Modern Artificial Intelligence

S. M. Ali Eslami

August 2015

Google DeepMind


THE UNIVERSITY
of EDINBURGH

Rare

UNIVERSITY OF
OXFORD

Microsoft
Research

Google DeepMind

Chappie (2015)

Ex Machina (2015)

Outline

1. What is machine learning?
2. What is deep learning?
3. What is artificial intelligence?
4. What are we working on at DeepMind?

Algorithms

$$\begin{array}{r} {}^{1}59 \\ +37 \\ \hline 96 \end{array}$$

Algorithms

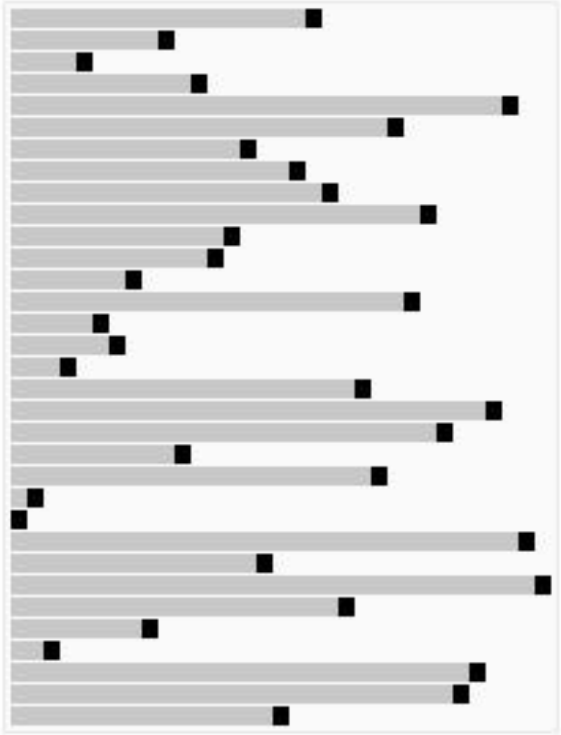

RolandH, Wikipedia (2006)

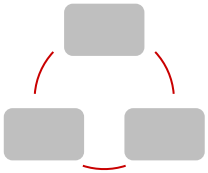

Algorithm

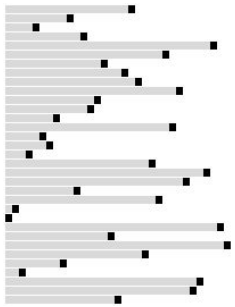

Input

Humans

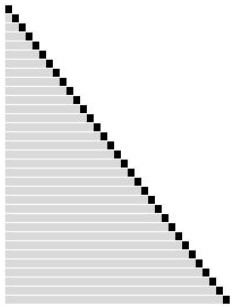

Output

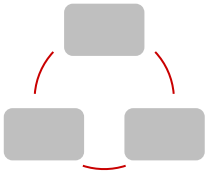

Algorithm

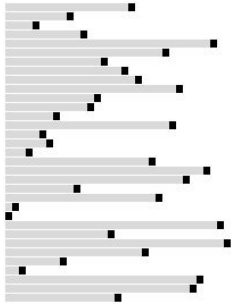

Input

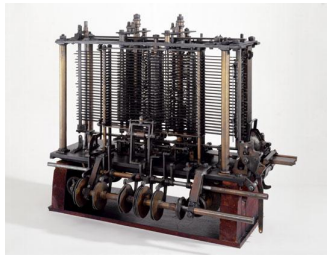

Programmable Computer

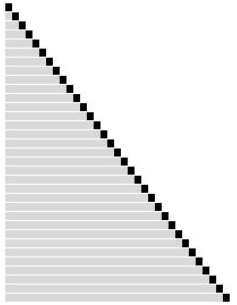

Output

Sorting

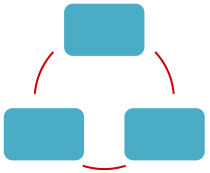

Algorithm

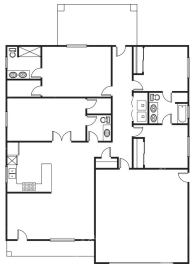

Input

Programmable Computer

Output

Rendering

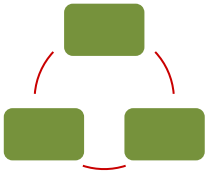

Algorithm

“Horse”

Input

Programmable Computer

Output

Search

?

Algorithm

Input

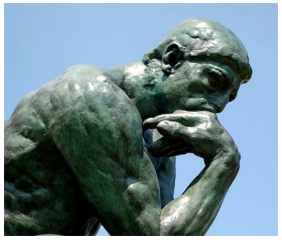

Human

Horse

Output

Image Classification

?

