Modern Artificial Intelligence

S. M. Ali Eslami

December 2015

THE UNIVERSITY of EDINBURGH

Rare

UNIVERSITY OF OXFORD

Microsoft Research

Google DeepMind

Elon Musk

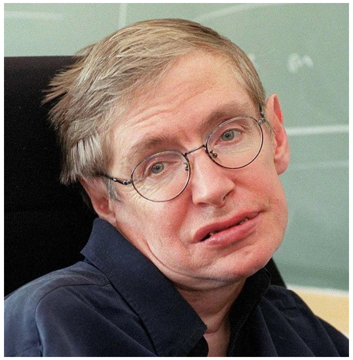

Stephen Hawking

Chappie (2015)

Ex Machina (2015)

Outline

1. Artificial General Intelligence
2. Deep Learning
3. Reinforcement Learning
4. Model-based Methods
5. Reinforced Variational Inference

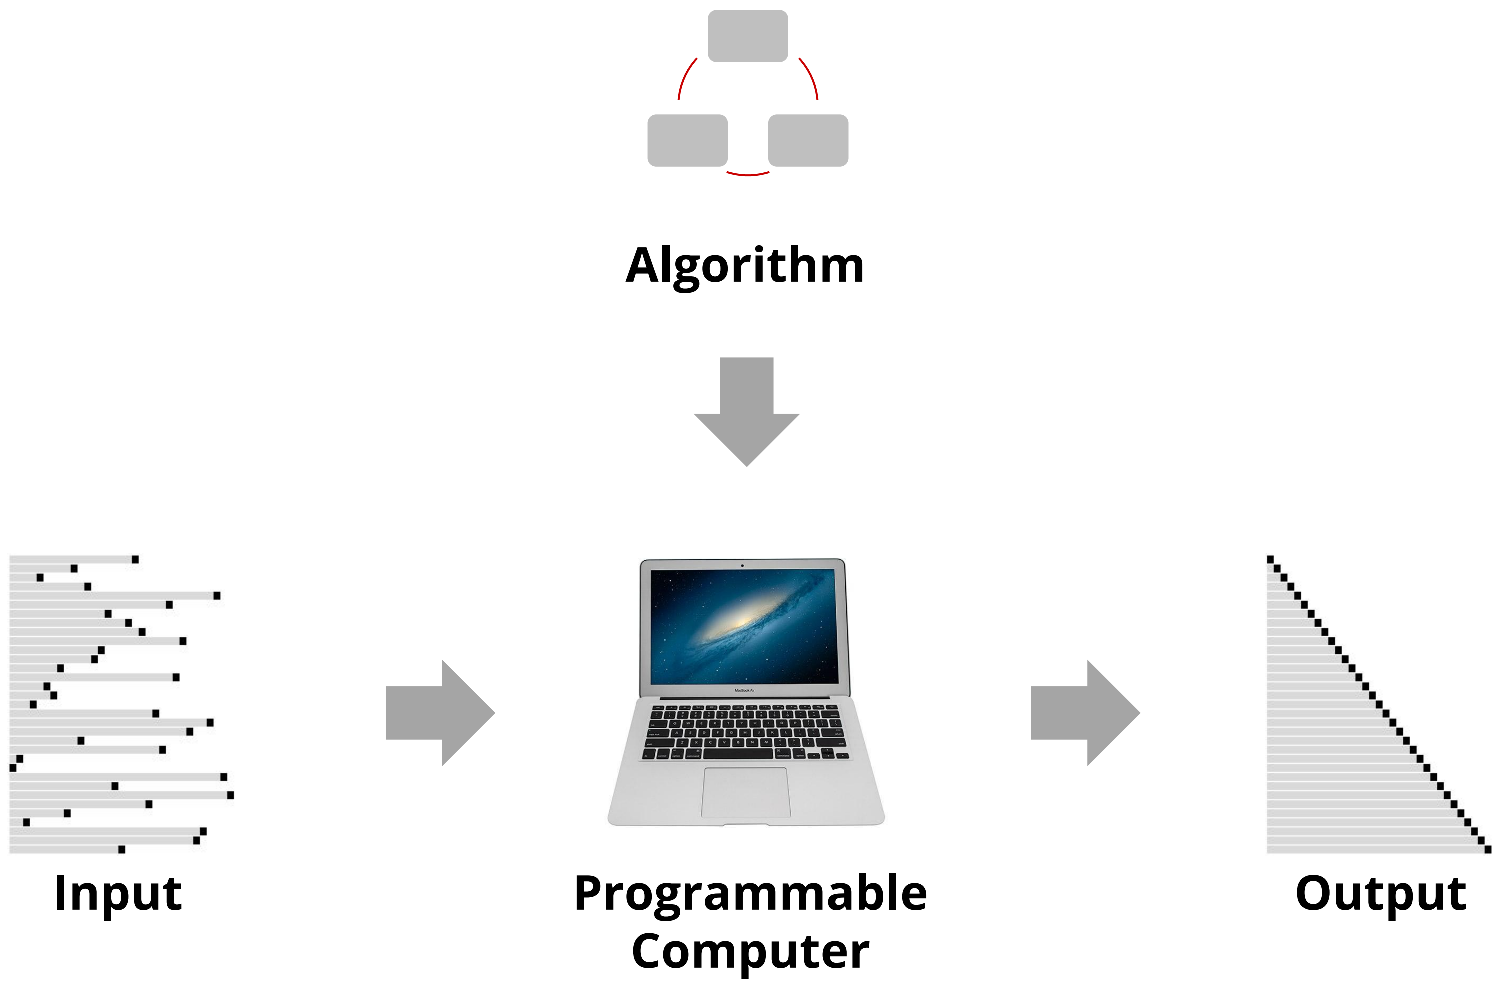
Algorithm
Input
Programmable
Computer
Output

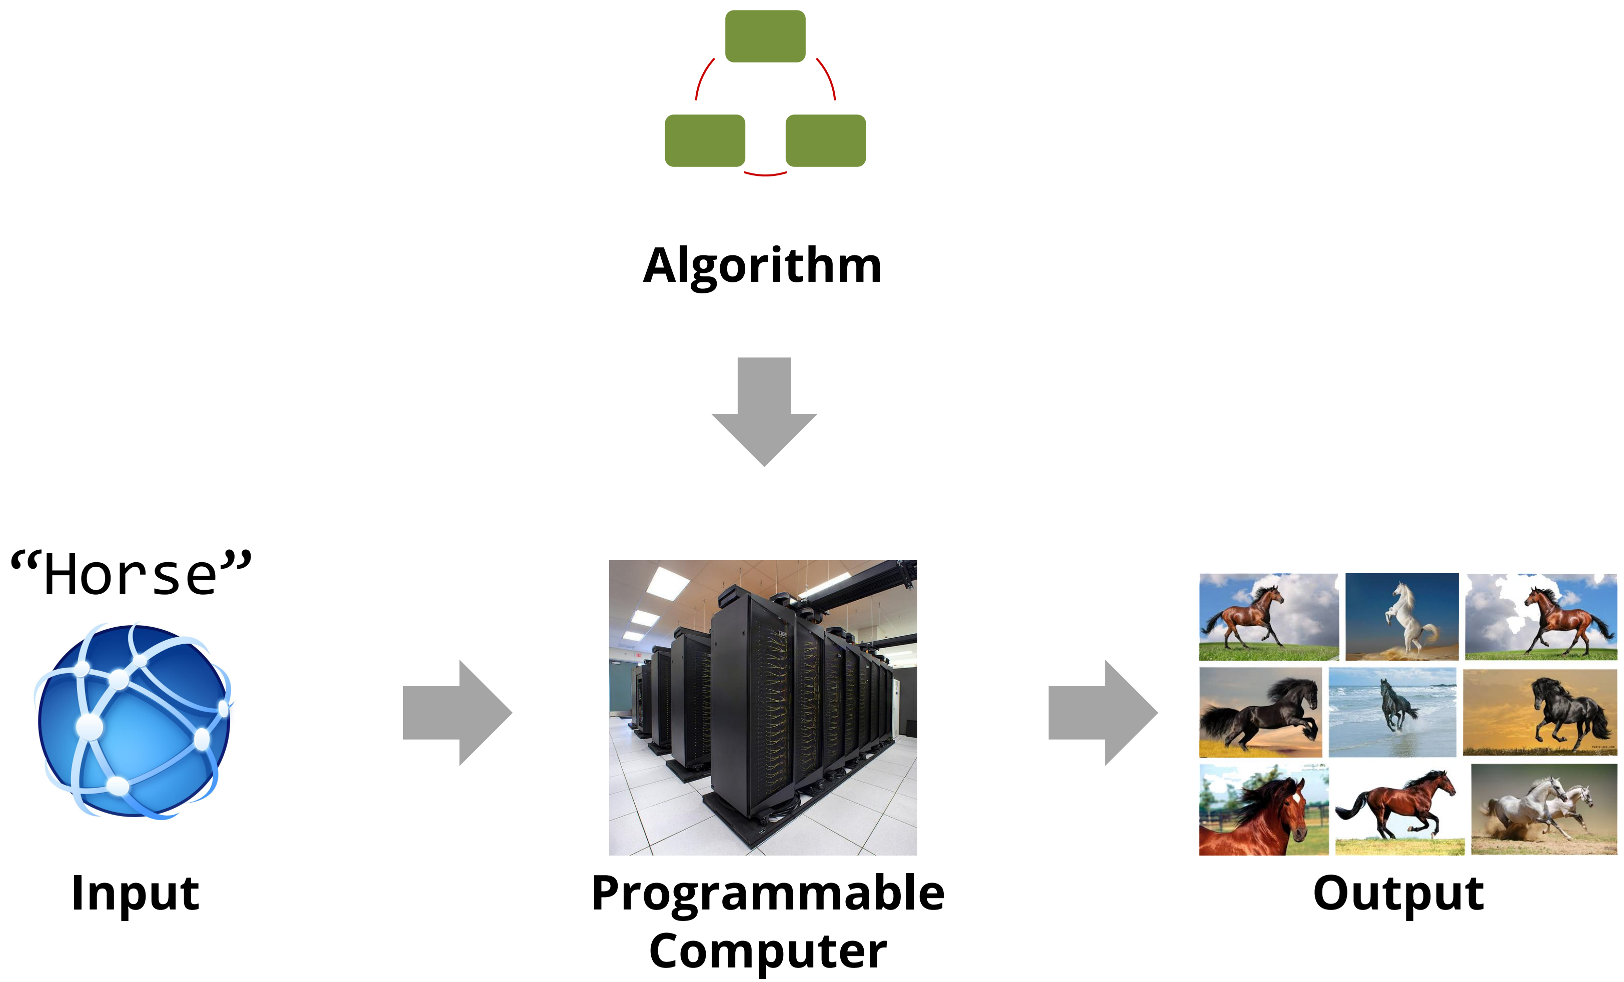

Search

?

Algorithm

Input

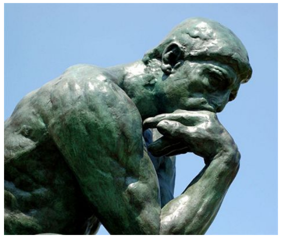

Human

Horse

Output

Image Classification

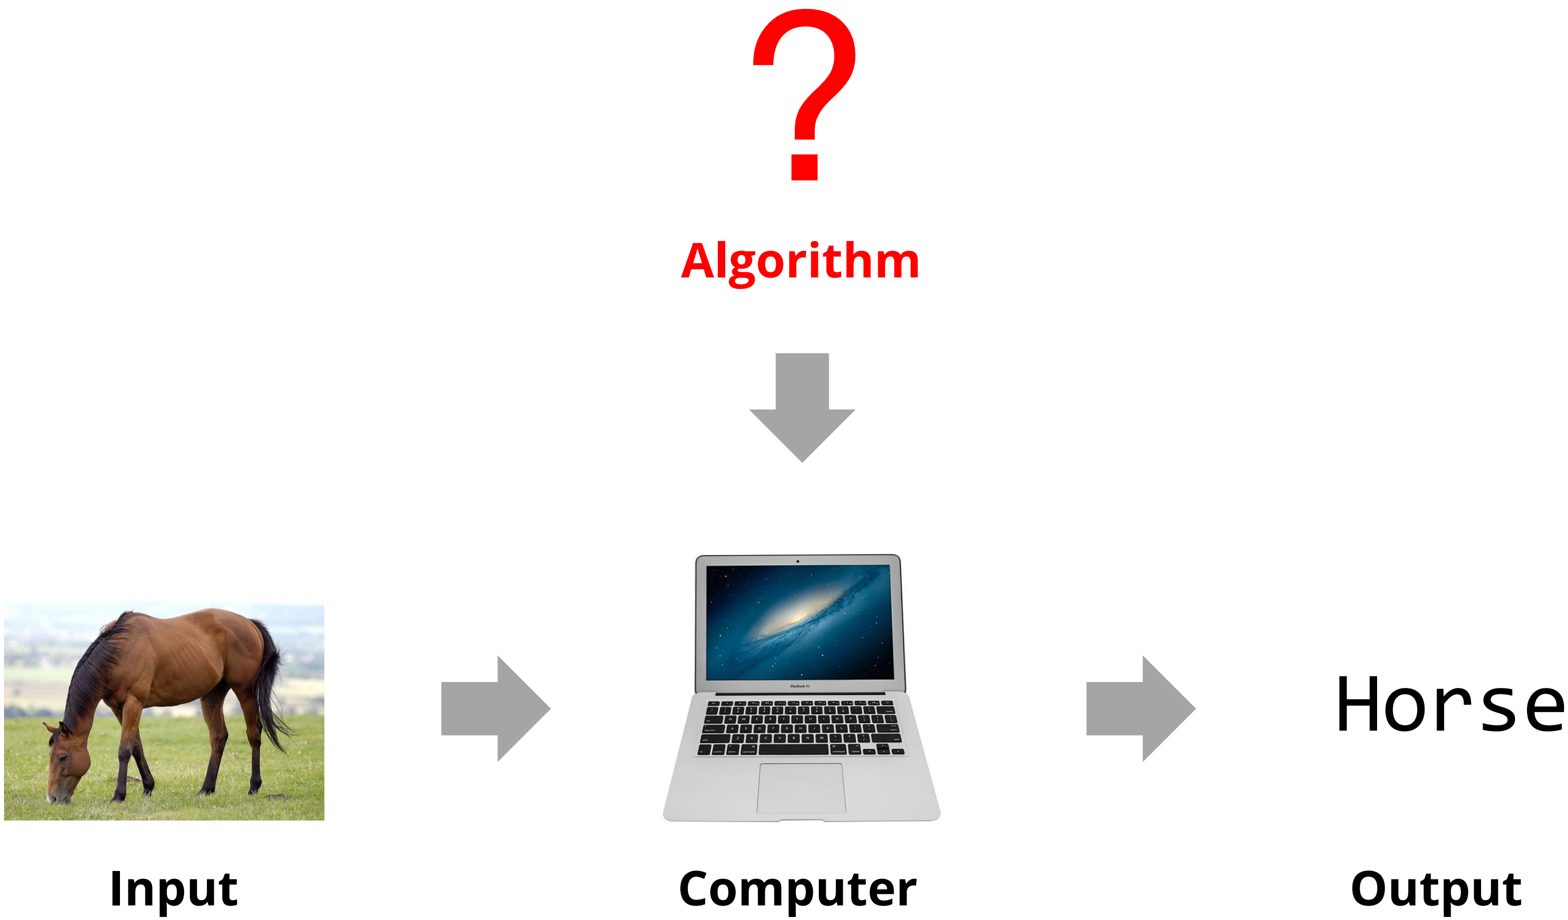

Image Classification

Tasks thought to require intelligence

- Image understanding,
- Natural language processing,
- Knowledge acquisition,
- Text understanding,
- Planning,
- Robotics,
- Forecasting,
- And many others.

Can a general system achieve all these tasks?

