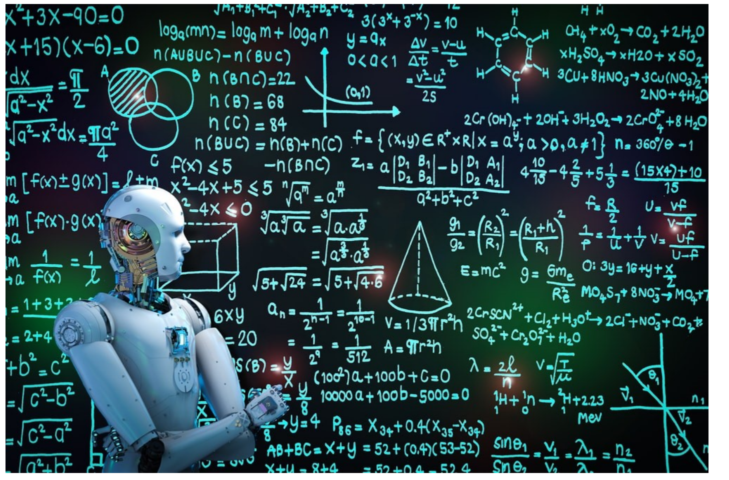

TDDE25 Seminar 15

Artificial Intelligence

Presentation originally by Patrick Doherty (modified)
Department of Computer and Information Science
Artificial Intelligence and Integrated Computer Systems Division
Linköping University

What is Artificial Intelligence?

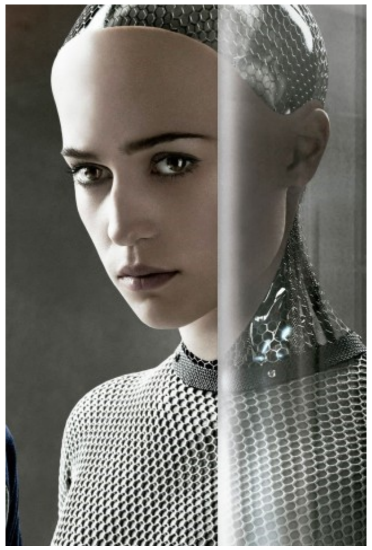
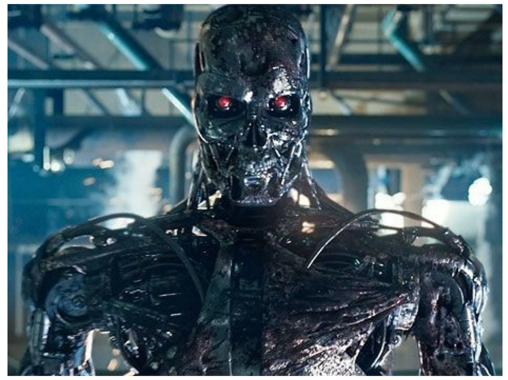



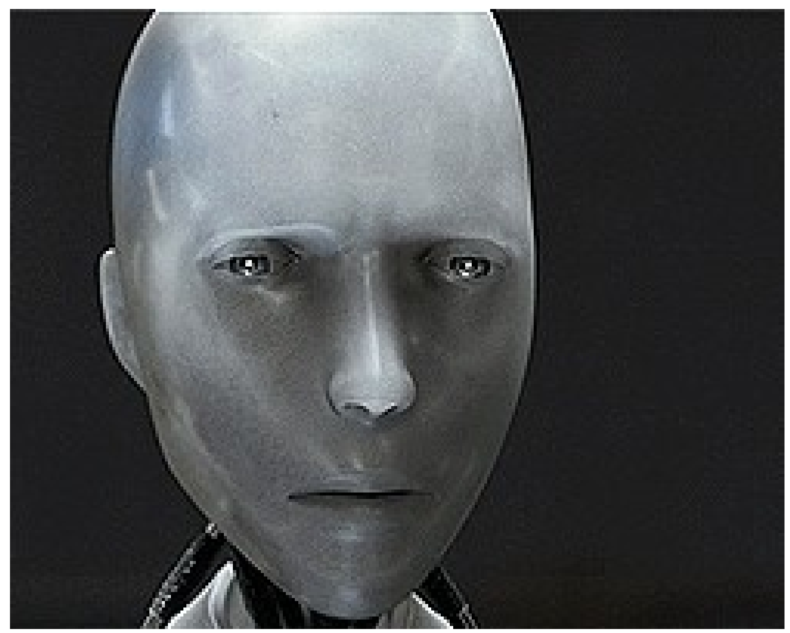


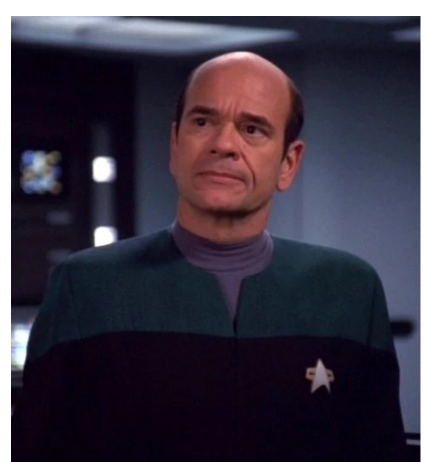
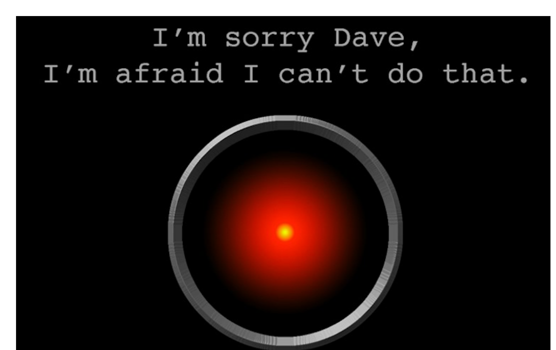

LINKÖPING UNIVERSITY

Hard or Easy?

Solving a system of linear equations may seem ”hard” to us…

System of Linear Equation

$$
\begin{aligned}
2.0x + 4.0y + 6.0z &= 18 \\
4.0x + 5.0y + 6.0z &= 24 \\
3.0x + 1y - 2.0z &= 4
\end{aligned}
$$

Matrix representation

$$
\mathbf{A} = \begin{bmatrix} 2.0 & 4.0 & 6.0 \\ 4.0 & 5.0 & 6.0 \\ 3.0 & 1.0 & -2.0 \end{bmatrix} \quad \mathbf{X} = \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \\ \mathbf{z} \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} 18.0 \\ 24.0 \\ 4.0 \end{bmatrix}
$$

…but distinguishing between cats and dogs is ”easy”!

Hard or Easy? (2)

Because solving equation systems was "hard",
we explicitly had to think about **all the details**
➔ can easily **program** a step-by-step solver

System of Linear Equation

$$2.0x + 4.0y + 6.0z = 18$$
$$4.0x + 5.0y + 6.0z = 24$$
$$3.0x + 1y - 2.0z = 4$$

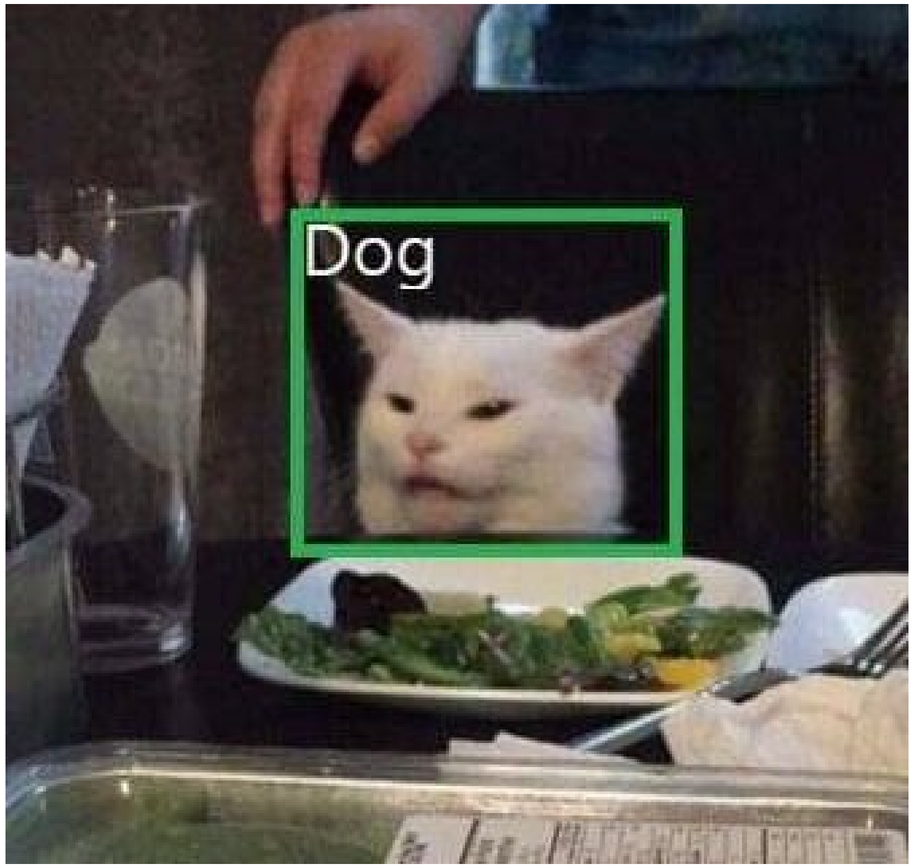

…but we don't really know **how**
we distinguish between cats and dogs!
➔ Hard to define an algorithm
➔ "Requires intelligence!"

LINKÖPING UNIVERSITY

AI History

Alan Turing (1912-1954)

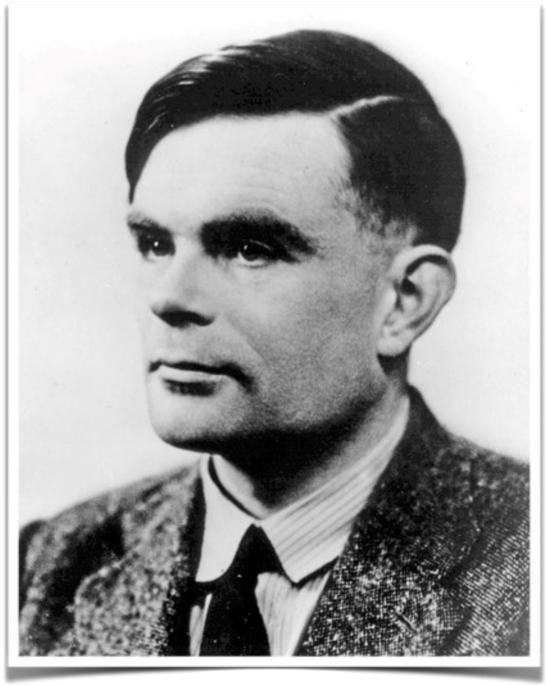

October, 1950

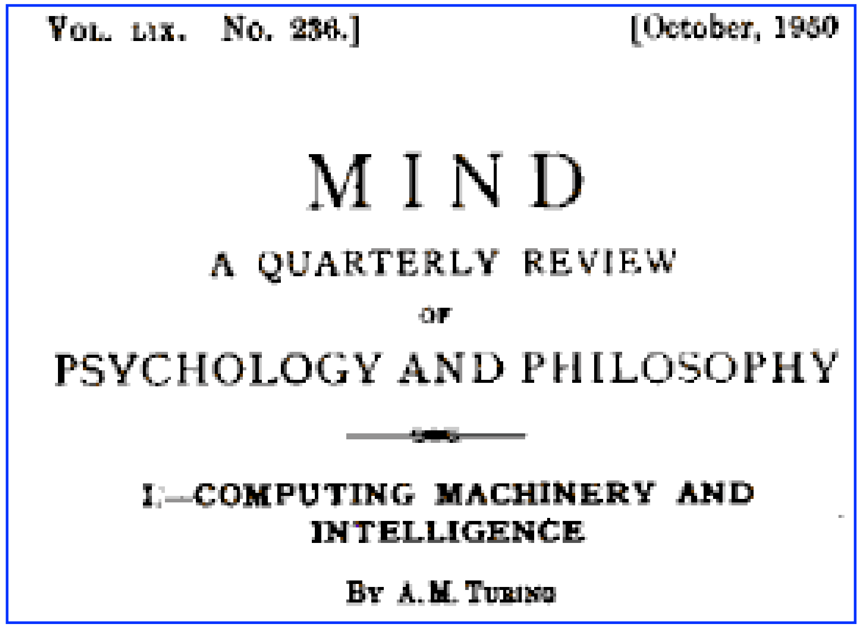

VOL. LIX. No. 236.] [October, 1950

MIND

A QUARTERLY REVIEW

OF

PSYCHOLOGY AND PHILOSOPHY

I.—COMPUTING MACHINERY AND INTELLIGENCE

BY A. M. TURING

I propose to consider the question, "Can machines think?"

Since the meaning of both "machine" and "think" is ambiguous, Turing replaces the question by another.

A behavioural test: The Turing test

Turing test

During the Turing test, the human questioner asks a series of questions to both respondents. After the specified time, the questioner tries to decide which terminal is operated by the human respondent and which terminal is operated by the computer.

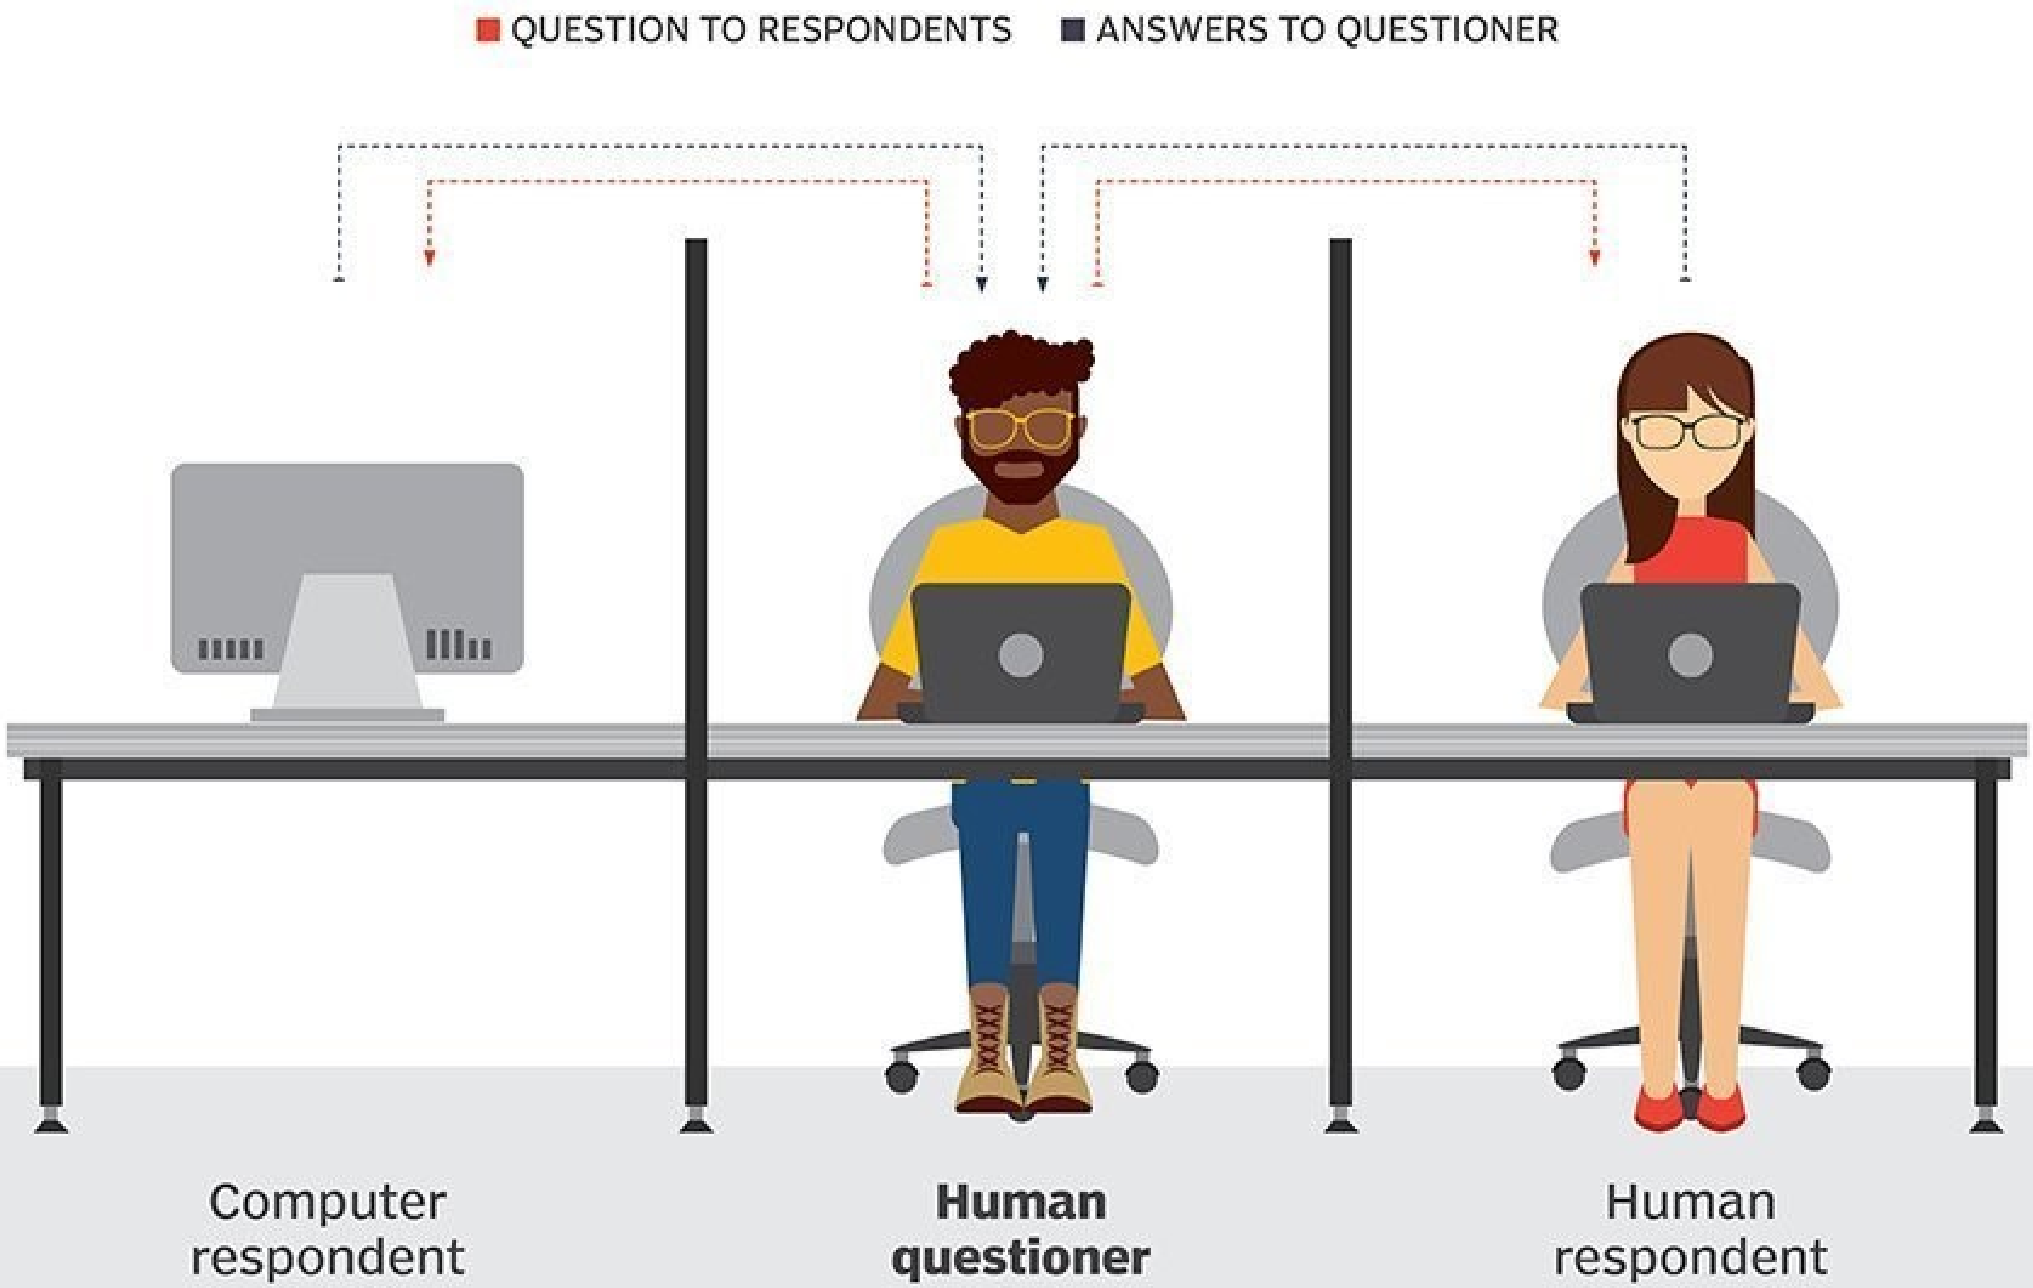

Controversial but still highly relevant

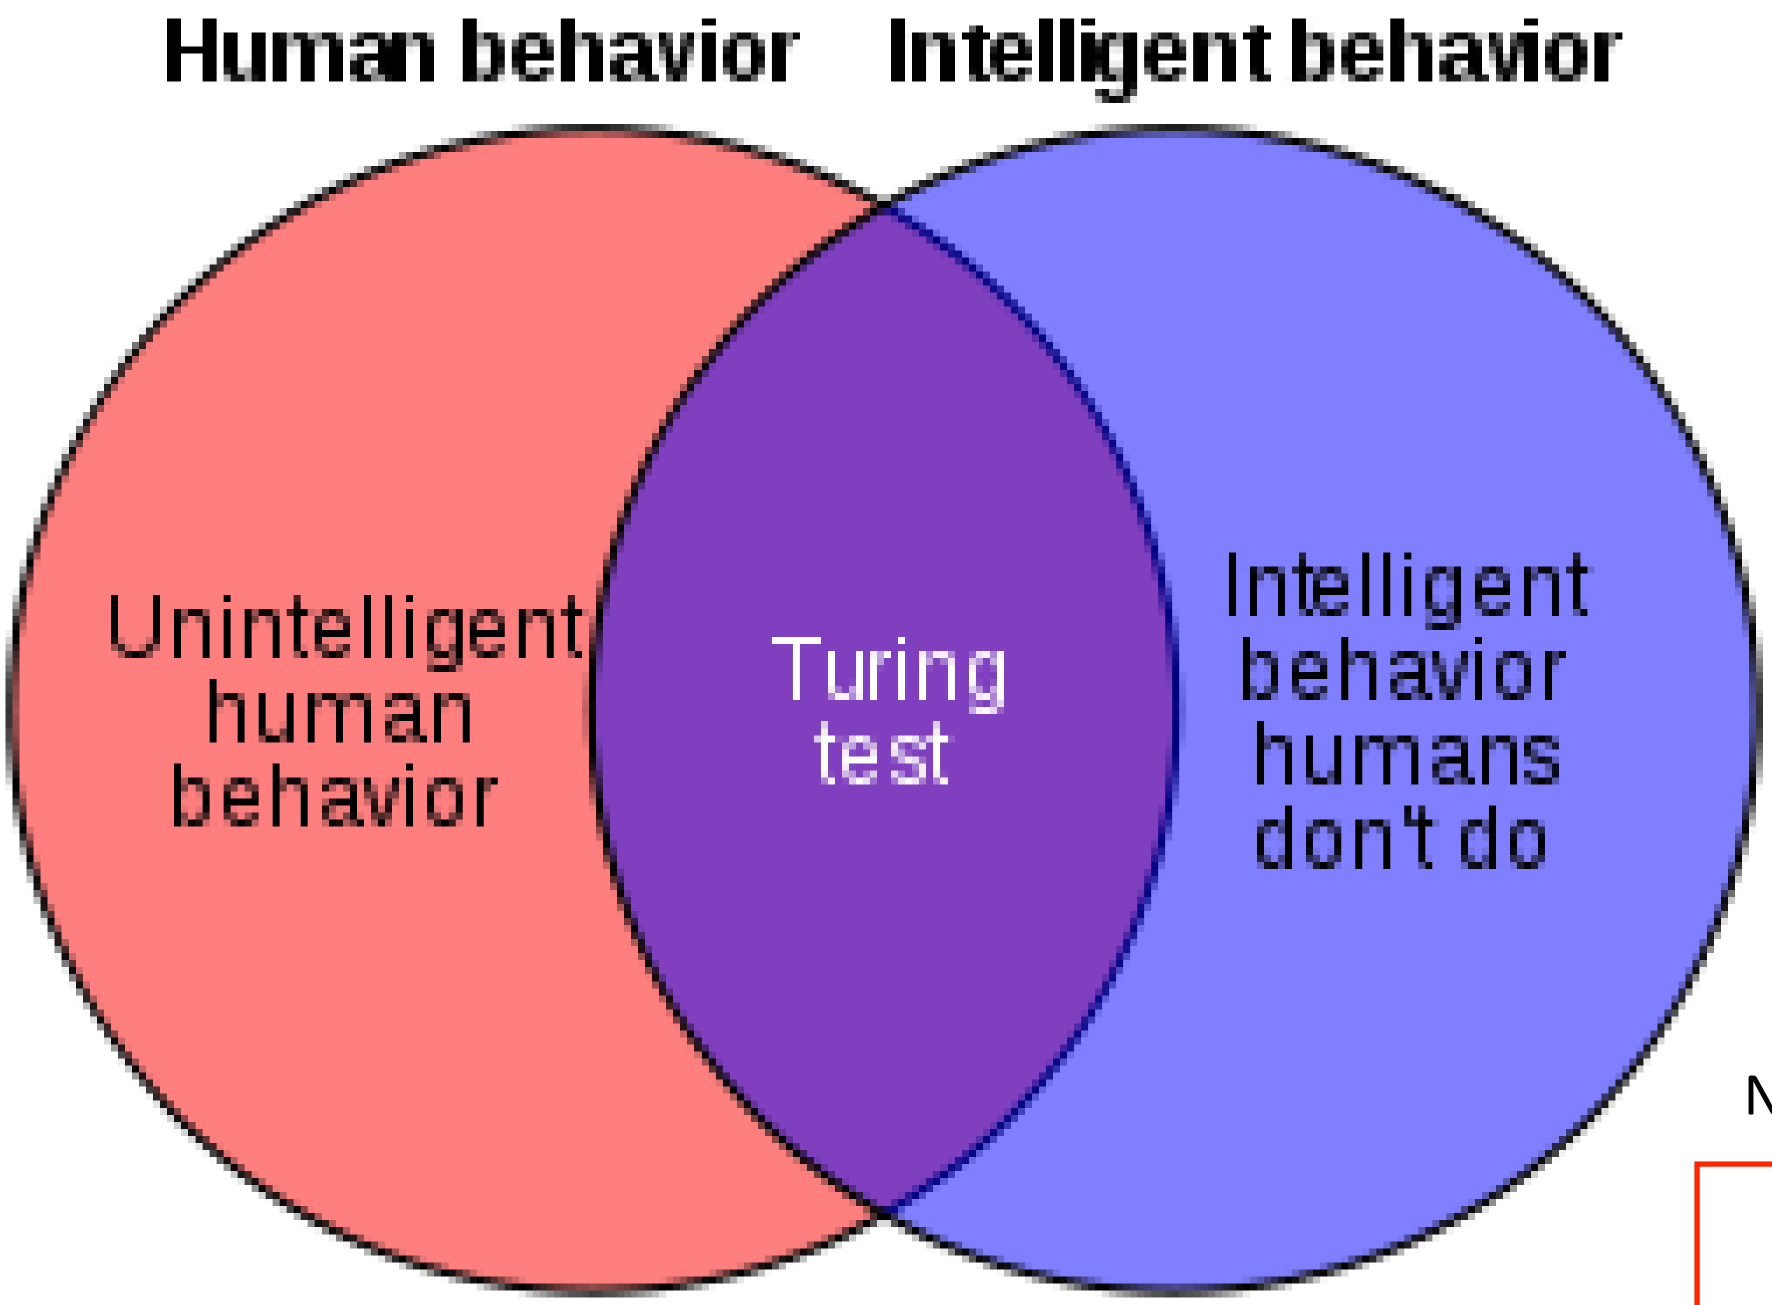

Intelligent behavior that humans do!

Not considered:

Vision
Motion
Movement
Perception

The Dartmouth Workshop (1956)

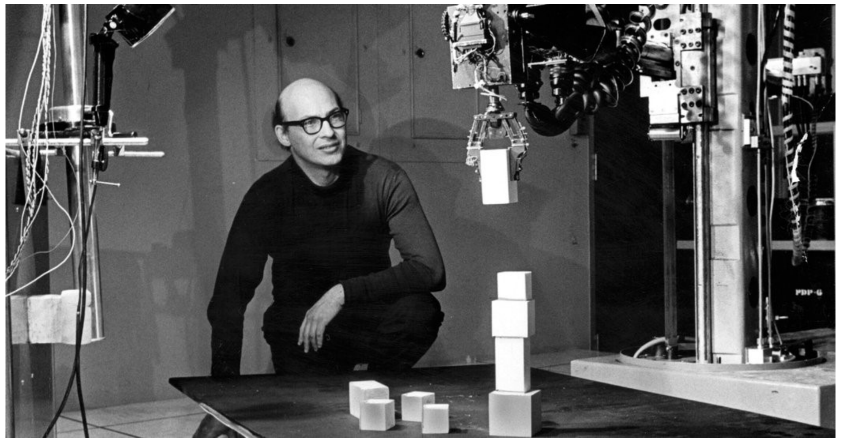

Marvin Minsky

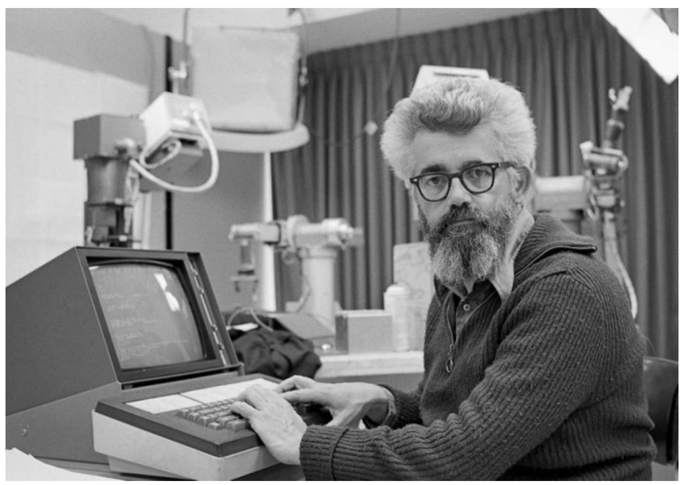

John McCarthy

We propose that a 2-month, 10-man study of **artificial intelligence** be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.

An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

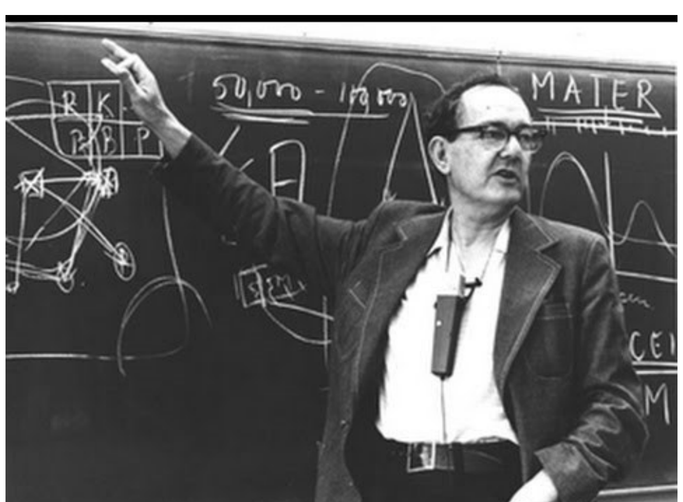

Herb Simon

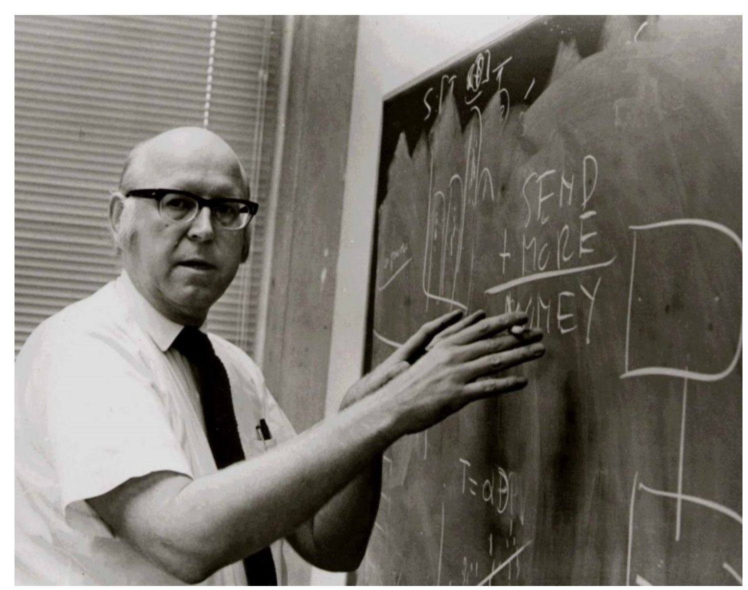

Allen Newell

plus more..

