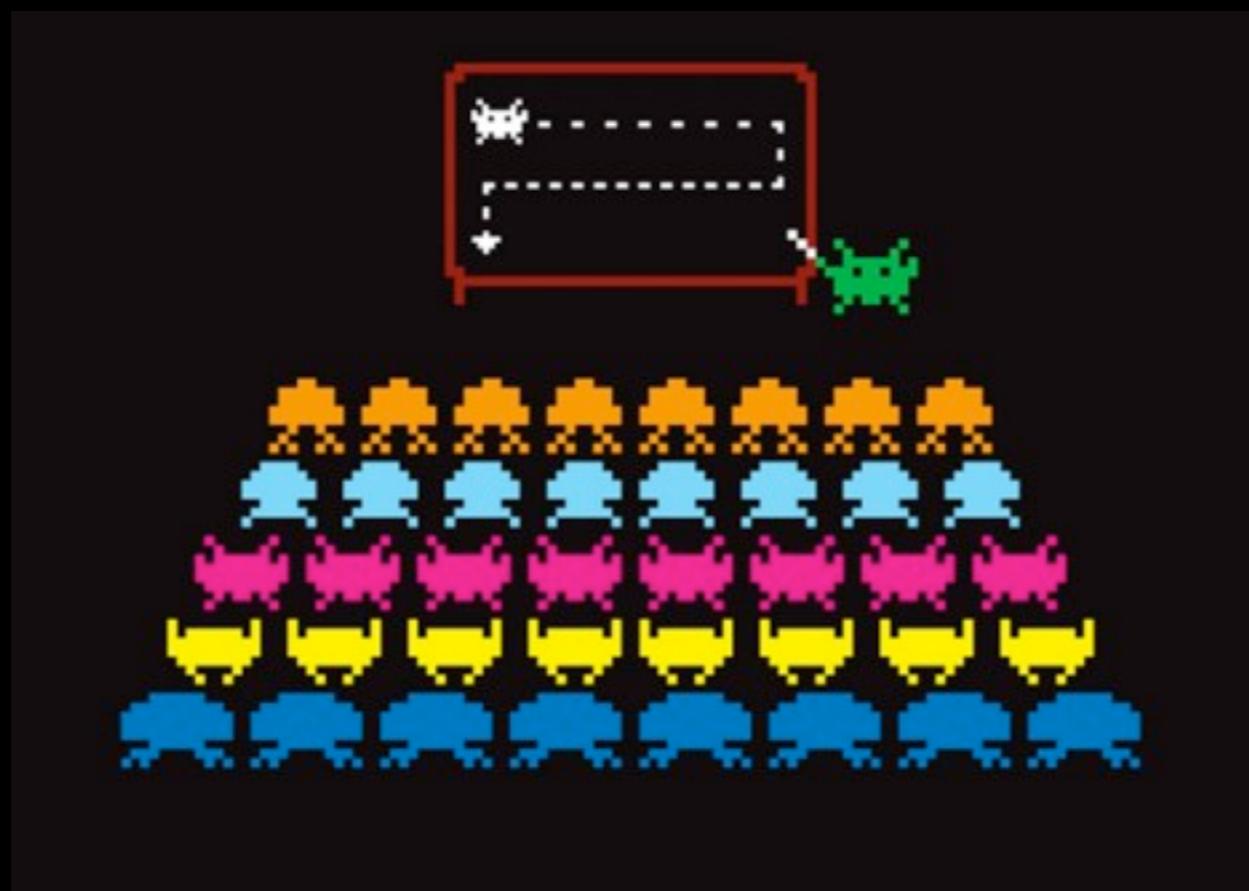


# Deep Learning



Raul Vicente

Institute of Computer Science, University of Tartu

There is a  
revolution in AI &  
machine learning

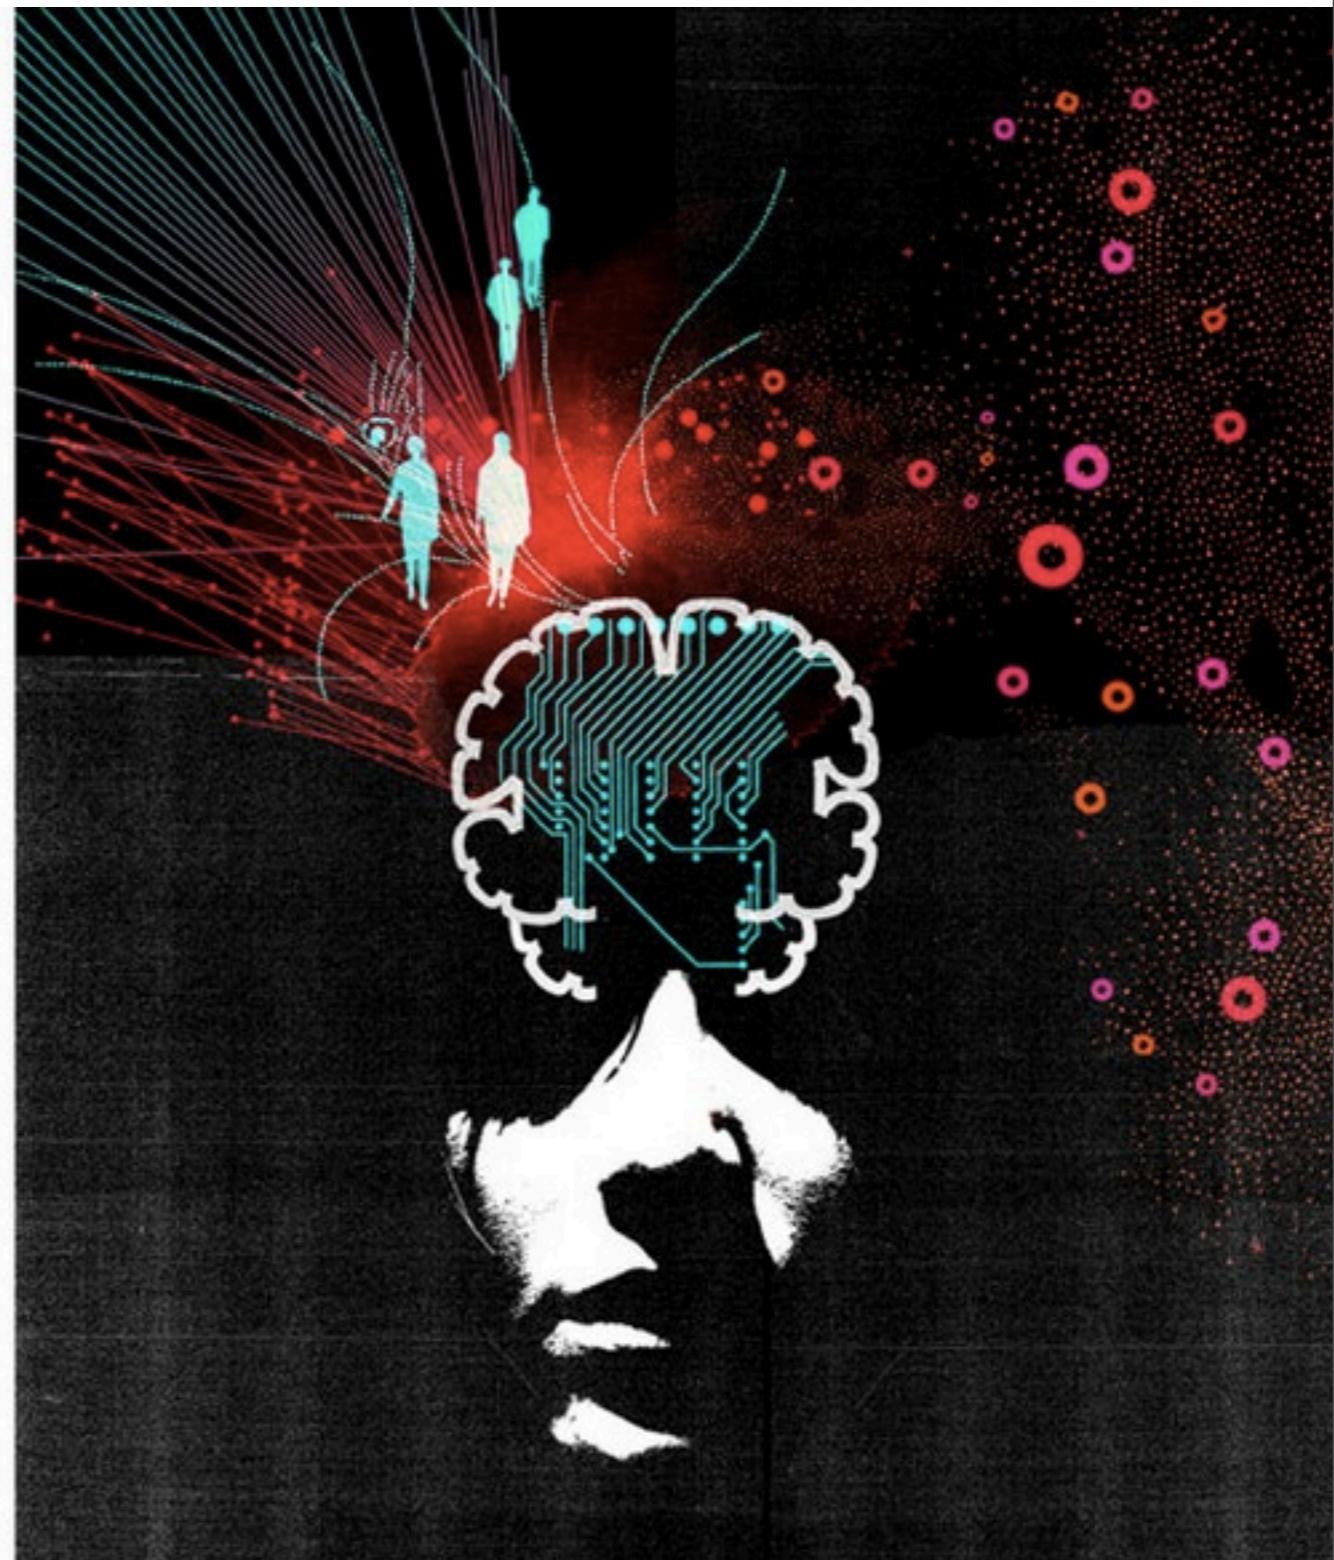
## Deep Learning



# 10 BREAKTHROUGH TECHNOLOGIES 2013

## Deep Learning

**With massive amounts of computational power, machines can now recognize objects and translate speech in real time. Artificial intelligence is finally getting smart.**



# **Is “Deep Learning” A Revolution in Artificial Intelligence?**

New Yorker 11/2012

## **Facebook taps ‘Deep Learning’ Giant for new AI Lab**

Wired 12/2013

## **Why Did Google Pay \$400 Million for DeepMind**

MIT Technology Reviews 1/2014

## **New Techniques from Google Are Taking Artificial Intelligence to Another Level**

MIT Technology Reviews 5/2013

## **Google Hires Brains that Helped Supercharge Machine Learning**

Wired 3/2013

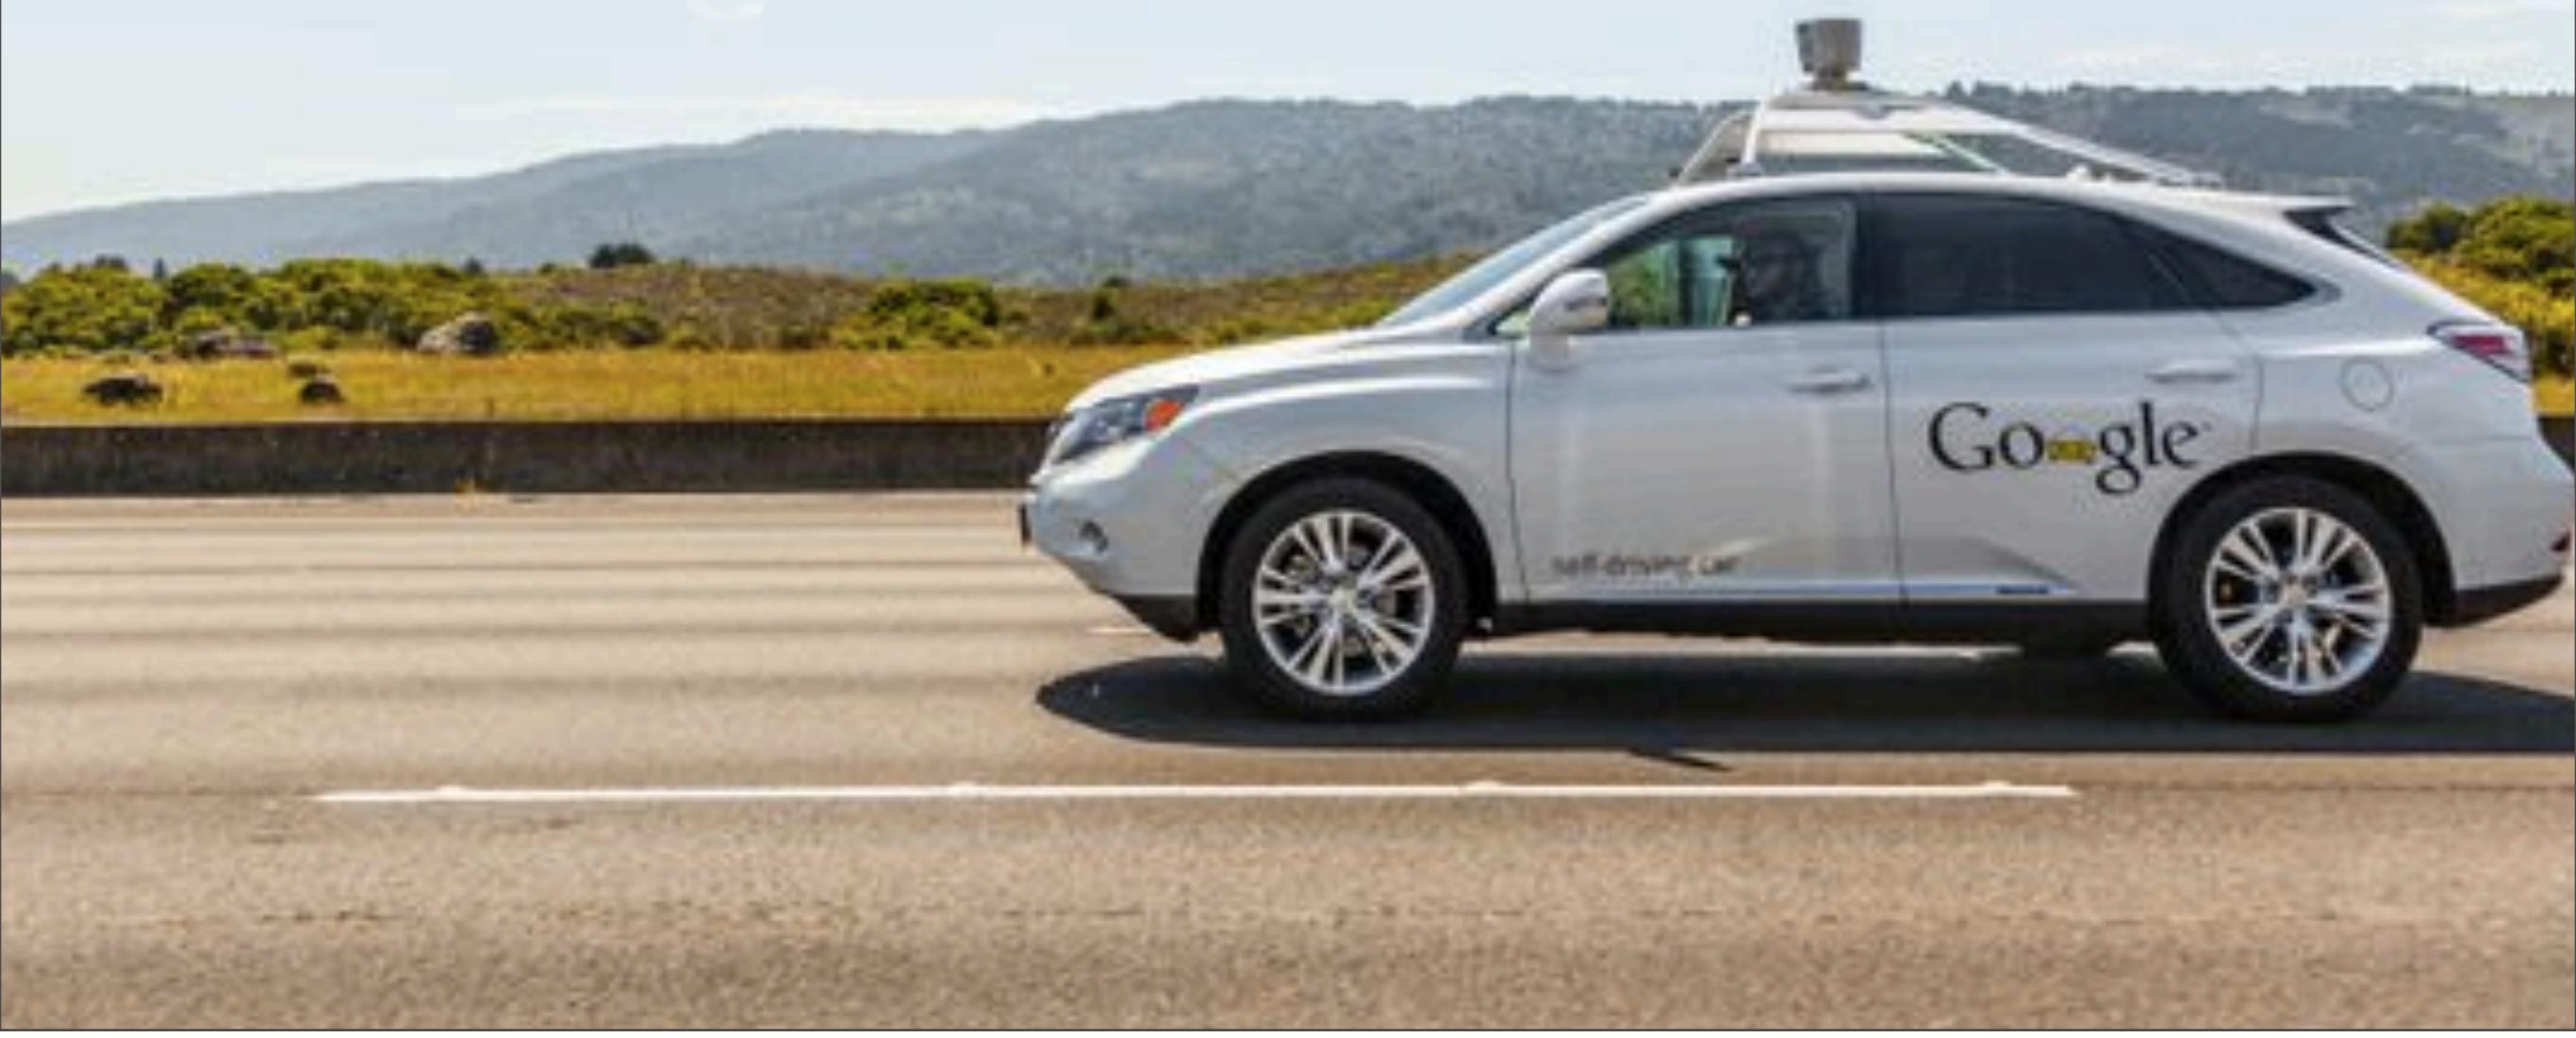


# Real-time speech translation

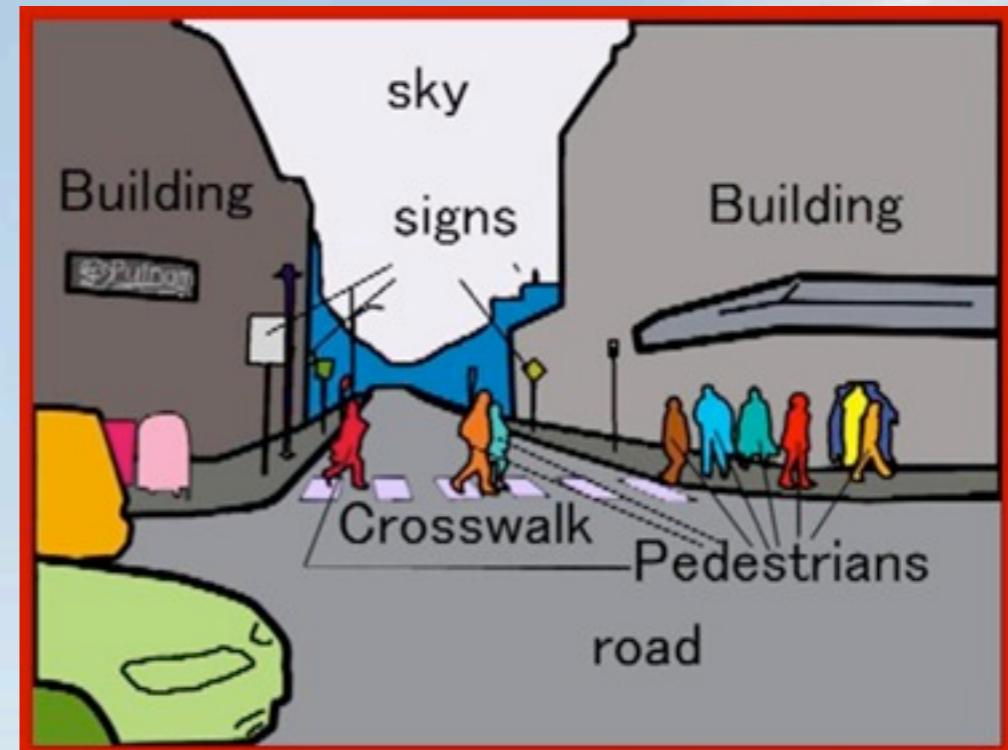
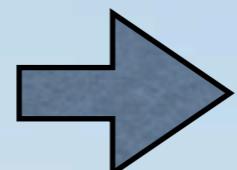


<https://www.youtube.com/watch?v=eu9kMleS0wQ>

# Computer vision





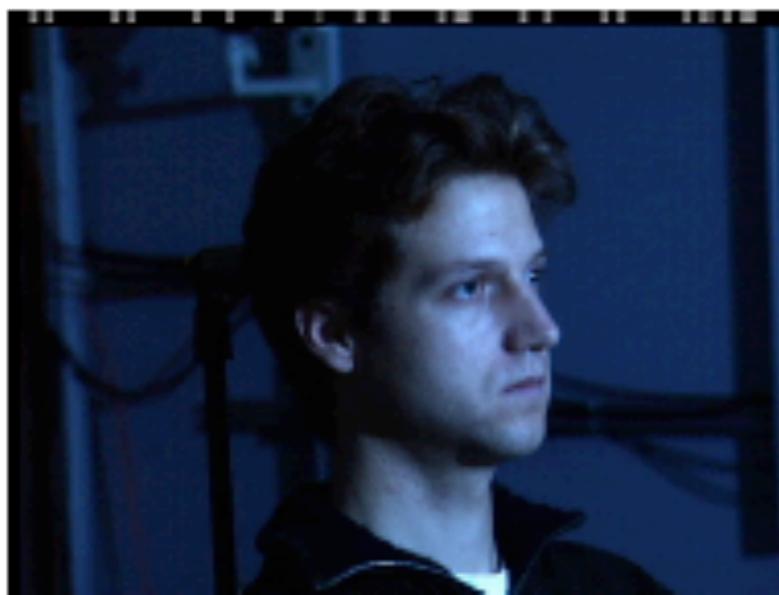


# You see this



**But the camera see this**

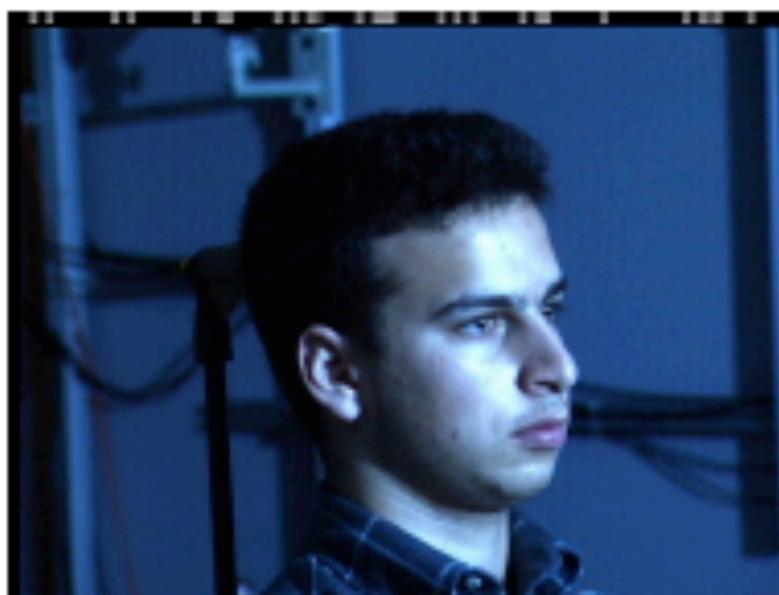
194	210	201	212	199	213	215	195	178	158	182	209
180	189	190	221	209	205	191	167	147	115	129	163
114	126	140	188	176	165	152	140	170	106	78	88
87	103	115	154	143	142	149	153	173	101	57	57
102	112	106	131	122	138	152	147	128	84	58	66
94	95	79	104	105	124	129	113	107	87	69	67
68	71	69	98	89	92	98	95	89	88	76	67
41	56	68	99	63	45	60	82	58	76	75	65
20	43	69	75	56	41	51	73	55	70	63	44
50	50	57	69	75	75	73	74	53	68	59	37
72	59	53	66	84	92	84	74	57	72	63	42
67	61	58	65	75	78	76	73	59	75	69	50



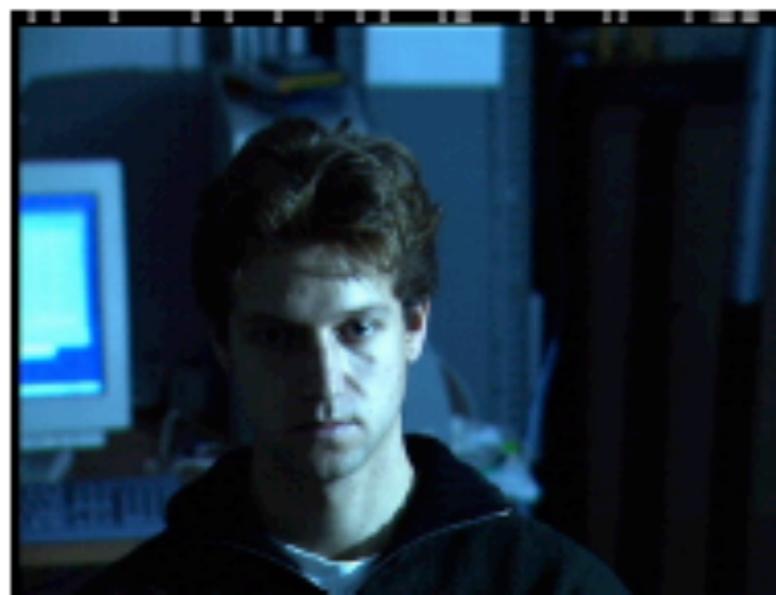
1.04



0.78



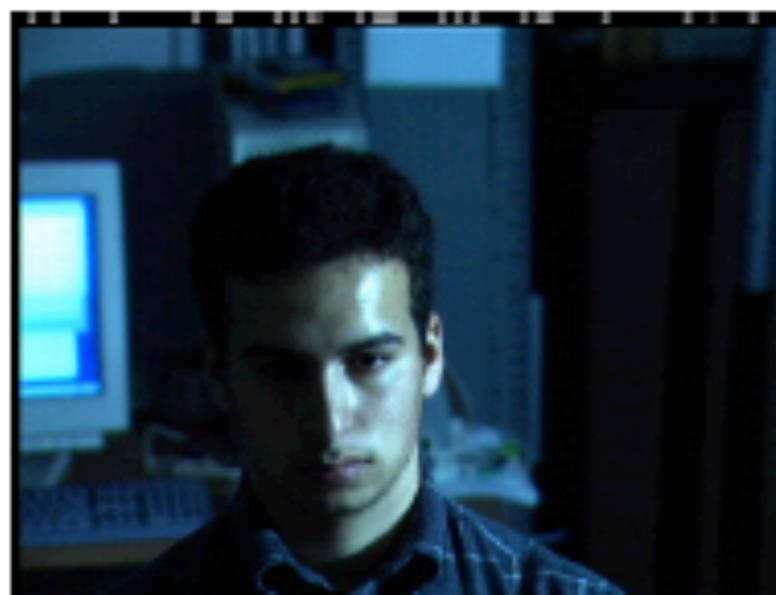
0.99



1.33



1.26





a man riding a horse drawn carriage down a street



a boat is docked in a canal with a large building in the background

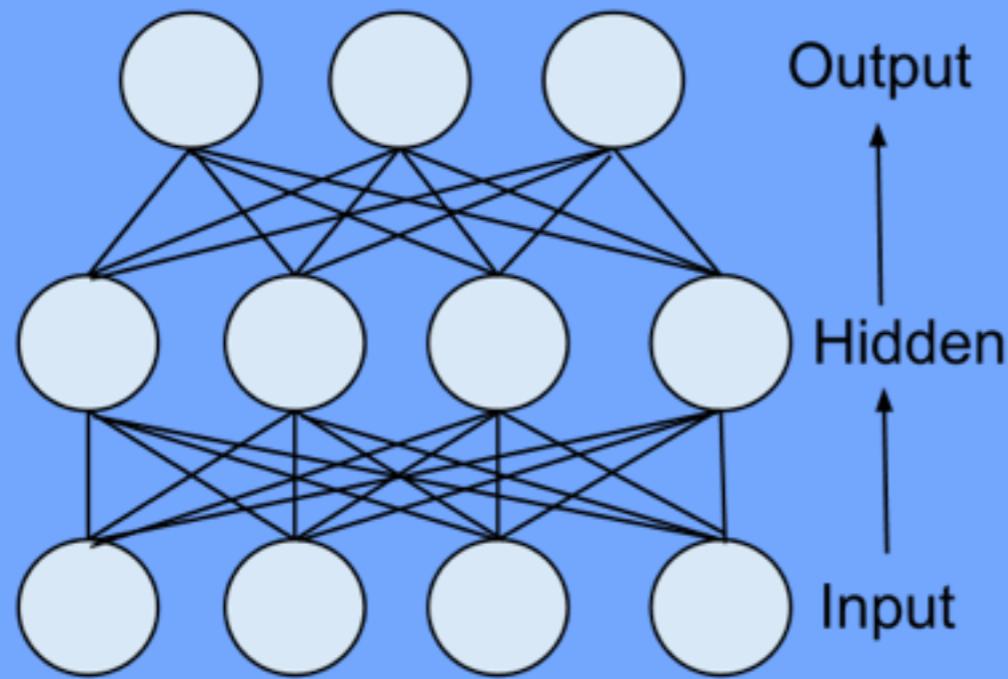


a young boy is holding a baseball bat



# 2 themes

# Deep Learning



# DeepMind



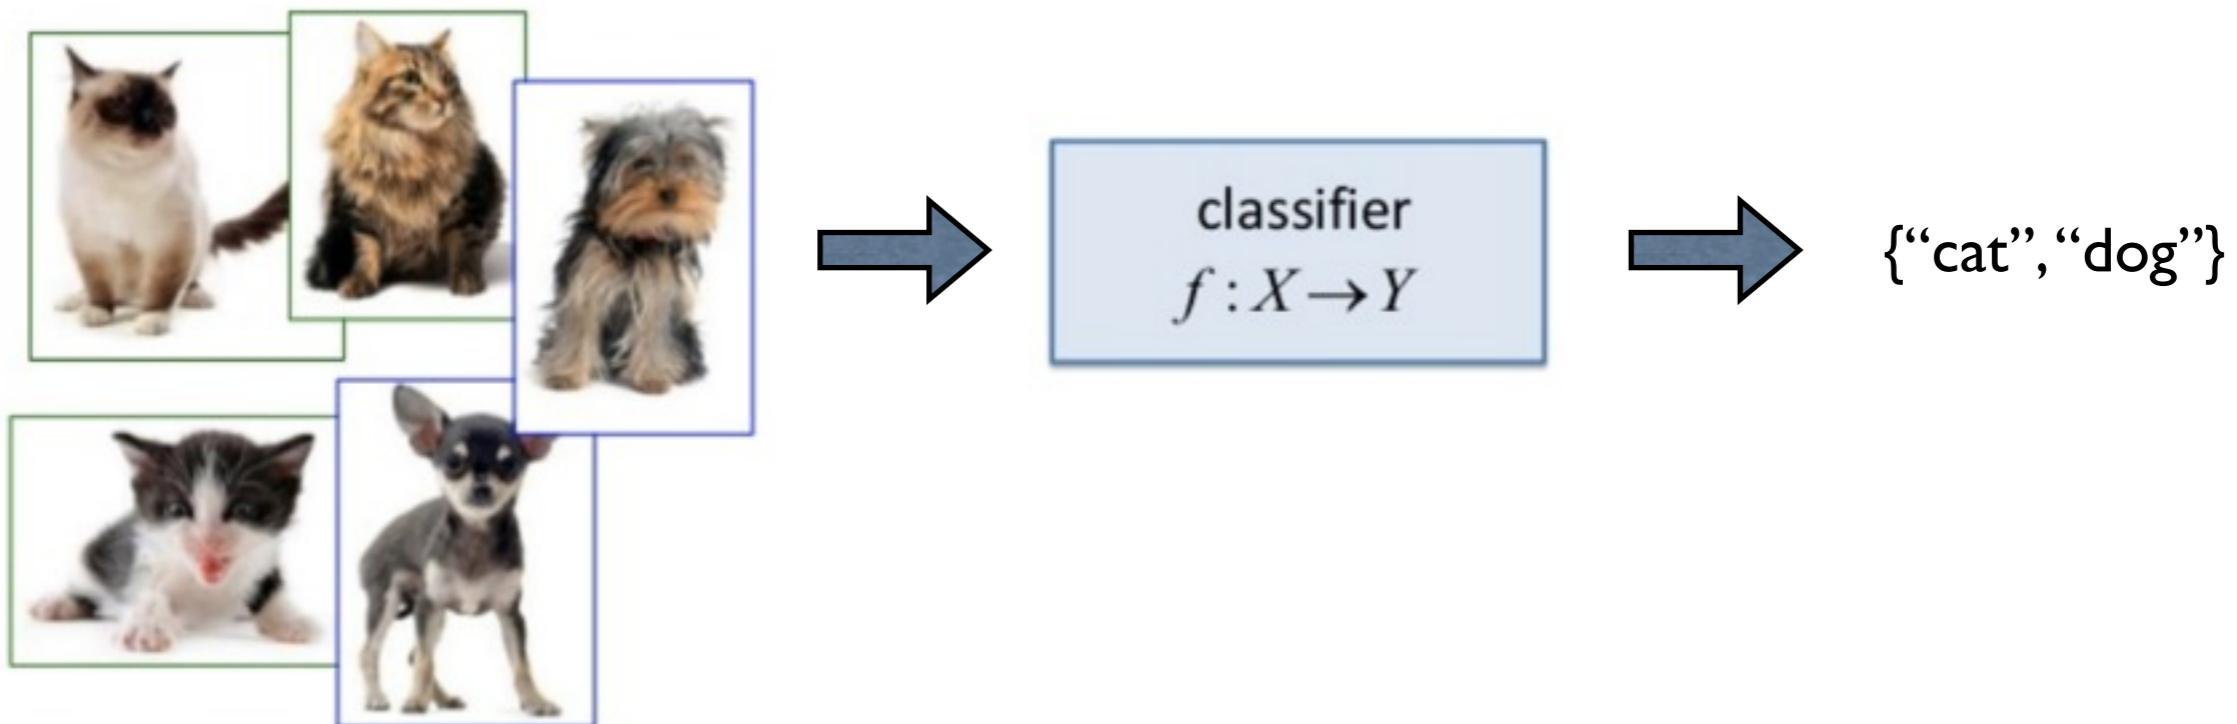
# Machine Learning: algorithms that learn from data

- Supervised



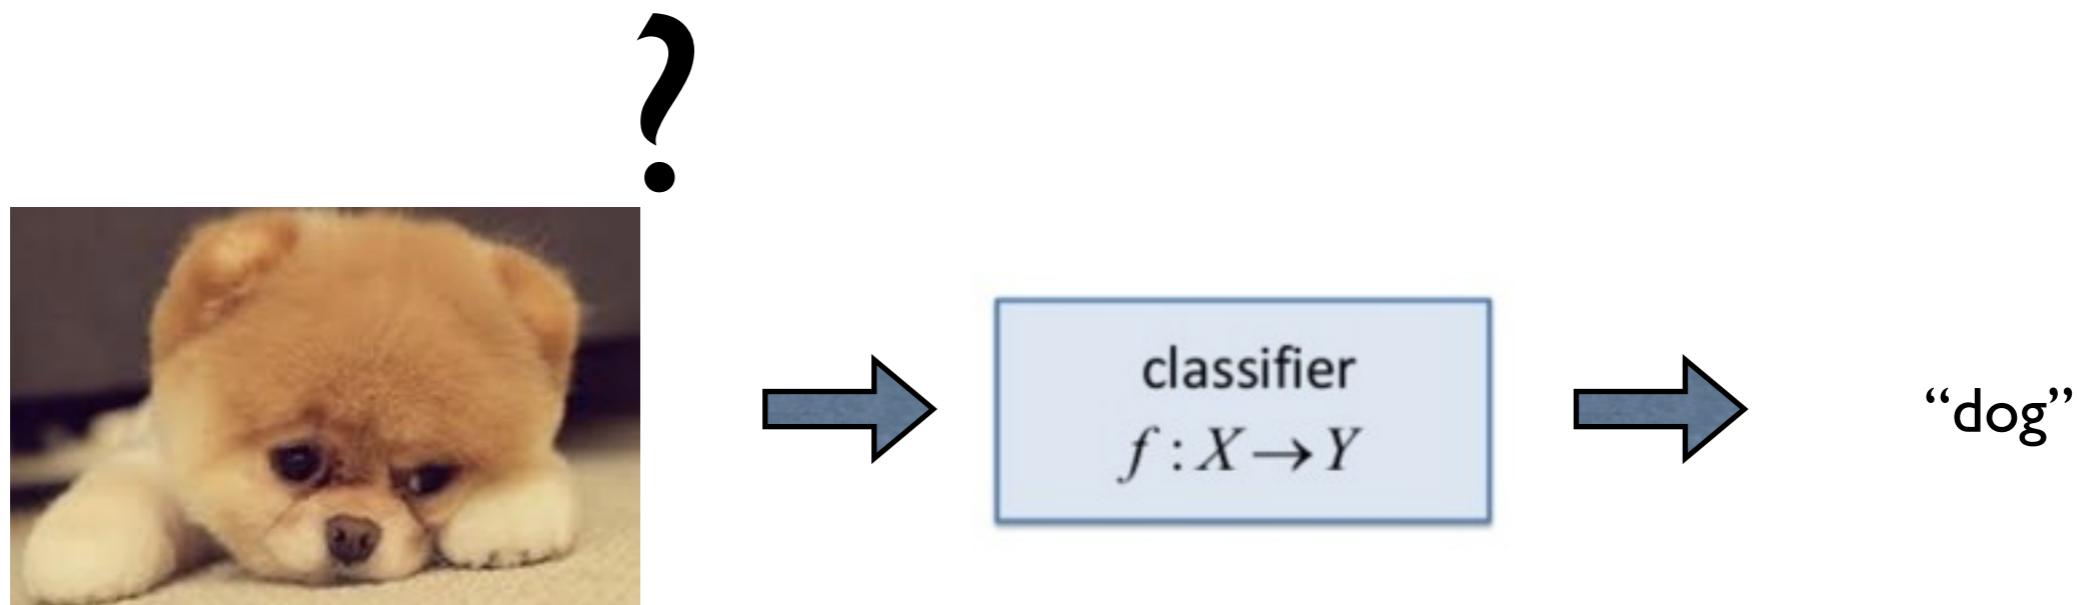
# Machine Learning: algorithms that learn from data

- Supervised



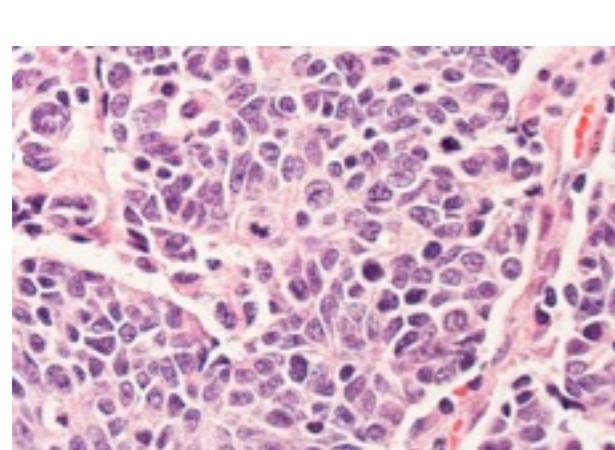
# Machine Learning: algorithms that learn from data