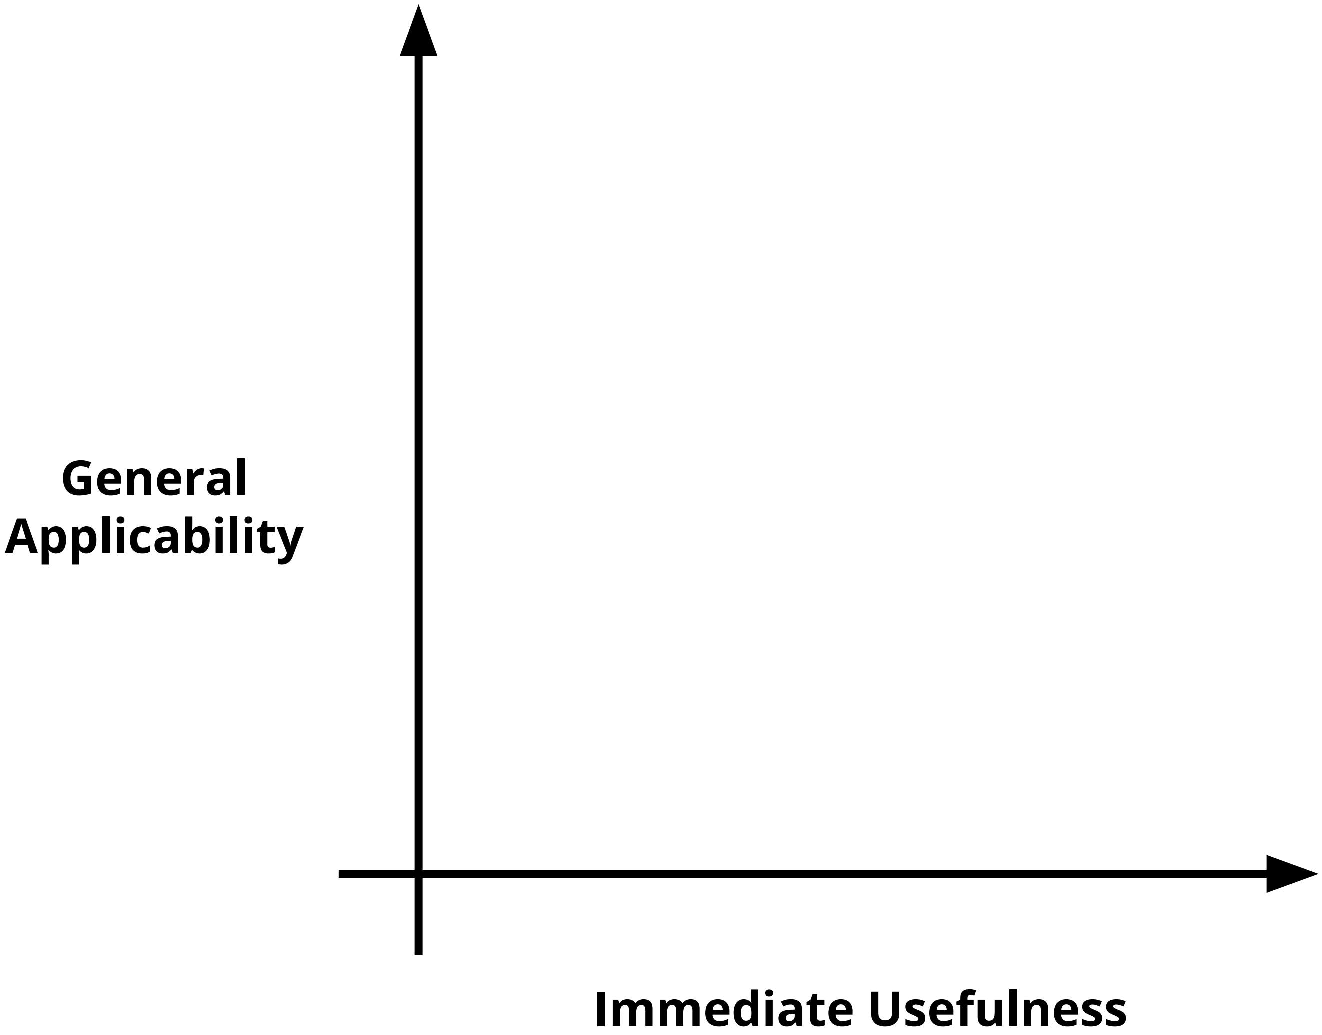
General
Applicability
Immediate Usefulness

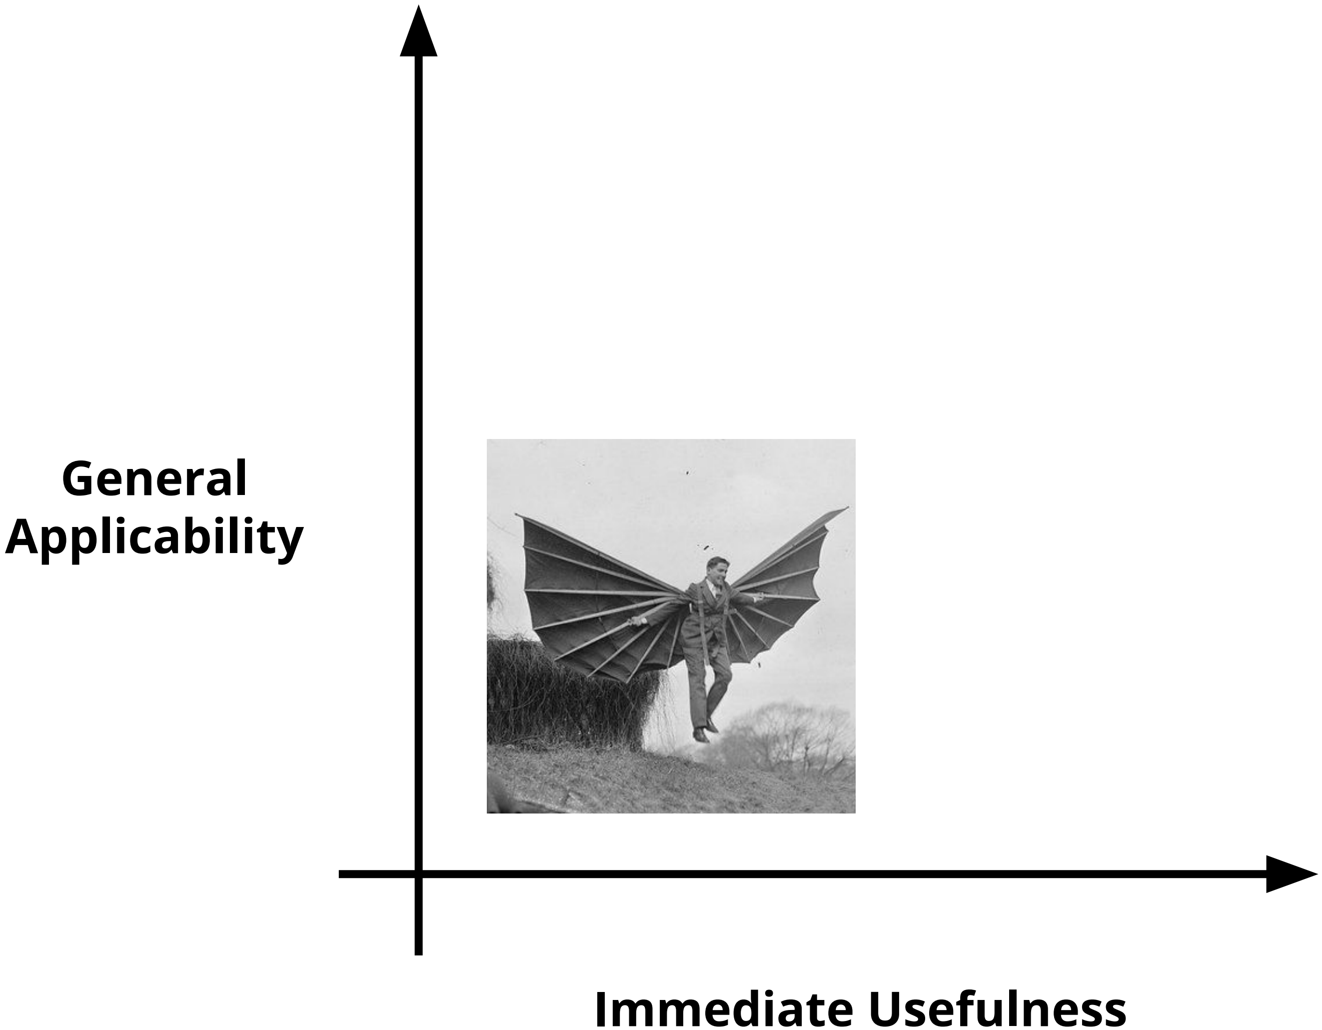
General Applicability
Immediate Usefulness

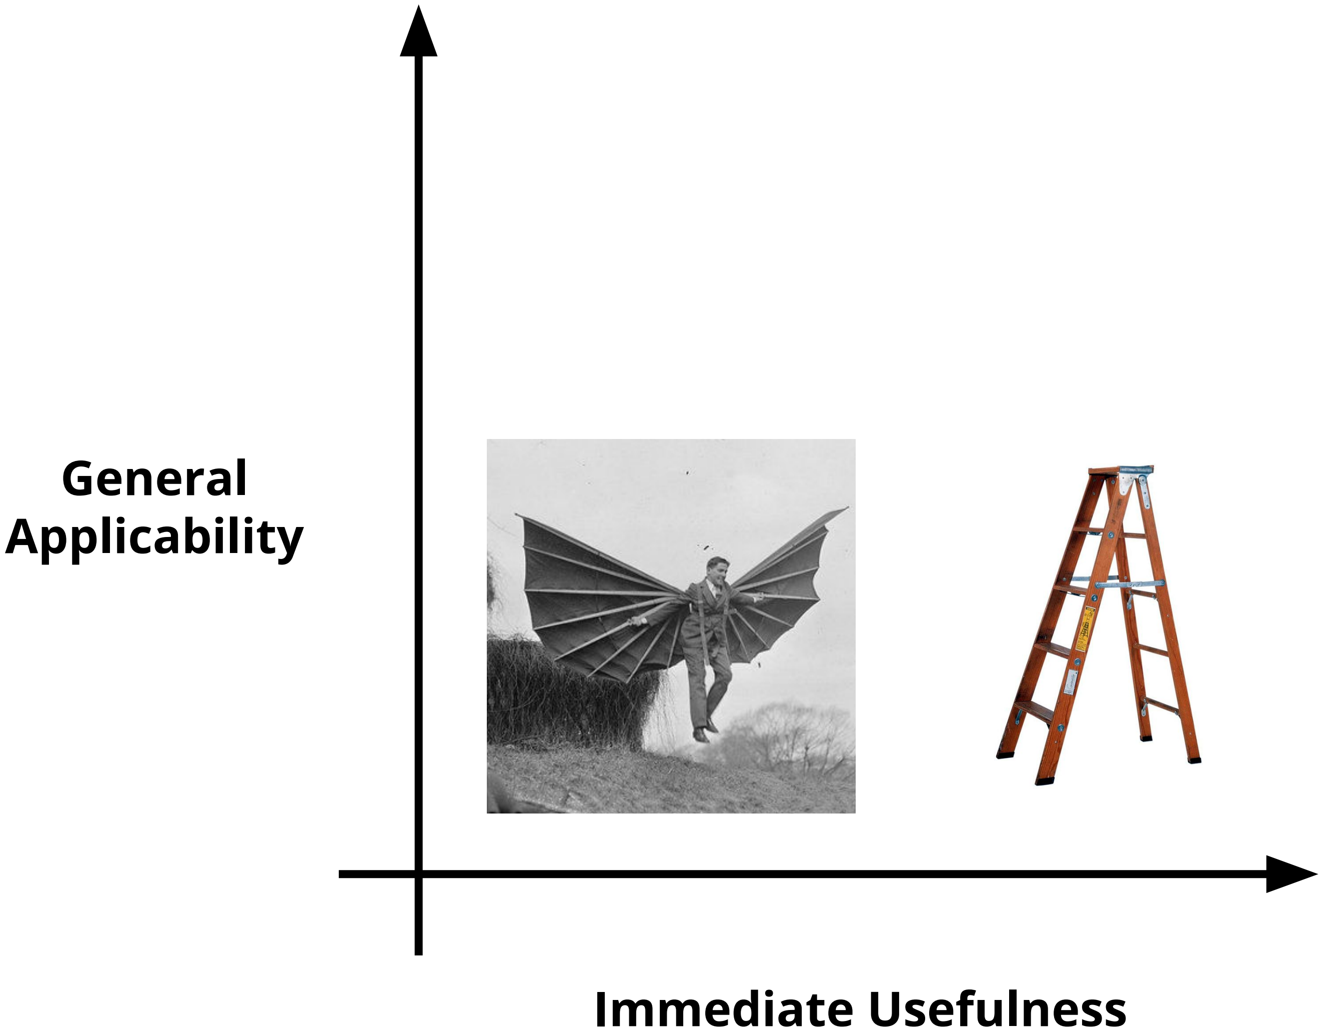
General Applicability
Immediate Usefulness

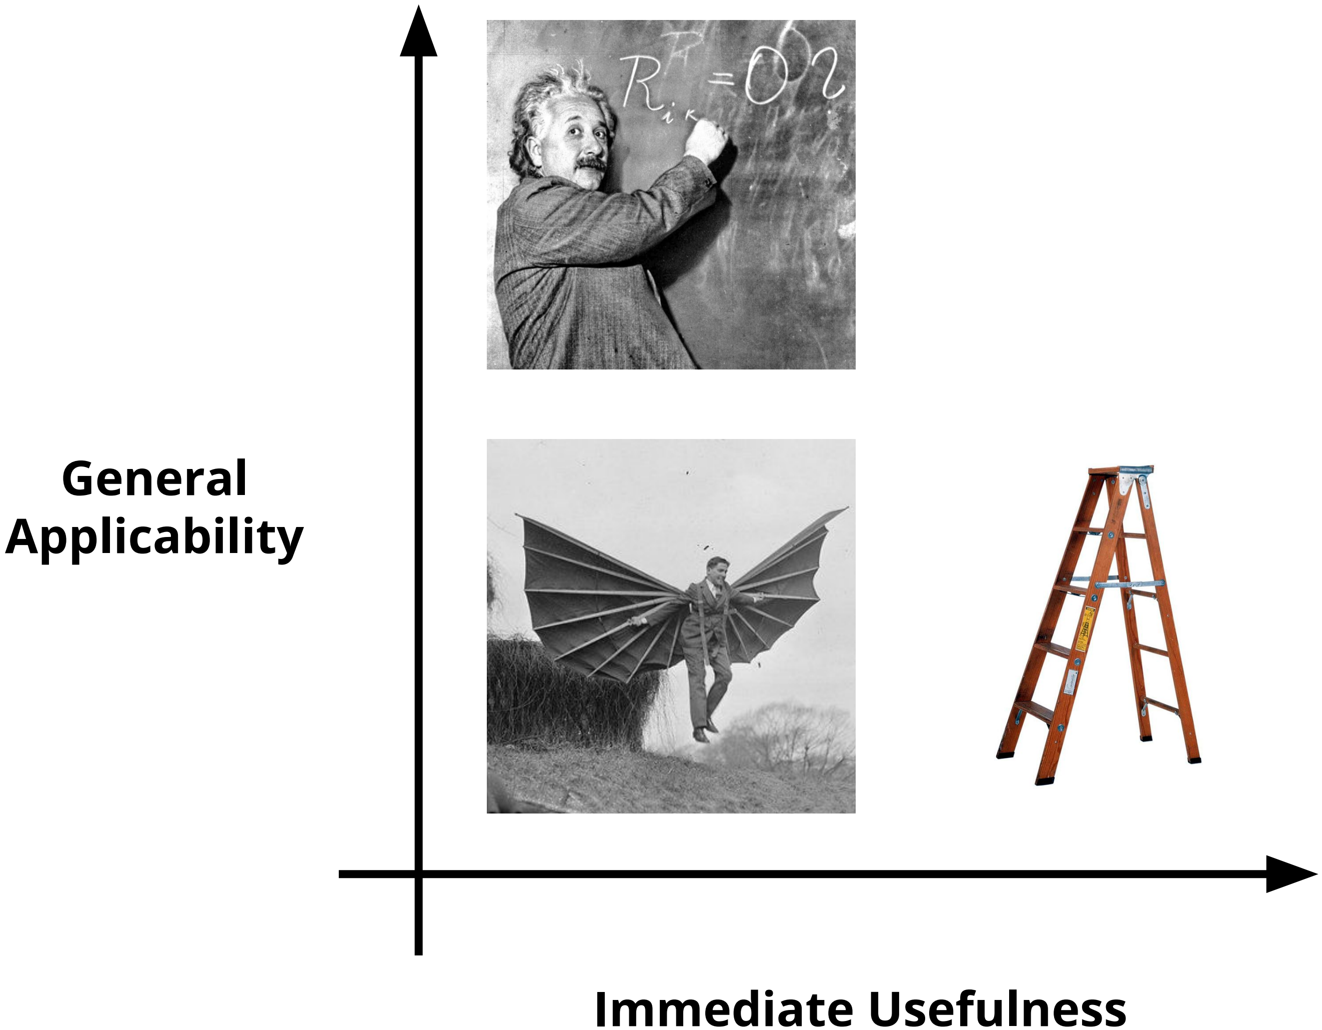
General Applicability
Immediate Usefulness