Did early AI solve any problems?

“As soon as it works, no one calls it AI anymore.”

John McCarthy

Different Views on Artificial Intelligence

Different Views on AI

Artificial Narrow Intelligence (Weak AI)

AI that specialises in one area – more and more successful

Example: Playing Chess

- Difficult for **us**, seemed to require **human intelligence**
  - 1951: First chess-playing program (Turing)
  - 1980s: Garry Kasparov: ”AI will never defeat human grandmasters”
  - 1990s: IBM’s Deep Blue won against Kasparov
  - 2020s:
    Far better than human grandmasters,
    running on a phone…

Different Views on AI

Example: Playing Chess

- Specialized AI techniques
  - Example: **<u>Search</u>** through the space of all possible moves
    - Succeeds through hardcoded adaptations to playing chess
    - Can't reason about anything else

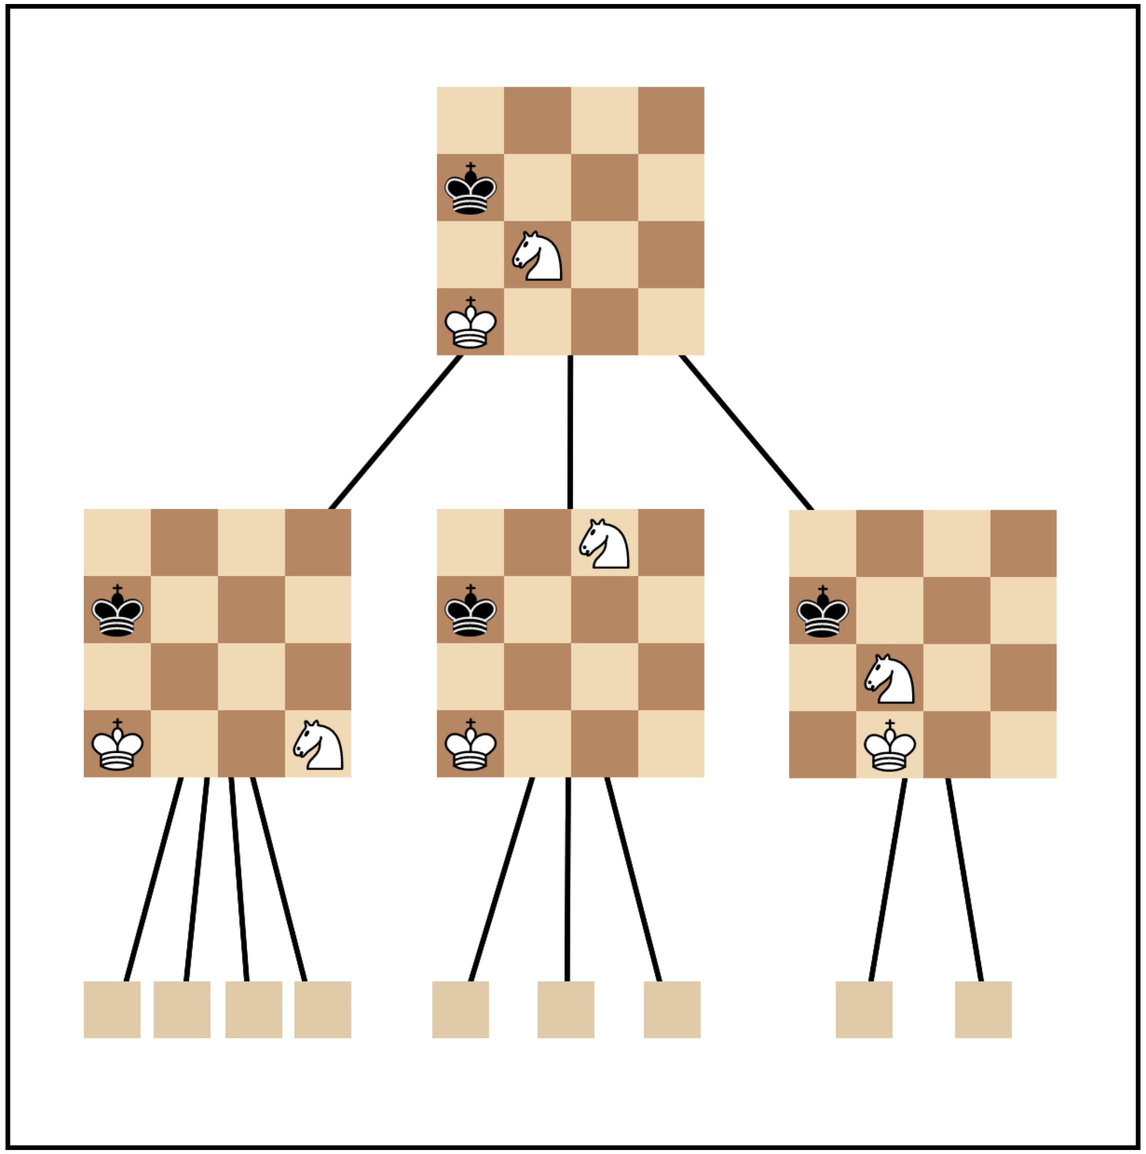

Different Views on AI

Example: IBM's Watson (2011)

LINKÖPING UNIVERSITY

200 million pages of info (incl Wikipedia) / 90 IBM Power 750 servers, 32 threads/server

Different Views on AI

Artificial General Intelligence (Strong AI)

Smart as a human across the board

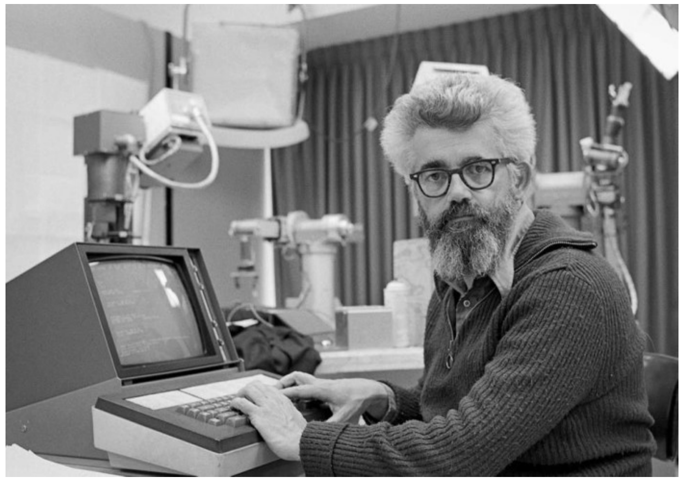

Human-level Intelligence with common sense
-- Passing the Turing Test
-- Obtaining a college degree
-- ...

“Human-level AI will be achieved, but new ideas are almost certainly needed, so a date cannot be reliably predicted—maybe five years, maybe five hundred years. I'd be inclined to bet on this 21st century.” -- John McCarthy

Considerably slower progress – far from the goal

Different Views on AI

Artificial Super Intelligence (ASI)

AI that surpasses humans

Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, **an ultraintelligent machine could design even better machines**; there would then unquestionably be an "intelligence explosion", and the intelligence of man would be left far behind. **Thus the first ultraintelligent machine is the last invention that man need ever make.** — I. J. Good [1965]

We are on the edge of change comparable
to the rise of human life on Earth. — Vernor Vinge [1993]

In his 2014 book Superintelligence: *Paths, Dangers, Strategies*, Nick Bostrom reasoned that a computer with near human-level general intellectual ability could initiate an **intelligence explosion** on a digital time scale with the resultant rapid creation of something so powerful that it might **deliberately or accidentally destroy humanity.**

Nick Bostrom

Artificial Intelligence: The Present

What is Artificial Intelligence?

A Definition:

"the scientific understanding of the mechanisms underlying **thought** and intelligent **behavior** and their **embodiment in machines**." (AAAI)

The Grand Goal:

"a freely moving machine with the intellectual capabilities of a human being." (Hans Moravec)

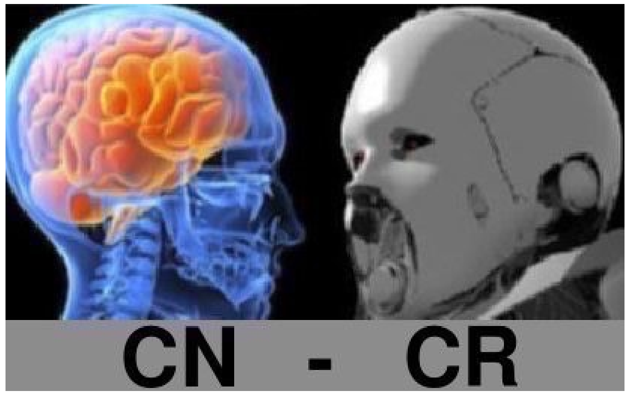

LINKÖPING UNIVERSITY

Modern Foundations of Artificial Intelligence

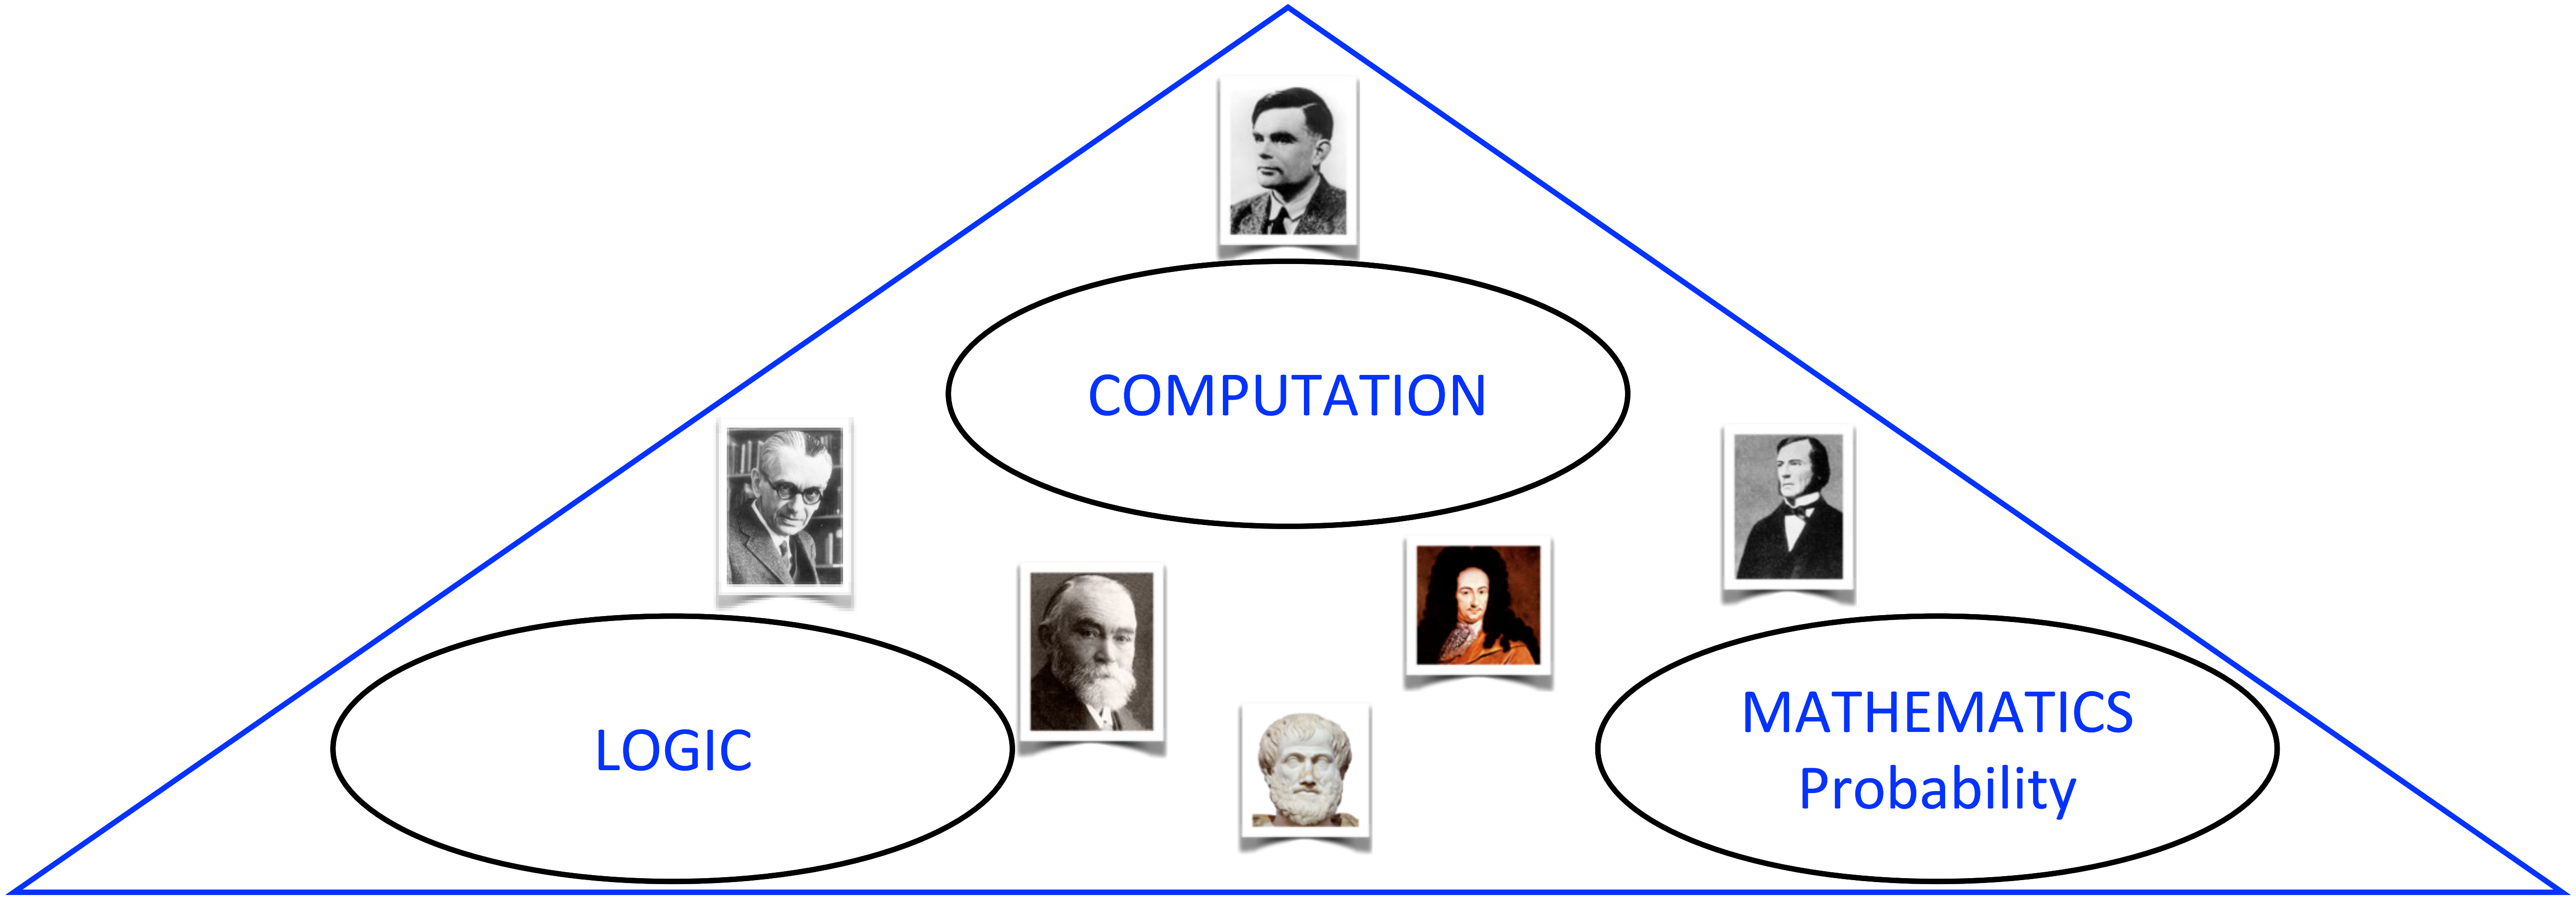

Physical Symbol System Hypothesis

Computer Science as Empirical Enquiry: Symbols and Search
Newell and Simon (1976)

- A **physical symbol system**, also called a **formal system**:
  - Takes physical patterns (symbols)
  - Combines them into structures (expressions)
  - Manipulates manipulating them (using processes)…
  - … to produce new expressions.

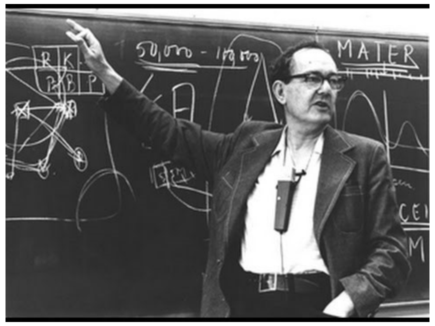

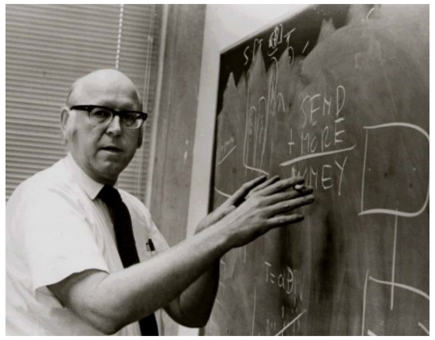

- Examples:
  - Formal logic (”and”, ”or”, ”not”, ”for all”; logic formulas; deduction)
  - Algebra (”+”, ”1”, ”2”; equations; rules of algebra)
  - Computers (values in memory; machine operations; the CPU)
  - Chess (pieces; chess board configurations; legal chess moves)
  - …

Physical Symbol System Hypothesis (2)

- The physical symbol system **<u>hypothesis</u>**:
  - "A physical symbol system has the necessary and sufficient means for general intelligent action."
  - ➔ **<u>Human thinking</u>** processes symbols
    (because the PSS is **<u>necessary</u>**)
  - ➔ Machines can be intelligent
    (because the PSS is **<u>sufficient</u>**)
  - Core part of AI research;
    also **<u>controversial</u>**, strongly criticized by some
    (do we need symbols for vision?)

Intelligent Agents

An **agent** is anything that can be viewed as **perceiving** its **environment** through **sensors** and **acting** upon that environment through **actuators.**

An agent's **behavior** can be described formally as an **agent function** which maps any percept sequence to an action

An **agent program** implements an **agent function**

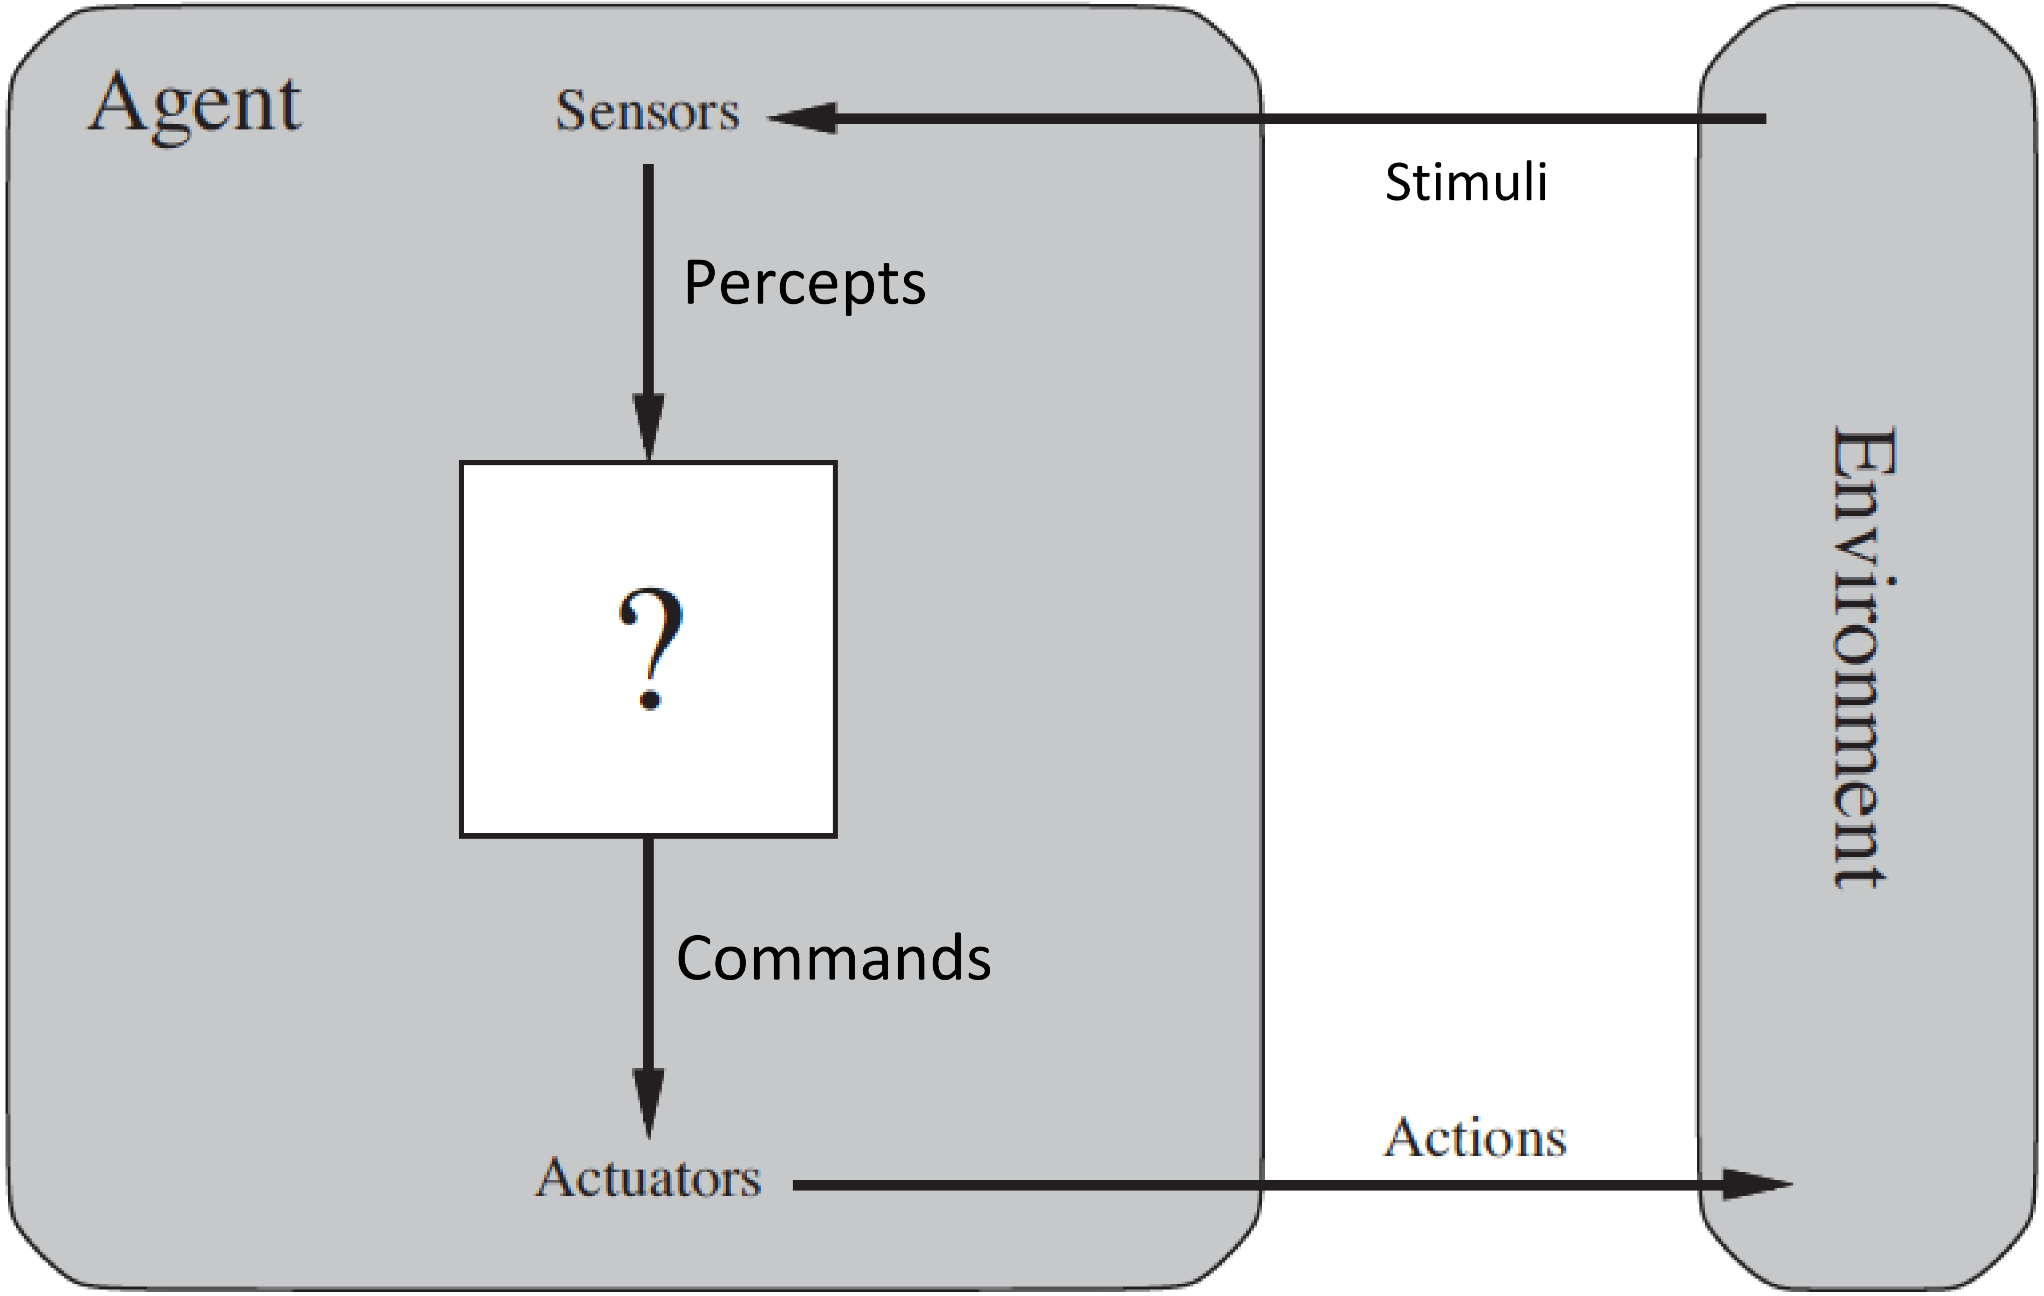

A **Rational Agent** is one that does the right thing relative to an external performance metric

Humans as Intelligent Agents

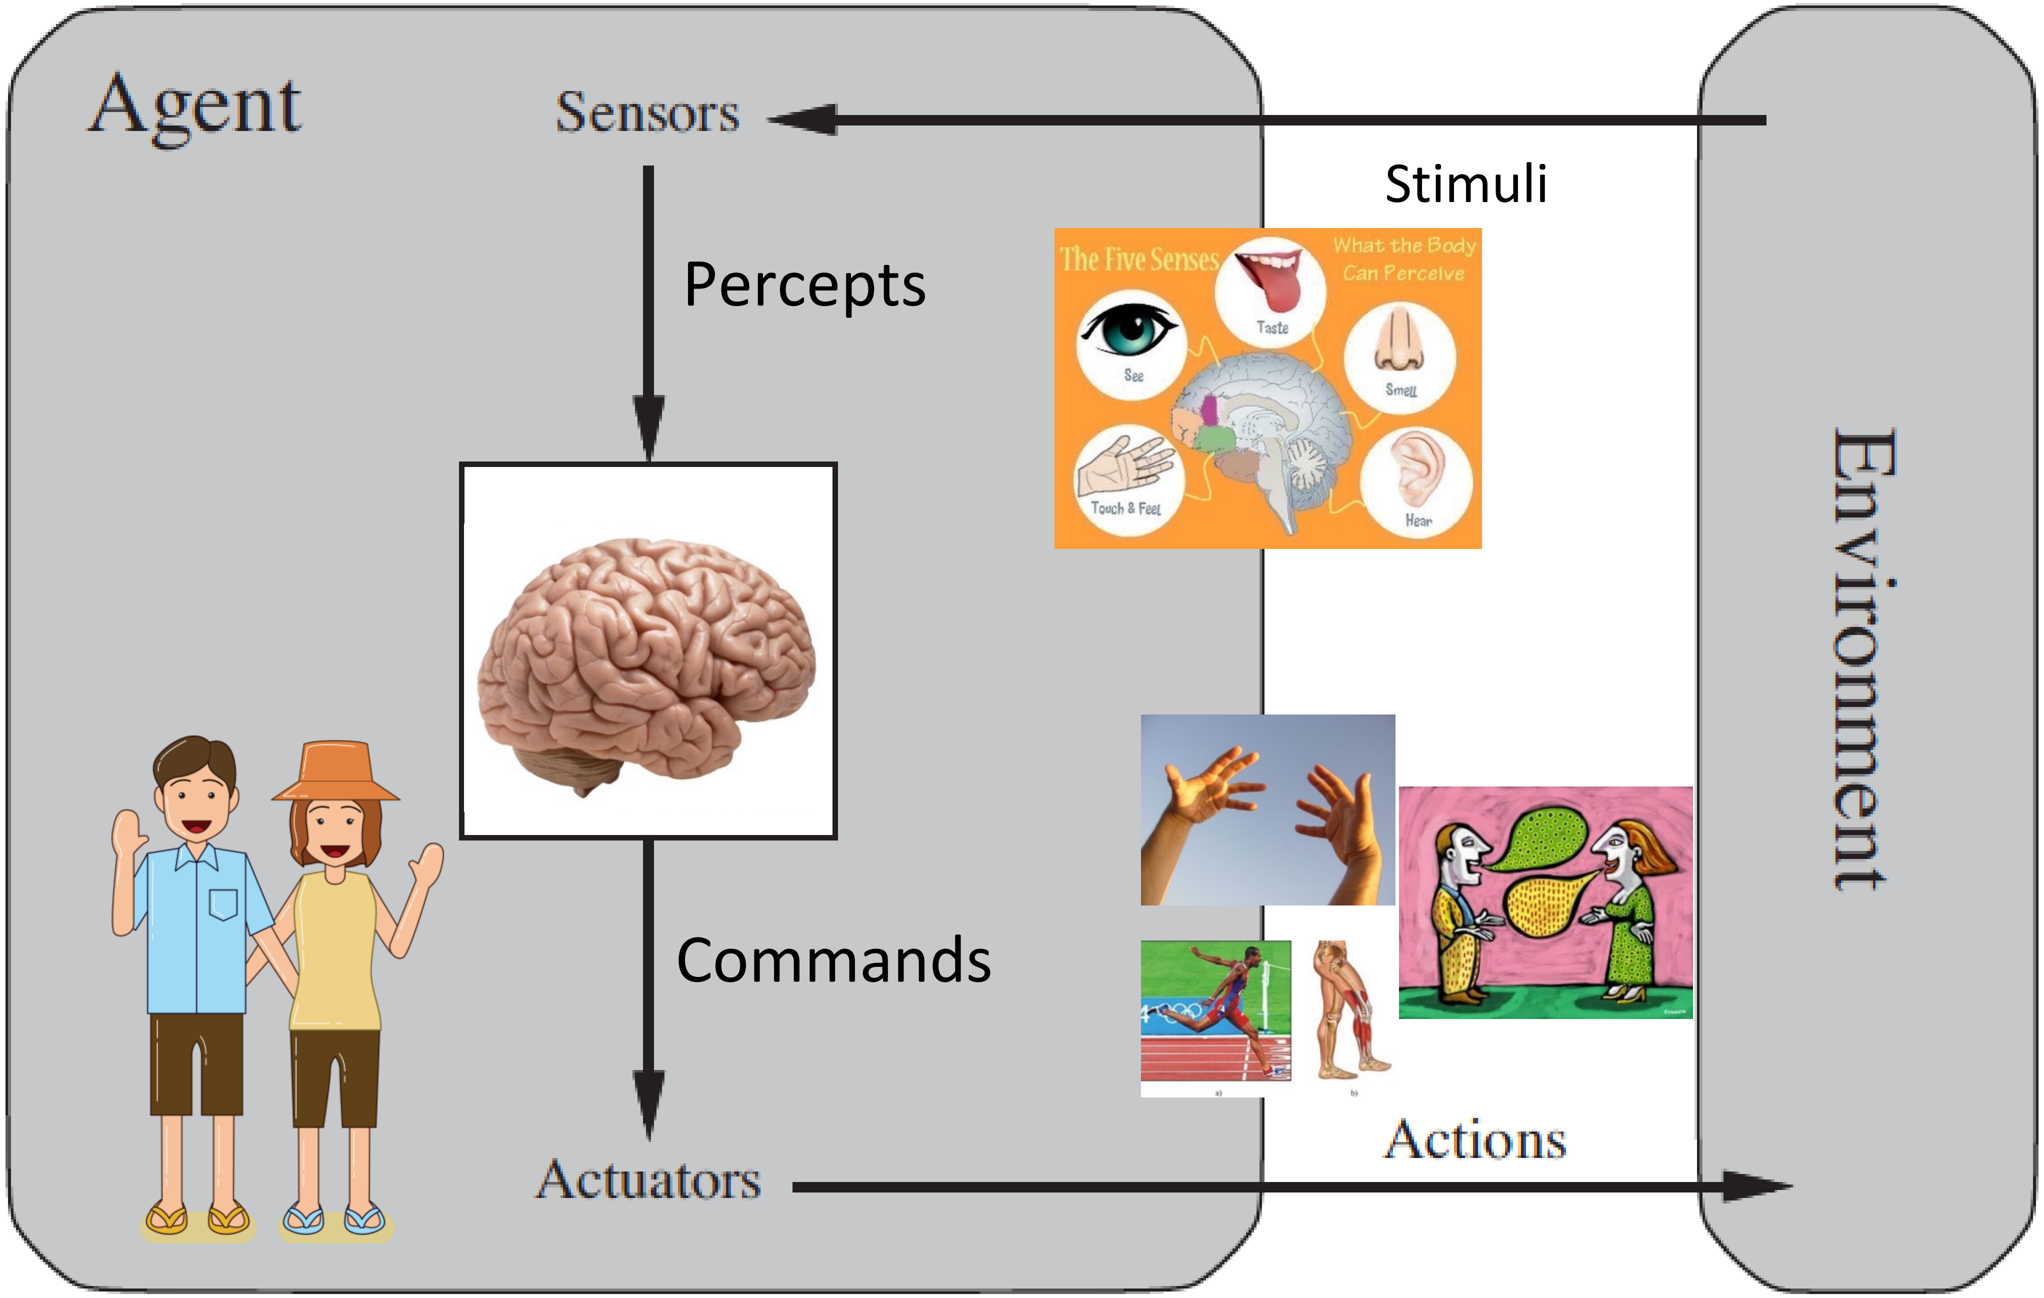

Humans interact with the environment through sensors and actuators

Humans use internalised models of the environment to reason and act intelligently