- Supervised

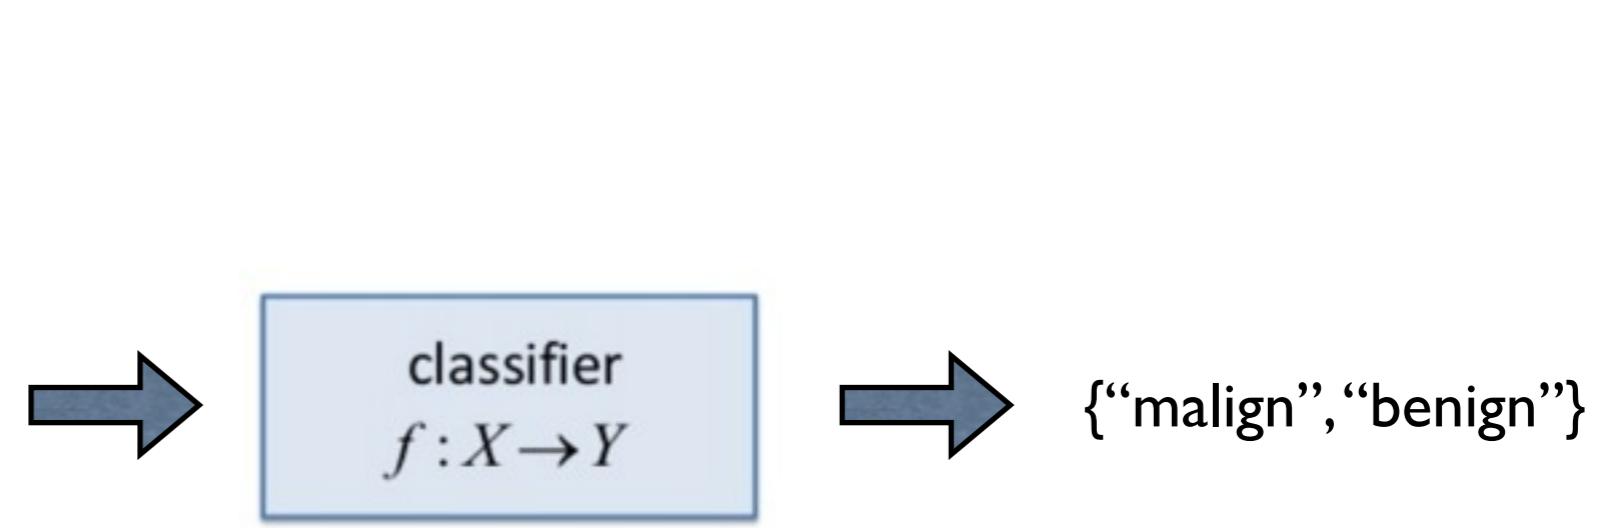


# Machine Learning: algorithms that learn from data

- Supervised

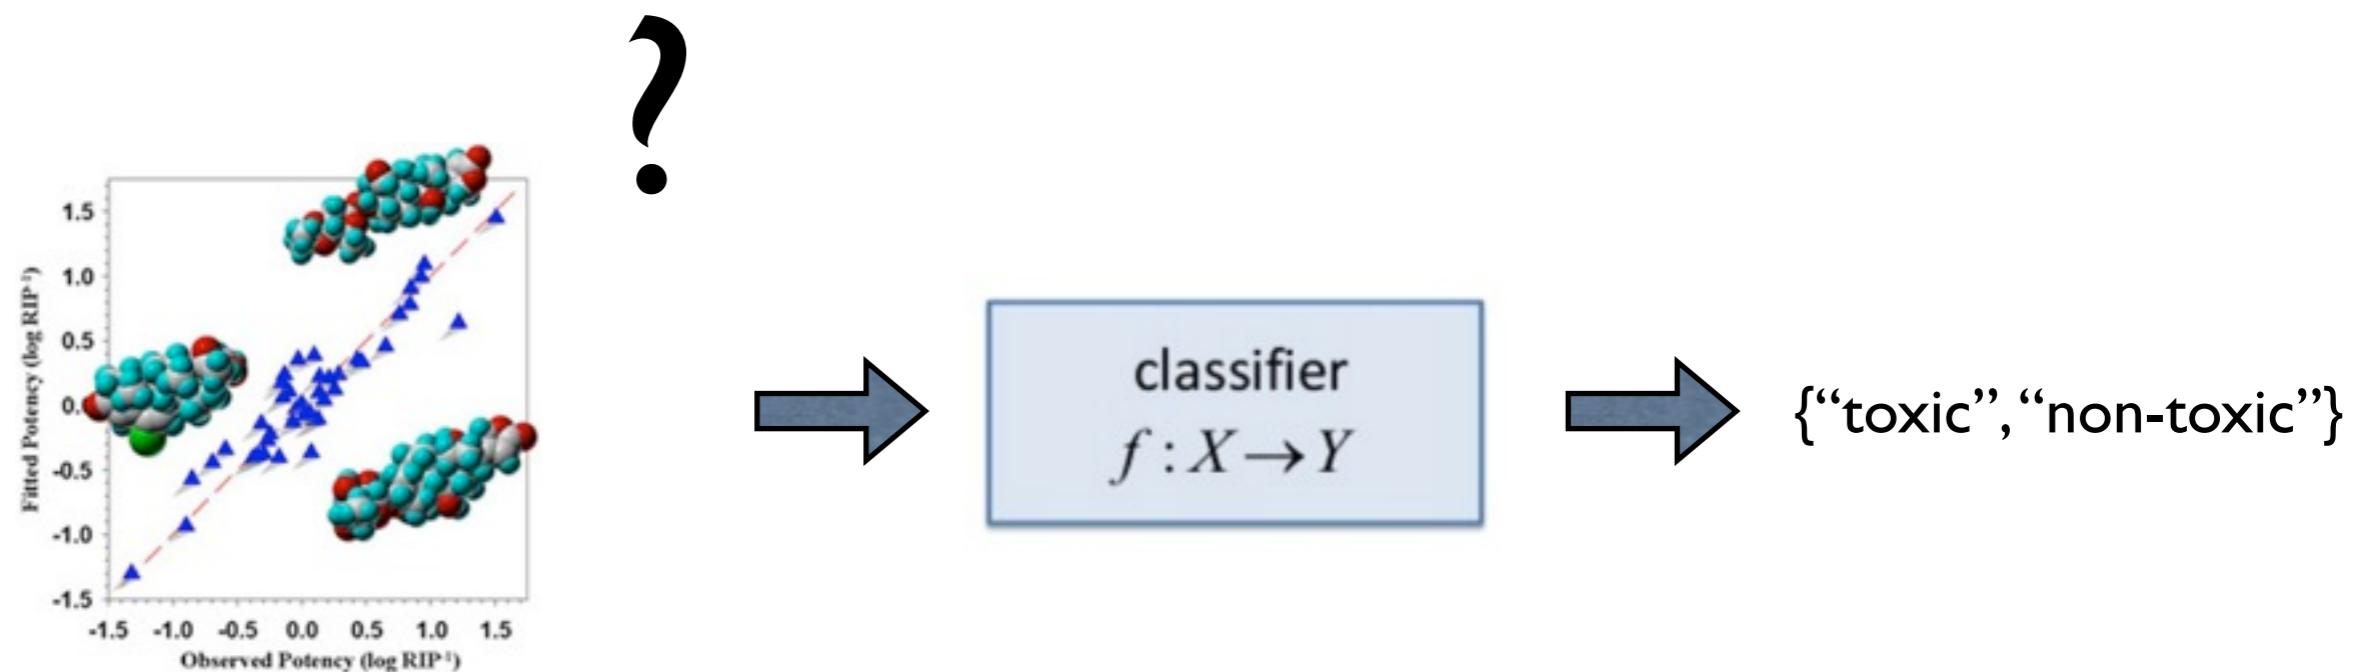


?



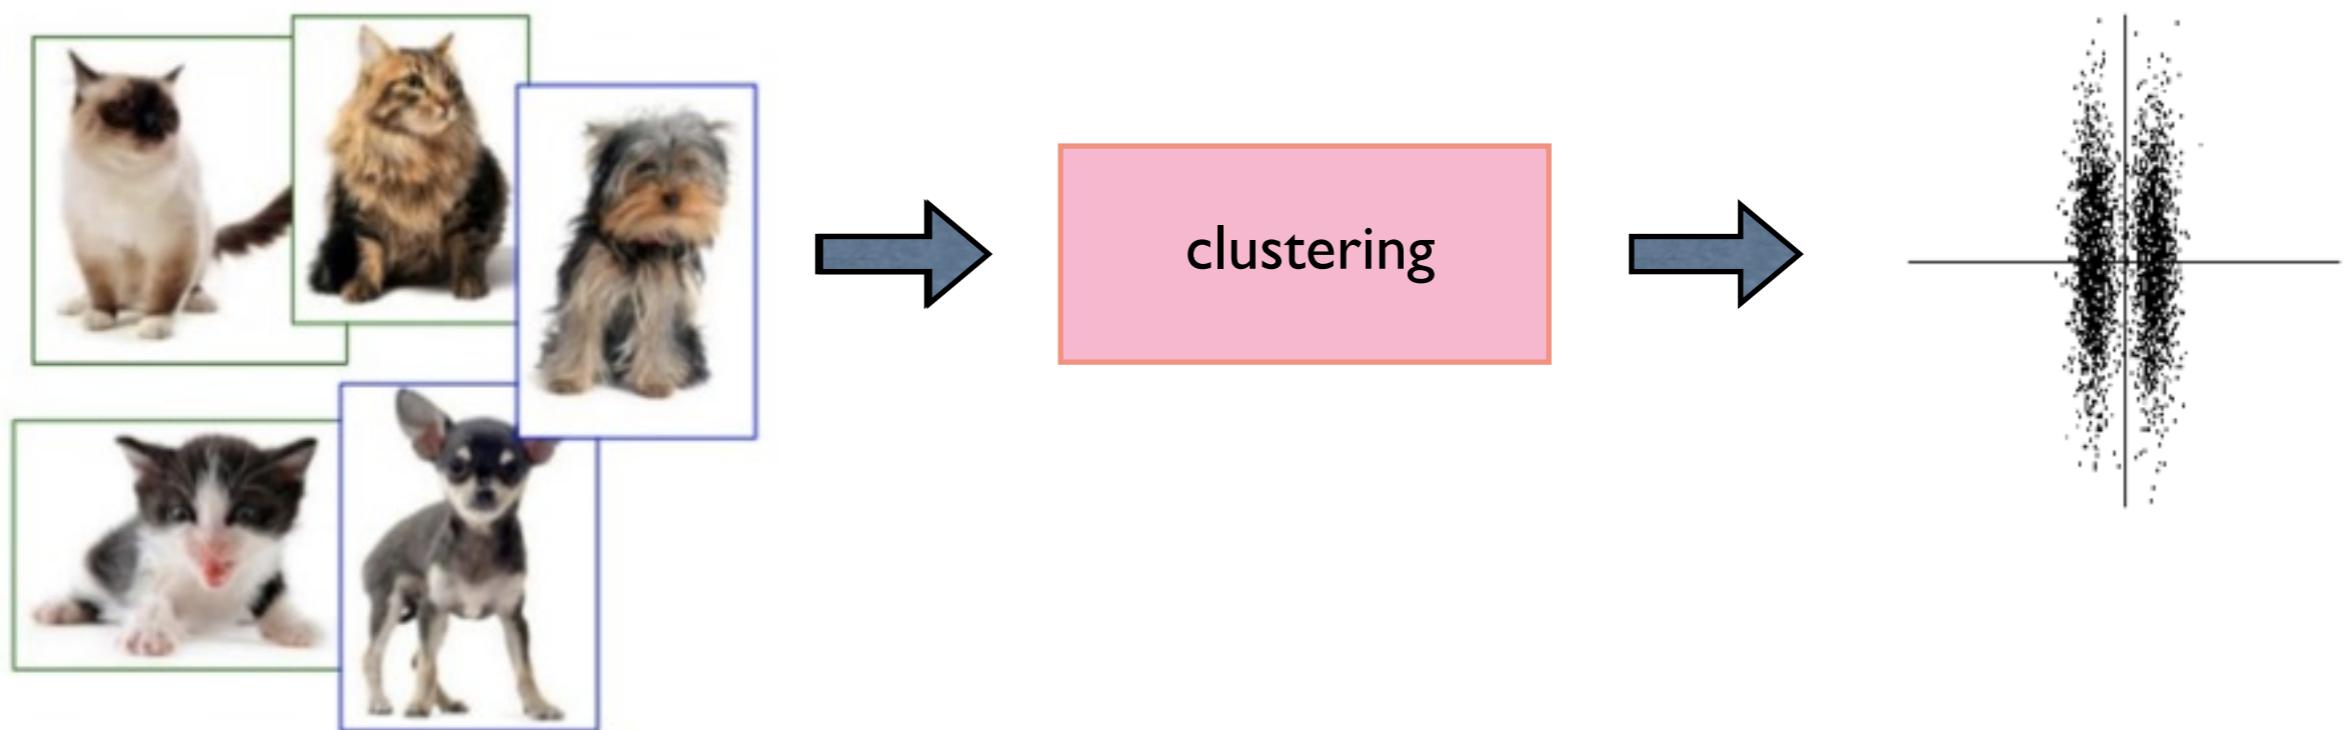
# Machine Learning: algorithms that learn from data

- Supervised



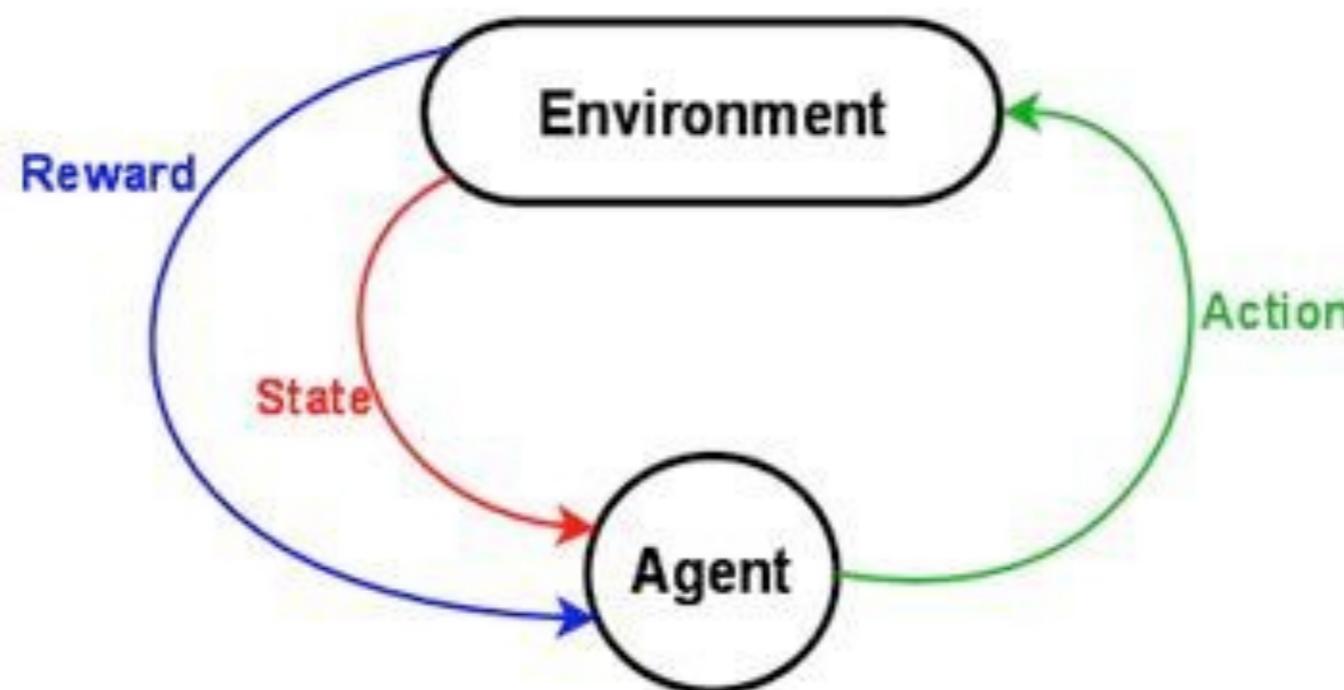
# Machine Learning: algorithms that learn from data

- Supervised
- Unsupervised



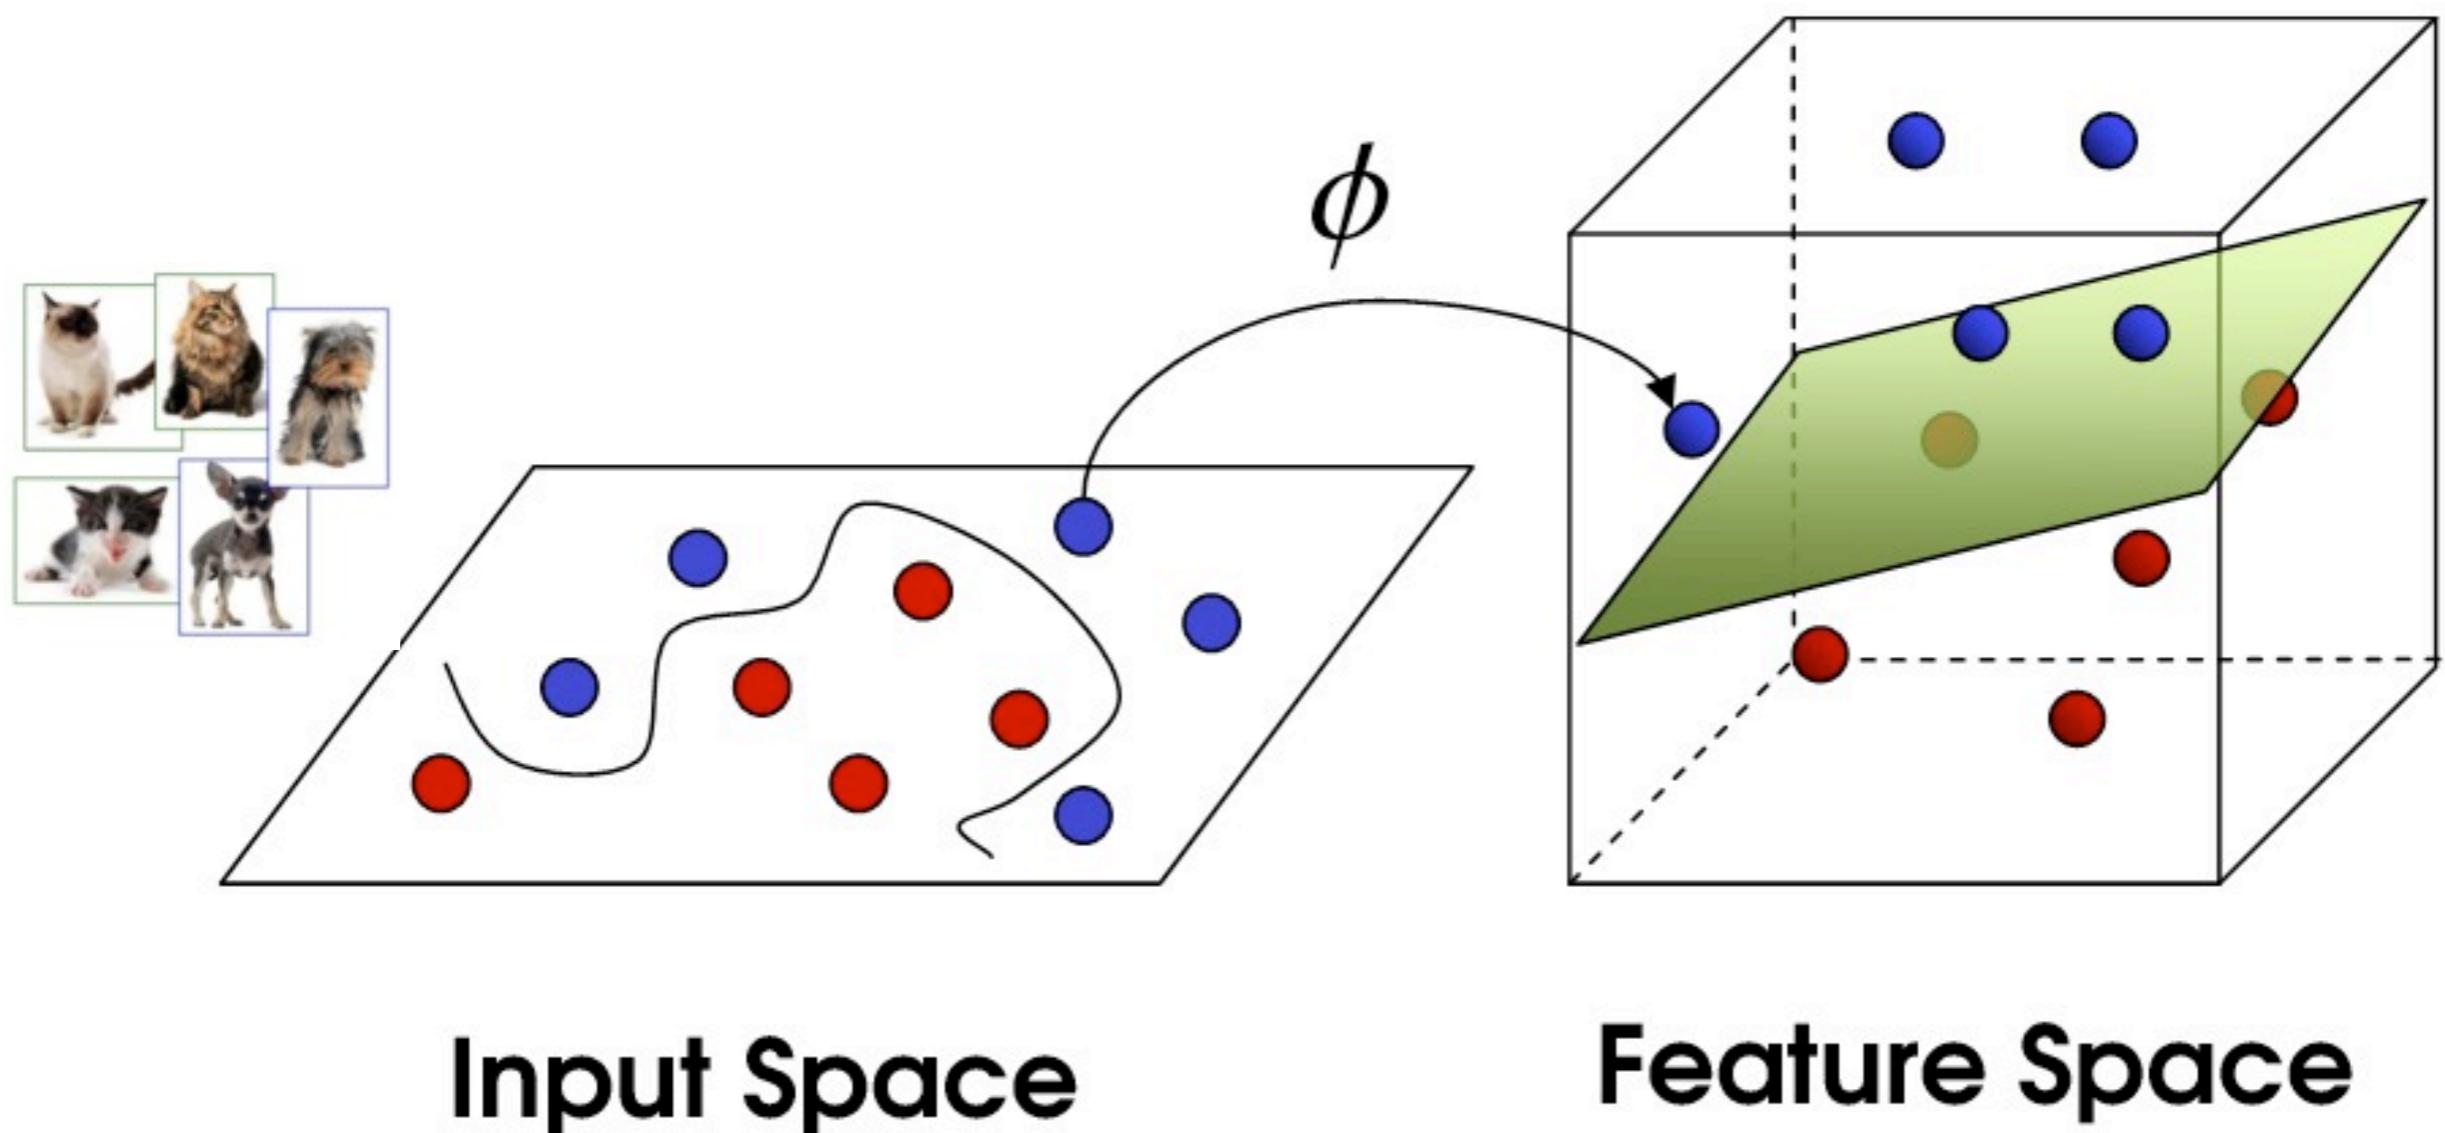
# Machine Learning: algorithms that learn from data

- Supervised
- Unsupervised
- Reinforcement

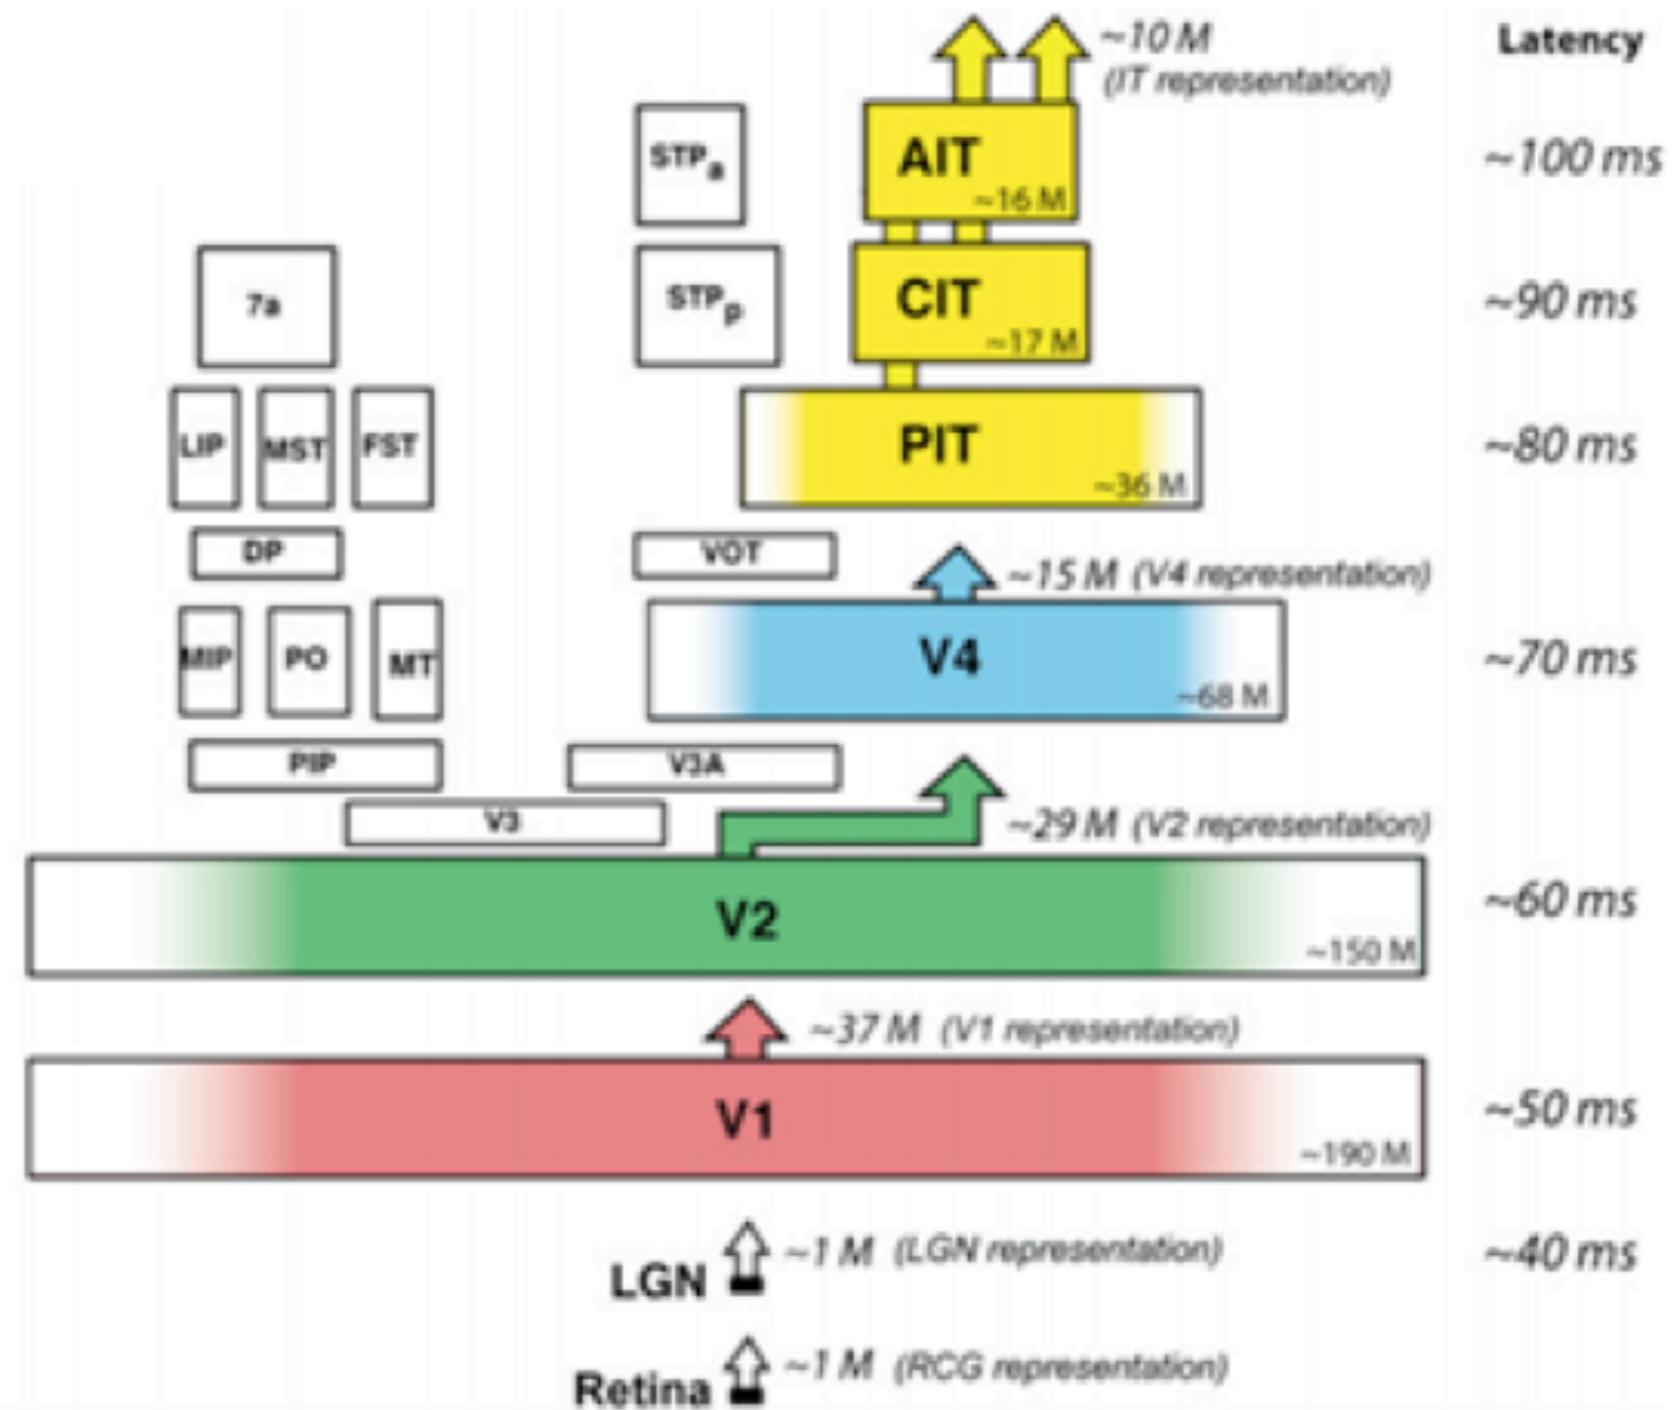
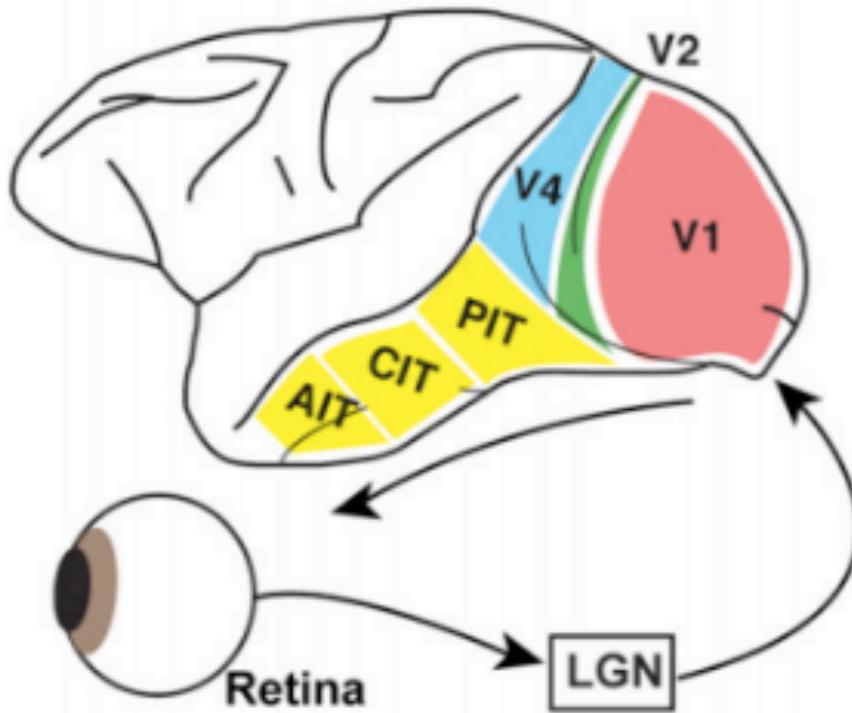


# Shallow Machine Learning

0 or 1 abstraction layer (feature transformation)



# Deep Human Learning

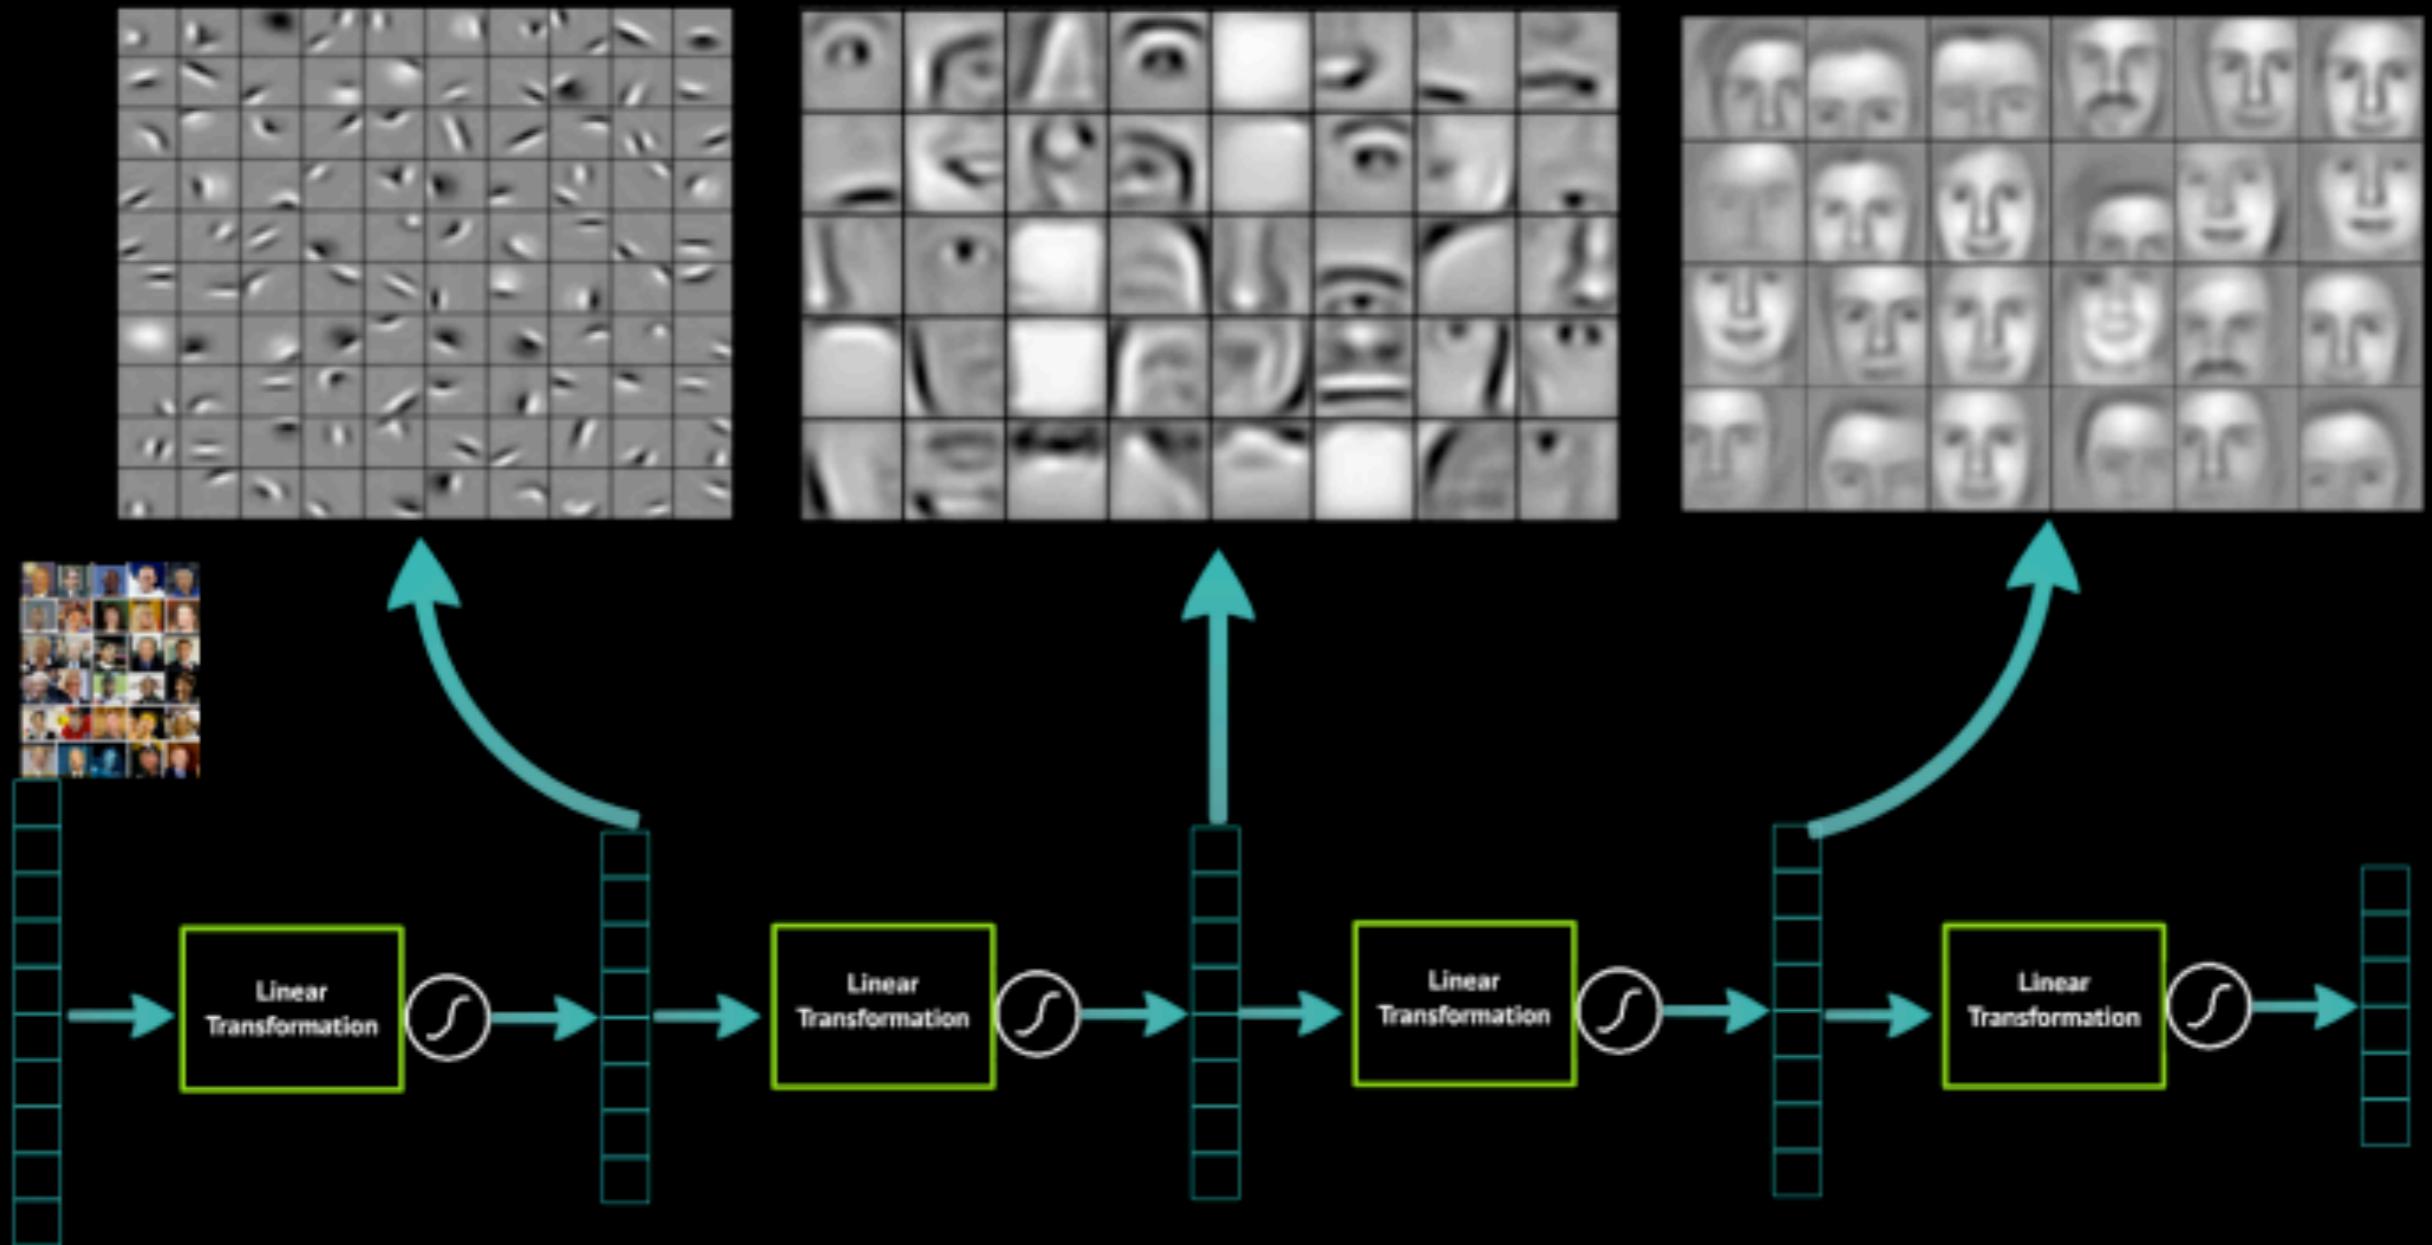


# What is deep learning?

many layers of non-linear processing to model complex relationships among data



# Deep Learning learns layers of features



Faces



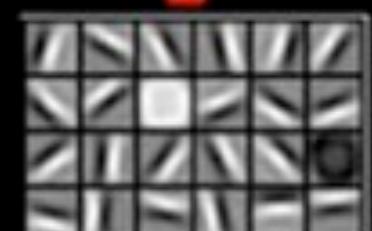
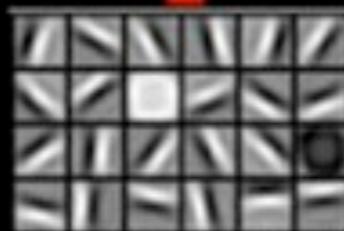
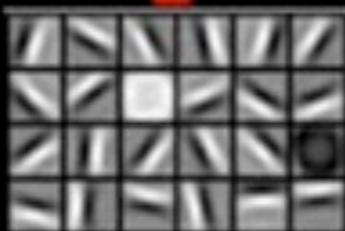
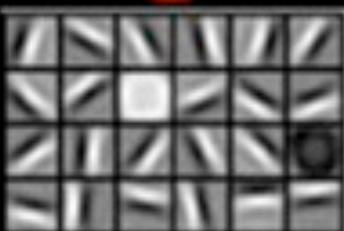
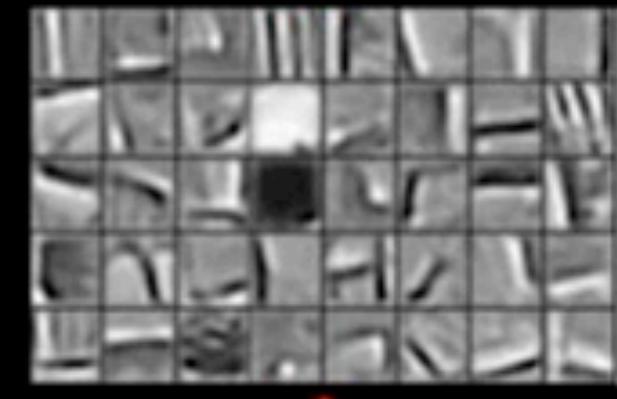
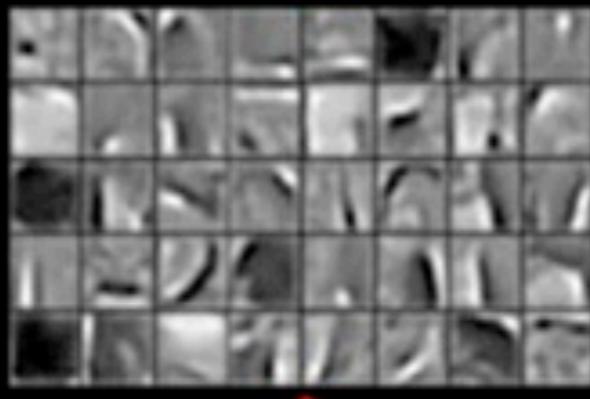
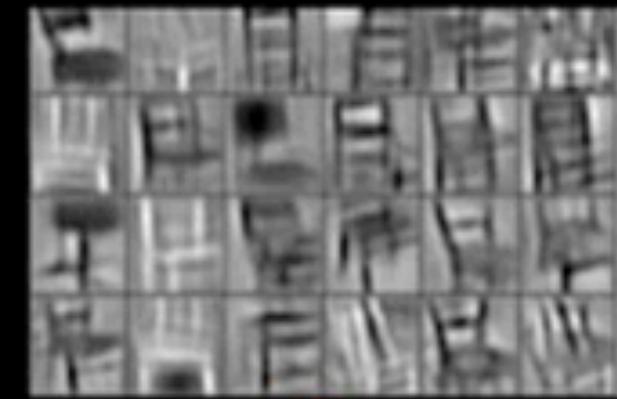
Cars