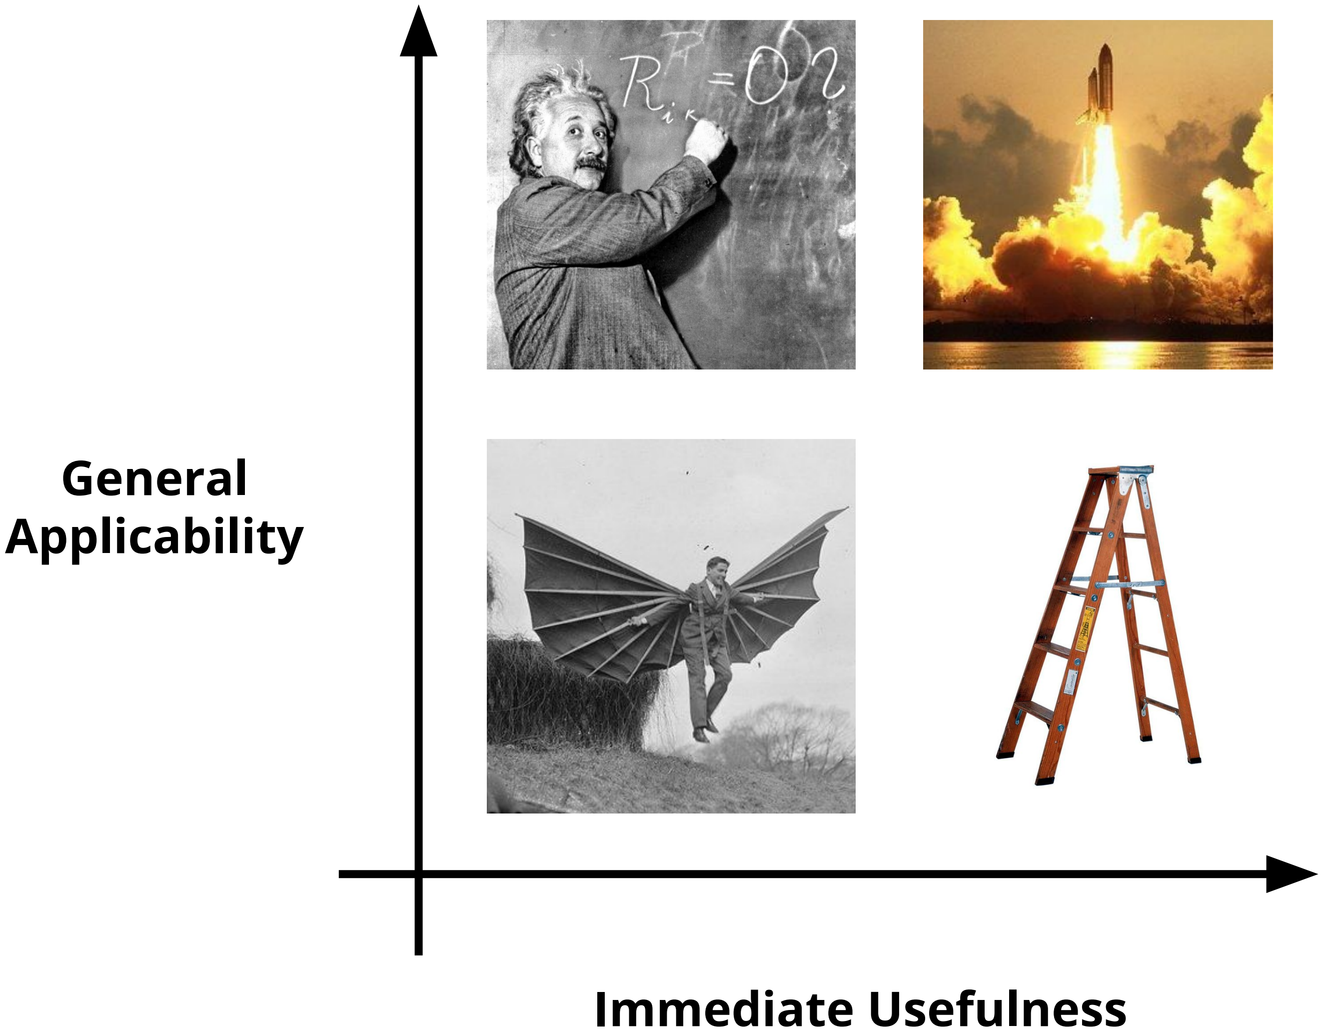
General Applicability
Immediate Usefulness

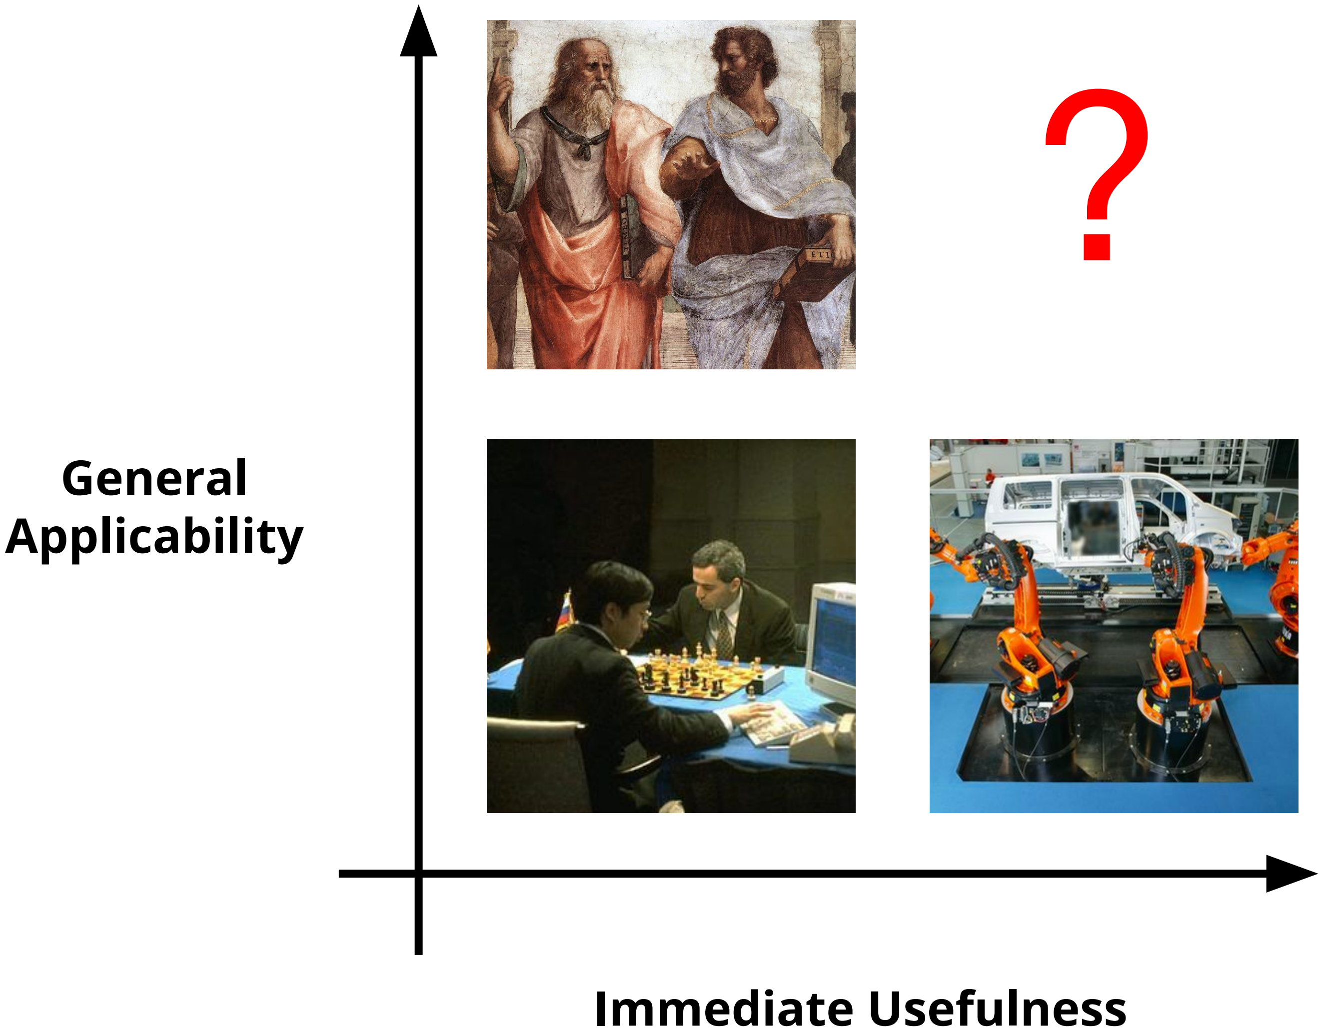
?
General Applicability
Immediate Usefulness

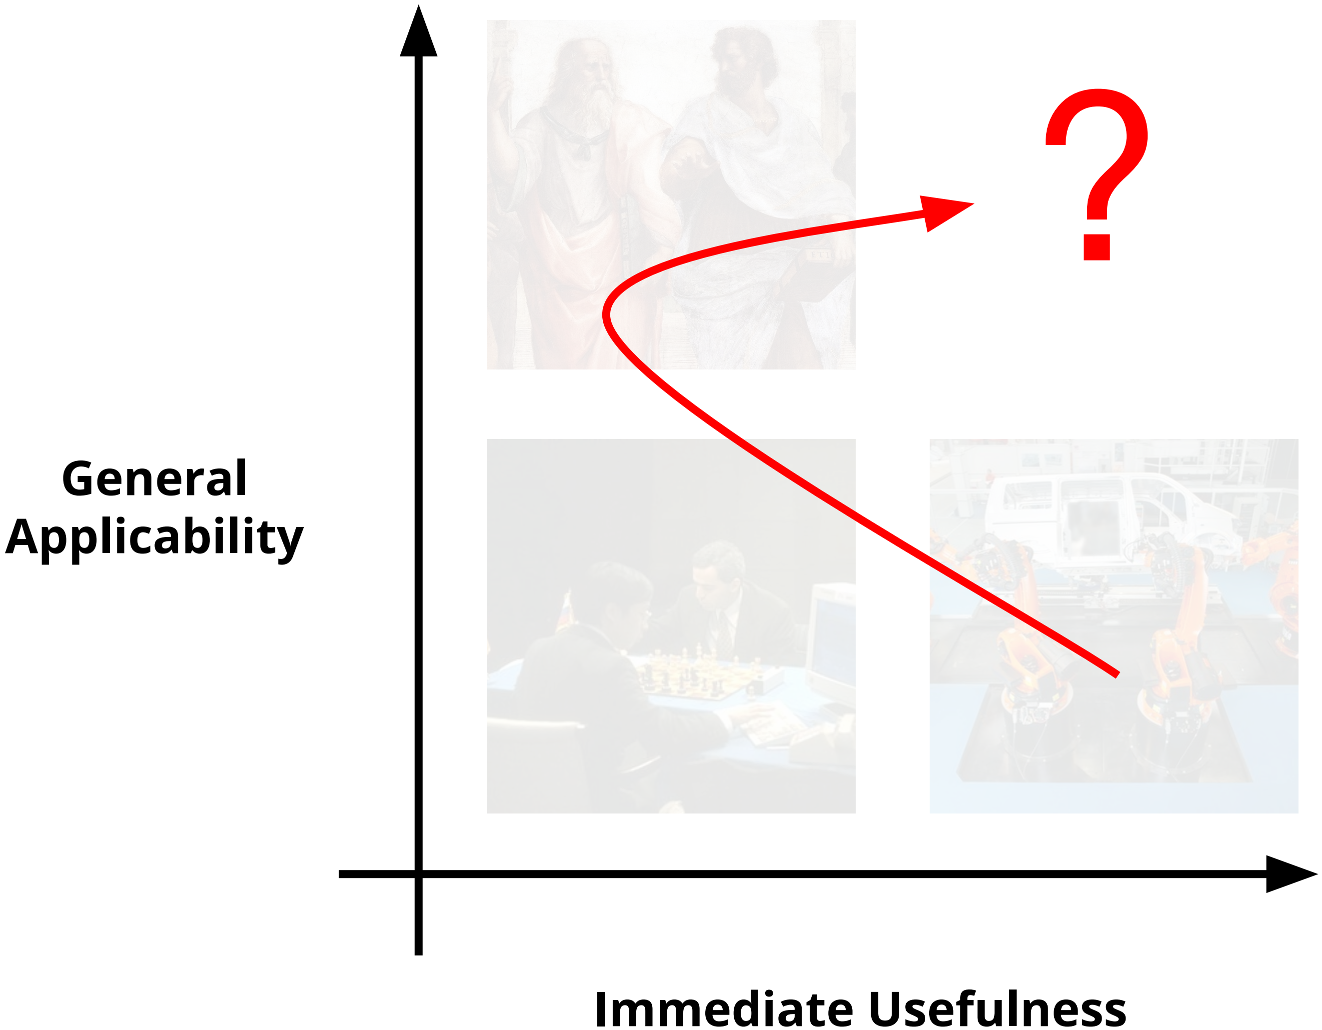
?
General Applicability
Immediate Usefulness

Deep Learning

Algorithm

?

Horse

Input

Computer

Output

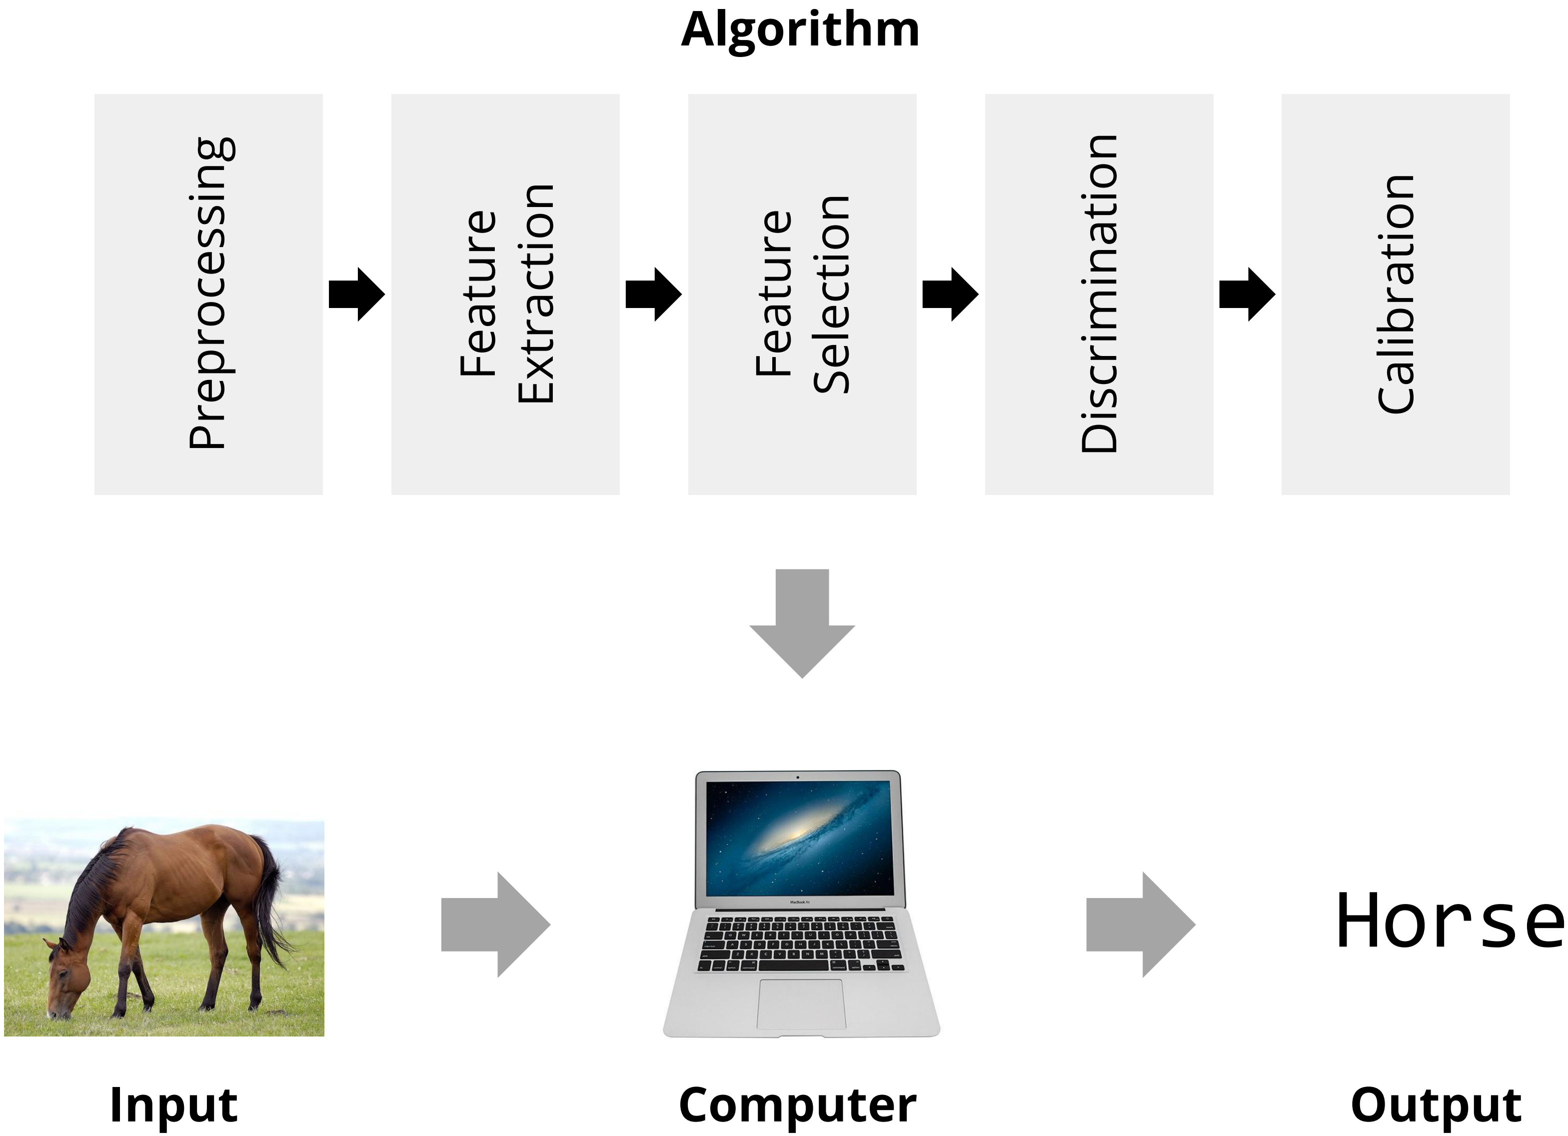
Algorithm
Preprocessing
Feature Extraction
Feature Selection
Discrimination
Calibration
Horse
Input
Computer
Output