Robots as Intelligent Agents

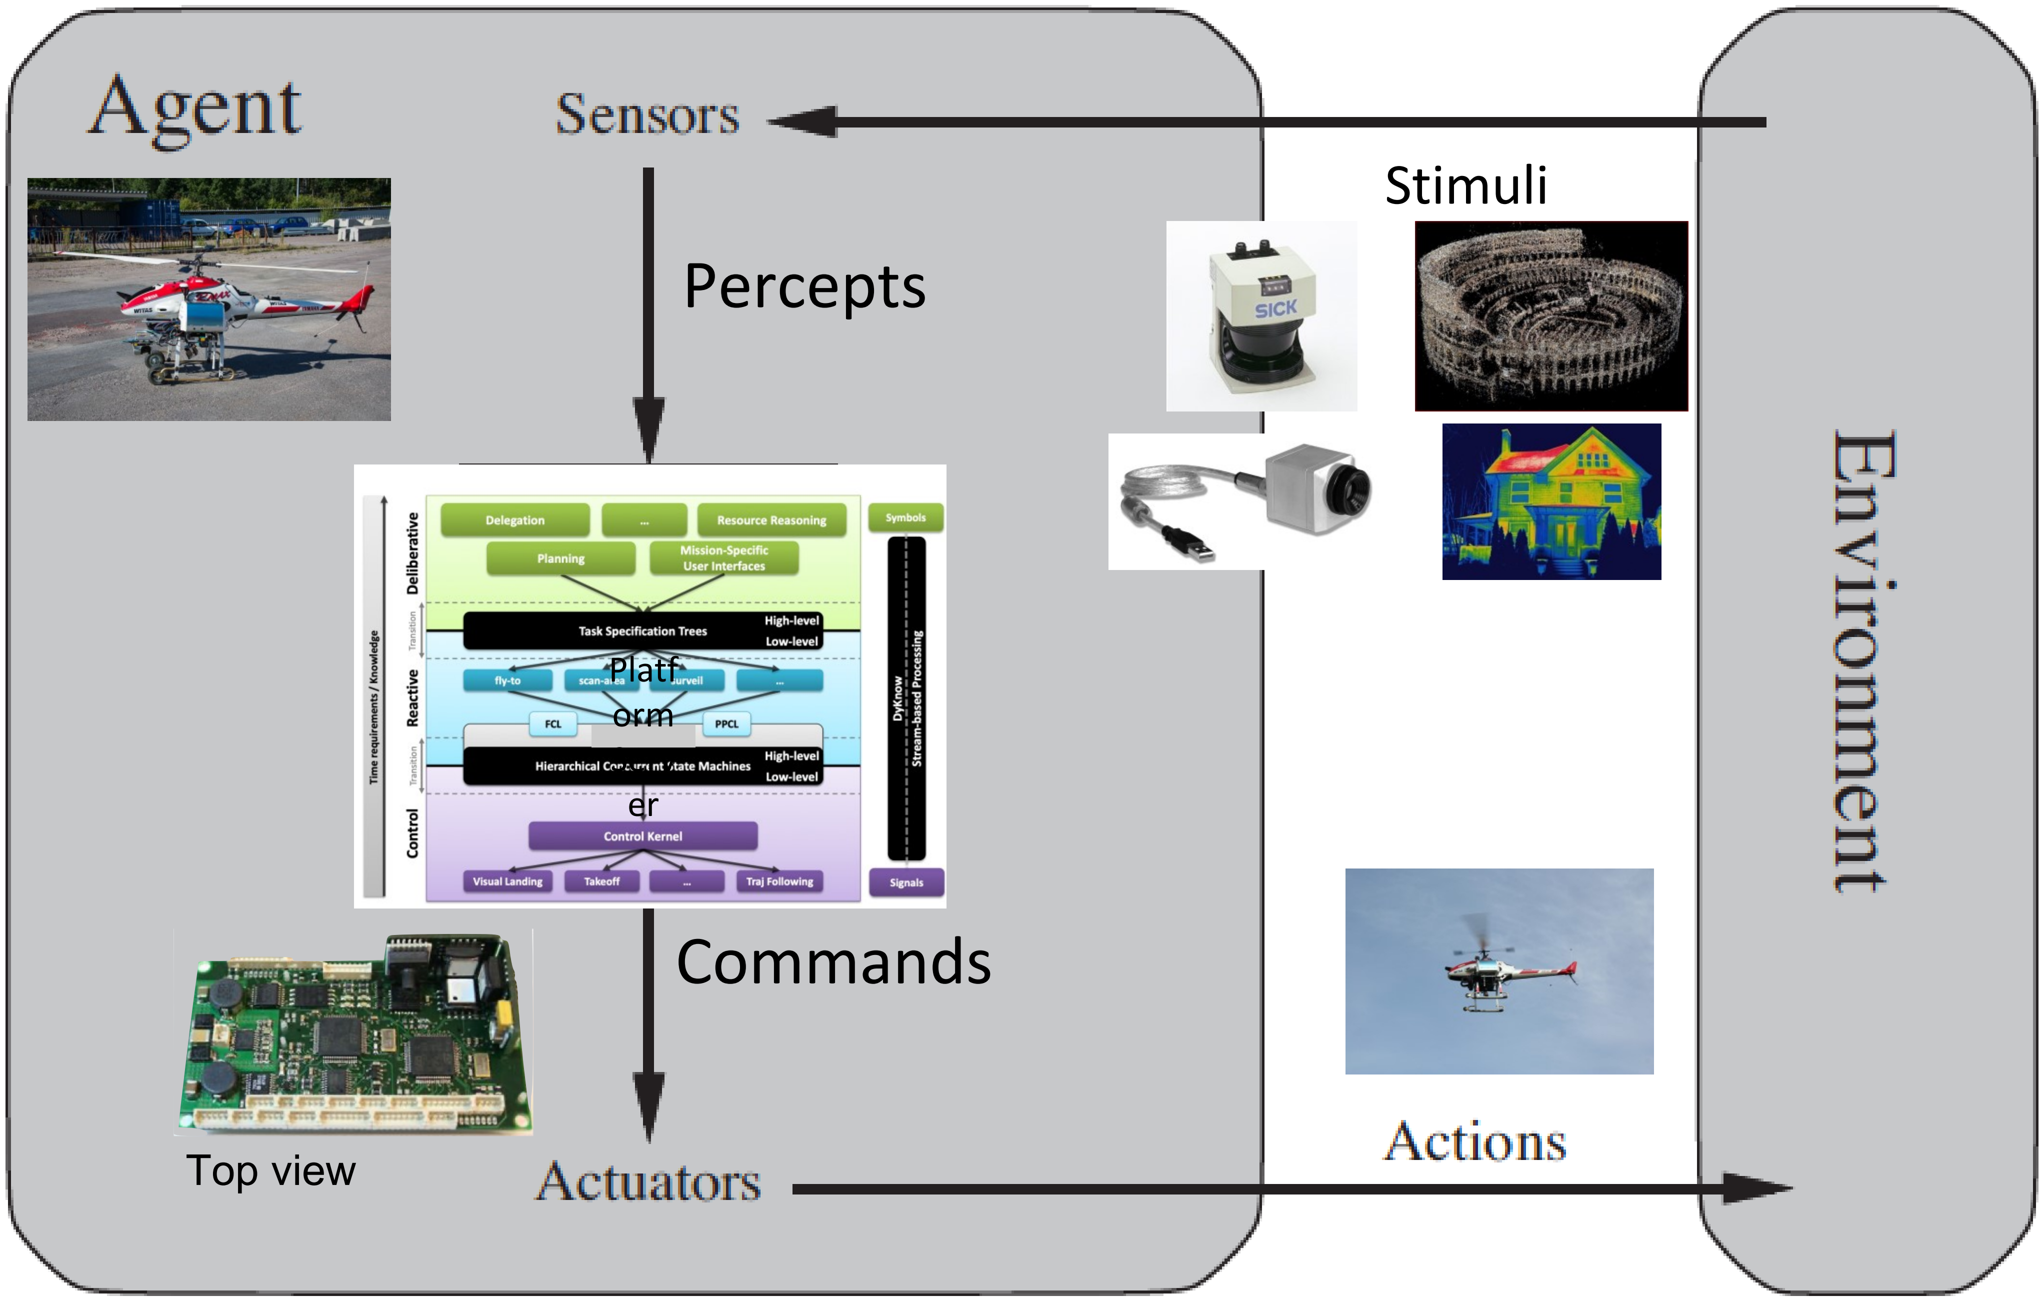

Robots interact with the environment through sensors and actuators

Robots use internalised models of the environment to reason and act intelligently

Agent Types

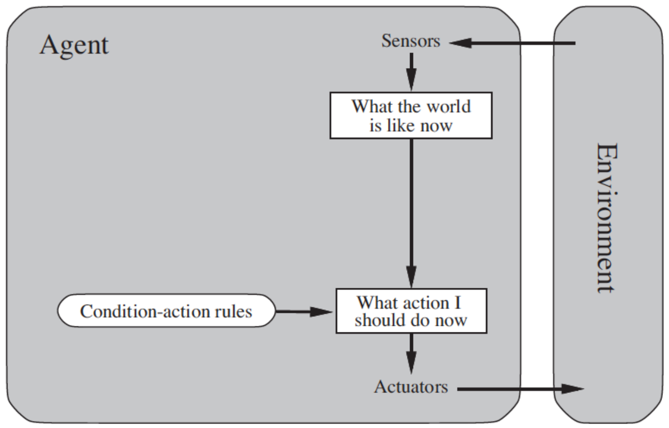

Simple reflex agent

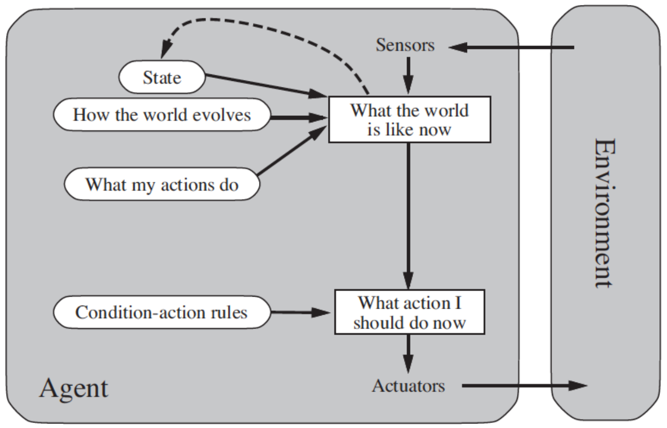

Model-based reflex agent

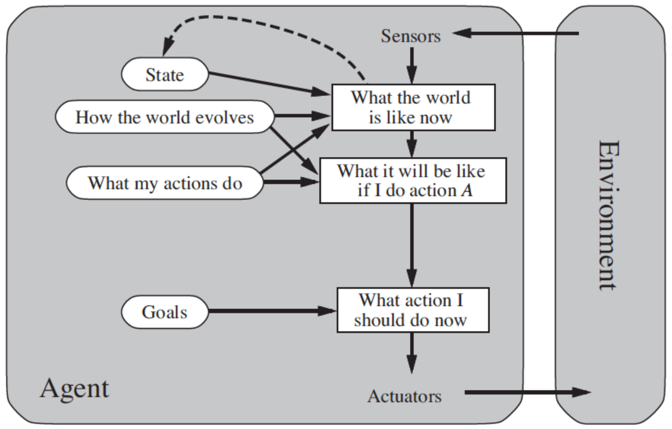

Goal-based agent

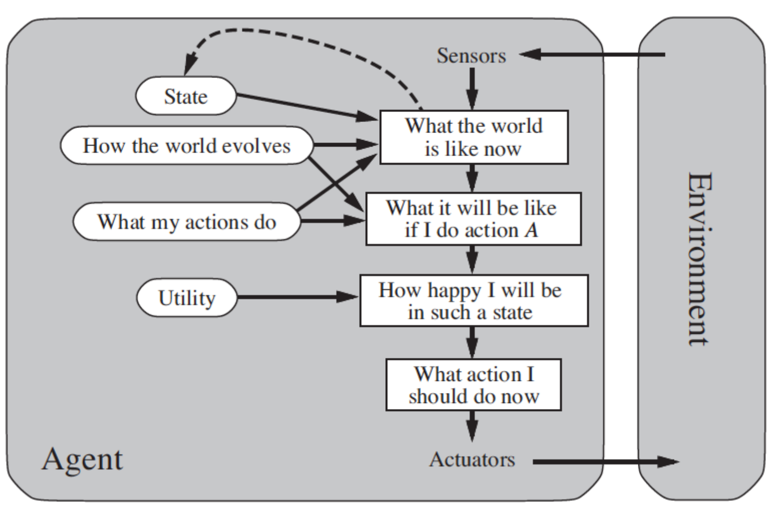

Utility-based agent

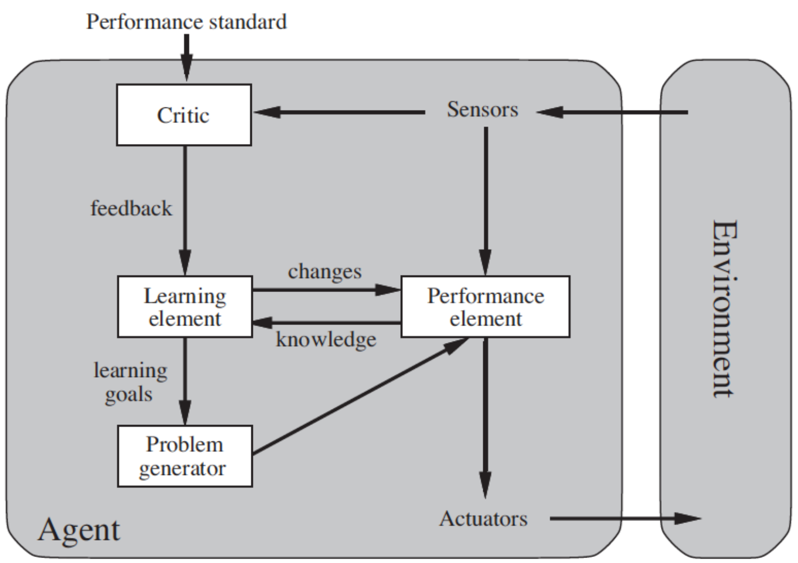

Learning agent

LINKÖPING UNIVERSITY

Simple Reflex Agent

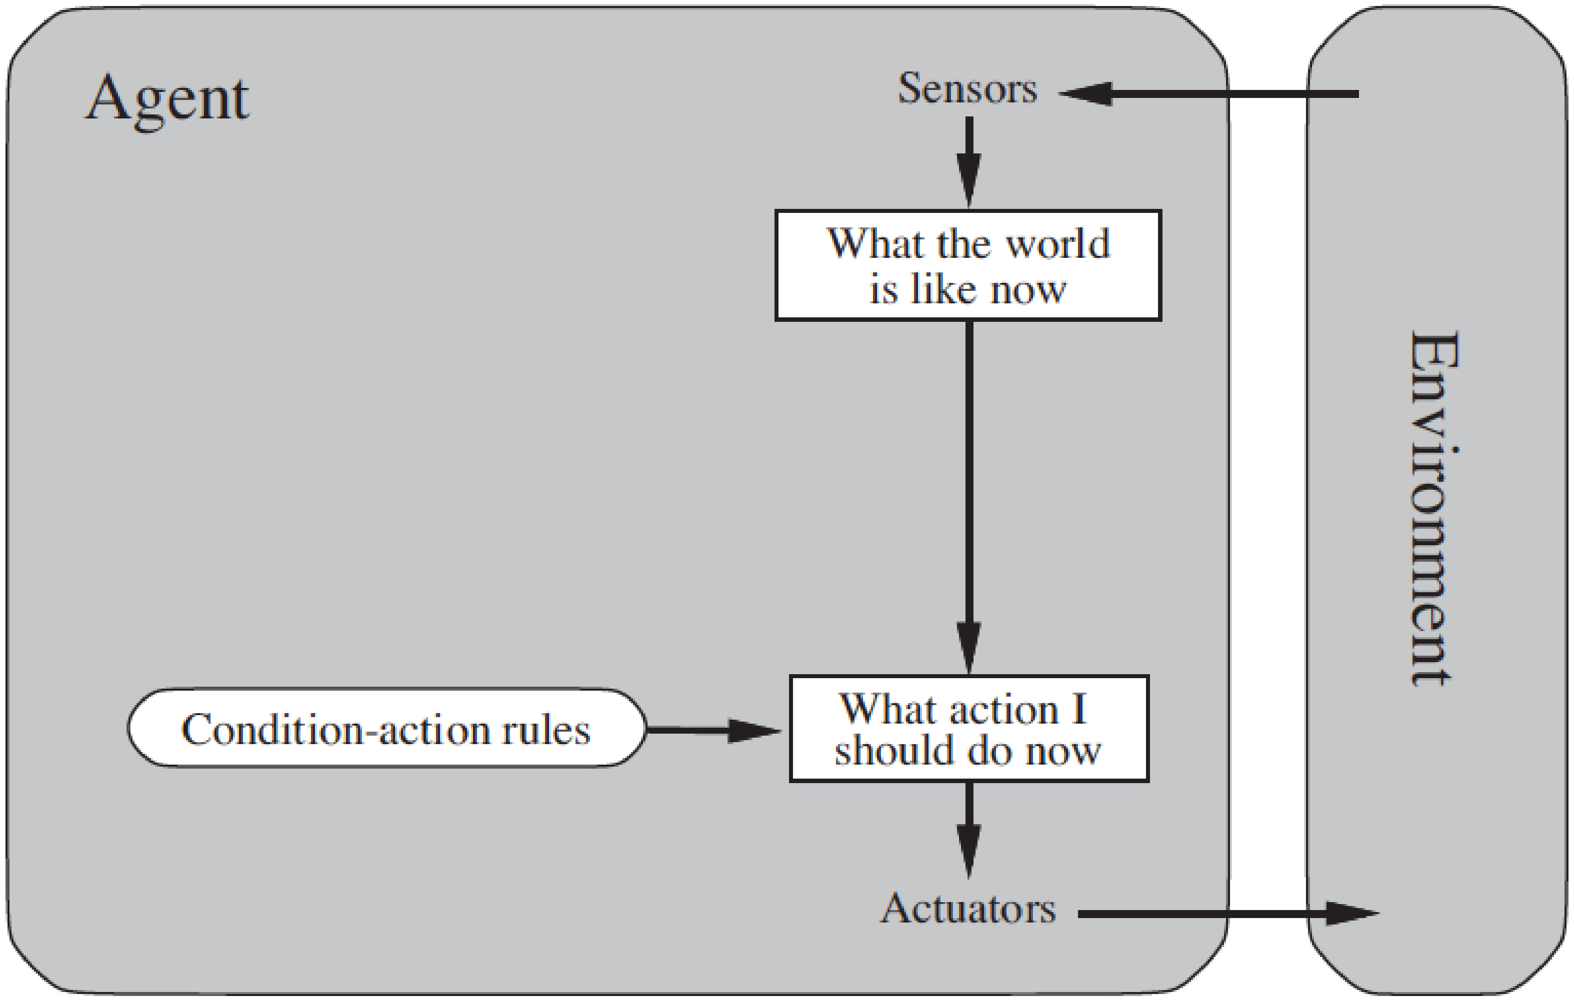

Stimulus-Response Agent

- Reacts to immediate stimuli in their environment
- No internal state
- Uses current state of the environment derived from sensory stimuli

Grey Walter's Tortoise

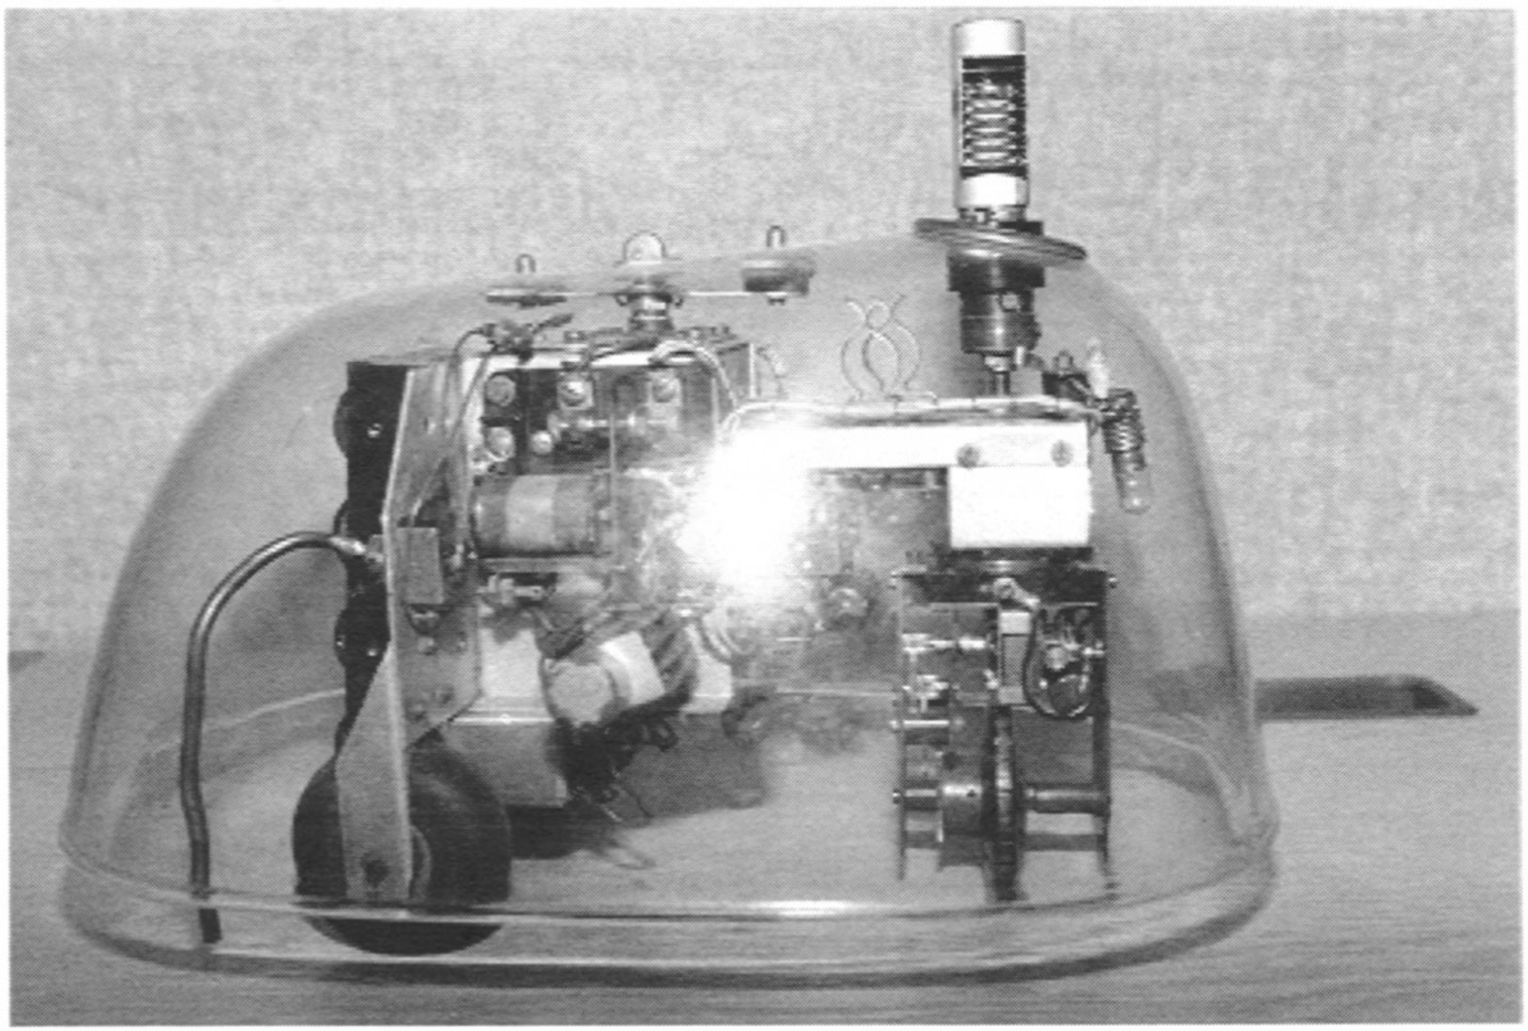

Figure 1.5
Grey Walter's tortoise, recently restored to working order by Owen Holland. (Photograph courtesy of Owen Holland, The University of the West of England.)

Analog Device

2 sensors:
- directional photcell
- bump contact sensor

2 actuators
2 nerve cells (vacuum tubes)

Skills:
- Seek weak light
- Avoid strong light
- turn and push (obstacle avoid.)
- Recharge battery

Gengis II: A Robot Hexapod

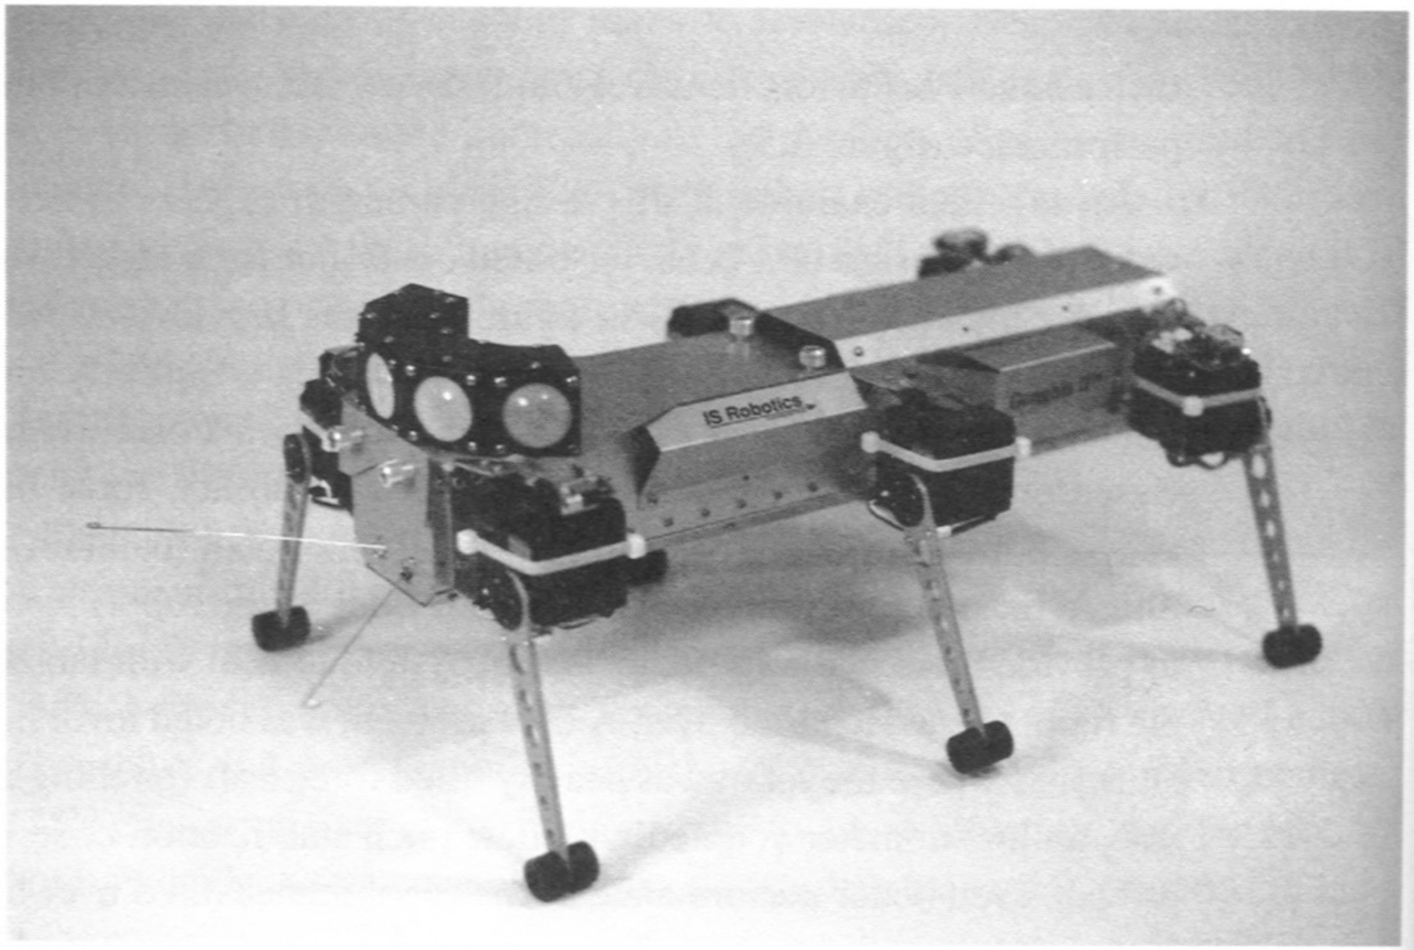

(B)

Figure 3.6
(A) Original Genghis. (Photograph courtesy of Rodney Brooks.) (B) Genghis II—a robotic hexapod, commercial successor to the original Genghis. (Photograph courtesy of IS Robotics, Somerville, MA.)

Brooks – Subsumption-Based Architectures.

