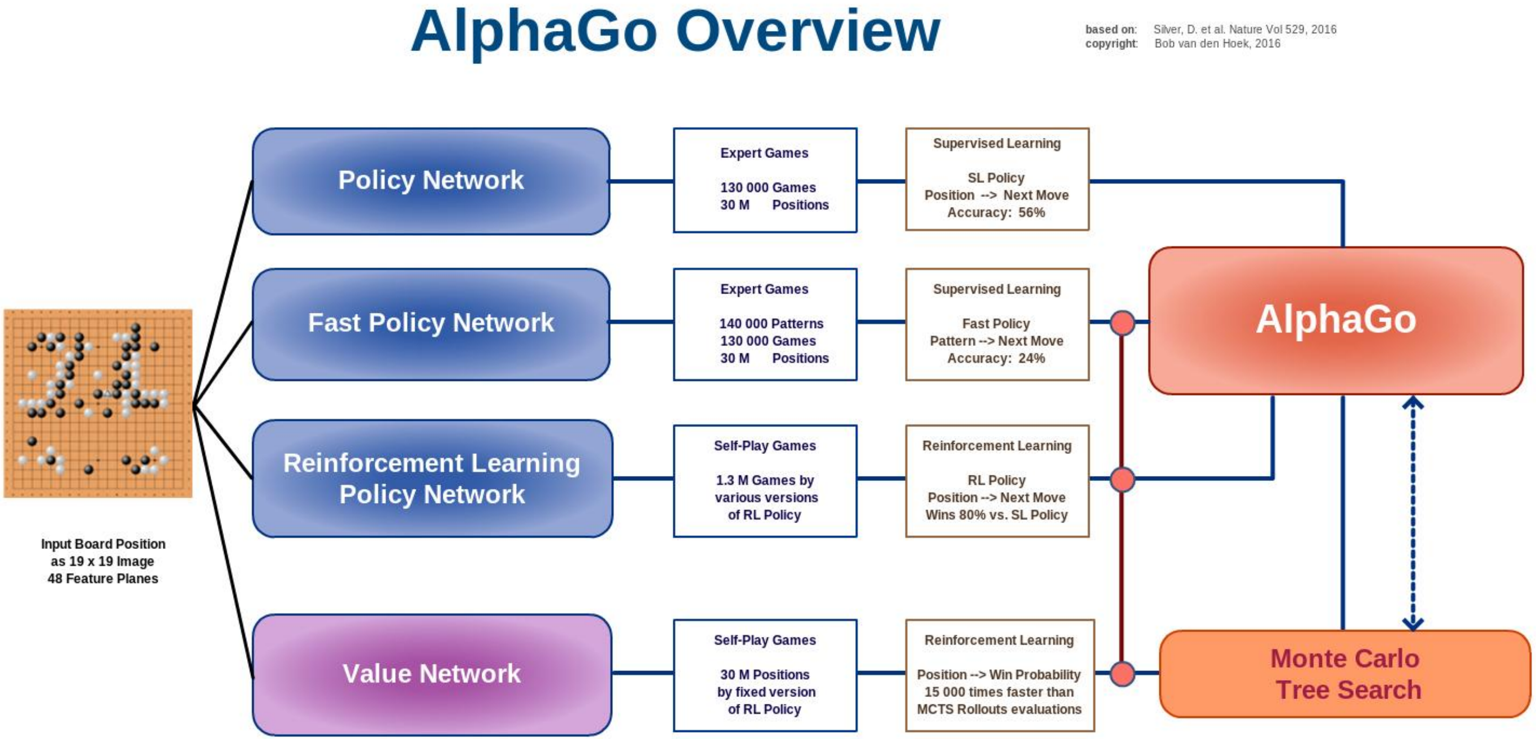

# Playing versus computers before beating humans

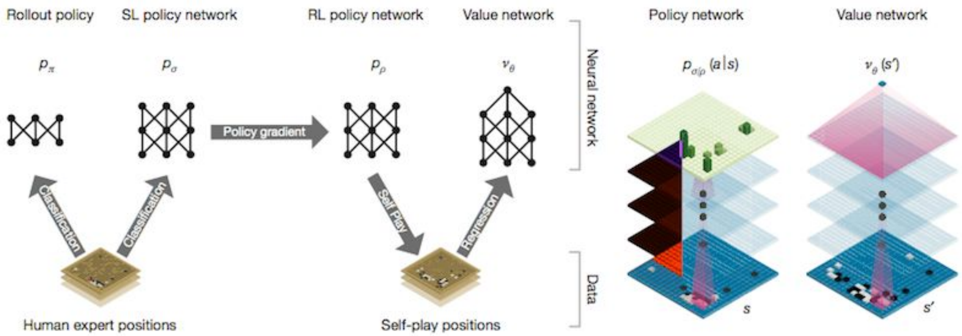


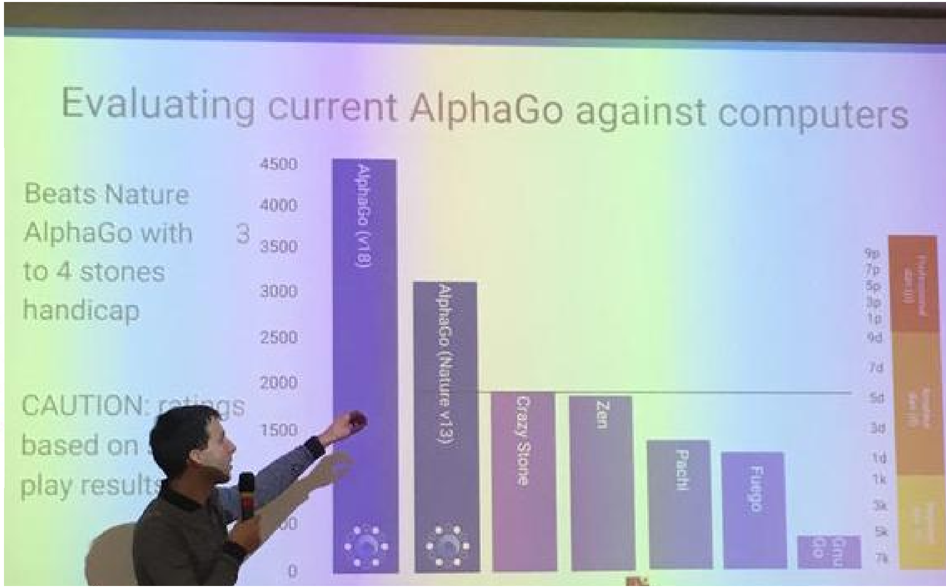

# Bootstrapping by interaction

Why would two androids casually chat one with another?