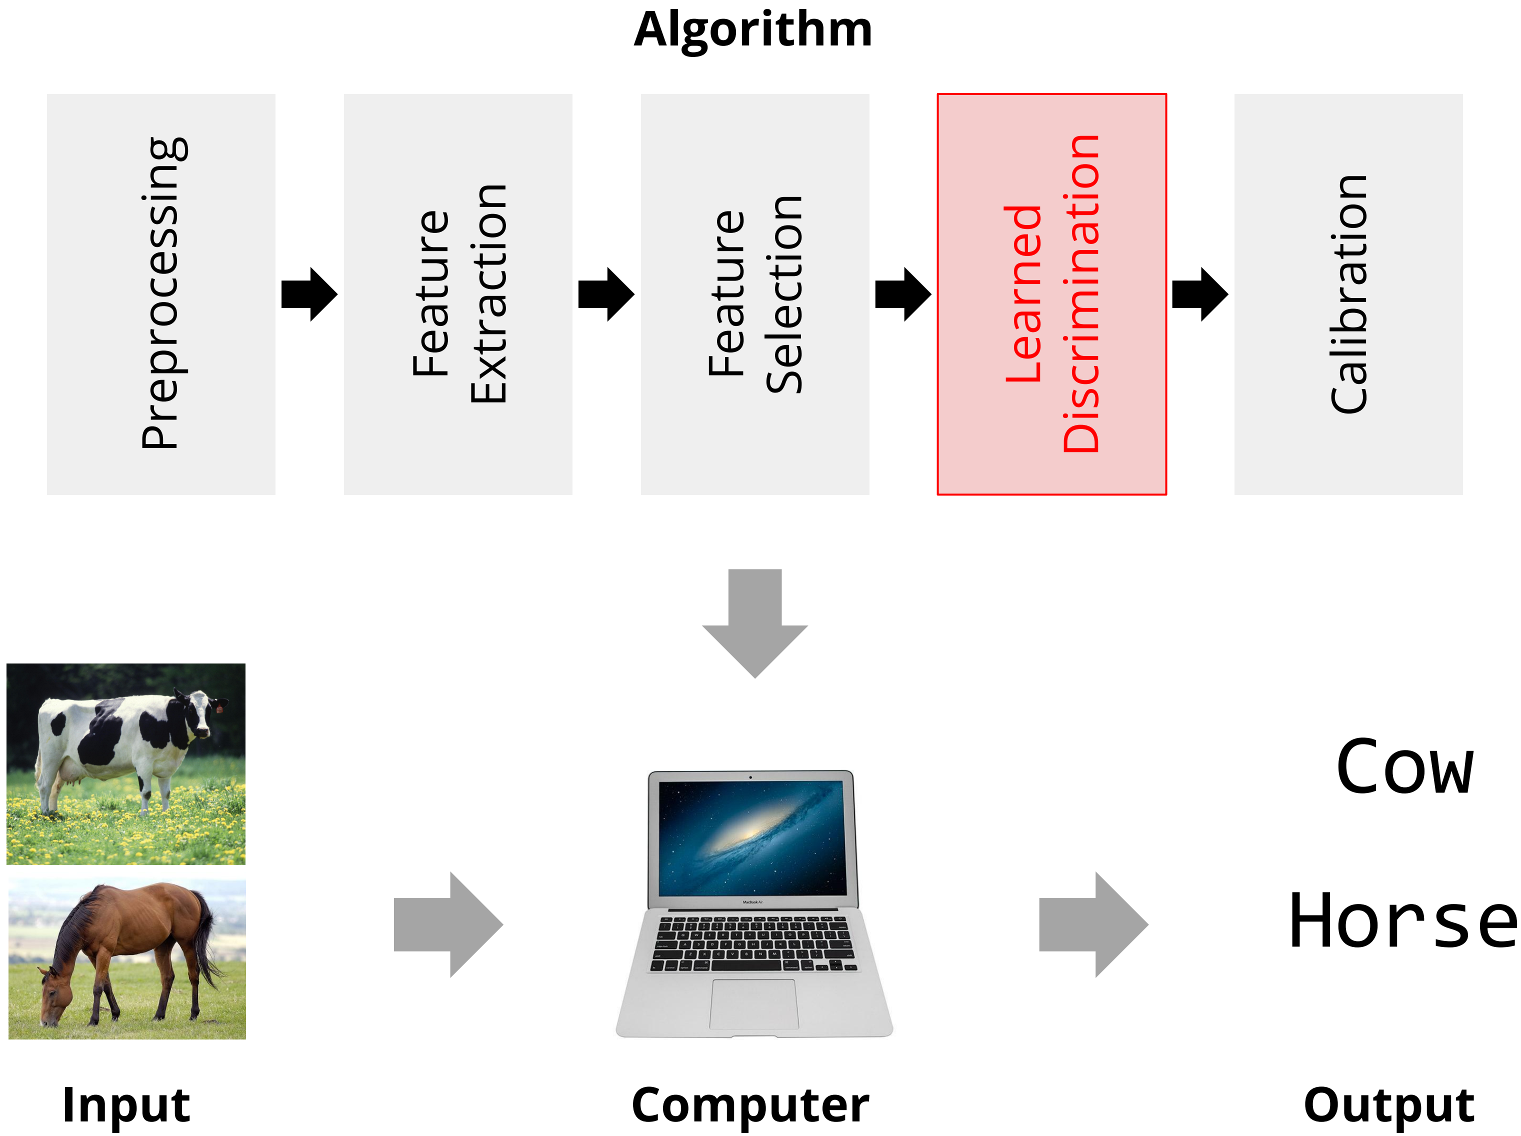
Algorithm
Preprocessing
Feature Extraction
Feature Selection
Learned Discrimination
Calibration
Cow
Horse
Input
Computer
Output

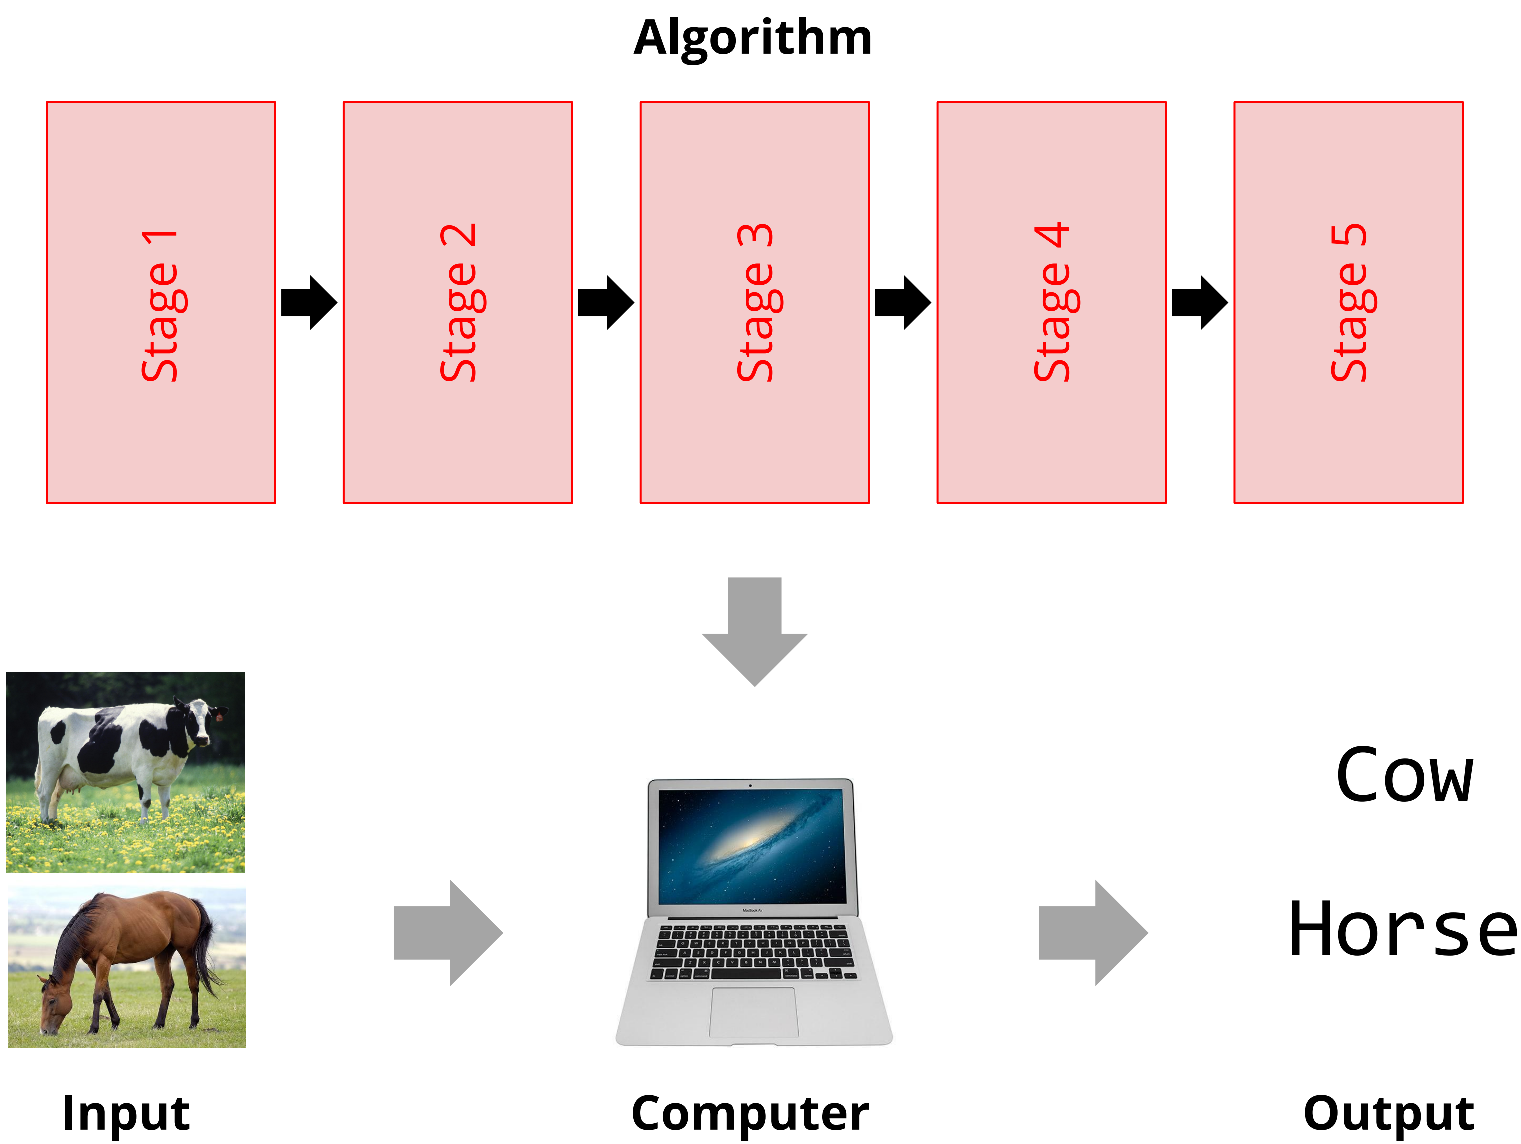
Algorithm
Stage 1
Stage 2
Stage 3
Stage 4
Stage 5
Cow
Horse
Input
Computer
Output

Deep Learning

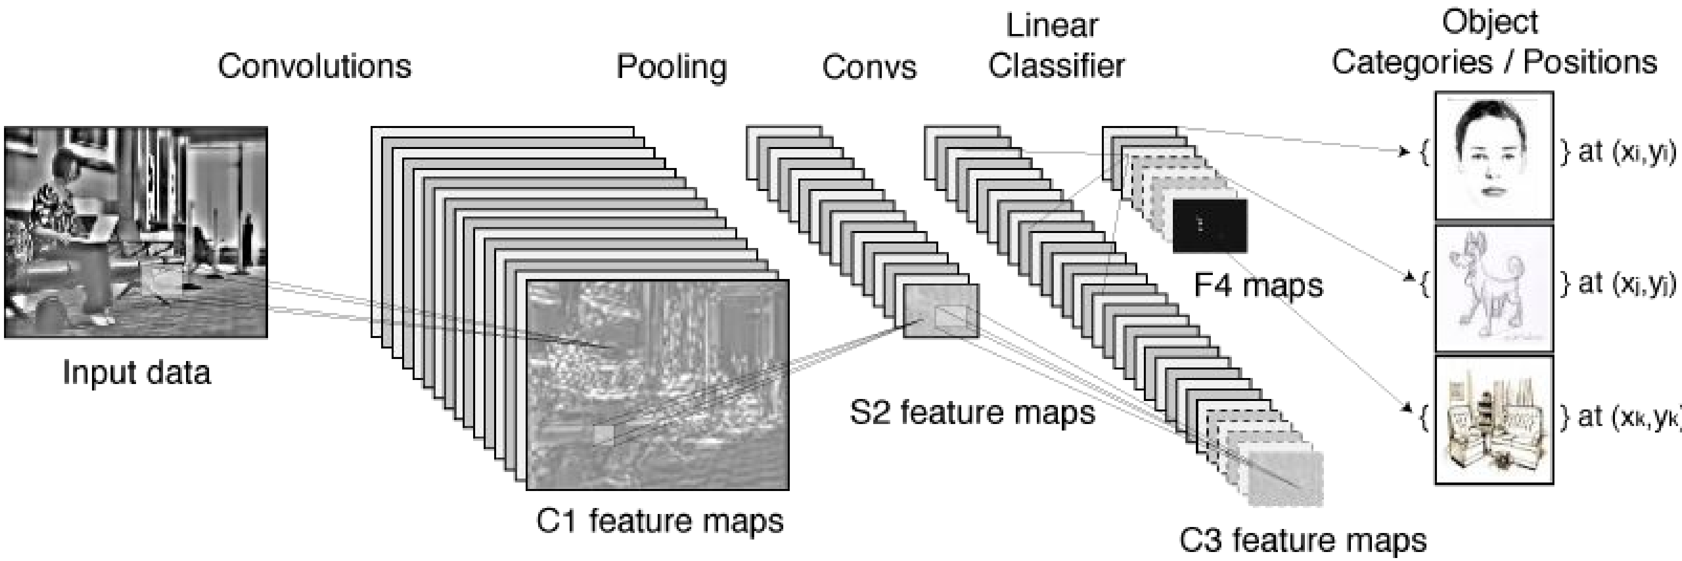

Convolutional neural networks for image classification
Torch (2015)