Founded iRobot

A Goal-Based Agent

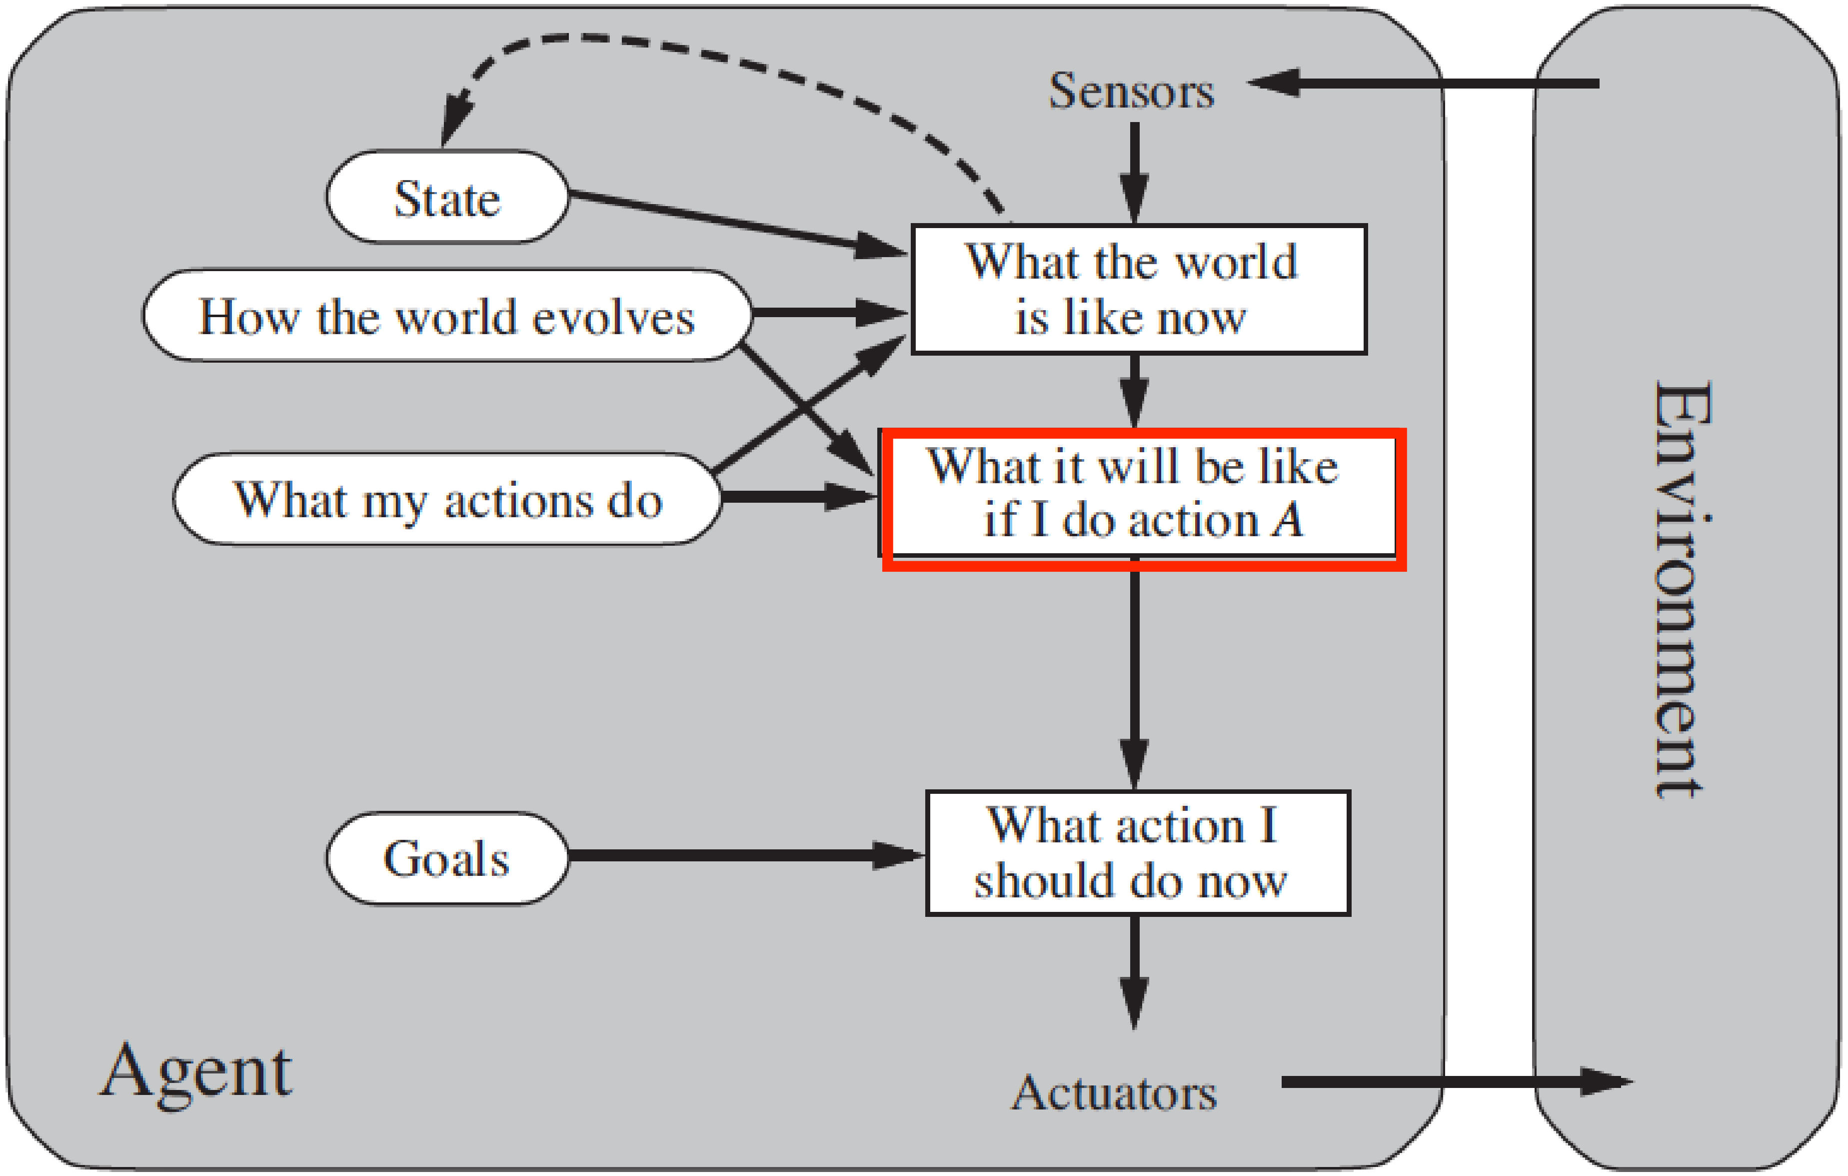

Planning and Reasoning Agents

Agents with Purpose!

Goal-based Agents:

- Rich internal state
- Can **anticipate** the effects of their actions
- Take those actions expected to lead toward achievement of goals
- Capable of **reasoning** and **deducing** properties of the world

Trade-offs between Deliberation and Reaction

Robot Control System Spectrum (Arkin)

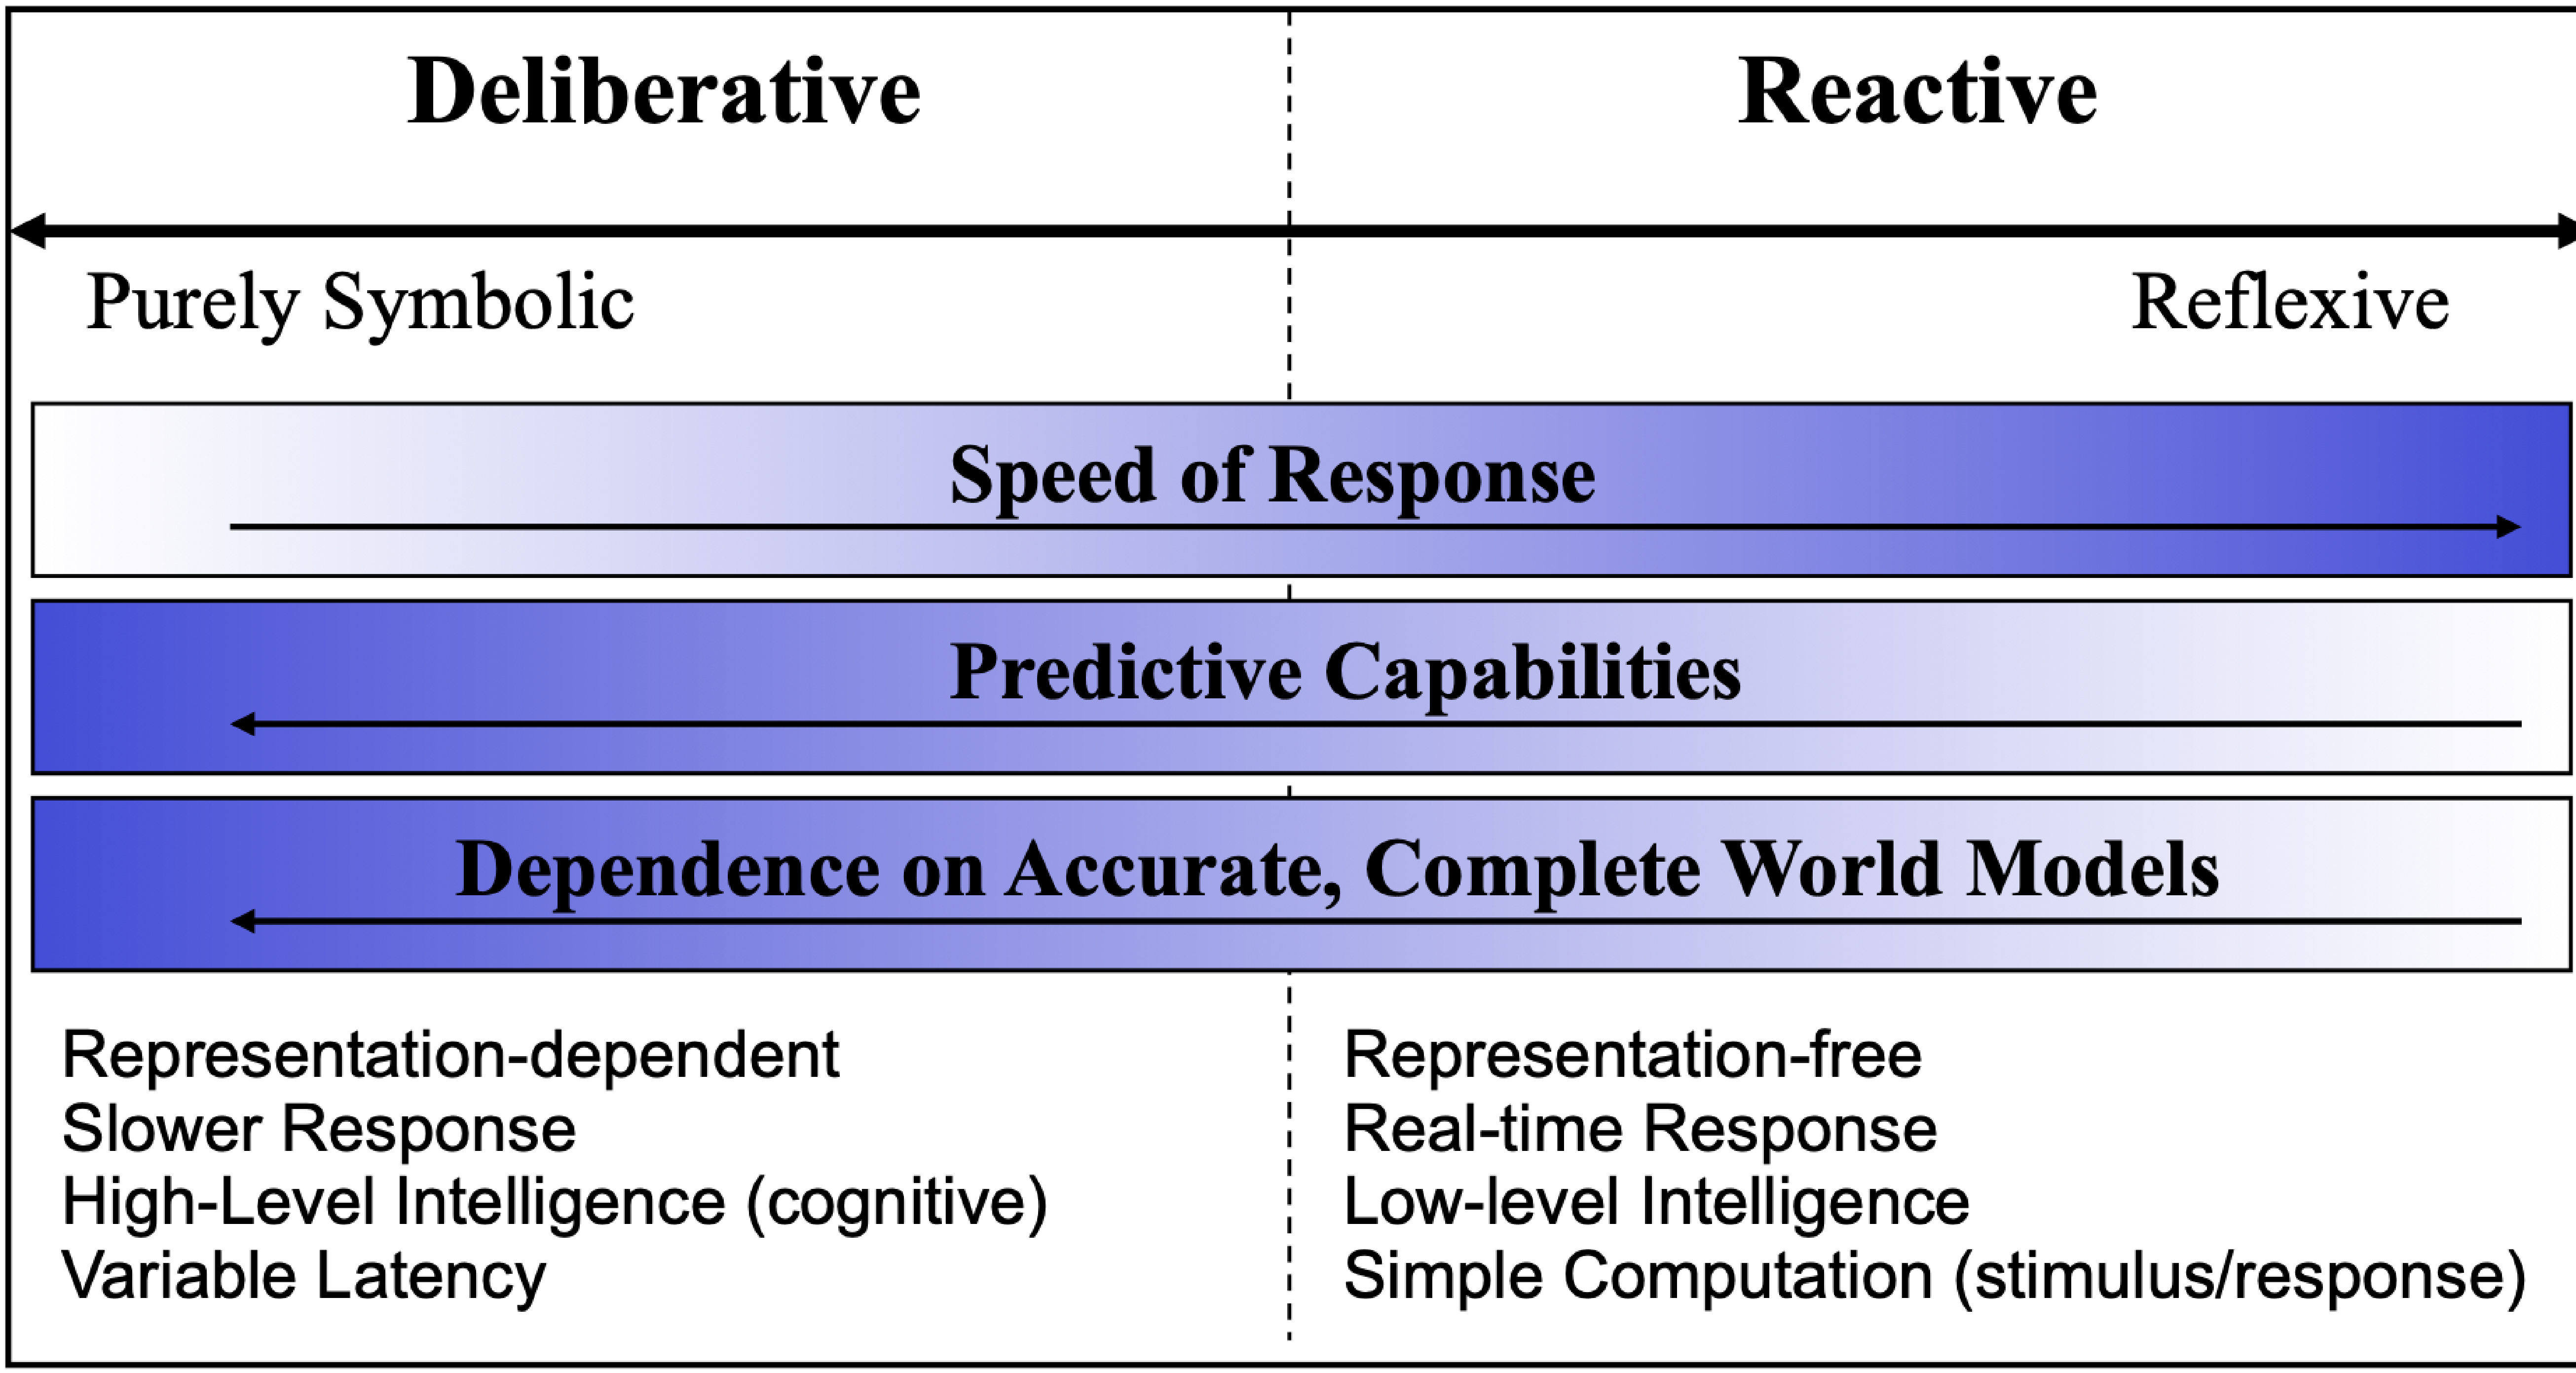

Integrated AI Architectures

Software Architectures that support efficient real-time/soft real-time interaction between control, reactive, and deliberative processes

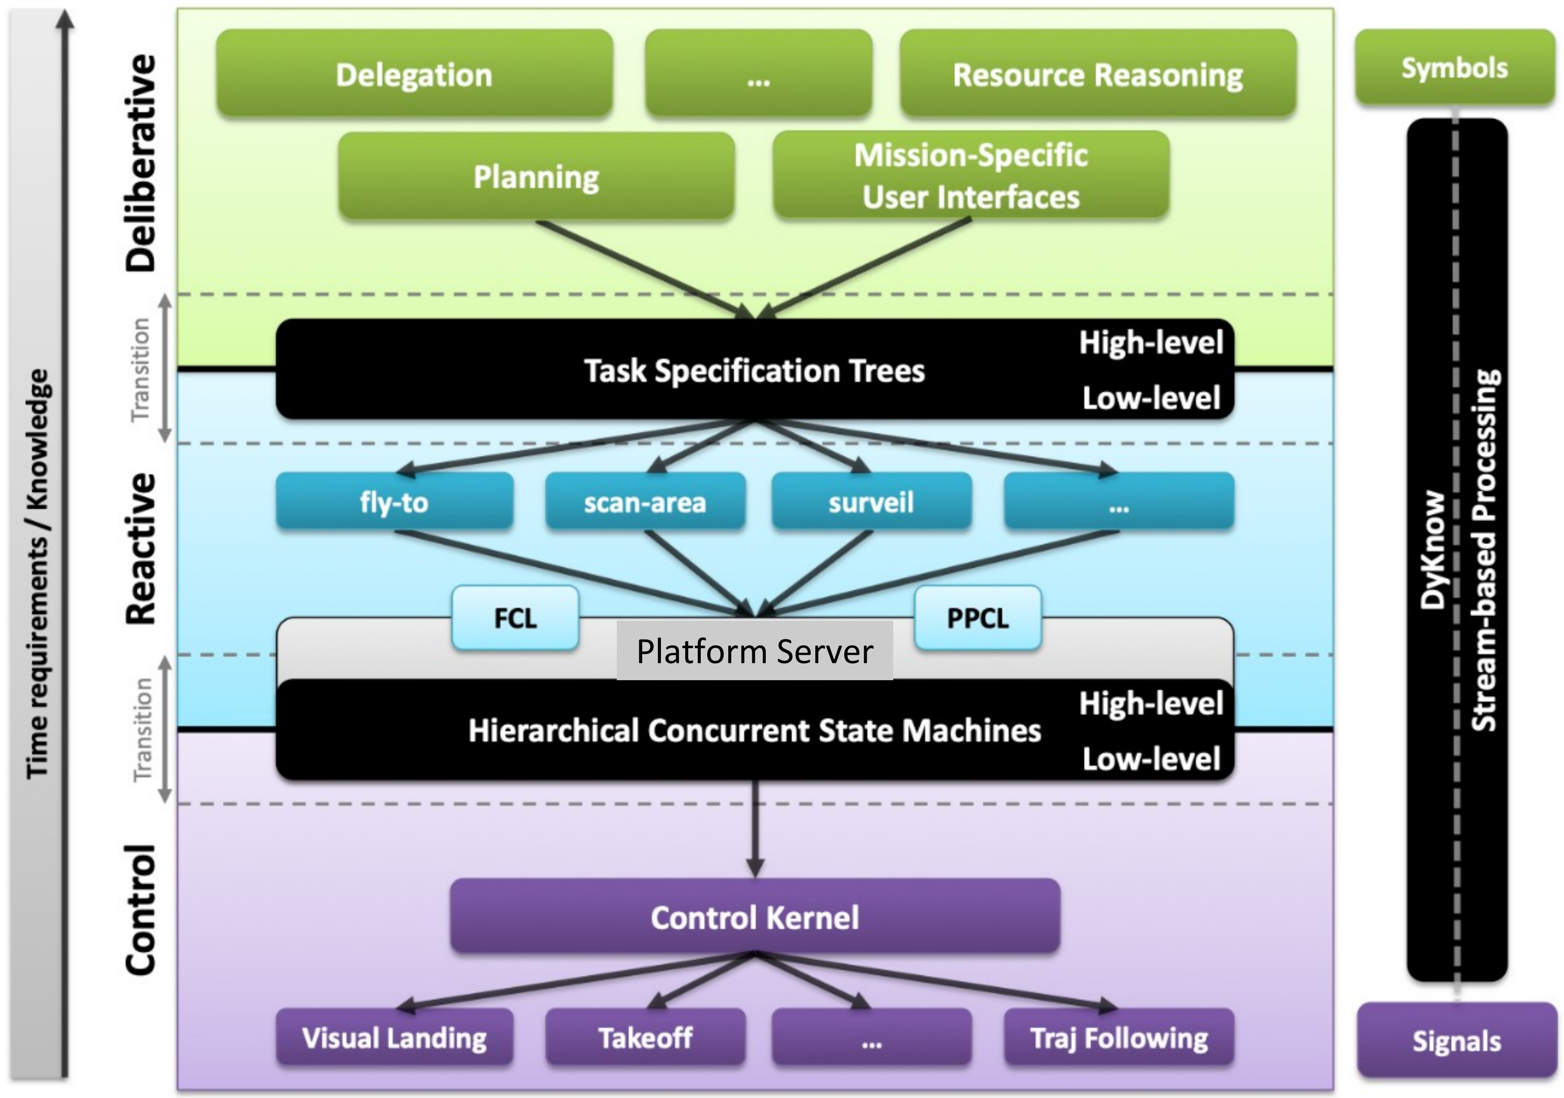

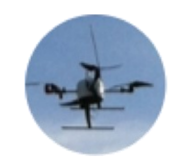

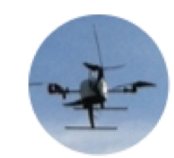

HDRC3: A Distributed Hybrid Deliberative/Reactive Architecture for Unmanned Aircraft Systems
P. Doherty, J. Kvarnström, M. Wzorek, P. Rudol, F. Heintz, G. Conte
Handbook of Unmanned Aerial Vehicles, K.P. Valavanis, G.J. Vachtsevanos (eds.), Springer Science
1st edition 2014, revised 2nd edition 2017.

LINKÖPING UNIVERSITY

Rationality/Rational Agents

Ideal Rational Agent is one that does the right thing!

For each possible percept sequence, an ideal rational agent should do whatever action is expected to [maximize its performance measure](), on the basis of the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

What is a good state?

How should I <u>act</u> to be in that good state?

How do I avoid bad states?

The Essence of Artificial Intelligence?

Building systems that learn to make, and do make, good decisions!

Let's begin with **making** good decisions

Given: The things we can do and their effects:

$$P(result(\text{action}) = state' \mid \text{action}, \mathbf{env})$$

What are the repercussions of acting in a context?

Given: An estimate on the utility or goodness of states:

$$Utility(state_i)\, for\ any\ state, state_i$$

What states are good for me?

LINKÖPING UNIVERSITY

The Essence of Artificial Intelligence

What is my expected utility/goodness when executing an action?

$$EU(\text{action} \mid \mathbf{e}) = \sum_{state'} P(result(\text{action}) = state' \mid \text{action}, \mathbf{e})\, U(state')$$

Take the weighted average of the utilities for states an action *can cause*

<u>Choose</u>: The action that maximises utility in any context!

$$best\text{–}action = argmax_{action} EU(action \mid \mathbf{e})$$

Decision Theoretic Agent

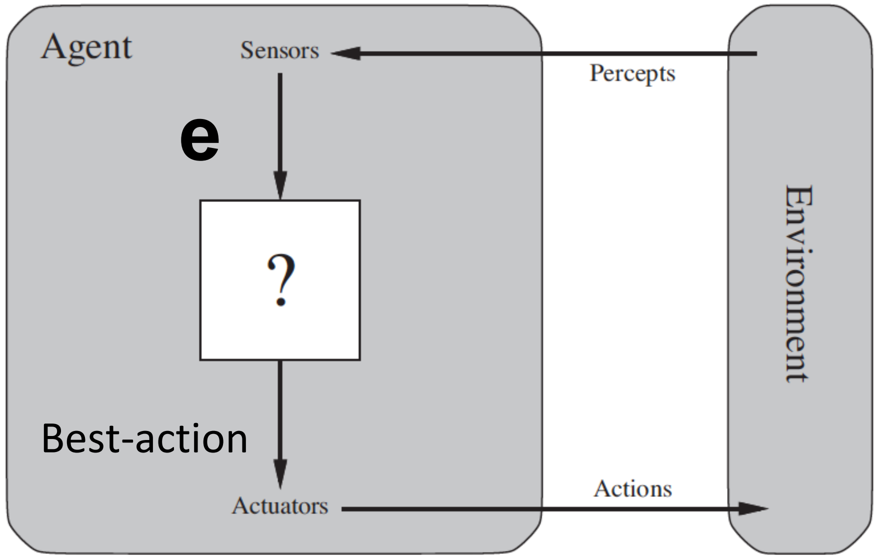

Utility-based Agent

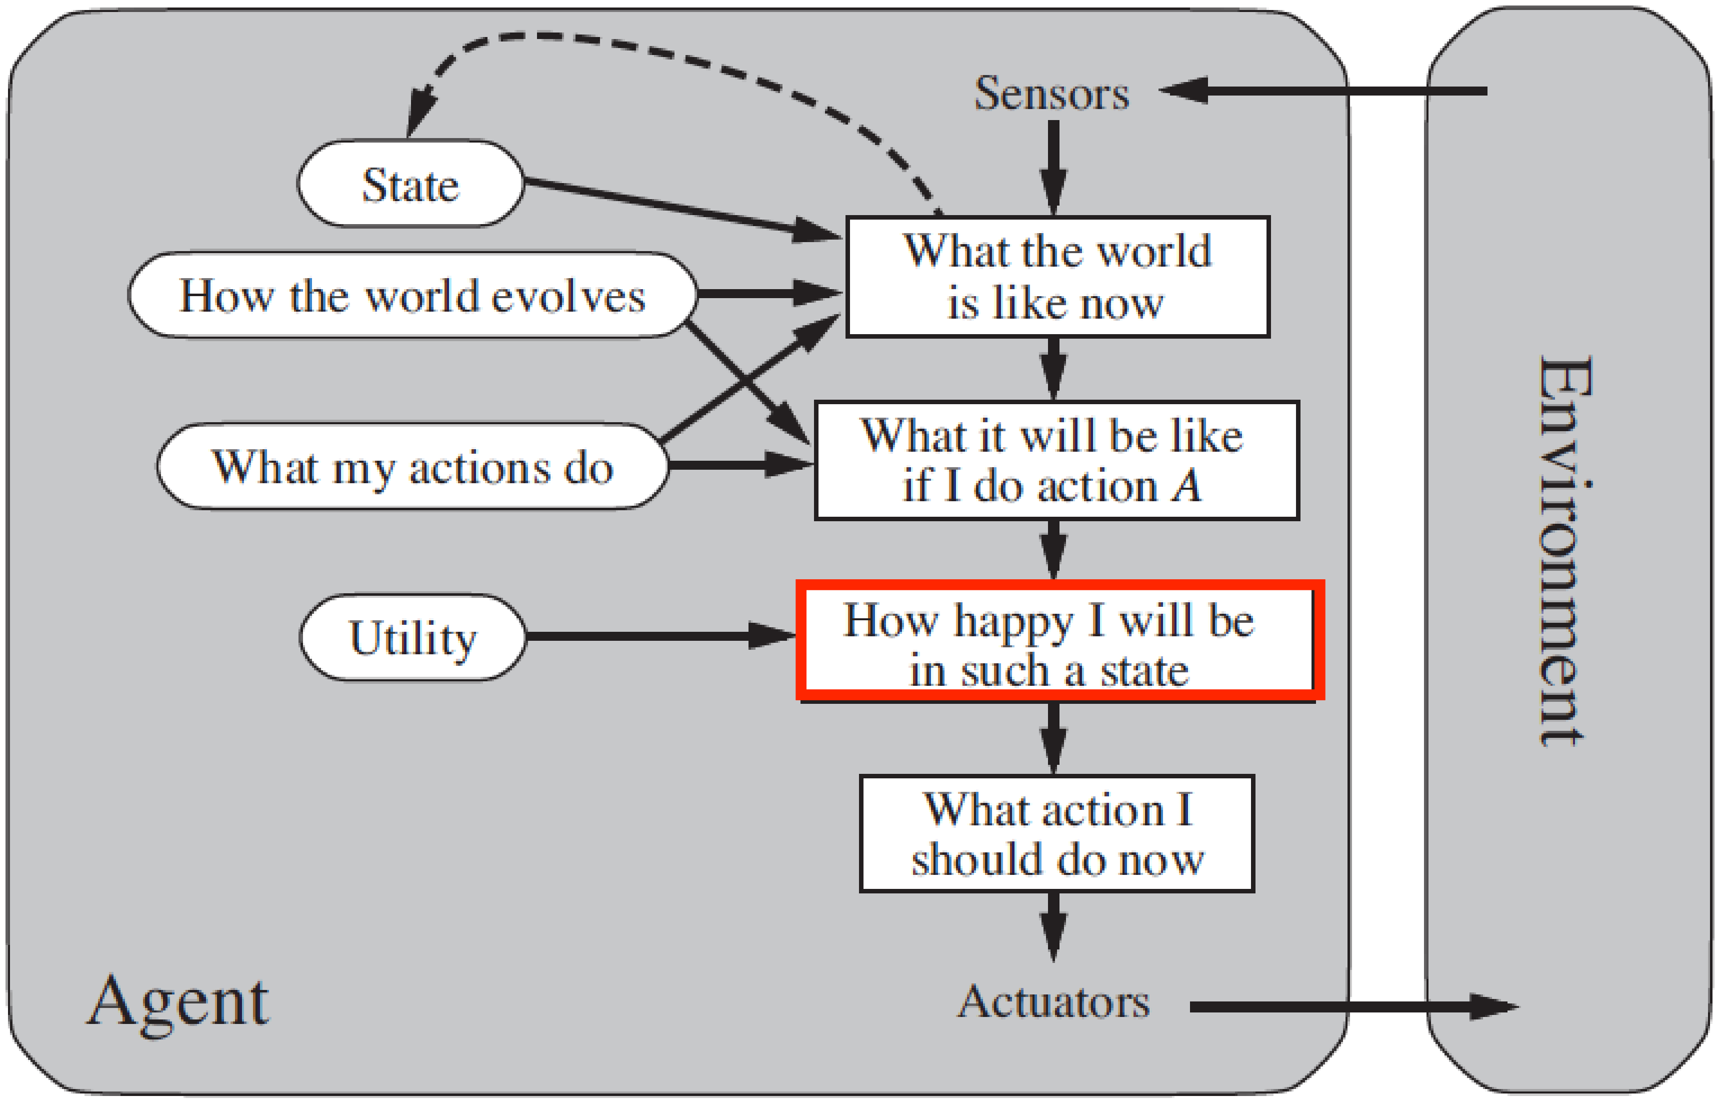

Decision Theory
+
Probabilities

Maximizing Expected Utility of an action

Internalization of Performance measure

Utility-based Agent

- Use of utility function that maps state (or state sequences) into real numbers
- Permits more fine-grained reasoning about what can be achieved, what are the trade-offs, conflicting goals, etc.