Elephants

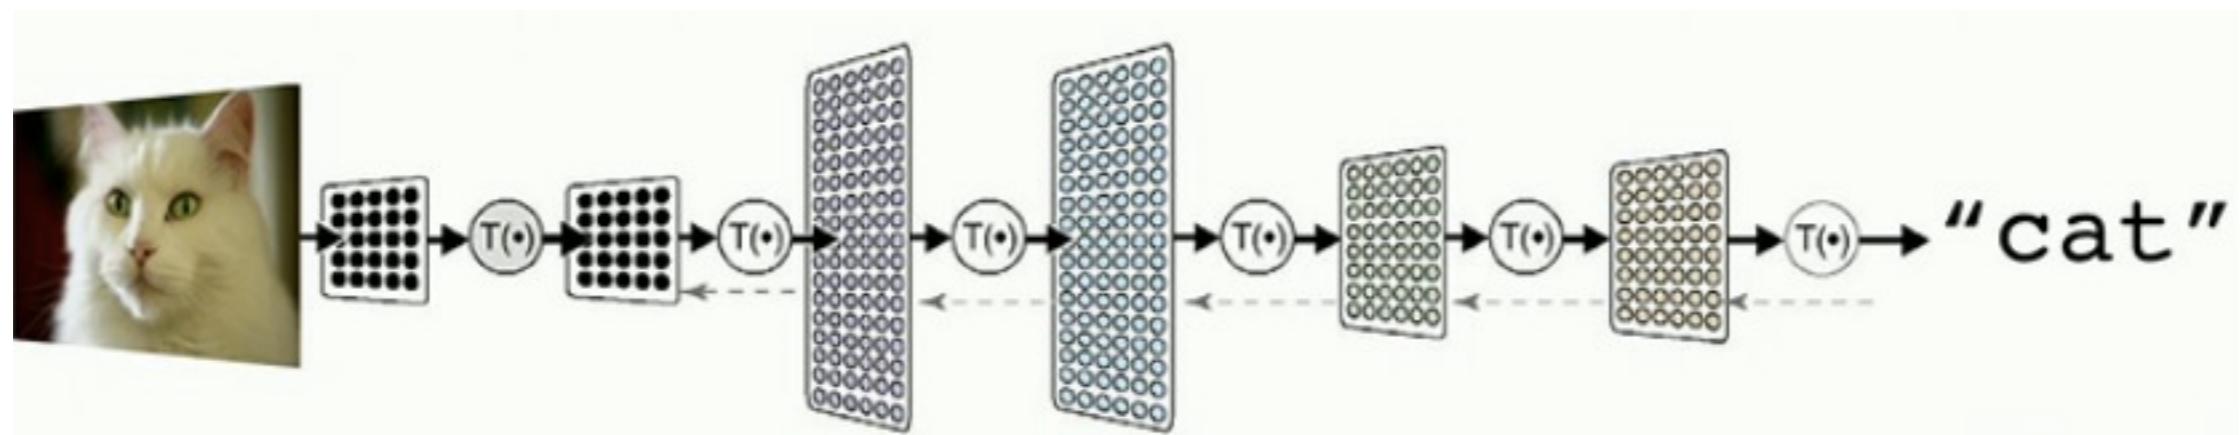


Chairs



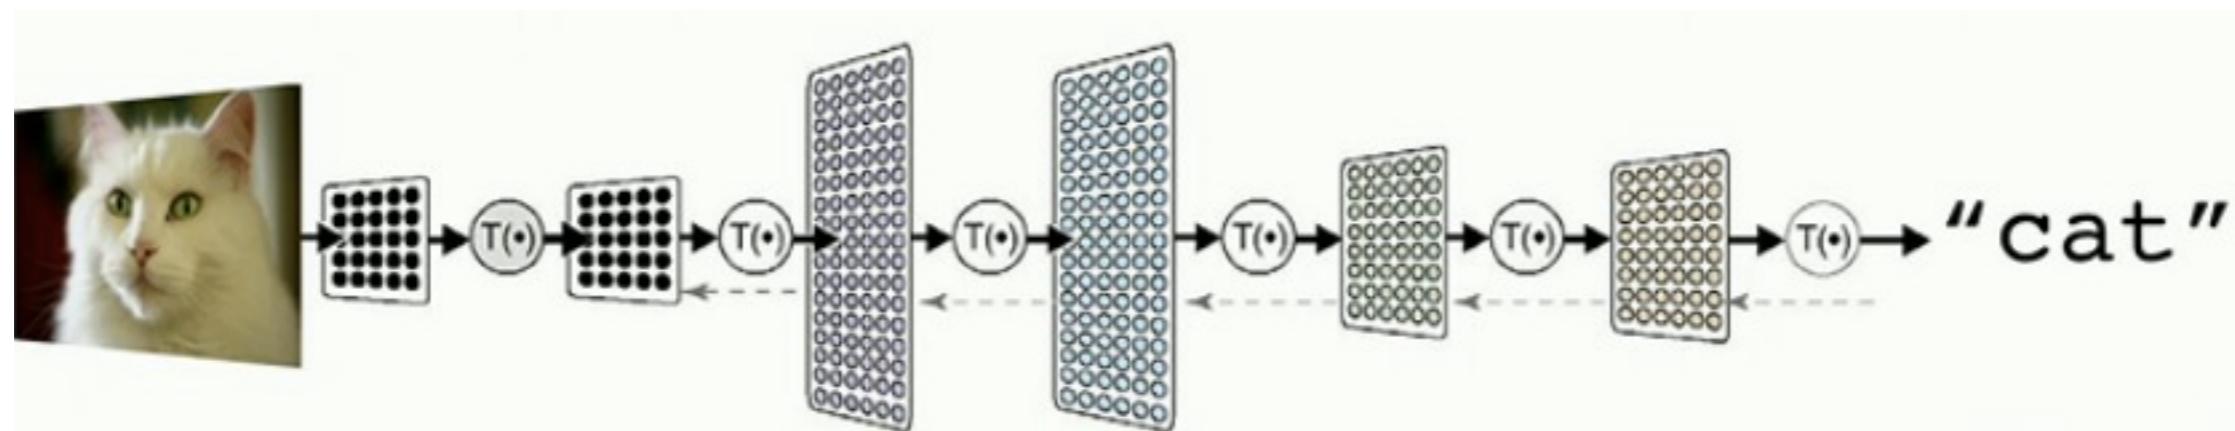
# What is deep learning?

**DL = Artificial Neural Networks** with many layers



# Artificial neural network

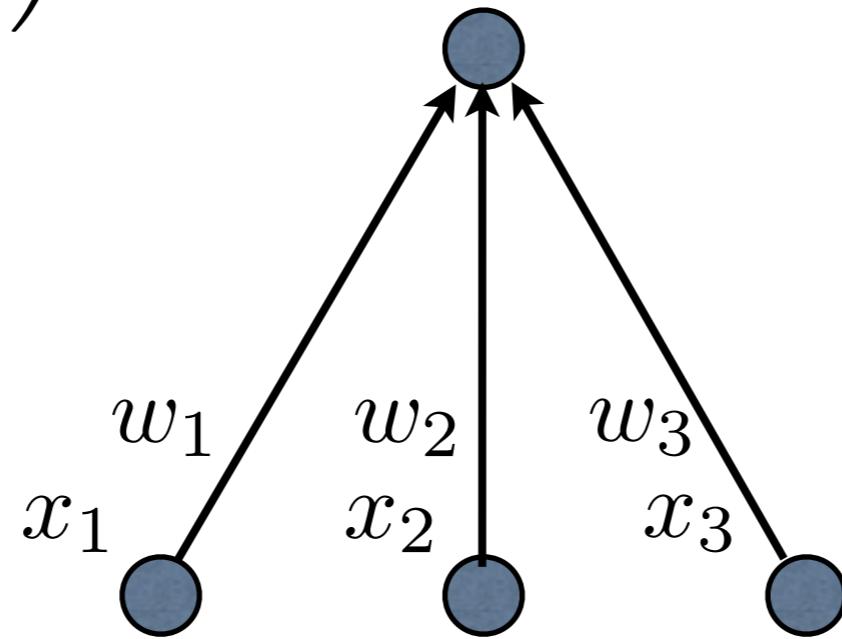
- A collection of simple trainable mathematical units, which collaborate to compute a complicated function
- Compatible with supervised, unsupervised, and reinforcement
- Brain inspired



# The neuron

- Different weights compute different functions

$$y_i = F \left( \sum_i w_i x_i \right)$$

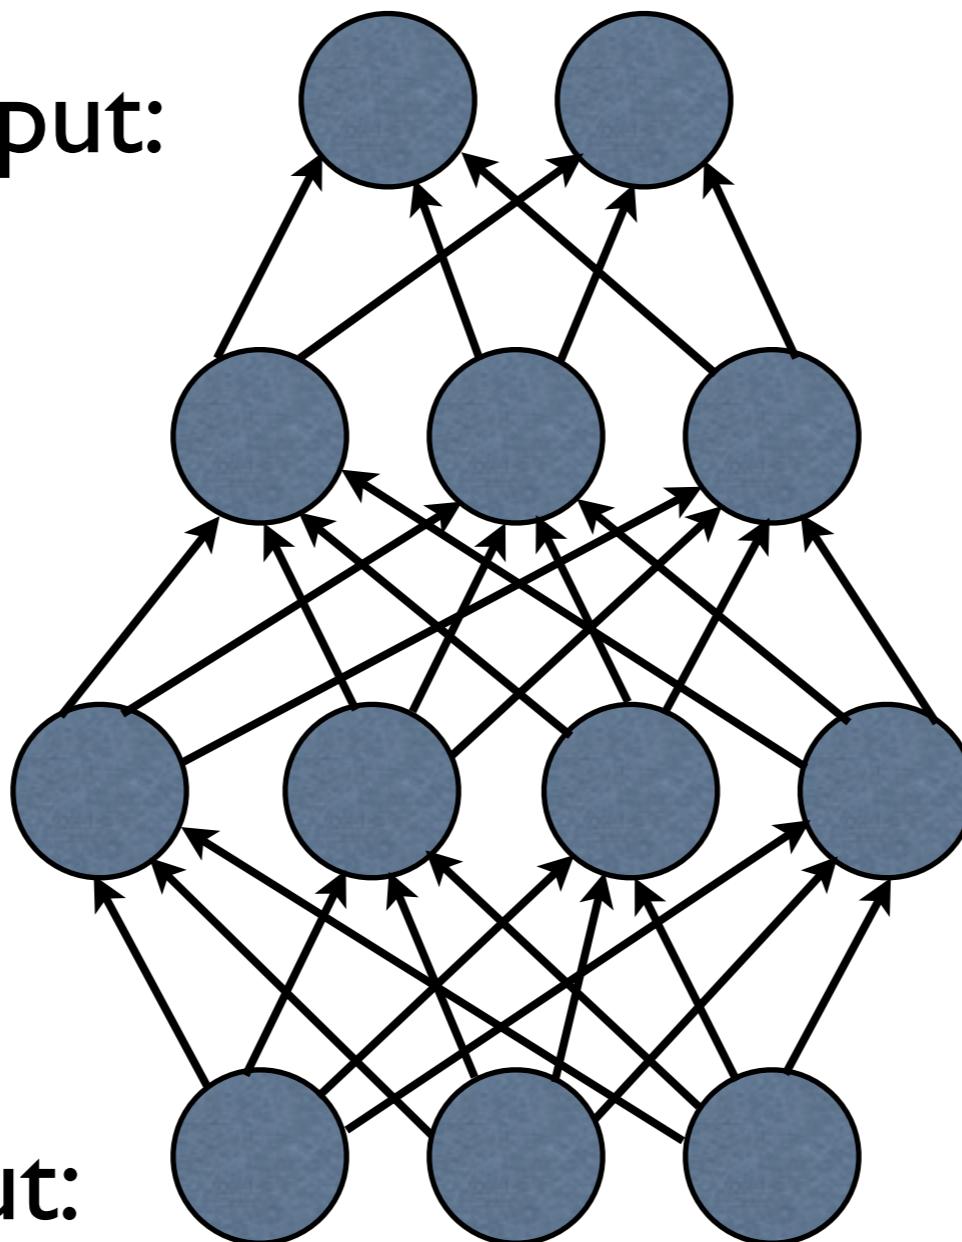


A graph of the sigmoid function, which is an S-shaped curve. The curve passes through the point (0, 0.5). It approaches a value of 0 as  $x$  goes to negative infinity and a value of 1 as  $x$  goes to positive infinity.

$$F(x) = \frac{1}{1 + \exp(-x)}$$

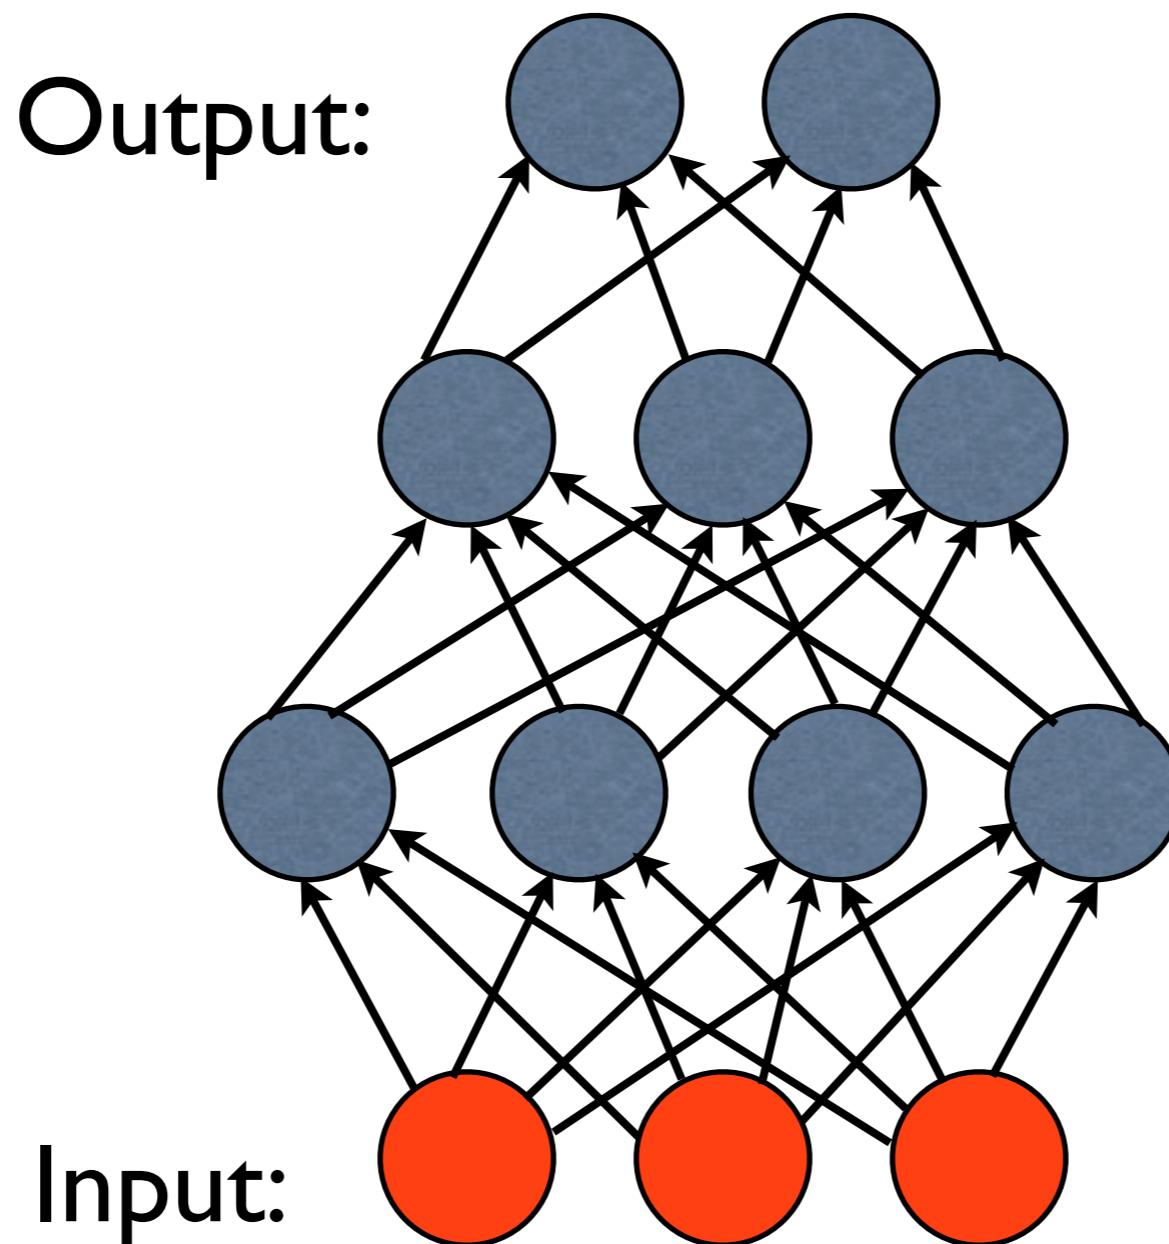
# Neural networks

Output:

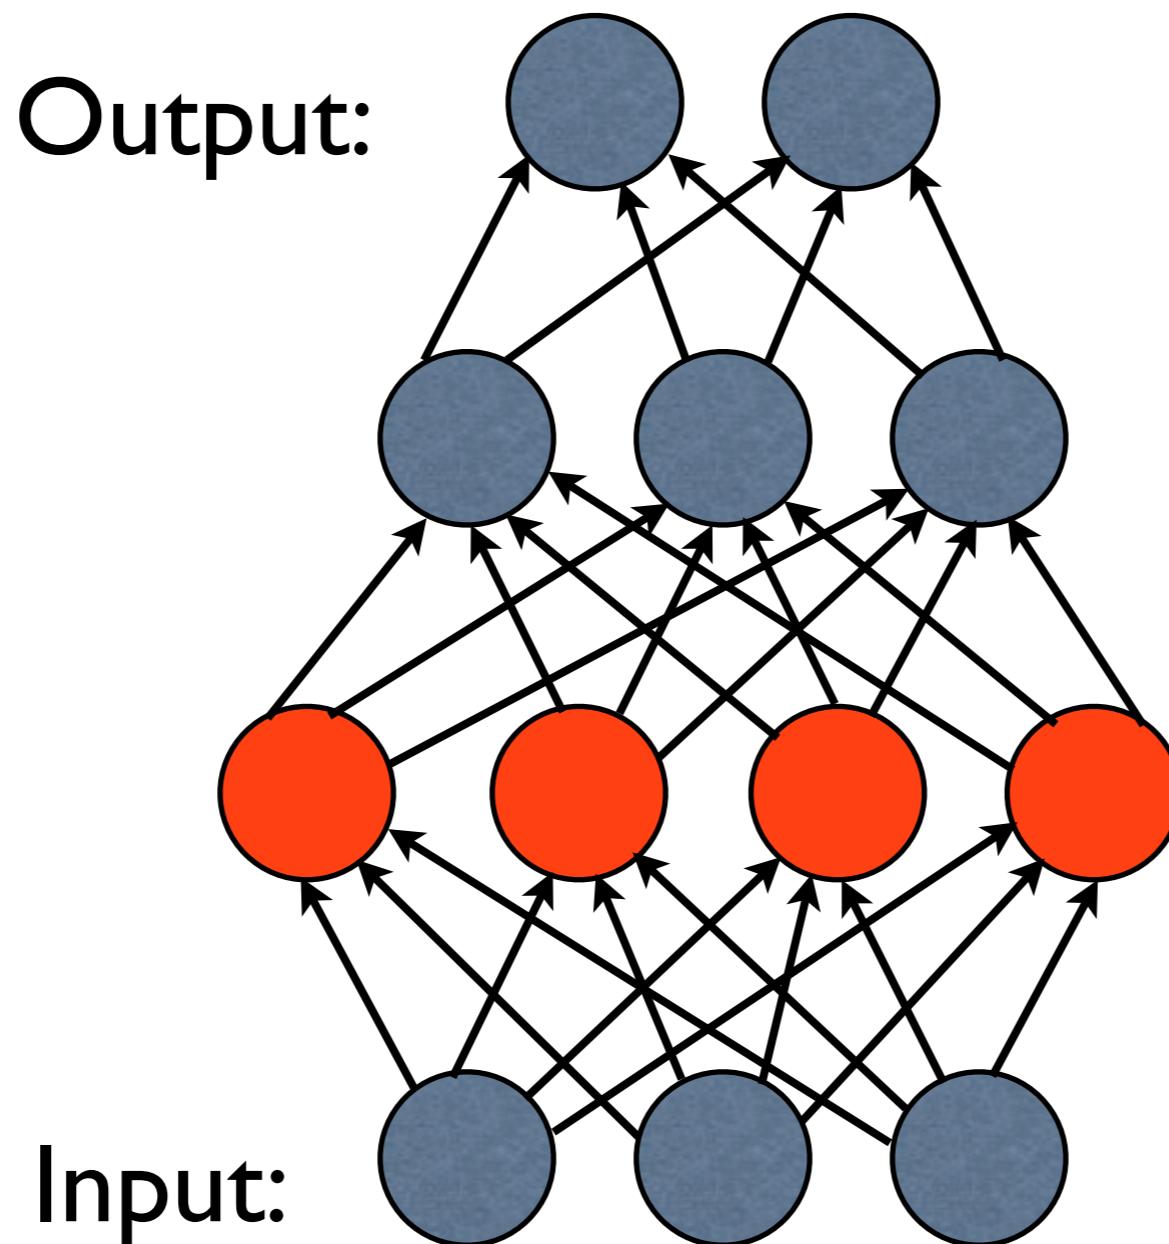


Input:

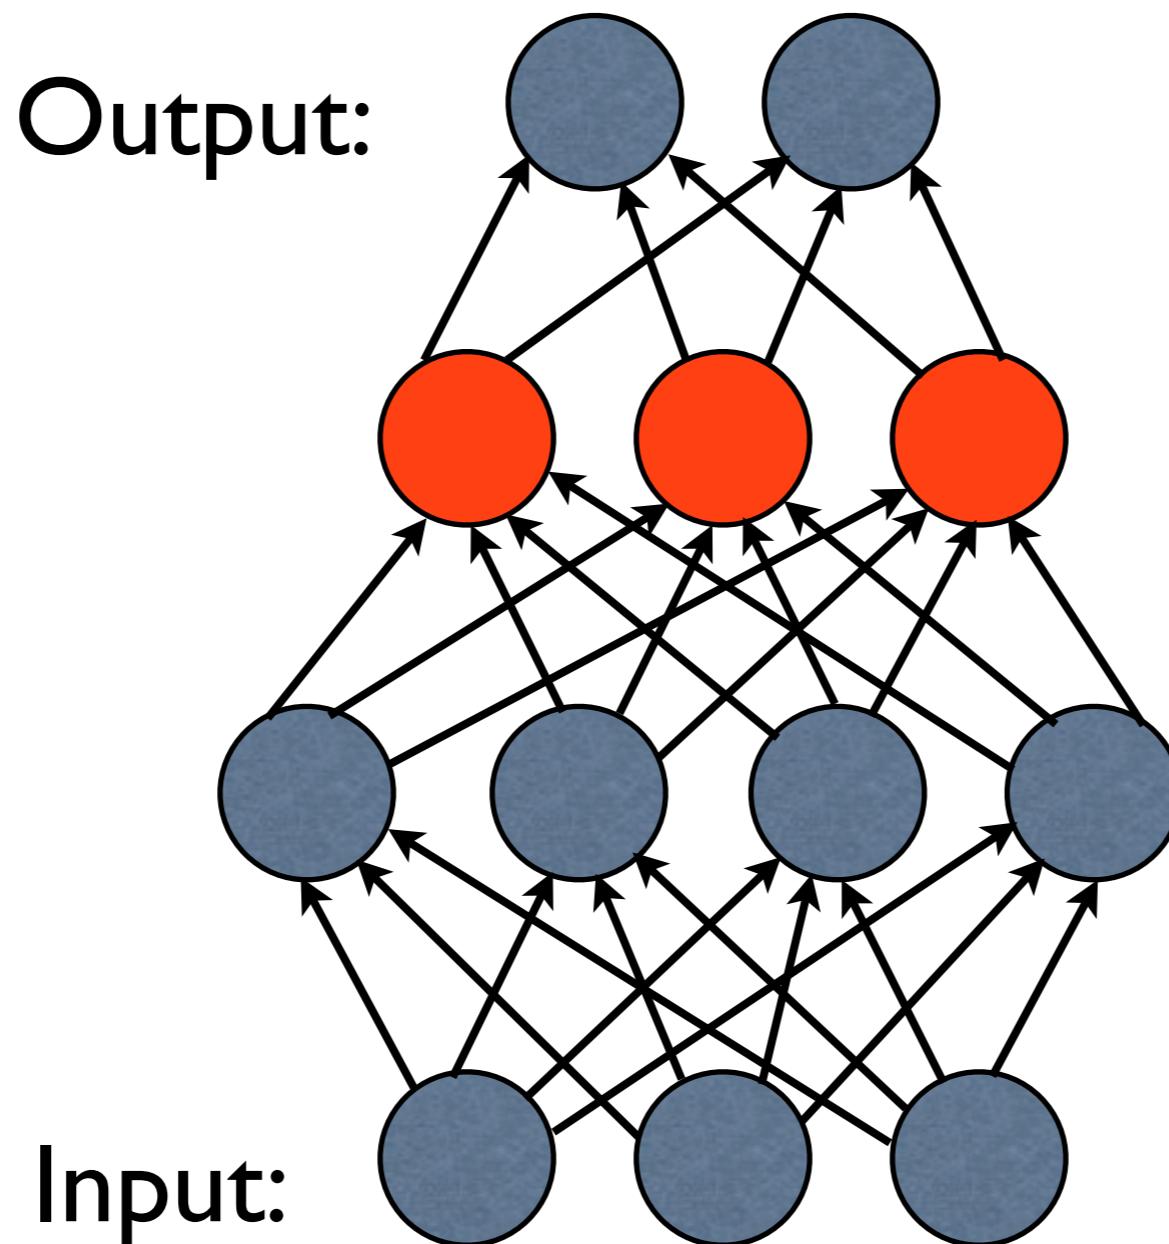
# Neural networks



# Neural networks

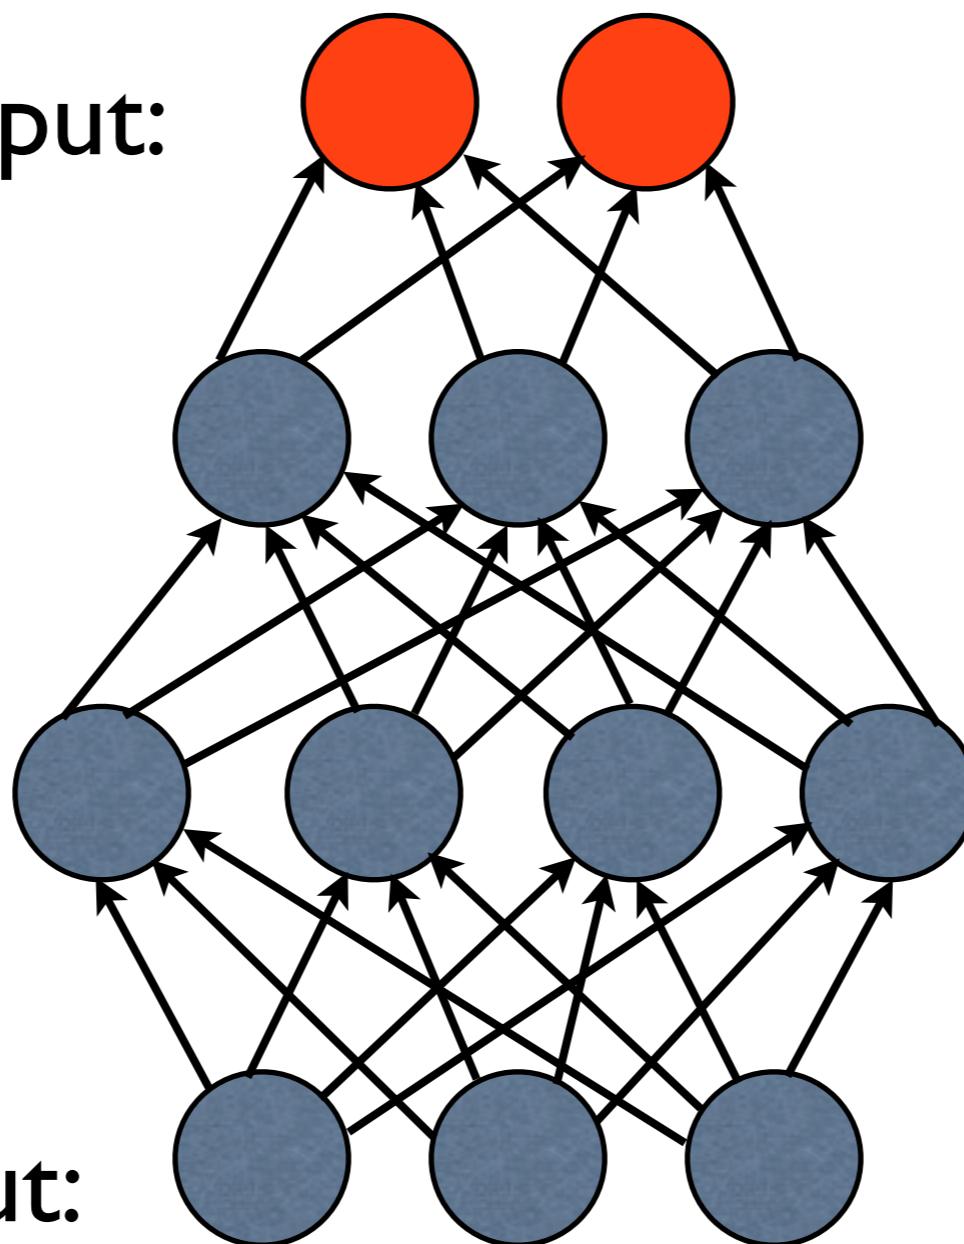


# Neural networks



# Neural networks

Output:



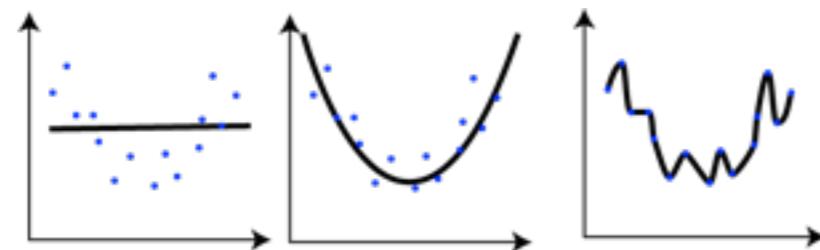
Input:

# Learning algorithm

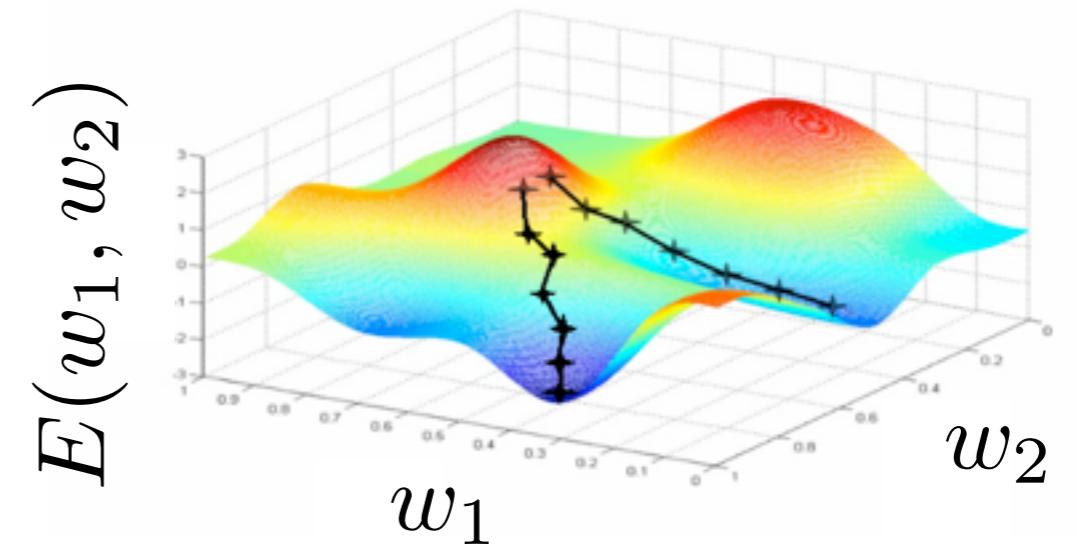
- **while** not done
  - pick a random training case (**x**, **y**)
  - run neuronal network on input **x**
  - modify connection weights to make prediction closer to **y**

# Very difficult to train

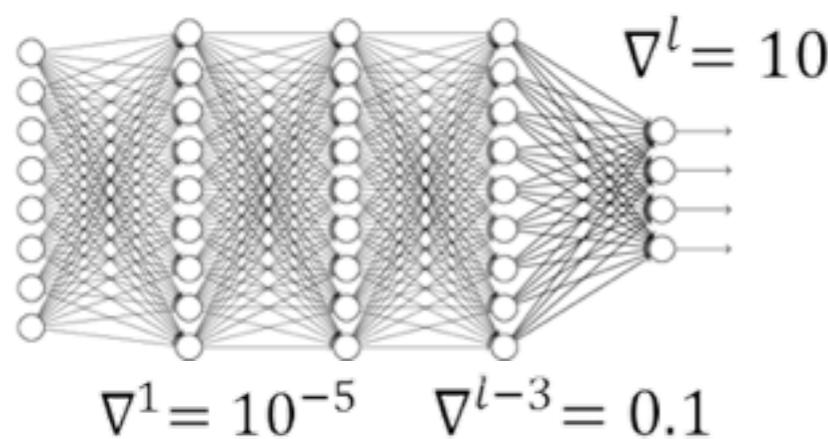
## Overfitting



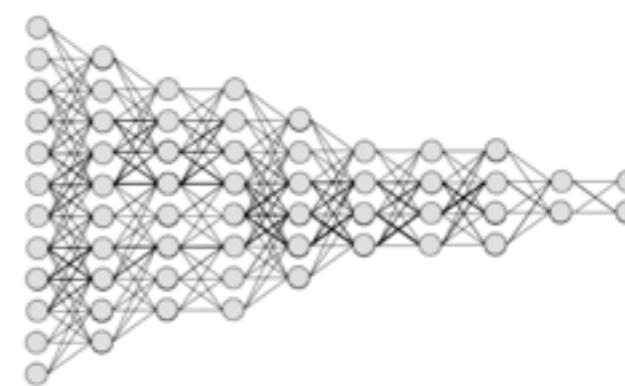
## Complex landscape

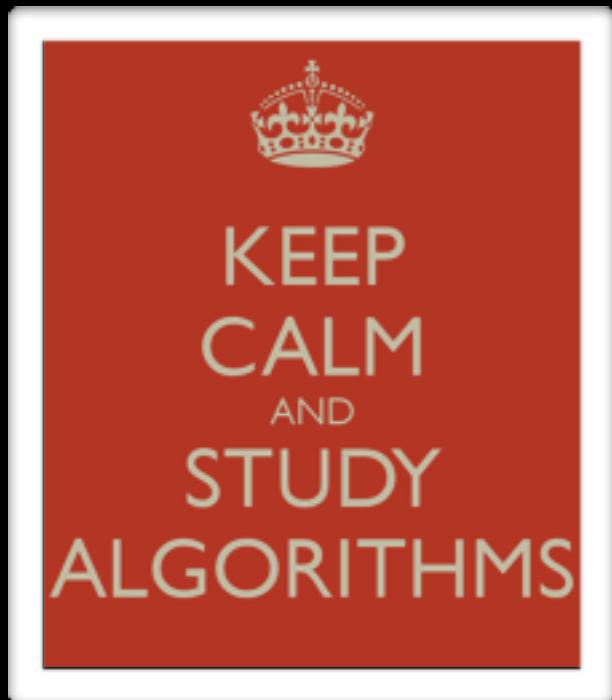


## Vanishing gradients



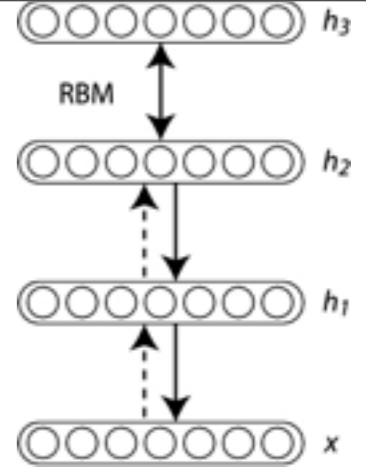
## Dimensionality





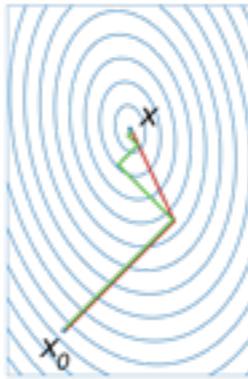
- Pre-training (weights initialization)

(complex landscape)