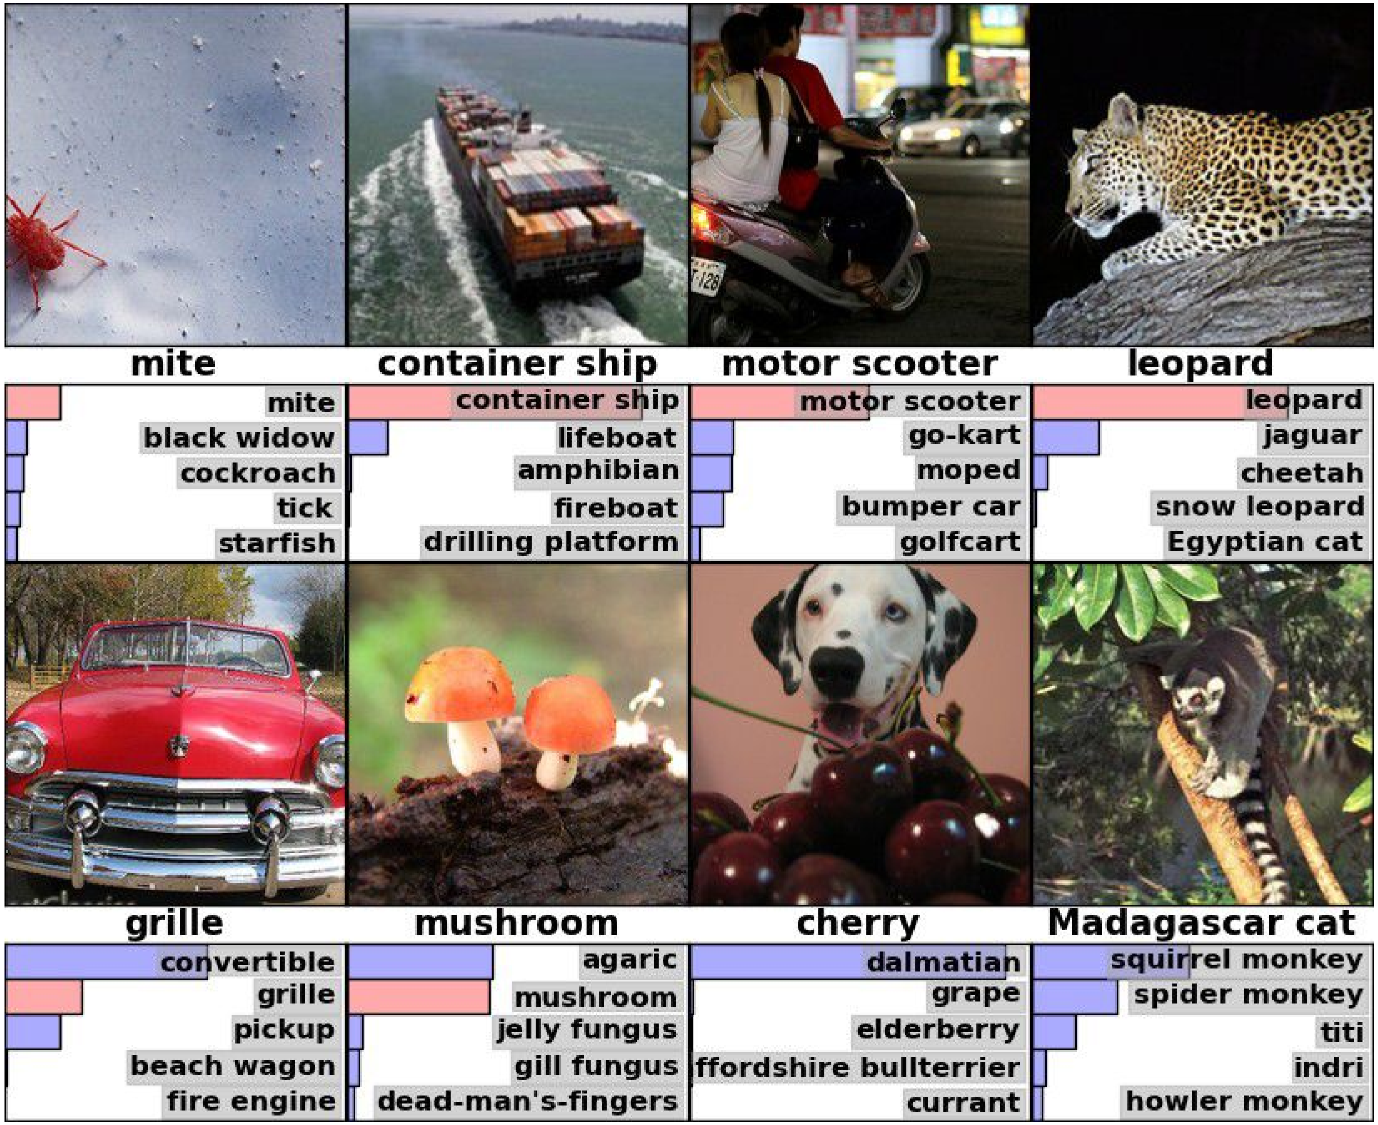

Krizhevsky et al. (2012)

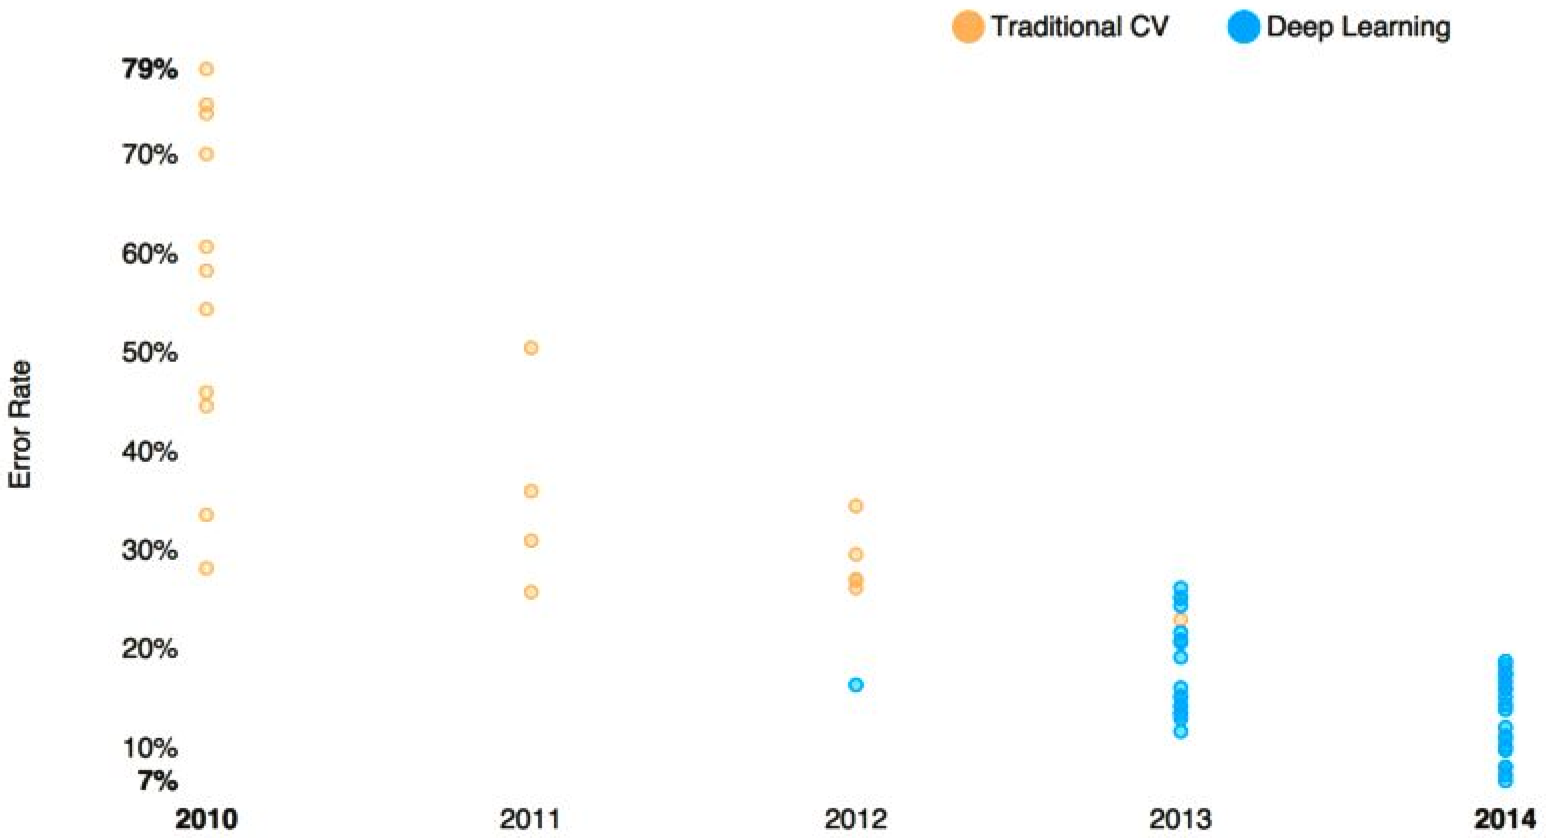

Clarifai (2014)

Where do the labels come from?

Reinforcement Learning

Architecture

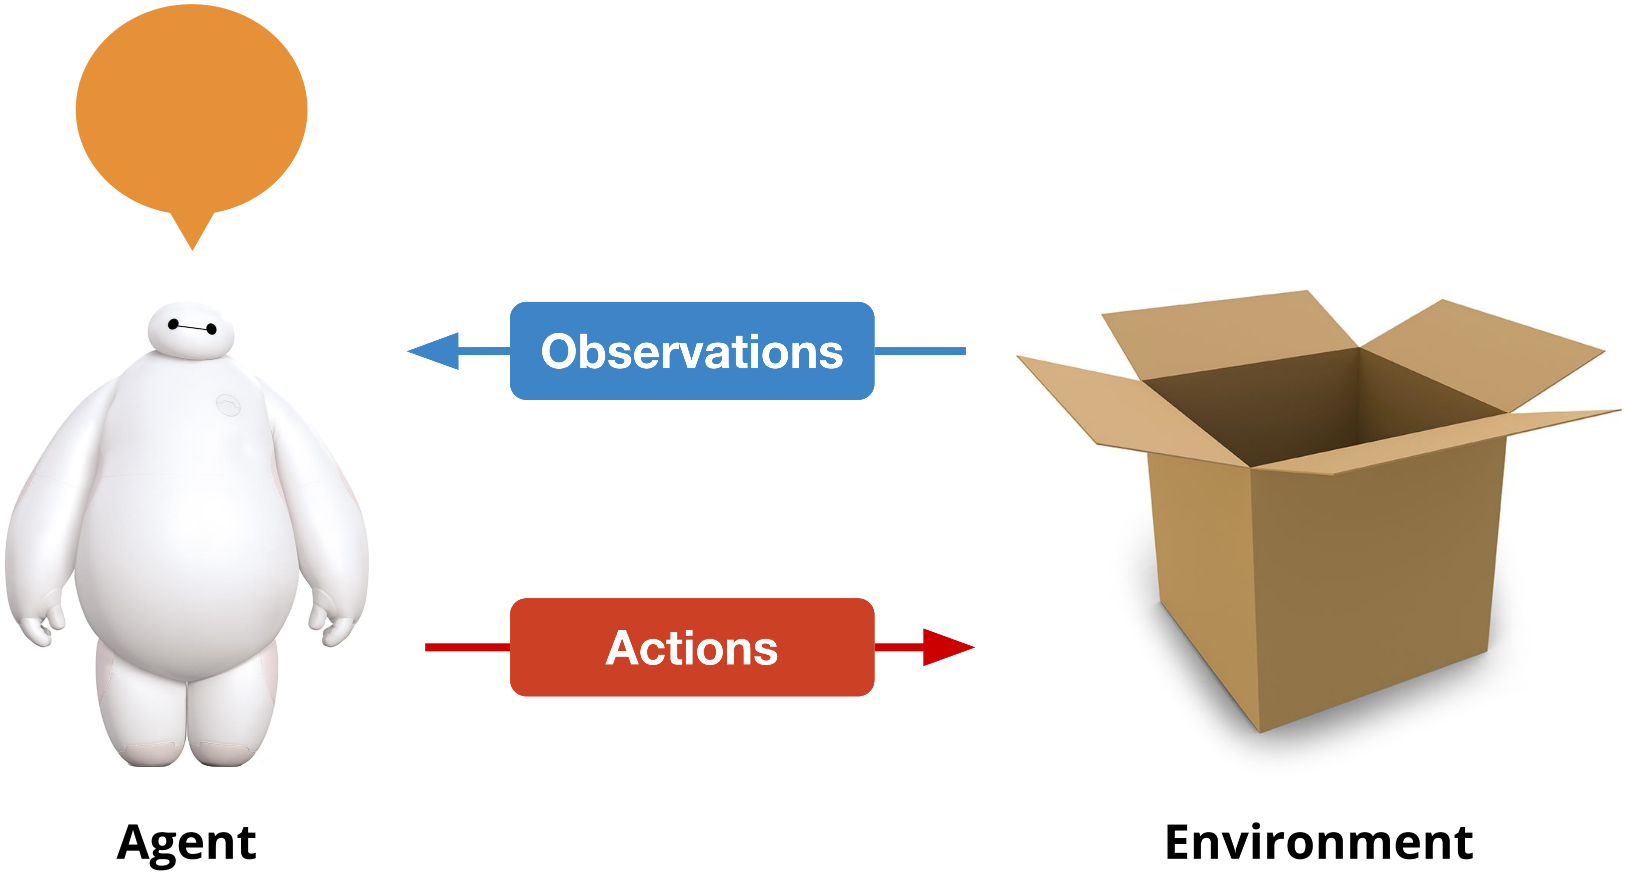

ATARI agents

100+ classic 8-bit Atari games

- 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

Human-level control through deep reinforcement learning

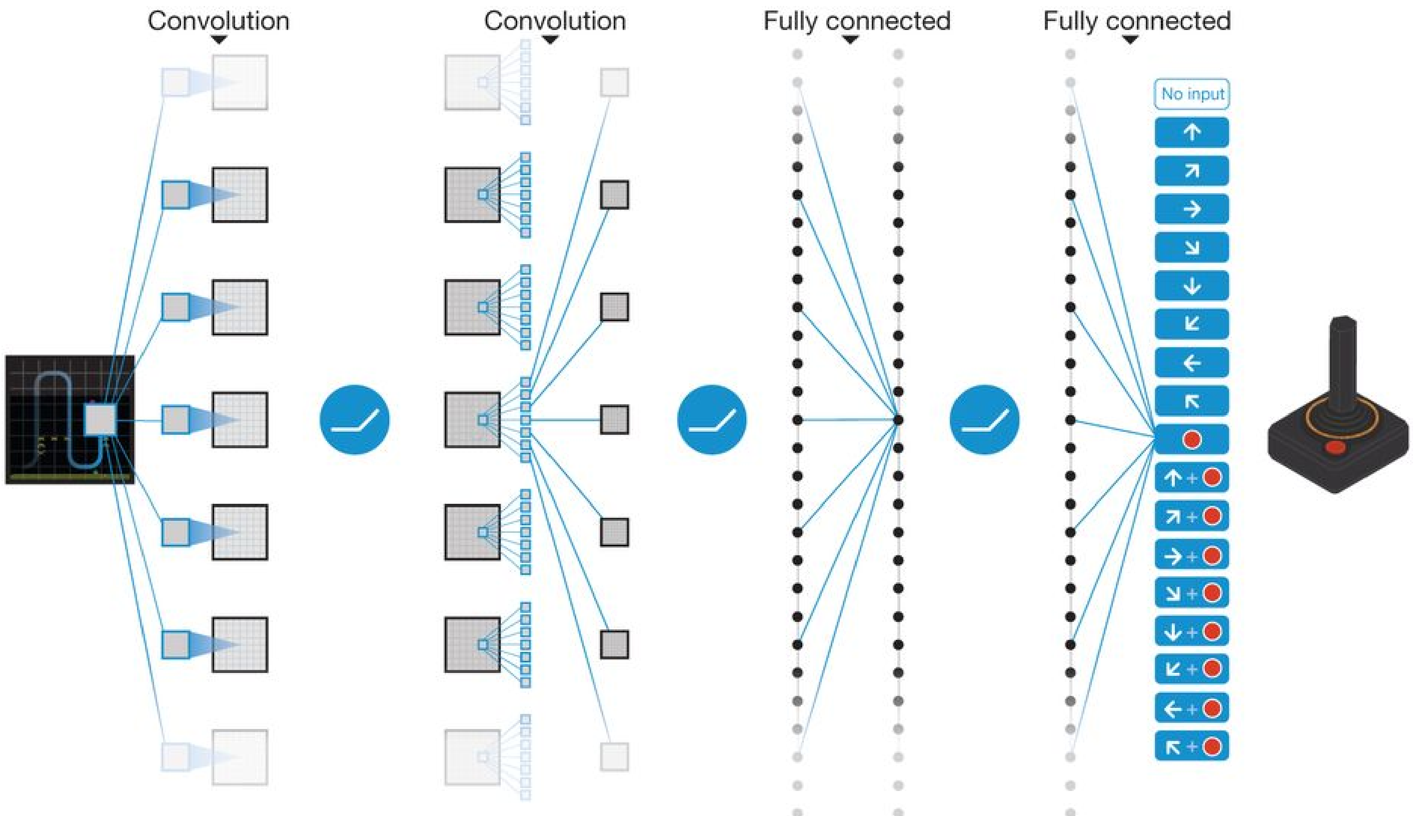

Mnih et al. (Nature, 2015)