# The distributed brain at the edge

Distributed RL is reminiscent of the philosophical omega point of knowledge

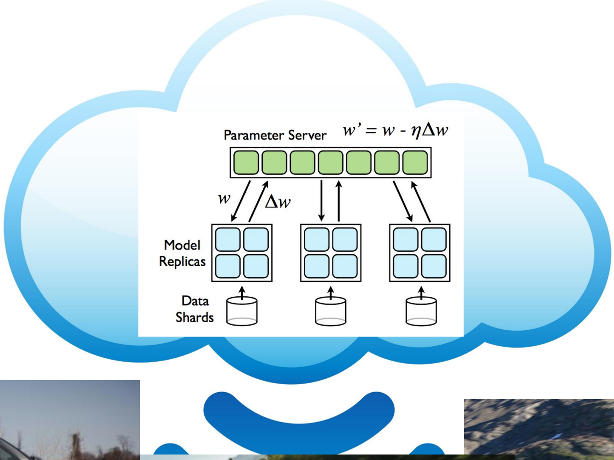


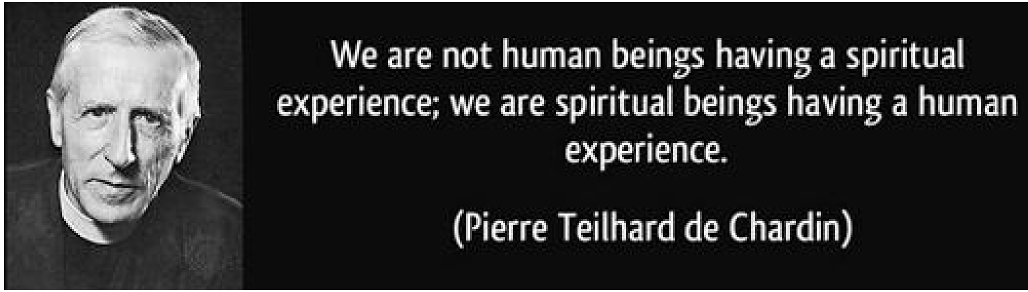


# Multiple Input Neural Networks

# Multi Inputs For Inference

Youtube Video Auto-encoding