Algorithm

Input

Computer

Horse

Output

Input

Algorithm

Computer

Horse

Cow

Horse

Output

Deep Learning

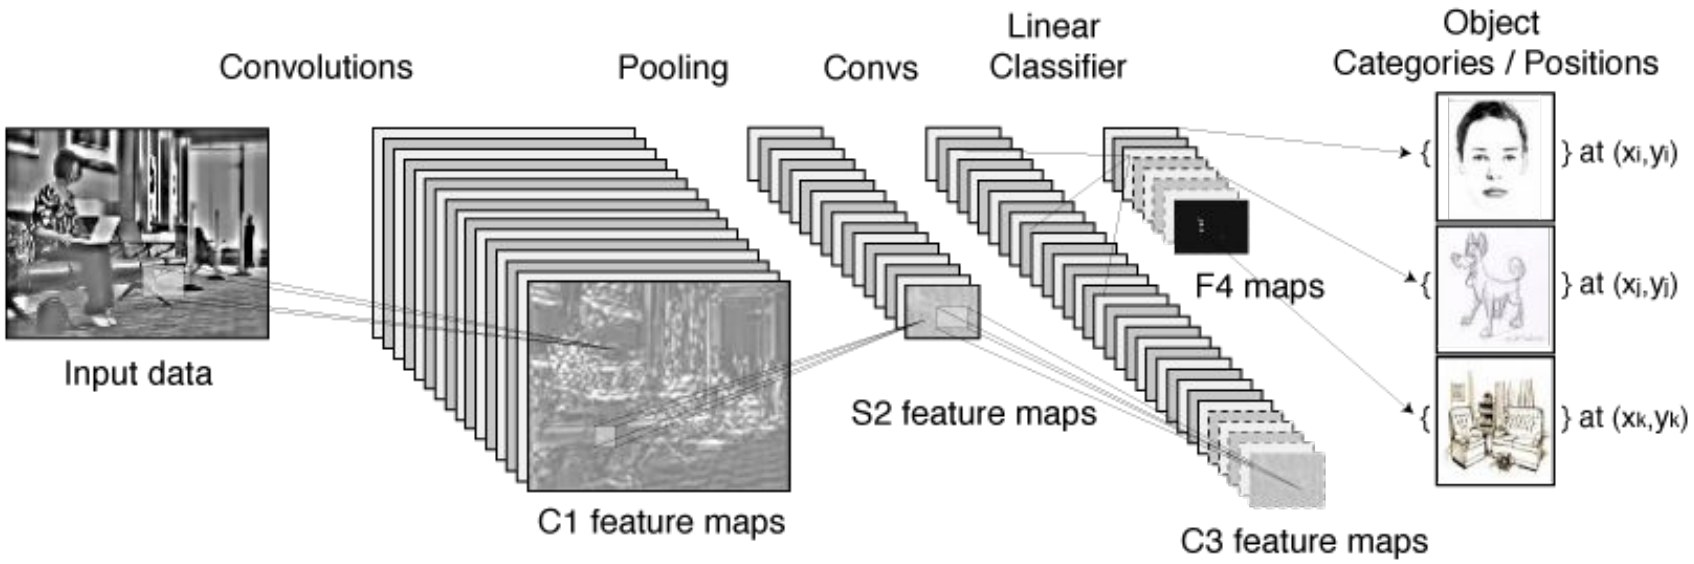

Convolutional neural networks for image classification
Torch (2015)

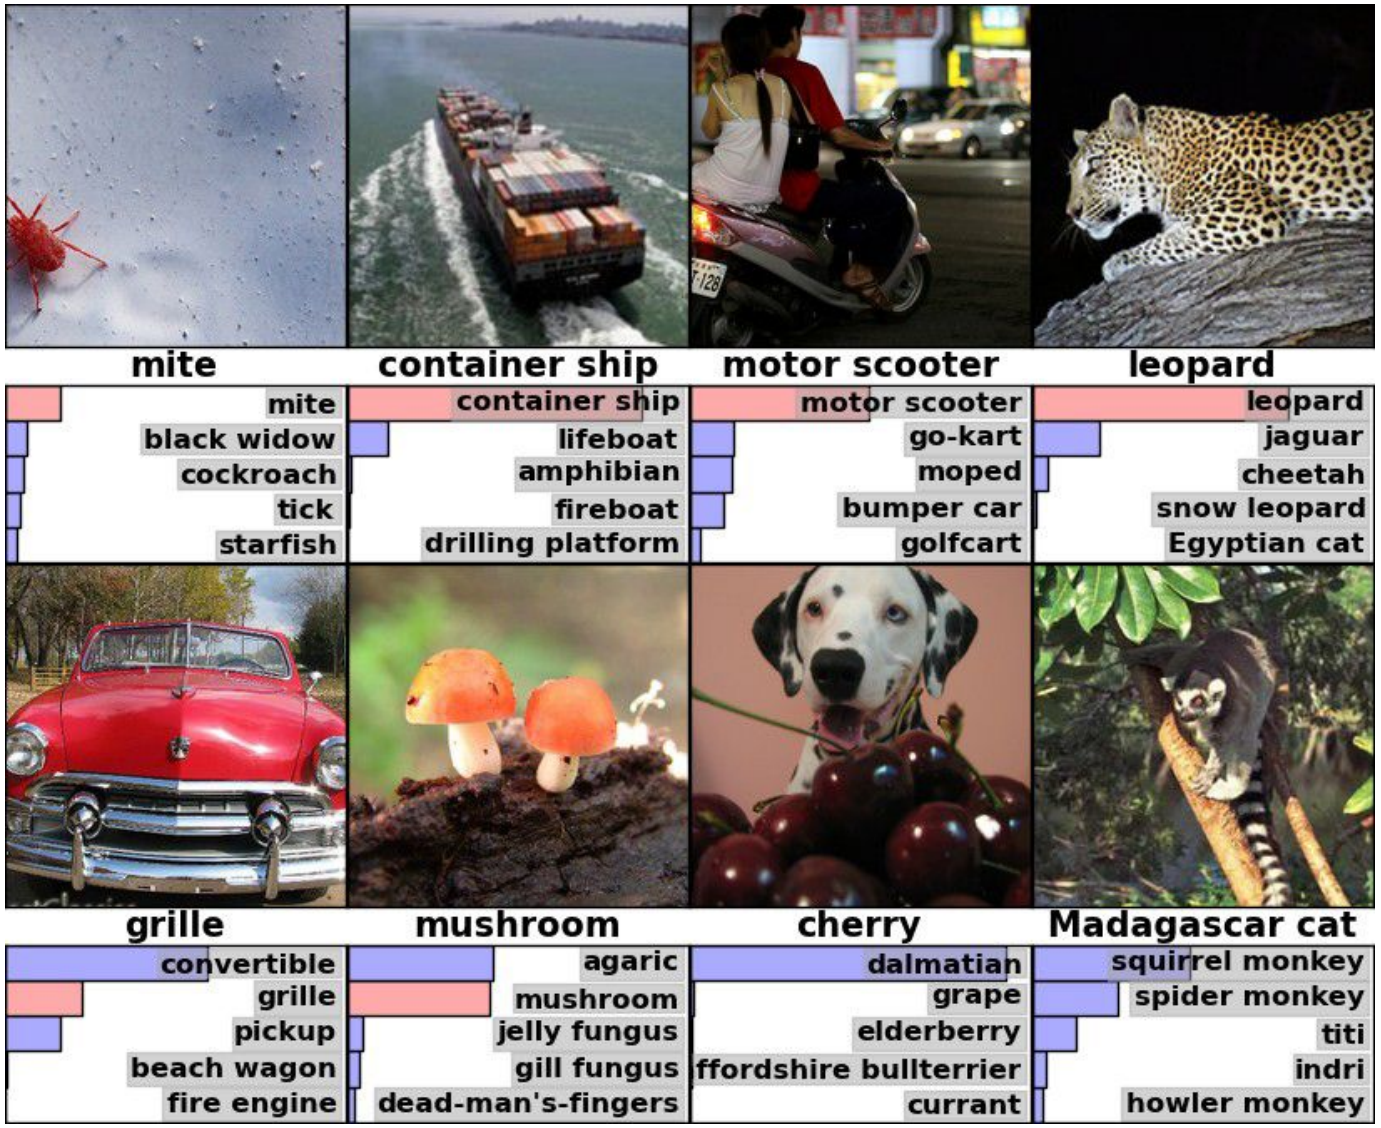

Krizhevsky et al. (2012)

Tasks thought to require intelligence

- Image understanding,
- Natural language processing,
- Knowledge acquisition,
- And many others.