Learning Agent

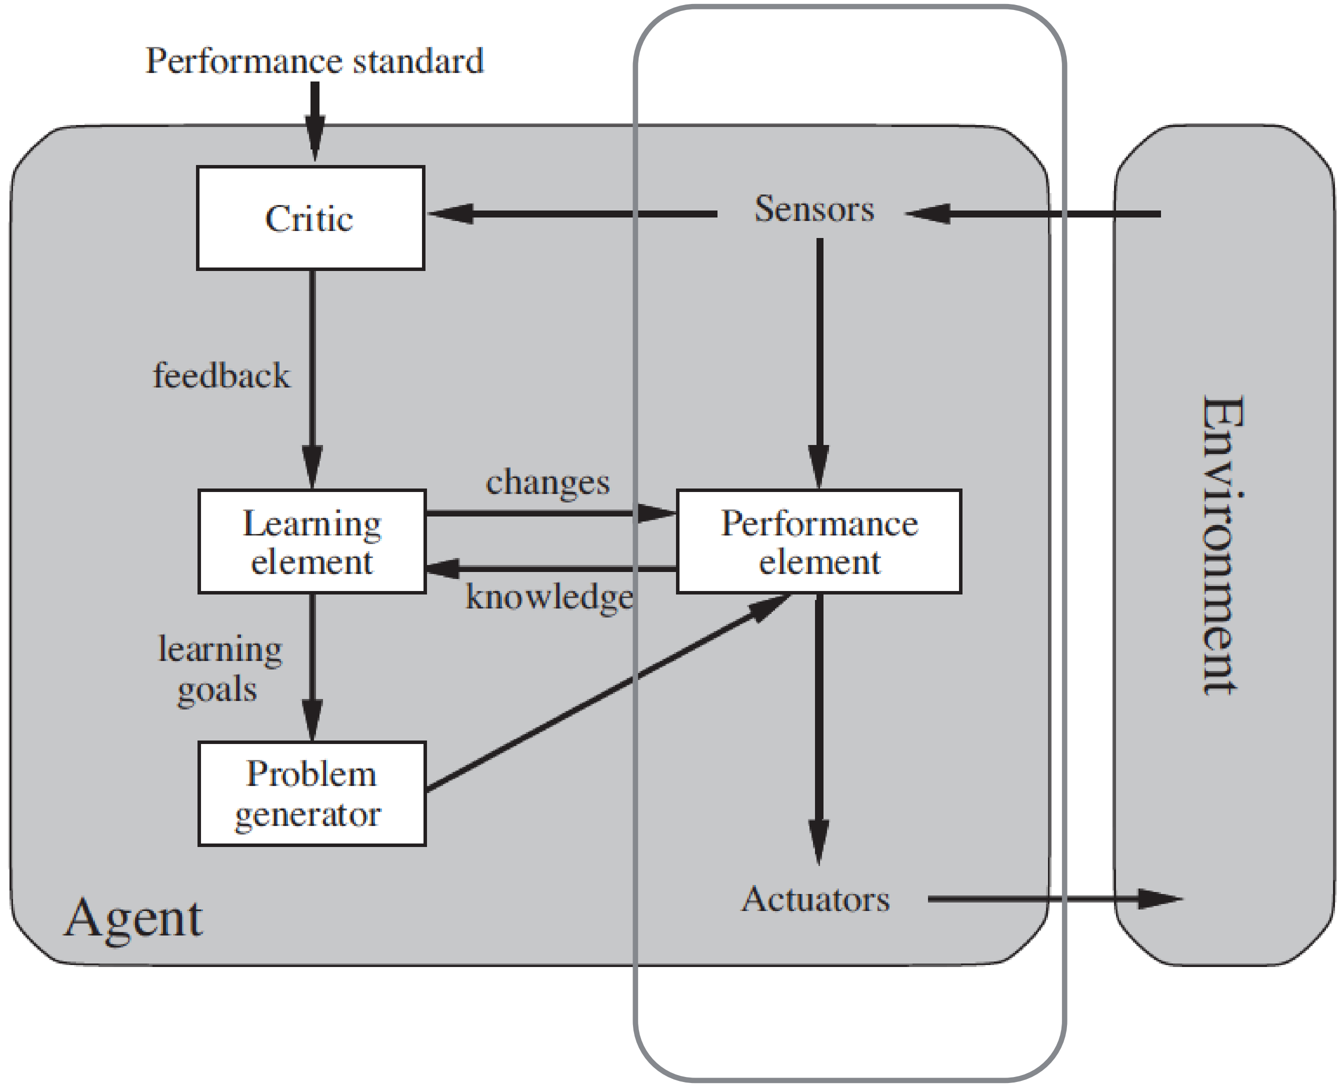

Previously the entire agent

- Bayesian Learning
- Clustering
- Classification
- Reinforcement Learning
- NN / Deep learning

Learning Agent:

- Has the ability to modify behavior for the better based on experience.
- It can learn new behaviors via exploration of new experiences

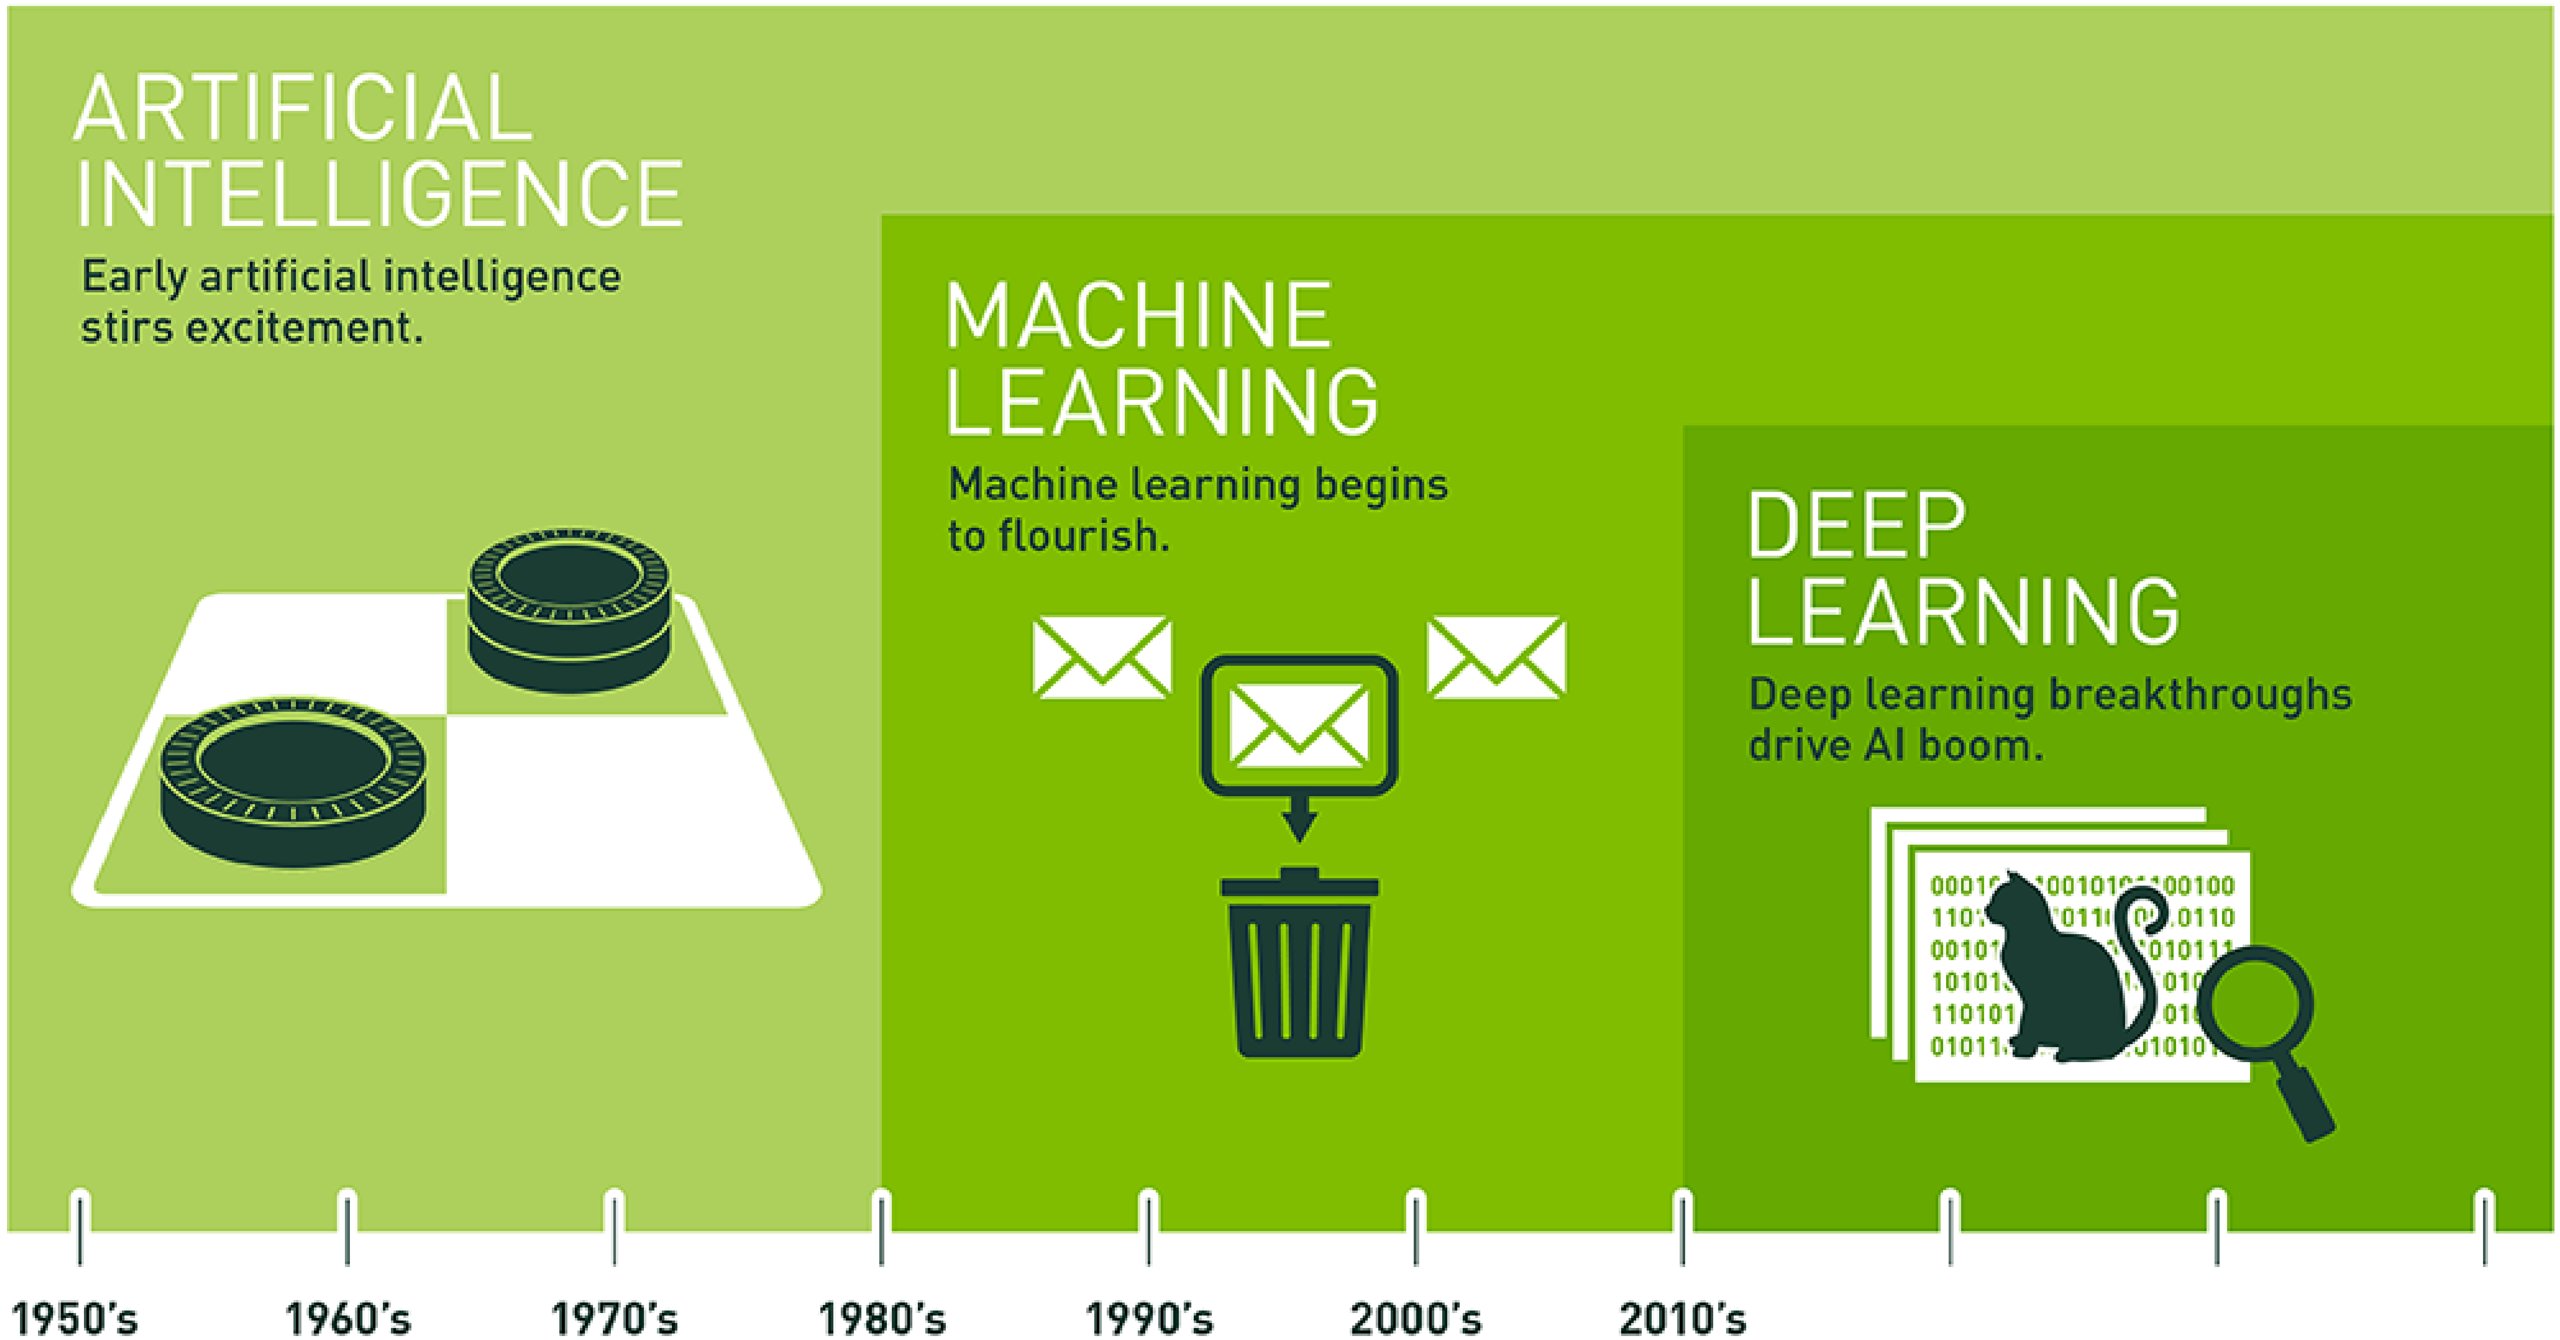

Since an early flush of optimism in the 1950s, smaller subsets of artificial intelligence – first machine learning, then deep learning, a subset of machine learning – have created ever larger disruptions.

The Essence of Learning?

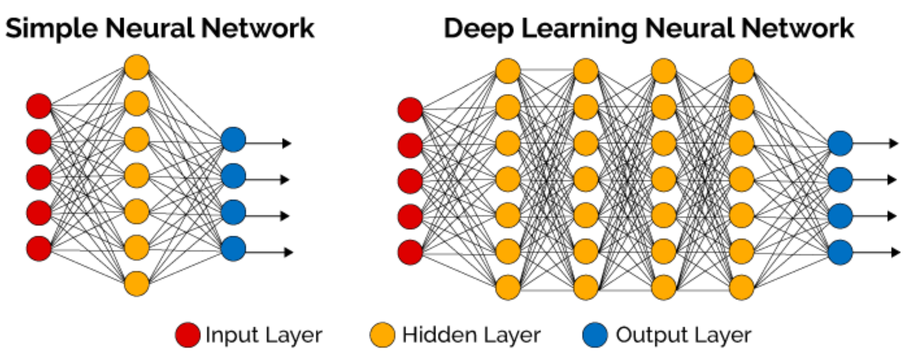

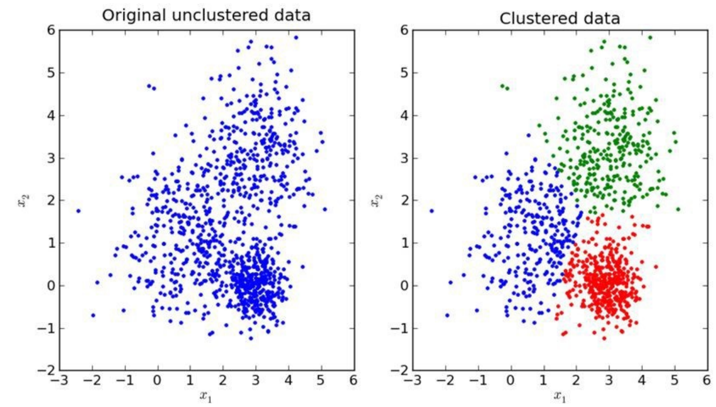

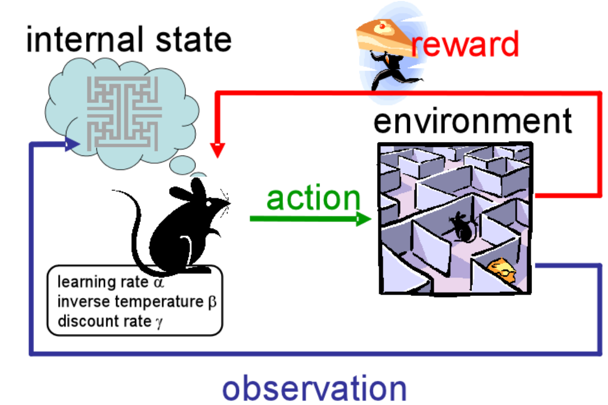

Function Approximation/ Classification

Identifying Patterns or Trends

Optimizing State-action Policies

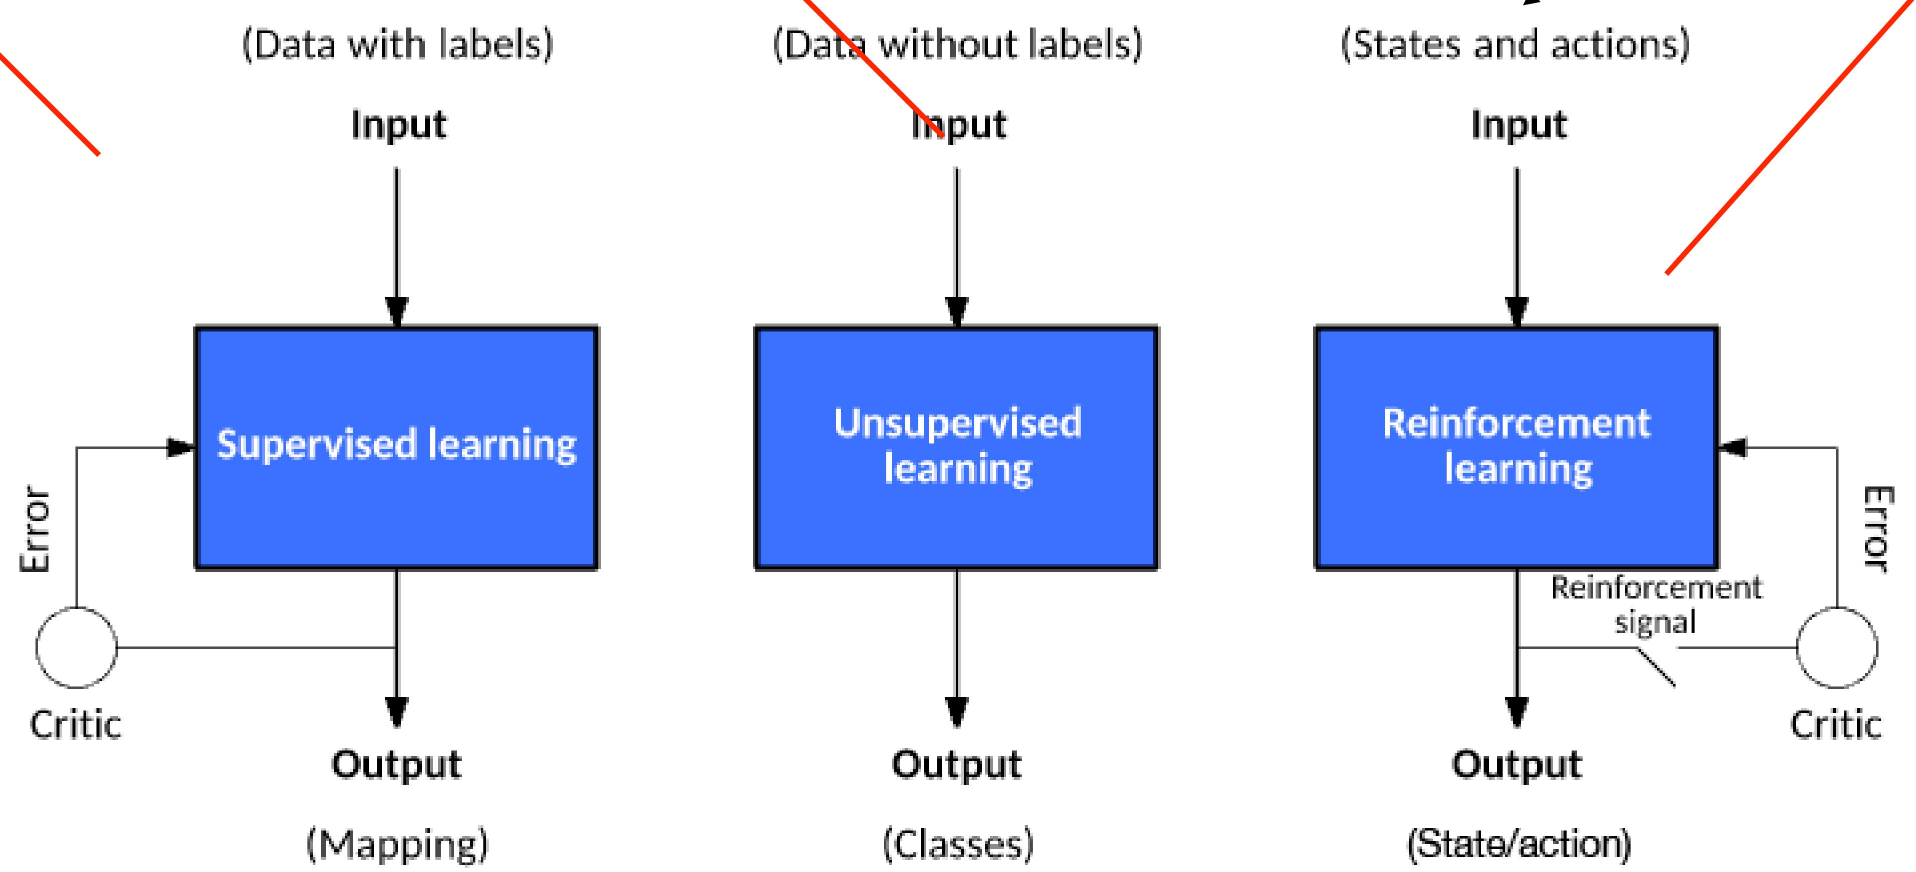

LINKÖPING UNIVERSITY

Did you ever watch a baby learn to walk?

The Emergence of Locomotion

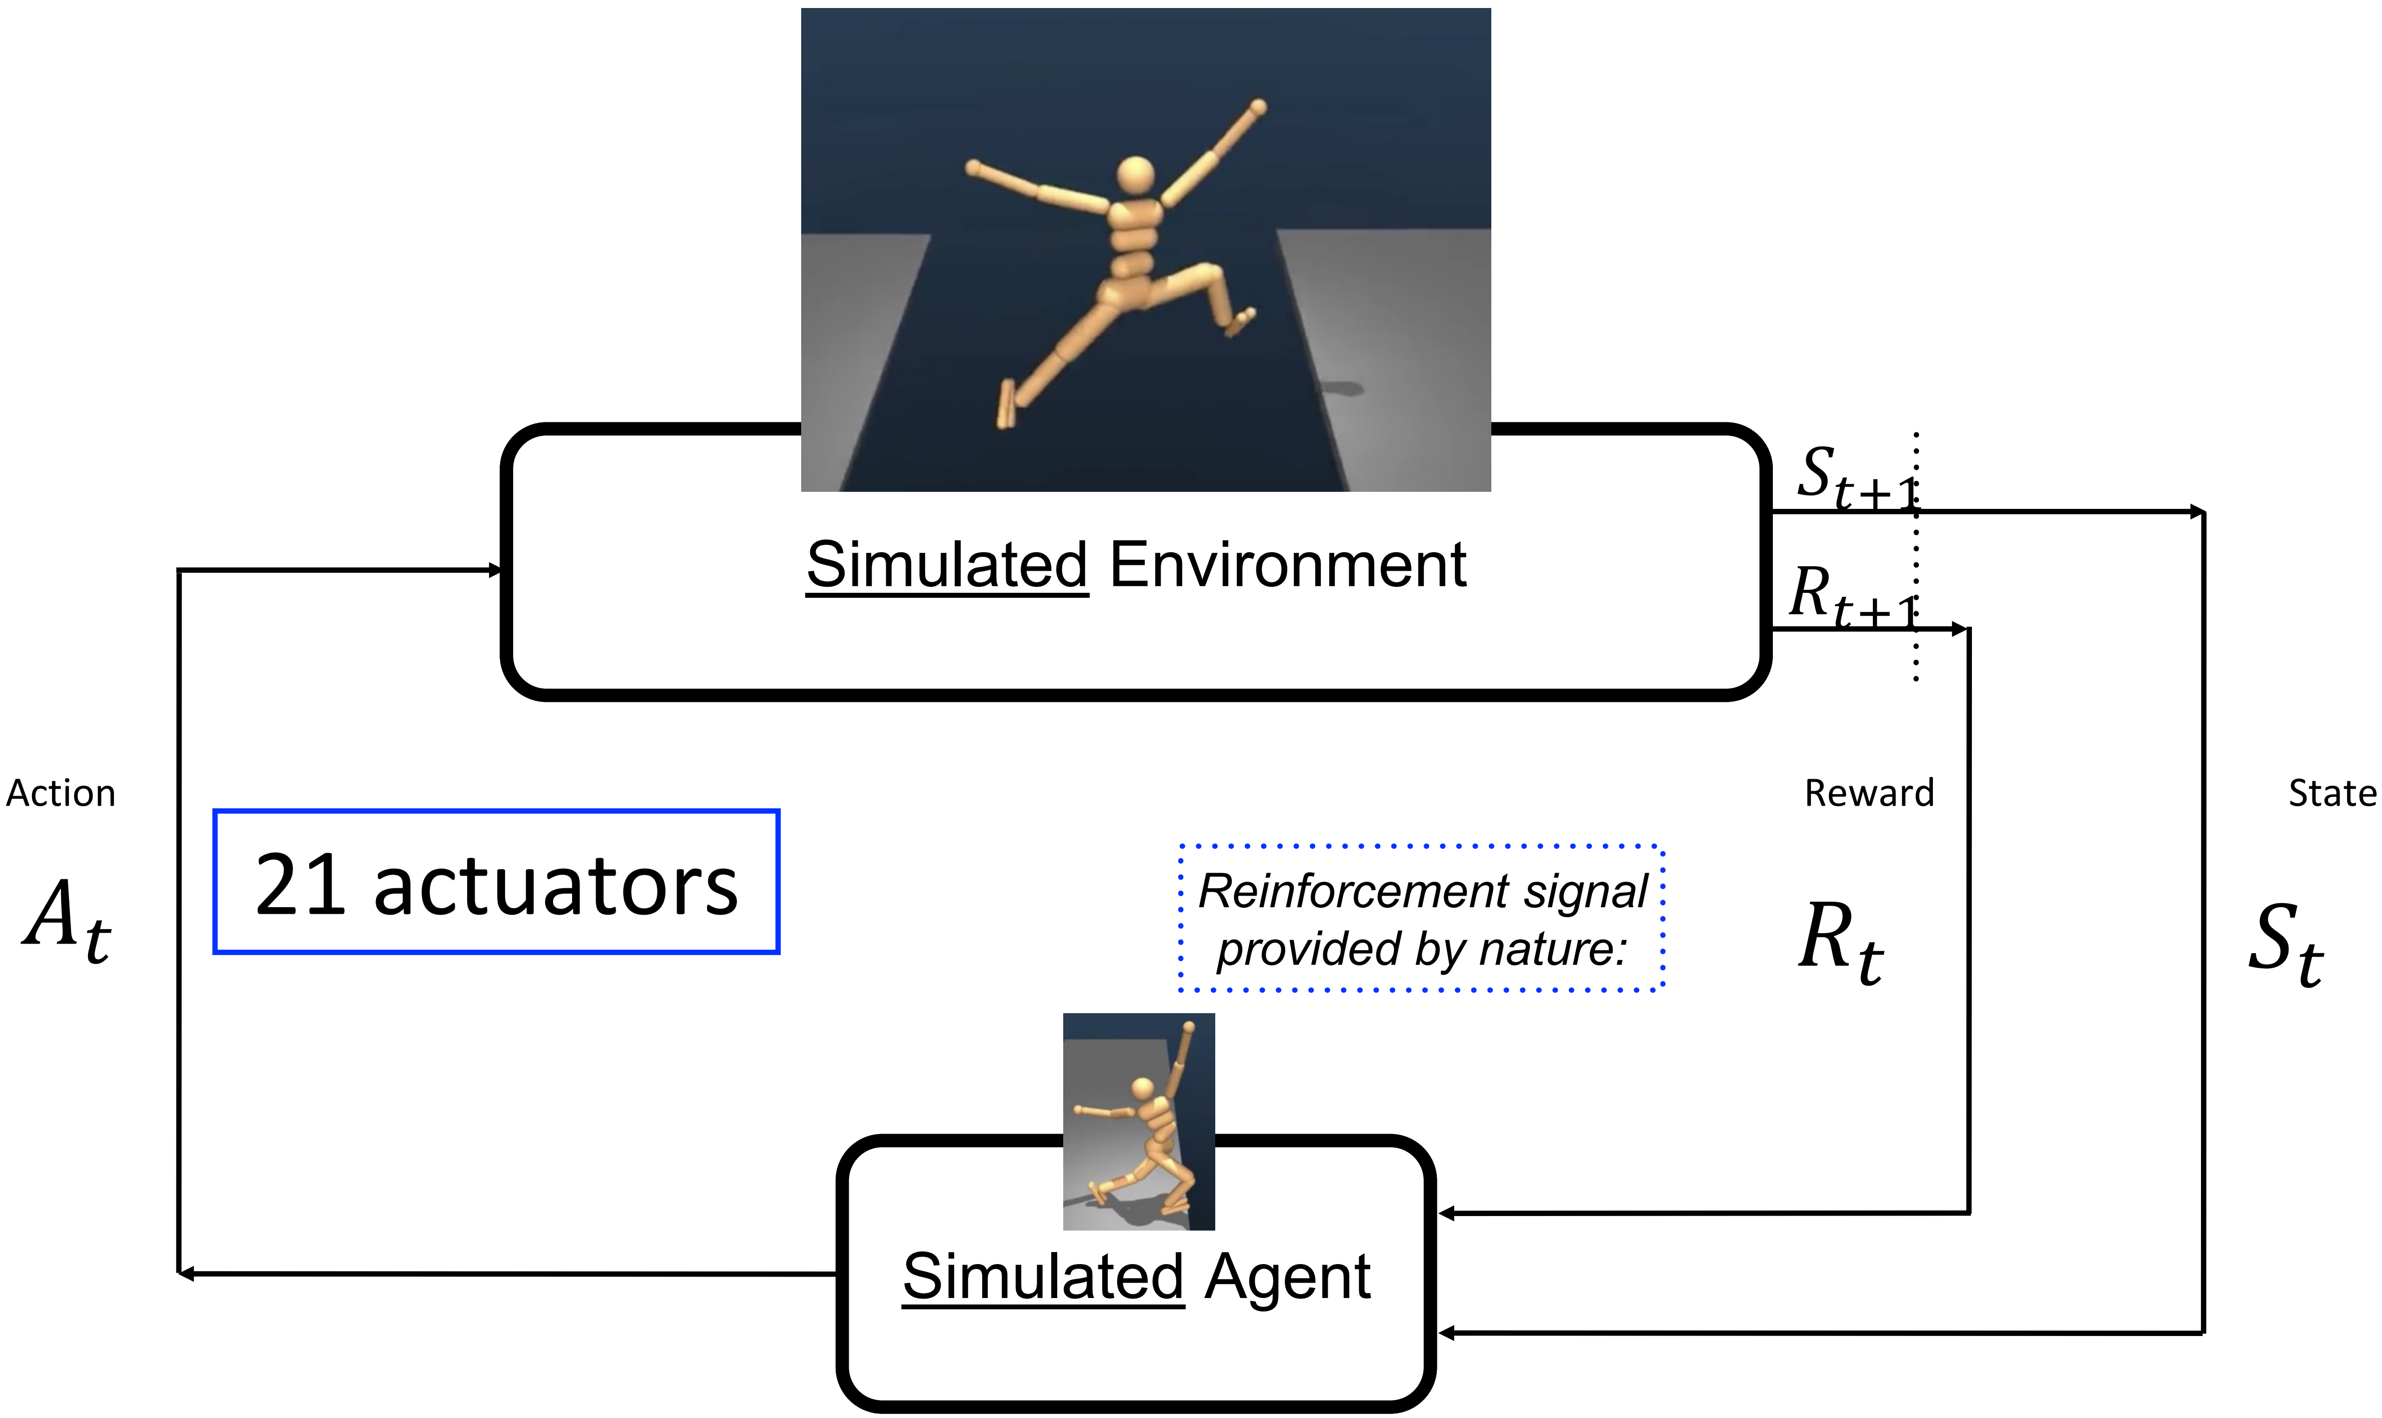

I only get a reward if I move forward!