Space Invaders agent

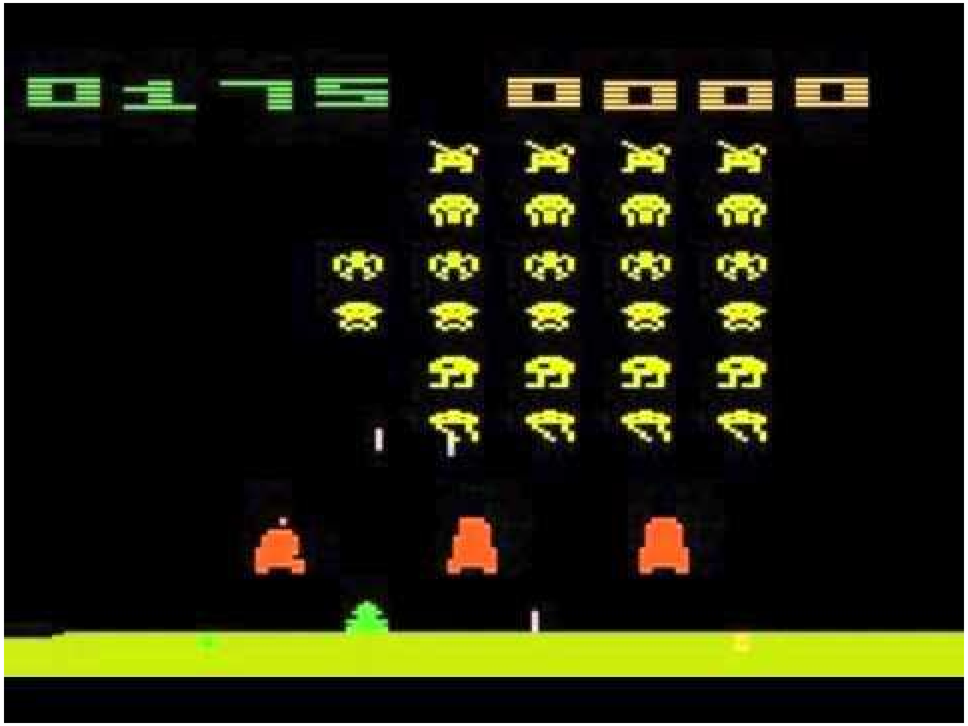

Breakout agent

General Atari agent

Human-level control through deep reinforcement learning

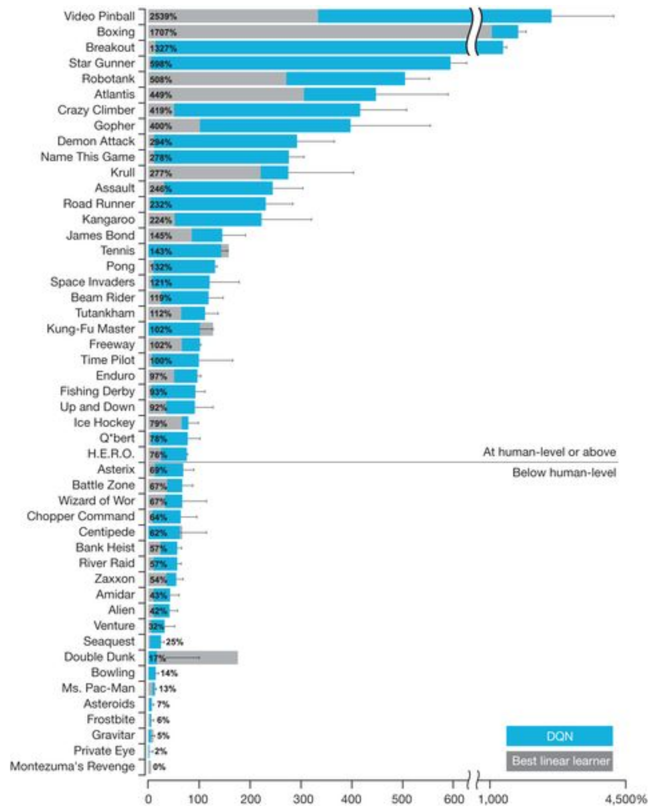

Mnih et al. (Nature, 2015)

Human-level control through deep reinforcement learning

Mnih et al. (Nature, 2015)

Deep Reinforcement Learning for Continuous Control

How many experiences do we need?

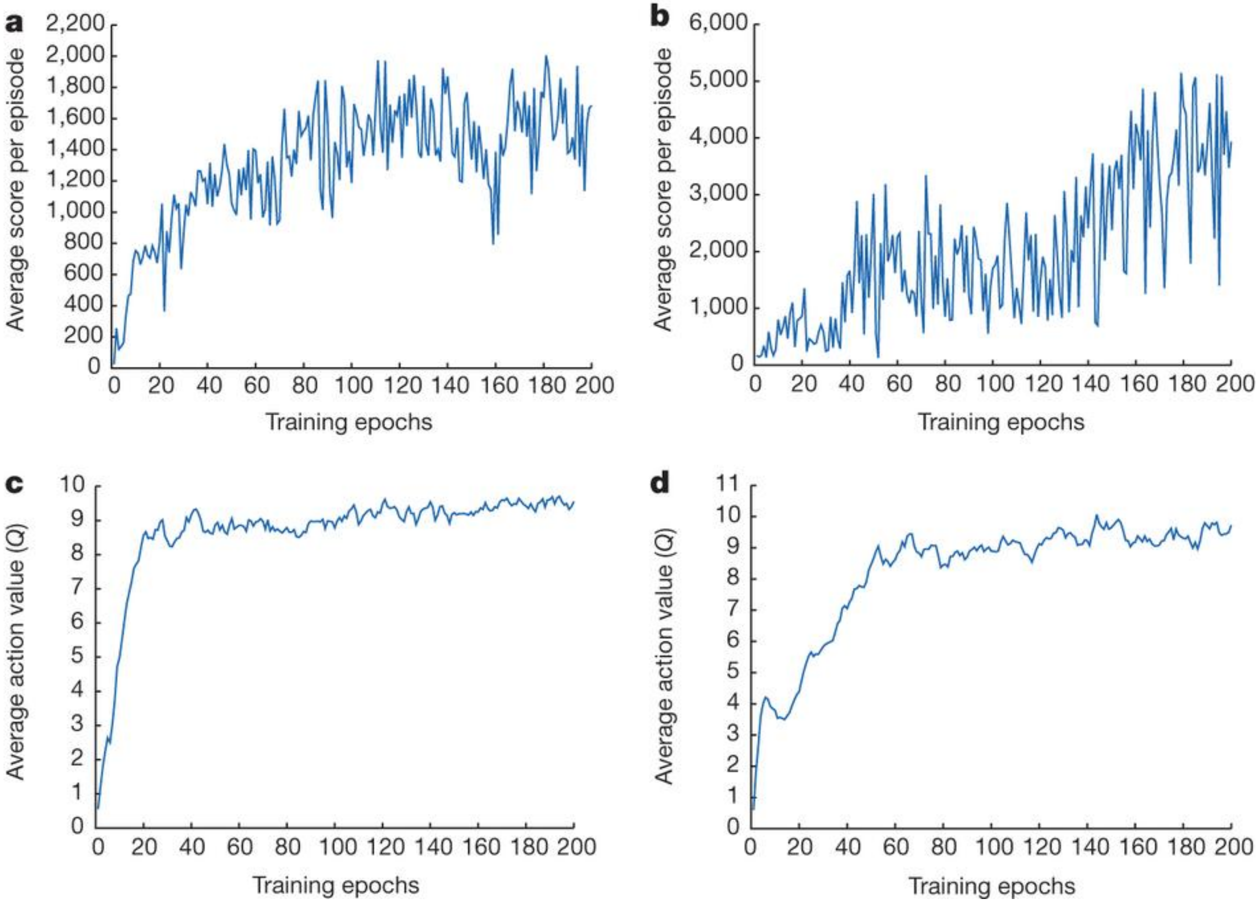

1 epoch = 50,000 interactions = 30 minutes of experience
Total experience: 10m interactions = 5 days