Machine learning has led to significant advances in almost all of these tasks.

Deep learning has made it possible to learn end-to-end without pre-programming.

Learning to draw shapes

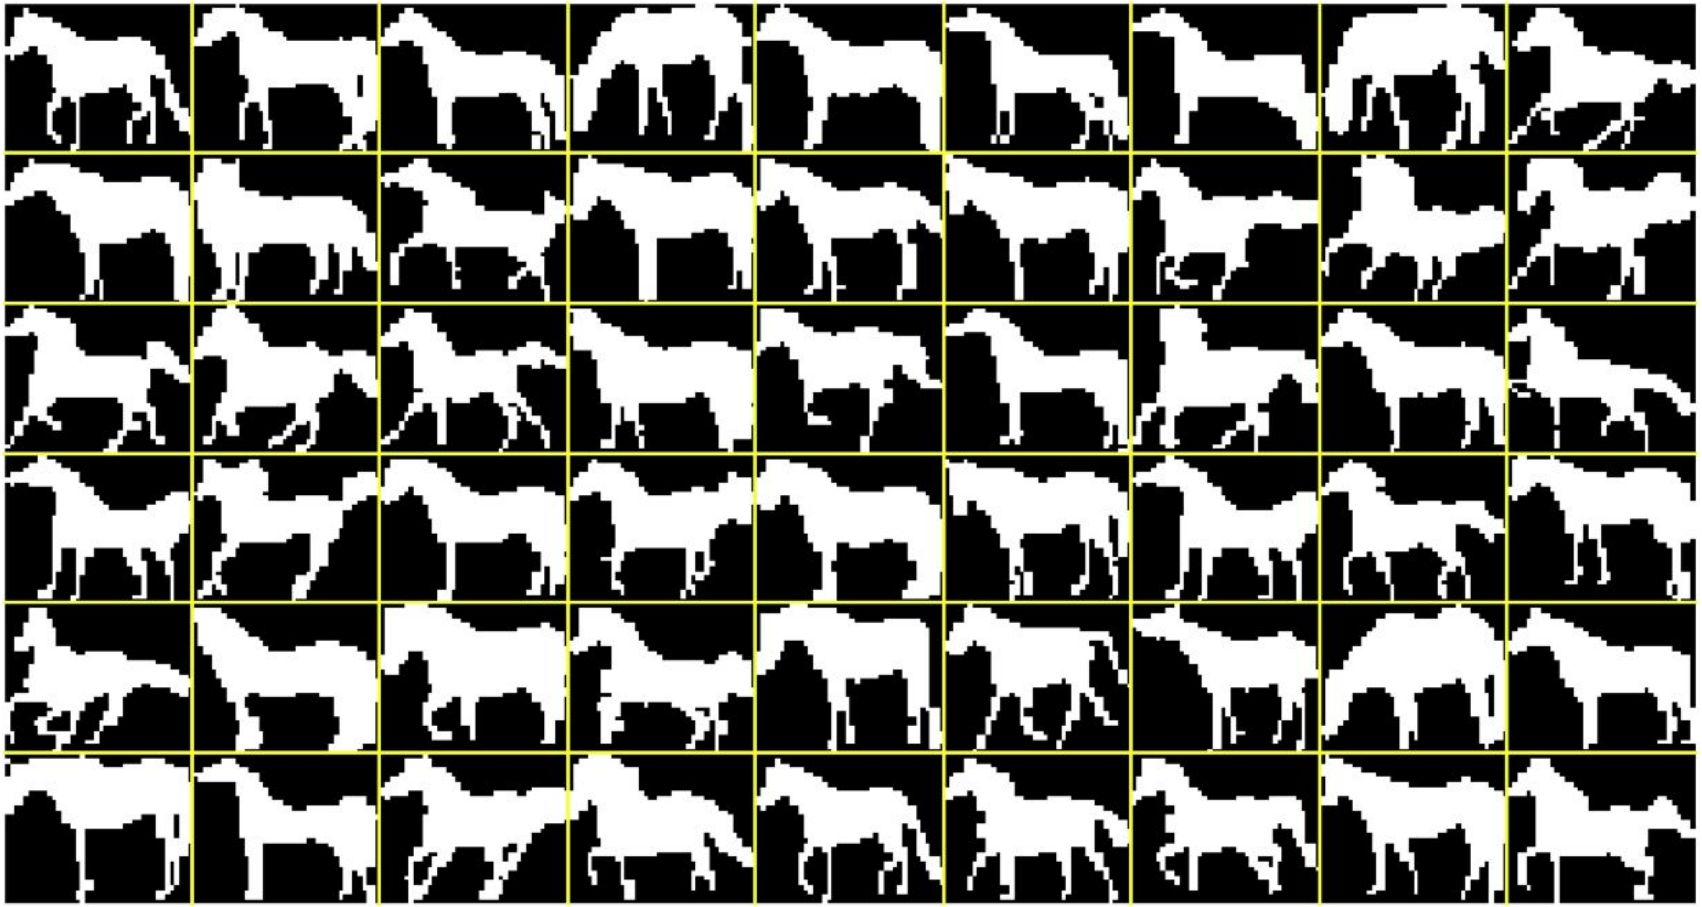

The Shape Boltzmann Machine

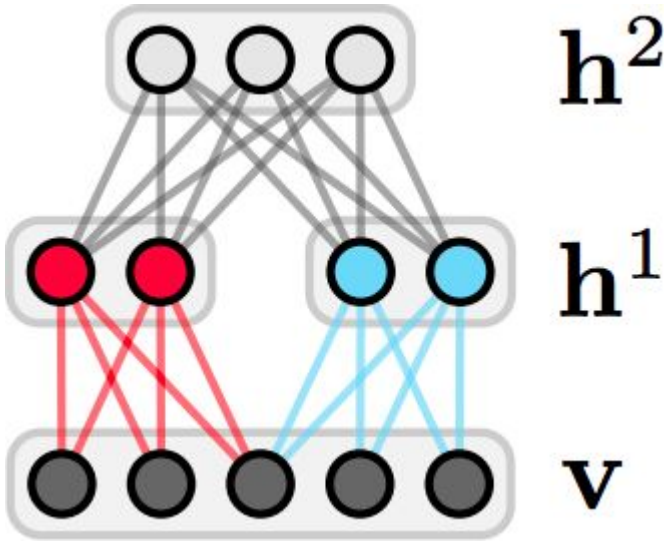

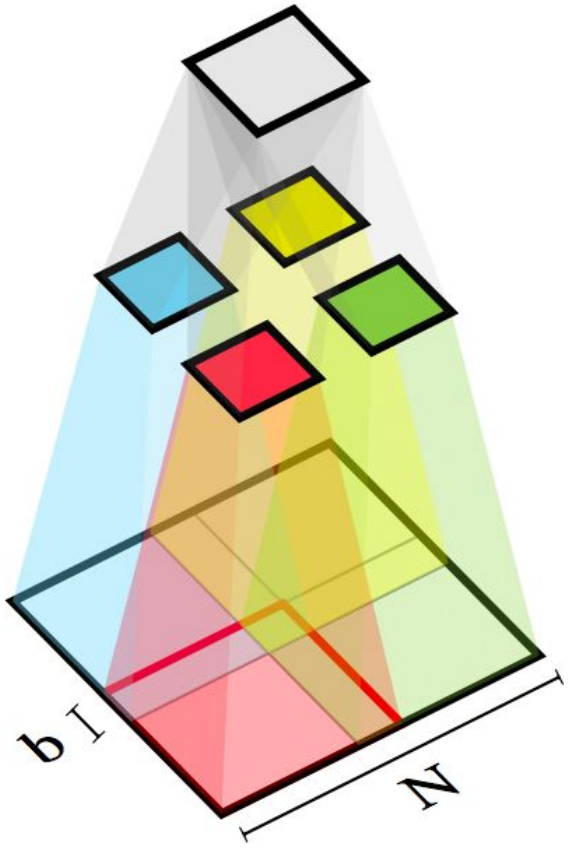

Sampling from an SBM

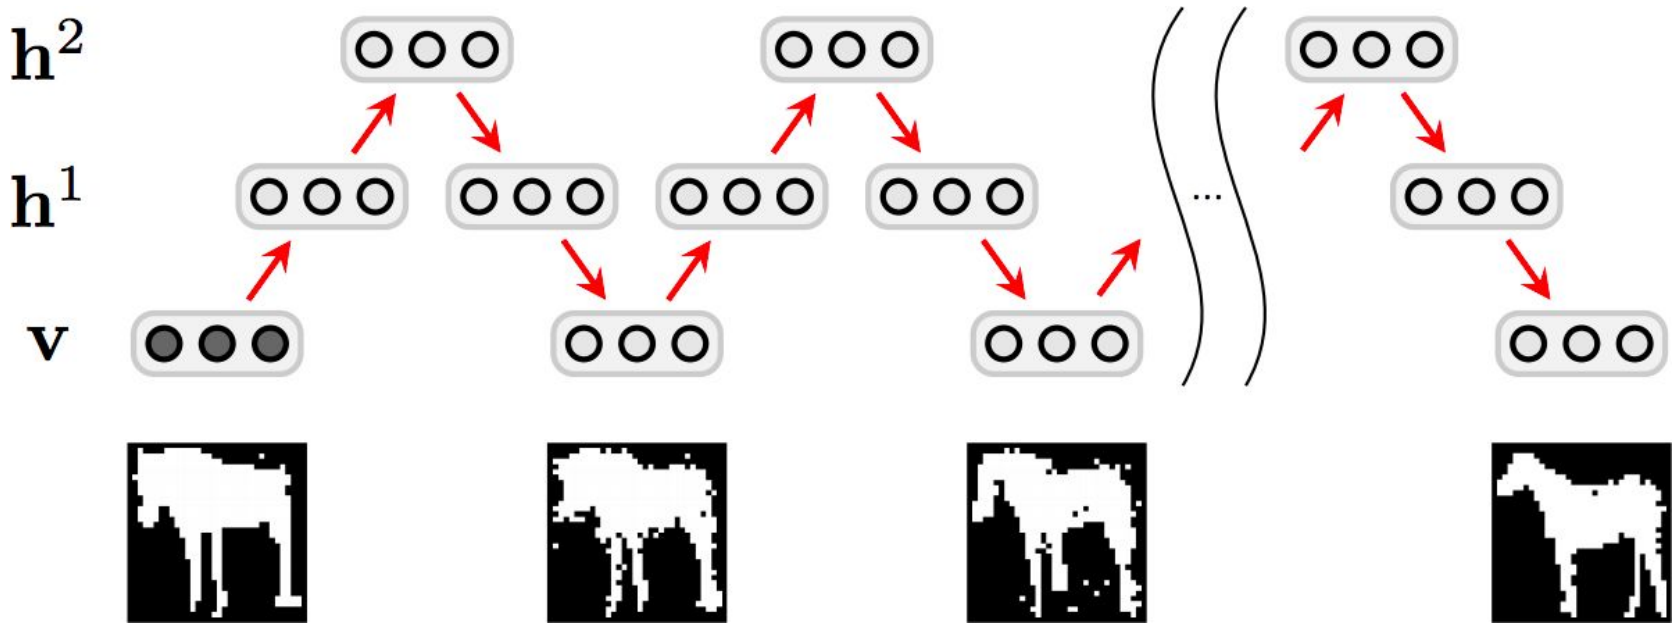