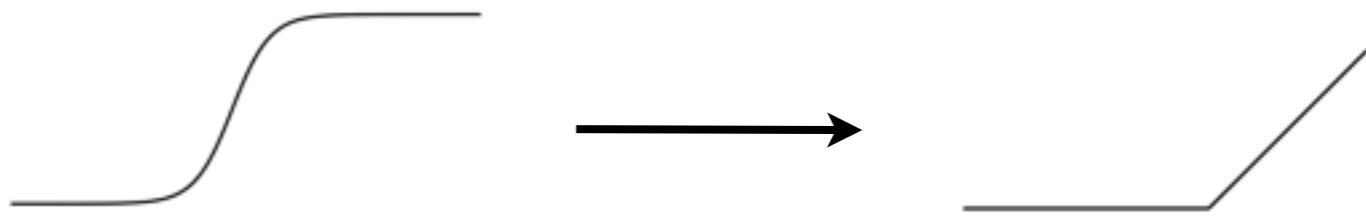
- Efficient descent algorithms

(complex landscape)



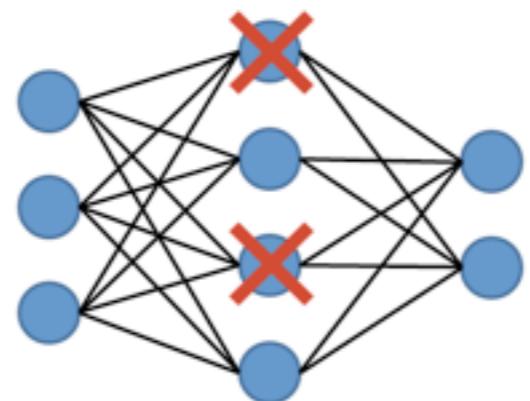
- Activation

(vanishing gradient)

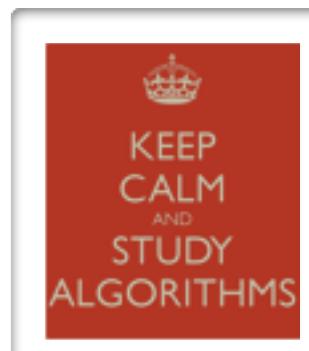


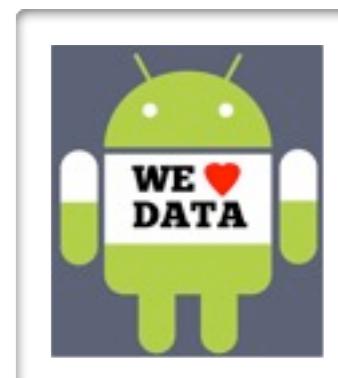
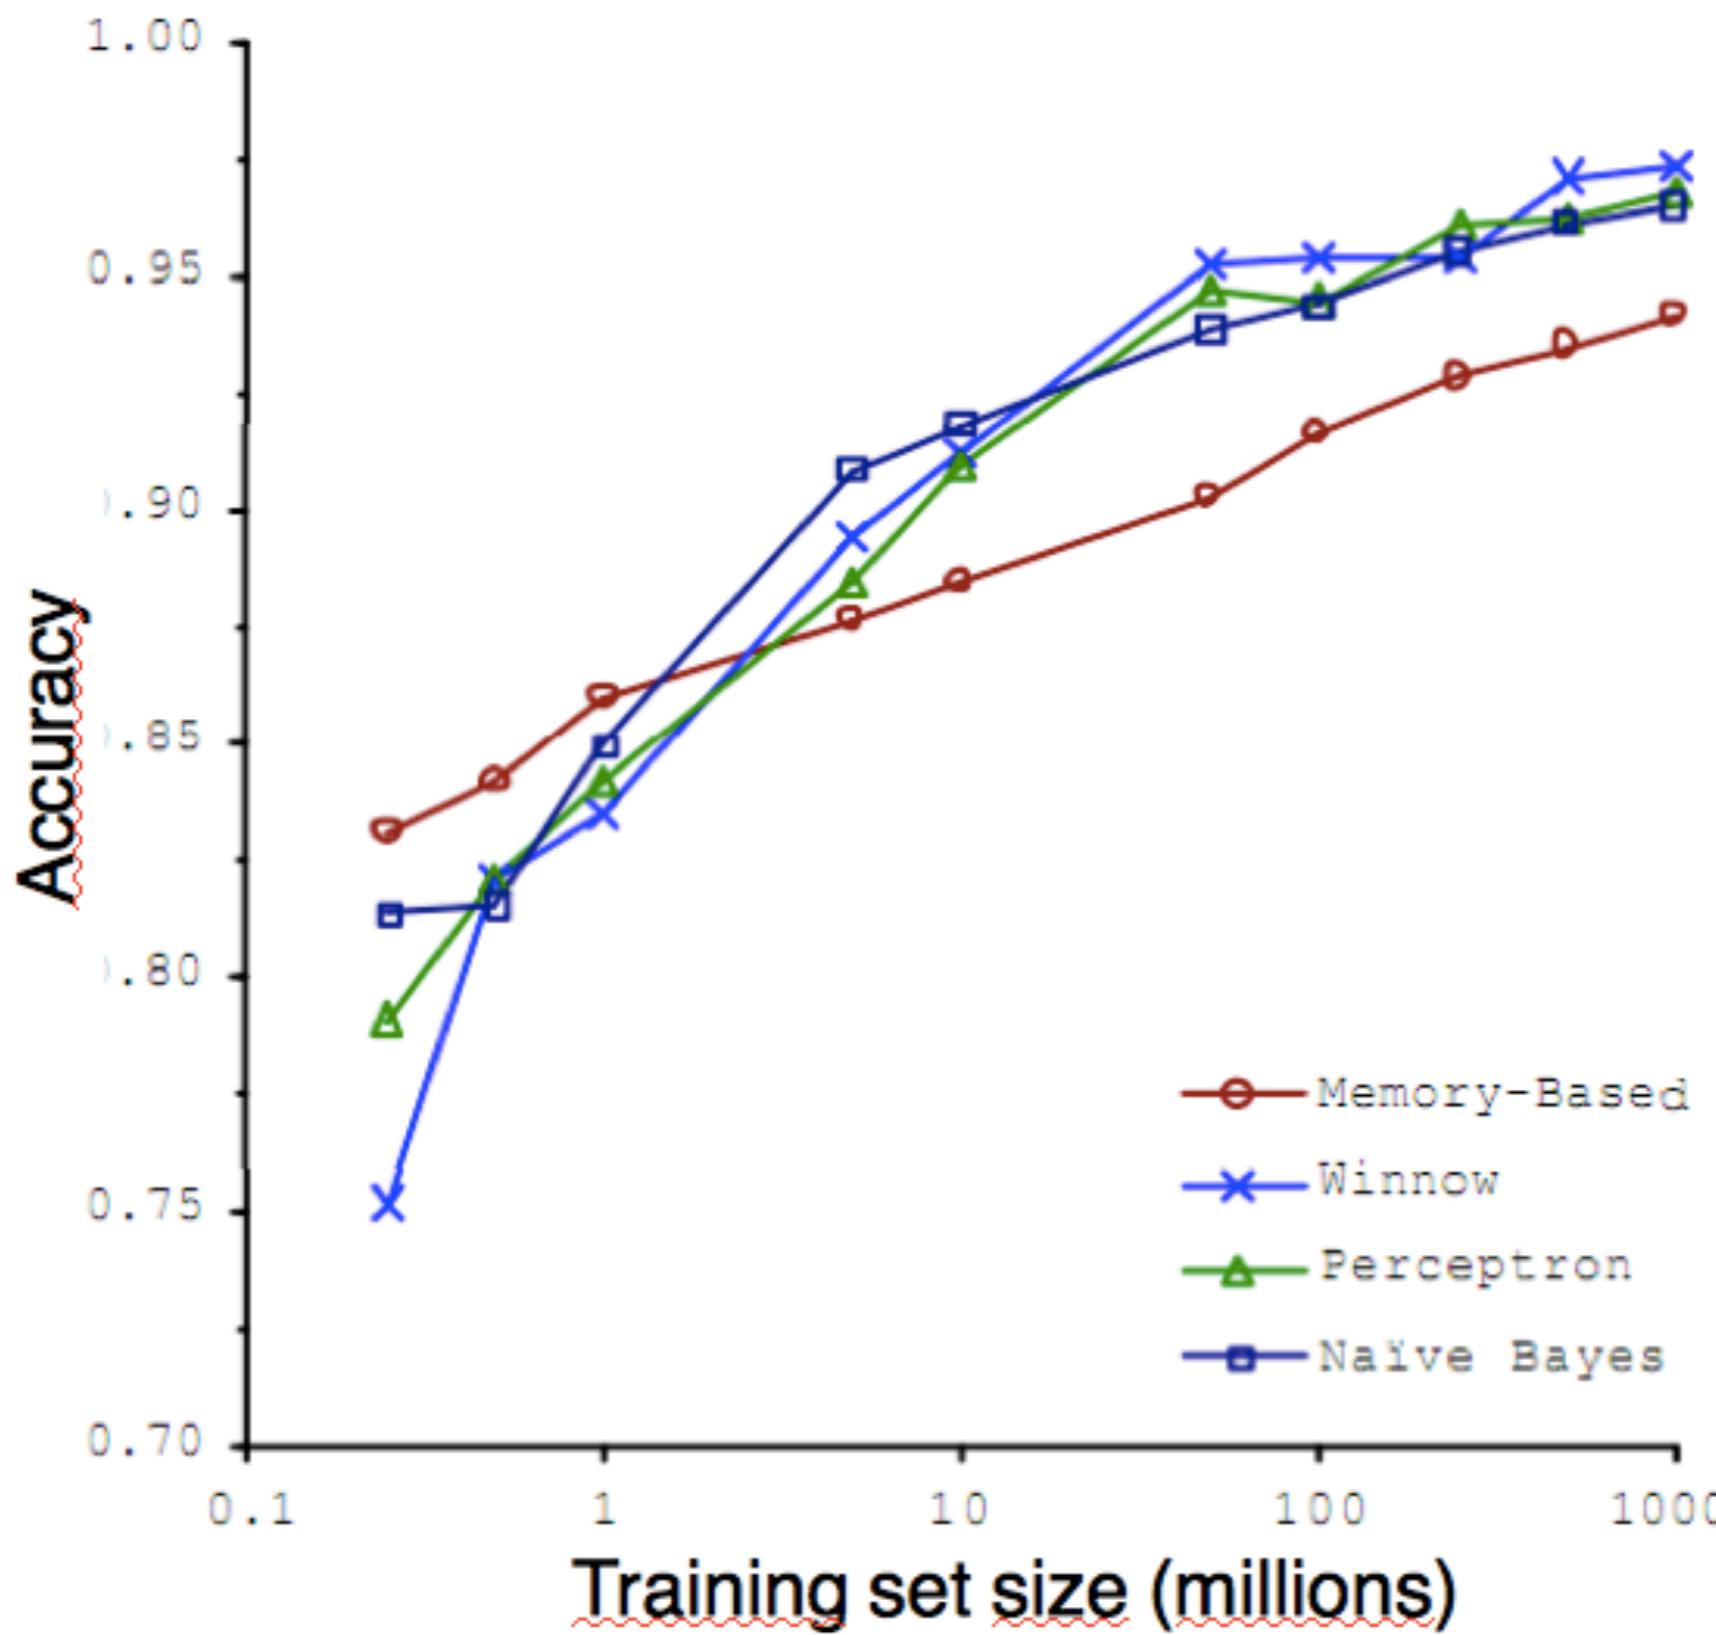
- Dropout

(overfitting)



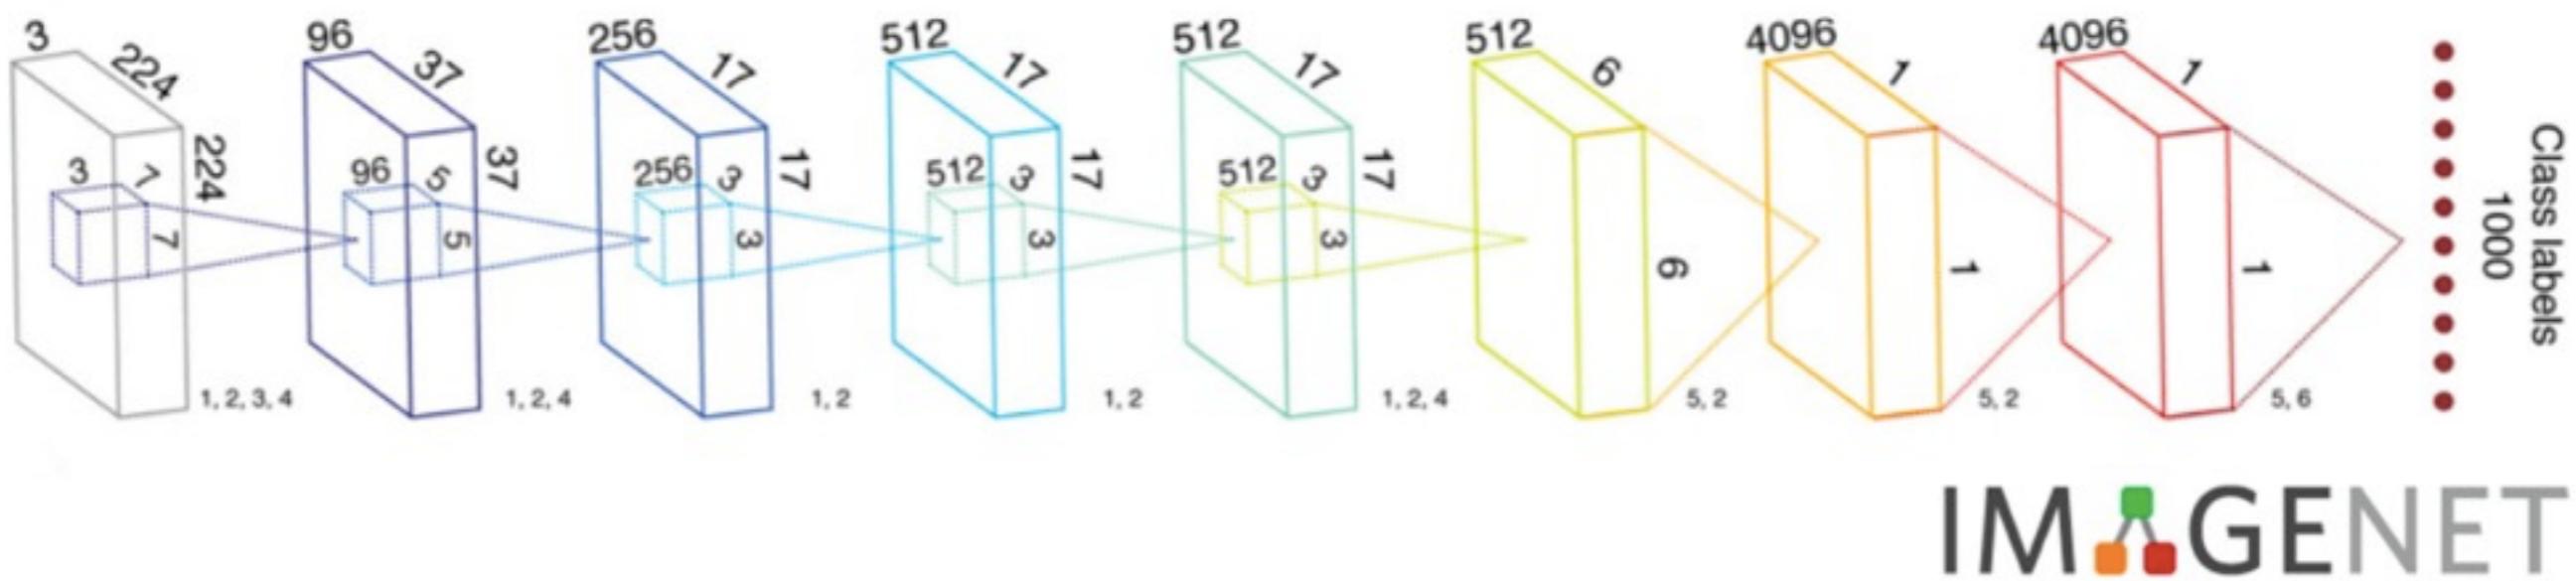
- Domain Prior Knowledge

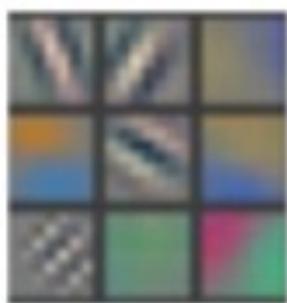
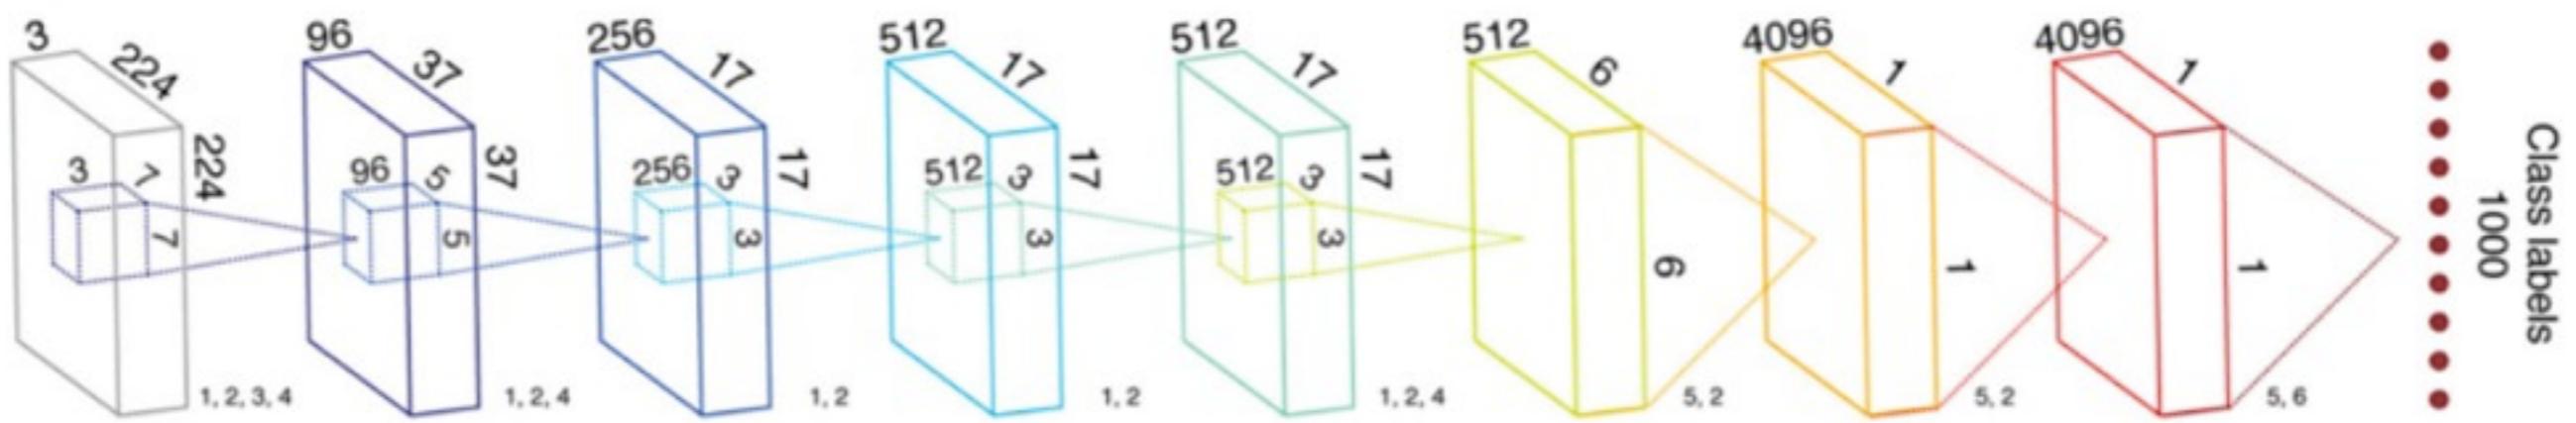




# Now that we are deep...

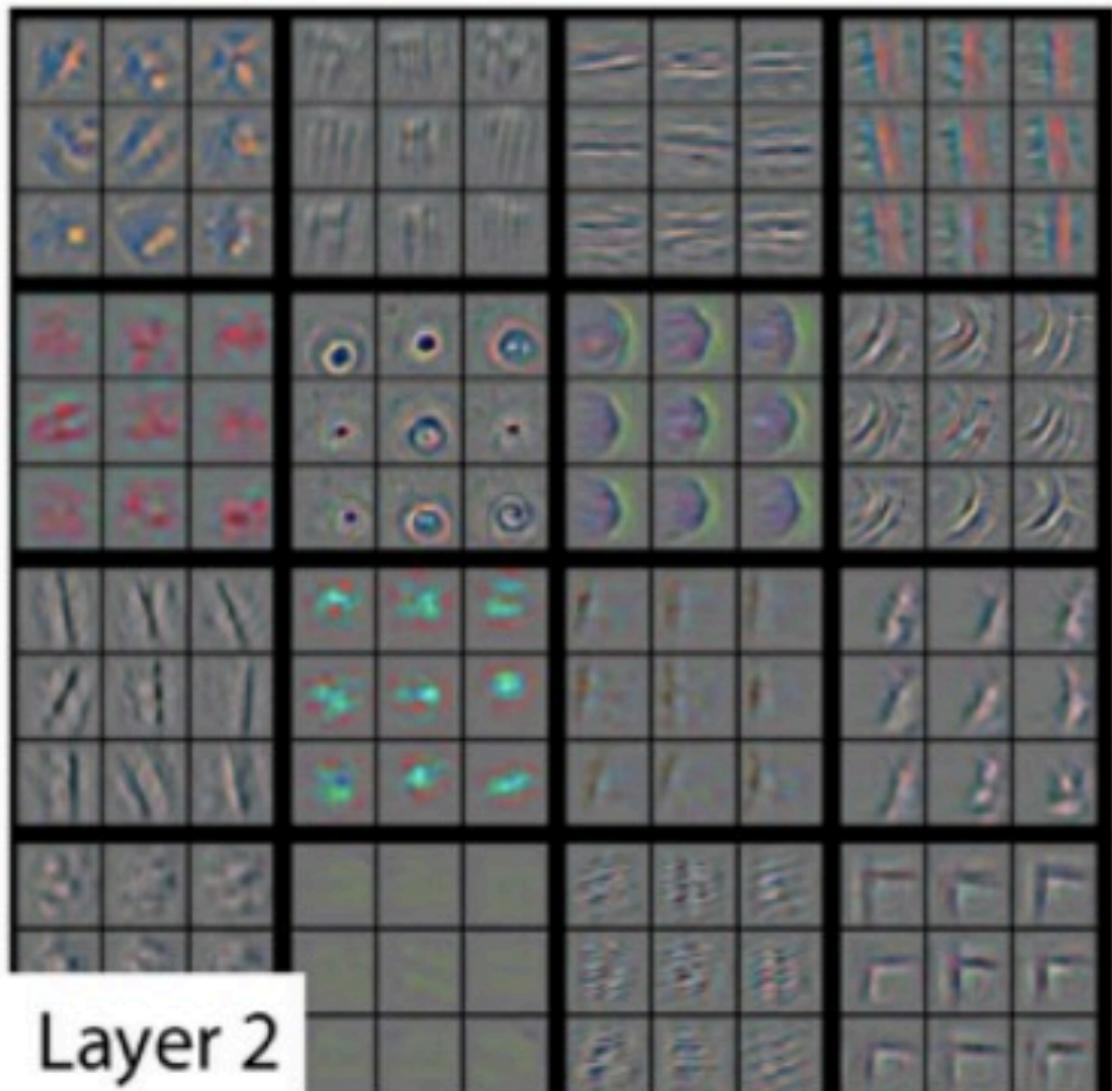
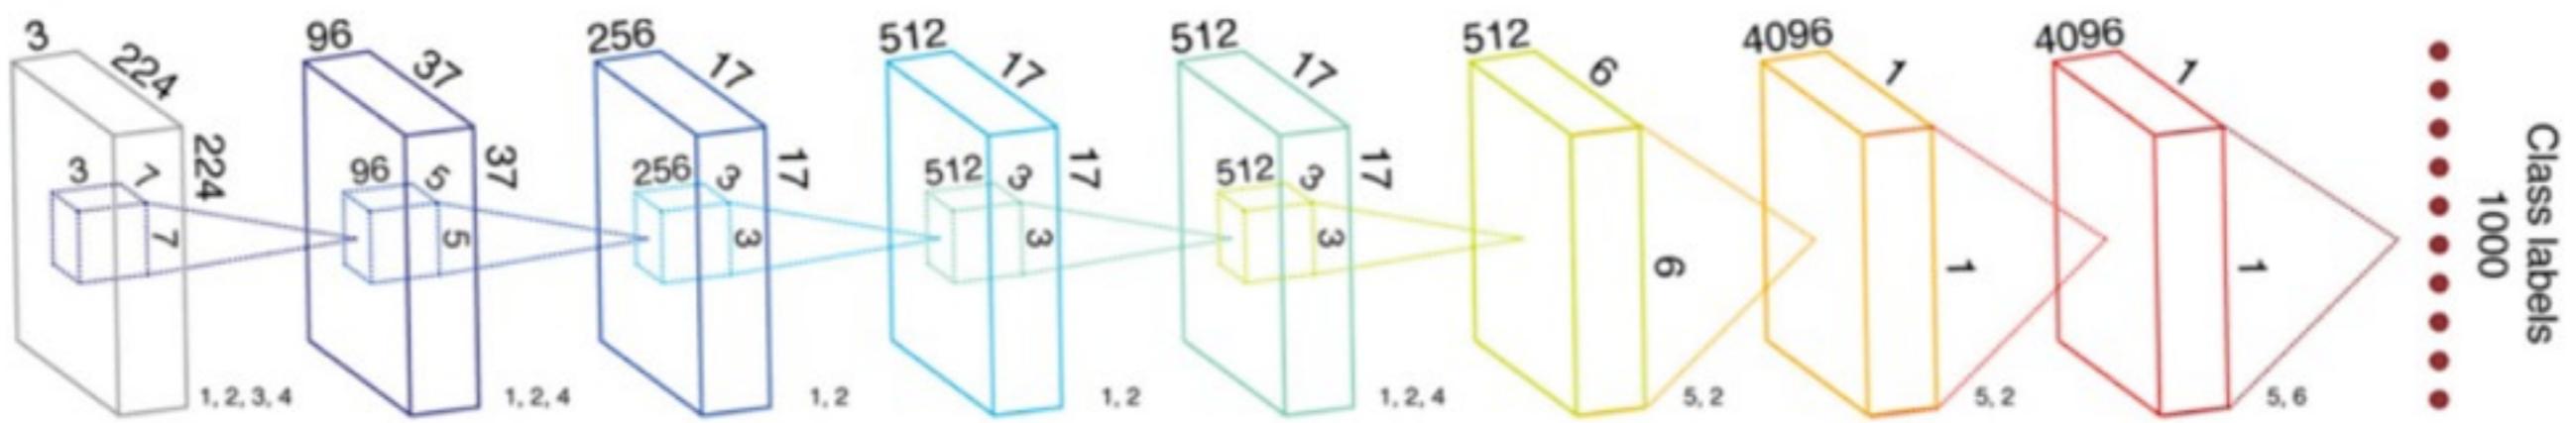
- Instead of hand-crafted features, let the algorithm build the relevant features for your problem
- More representational power
- Powerful function approximator



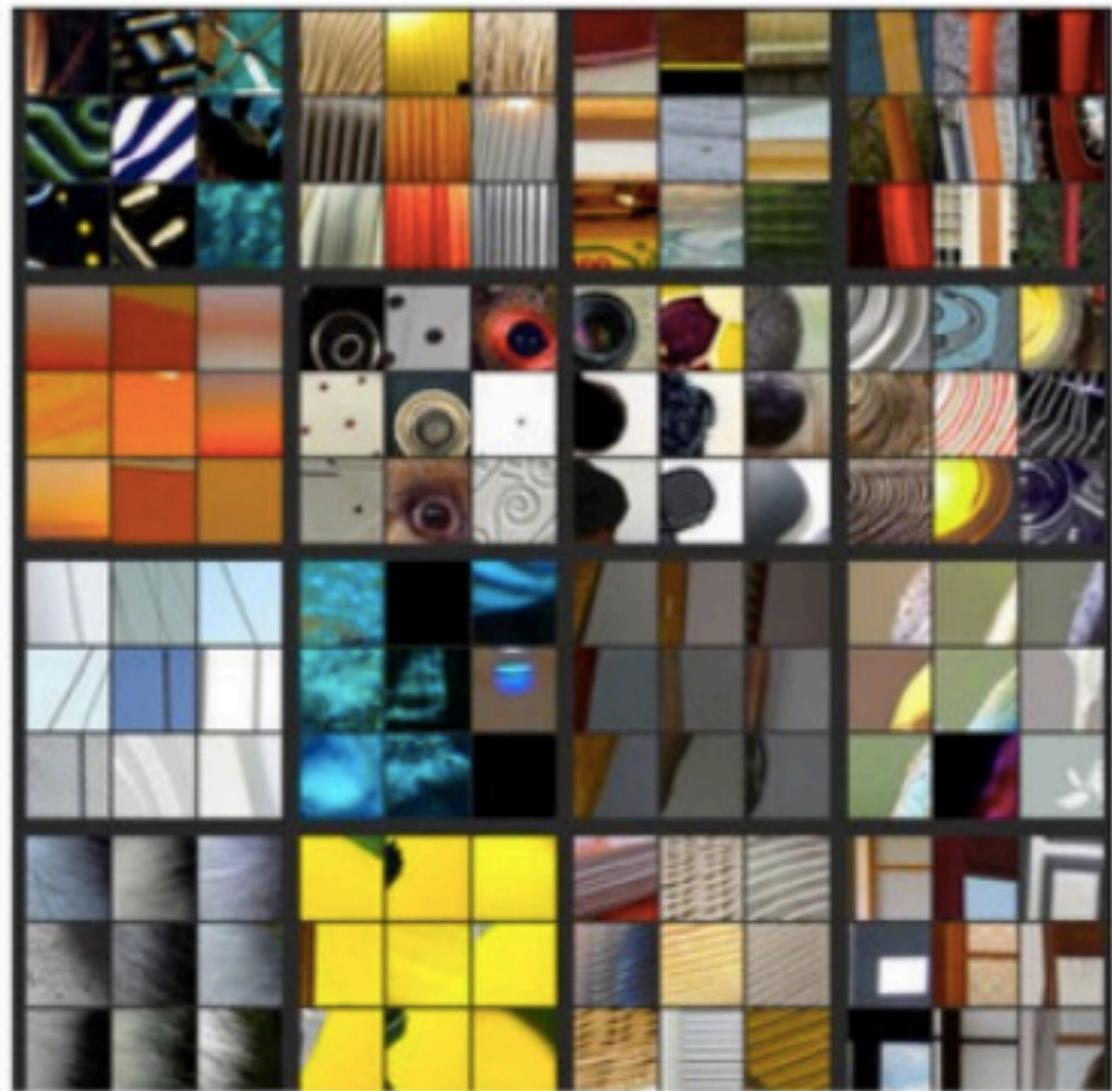


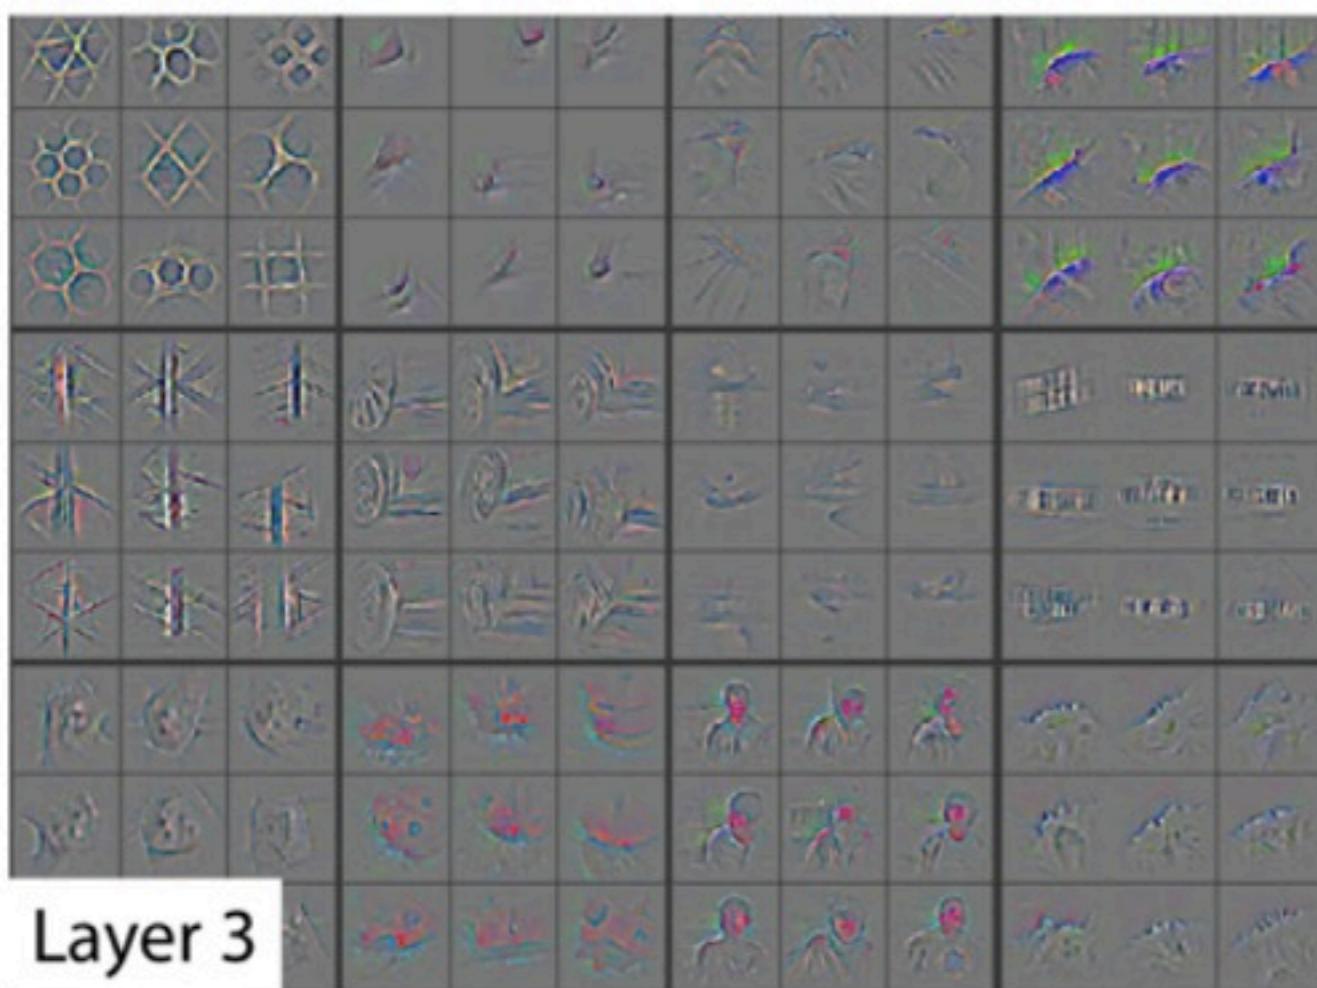
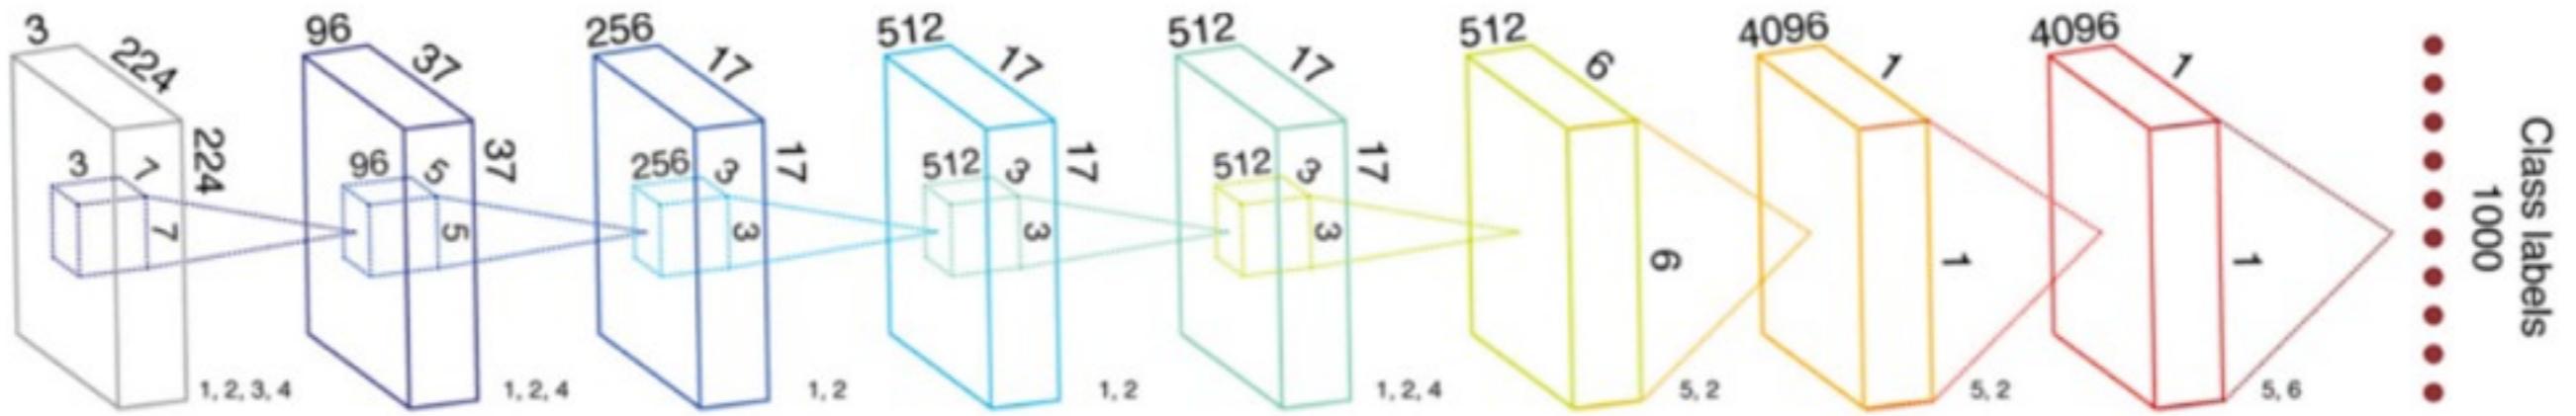
Layer 1