Model-based Methods

Three learning paradigms

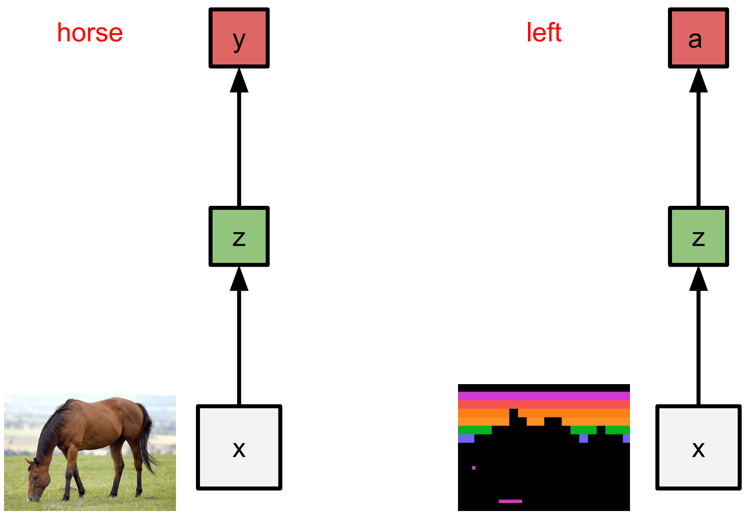

Supervised Learning

Reinforcement Learning

Three learning paradigms

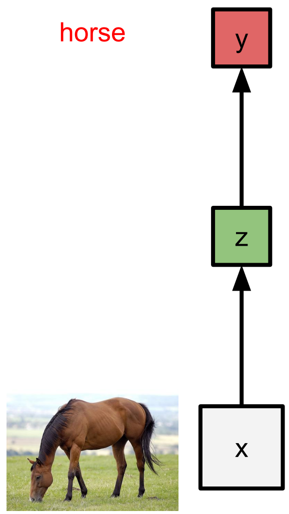

Supervised Learning

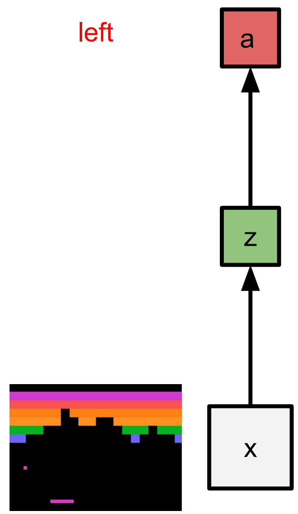

Reinforcement Learning

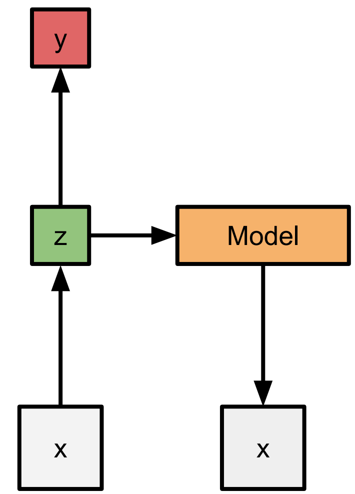

Generative Modelling

Three learning paradigms

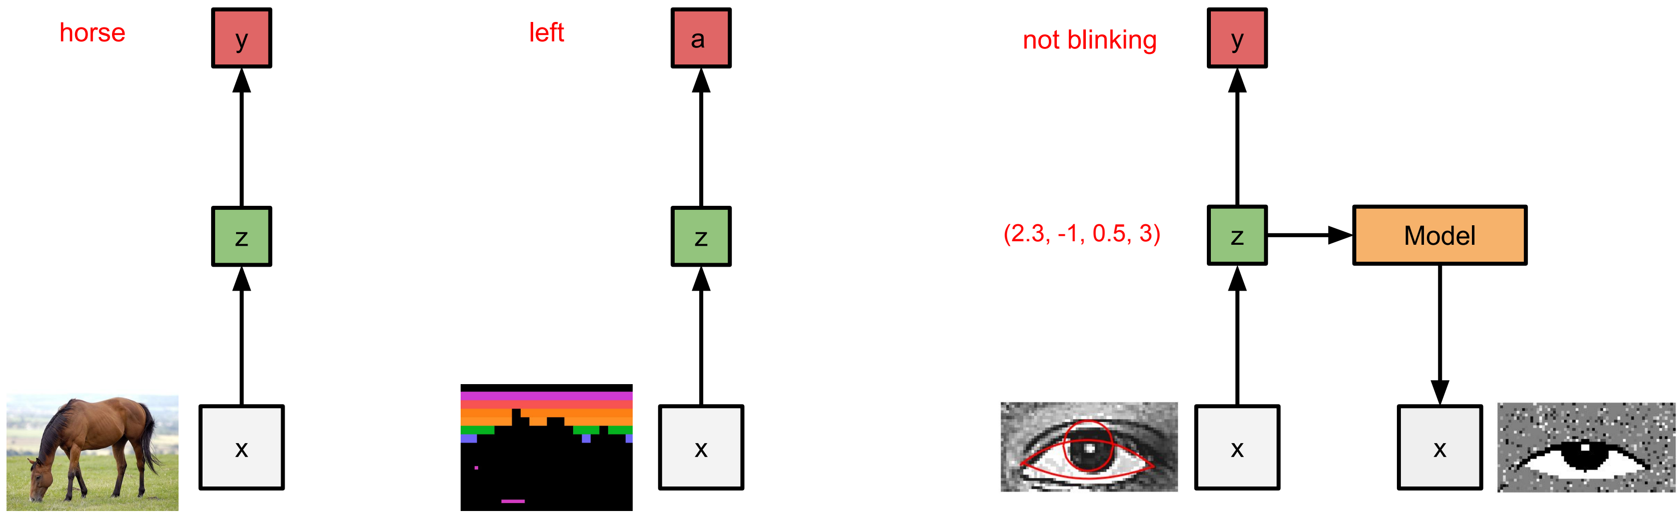

Supervised Learning

Reinforcement Learning

Generative Modelling

Learning to draw shapes

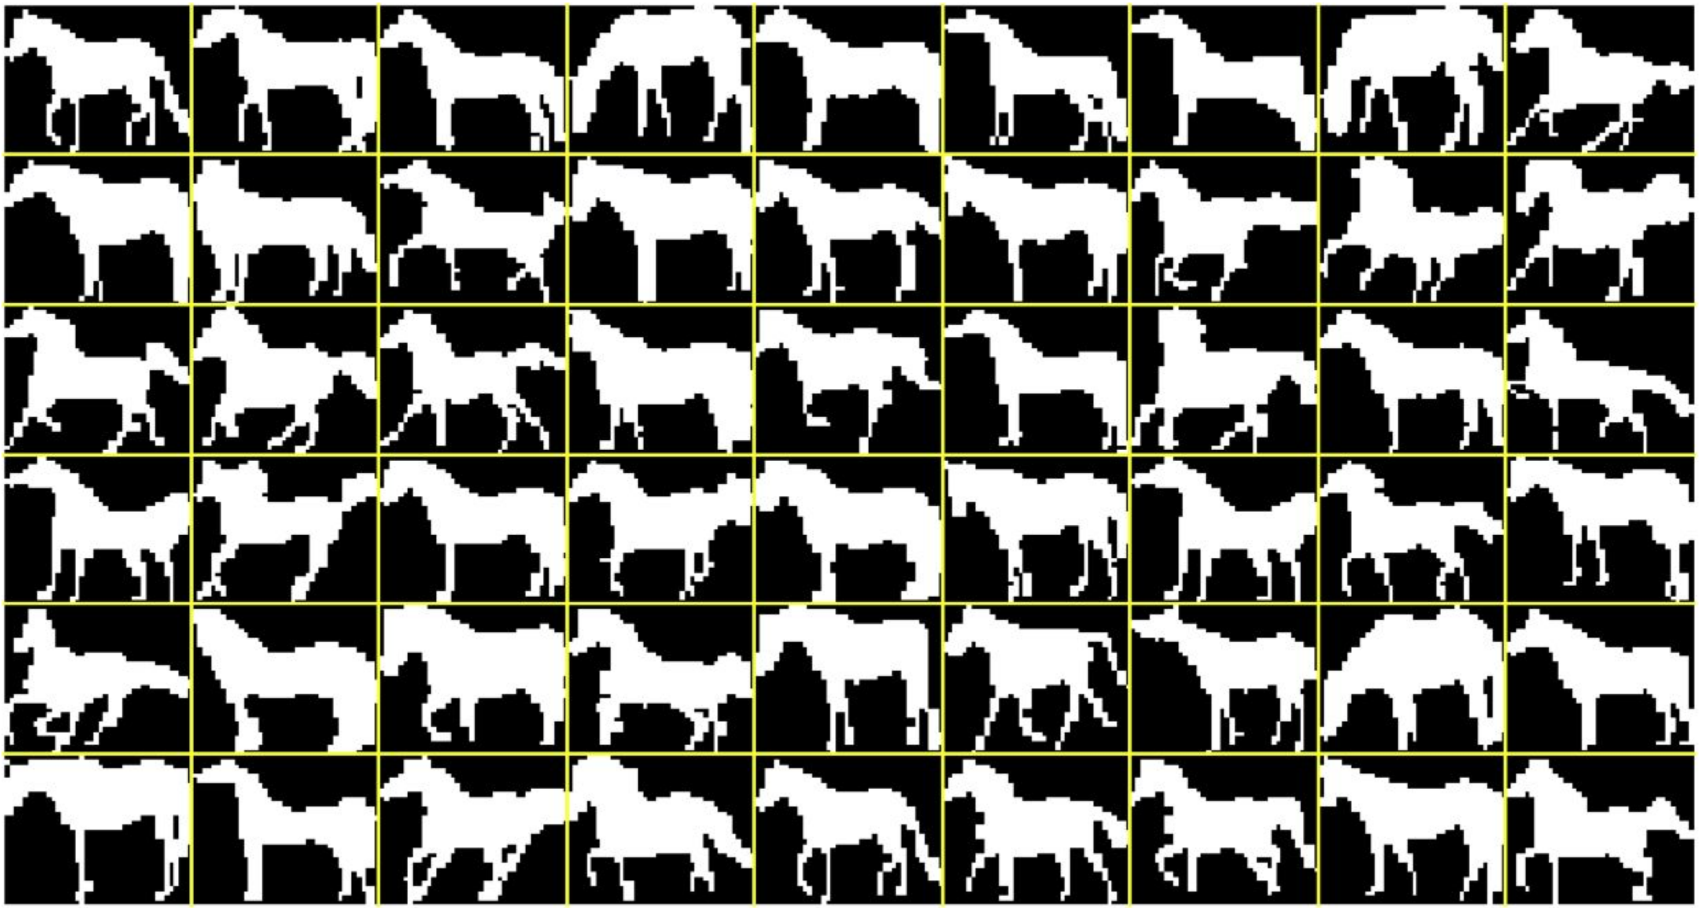

The Shape Boltzmann Machine

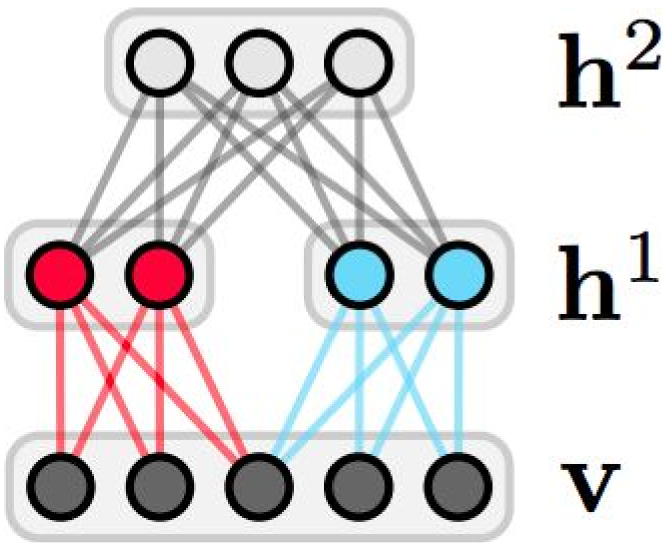

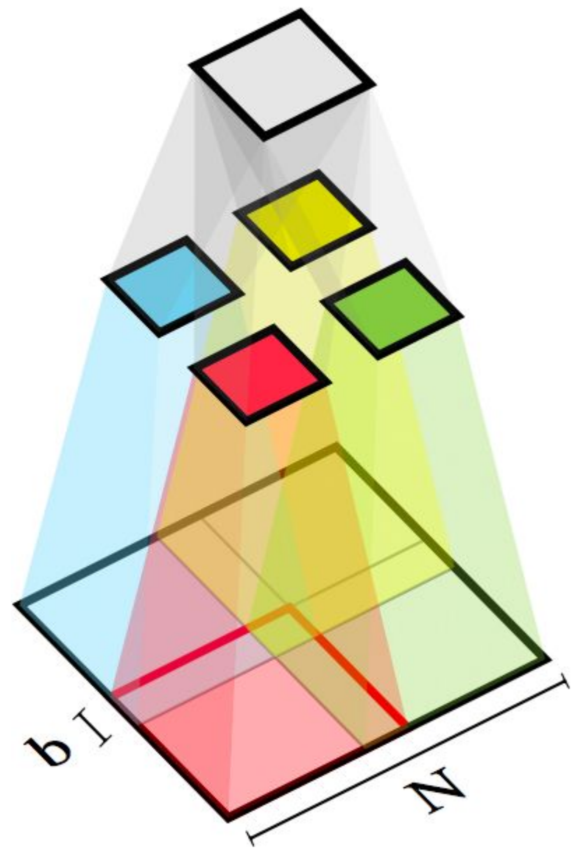

Sampling from an SBM

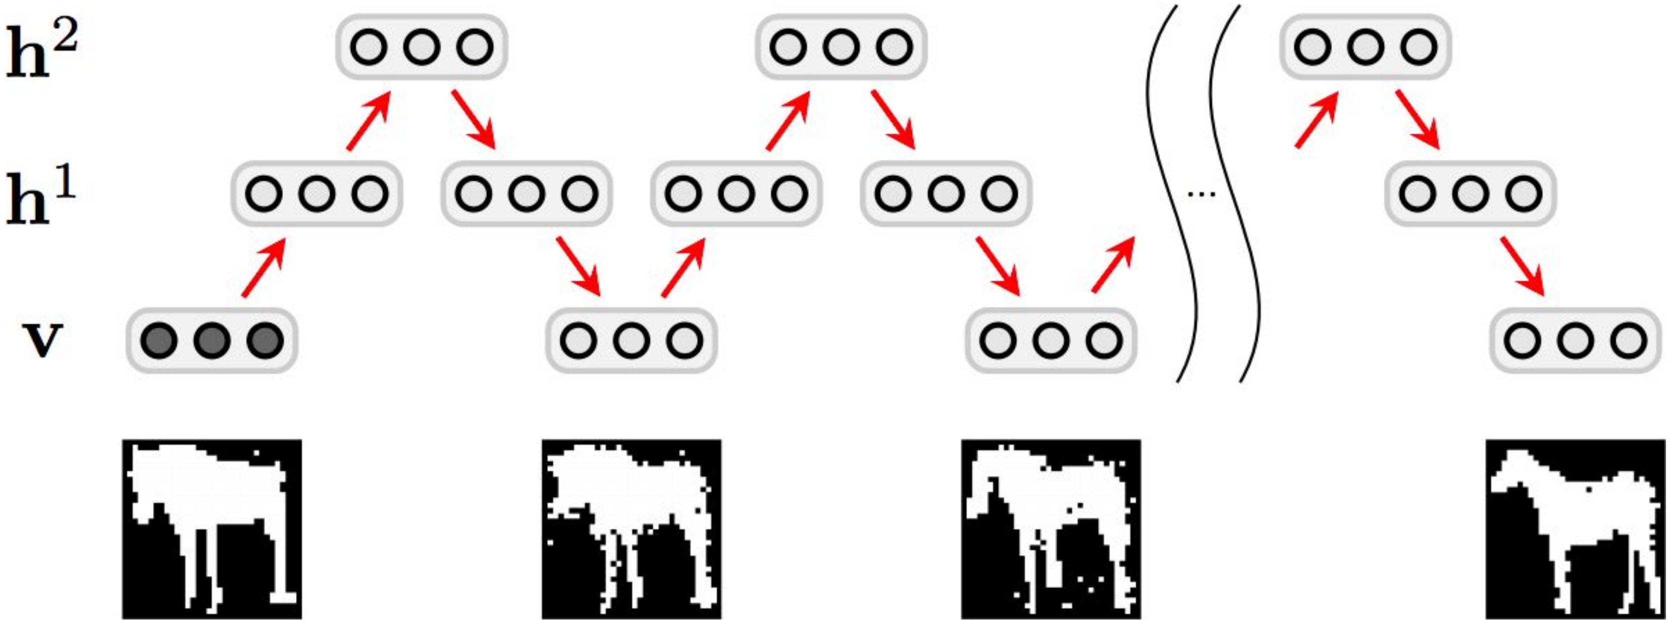

Learning to draw shapes

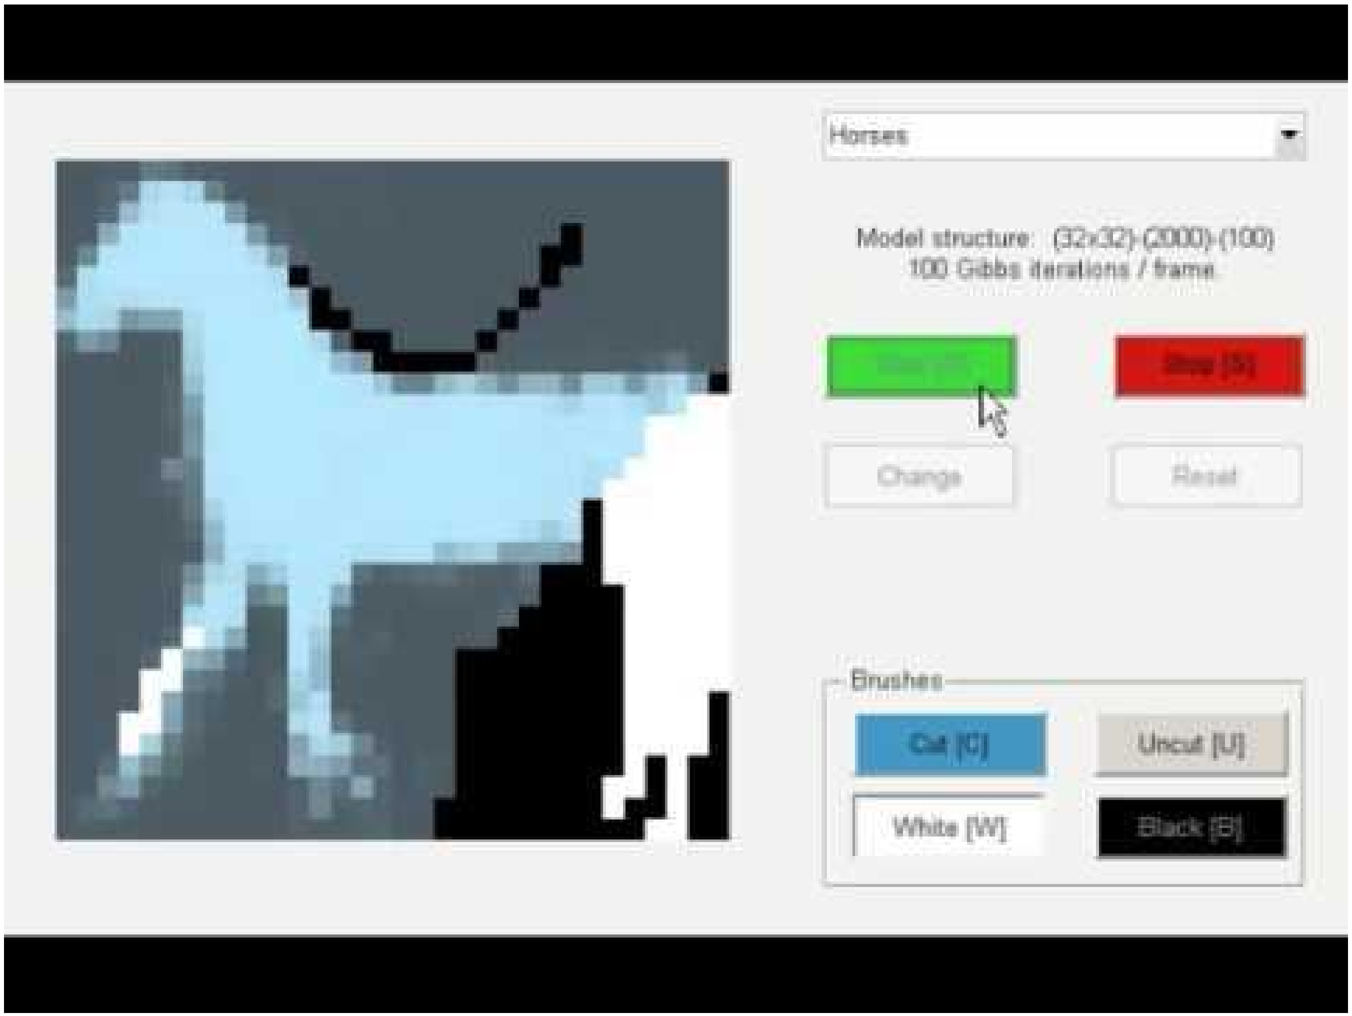

Learning to draw shapes

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 (**IJCV, 2013**)

Factored Shapes and Appearances

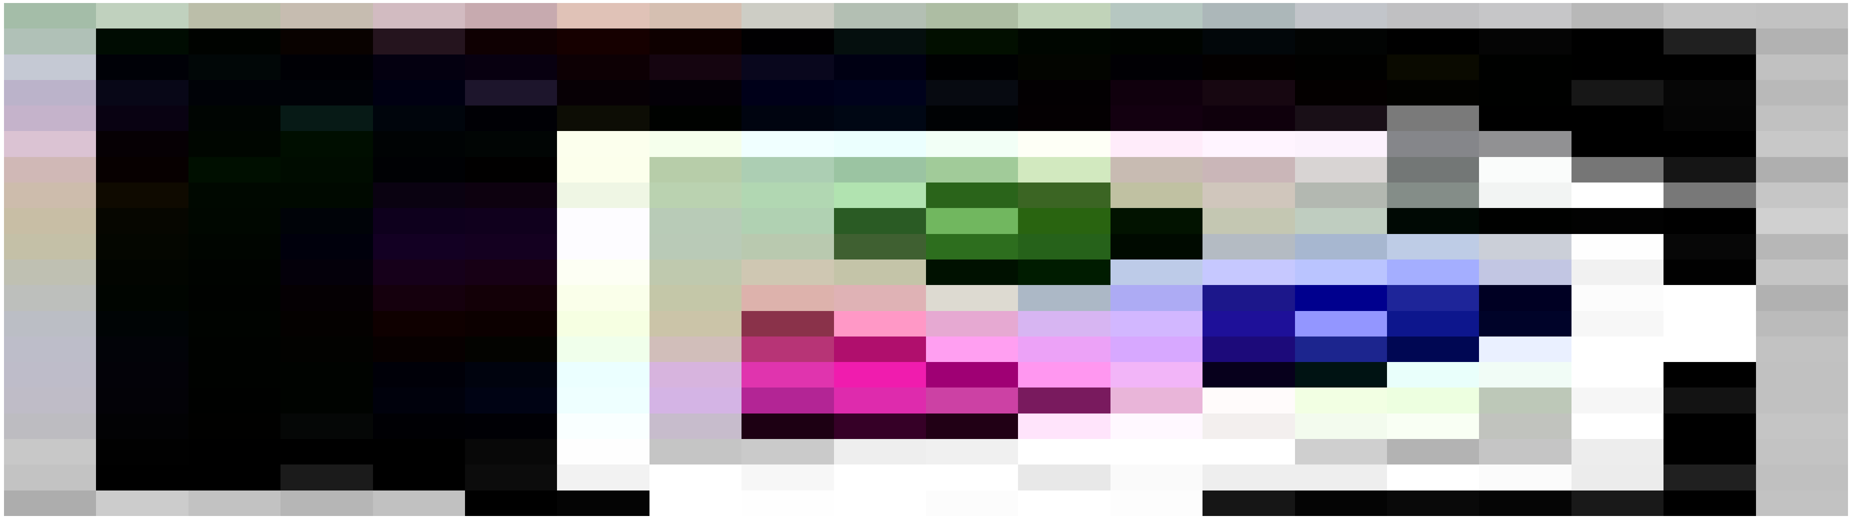

A Generative Model for Parts-based Object Segmentation

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

Neural Information Processing Systems (**NIPS, 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 (**NIPS, 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 (**NIPS, 2012**)