The Emergence of Locomotion

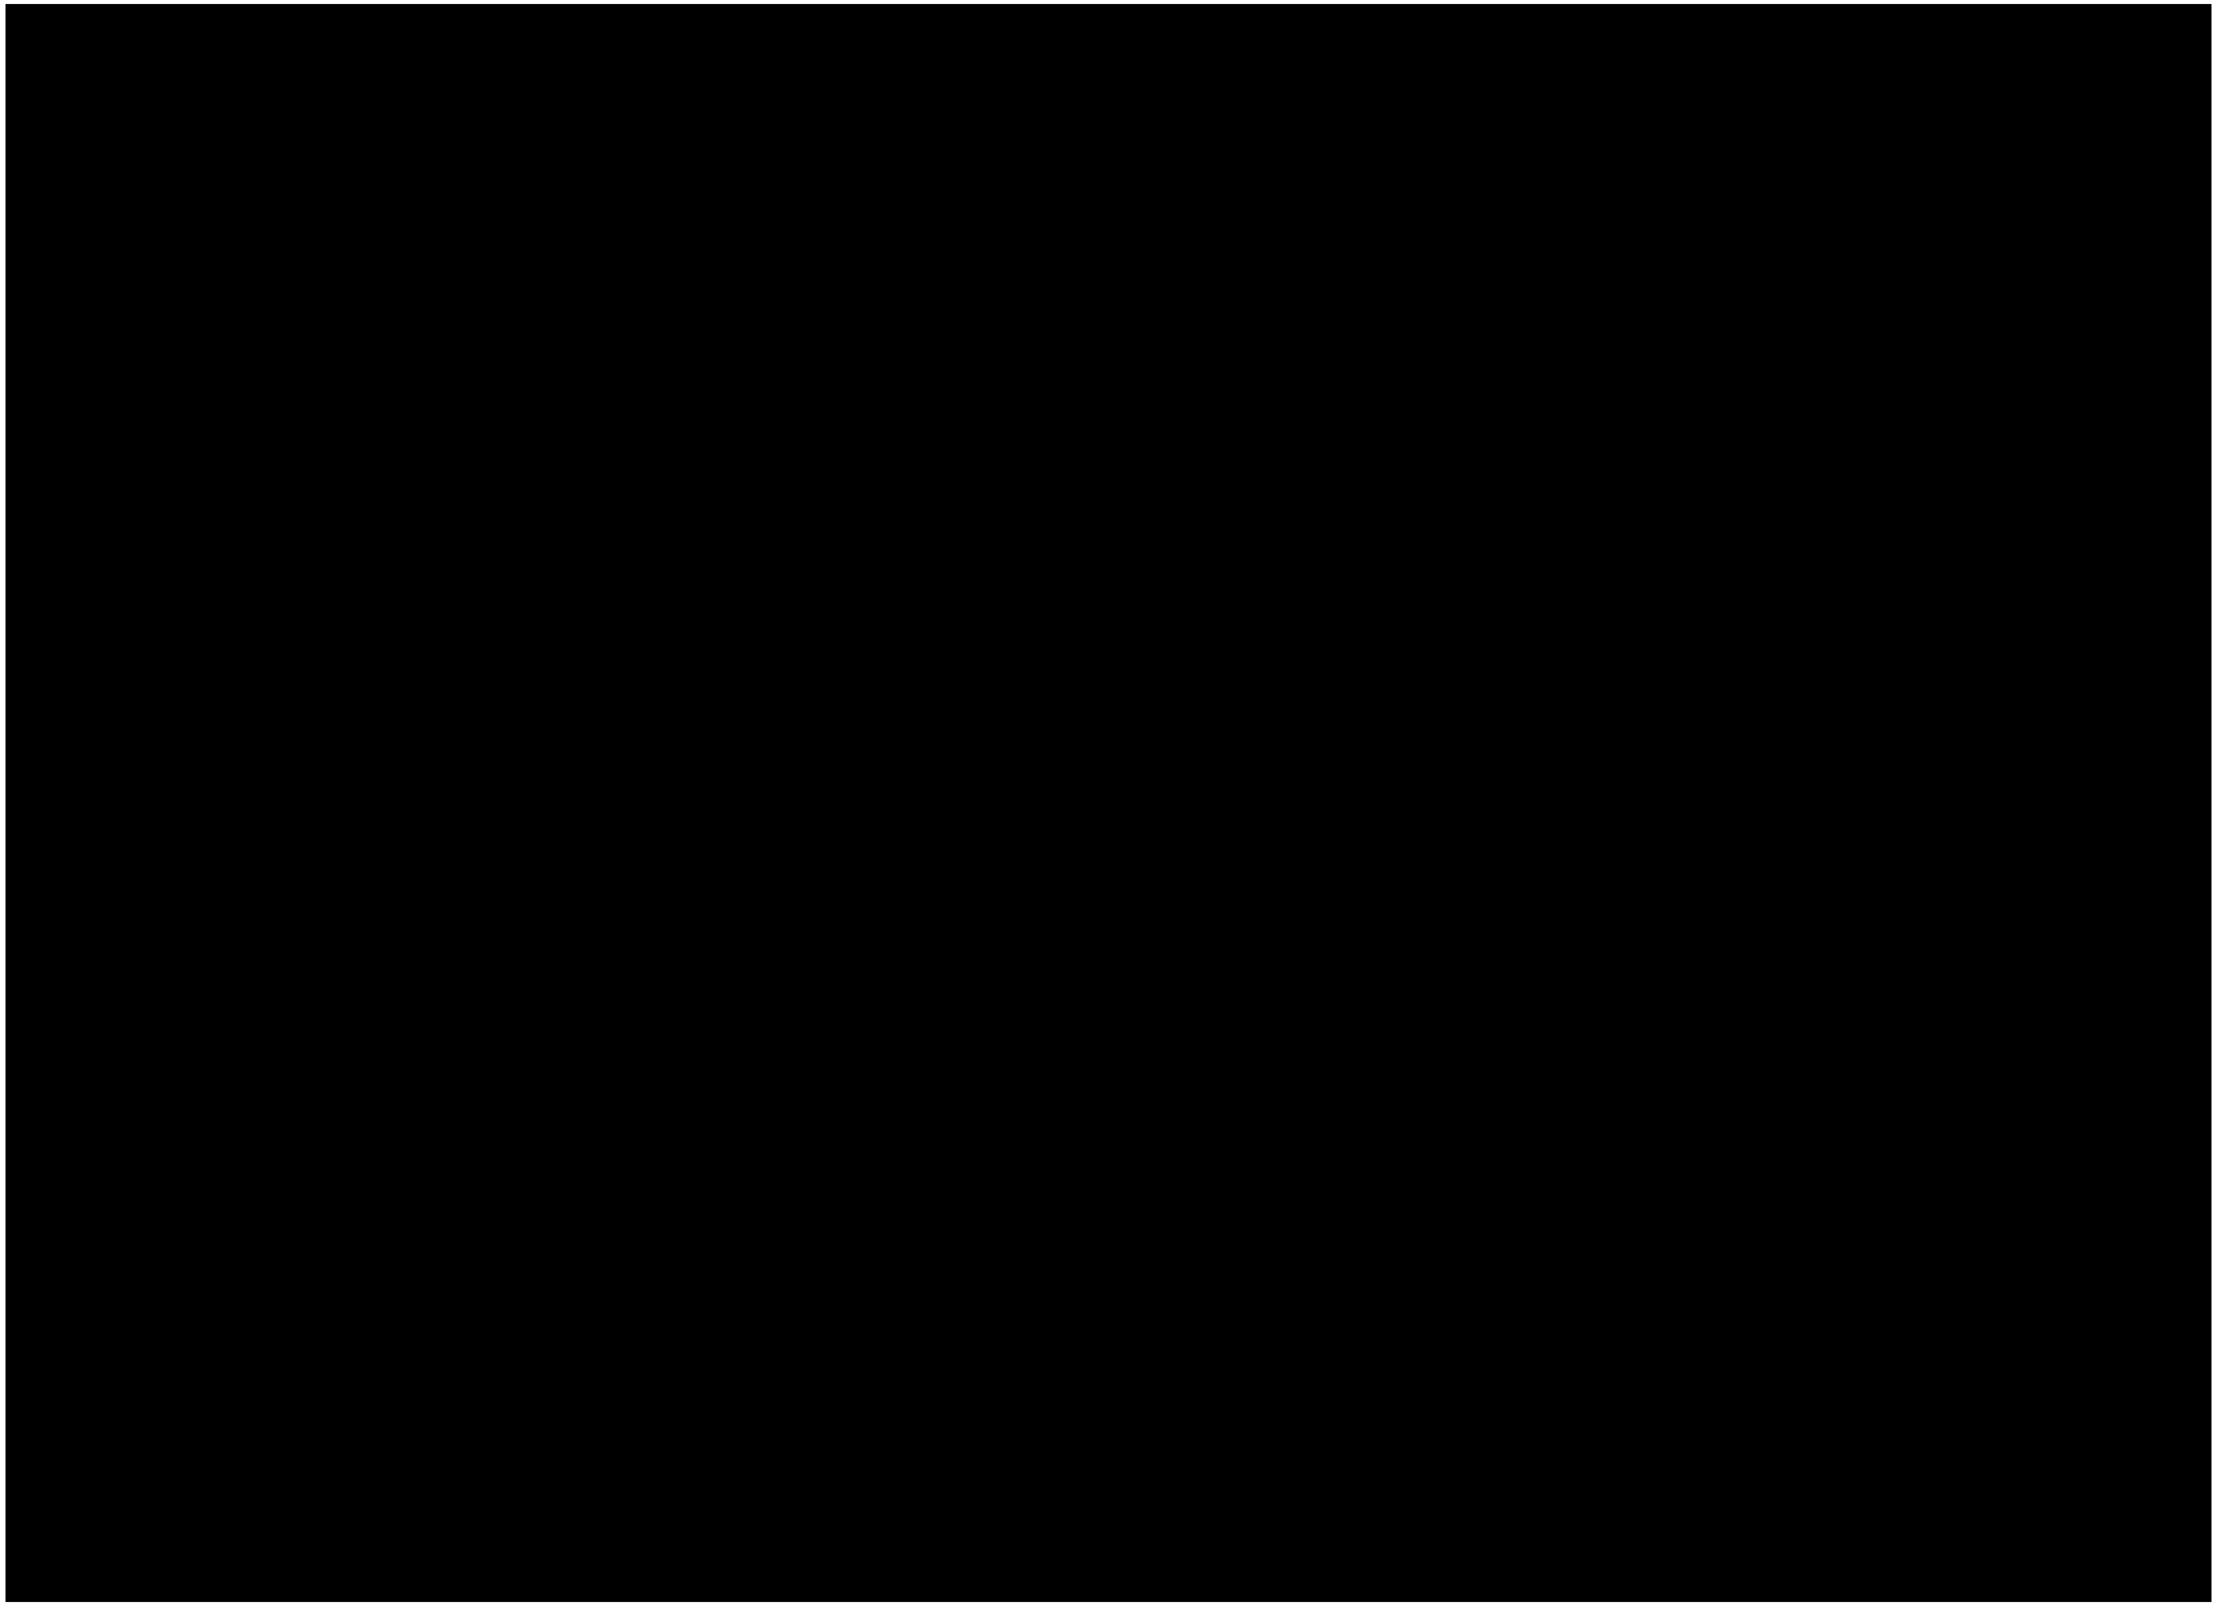

Google Deep Mind

Flipping Pancakes

RL and Utility: Learning and Maximizing Expected Utility

RL results in a trajectory:

$$S_0,$$
$$A_0, R_1, S_1,$$
$$A_1, R_2, S_2,$$
$$A_2, R_3, \ldots$$

Goal: The agent should learn an optimal policy from its experience that determines which action to execute in each state to maximise its accumulated rewards.

Given an experience: $s_t, a_t, r_{t+1}, s_{t+1}$ Update its value to the agent

$$Q[s_t, a_t] = Q[s_t, a_t] + \alpha * (r_{t+1} + \gamma * \max_{a \in \mathcal{A}(s_{t+1})} Q[s_{t+1}, a]) - Q[s$$

New Utility | Previous Utility | New Reward | Utility of best action to fire in new state - Previous Utility

Learning Rate

Discount Factor

Temporal Difference

"Deep" Learning?

- A **deep neural network** has more than one hidden layer
  - Basic ingredients described in 1962
  - "**Deep learning**" term introduced 1986 (Rina Dechter)
  - Problem: Computational power!
    - Progress: GPU-based systems
    - Later: Explicit hardware for neural network ops
    - Apple A11: 0.6 TOPS (trillion ops/second)
    - Apple A16: 17 TOPS

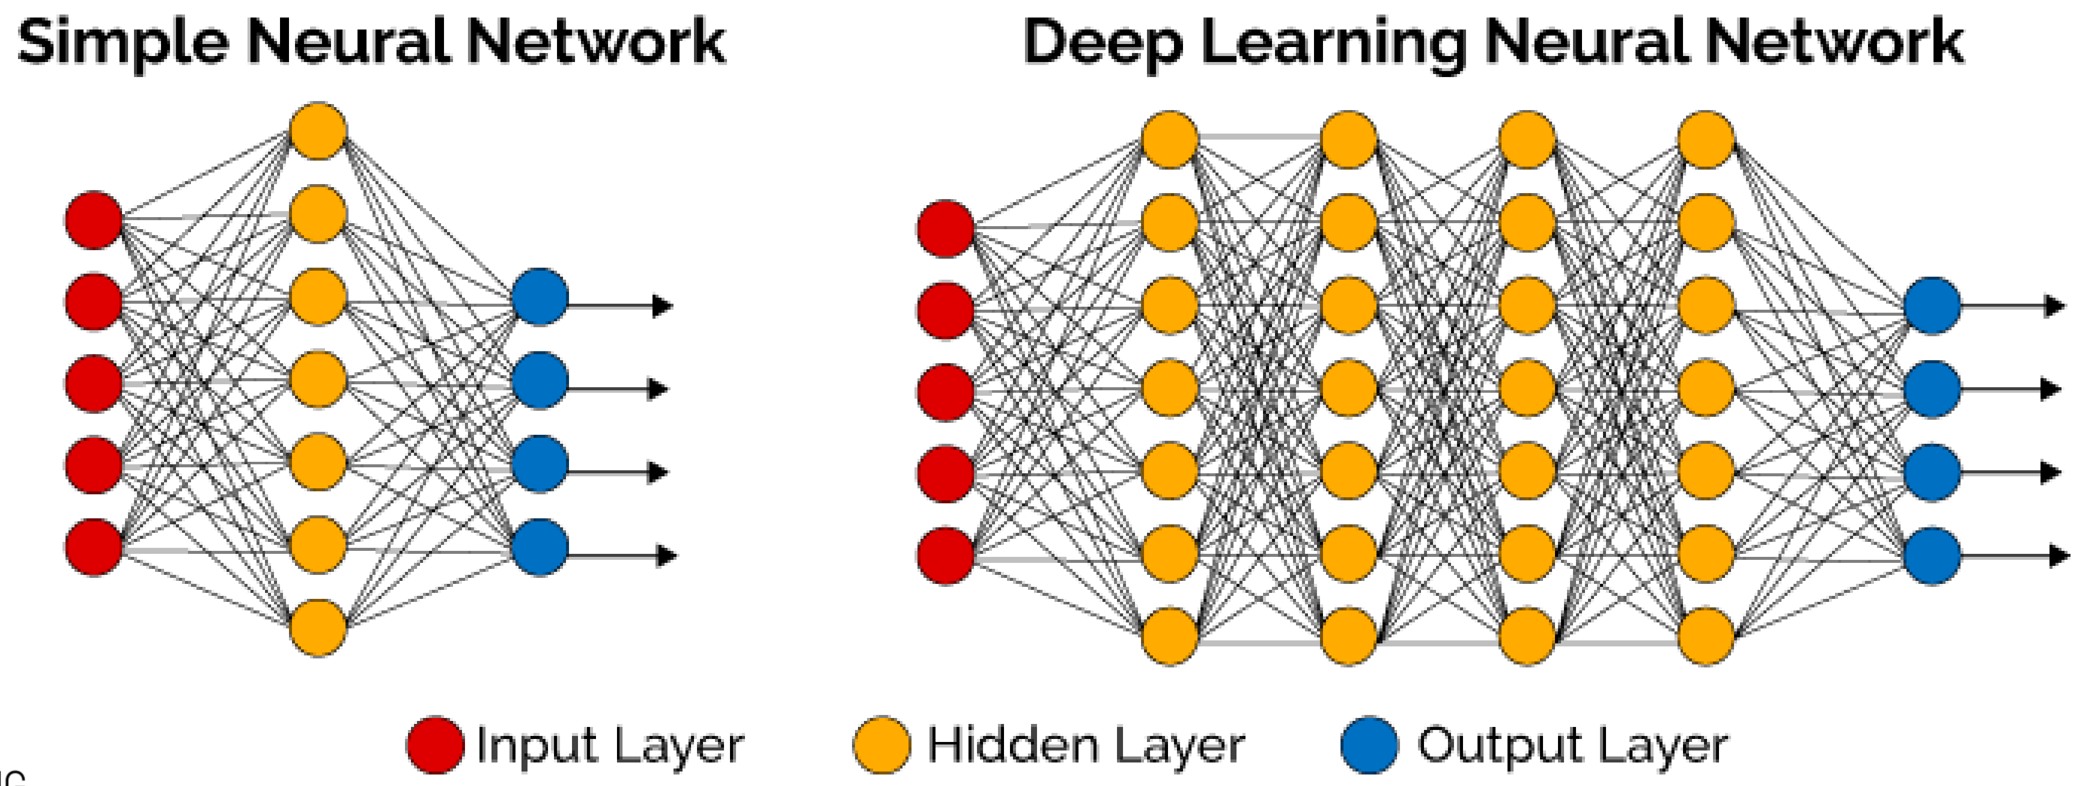

Advances and Advantages of Deep Learning

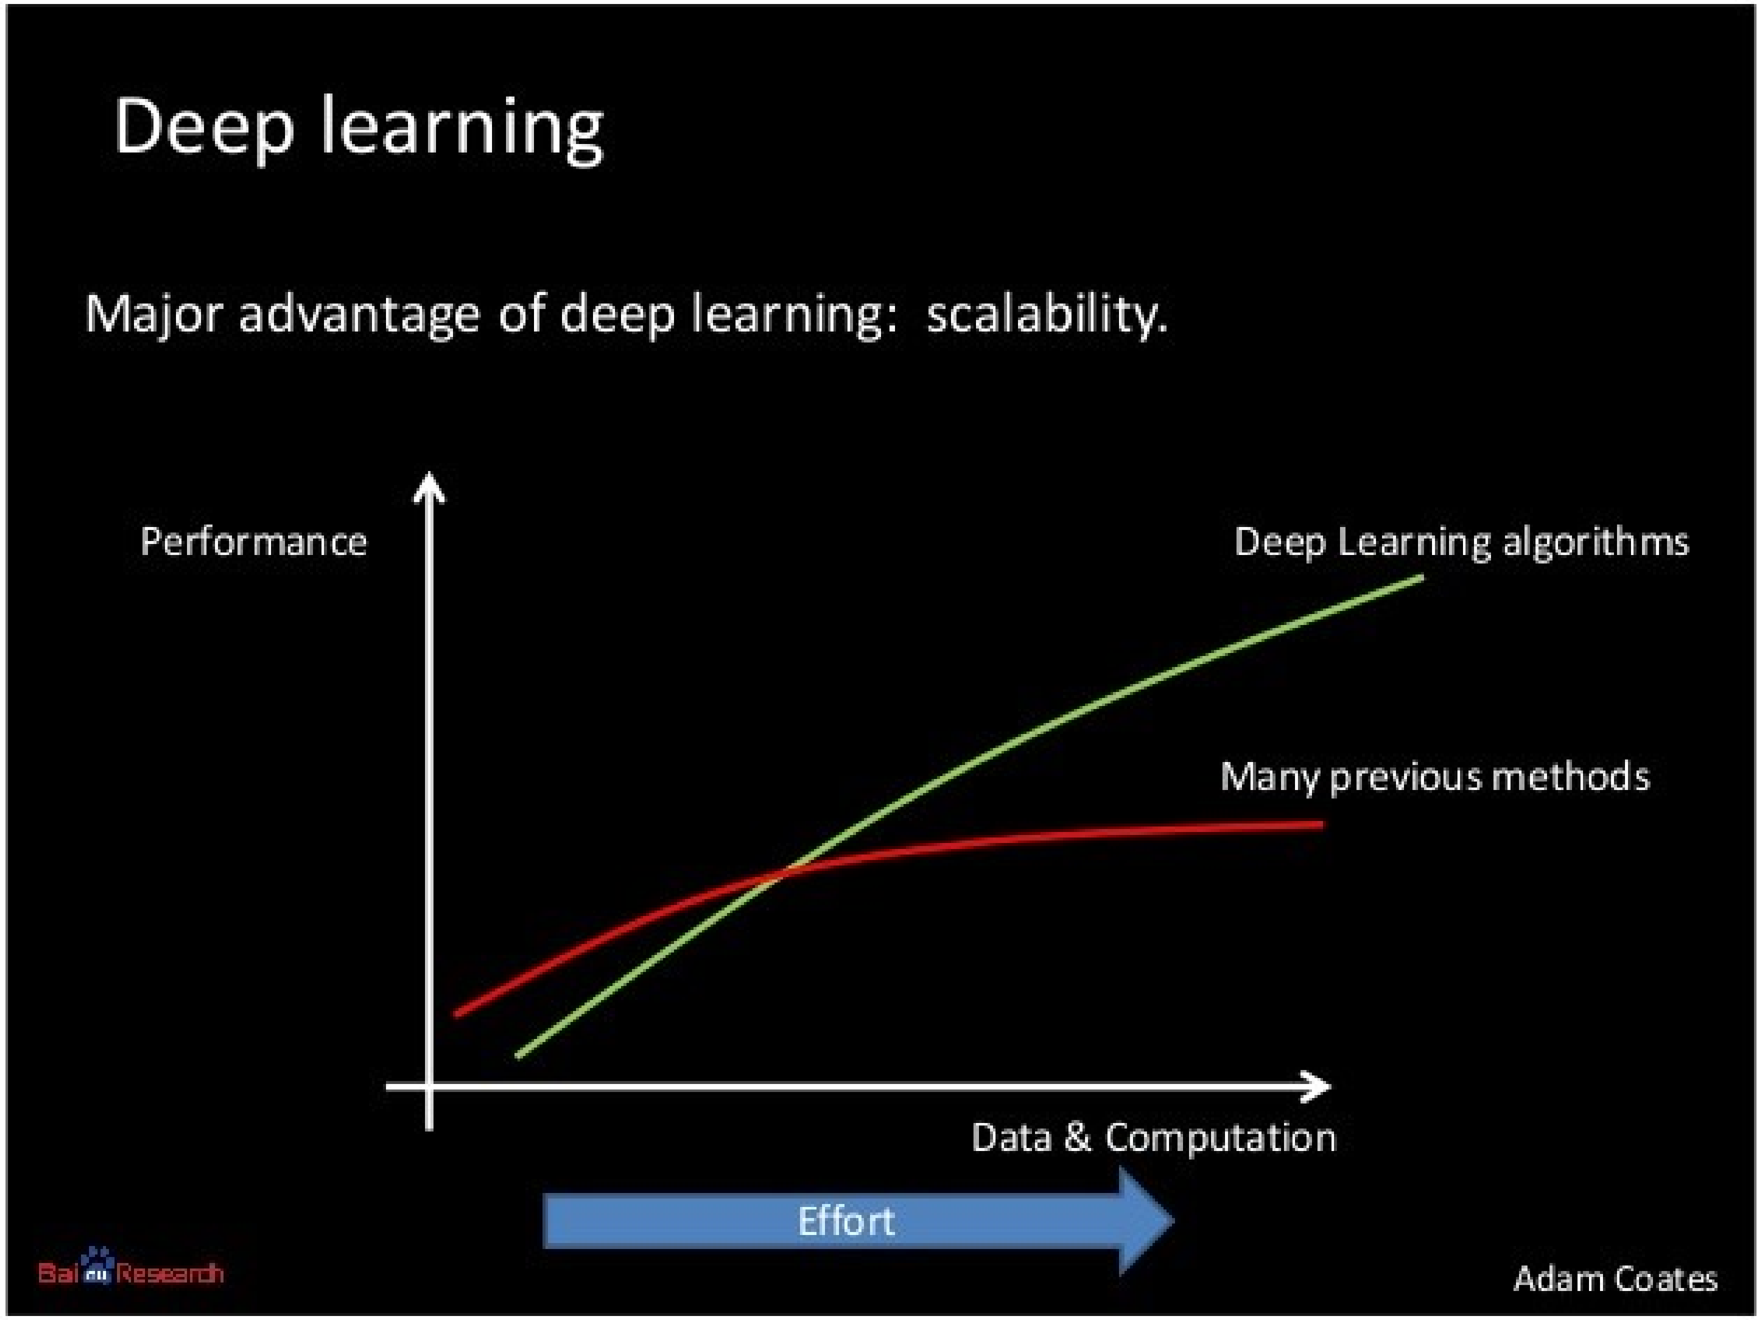

Large amounts
Data

Large amounts
Processing
Power

LINKÖPING
UNIVERSITY

Major Applications: Image Processing

The Deep Learning “Computer Vision Recipe”

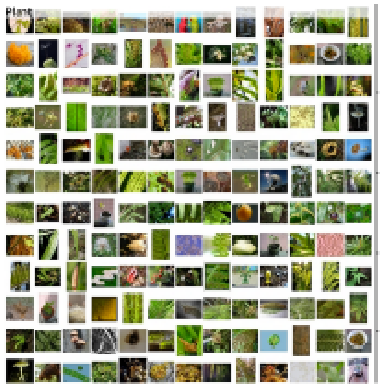

\+

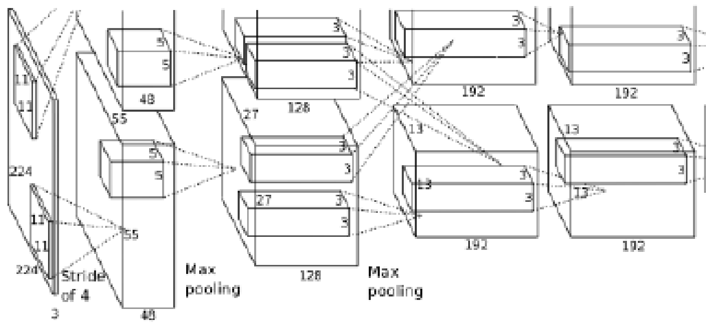

\+

=

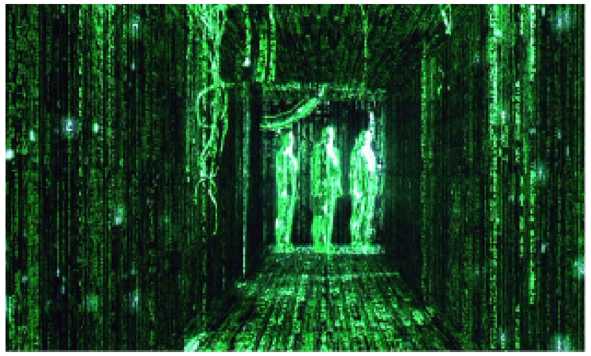

Big Data: ImageNet | Deep Convolutional Neural Network | Backprop on GPU | Learned Weights