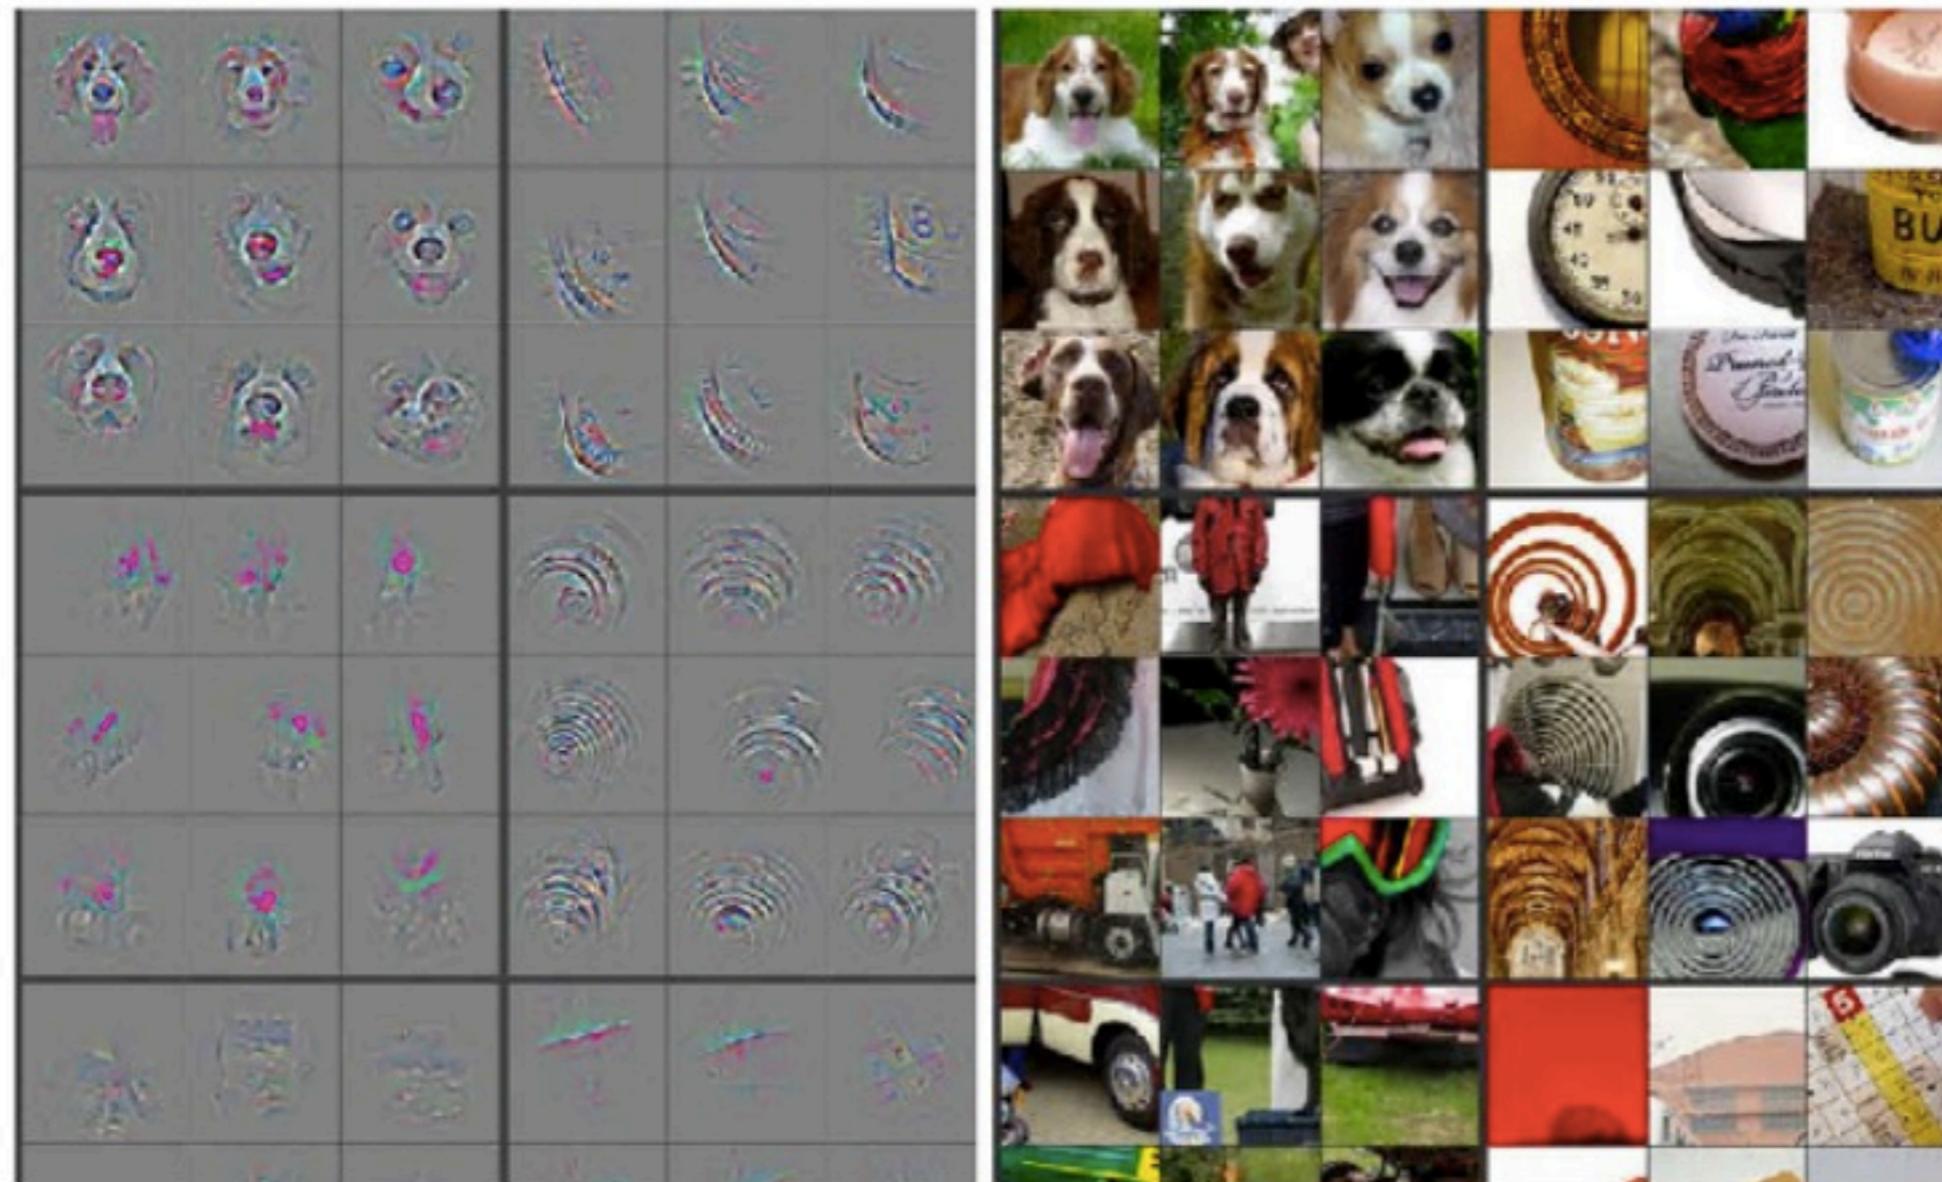
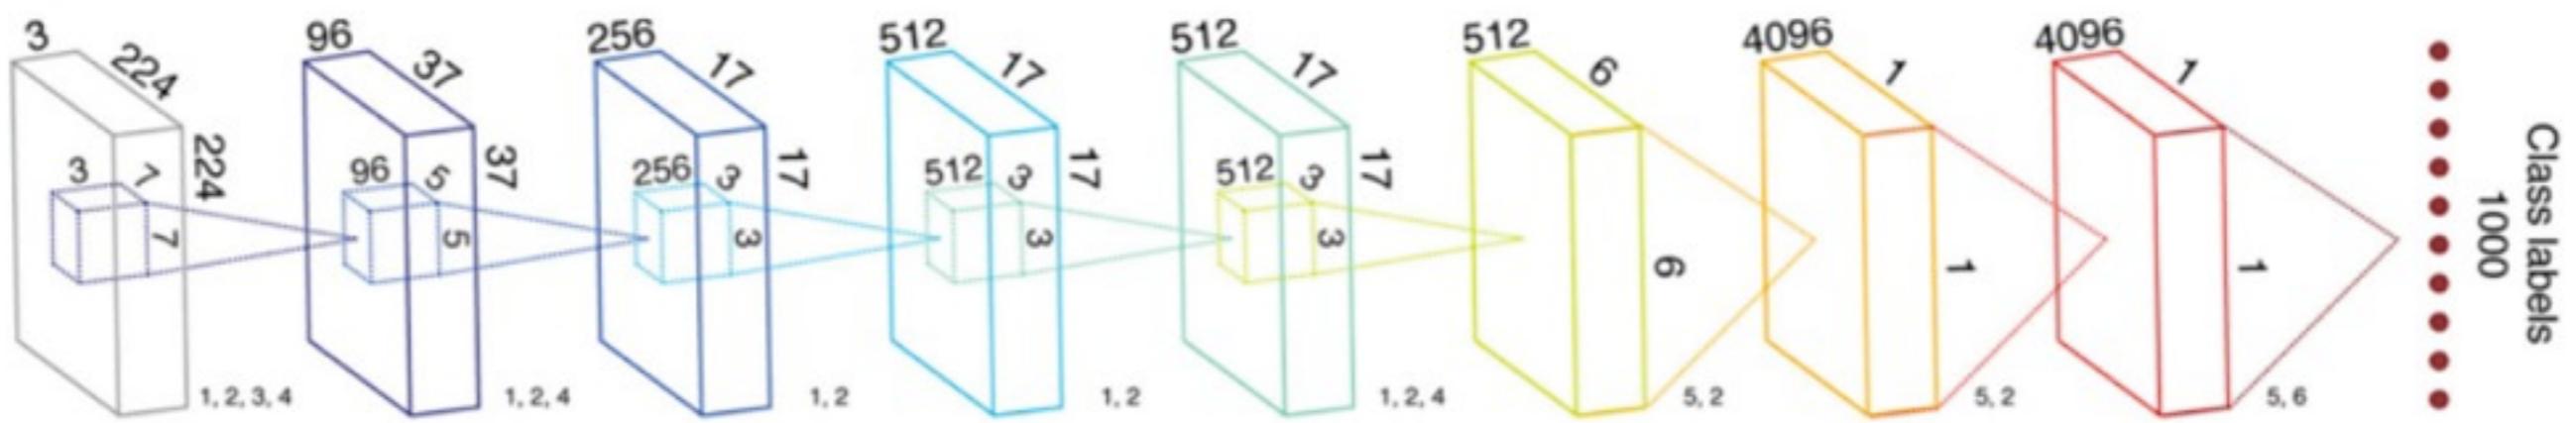
Layer 2

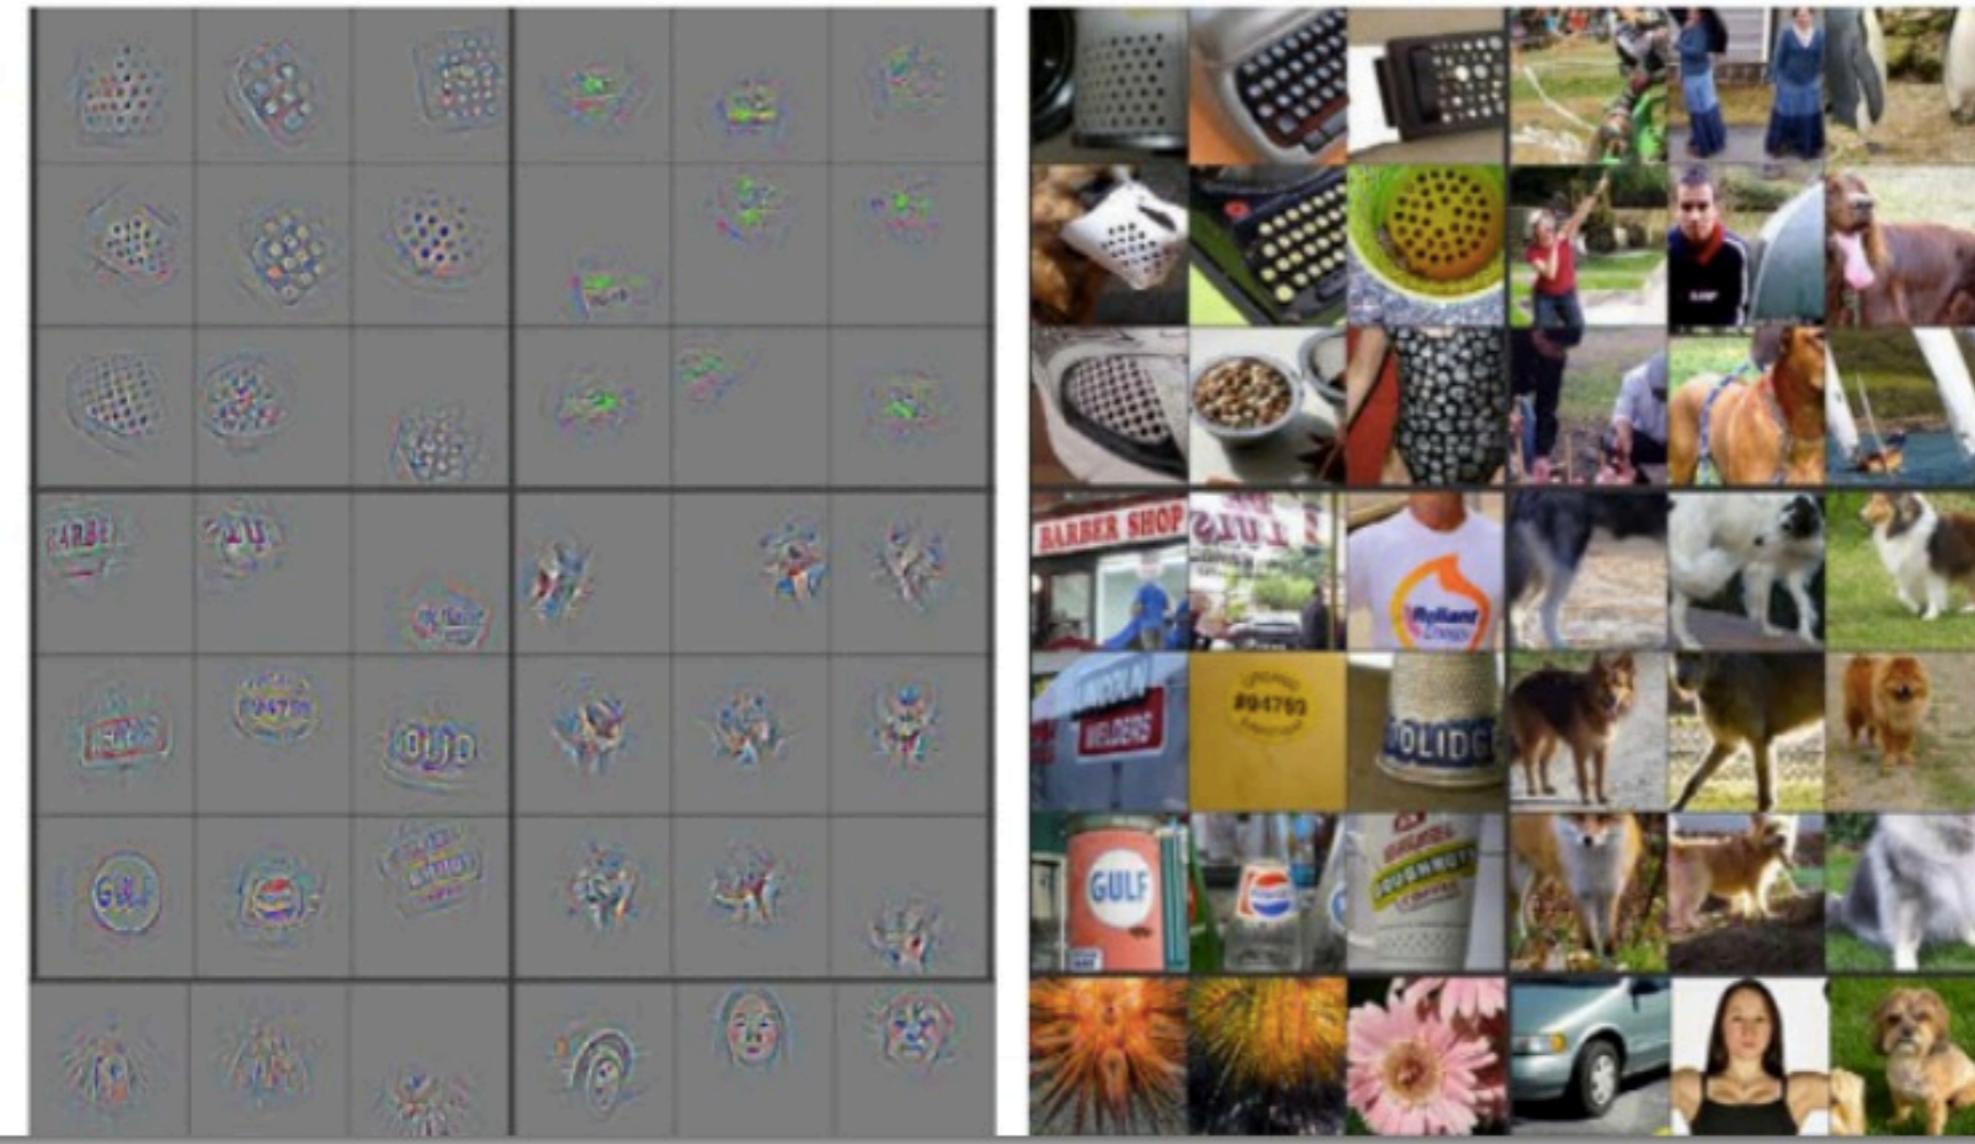
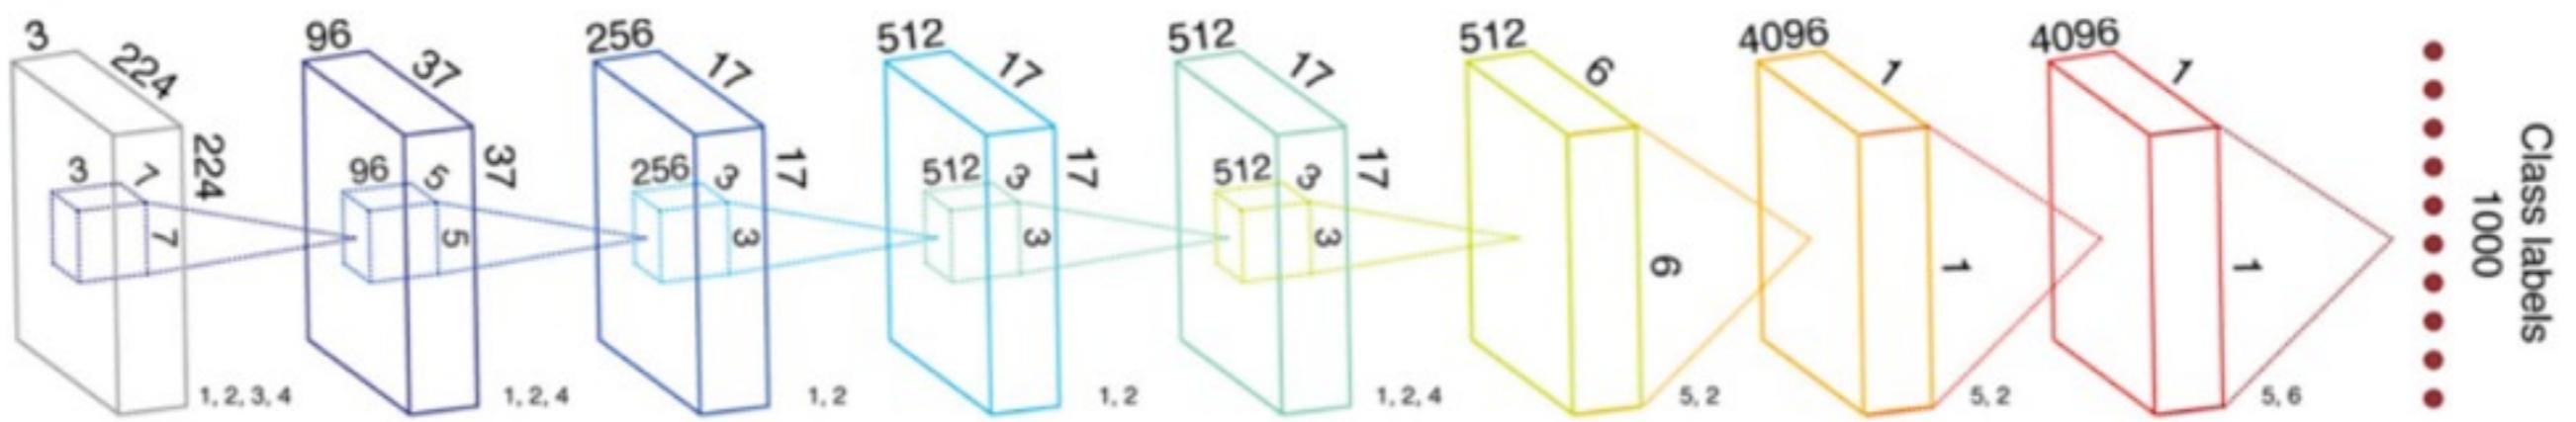




Layer 3



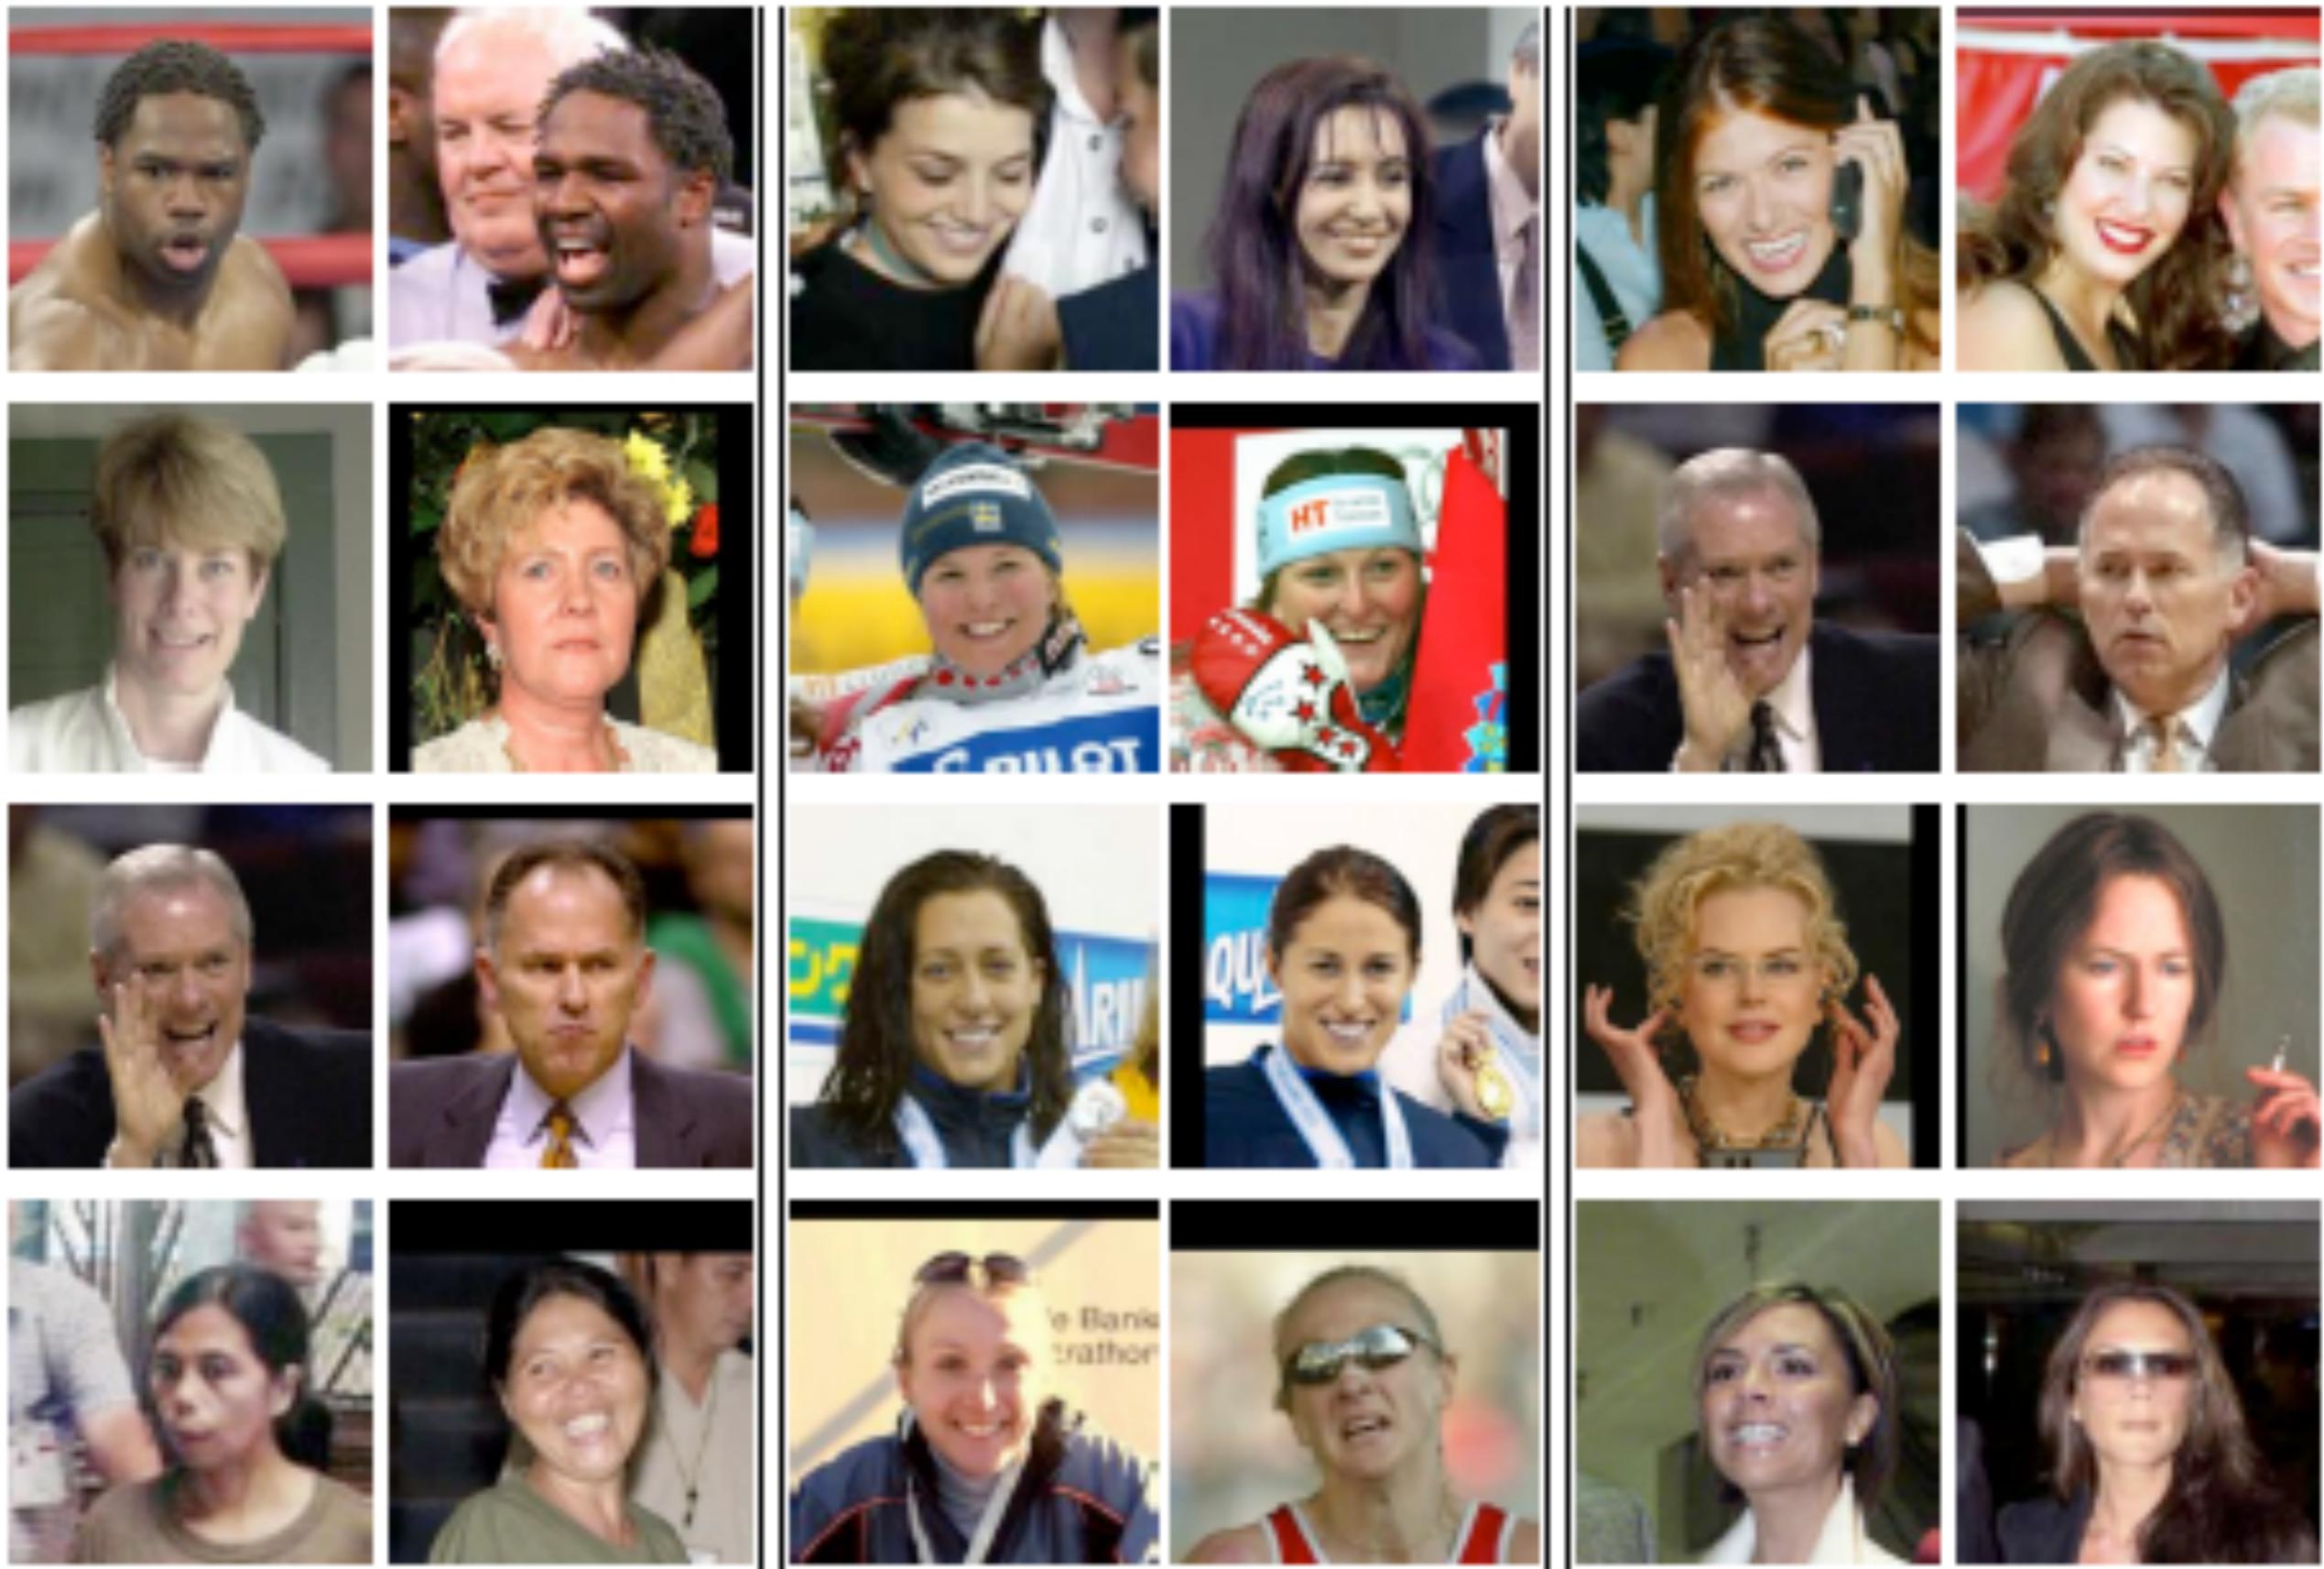


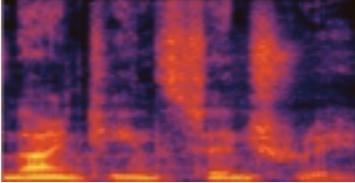
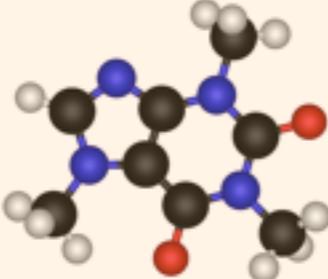


# False accept



# False reject



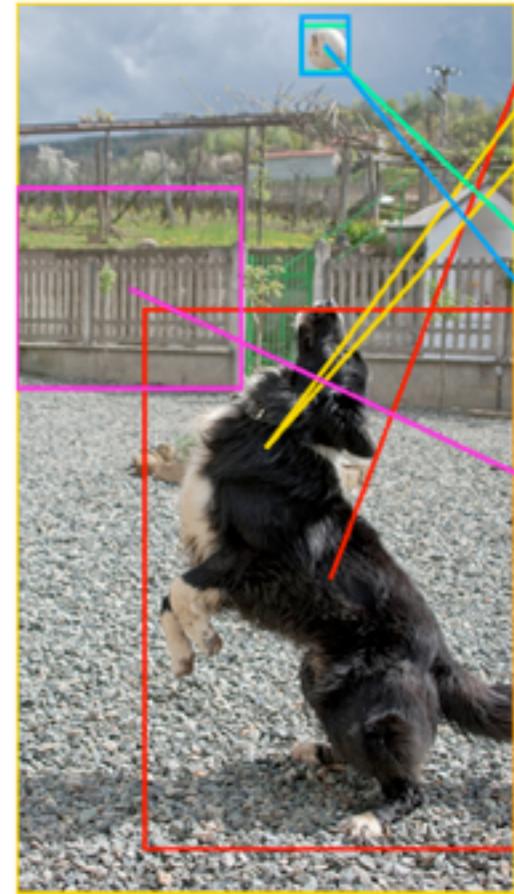
<b>Input</b>	<b>Output</b>
Pixels: 	“frog”
Audio: 	“na in tiz sen tu rii”
“Buenos dias, que tal estas?”	“Guten morgen, wie geht es dir?”
	“Toxic”
* <pre>j=8584 for x in range(8):     j+=920 b=(1500+j) print((b+7567))</pre>	25011.



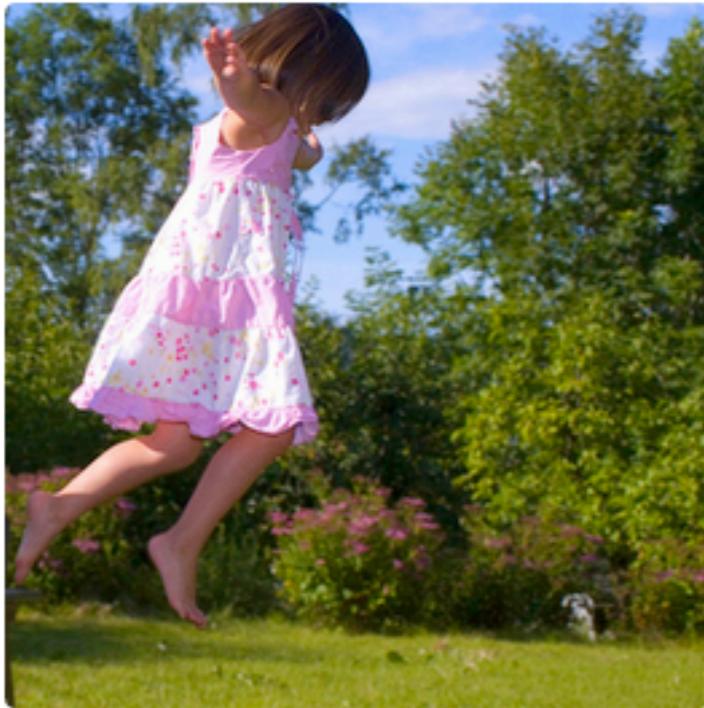
"man in black shirt is playing guitar."



"construction worker in orange safety vest is working on road."



1.31 dog
0.31 plays
0.45 catch
-0.02 with
0.25 white
1.62 ball
-0.10 near
-0.07 wooden
0.22 fence

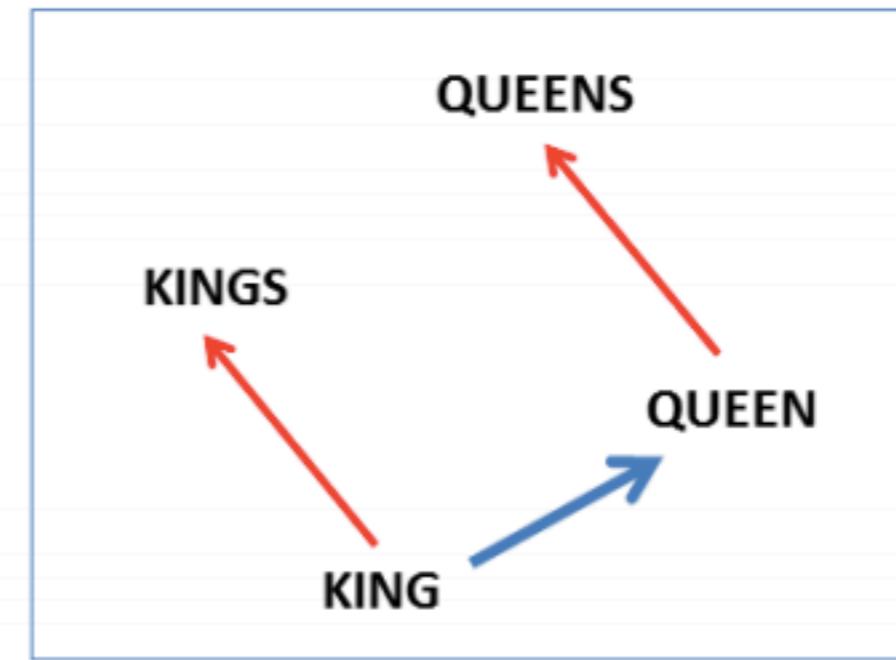
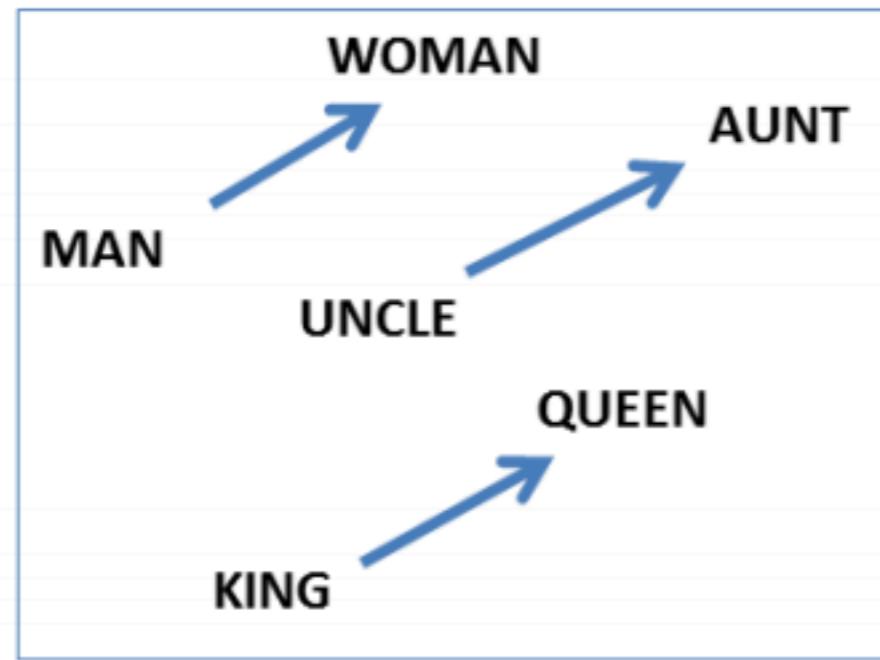


"girl in pink dress is jumping in air."



"black and white dog jumps over bar."

Karpathy, Fei-Fei, "Deep Visual-Semantic Alignments for Generating Image Descriptions" (2014)



- blue + red =



- blue + yellow =



- yellow + red =



- white + red =



Nearest images



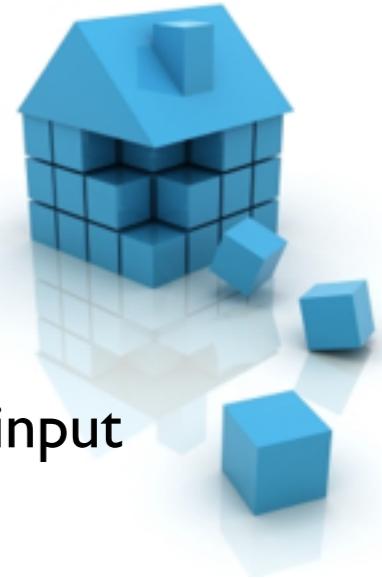
If all you have is a hammer  
in the toolbox, everything  
looks like a nail.”