Learning to draw shapes

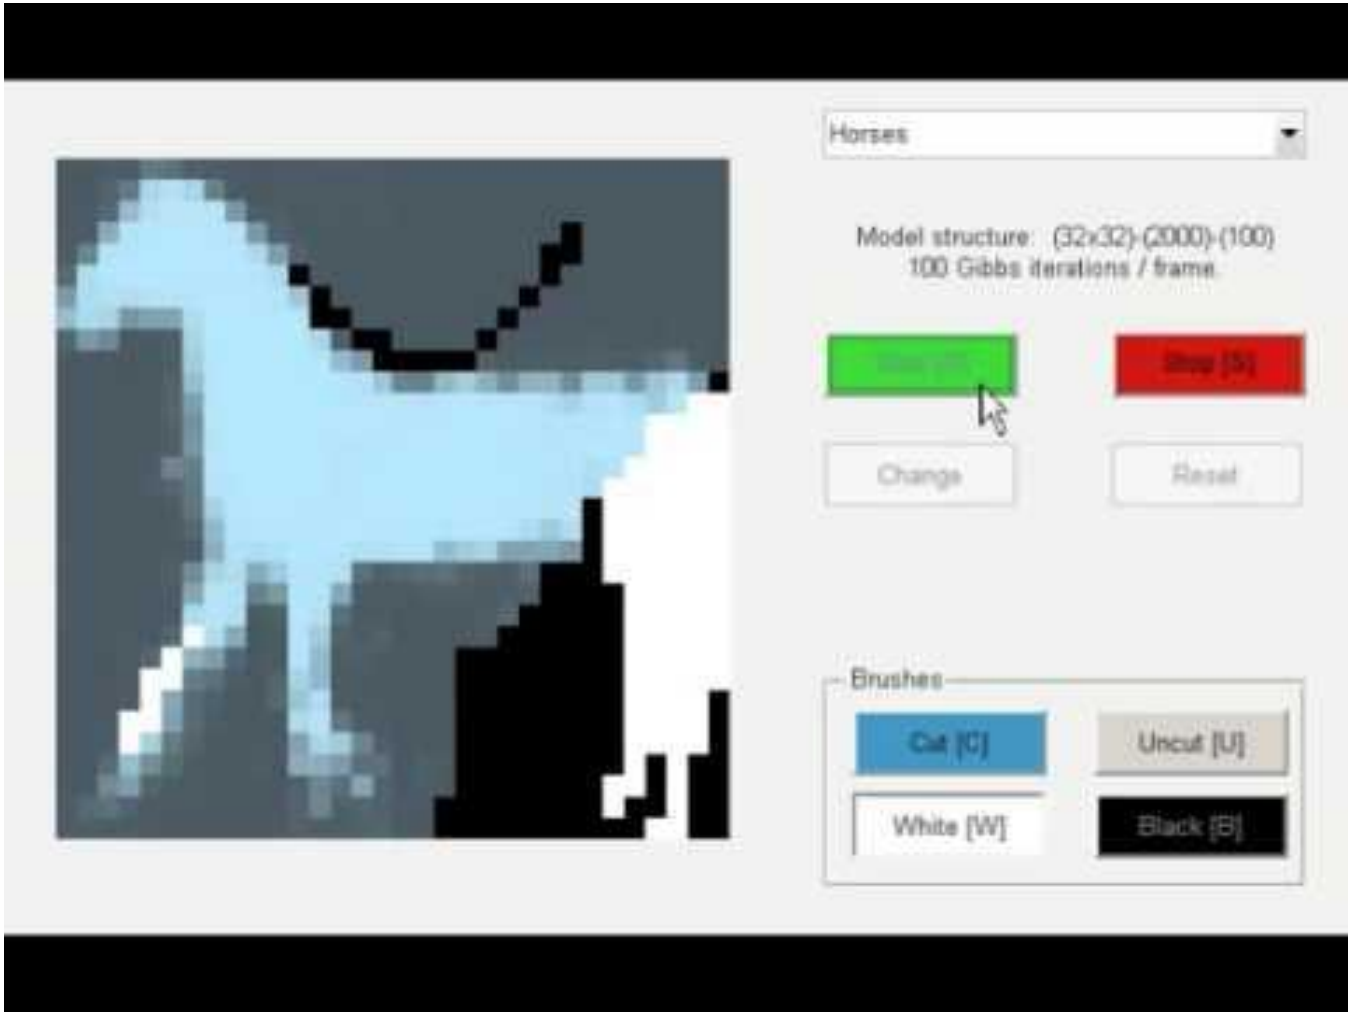

SBM generalisation

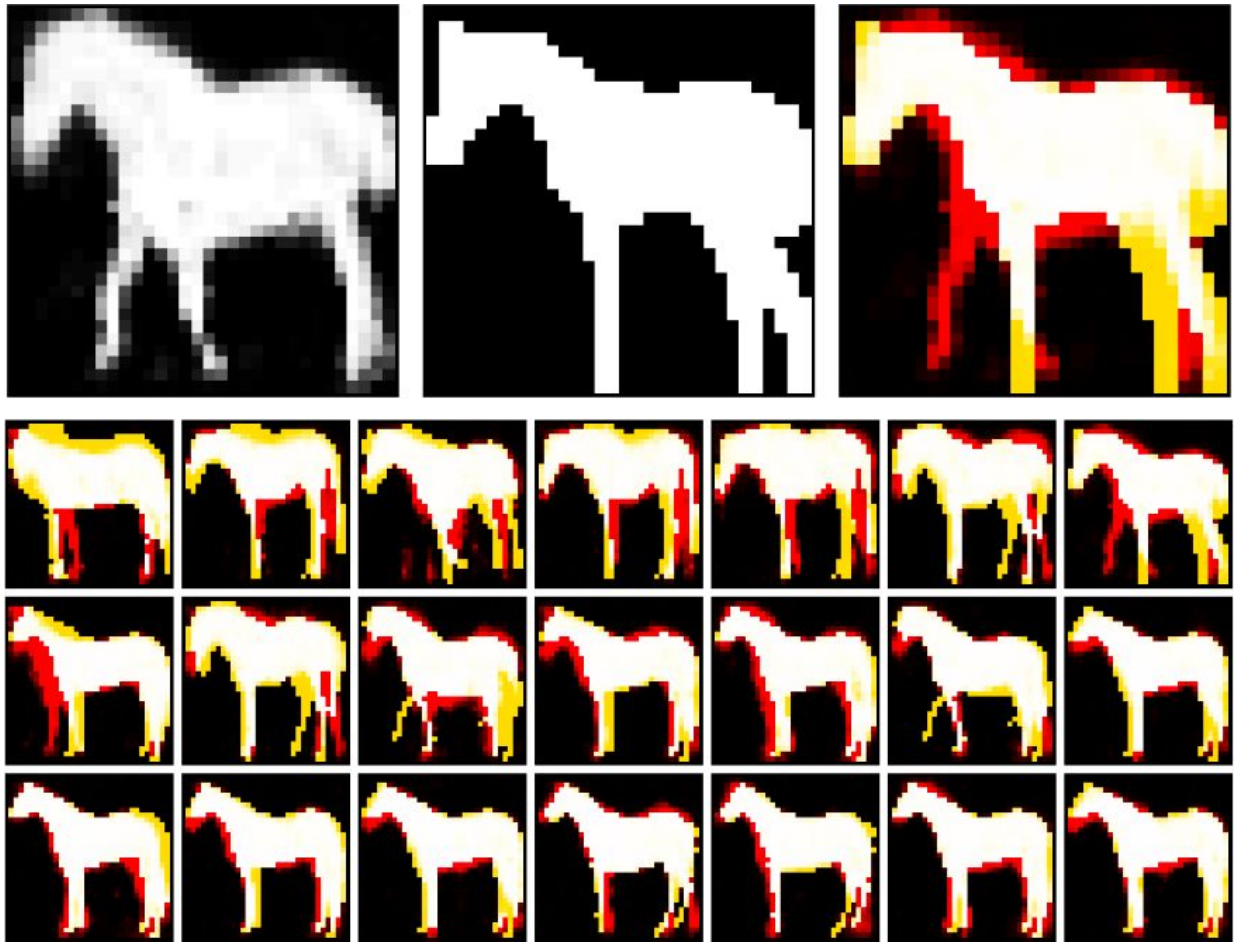

Learning to draw shapes

- **12 million** parameters
- **4 hours** to train

The Shape Boltzmann Machine: a Strong Model of Object Shape
S. M. Ali Eslami, Nicolas Heess, Christopher K. I. Williams, John Winn
International Journal of Computer Vision, Springer (2013)

Learning to segment objects