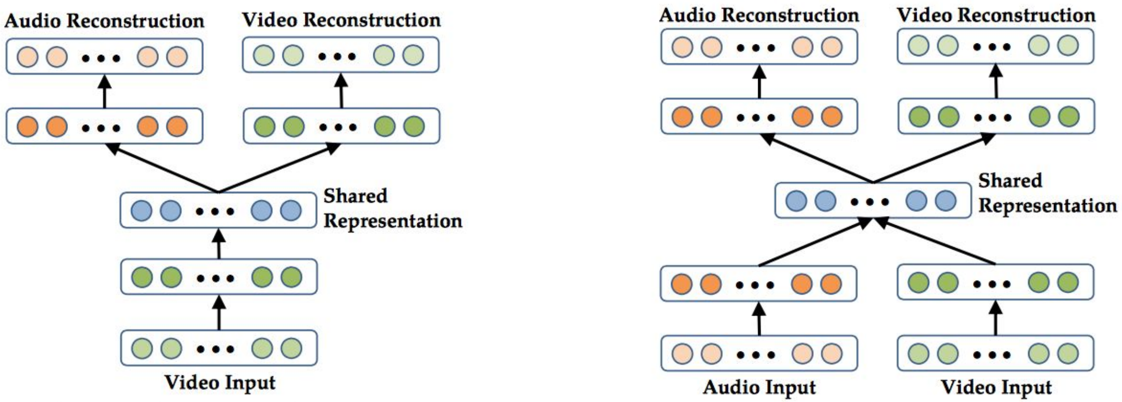


(a) Video-Only Deep Autoencoder

(b) Bimodal Deep Autoencoder

# Multiple Input Control

Multiplexing Inputs

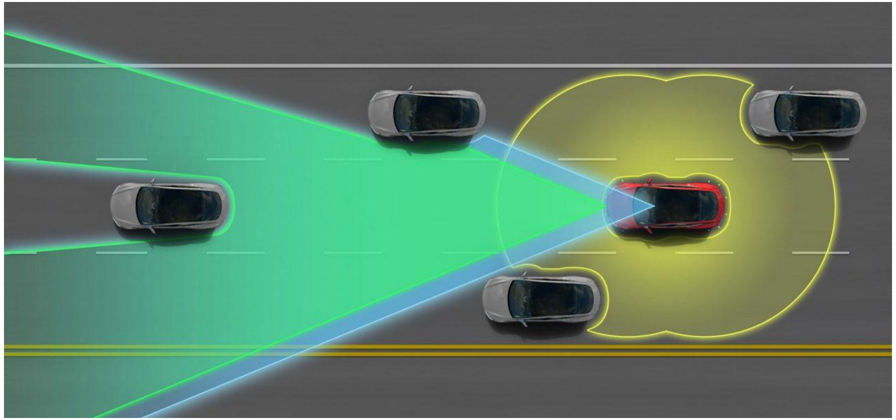

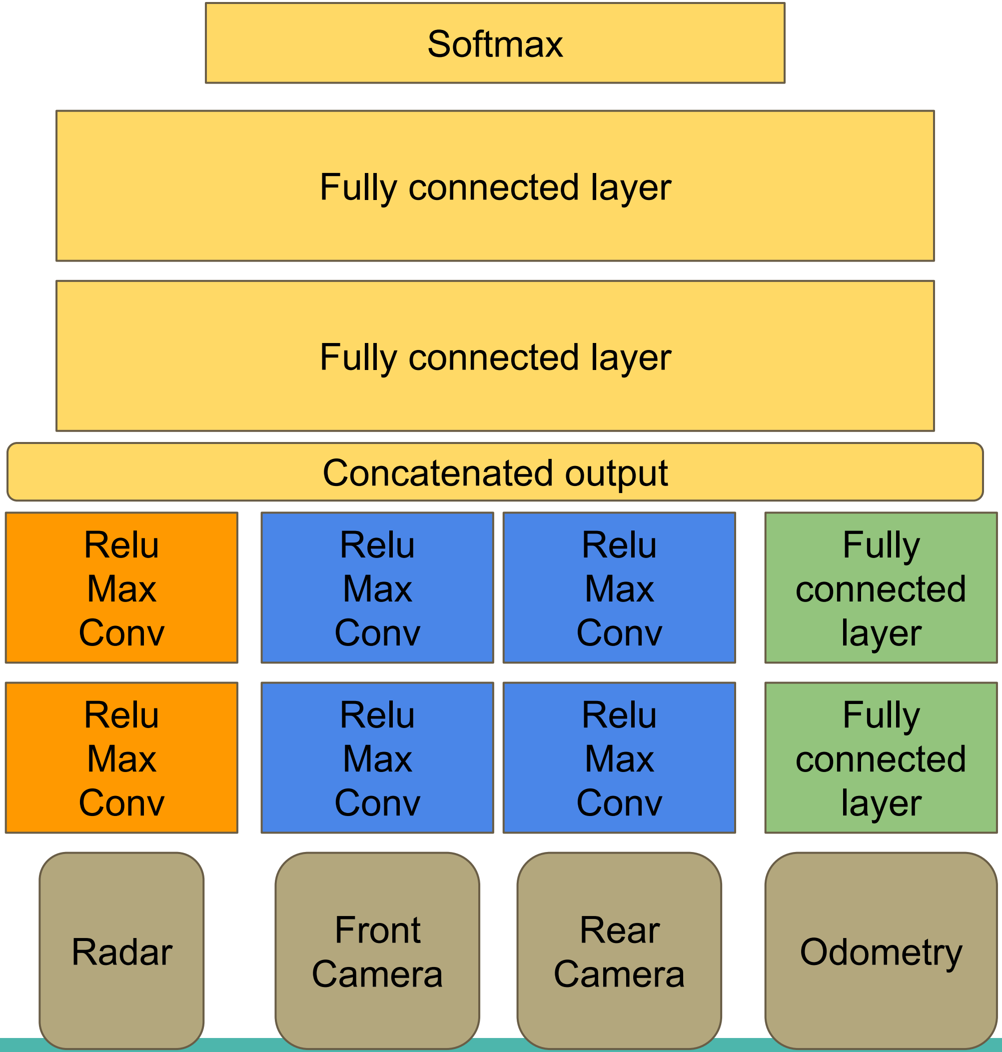

# Multiplexing In The Human Brain