- Bernard Baruch



# Best potential when:

- the structure of your problem/data naturally maps to a multilayer architecture
  - hierarchy of abstract features derivable from non-linear transformations of input
- enough data to learn features
  - unlabeled data can also be used for learning features



# Playing Atari with Deep Reinforcement Learning



# Playing Atari with Deep Reinforcement Learning



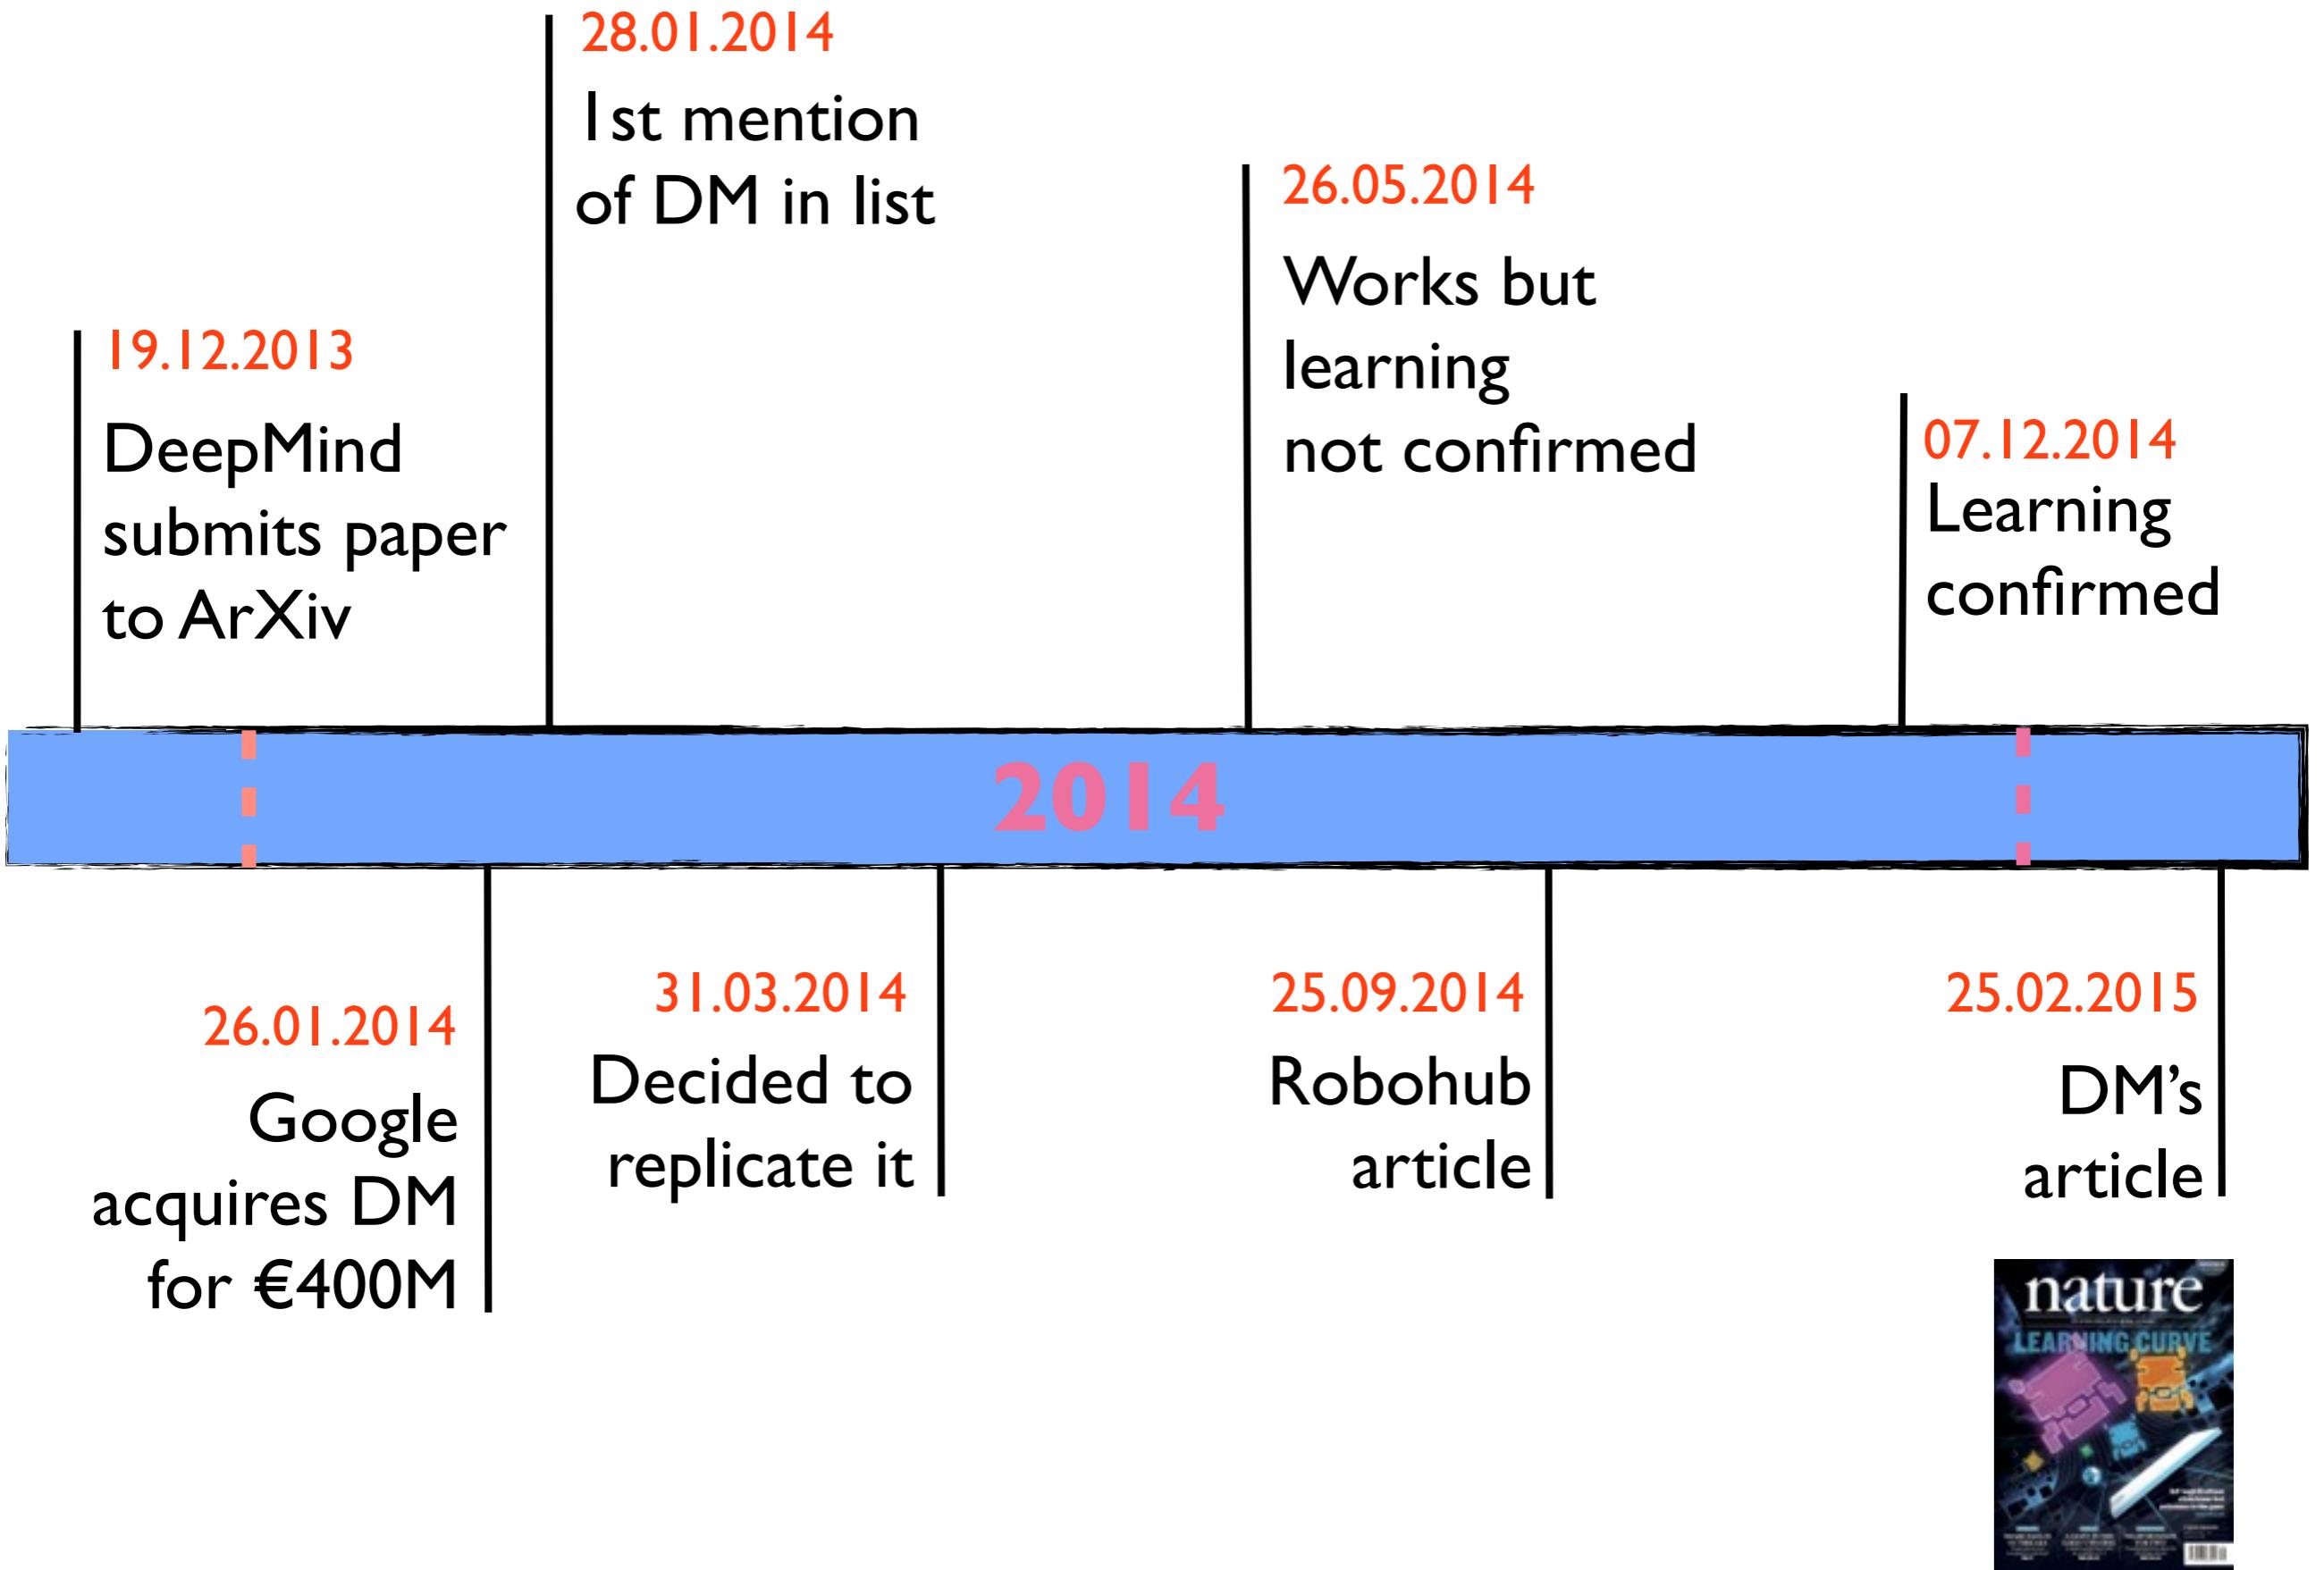
# Playing Atari with Deep Reinforcement Learning

Self-taught AI that learns to play better than humans

# Playing Atari with Deep Reinforcement Learning

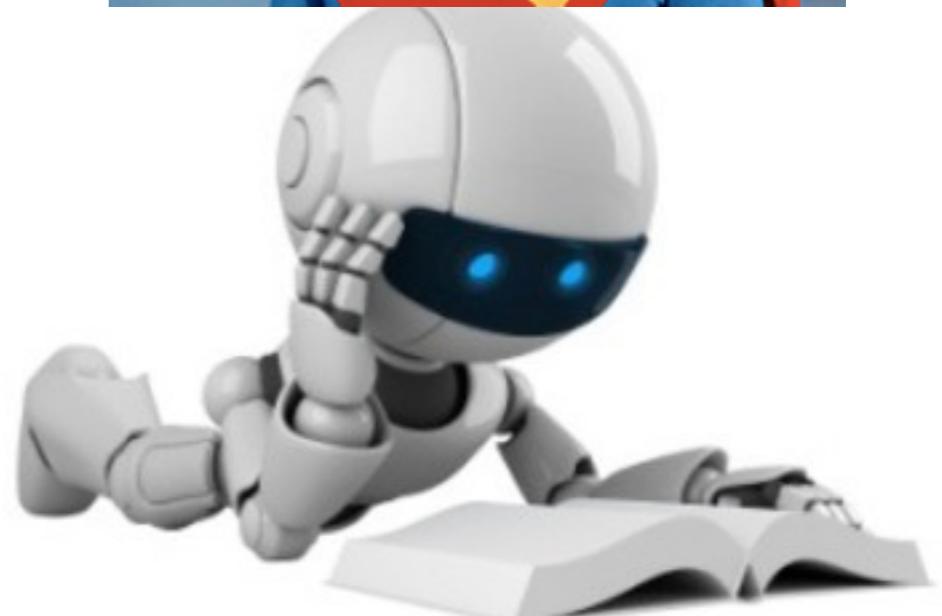


Self-taught AI that learns to play better than humans



# Why all the buzz...?

- Super-human level control
- Same algorithm learnt different games
- Learn from raw input



# Playing Atari with Deep Reinforcement Learning

Playing Atari with Deep Reinforcement Learning  
Mnih, et al Dec. 2013, arXiv:1312.5602

# Playing **Atari** with Deep Reinforcement Learning

Playing Atari with Deep Reinforcement Learning  
Mnih, et al Dec. 2013, arXiv:1312.5602



# Playing **Atari** with Deep Reinforcement Learning

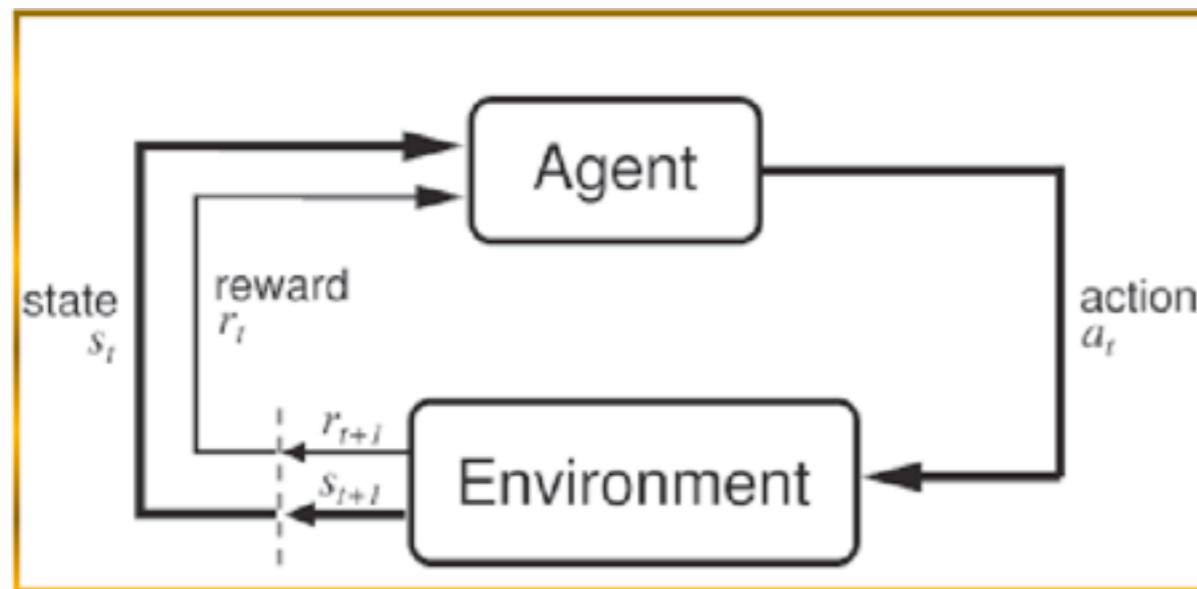
Playing Atari with Deep Reinforcement Learning  
Mnih, et al Dec. 2013, arXiv:1312.5602

# Playing Atari with Deep Reinforcement Learning

Playing Atari with Deep Reinforcement Learning  
Mnih, et al Dec. 2013, arXiv:1312.5602



# Playing Atari with Deep Reinforcement Learning



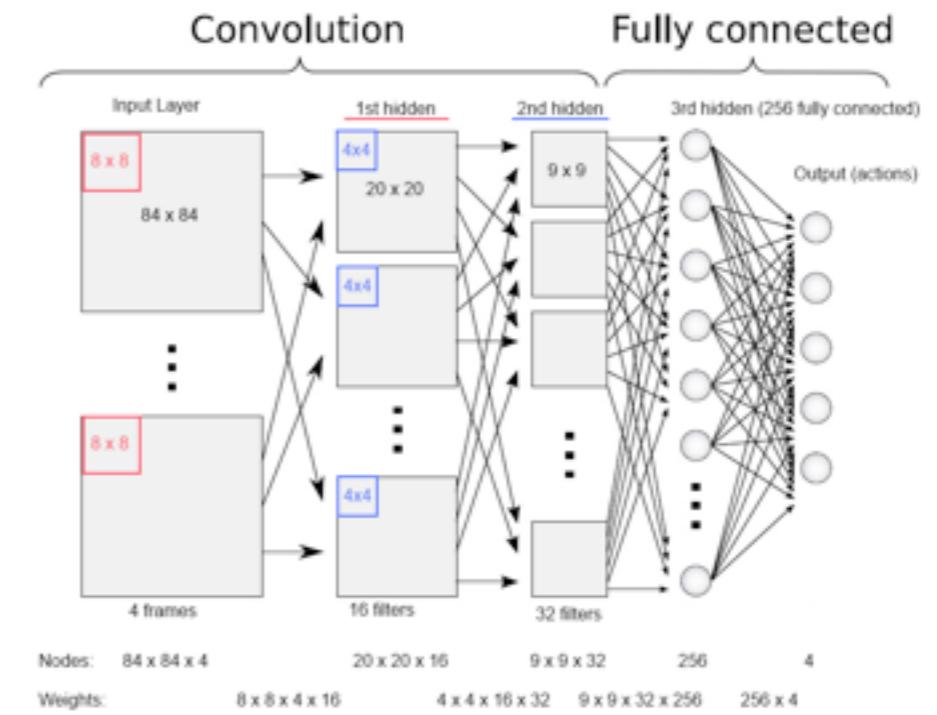
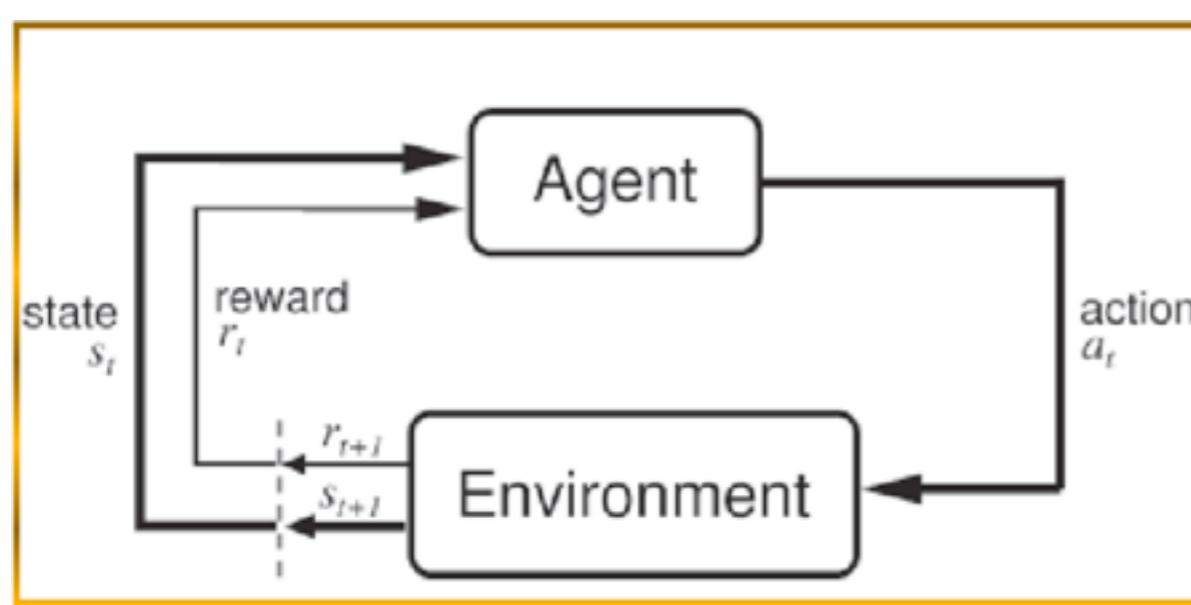
Playing Atari with Deep Reinforcement Learning  
Mnih, et al Dec. 2013, arXiv:1312.5602

# Playing Atari with **Deep** Reinforcement Learning

Playing Atari with Deep Reinforcement Learning  
Mnih, et al Dec. 2013, arXiv:1312.5602



# Playing Atari with Deep Reinforcement Learning



Playing Atari with Deep Reinforcement Learning  
Mnih, et al Dec. 2013, arXiv:1312.5602



Pong



Breakout



Seaquest



Beam Rider



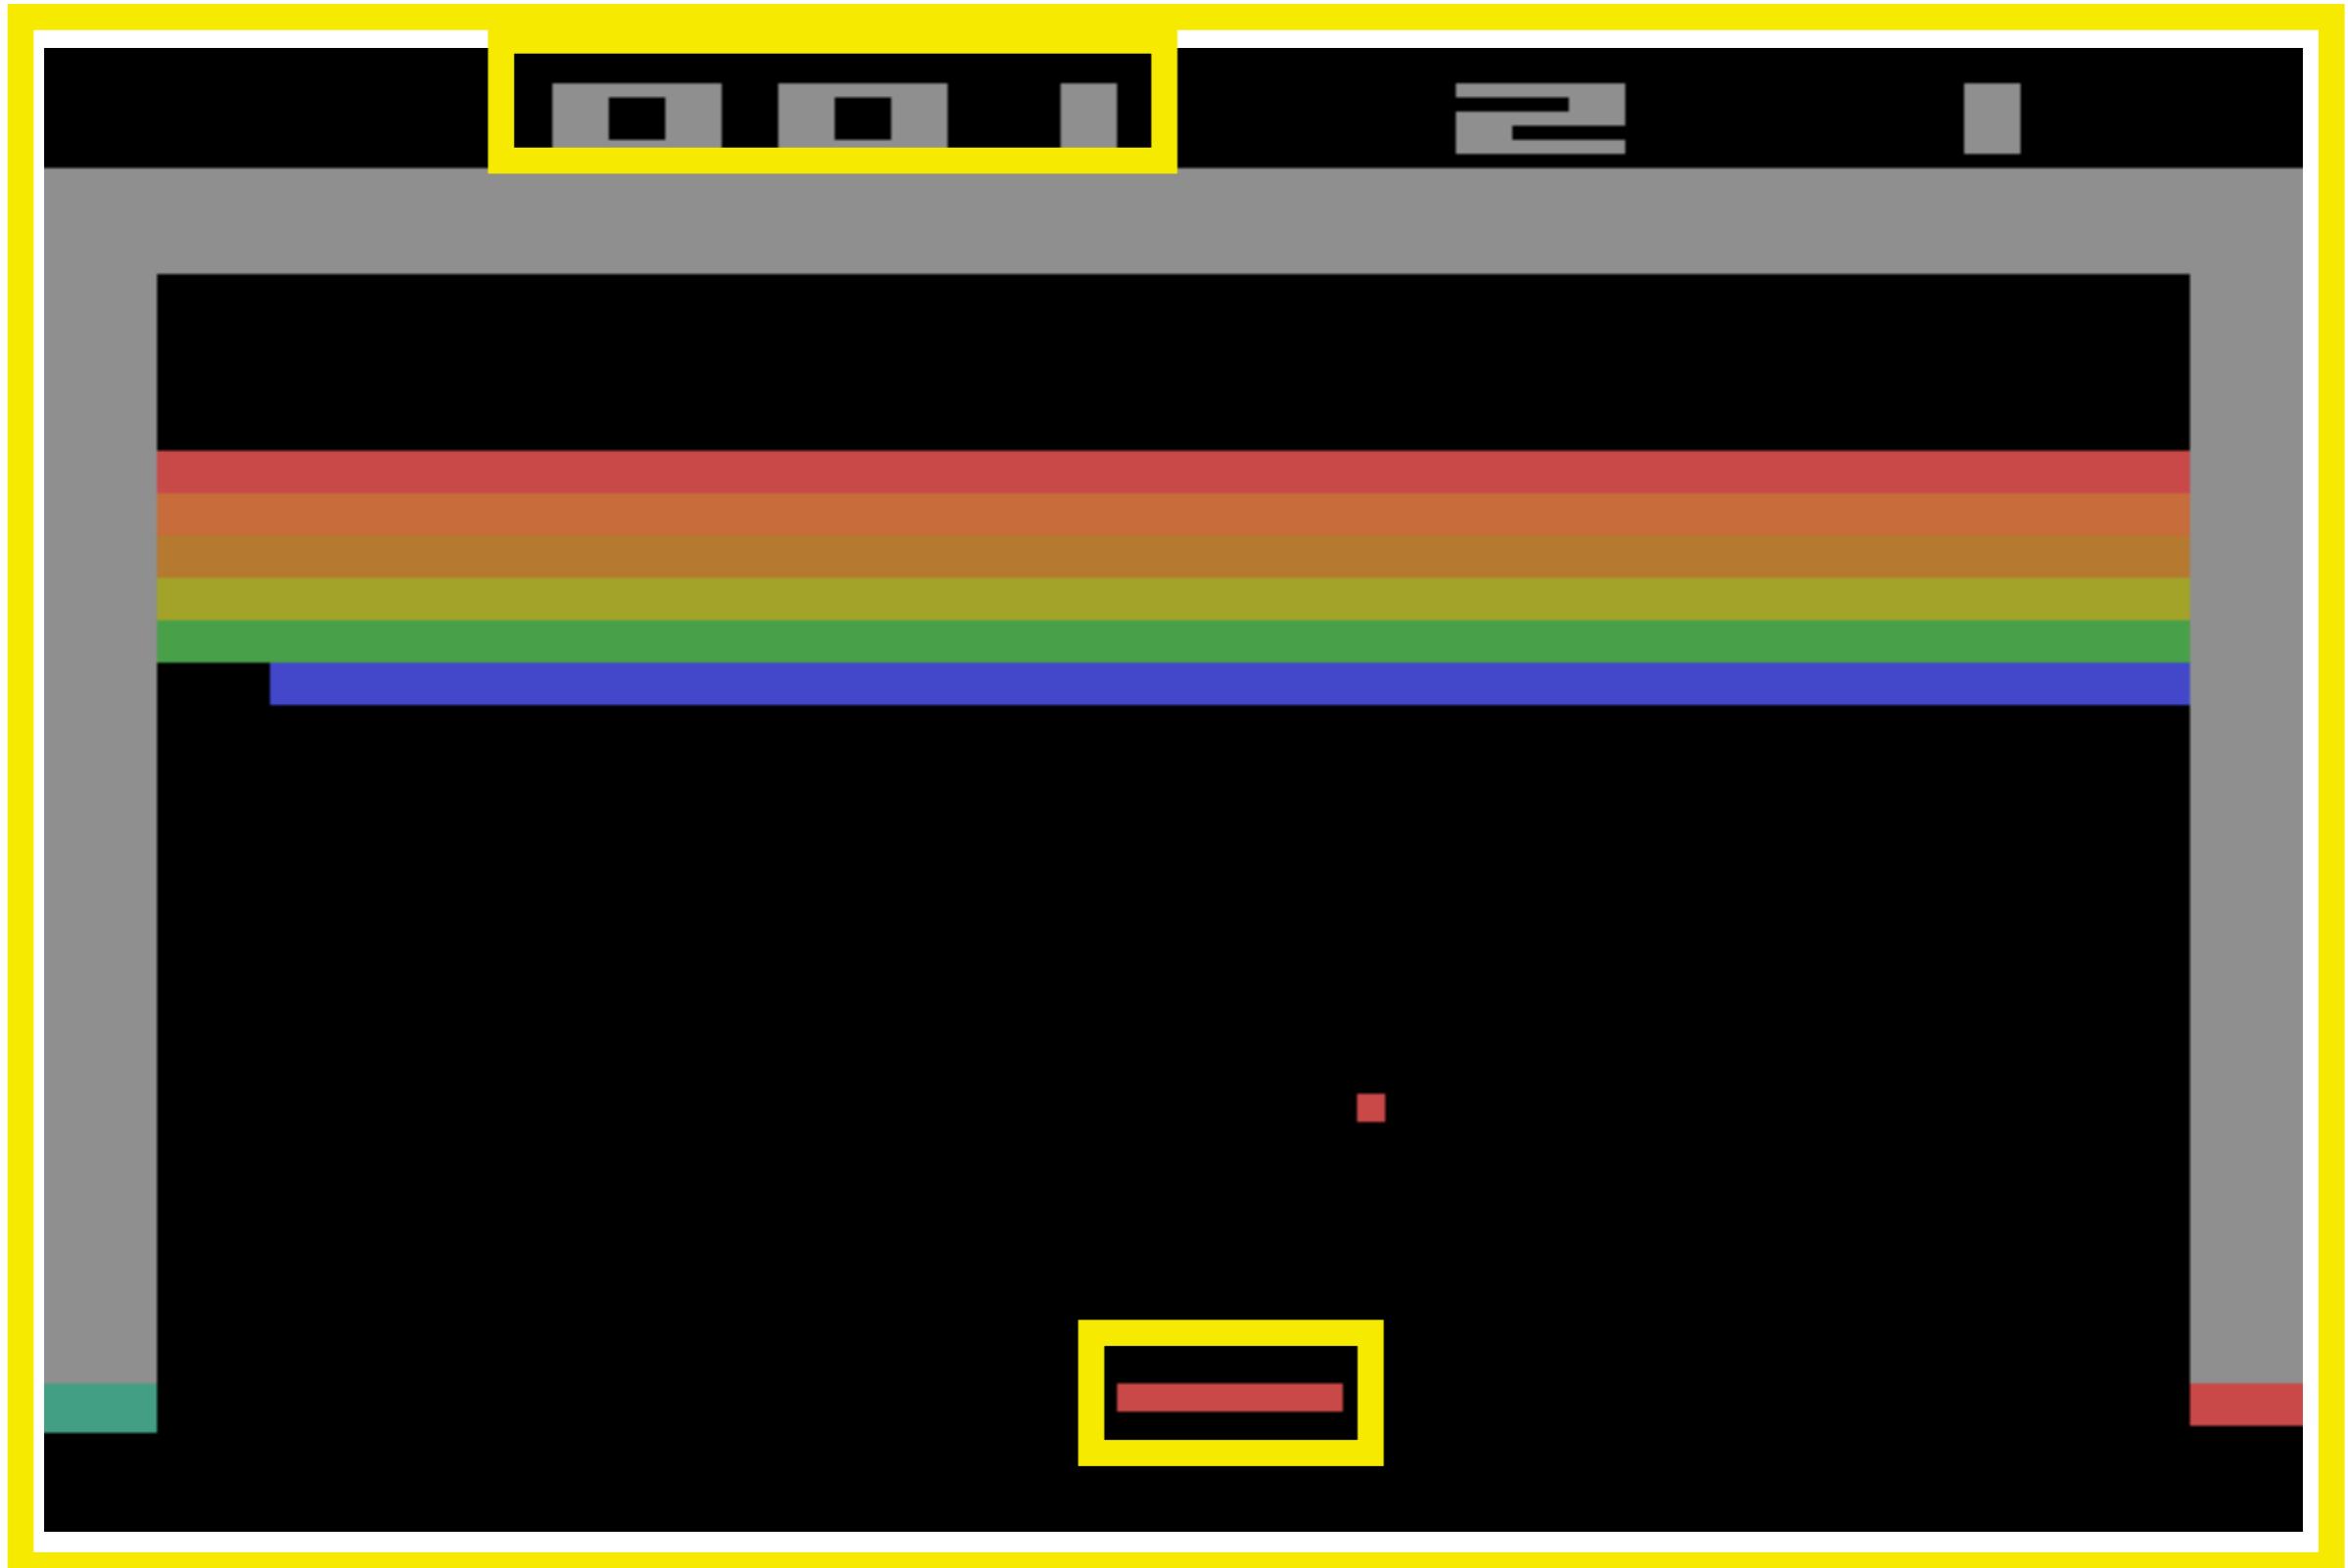
Space Invaders



Enduro

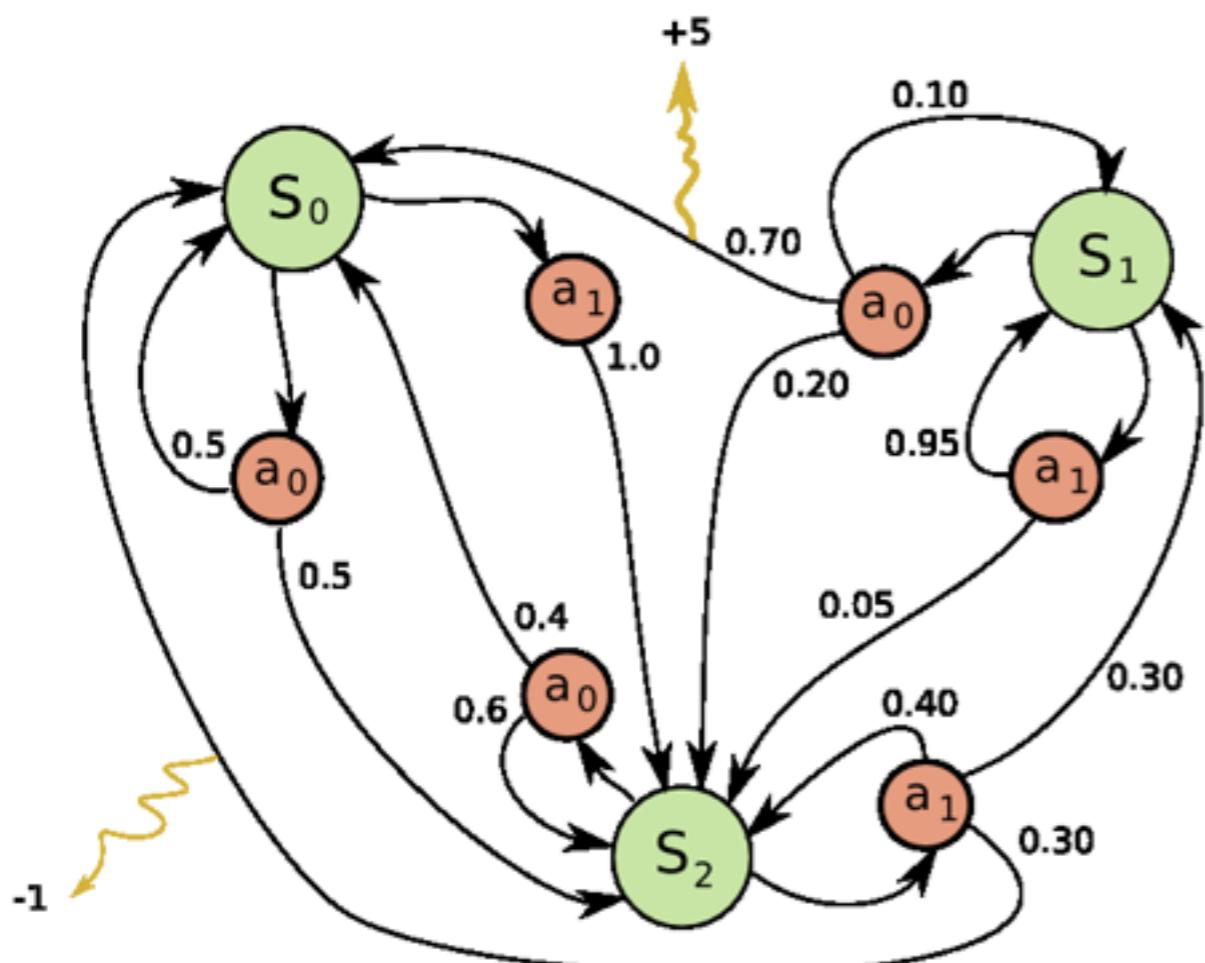
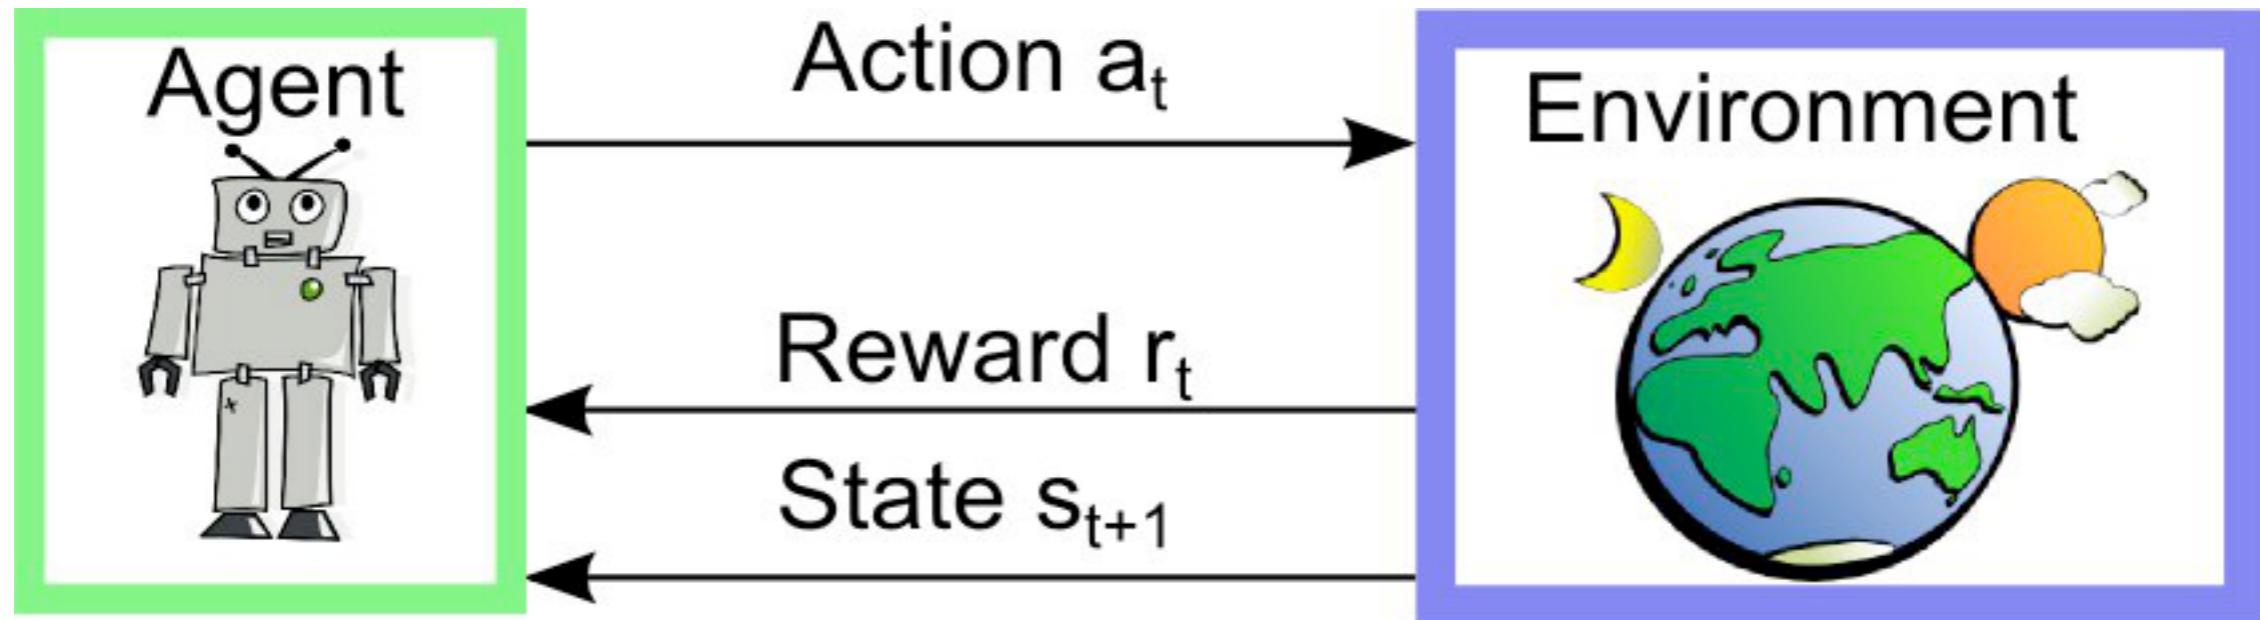
...and 43 more games

# Reward



State

Action



Markov decision process

$s, a, r, s, a, r, \dots$

# Why is hard?



“I must have done something bad, but what is it?”

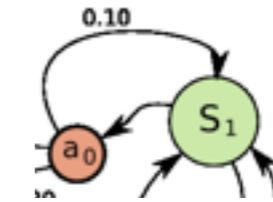
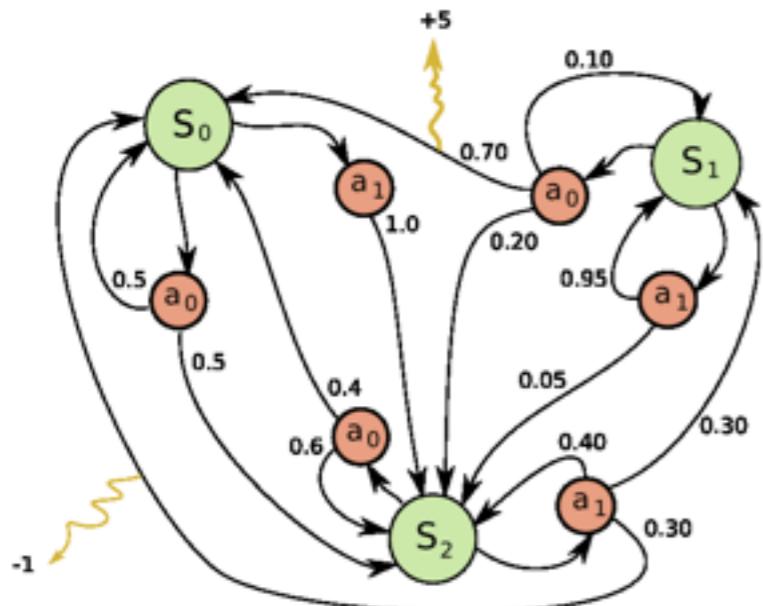
“Shall I continue with this strategy or try others?”

# First, it has to learn to see!



It starts from scratch (pixels) with no high-level concept such as ball, or block, or anything!

# Environment is not known!



s, a, r, s, a, r, ...



84x84 pixels

256 gray

4 frames

...and it is huge!