A Generative Model for Parts-based Object Segmentation

S. M. Ali Eslami, Christopher K. I. Williams

Neural Information Processing Systems (2012)

Learning to segment objects

A Generative Model for Parts-based Object Segmentation

S. M. Ali Eslami, Christopher K. I. Williams

Neural Information Processing Systems (2012)

Recurrent Neural Networks for Image Generation

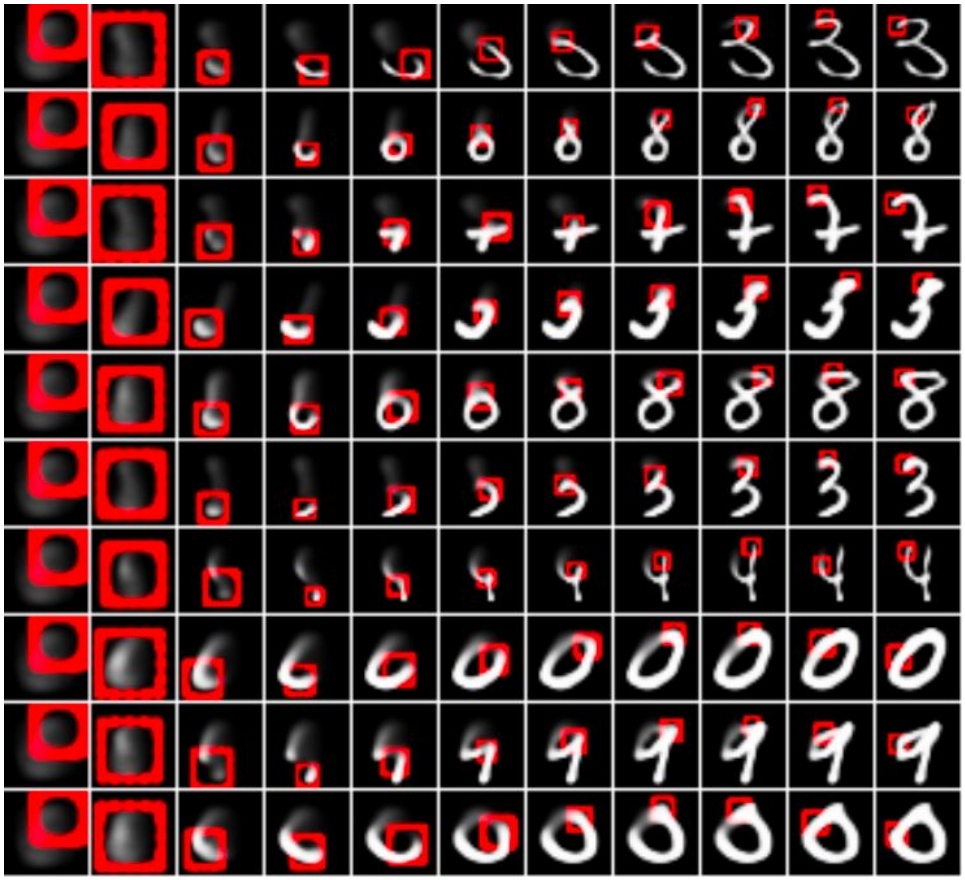

Gregor et al. (2015)

Recurrent Neural Networks for Image Generation

Gregor et al. (2015)

Artificial General Intelligence (AGI)

Machine learning has led to significant advances in many specific tasks.

Deep learning has made it possible to learn end-to-end without pre-programming.

Artificial General Intelligence is looking for agents that successfully operate across a wide range of tasks.