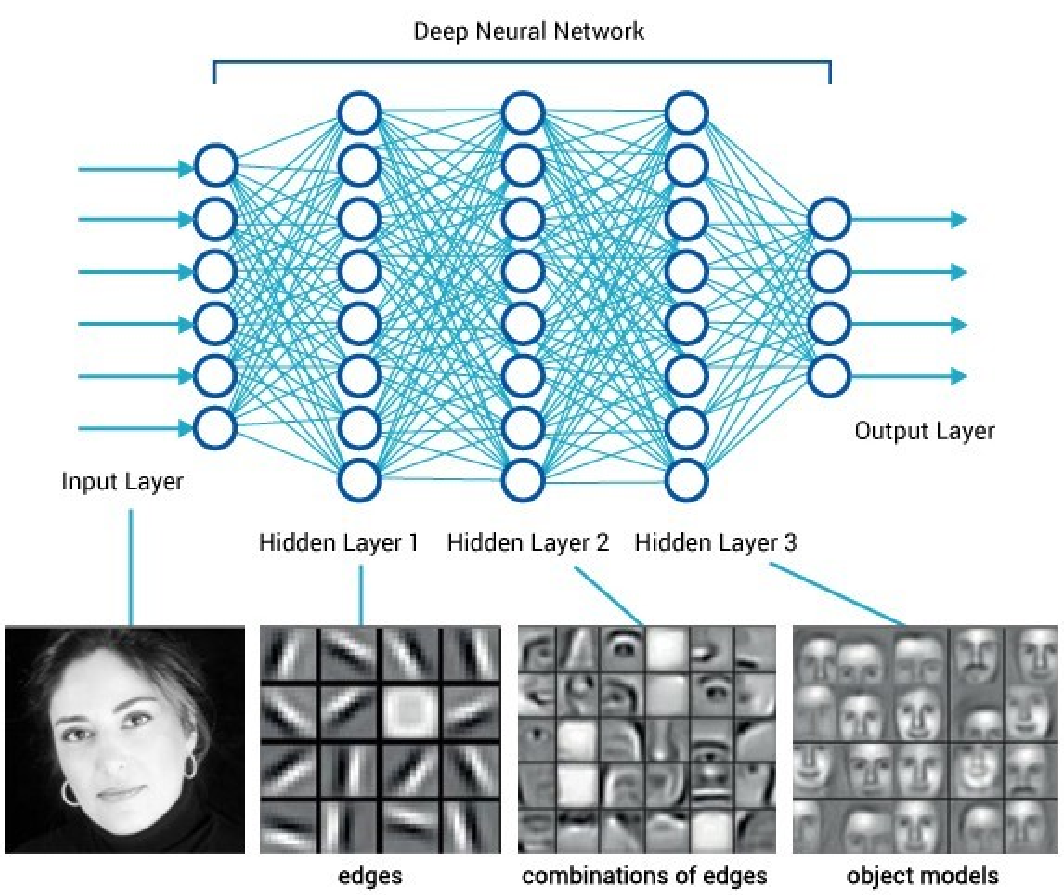

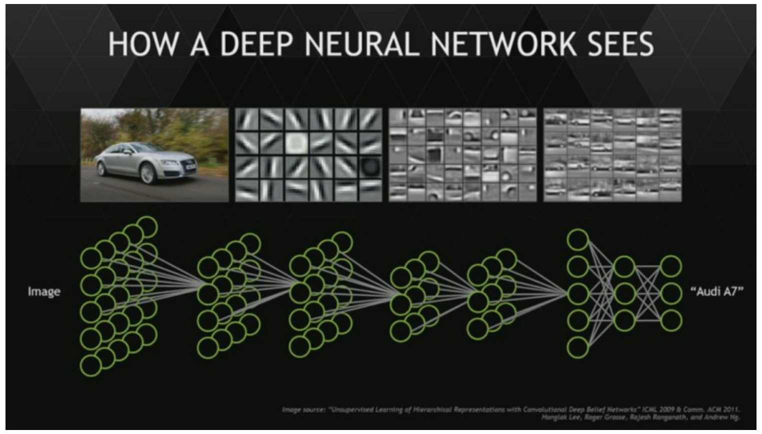

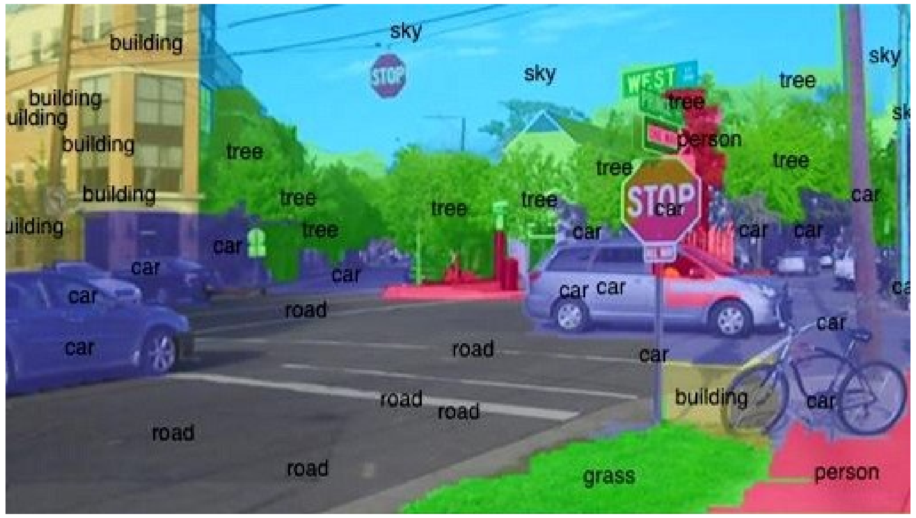

LINKÖPING UNIVERSITY

Major Applications: Natural Language Processing

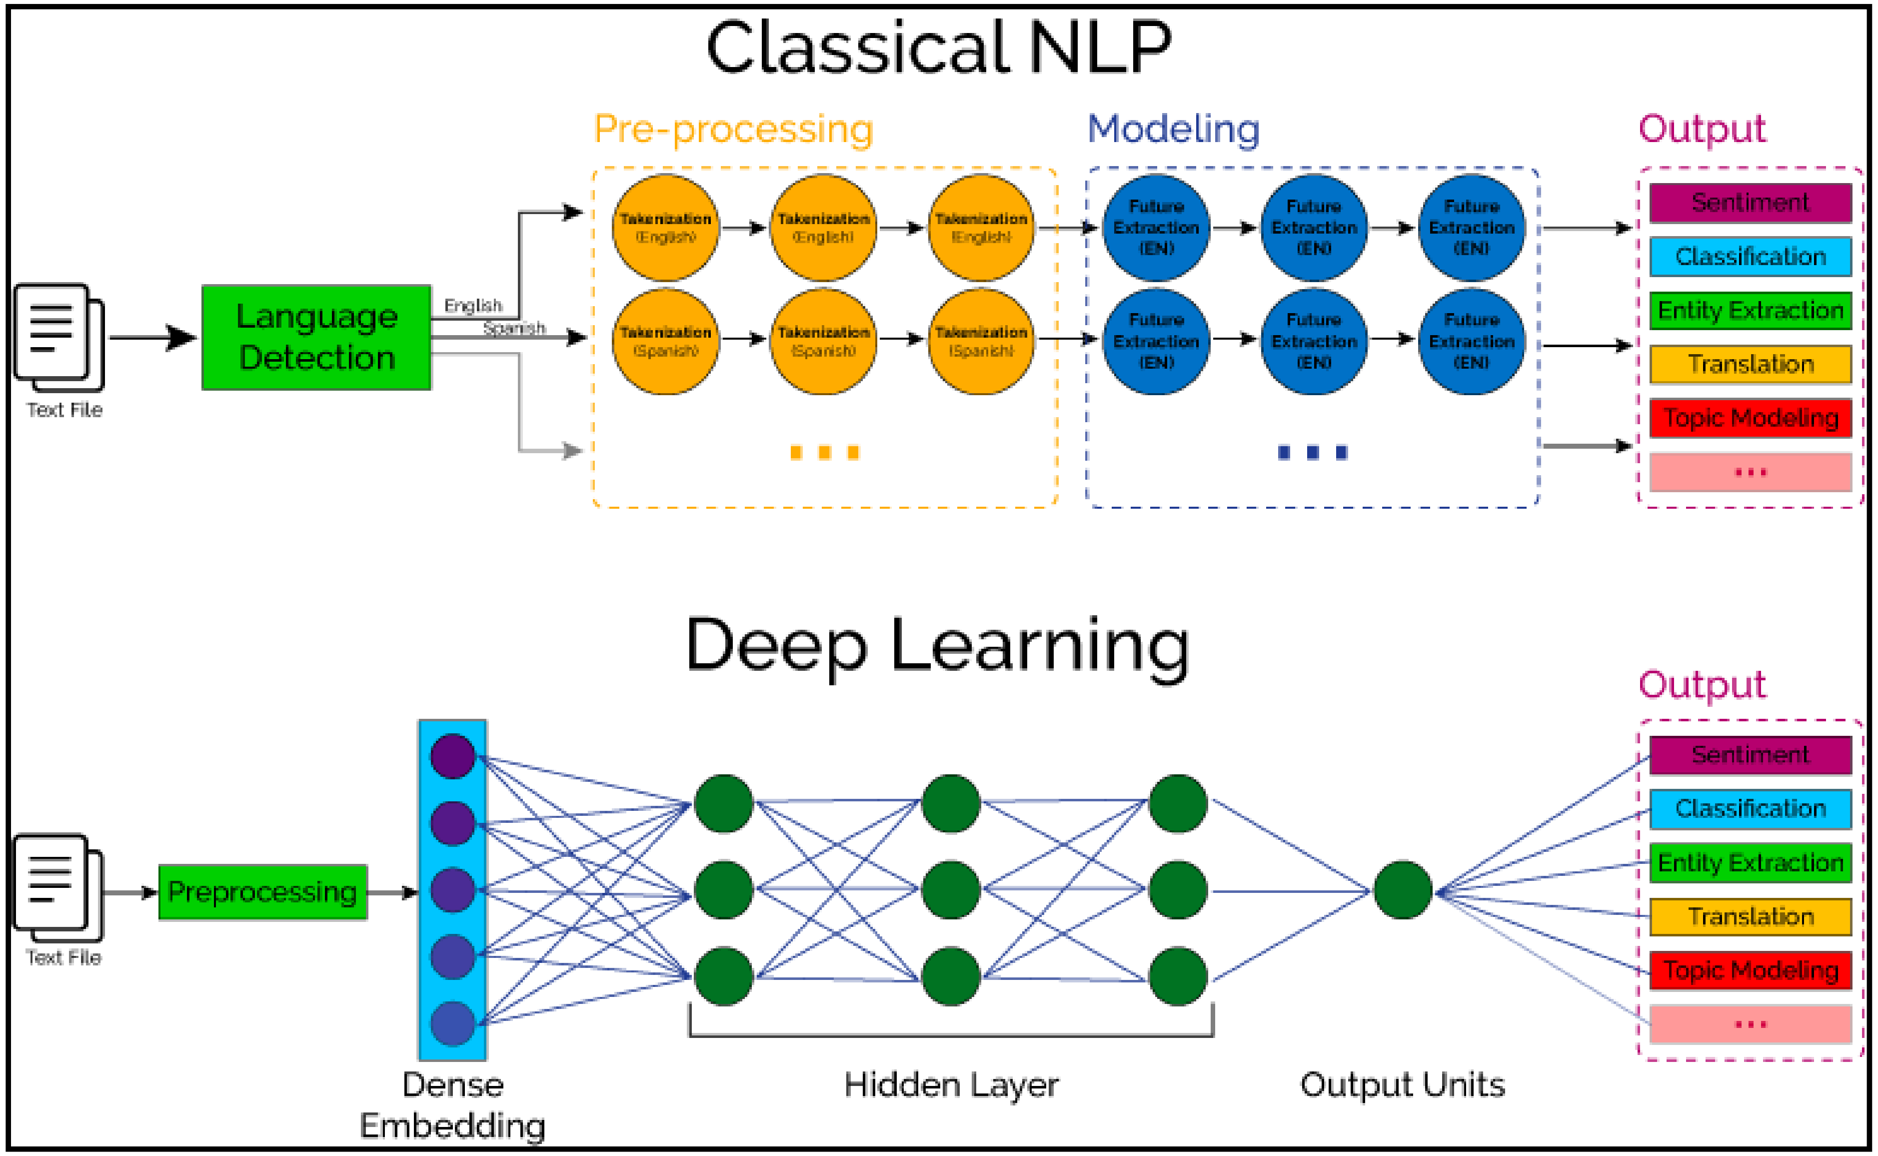

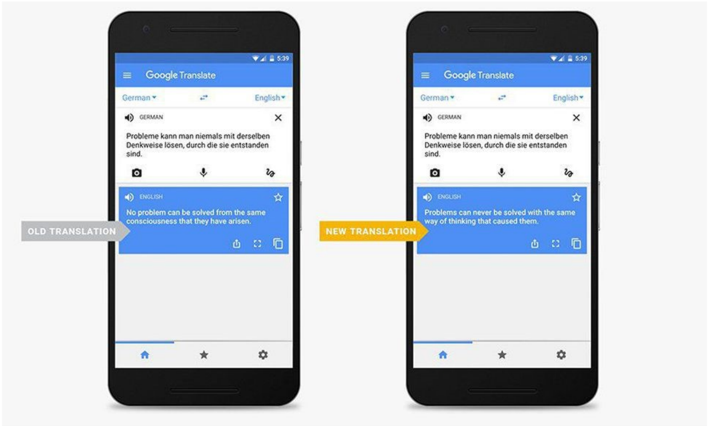

Google Translate

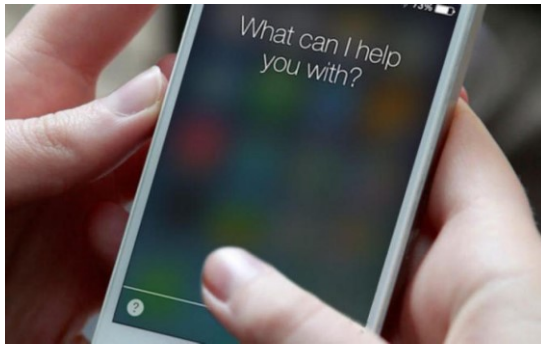

Siri

Alexa

LINKÖPING UNIVERSITY

Major Applications: Game Playing

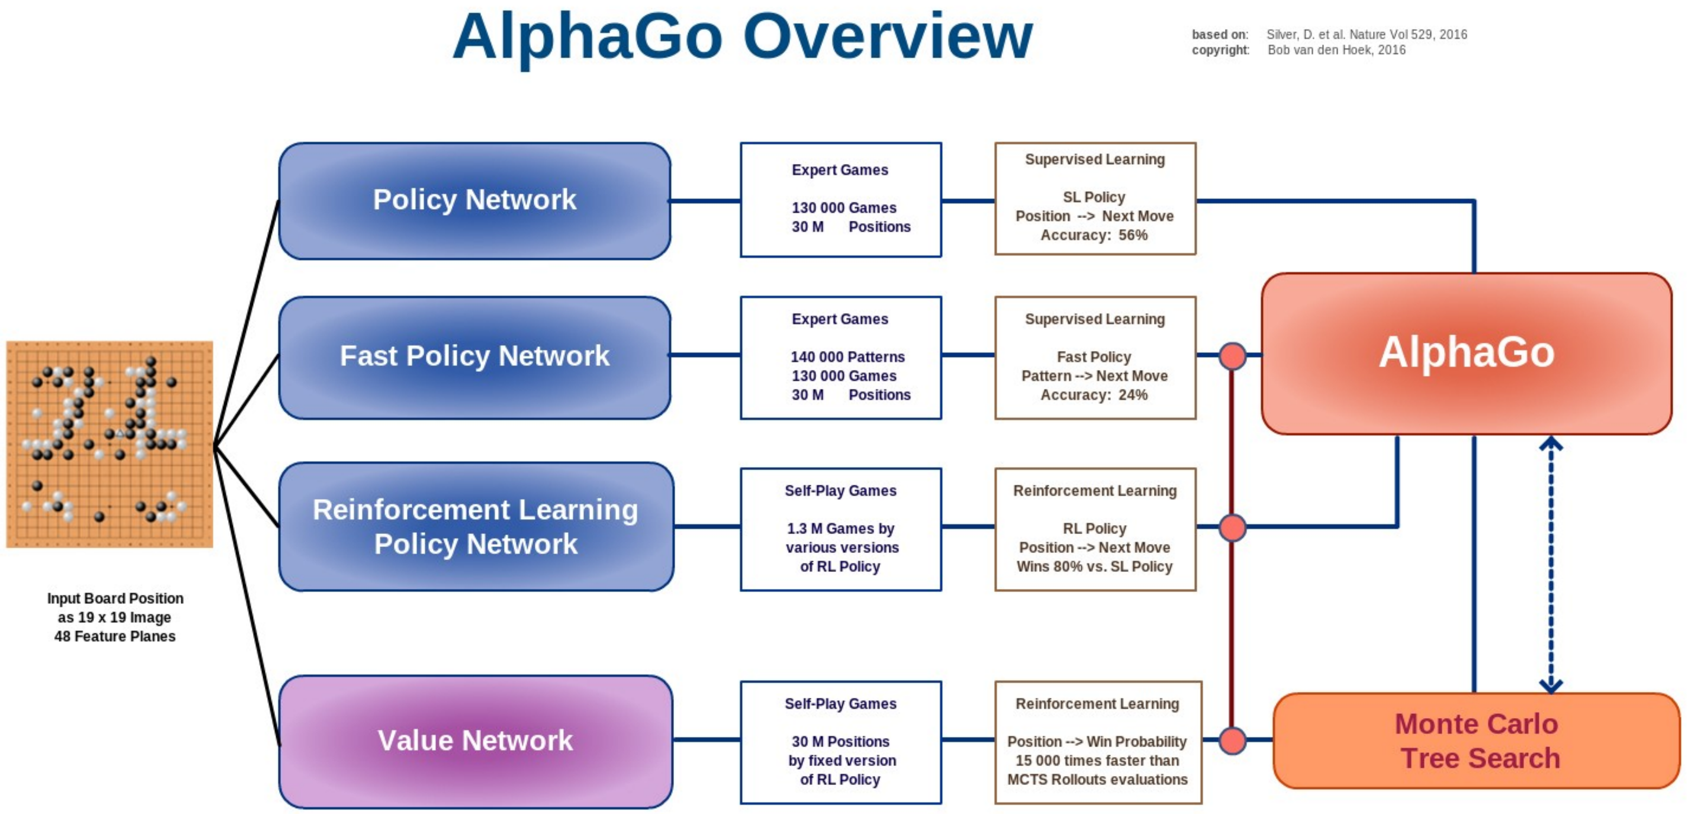

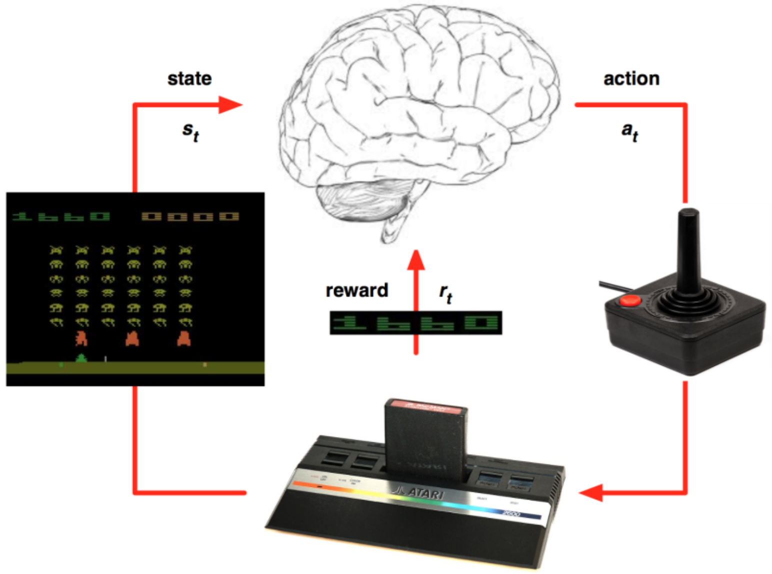

AI and Robotics

Robotics- Boston Dynamics: ATLAS & HANDLE

Integration

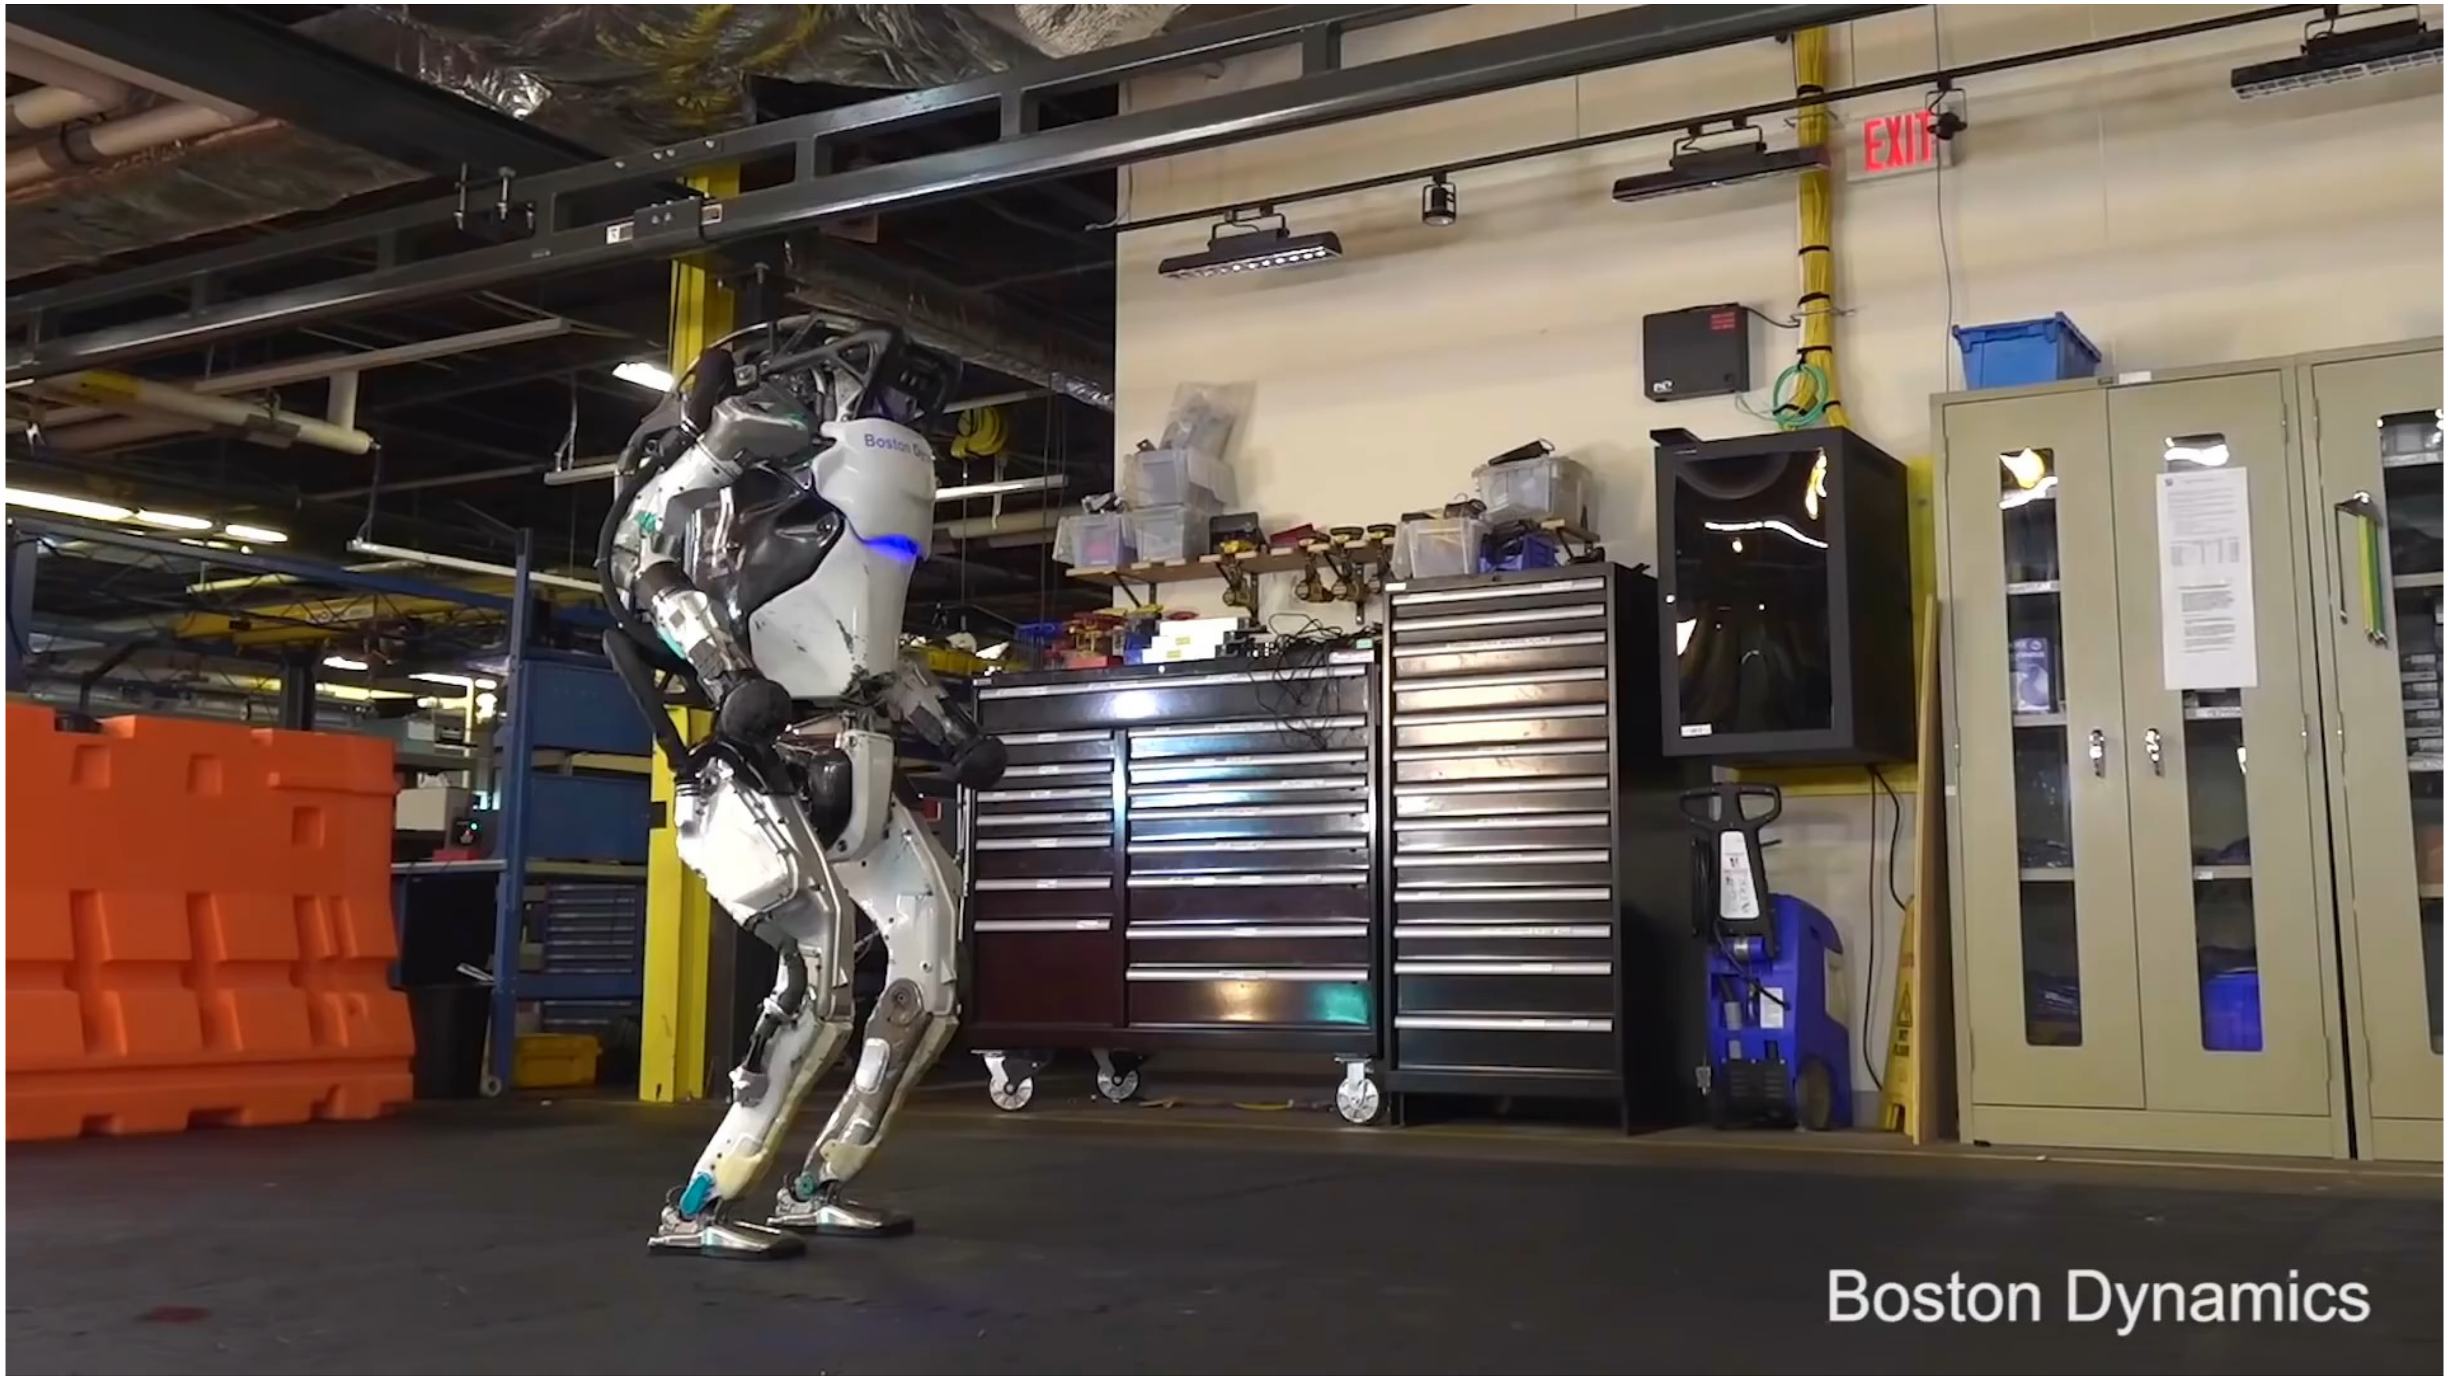

LINKÖPING UNIVERSITY

SPOT's got an Arm!

A Possible Future

Ray Kurzweil - Law of Accelerating Returns

"21st century will achieve 1000 times the progress of the 20th century"