# Q-learning

- **Total reward** at time step t:

$$R_t = \sum_{i=0}^{\infty} \gamma^i r_{t+i} = r_t + \gamma R_{t+1}$$

- $\gamma$  is the **discount factor**  $0 \leq \gamma < 1$ , that makes future rewards worth less than current
- Define **optimal action value function**  $Q^*(s, a)$  as maximum expected return after taking an action  $a$  in state  $s$ , and then following the most optimal policy

$$Q^*(s, a) = \max_{\pi} E [R | s_t = s, a_t = a, \pi]$$

- The **optimal strategy** in state  $s$  is to select an action  $a$  for which  $Q^*(s, a)$  is maximum

# Q-learning (continued)

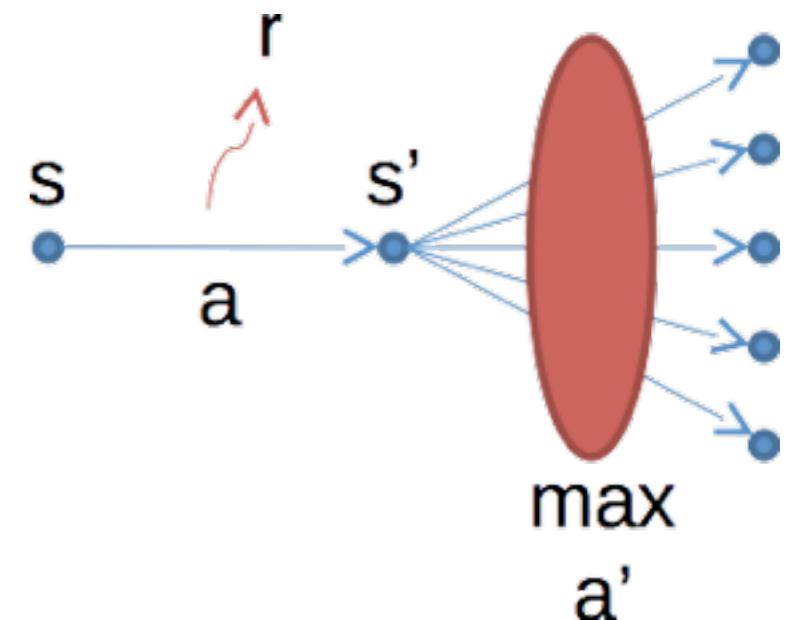
- **Bellman equation:**

$$Q^*(s, a) = E_{s'} [r + \gamma \max_{a'} Q^*(s', a')]$$

- **Iterative approximation**

$$Q_{i+1}(s, a) = Q_i(s, a) + \alpha (r + \gamma \max_{a'} Q_i(s', a') - Q_i(s, a))$$

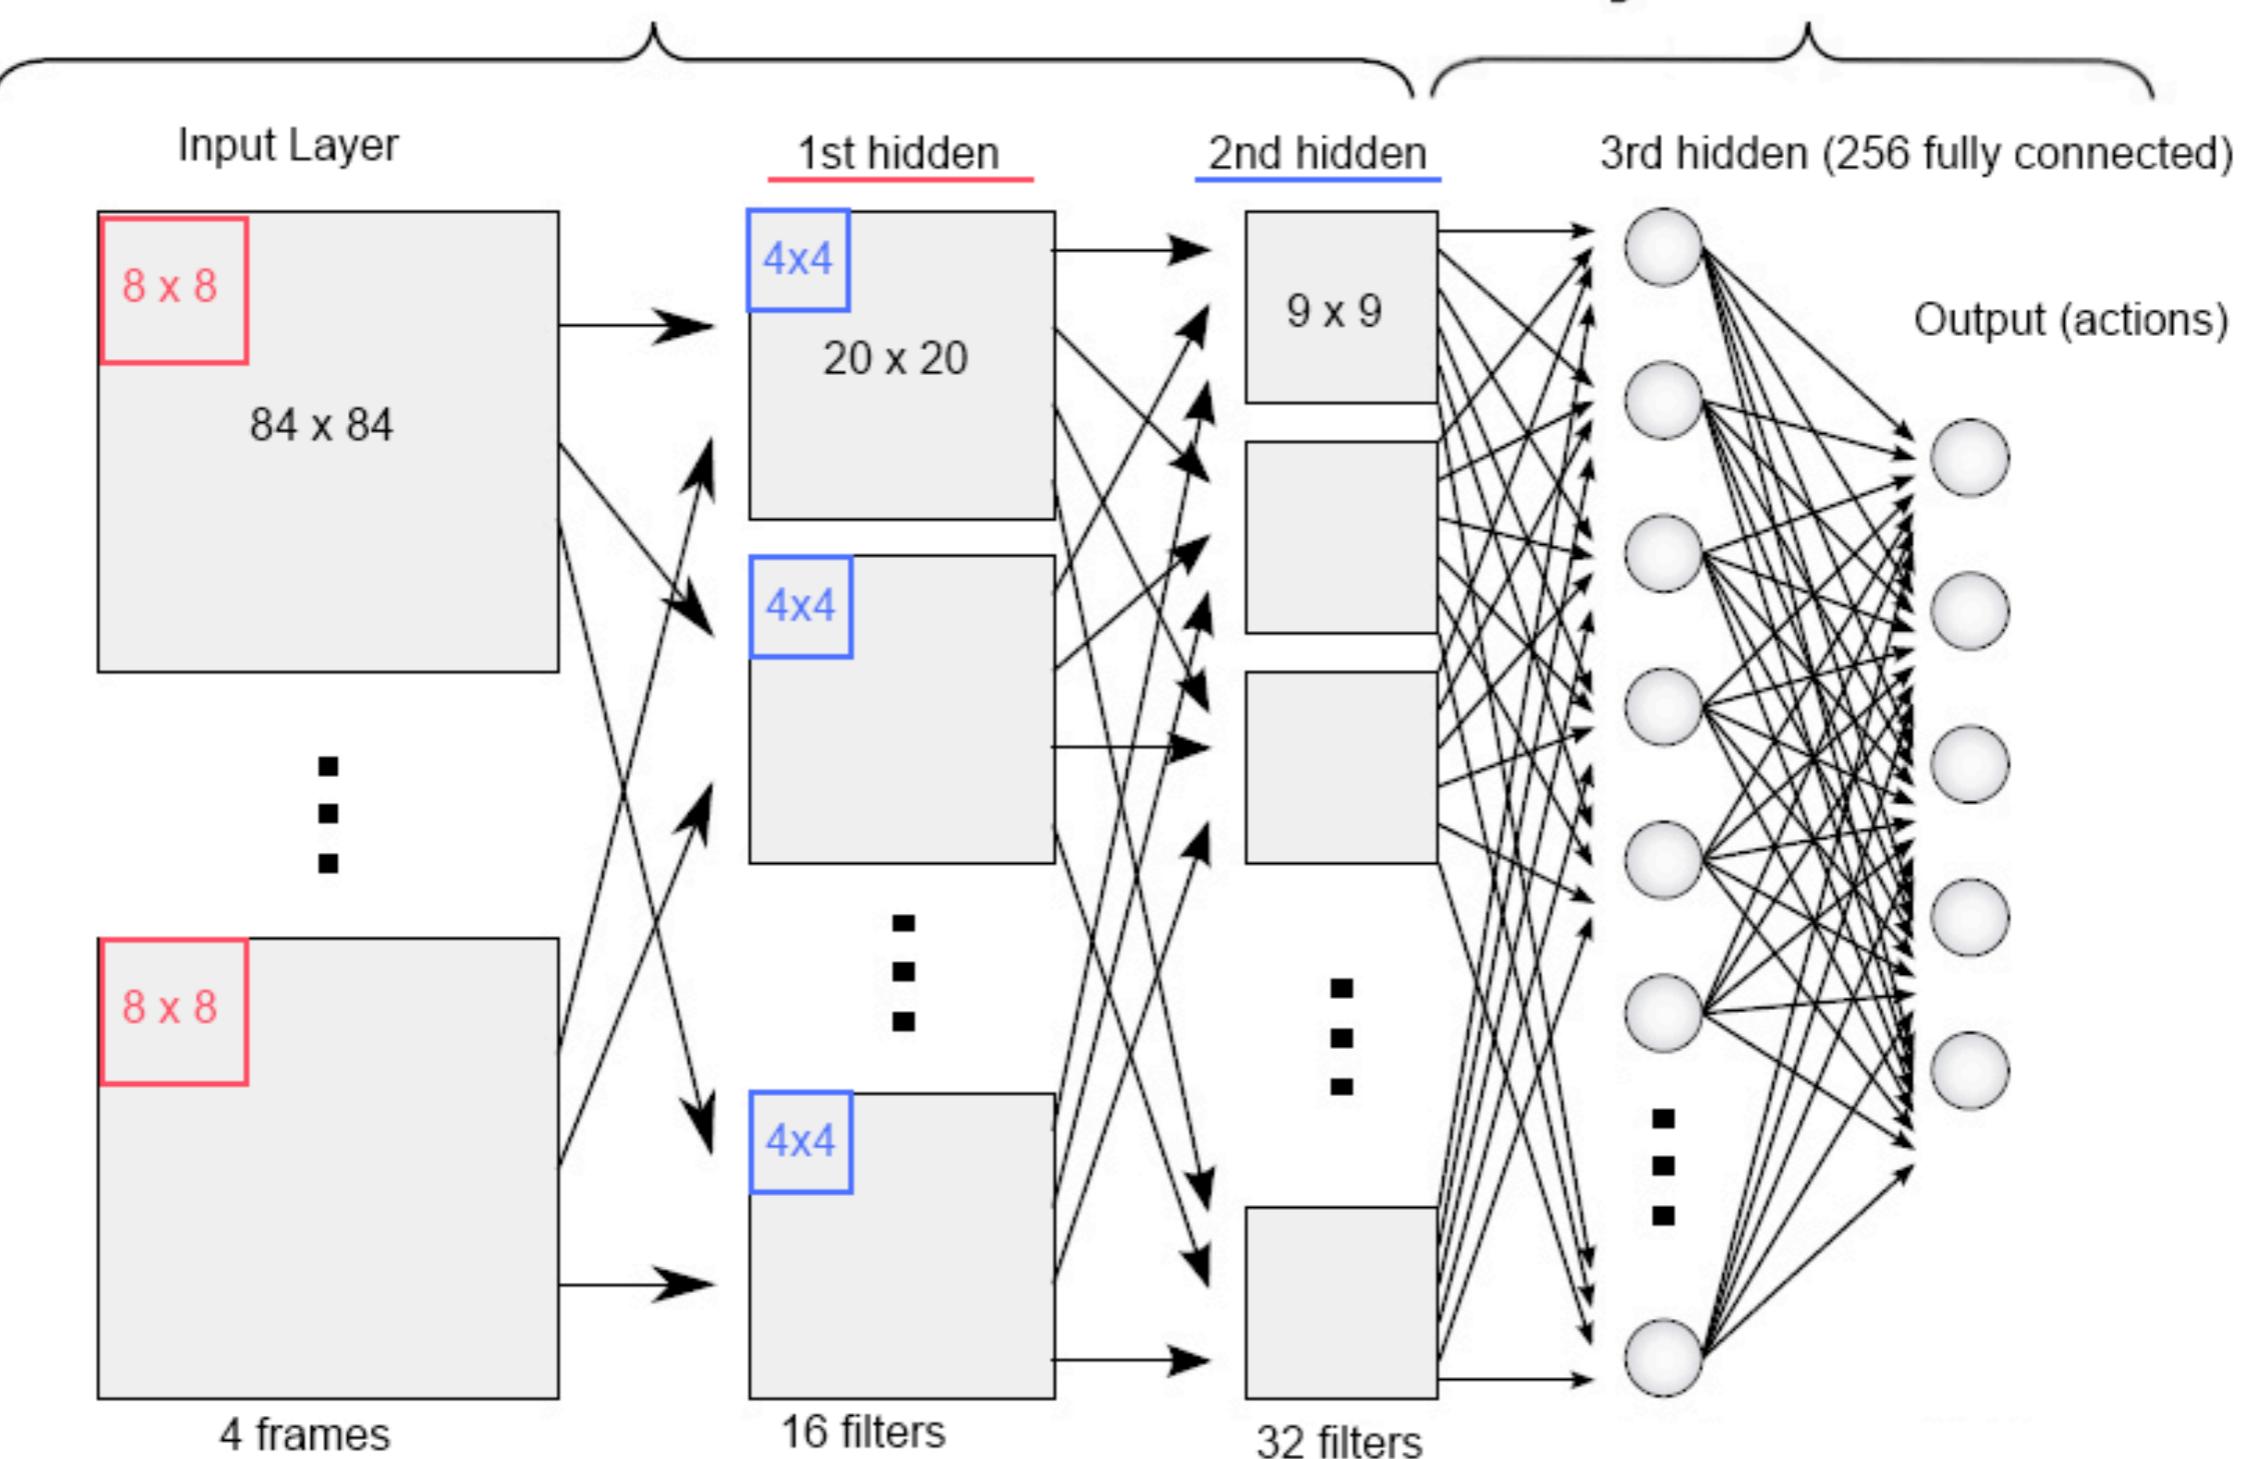
$$\lim_{i \rightarrow \infty} Q_i = Q^*$$



- If state and action spaces are finite, Q-function can be implemented as a table. But for large state spaces this table might be huge
- **Use convolutional neural networks to approximate the Q-function!**

# Convolution

# Fully connected



Nodes:  $84 \times 84 \times 4$

$20 \times 20 \times 16$

$9 \times 9 \times 32$

256

4

Weights:

$8 \times 8 \times 4 \times 16$

$4 \times 4 \times 16 \times 32$

$9 \times 9 \times 32 \times 256$

$256 \times 4$

# Learning algorithm

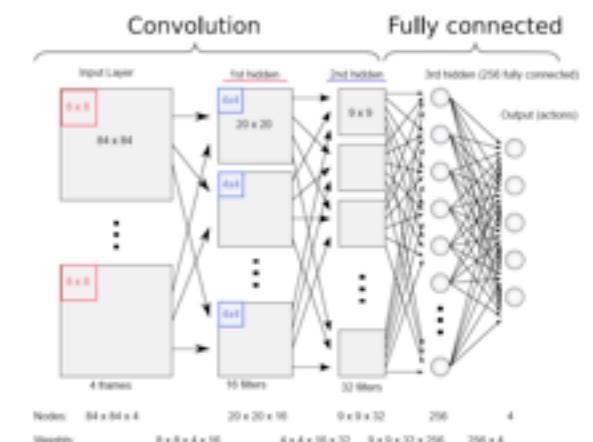
- Play a lot of games
  - Play a lot of frames
    - Get a frame and **pre-process it**

# Pre-processing

- Image
  - 1.  $210 \times 160$  RGB
  - 2.  $110 \times 84$  gray-scale
  - 3.  $84 \times 84$
- Every 4th frame from the emulator
- State contains 4 frames
- Reward is binarized

# Learning algorithm

- Play a lot of games
- Play a lot of frames
  - Get a frame and pre-process it
  - Choose action using **Deep Neural Net (DNN)** and current state
    - Do it randomly with probability  $\epsilon$



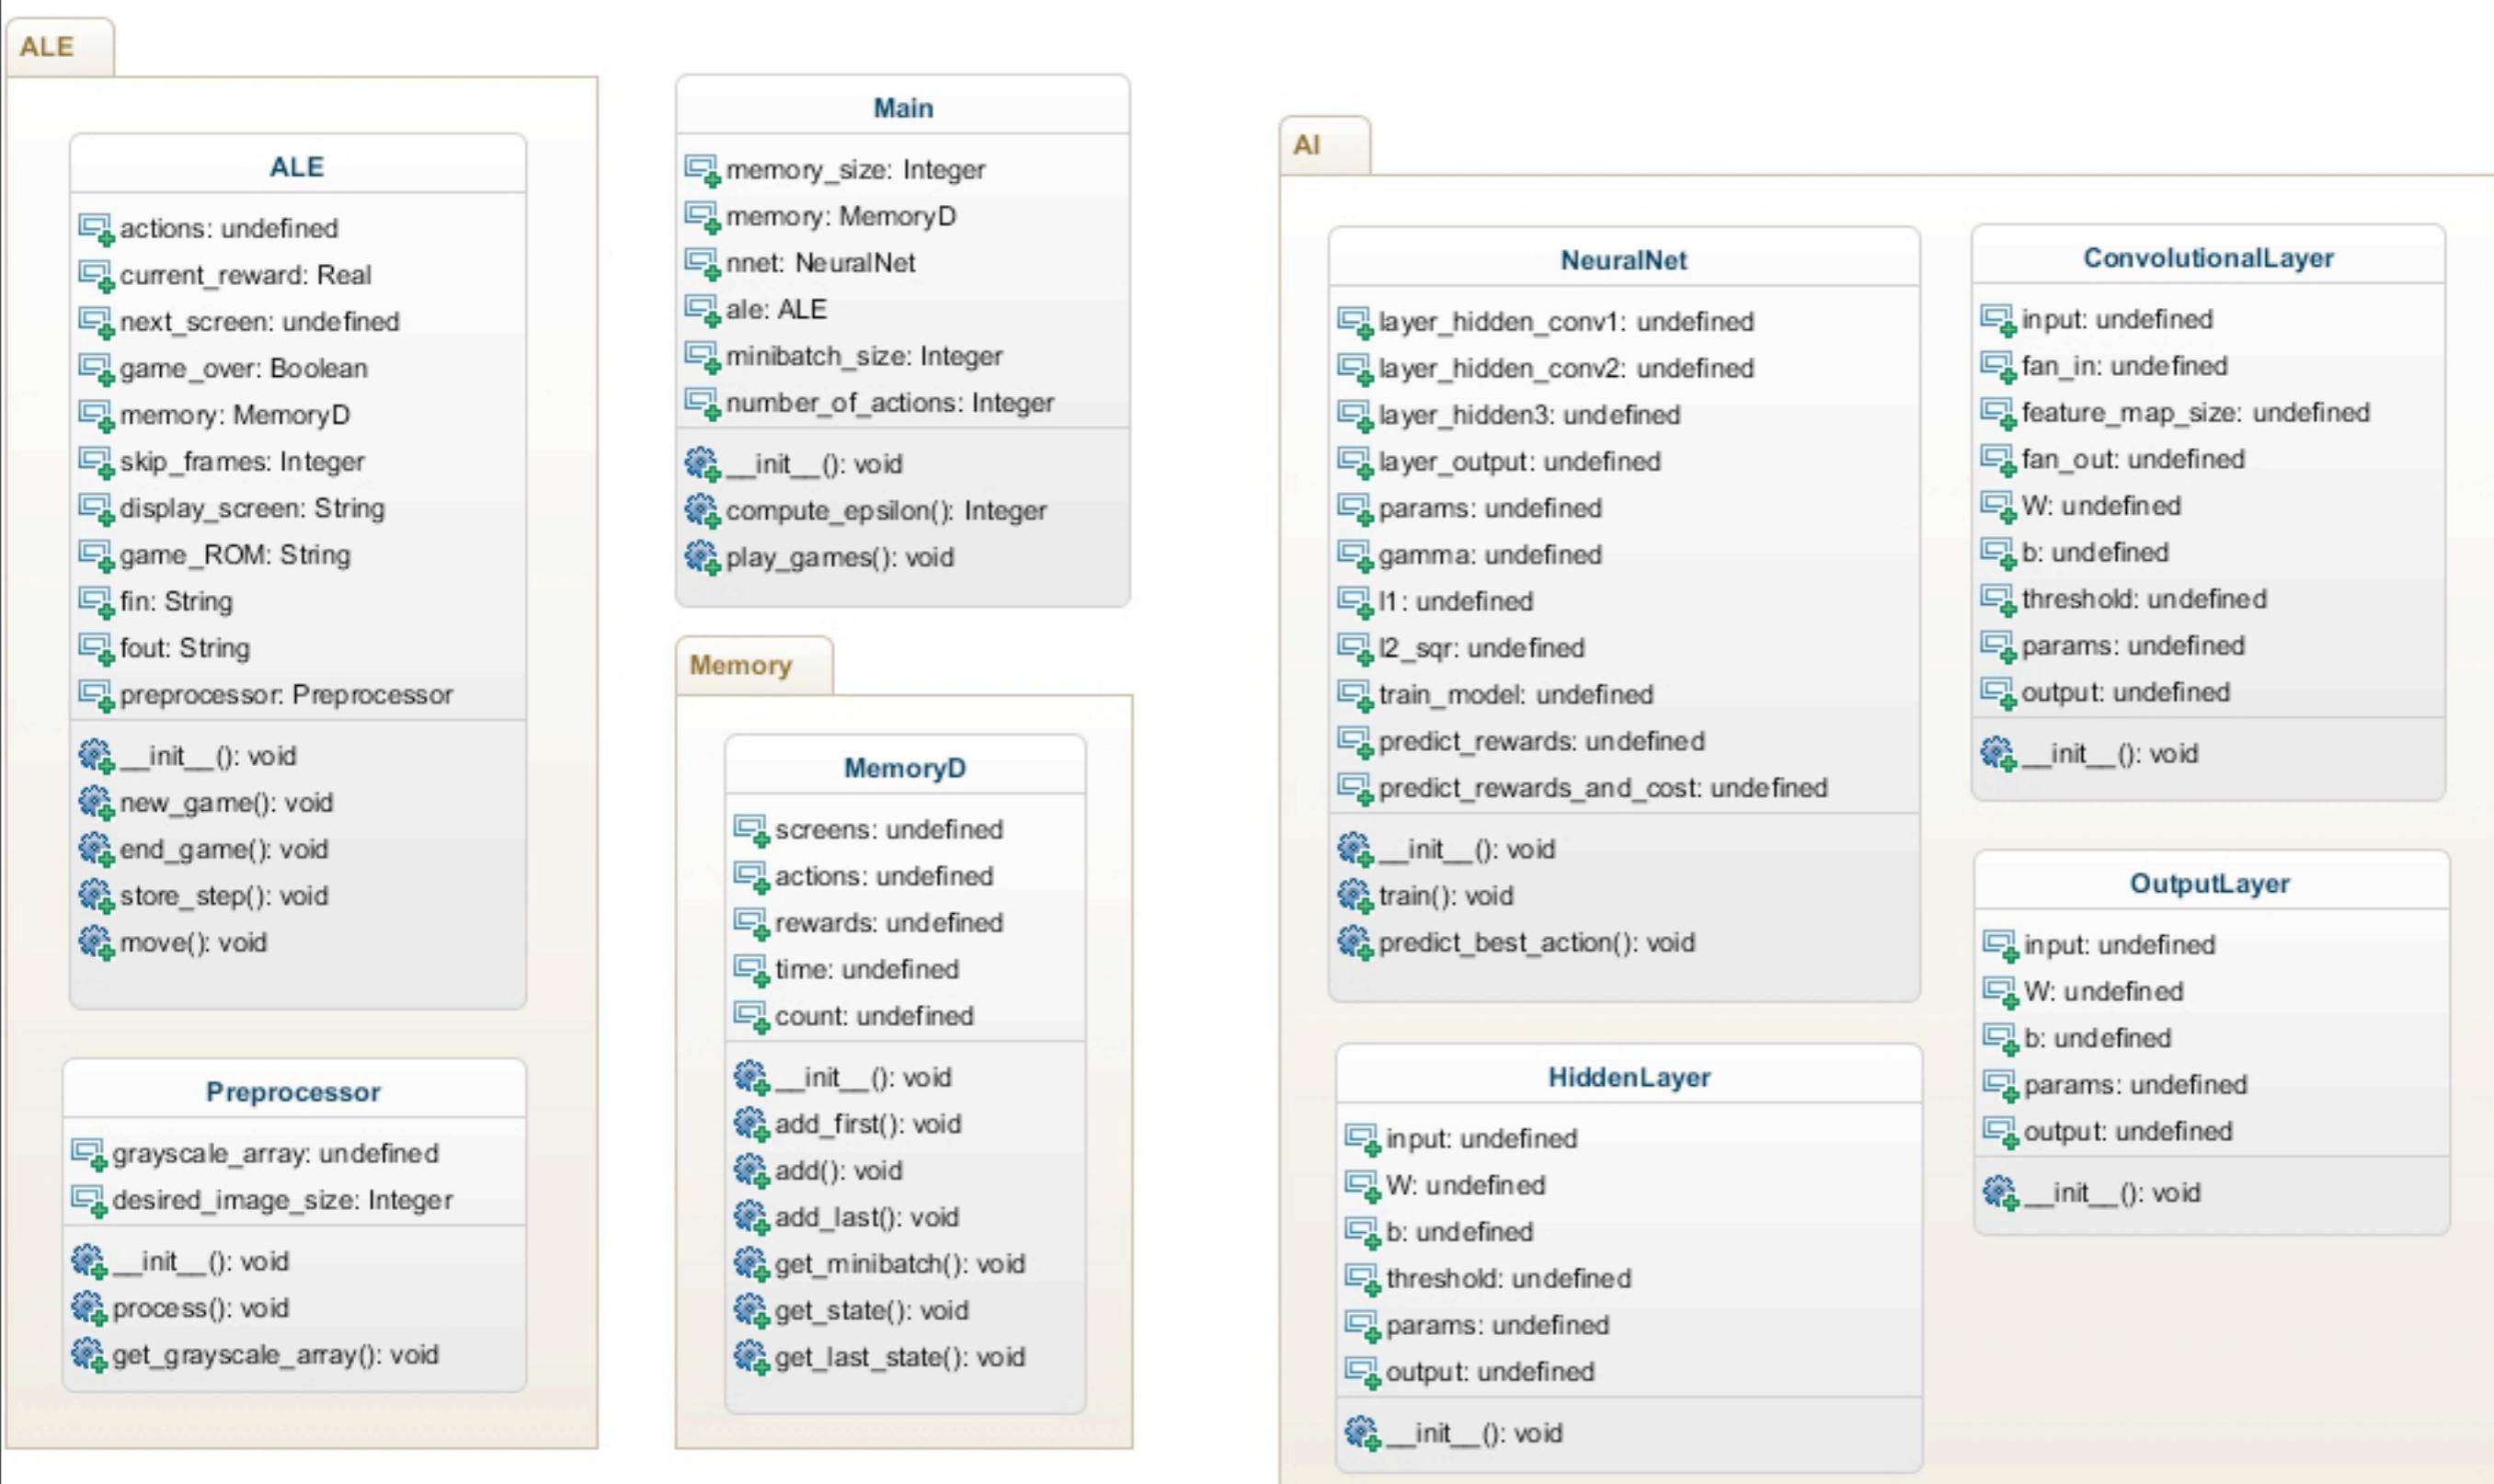
# Learning algorithm

- Play a lot of games
  - Play a lot of frames
    - Get a frame and pre-process it
    - Choose action using Deep Neural Net (DNN) and current state
      - Do it randomly with probability  $\epsilon$
      - Observe reward and next image and save everything to the memory
    - Train the network
      - 32 random transitions:  $32 \times \{\text{state}, \text{action}, \text{reward}, \text{state}\}$
      - Q-learning
$$Q(s_t, a_t) = r_t + \gamma \max_a DNN(state_t + 1)$$
      - Gradient descent

# In a nutshell...

1. At each step choose action with **maximum total reward**
2. Use **convolutional neural network** to estimate this total reward
3. Use **Q-learning rule** to iteratively train this network to estimate better
4. Use **experience replay memory** to stabilize the learning

# In real life...



<https://github.com/kristjankorjus/Replicating-DeepMind/>

# Let's play

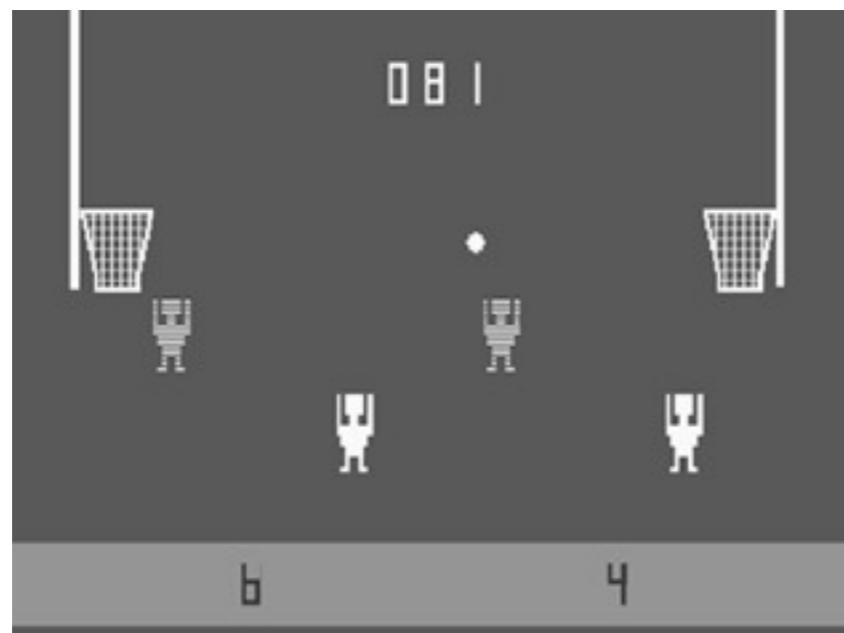
# Why is it important?

- Complex and different tasks
  - challenging even for humans
- Model free
  - same architecture for all games
- Beats all other results so far
  - benchmark for future research
- Easy to understand
  - sparked a lot of interest
- Parallels with Neuroscience