Architecture

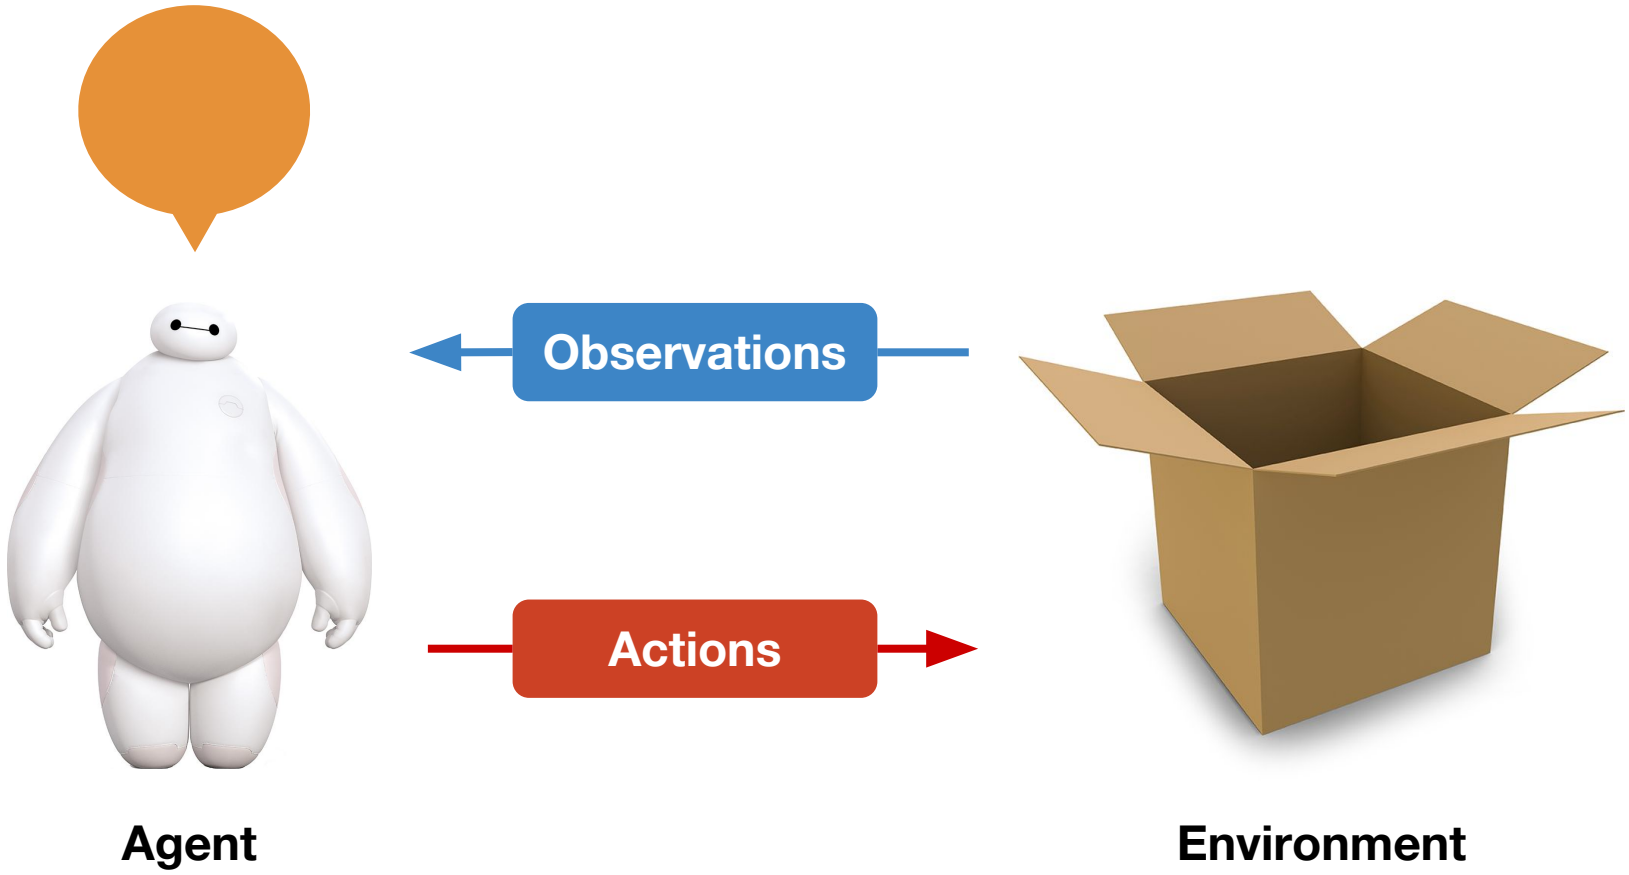

Learning to move

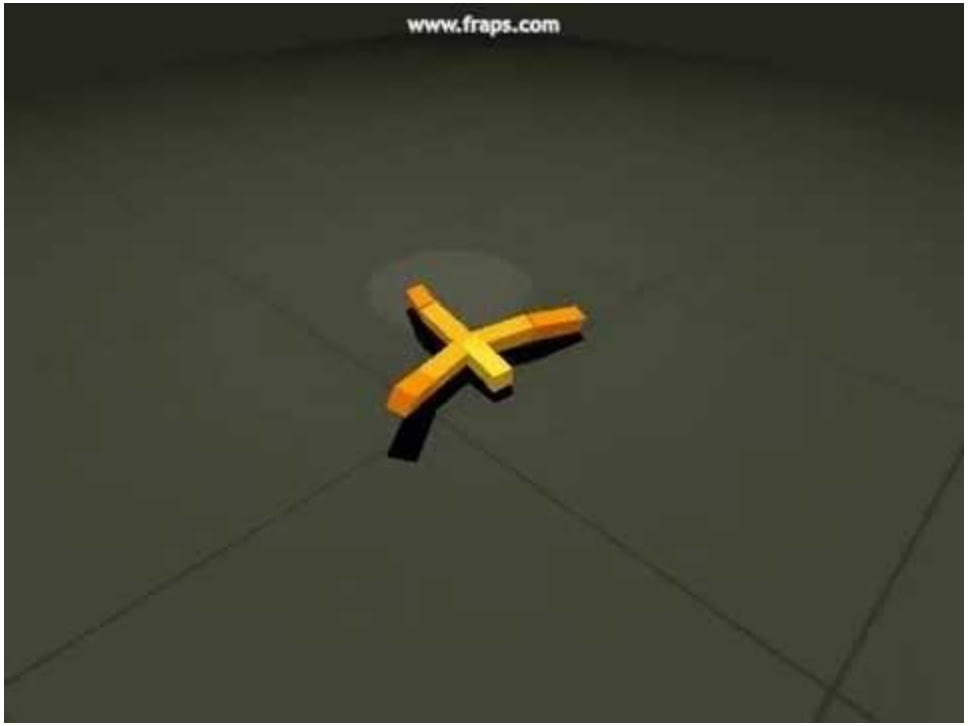

Learning to move

From pixels to actions

Games are the perfect platform for developing and testing AI algorithms

- Embodied cognition
- Right level of complexity
- Measurable progress
- Unlimited training data
- No testing bias
- Faster than real-time
- Thousands of tests in parallel

ATARI agents

ATARI 2600 testbed: 100+ classic 8-bit Atari games from the 70/80s

- Observations: Raw video (**~30k dimensional**)
- Actions: 18 buttons but **not** told what they do
- Goal: Simply to maximize score

- Everything learnt from **scratch**
- **Zero** pre-programmed knowledge
- **One** algorithm to play **all** the different games