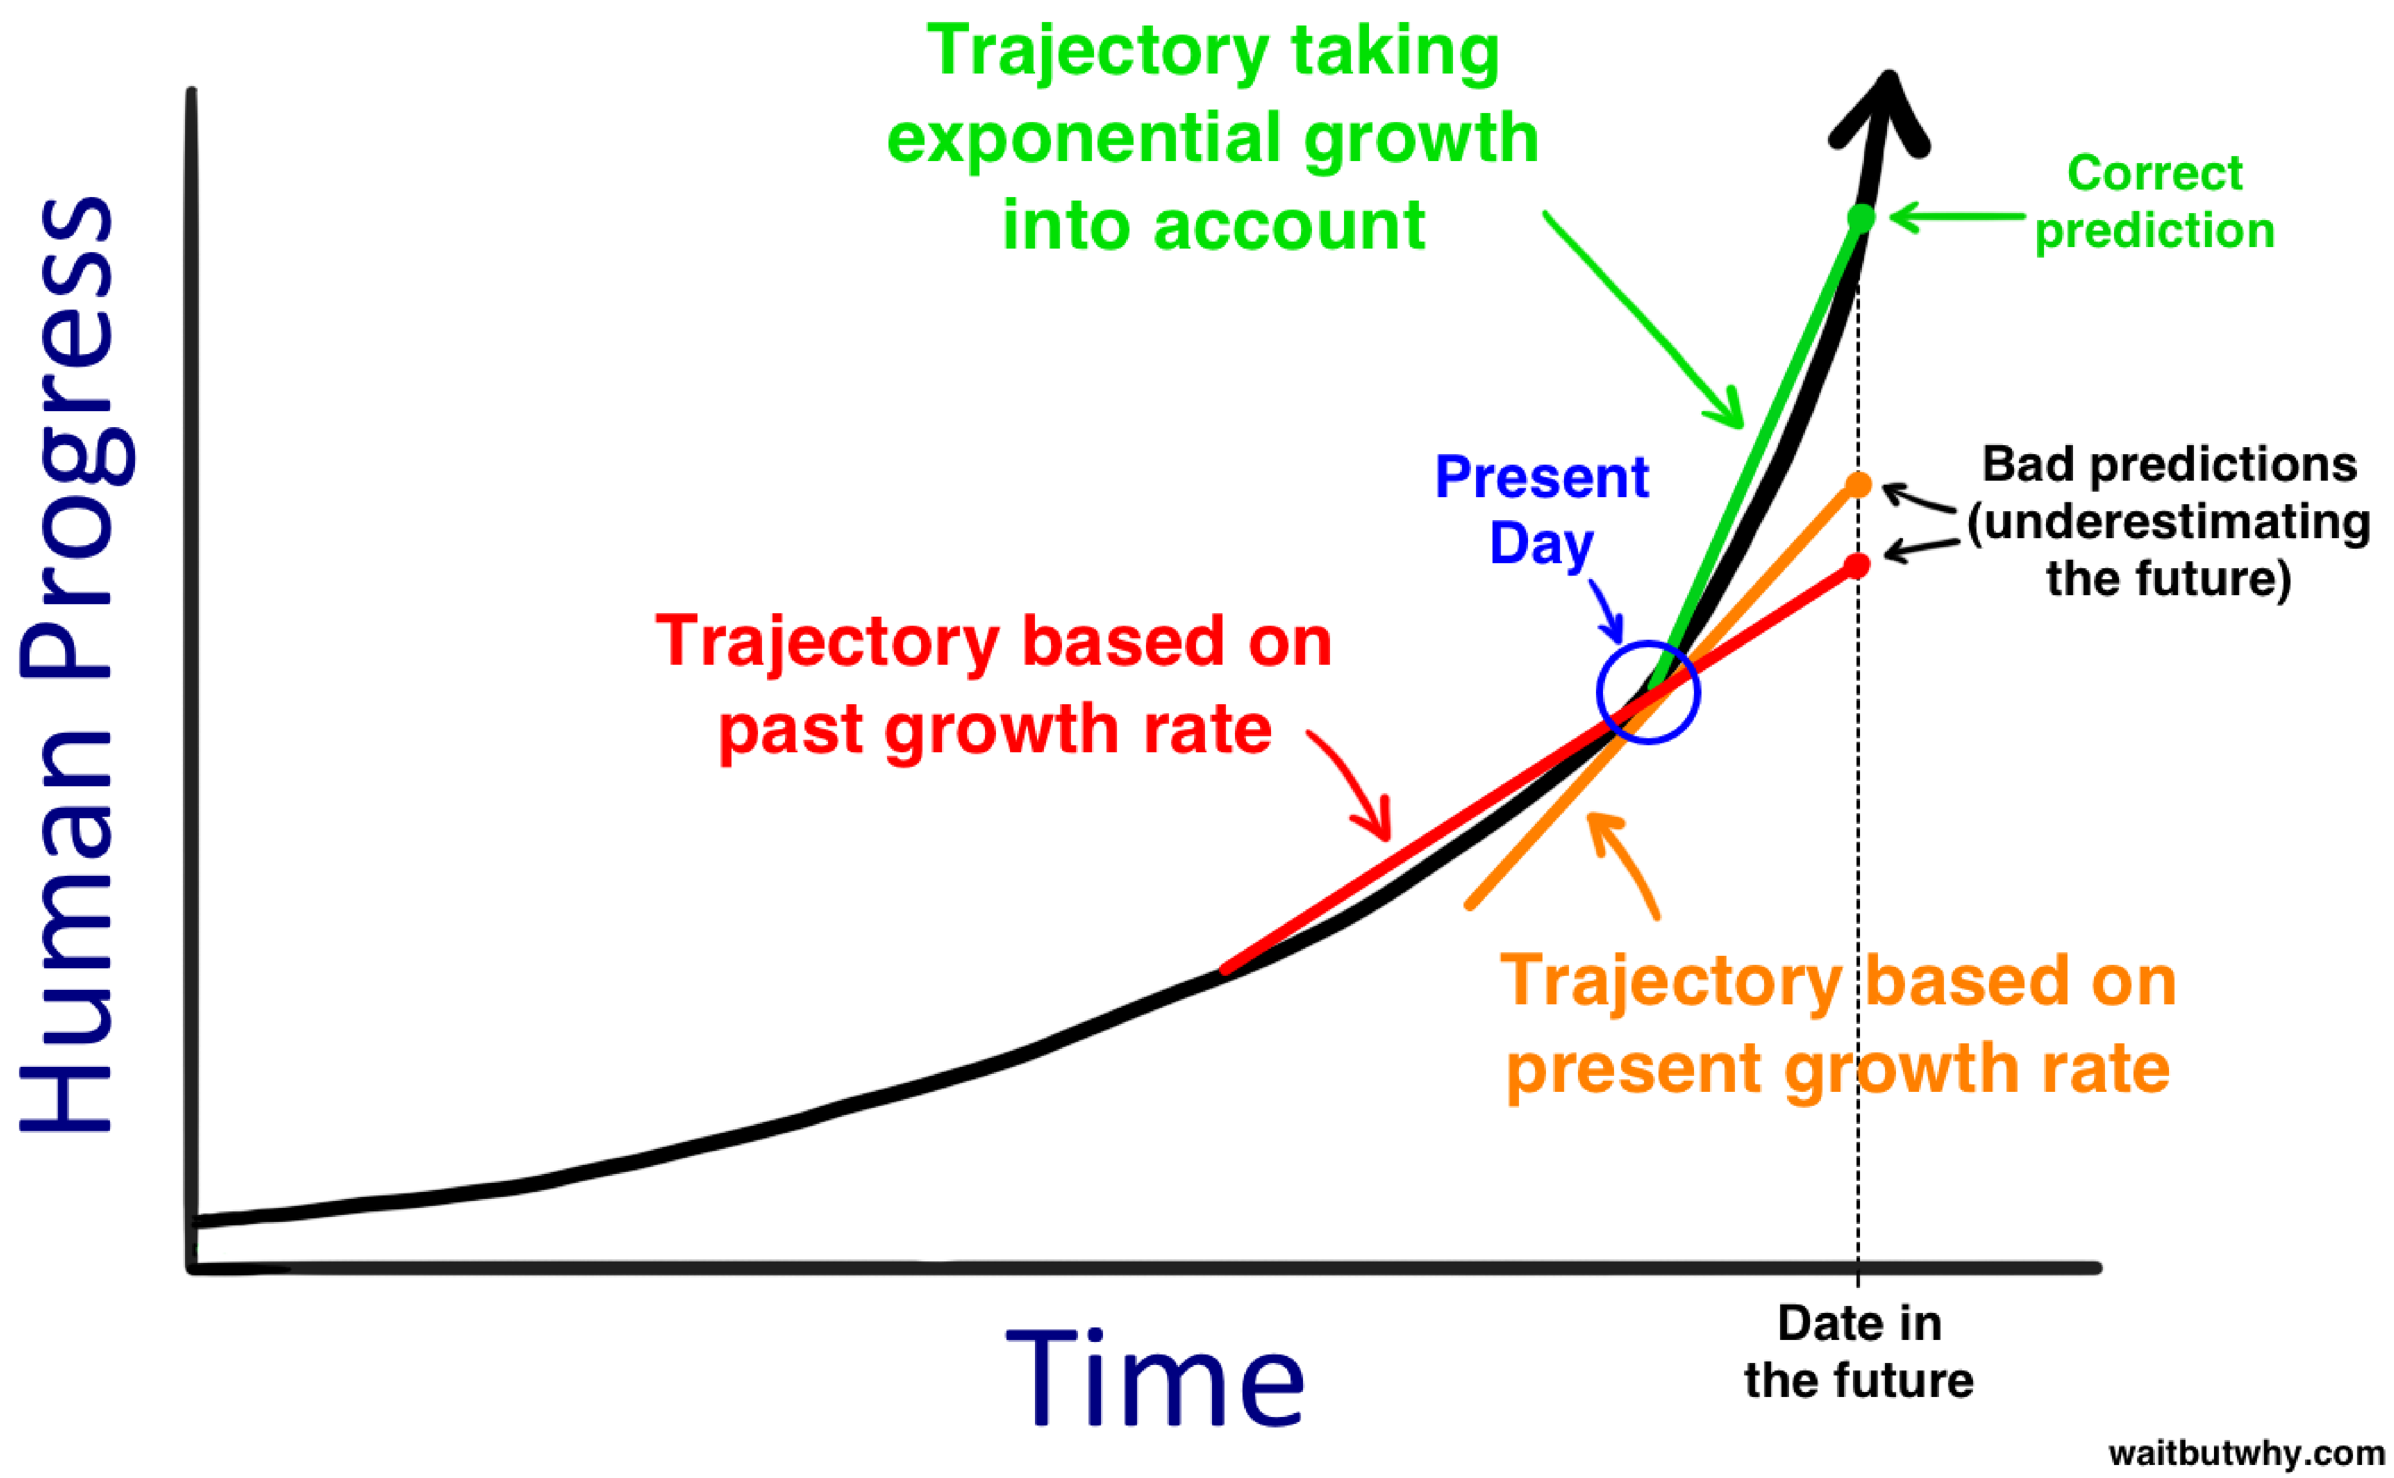

Kurzweil's View of Computational Power

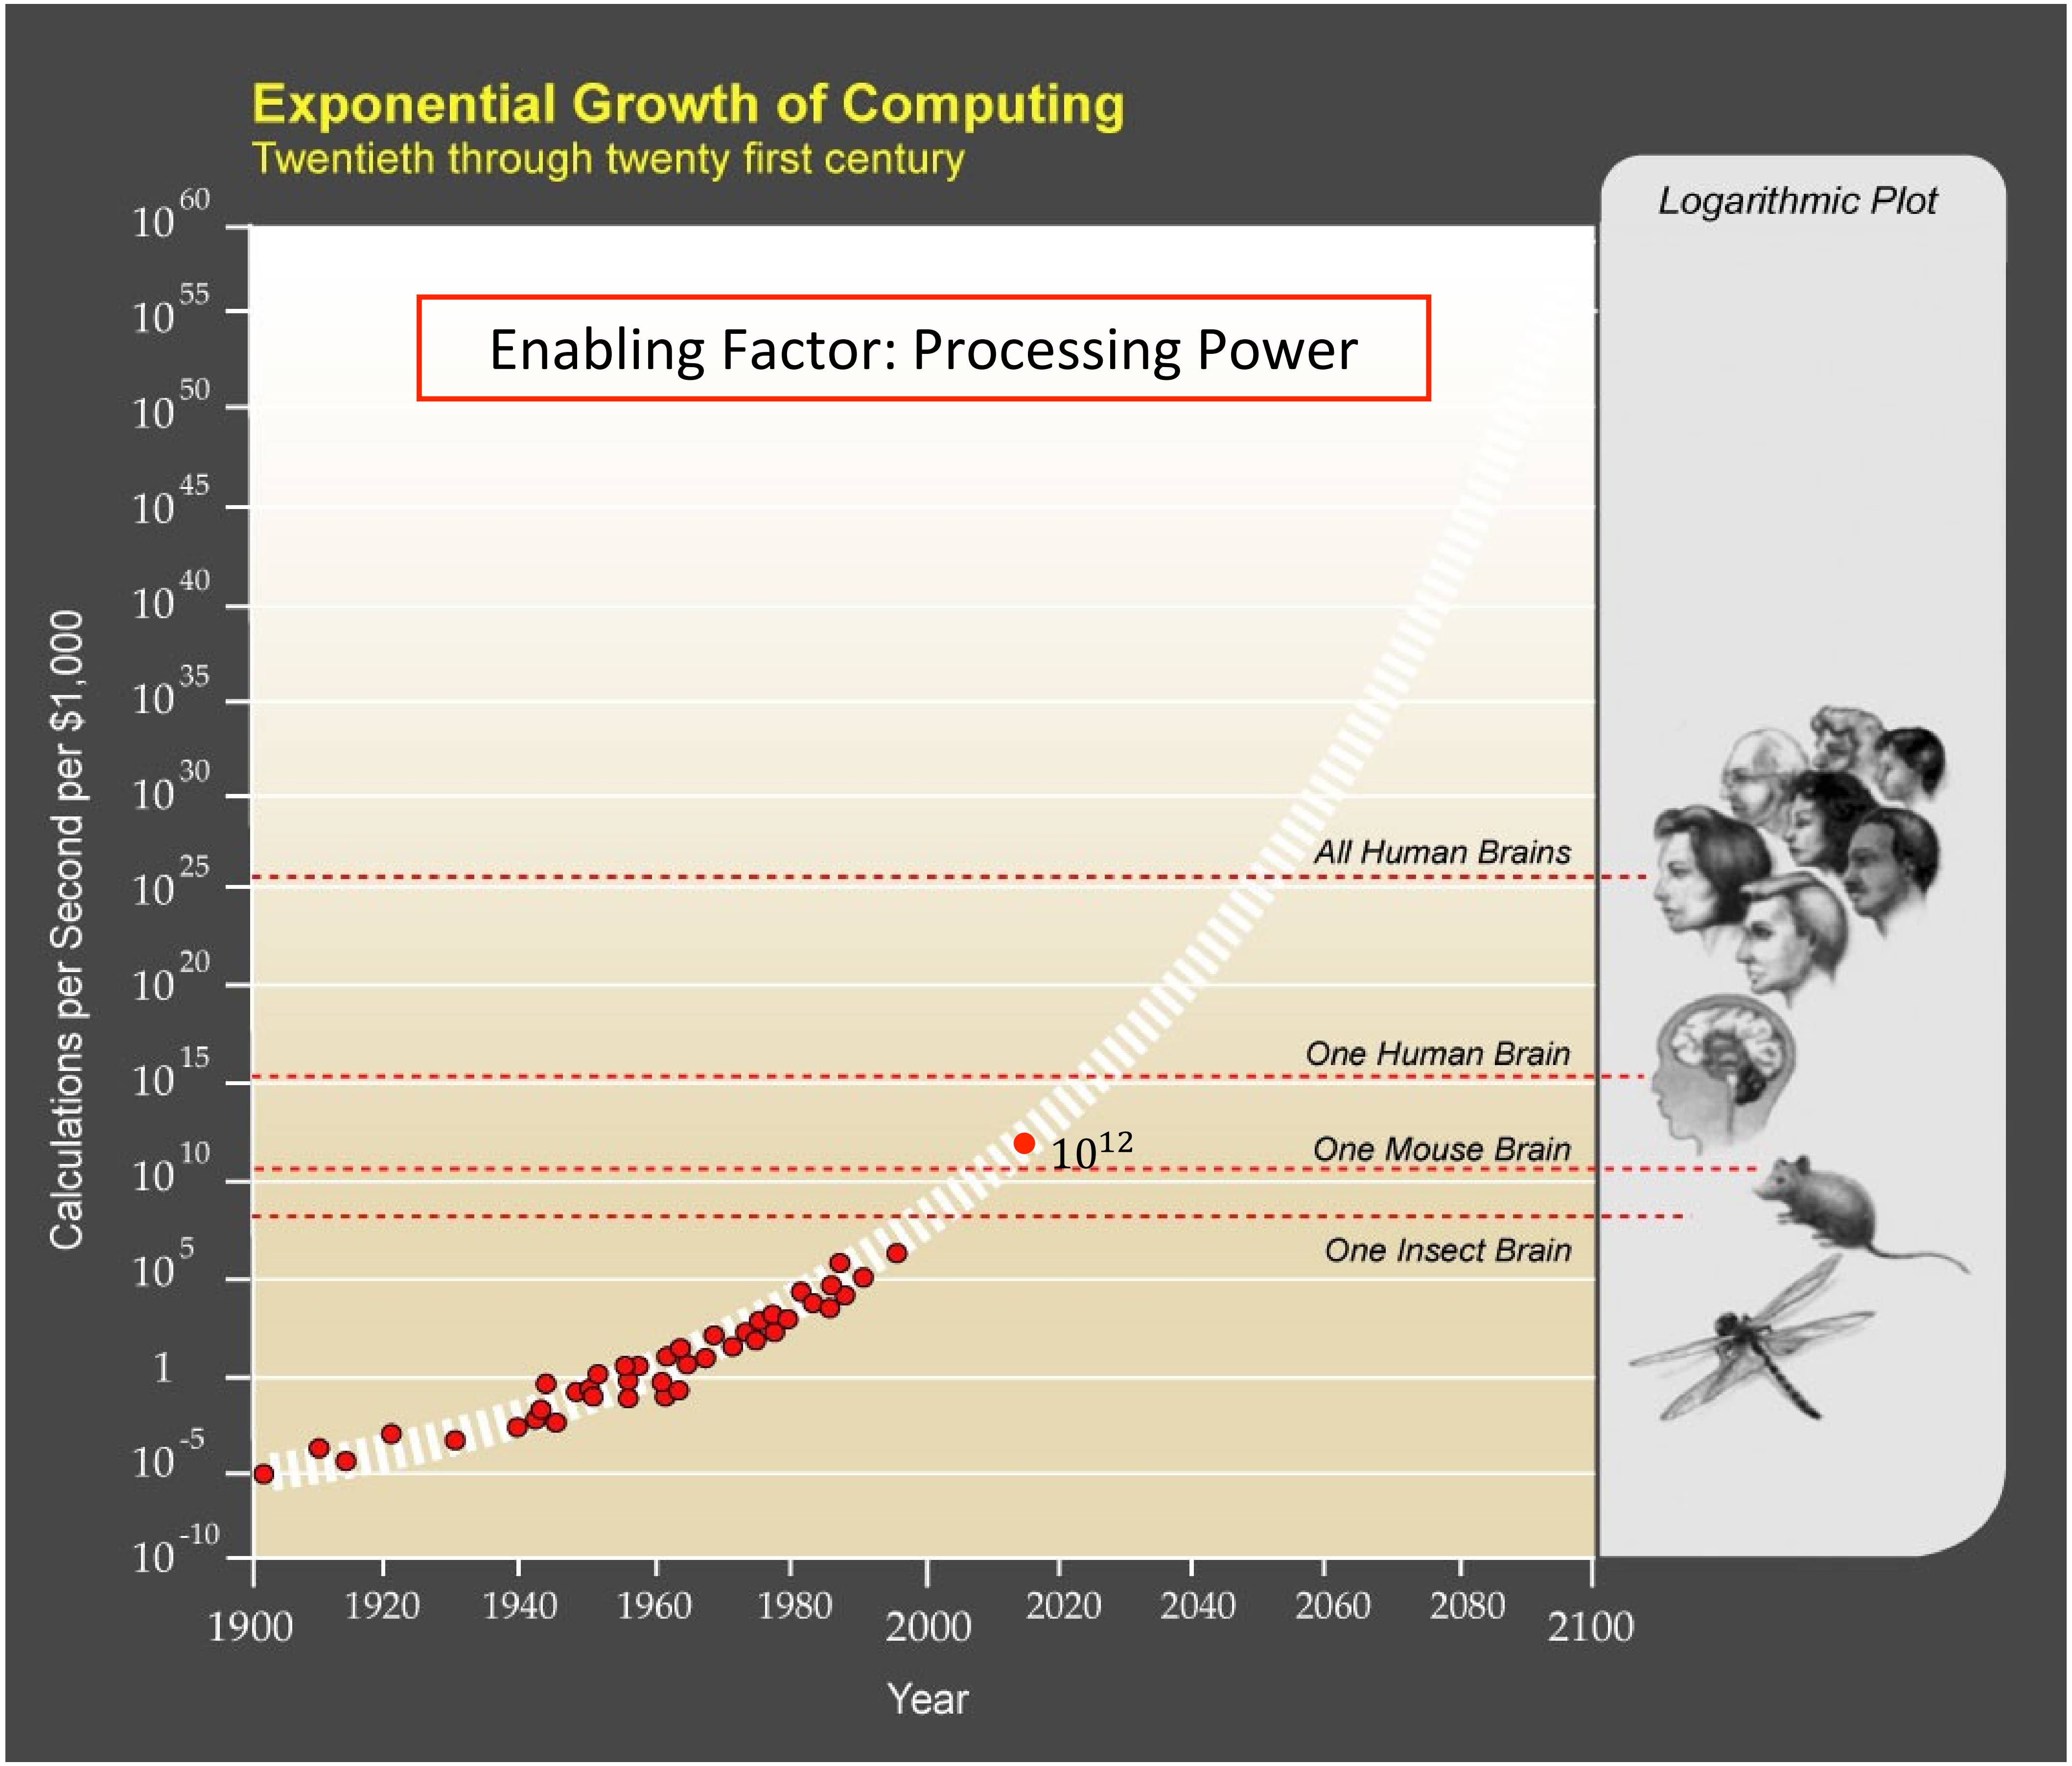

Enabling Factor: Lots of Data

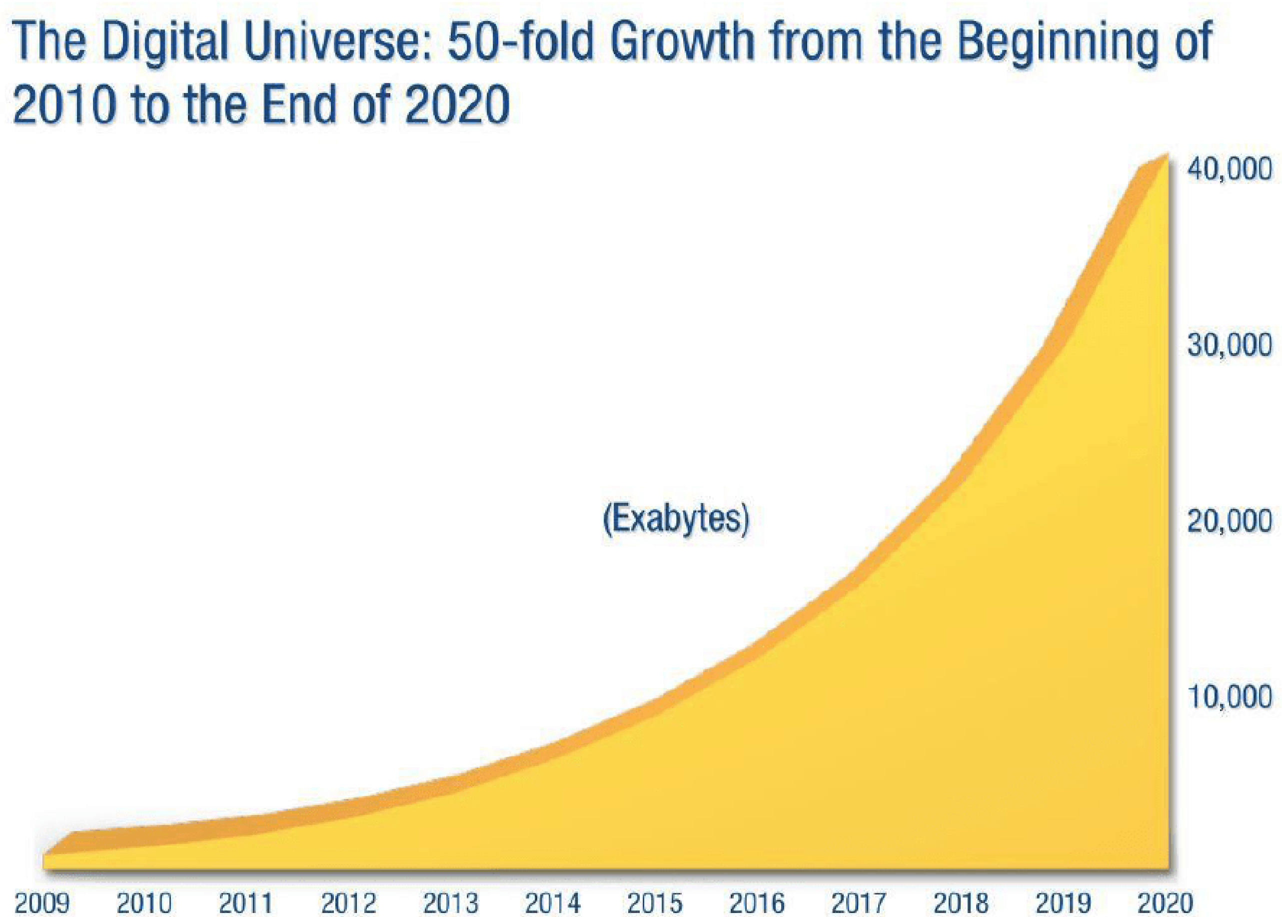

This IDC graph predicts exponential growth of data from around 3 zettabytes in 2013 to approximately 40 zettabytes by 2020. An exabyte equals 1,000,000,000,000,000,000 bytes and 1,000 exabytes equals one zettabyte. Source: IDC's Digital Universe Study, December 2012, http://www.emc.com/collateral/analyst-reports/idc-the-digital-universe-in-2020.pdf.

$$1\ Exabyte = 10^{18}\ Bytes$$
$$= 10^{9}\ Gigabytes\ (GB)$$

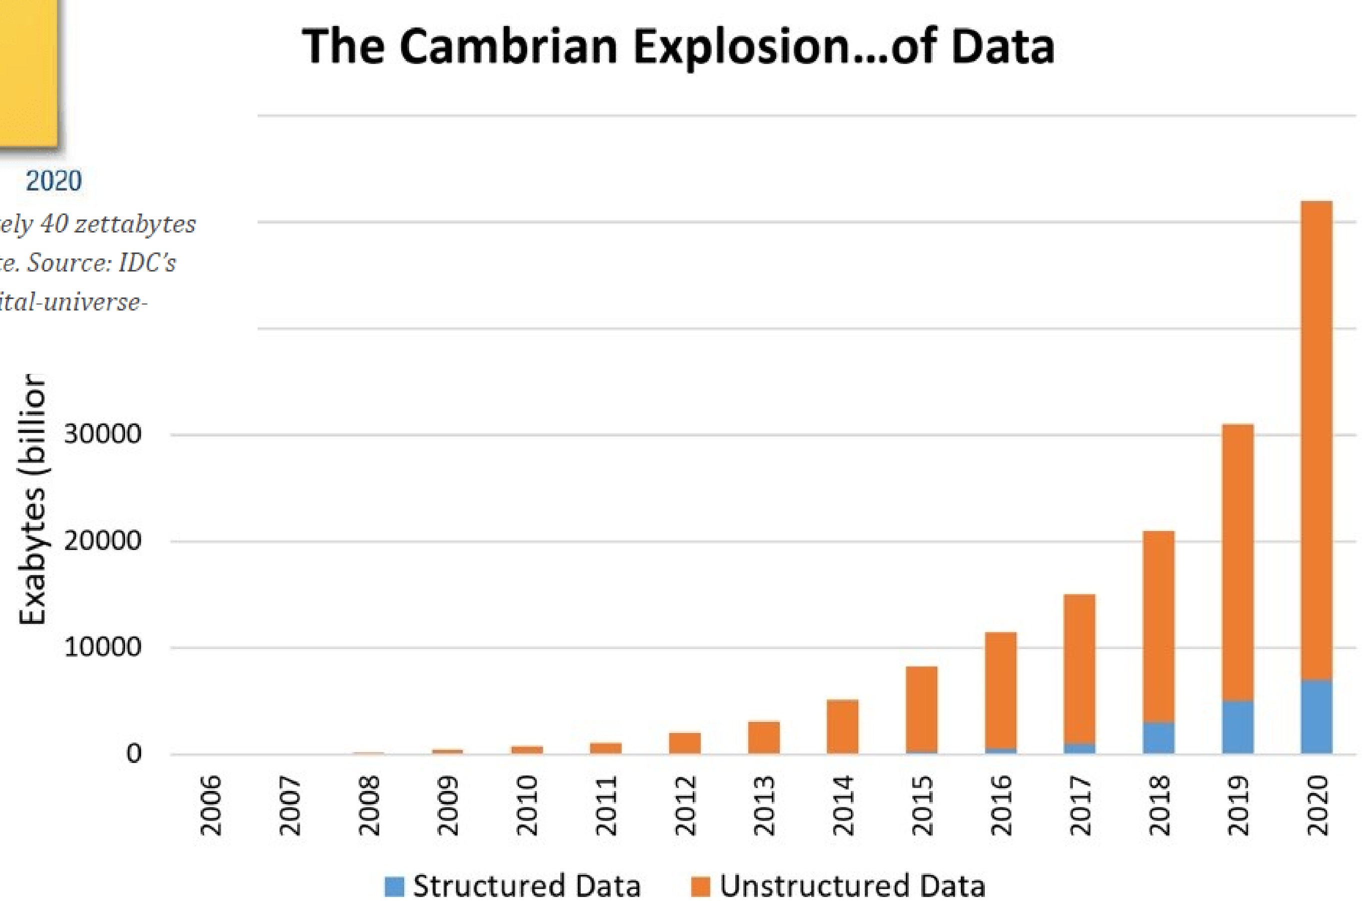

Summary / Overview

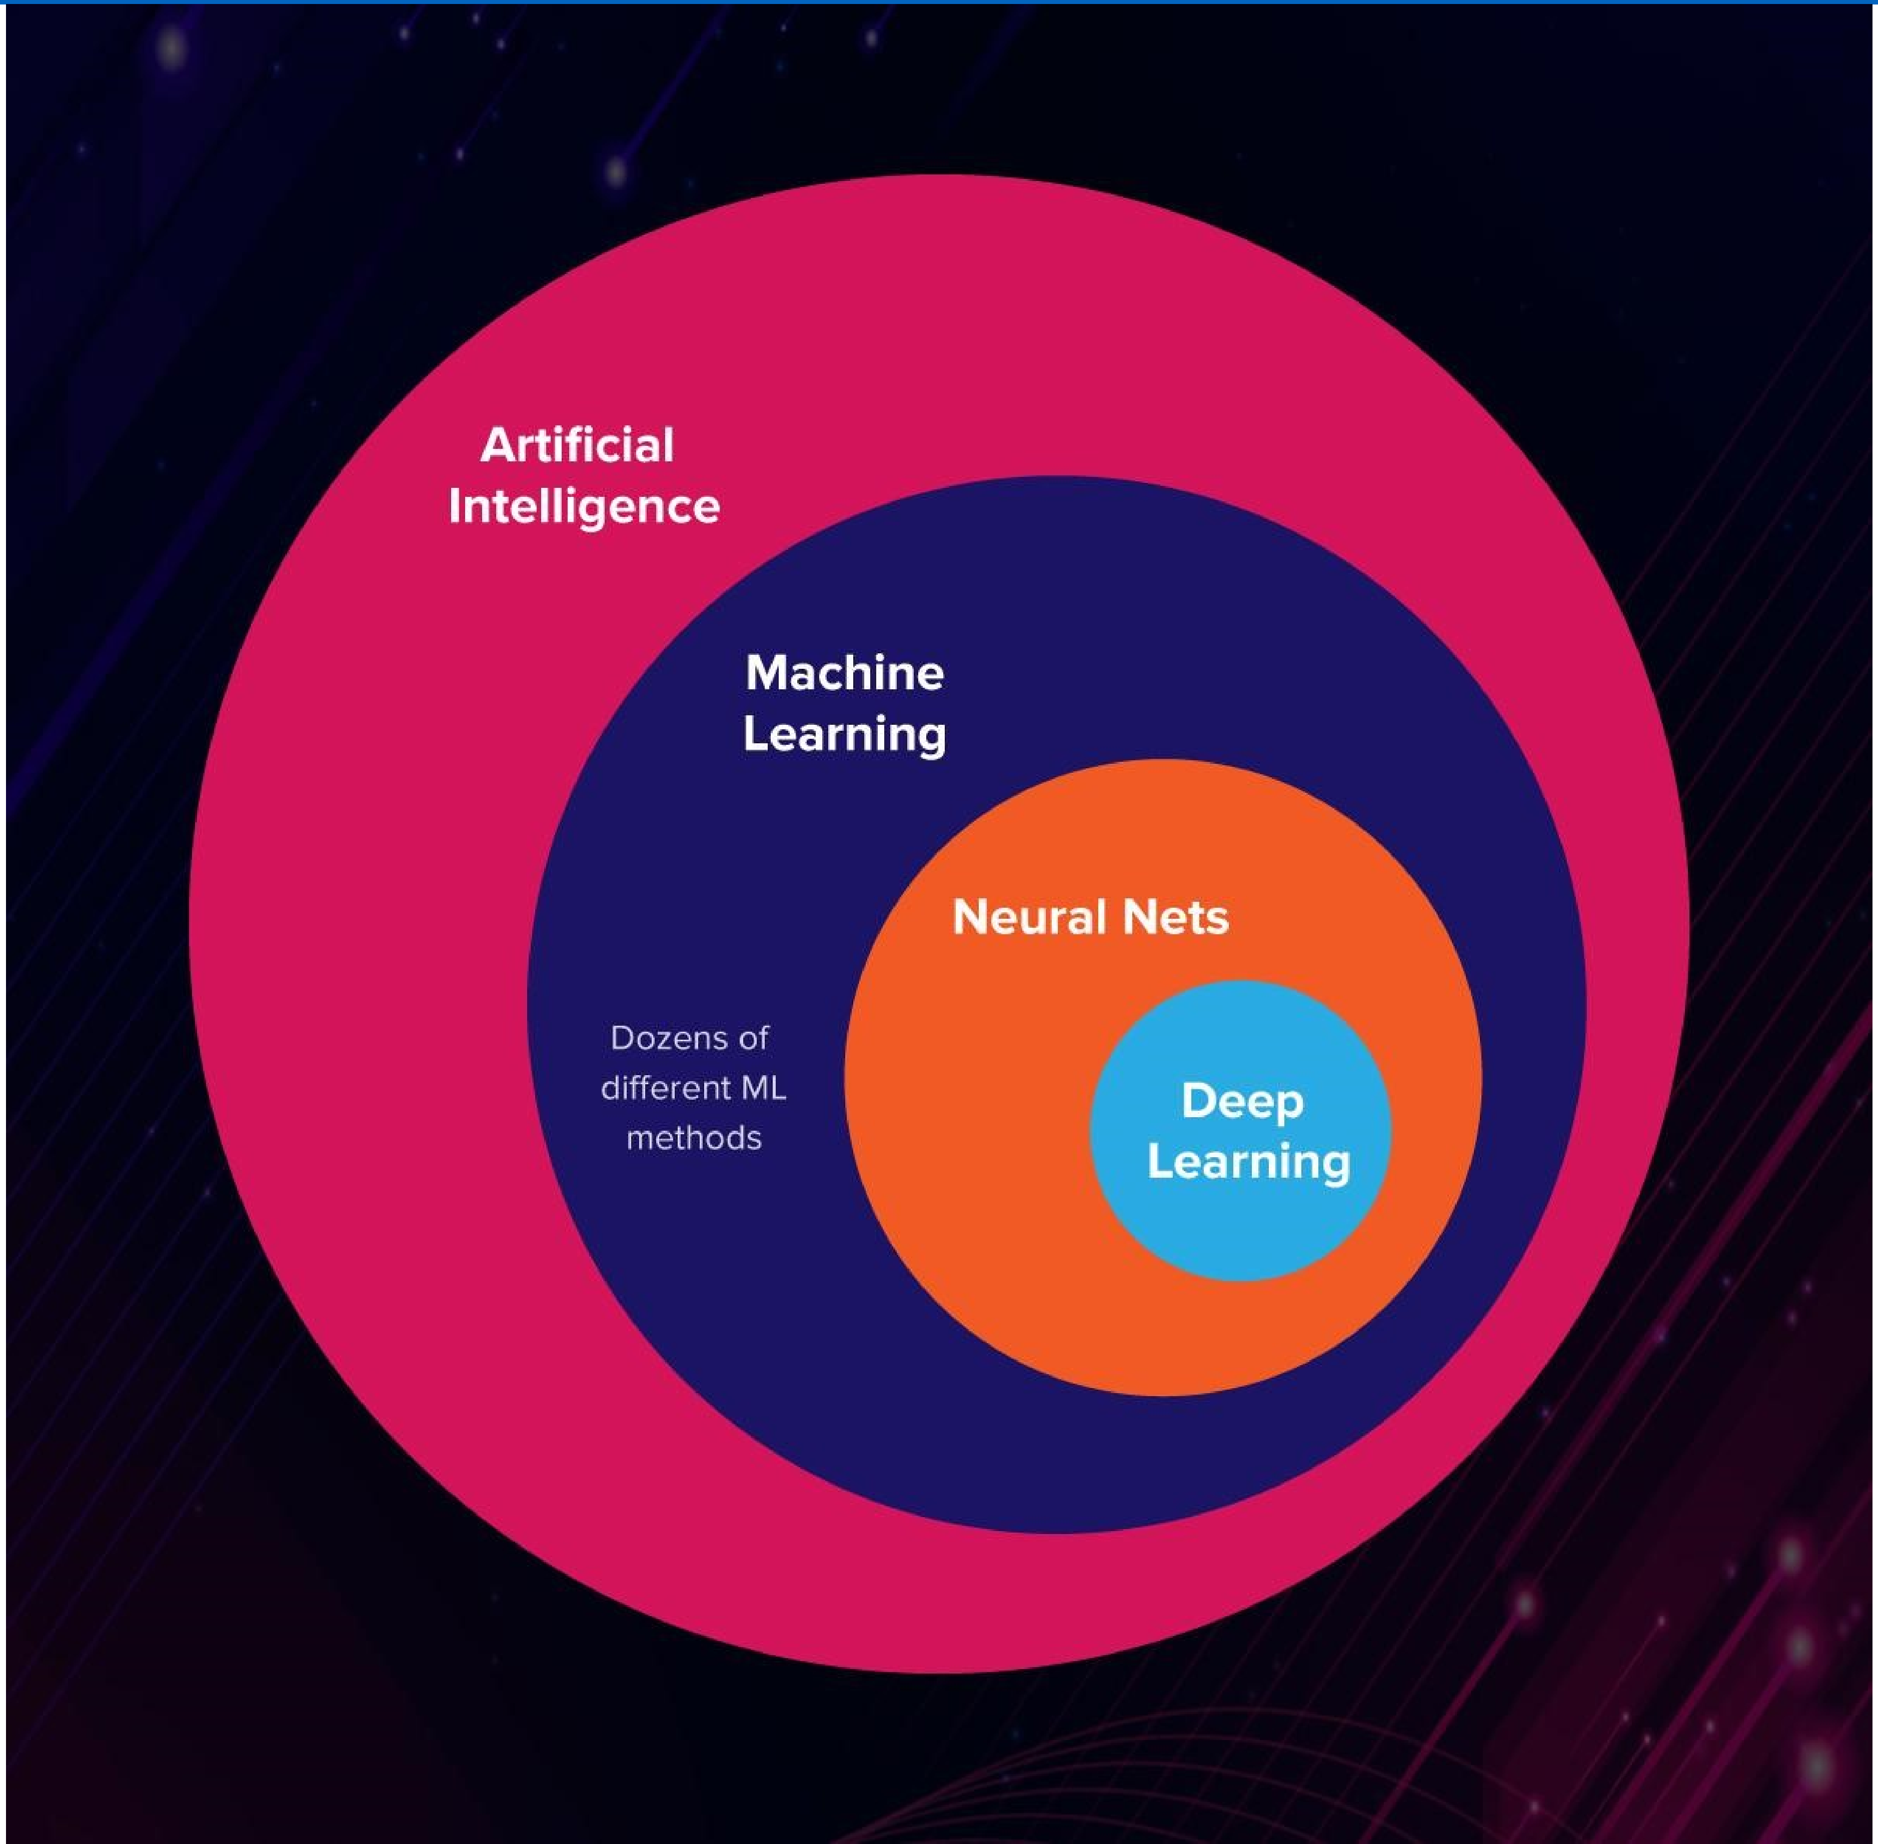

LINKÖPING
UNIVERSITY