# Ongoing projects

**Collective phenomena:** cooperation/competition in multi-agent learning systems

- Activate 2nd player (or more)
- Engineer the reward



9

0

1

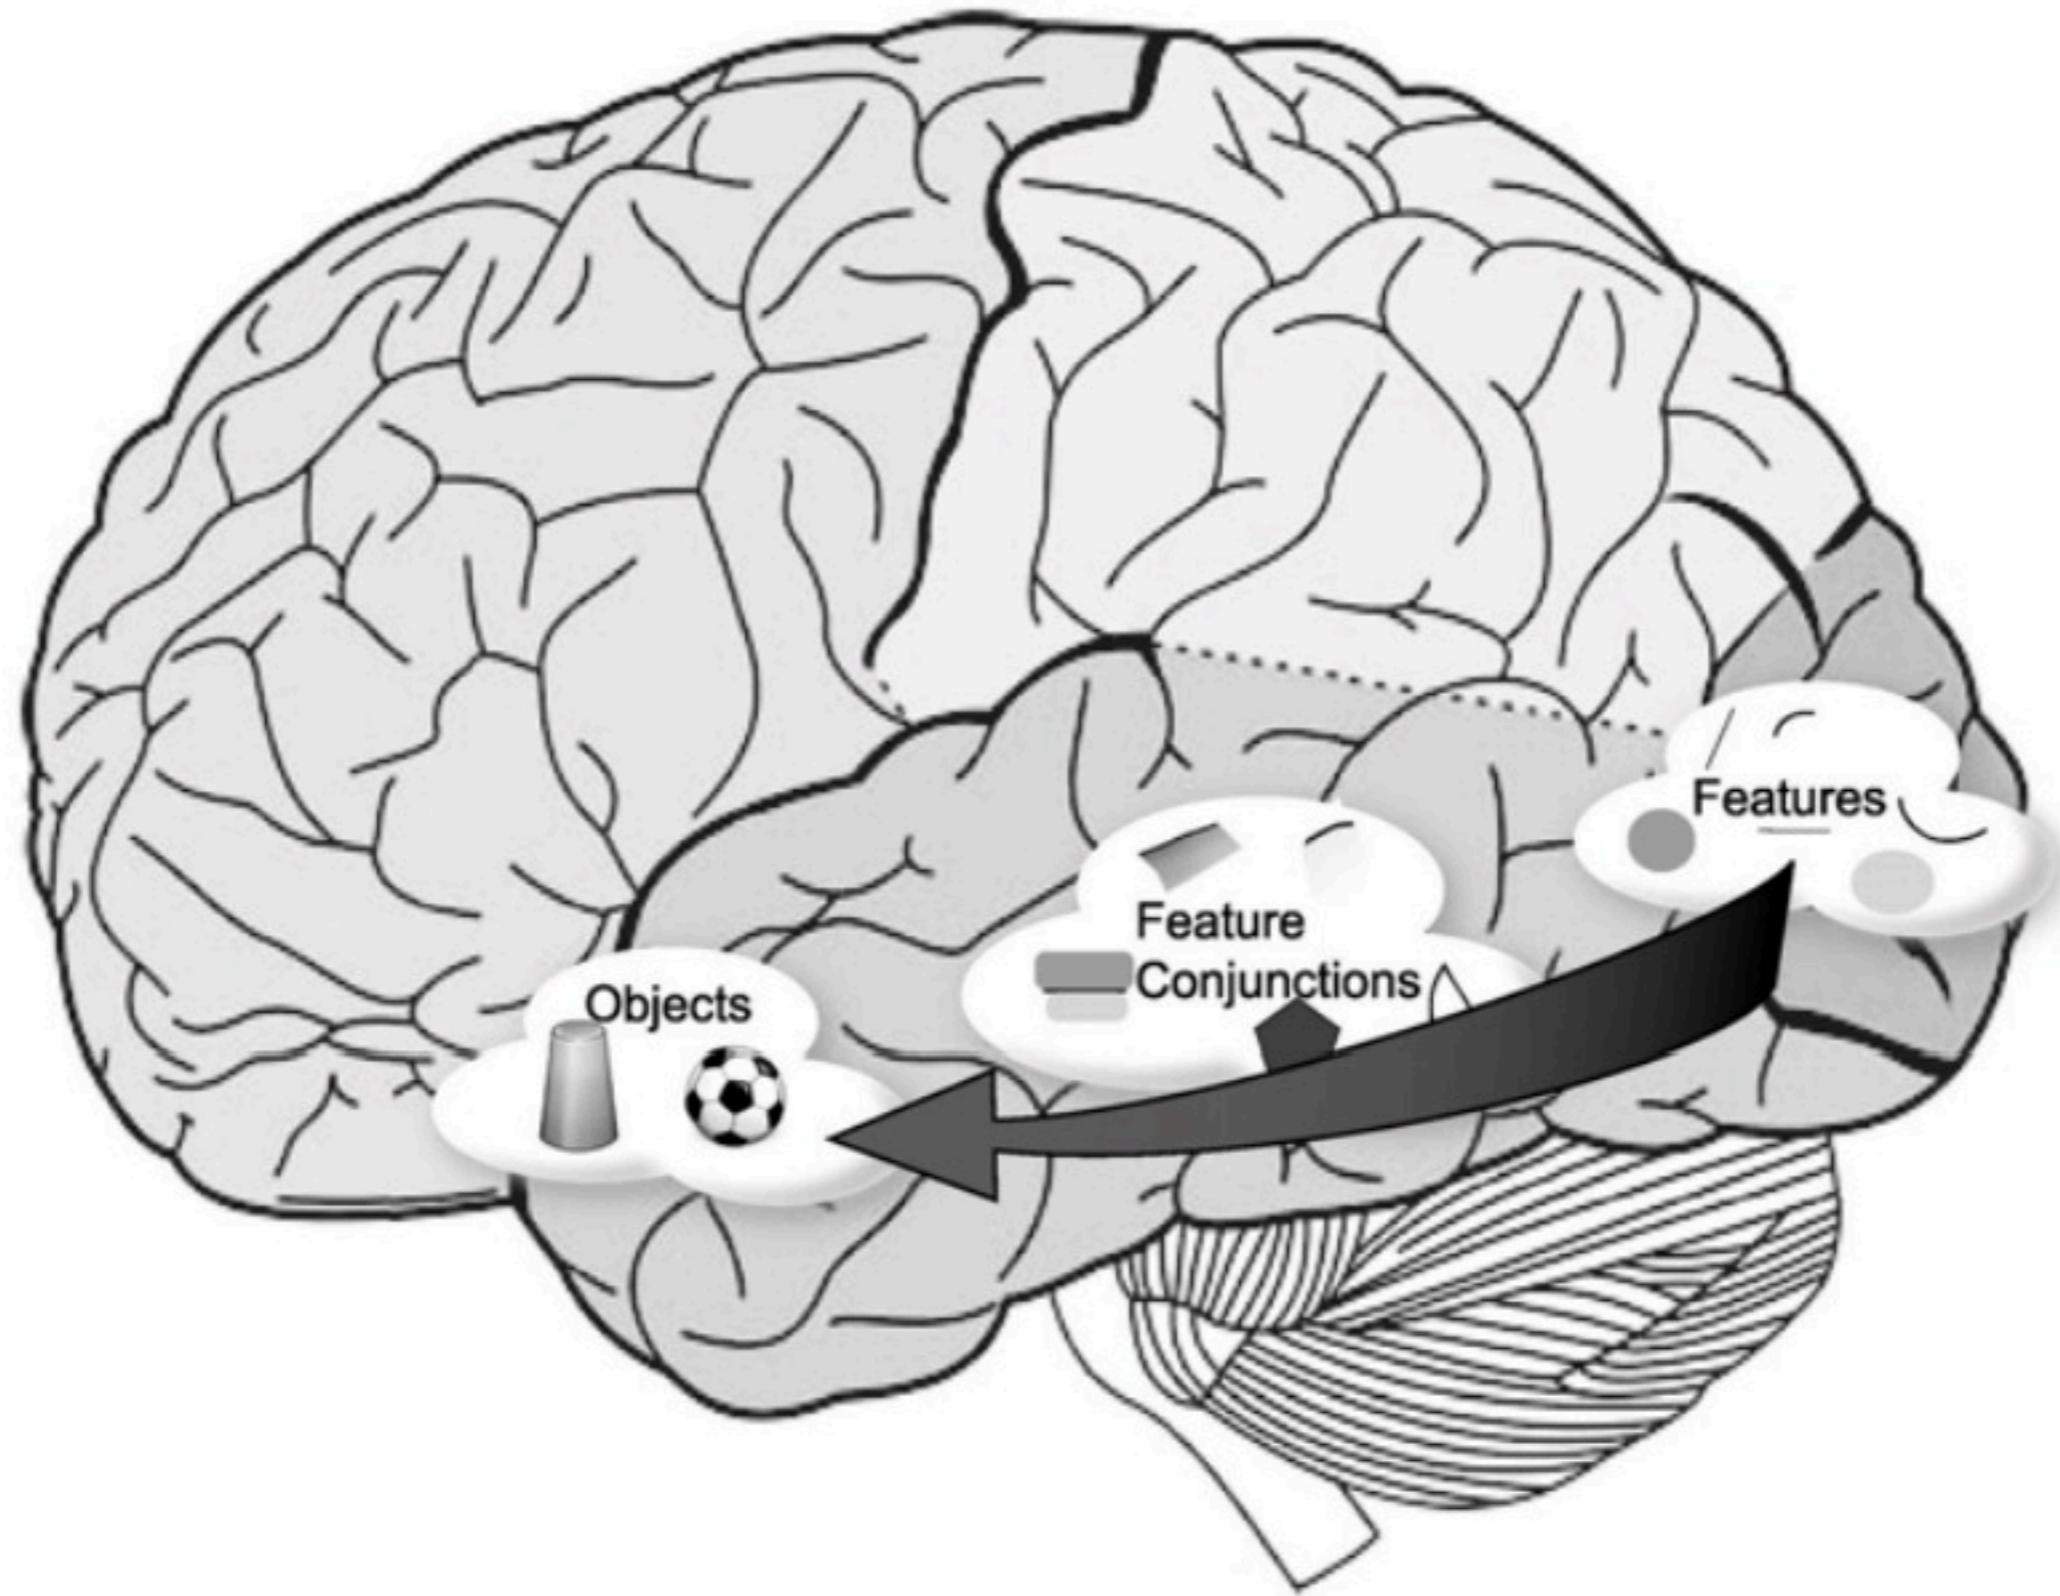
2

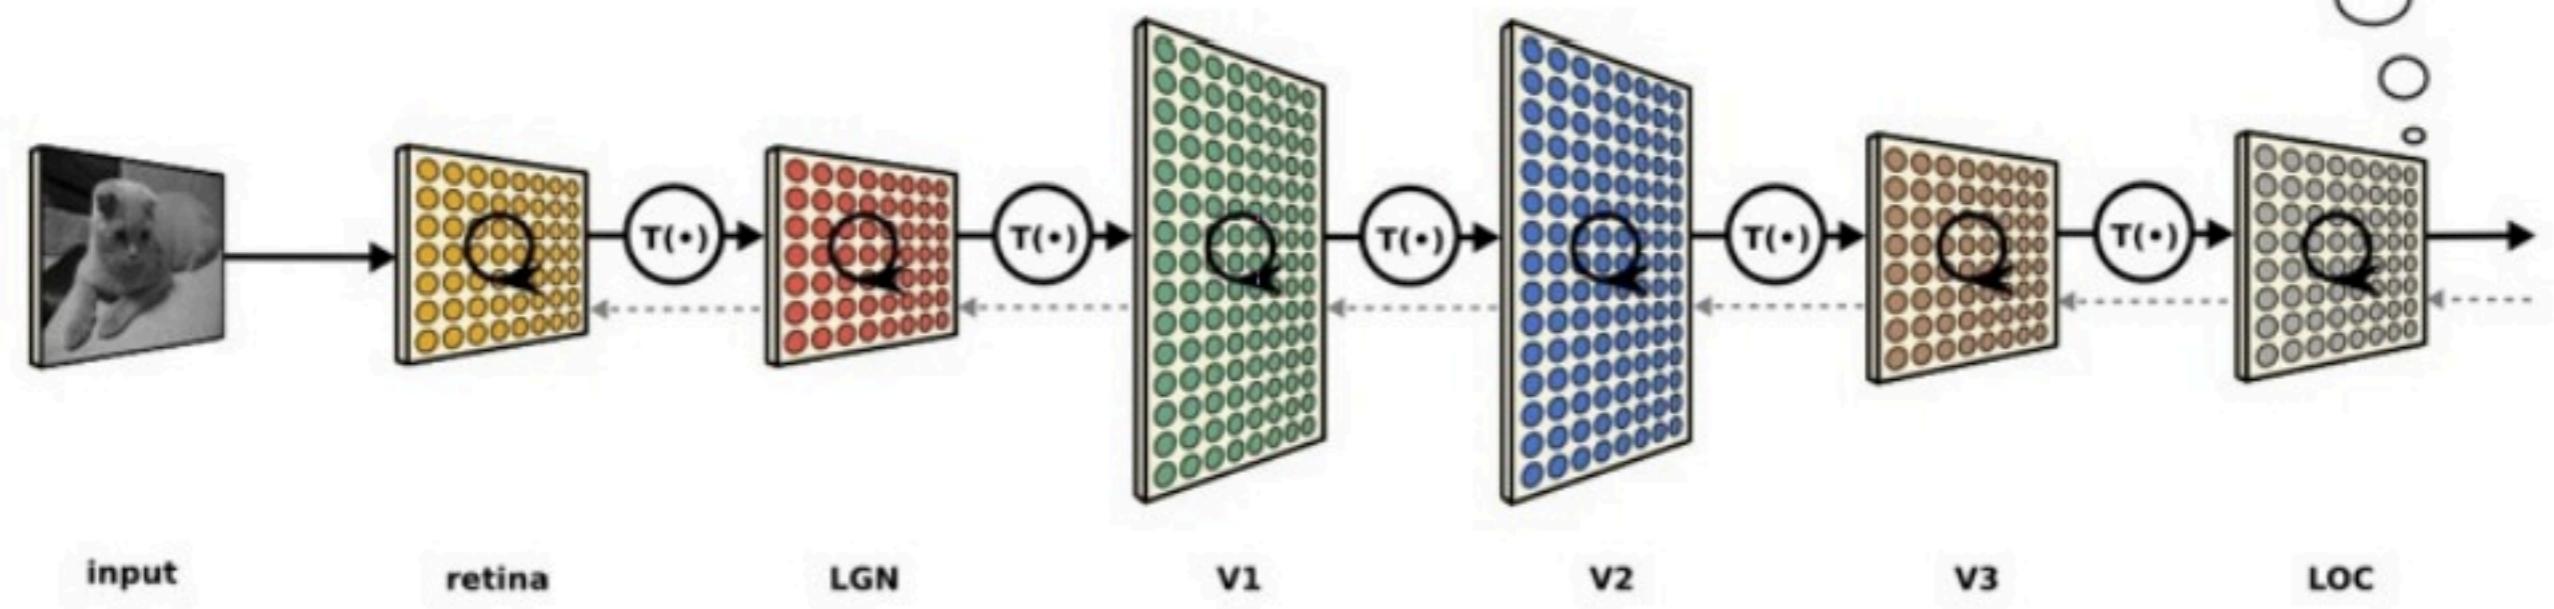
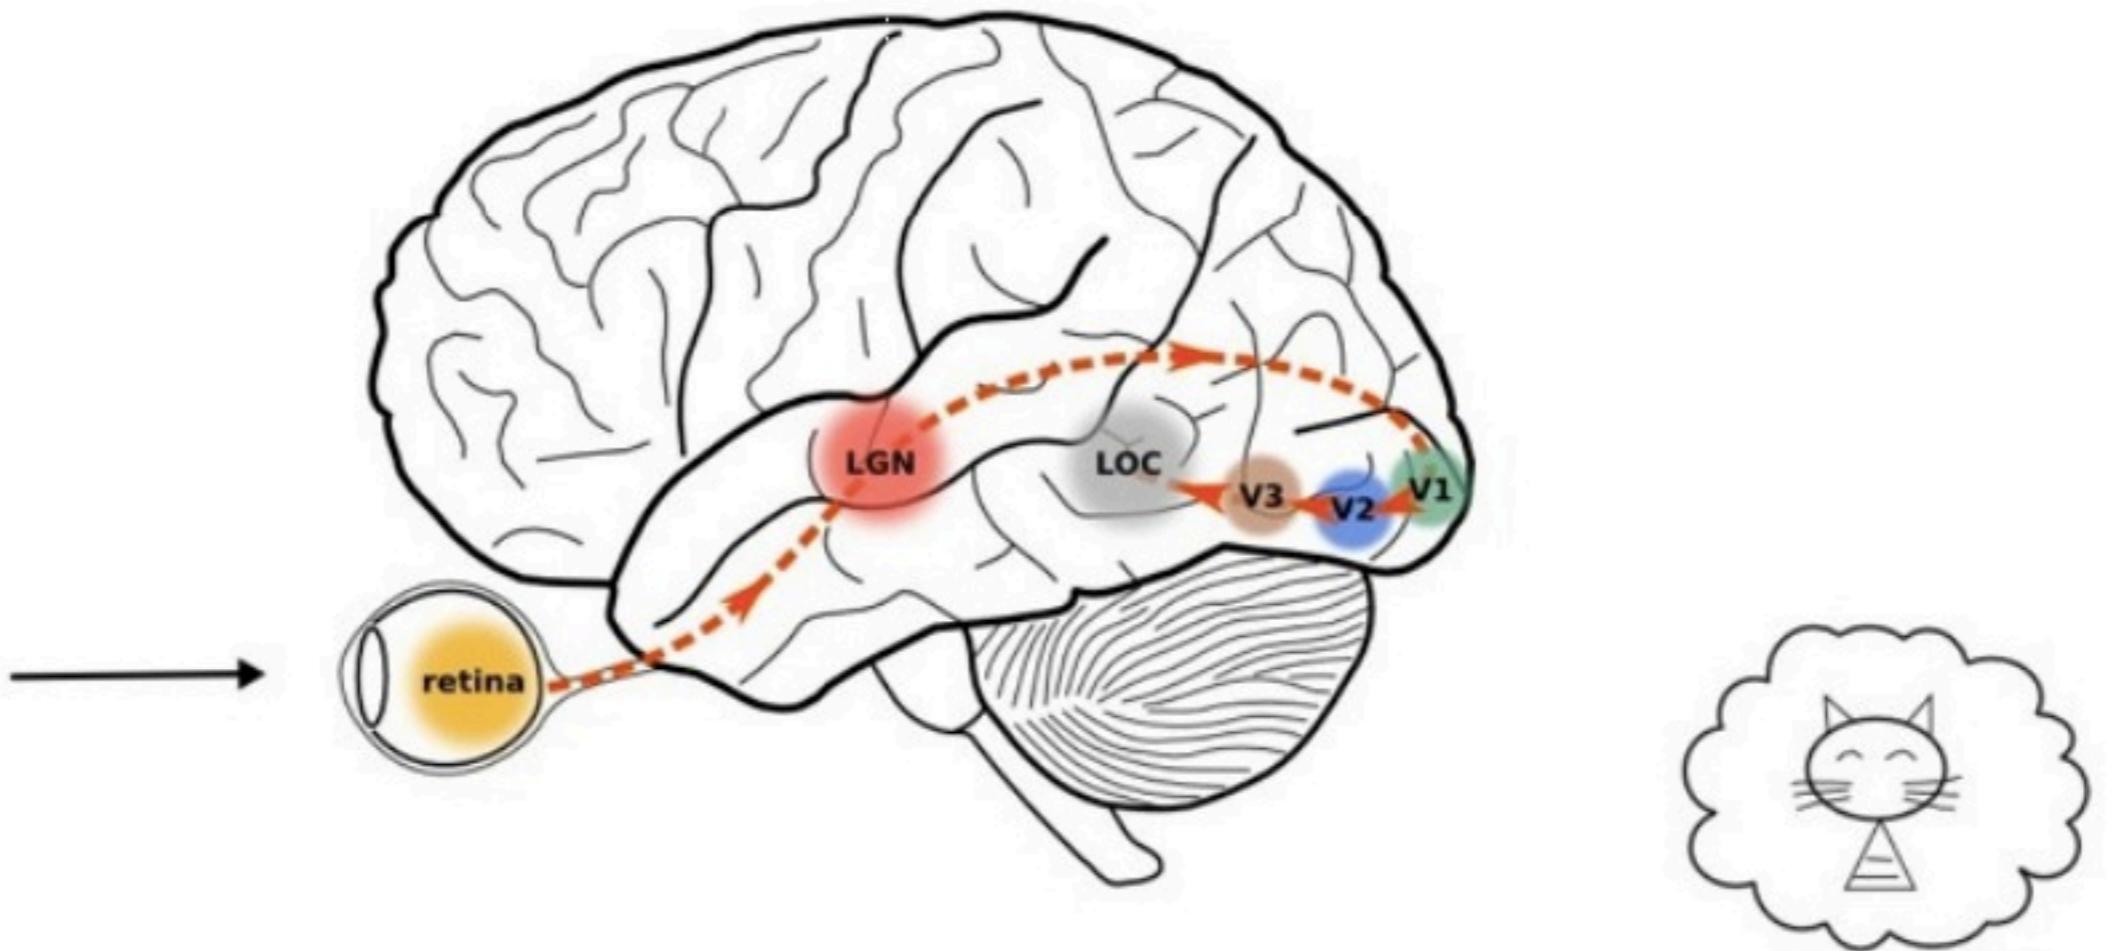
3

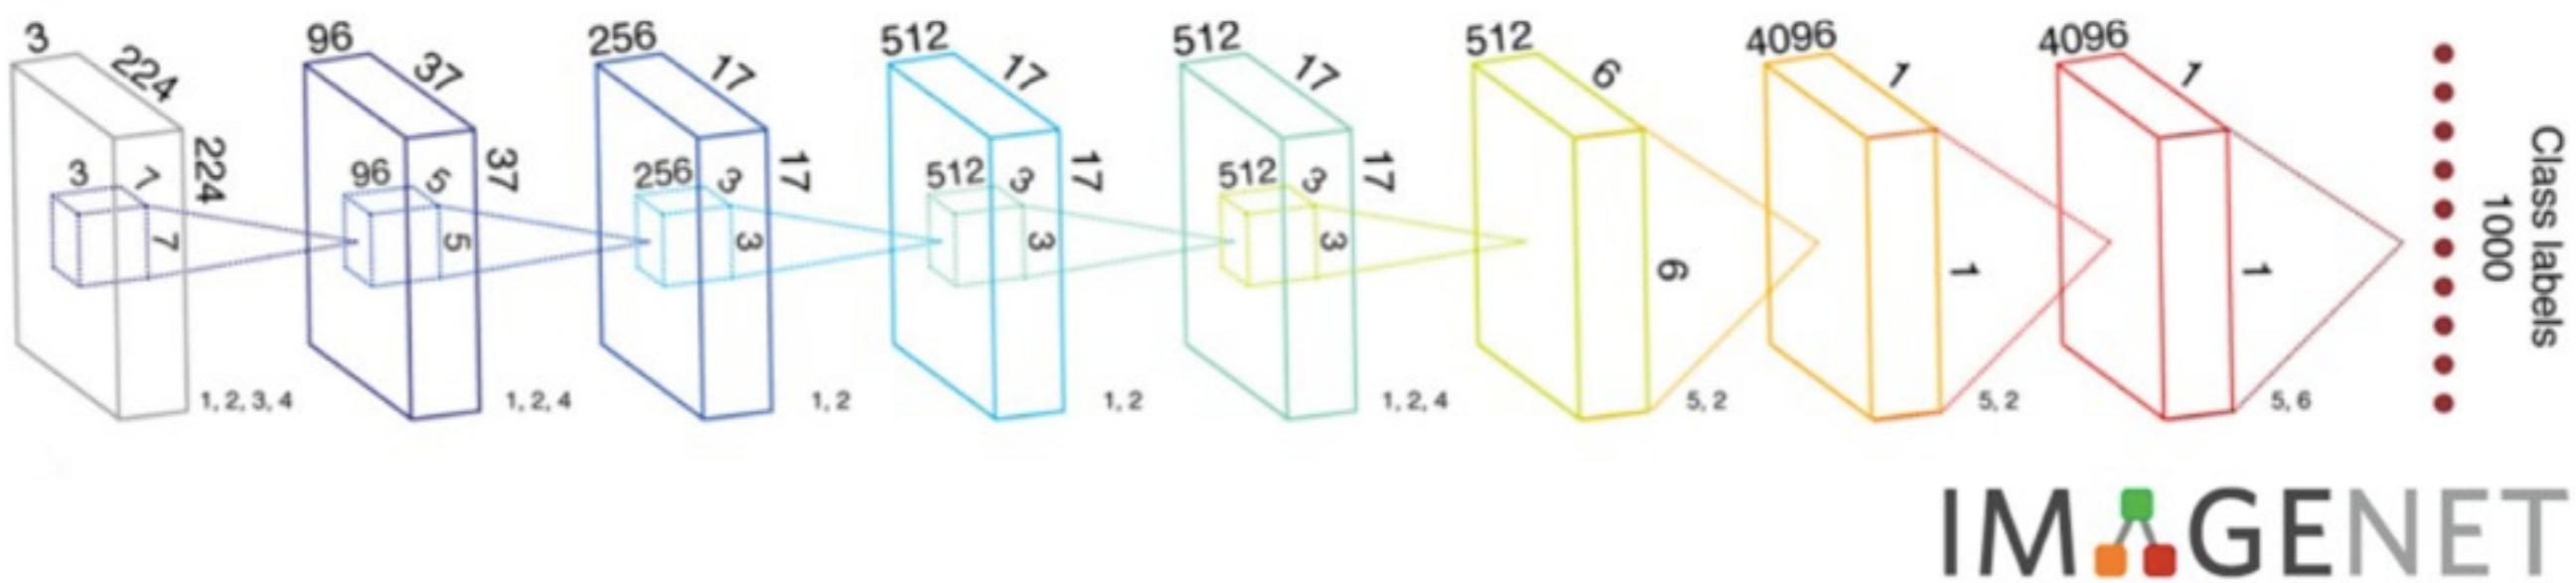


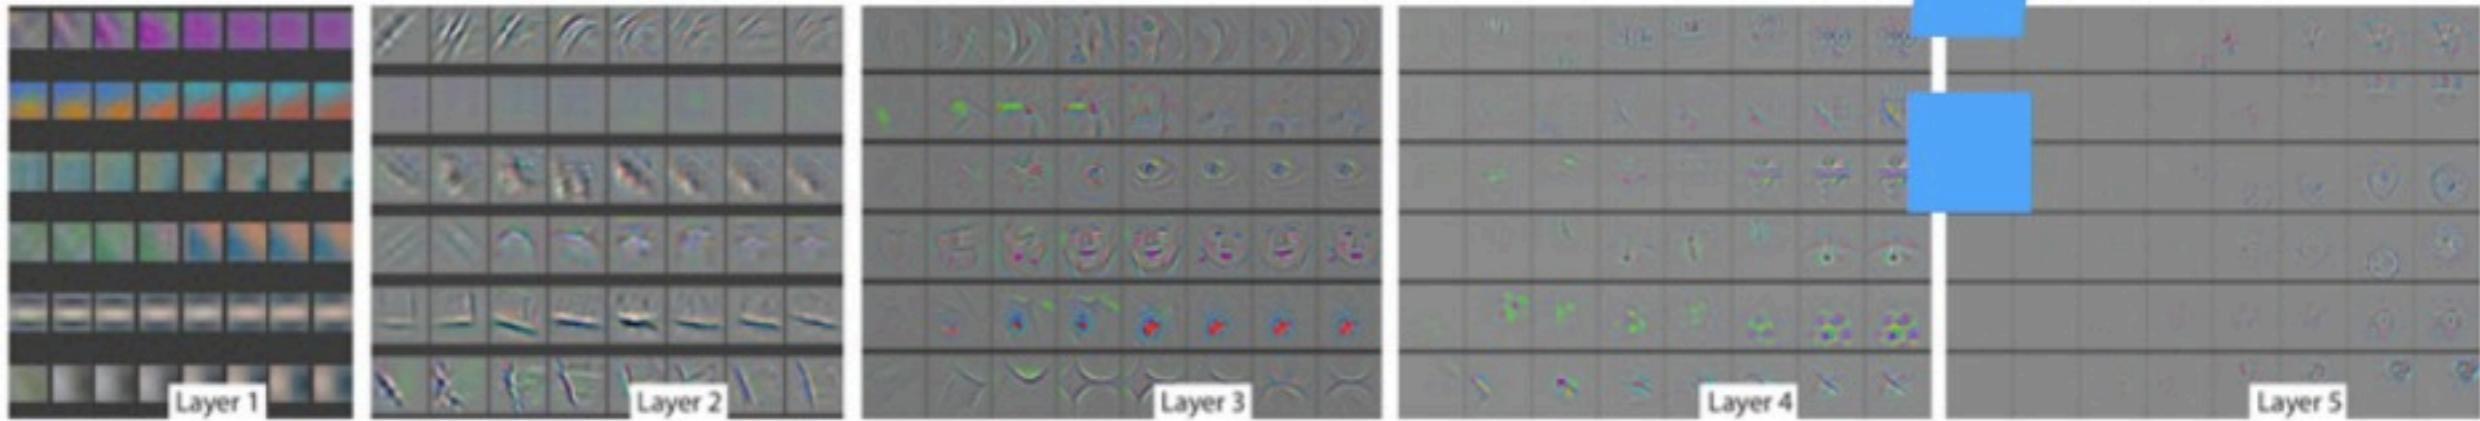
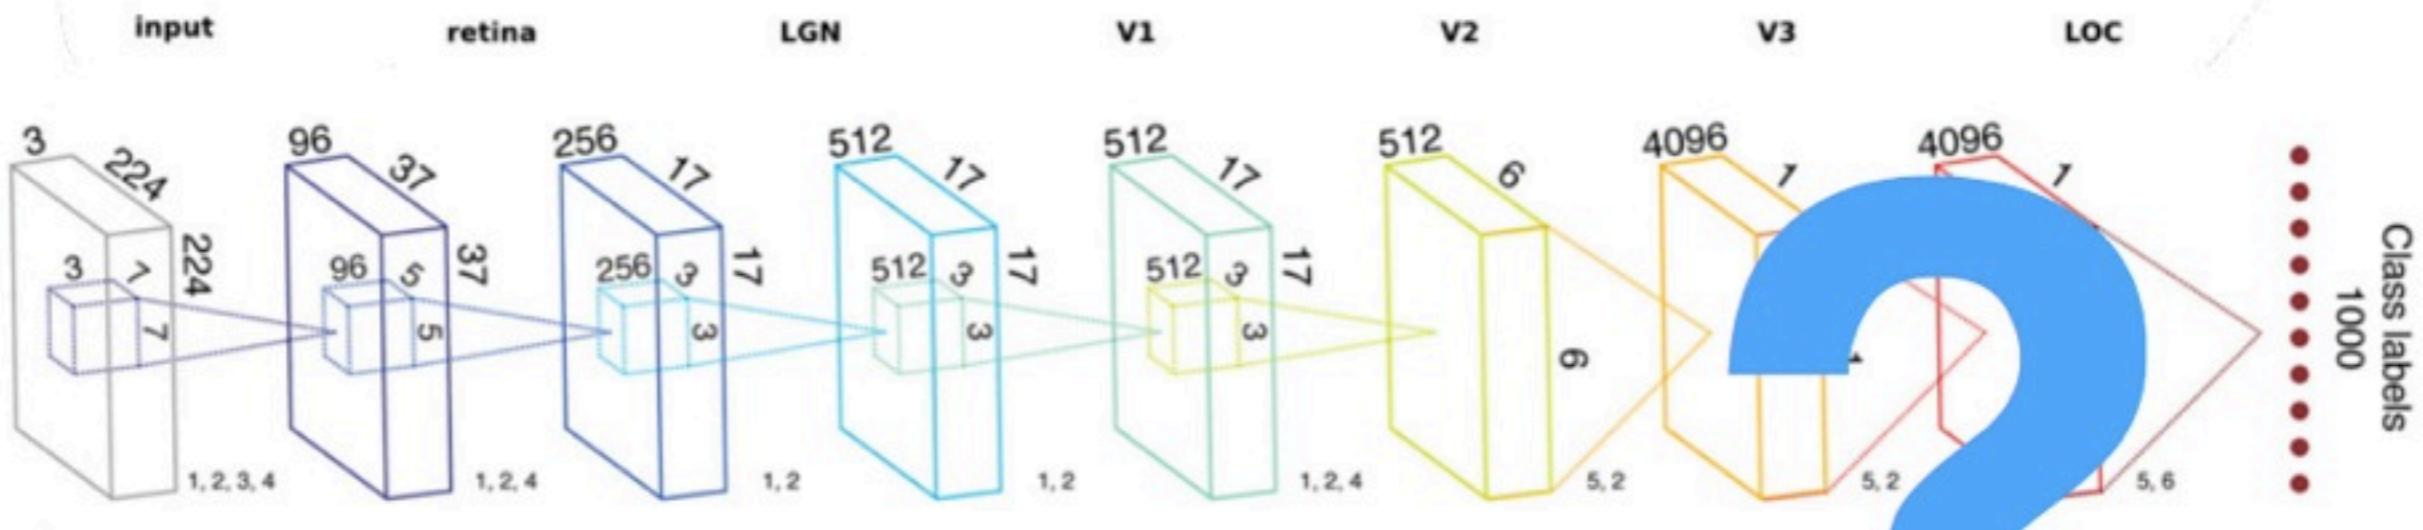
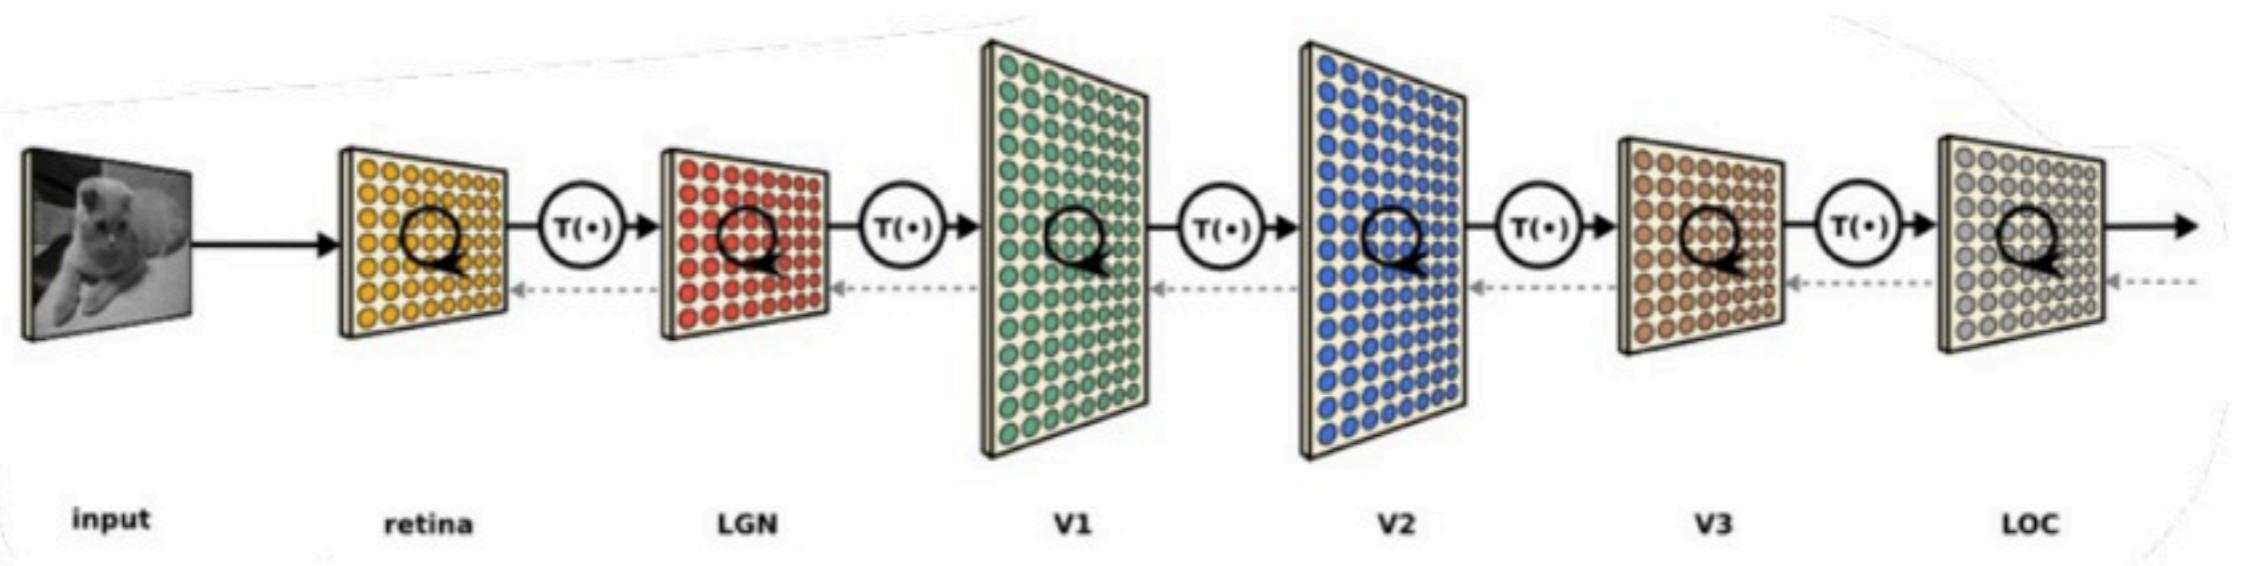
# Ongoing projects

**Deep Networks vs brain:** compare CNN vs brain responses to natural images using intra-craneal recordings

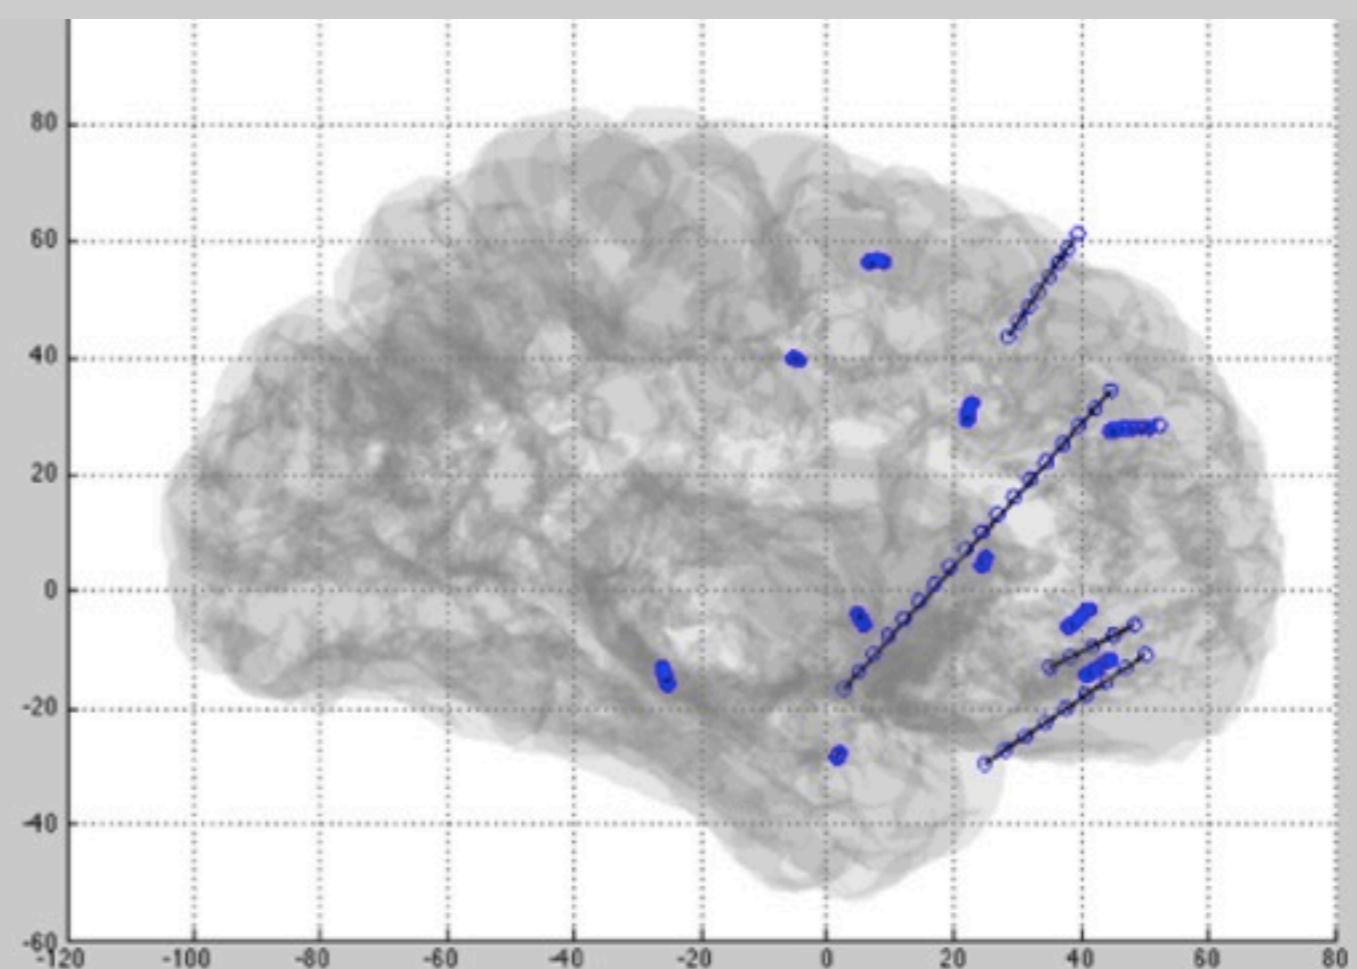
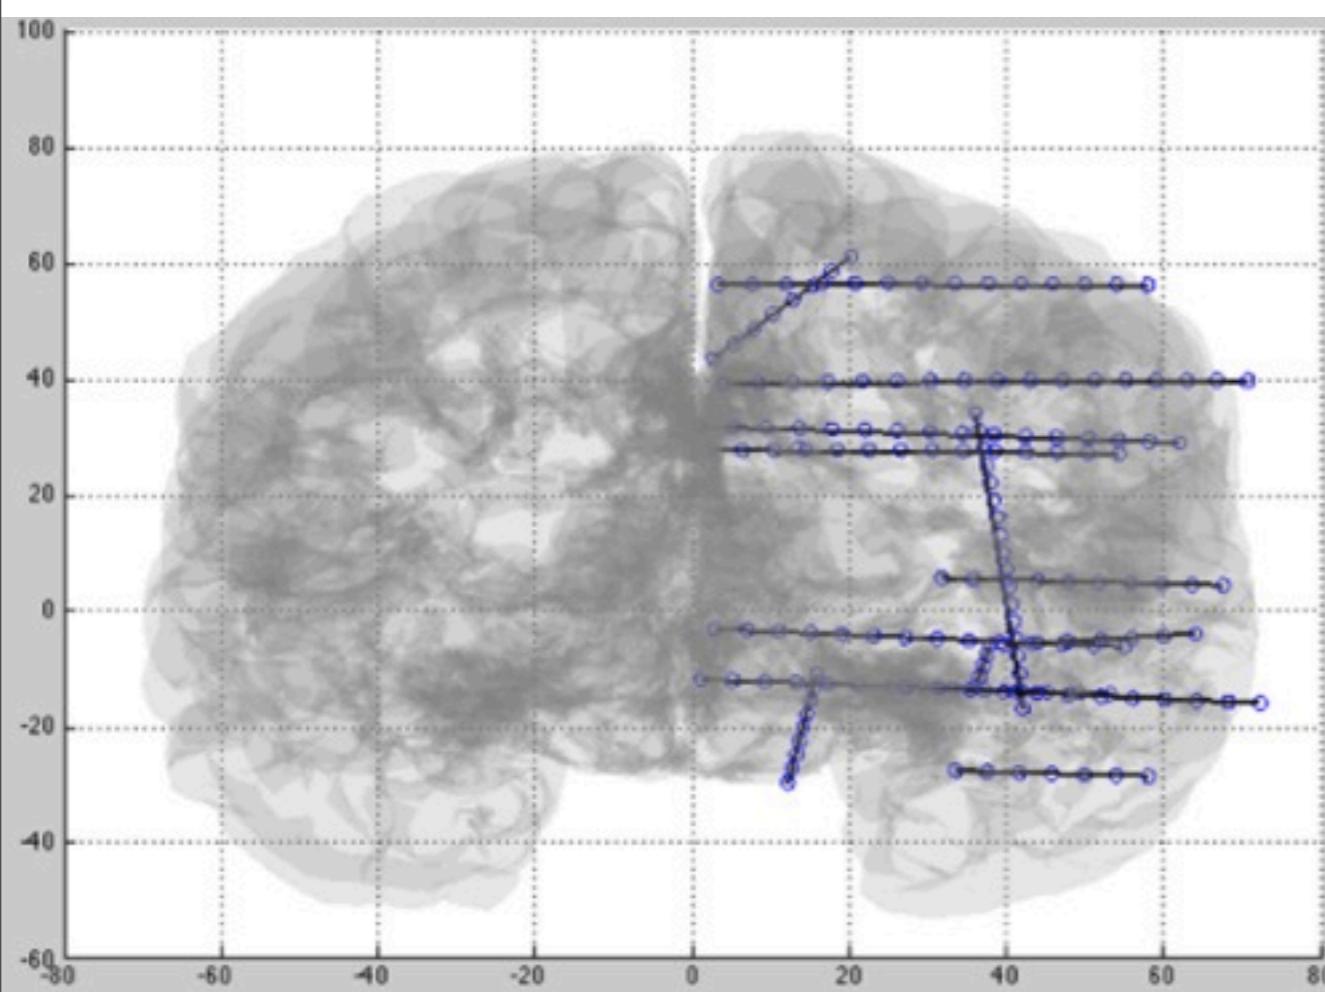
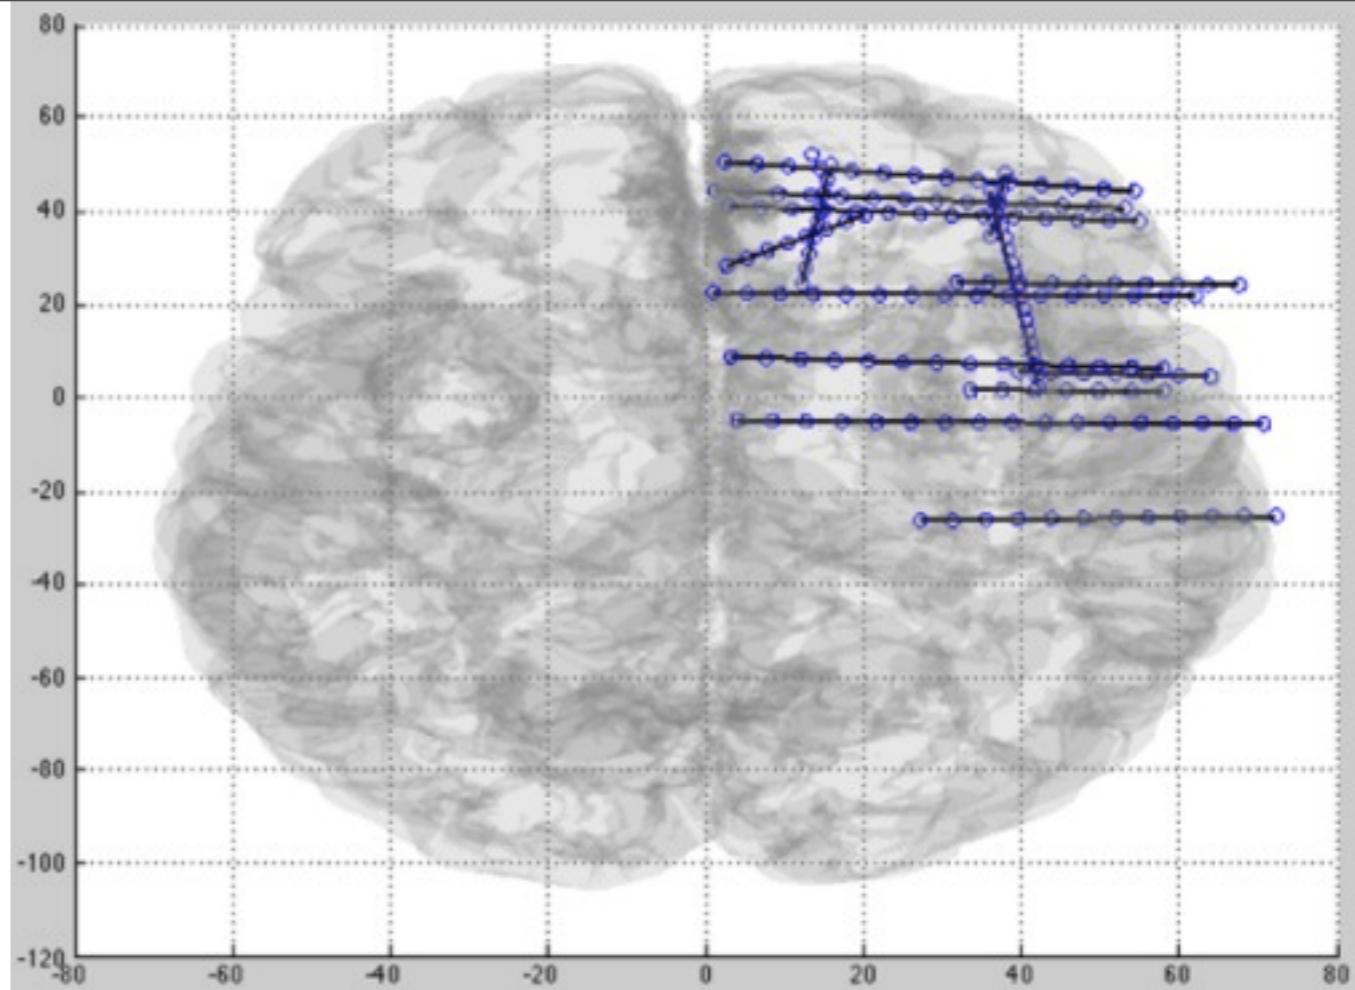






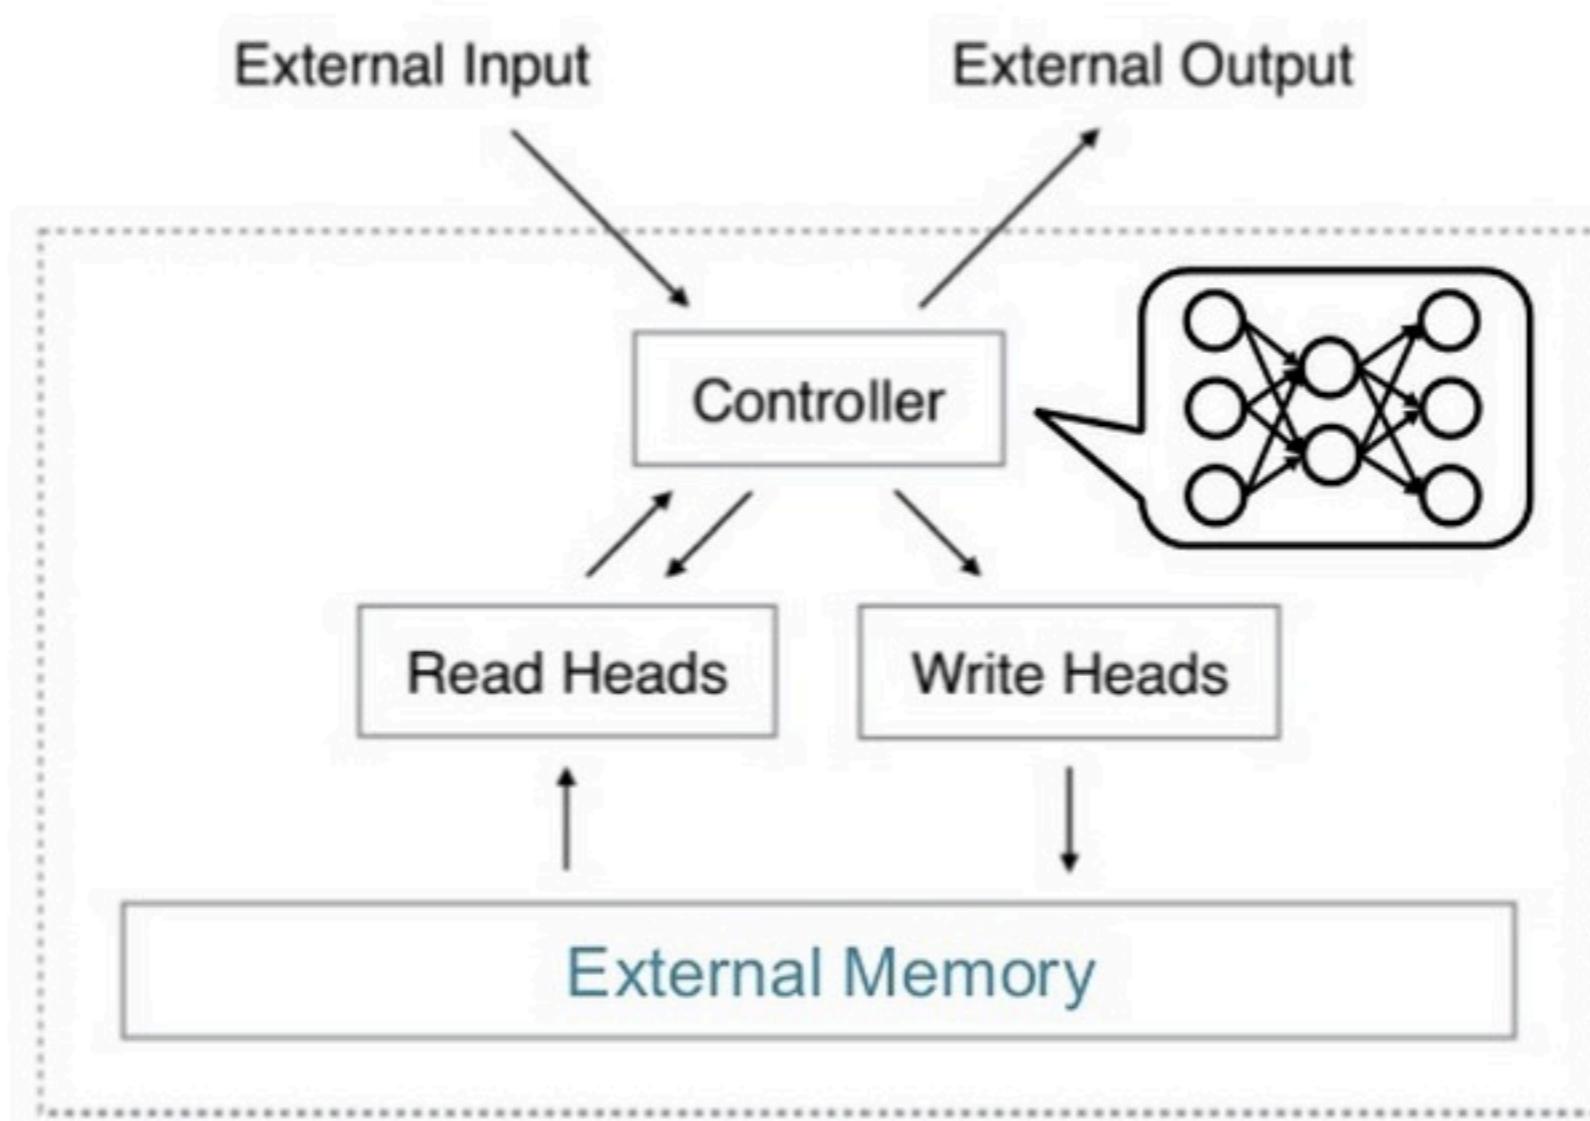


109 patients (>1000 electrodes)  
Hospital de Lyon (France)



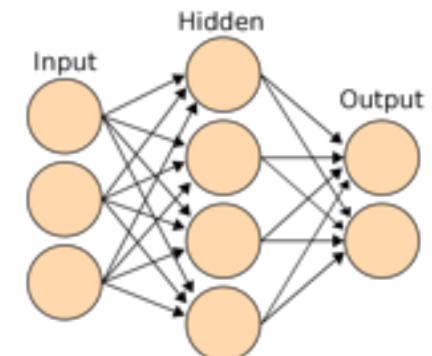
# What about memory?

Neural Turing Machines learn to approximate algorithms from examples!



# Take home message

- Deep Learning = Neural Networks 2.0



- Feature learning & approximate wild functions



- Pushing ML and AI to unthinkable applications

Jaan



# Thanks!