Space Invaders agent

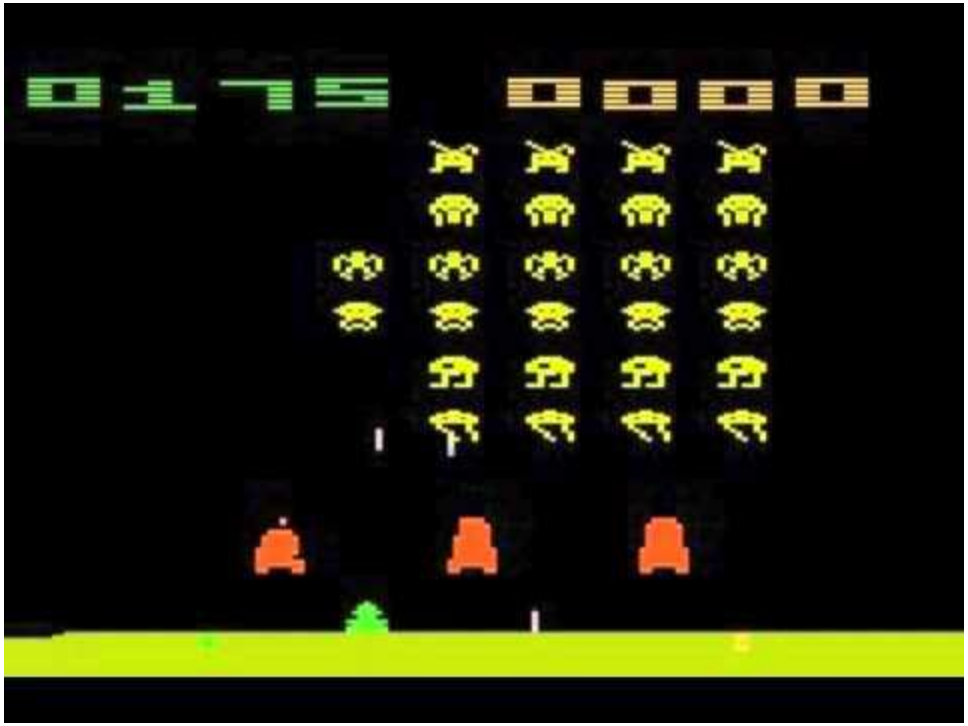

Breakout agent

General Atari agent

Human-level control through deep reinforcement learning

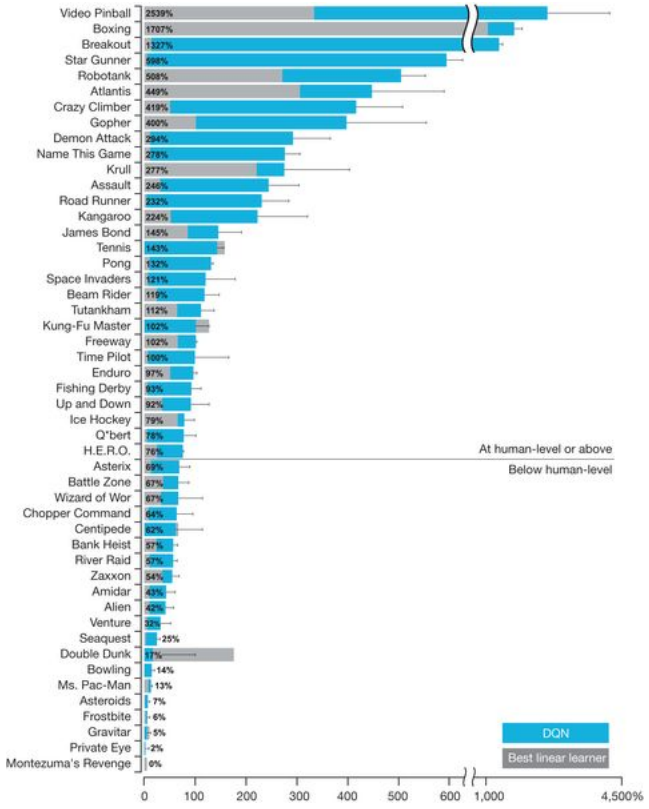

Mnih et al. (2015)

nature.com/articles/nature14236

Human-level control through deep reinforcement learning

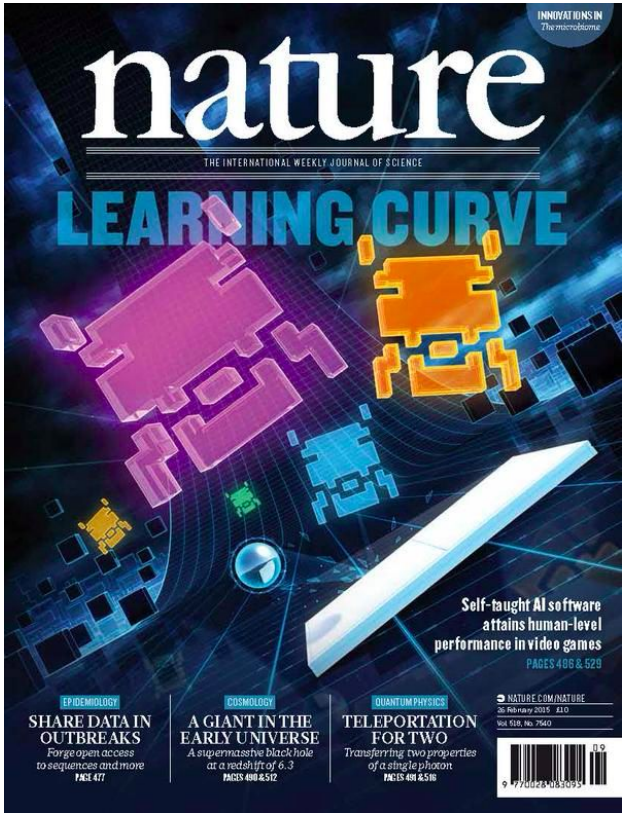

Mnih et al. (2015)

nature.com/articles/nature14236

Transfer learning

Applying previously learnt knowledge to a new situation

1. Perception

Identify the salient features in an environment

2. Conceptualisation

Re-represent those features as suitably abstract concepts

3. Action

Select and apply knowledge to help performance in new environments

Behind the scenes

Lucky not to have to ‘publish or perish’

ICML 2014 – 6 papers

NIPS 2014 – 8 papers

ICML 2015 – 6 papers

Mix of deep learning, reinforcement learning, optimisation and search:

- Towards End-to-End Speech Recognition with Recurrent Neural Networks
- Deep AutoRegressive Networks
- Stochastic Backpropagation and Approximate Inference in Deep Generative Models
- Neural Variational Inference and Learning in Belief Networks
- Skip Context Tree Switching
- Deterministic Policy Gradient Algorithms

Society's biggest challenges

Information overload

Automatic extraction of actionable knowledge

Help humans be faster

System complexity

Climate, biology, energy, macroeconomics

Help humans be smarter

AGI *could* be a meta-solution to these problems

http://arkitus.com

http://deepmind.com

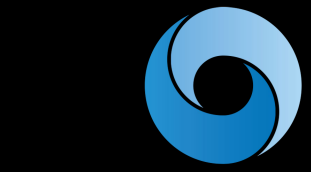

Google DeepMind