Recurrent Neural Networks for Image Generation

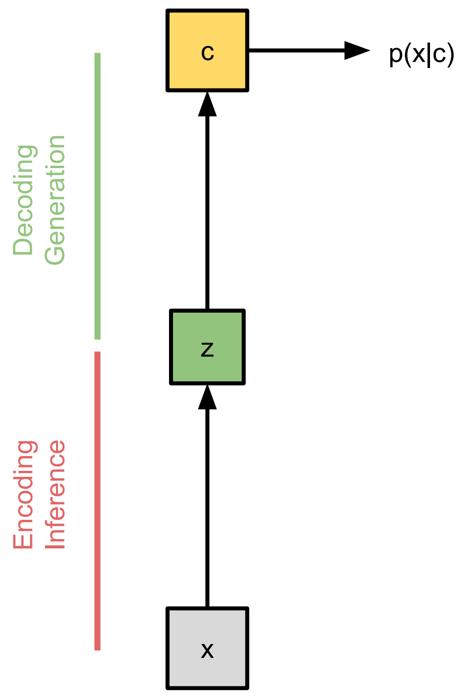

Gregor et al. (2015)

Recurrent Neural Networks for Image Generation

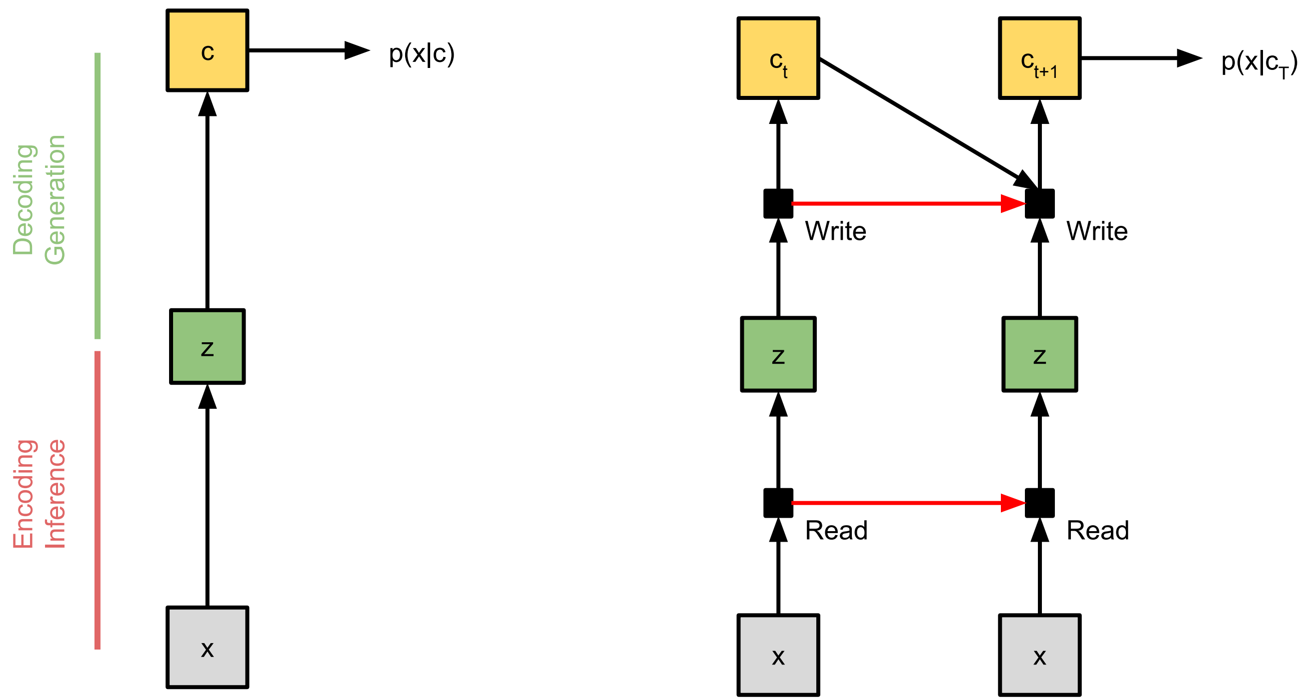

Gregor et al. (2015)

Recurrent Neural Networks for Image Generation

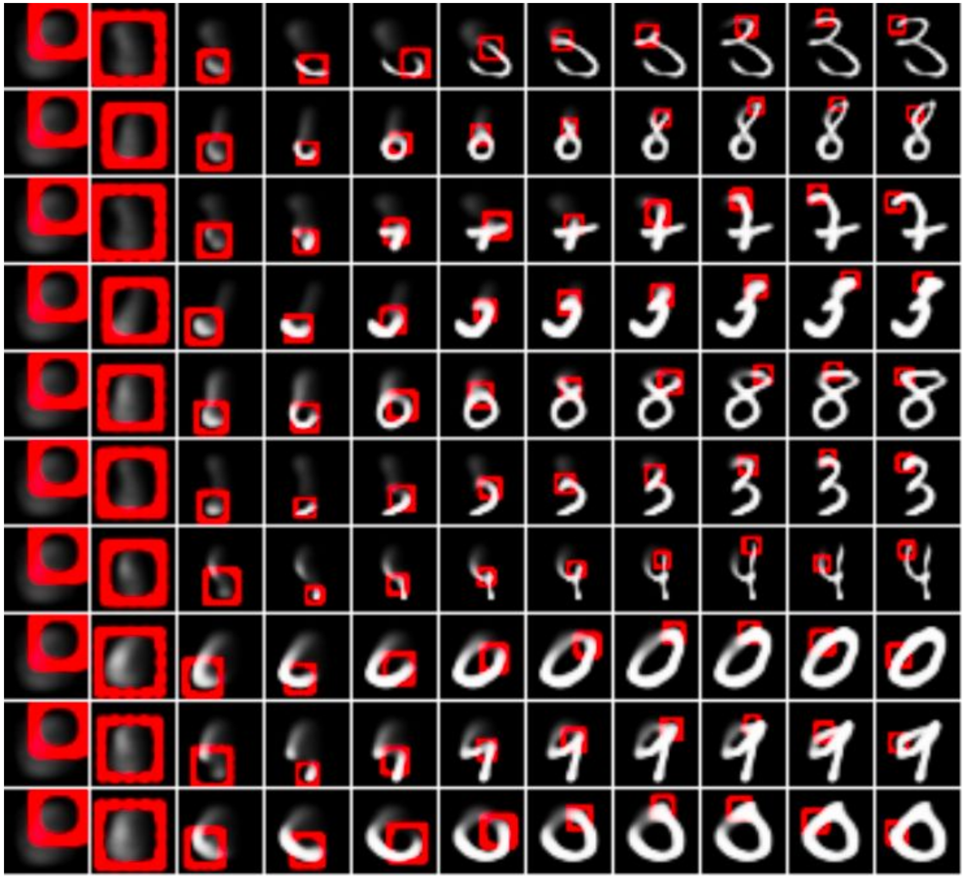

Gregor et al. (2015)

Recurrent Neural Networks for Image Generation

Gregor et al. (2015)

Generative Modelling

Model **p(x|z)** can be:

- Fully learned (e.g. autoencoders)
- Partially specified
- Fully specified (e.g. renderers)

What models should we use?

Choice of model **p(x|z)**
almost always constrained
by our ability to compute **p(z|x)**

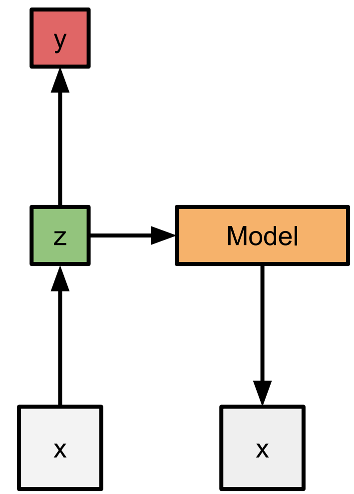

How should we do inference?

Reinforced Variational Inference

Modern Variational Inference

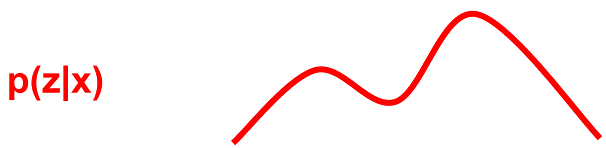

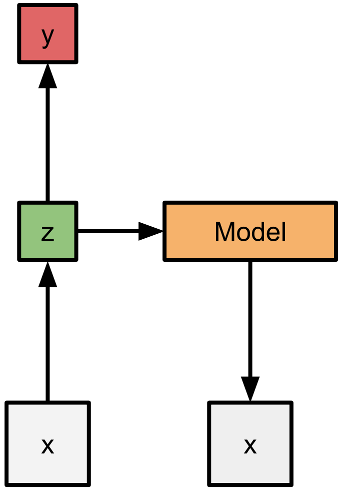

Consensus Message Passing for Layered Graphical Models
Varun Jampani, S. M. Ali Eslami, Daniel Tarlow Pushmeet Kohli, John Winn
Artificial Intelligence and Statistics (**AISTATS**, **2015**)

Modern Variational Inference

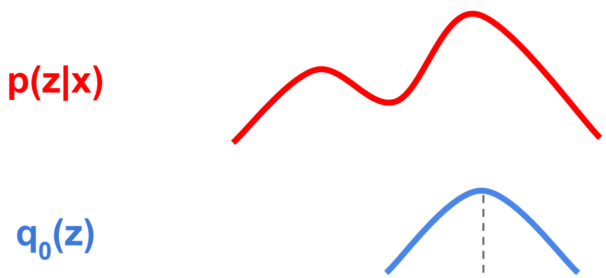

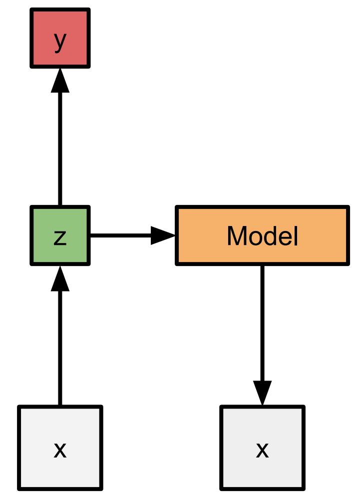

Consensus Message Passing for Layered Graphical Models

Varun Jampani, S. M. Ali Eslami, Daniel Tarlow Pushmeet Kohli, John Winn

Artificial Intelligence and Statistics (**AISTATS**, **2015**)

Modern Variational Inference

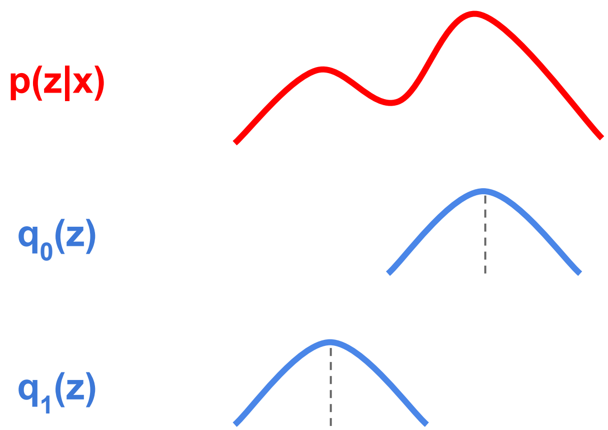

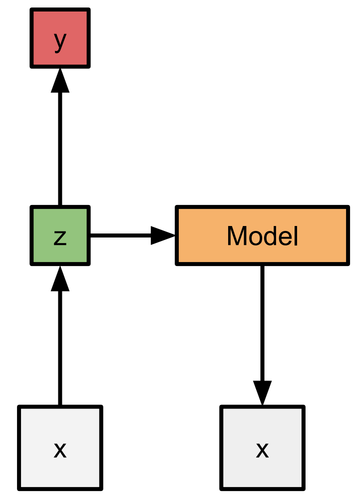

Consensus Message Passing for Layered Graphical Models

Varun Jampani, S. M. Ali Eslami, Daniel Tarlow Pushmeet Kohli, John Winn

Artificial Intelligence and Statistics (**AISTATS**, **2015**)

Modern Variational Inference

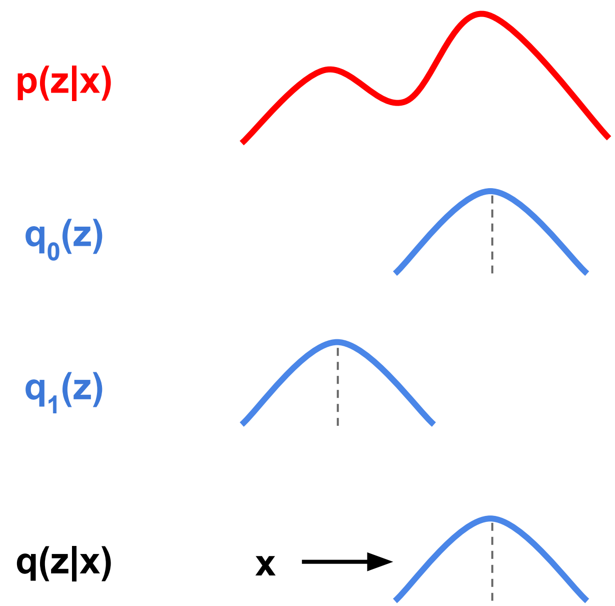

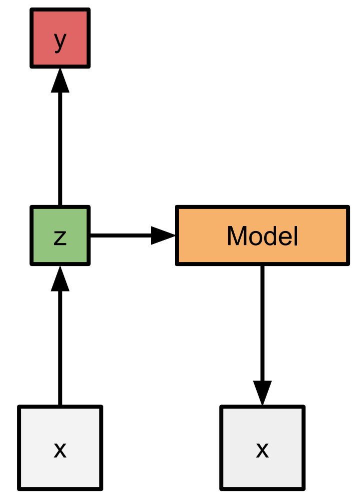

Consensus Message Passing for Layered Graphical Models

Varun Jampani, S. M. Ali Eslami, Daniel Tarlow Pushmeet Kohli, John Winn

Artificial Intelligence and Statistics (**AISTATS**, **2015**)

Modern Variational Inference

- Approximate **p(z|x)** using **q(z|x)**
- Parameterise **q(z|x)**
- Minimise KL[**q(z|x)** | **p(z|x)**]
- Samples from **q(z|x)** can be used as codes representing the image **x**

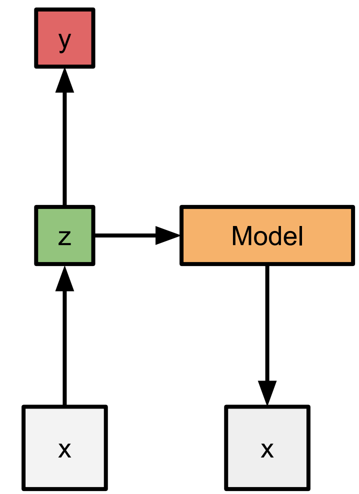

Consensus Message Passing for Layered Graphical Models
Varun Jampani, S. M. Ali Eslami, Daniel Tarlow Pushmeet Kohli, John Winn
Artificial Intelligence and Statistics (**AISTATS**, **2015**)

Minimising the KL

- Minimise KL[**q(z|x)** | **p(z|x)**]

- Maximise $L(q) = E_q [\log p(x, z) - \log q(z|x)]$

  - Potentially high variance
  - Can require knowledge of $\nabla p(x, z)$

- RL objective to maximise $J(p) = E_p [\sum_t r(s_t, a_t)]$

- Connection between VI and RL hinted at by many (e.g. VAE, DLGM, NVIL, etc.)

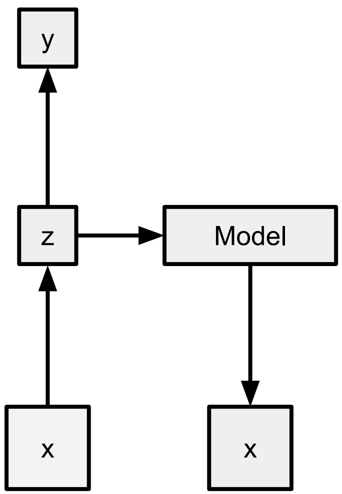

Variational Inference as Reinforcement Learning

maximise $\int p_\theta(y) f(y) \mathrm{d}y$

Generic expectation		RL		VI	
Optimization var.	θ	Policy param.	θ	Variational param.	θ
Integration var.	y	Trajectory	τ	Latent trace	z
Distribution	$p_\theta(y)$	Trajectory dist.	$p_\theta(\tau)$	Posterior dist.	$q_\theta(z\|x)$
Integrand	$f(y)$	Total return	$R(\tau)$	Free energy	$\log\left(\frac{p(x,z)}{q_\theta(z\|x)}\right)$

Reinforced Variational Inference

Theophane Weber, Nicolas Heess, S. M. Ali Eslami, John Schulman, David Wingate, David Silver. Neural Information Processing Systems, Workshop on Advances in Approximate Bayesian Inference (**NIPS, 2015**)

Summary

Prediction as a subset of **inference**

Inference as a **reinforcement learning** problem

Reinforcement learning as a **deep learning** problem

http://arkitus.com

ali@arkitus.com